

Mutual Information for the Detection of Crush Conditions



Manchester
Metropolitan
University

Thesis presented to the Faculty of Science and
Engineering, in partial fulfillment of the Requirements for
the Degree of Doctor of Philosophy

September, 2012

Author:

Peter HARDING

Supervisor:

Prof. Martyn AMOS

Ind. Collaborator:

Dr. Steven GWYNNE

Contents

Table of Contents	i
List of Tables	vi
List of Figures	vii
List of Equations	xv
Declarations	xvii
Acknowledgements	xviii
Abstract	xix
1 Introduction	1
1.1 Scope of Study	3
1.2 Identification of Crush Conditions in <i>in silico</i> Simulations	3
1.3 Contributions	4
1.4 Thesis Outline	5
2 Evacuation and Crush	7
2.1 Introduction	7
2.2 What is Evacuation?	7
2.3 The Behaviour of Evacuating Crowds	9
2.3.1 Fallacies of Crowd Behaviour	9
2.4 What are Crush Conditions?	13
2.5 What Causes Crush?	14
2.5.1 Spatial	14
2.5.2 Temporal	15
2.5.3 Perceptual and Cognitive Factors	17
2.5.4 Procedural	18
2.5.5 Structural	19
2.6 Types of Force in Evacuation Scenarios	19
2.6.1 Pushing	19
2.6.2 Stacking	20

2.6.3	Leaning	23
2.7	Historical Examples of Crush	25
2.7.1	Hillsborough	25
2.7.2	Rhode Island Nightclub	32
2.7.3	Gothenburg Dancehall	37
2.7.4	E2 Nightclub Incident	42
2.7.5	Mihong Bridge Spring Festival Disaster	46
2.8	A Diagnosis Issue in Crowd Crush Situations	52
2.9	Detecting Crush Conditions via Phase Transitions - An Initial Idea	54
2.10	Scope of this Study	55
2.11	Summary	56
3	Computational Studies of Evacuation	57
3.1	Introduction	57
3.2	Computational Evacuation Modelling	57
3.2.1	Representations of Physical Environment	58
3.2.2	Behavioural/Movement Modelling	65
3.3	Recent Trends in Research Activity	67
3.3.1	Trends in Model Usage	69
3.4	Modelling of Crush	72
3.4.1	Implicit	72
3.4.2	Explicit	74
3.5	Difficulties in Modelling Injuries Caused by Force	75
3.6	Our Proposed Approach	76
3.7	Summary	77
4	Social Forces Model	79
4.1	Introduction	79
4.2	Background	79
4.3	Description	80
4.4	Mathematical Definition	81
4.4.1	Movement	82
4.4.2	Social Force	82
4.4.3	Boundary Force	85
4.4.4	Goal Finding	85
4.5	Visual Example	85

4.6	Critical Analysis	87
4.6.1	Limitations	88
4.7	Time-lapsed Visual Example	95
4.8	Differences in the FDS+Evac Implementation	95
4.9	Summary	98
5	Mutual Information	100
5.1	Introduction	100
5.2	Background	100
5.3	Joint Entropy	102
5.3.1	Conditional Entropy	102
5.4	Mutual Information	103
5.4.1	Conditional Mutual Information	104
5.4.2	Multi-variate Mutual Information	105
5.4.3	The Logarithmic Base	106
5.5	MI Example	107
5.6	Information Theory	108
5.7	Biological Stimuli-response Systems	110
5.8	Chaotic Systems	111
5.9	Medical Image Processing	112
5.9.1	Normalised Mutual Information	113
5.9.2	Higher Dimensional Mutual Information	114
5.10	Complex Systems	116
5.11	Summary	121
6	Detection of a Phase Transition	122
6.1	Introduction	122
6.2	Scalar Noise Model	122
6.2.1	Phase Transition Point	124
6.2.2	Susceptibility	125
6.2.3	Methodology	126
6.3	Development	128
6.4	Experimentation	130
6.4.1	Multivariate Mutual Information	130
6.4.2	Conditional Mutual Information	132
6.4.3	Normalised Mutual Information	133
6.4.4	Studholme's Mutual Information	134

6.4.5	Wicks' Mutual Information	135
6.5	Correlation Analysis	137
6.6	Conclusions	138
7	Initial Evacuation Simulations	140
7.1	Introduction	140
7.2	Hypotheses	140
7.3	Experimental Aims	141
7.4	Expected Outcomes	141
7.5	Methodology	142
7.6	Experimentation	145
7.6.1	Single Run	148
7.6.2	Partial Data Analysis	149
7.7	Calibration	150
7.7.1	X and Y Discretisation	151
7.7.2	Orientation Discretisation	151
7.8	Negation of False Positives	152
7.9	Results	154
7.10	Hypothesis Testing	155
7.11	Distribution of Force Across Agents	155
7.12	Comparison with the Crowd Density Measure	157
7.13	Summary	158
8	Analysis of a Historic Event	160
8.1	Introduction	160
8.2	Fire Dynamics Simulator	160
8.3	The Station Nightclub Disaster	161
8.4	Methodology	161
8.4.1	Experimental Setup	162
8.4.2	Validation	165
8.4.3	Detection of Crush	168
8.4.4	Idealised Scenario	171
8.4.5	Realistic Scenario	172
8.4.6	Correlation Analysis	172
8.5	False Positives	173
8.5.1	Specification	174
8.5.2	Baseline	174

8.6	Summary	178
9	Conclusions and Recommendations	179
9.1	Conclusions	179
9.2	Recommendations for Further Research	181
	Bibliography	182
A	Publications	198
	Mutual Information for the Detection of Crush Conditions - Jour- nal article	198
	Mutual Information for the Detection of Crush Conditions - Con- ference Publication	209
	Prediction and Mitigation of Crush Conditions	215

List of Tables

2.1	Prisoners' Dilemma payoff matrix	11
2.2	Pedestrian Levels of Service (LoS). Available space (s) is measured in m^2 per person [37].	15
3.1	List of Evacuation models considered during literature review. Note: A model's inclusion in this table does not imply that it is widely used or actively developed, just that research papers exist which have explicitly referenced the model in question.	70
5.1	The inputs, X and Y , and the corresponding output, Z , of the dual input binary XOR gate, as shown in Figure 5.2	106
5.2	Probability distributions for variables X and Y	107
5.3	Joint probability distributions for variables X and Y	107
6.1	Variable values used during the experimentation with the SNM.	128
6.2	The absolute values of the Spearman's rank correlation coefficient, and the associated p-values, of the MI metrics and χ , the susceptibility of the system, across all test runs of the SNM.	138
8.1	The probabilities of agents starting at each area of the building knowing of the existence of of each of the possible exit routes.	164
8.2	Comparison of valid exit metrics for the two NIST simulations [44], using Simulex [139] and EXODUS [46], and our simulations using FDS.	167

List of Figures

2.1	Left: A visual example of stacking, in which the force of the person above is distributed across the body of those below. This situation may well lead to compressive asphyxia, but is unlikely to result in diagnosable trample injuries as the weight is more evenly distributed about the body. Right: A visual example of trampling, in which the entire weight of the person on top is channelled down through the feet to small portion of the body underneath. This situation may well lead both trample injuries and compressive asphyxia, as the person above transmits force to a highly localised area.	22
2.2	A visual example of the values used in Smith and Games[125] leaning crowd model. The figure on the right displays the lean angle, θ , the centre of mass height, H , and the gravitational force acting on the individual, mg . The figure on the right shows the push height, H' , and how the push force, $P_{0...N}$ propagates through a line of people.	24
2.3	Bottom: Safe occupancy levels of Pen 3 at Hillsborough stadium. Top: Estimated actual occupancy and density of Pen 3 at the time of the disaster. All internal sections represent the crowd barriers in place at the time.	28
2.4	Diagram showing the route that the influx of supporters took into the ground, and the routes at which fans already in pens 3 and 4 tried to take to alleviate the crush inside the pens [5].	29

2.5	Barrier 124A at the Lepping’s Lane terrace of Hillsborough stadium after the crowd crush. Figures estimate the force required to cause this crush barrier to collapse was in excess of $8kNm^{-1}$, or $8000Nm^{-1}$ [99].	29
2.6	Floor plan of the Station Nightclub. Ignition points mark the area in which the pyrotechnics were discharged [45]. .	33
2.7	Diagram showing the location of the deceased in the aftermath of the Station Nightclub fire. This diagram was compiled by the West Warwick, Police Department (Rhode Island), and was released during the investigation into the incident.	34
2.8	Top: Representation of the dancehall in Gothenburg [15]. Bottom: More detailed floorplan showing the first floor of the building, the stairwell at the south-west of the structure lead down to the only available exit. The similar stairwell at the east of the structure marks the ignition point of the fire, and was also blocked as it had been used to store furniture not in use during that night’s festivities.	38
2.9	A representation of the stairway leading from the E2 Nightclub down to South Michigan Avenue. The first floor doorway (left) opens into a landing at the top of the stairway, whilst the ground floor door (right) opens into a similar landing at the bottom of the stairway.	44
2.10	Top: The Miyun county “Rainbow Bridge”, the scene of the Miyun bridge disaster of 5 th February 2004. Bottom: Diagram showing dimensions of Miyun bridge (not to scale).	48
2.11	Top down density maps of Mihong bridge at different times from the start of the incident [155]	50
3.1	An atrium as may be found in large buildings, during evacuation inflow will occur at North, East, and West intakes, whilst exit will occur at the South.	61
3.2	Example of a cognitive behavioural model	66

3.3	Example Movement Model. Based on Helbing’s pedestrian model	67
3.4	Publications on evacuation modelling over the last twenty year period.	68
3.5	Number of papers published each year in which the evacuation model has been directly referenced. Data taken from <i>Science Direct</i> [32] (top), and from <i>Thompson Reuters Web of Knowledge</i> [141] (bottom).	71
3.6	Time-line showing the release of evacuation models included in our survey.	73
3.7	Figure A shows <i>laminar flow</i> , where the exit capacity appears sufficient for egress. Figure B shows the reduction of exit capacity inducing a transition into <i>turbulent flow</i>	77
4.1	Visualisation of interaction between four pedestrians. Forces f_1 and f_2 show the social force exerted on p_0 by p_1 and p_2 respectively. Pedestrian p_3 does not contribute to movement, as they lie outside the interaction area R_0 of p_0 . Vector m_0 shows the final movement of p_0 , away from p_1 and p_2 and a rate inversely proportional to their proximity.	81
4.2	The changing desire of agent i to decrease their proximity to agent j , shown across the entire interaction radius ($0m \geq R \leq 2m$).	84
4.3	The effect of the force co-efficient κ on agent movement at very high density. At proximities less than zero, the agents body representations are effectively overlapping.	84
4.4	Graphical example of movement vectors within the SFM. Figure One shows the component vector leading an agent away from walls and towards the goal, Figure Two shows the resultant movement vector.	86
4.5	The arching behaviour about the exit can clearly be seen during this simulation.	87

4.6	Body representation similar to that suggested by Thompson and Marchant [140], where R_t and R_s represent the radius of the torso and shoulders respectively, whilst R_d represents the radius of the body at its greatest point. . .	89
4.7	SFM agent distributions at differing values of desired velocity (V_0). Top: $V_0 = 0.5ms^{-1}$. Bottom: $V_0 = 5.0ms^{-1}$.	91
4.8	Barricading behaviour of the SFM. Injured pedestrians (Red) become immovable obstacles, which the active pedestrians (Green) are unable to circumvent. The system will remain in this state permanently once this behaviour has been established.	94
4.9	Example visual output at different times during a single evacuation using the social forces model, $V_0 = 5ms^{-1}$, $N = 200$. Force listed as per the traditional Helbing model, i.e. radial force about the specific circumference of each evacuee.	96
4.10	A graphical representation of an FDS+Evac “flow field” [72]. In this case all evacuees would head towards the left exit, using the flow field (blue arrows) to find their way towards this exit.	98
5.1	Venn diagram displaying a quantitative representation of the Shannon entropies, joint entropy, Conditional entropies, and the mutual information of a pair of variables X and Y	104
5.2	A dual input, single output XOR gate. X and Y are binary inputs, whilst Z forms the binary output. Table 5.1 contains the input and output data	106
5.3	Schematic diagram of general communication system. Information A is transmitted across the channel, where it is modified by with the addition of noise (source C), the resulting information content received at point B will therefore only contain a subset of the original content ($B \subset A$). 109	109

5.4	Visual example of Fraser-Swinney recursive subdivision. Using this method the calculation space may be discretised according to the amount of information they contain. Referencing is achieved by <i>tree notation</i> , therefore R_x is referenced as $R(0, 3, 1)$	112
5.5	An example of the type of error that can cause a rise in MI even though the images are becoming increasingly misregistered.	114
5.6	Comparative depictions of the level of shared information measured during a) McGill's [94] traditional MMI technique, and b) the technique used by Studholme [133]. In diagram B, the darker shaded area in the centre of the Venn diagram denotes that this central union is counted twice when calculating the final MI value	116
5.7	Visualisation of the heading update for particle i , utilising all particles within radius R , corresponds to term $\langle \theta_n^{NR} \rangle$ in Equation 6.3.	118
5.8	Three snapshots of the behaviour of the SNM under varying values of η . A) $\eta = 0$, B) $\eta = \frac{\pi}{2}$, and C) $\eta = \pi$	118
5.9	Output from the SNM across different values of the Noise value (η), showing the susceptibility, the MI, and the Binder cumulant. Error bars represent minimum and maximum values, and are shown for MI only.	120
6.1	Visualisation of the heading update for particle i , utilising all particles within radius R , corresponds to term $\langle \theta_n^{NR} \rangle$ in Equation 6.3.	123
6.2	The changing values of the order parameter, ϕ , for different values of the noise parameter, η . The error bars represent the standard deviation of the values across the 64 simulation runs.	125
6.3	The value of the Binder cumulant, β , in the scalar noise model, calculated for different values of the noise parameter, η	126

6.4	The susceptibility of the scalar noise model, χ , calculated for different values of the noise parameter, η	127
6.5	Computation times, in seconds, of 60,000 time-steps of the SNM for varying values of the population size, N . The timings of both the Serial C++ implementation and the openMP implementation are shown. The y-axis in this figure is logarithmically scaled.	130
6.6	The multivariate mutual information, $I(X, Y, \Theta)$, measured across 40 runs of the scalar noise model across different values of the noise parameter, η . The error bars represent the standard deviation of the data.	131
6.7	The standard deviation of $I(X, Y, \Theta)$ measures across all experiment runs.	132
6.8	The conditional mutual information of X and Y conditional on Θ , $I(X, Y \Theta)$, across changing values of the noise parameter, η . The error bars represent the standard deviation of values across all test runs.	133
6.9	The Trivariate Normalised Mutual Information of the SNM measured for changing values of η . Error bars represent the absolute standard deviation from the mean.	134
6.10	Studholme's Multivariate Mutual Information of the SNM measured across changing values of the noise parameter, η . Error bars represent the absolute standard deviation from the mean.	135
6.11	The Wicks Mutual Information, $I(X, Y \Theta)$, measured across varying values of the noise parameter, θ . The error bars represent the standard deviation across all test runs.	136
6.12	The variance of the WMI, as calculated from the entirety of the test runs.	137
6.13	The values of Wicks' MI (blue), susceptibility (green), and Binder cumulant (red) calculated from the scalar noise model for different values of the noise parameter, η . Error bars present represent the standard deviation of the data.	139

7.1	Representation of the evacuation topology used during the simulations in this Chapter. All evacuees are distributed evenly across the central 15m-by-15m room, and effect egress through the 2m wide door.	143
7.2	Image showing the barricading behaviour of the SFM. Red agents represent those incapacitated due to injury, whilst green agents are free to move, but unable to escape past the injured parties. In these situations the simulation will continue indefinitely with no resolution.	144
7.3	Results from the MI analysis of the SFM, MI was calculated according to Equation 7.1	146
7.4	MI for this result was calculated according to Equation 7.2. Significant improvements can be seen, most importantly the improved definition of the phase-transition displayed at $0 > t < 4$. This data represents the average of 64 identical experimental runs.	147
7.5	The <i>smoothed</i> output from a single run of the SFM. . . .	148
7.6	Results of the partial data MI analysis of the SFM. . . .	150
7.7	Absolute correlation (both Pearson's and Spearman's) between force and MI, across different binning values for variables b_x and b_y	152
7.8	Absolute correlation (both Pearson's and Spearman's) between force and MI, across different binning values for variable b_θ	153
7.9	Simulation showing the changing values of Force and $f(I, \rho_{max})$	154
7.10	An example of the forces measured from each agent at multiple points during a simulation identical to that presented in this Chapter.	156
7.11	Average (arithmetic mean) values found during of 64 simulations in which the maximum crowd density (pm^{-2}) and force (Nm^{-1}) were recorded. The crowd density and the force recorded during the entire milling period ($t < 0$) were $1pm^{-2}$ and $0Nm^{-1}$ respectively.	157

8.1	(Top) Floorplan of Station nightclub, taken from official report. (Bottom) Rendering in FDS+Evac. [44]	163
8.2	Results of NIST recreation (left) and simulation (right) of the fire and smoke spread during the initial 90 seconds of the Station nightclub fire. [44]	166
8.3	Comparison of leaving profiles between our simulation (FDS) and official NIST results.	167
8.4	Comparison of average force between real and idealised scenarios.	168
8.5	Screenshot of our fire scenario simulation after 65 elapsed seconds.	170
8.6	Comparison of Mutual Information between idealised and actual scenarios.	172
8.7	Scatterplot of Force versus Mutual Information.	173
8.8	Topography of test configuration. Position A marks the centre of the entry point for pedestrians, position B marks the centre of the exit.	175
8.9	Images showing the false positives tests after a sufficient amount of time for the system to settle into a representative state. Top: First test, showing low usage. Bottom: Higher usage test.	176
8.10	MI (green) and Average Force (red) against time for the simulation described previously.	176
8.11	MI (green) and Average Force (red) against time for the simulation in which the evacuation capacity of the structure is overwhelmed.	177

List of Equations

4.1	SFM Movement Equation	82
4.2	SFM Social Force Equation	82
4.3	SFM Boundary Force	85
4.4	SFM Goal Force	85
4.5	SFM Physical Force with Damping	90
5.1	Hartley Entropy	101
5.2	Shannon Entropy	102
5.3	Joint Entropy	102
5.4	Conditional Entropy	103
5.5	Conditional Entropy Equivalence	103
5.6	Bi-variate Mutual Information	103
5.7	Bi-variate Mutual Information Equivalences	104
5.9	Multi-variate Conditional Mutual Information	104
5.10	Multi-variate Mutual Information	105
5.11	N-variable Mutual Information	105
5.12	Time-shifted Mutual Information	111
5.13	Bi-variate Normalised Mutual Information	113
5.14	Studholme’s Tri-variate Mutual Information	115
5.15	McGill’s Tri-variate Mutual Information Equivalences	115
5.16	Studholme’s Tri-variate Mutual Information Equivalences	115
5.20	SNM Coordinate Update	117
5.20	SNM Heading Update	117
5.20	SNM Velocity Update	117
5.21	Wicks’ MI Calculation	119
5.22	Binder Cumulant	119
5.23	Order Parameter	119
5.24	Susceptibility	119
6.4	SNM Coordinate Update	123

6.4	SNM Heading Update	123
6.4	SNM Velocity Update	123
6.5	Order Parameter	124
6.6	Binder Cumulant	124
6.7	Scalar Noise Model Susceptibility	126
6.8	Tri-variate Mutual Information	130
6.9	Tri-variate Conditional Mutual Information	132
6.10	Tri-variate Normalised Mutual Information	133
6.11	Studholme's Tri-variate Mutual Information	135
6.12	Wicks' Mutual information	136
7.1	Geographic Mutual Information	145
7.2	Geo-directional Mutual Information	147
7.3	MI Scaling Equivalence	149
7.4	Density Normalised Mutual Information	153
8.1	MI calculation for the Station nightclub analysis	171

Declarations

This thesis is submitted to Manchester Metropolitan University in support of my application for admission to the degree of Doctor of Philosophy. No part of it has been submitted in support of an application for another degree or qualification of this or any other institution of learning. Parts of the thesis appeared in the following refereed papers in which my own work was that of a full pro-rata contributor:

P. J. Harding, M. Topsom, and N. Costen; **(2011)** *Mutual Information Based Gesture Recognition*; Computer Analysis of Images and Patterns; A. Berciano, D. Diaz-Pernil, W. Kropatsch, H. Molina-Abril, and P. Real (*Eds.*); Springer, Germany; p. 571-578

P.J. Harding, M. Amos, and S. Gwynne; **(2011)** *Mutual Information for the Detection of Crush*; PLoS ONE 6(12); e28747. doi:10.1371 / journal.pone.0028747

P.J. Harding, M. Amos, and S. Gwynne; **(2010)** *Mutual Information for the Detection of Crush*; Pedestrian and Evacuation Dynamics; R.D. Peacock, E.D. Kuligowski, and J.D. Averill (*Eds.*), Springer, U.S.A; p. 779-784

P.J. Harding, M. Amos, and S. Gwynne; **(2008)** *Prediction and Mitigation of Crush Conditions in Emergency Evacuations*; Pedestrian and Evacuation Dynamics 2008; W.W.F. Klingsch, C. Rogsch, A. Schadschneider, and M. Schreckenberg (*Eds.*); Springer, Germany; p. 233-246

Acknowledgments

I would first like to thank my supervisor, Dr Martyn Amos, for the help and guidance that he has given during the completion of this work. His input to the specifics of the project, and advice and experience of general research practice has been invaluable to this work, and will remain with me for the rest of my career.

Secondly, my industrial collaborator Dr Steve Gwynne, who has been of immense help with the specifics of crowd psychology and simulation, and an inspiration with respect to his constant enthusiasm for the project, and the field of evacuation science in general.

I would like to thank my parents, Anne and Barrie, my sister, Sarah, and my girlfriend, Emily, for the help, support, love and understanding that they have shown during my time completing this thesis, and throughout my life.

Monica de Carlos deserves specific thanks, as her help and guidance during my formative years were responsible for my choice of a career in the sciences.

Many office mates I've had over the years, for the time they have spent listening to my problems, and the useful advice and comments that they have given during my time spent with them. Angel, Ben, Chema, Hui, Jesus, John, Matt, Mike, Naomi, Sergei, and Tom have been an invaluable help during my time as a researcher, and also been responsible for many good times outside of work.

Finally I would like to thank the Dalton Research Institute and Manchester Metropolitan University for funding this research.

Abstract

Fatal crush conditions occur in crowds with tragic frequency. Event organizers and architects are often criticised for failing to consider the causes and implications of crush conditions, but the reality is that the prediction of such conditions offers significant challenges. This thesis investigates the use of crush metrics in simulation environments, which can be used to help quantify the danger of crush conditions forming during *real life* evacuations.

An investigation is carried out in the use of computer models for the purpose of simulating building evacuation. From this review we identify the most suitable methodologies for modelling crowd behaviour, and we detail the specific areas of functionality which must be in place before modellers can incorporate crush analysis into an evacuation simulation. We find that full treatment of physical force within crowd simulations is precise but computationally expensive; the more common method, human interpretation of simulation output, is computationally “cheap” but subjective and time-consuming.

A technique which admits a low computational cost alternative to the explicit modelling of physical force, yet still offers a quantitative metric for the level of force present during an *in silico* evacuation is proposed. This technique and the precise manner in which we apply it to the problem of crush detection is shown and we present the results of initial experiments.

To further test the ability of our technique to identify dangerous evacuation conditions, we recreate a well-known historical evacuation. Results of these experiments show that we do offer an effective and efficient route towards the low cost automatic detection of crush, and an alternative approach to traditional methods.

Chapter 1

Introduction

Overloading pedestrian routes can quickly lead to the development of *crush conditions* (the formation of dangerous levels of physical force within a crowd), as observed in the Hillsborough disaster [136], Station Nightclub [44] and Saudi Arabian Hajj [52] incidents, as well as the recent Love Parade tragedy in Germany [6]. Some suggest that approximately two-thousand individuals per year die as a direct result of crush conditions [56], and that this figure continues to rise [155]. A more sophisticated understanding of how crush conditions form is therefore critical for the design of tall buildings and other highly-populated, contained areas (such as ships, nightclubs and stadia), as well as for the planning of events and formulation of incident management procedures. A first step towards this deeper understanding is a method for *detecting* the early-stage formation of crush, which is the problem we address here.

The study of crowd evacuation scenarios has taken on additional significance in the light of events such as 9/11 . Many tall buildings (such as the World Trade Centre towers) were designed alongside the assumption that any necessary evacuation could and would be conducted in a *phased* manner (e.g. floor-by-floor). One significant factor in building design is the capacity of exit routes (such as corridors and stairwells). Capacities are calculated based on projections of *controlled* population movement in phased evacuations. If the phased evacuation assumption breaks down (if, for example, occupants of a specific floor refuse to wait their “turn” for fear of catastrophic building failure) then this will have severe implications for overall safety, as exit routes can become overloaded.

Computer-based simulation studies are often used to analyse the movement of individuals in various scenarios. Such work encompasses, for example, the study of historical events [52], the examination of evacuation procedures [47], and the design of aircraft [39]. Existing simulation frameworks include EXODUS [102], PEDFLOW [77] and EVACNET [67] (see [78] for an extensive review), and these offer a range of “real world” features, including exit blockage/obstacles, occupant impatience and route choice [48]. However, the phenomenon of *crush* is one that has received relatively little attention from the designers of evacuation simulations. Many simulations do not *explicitly* consider the effects of crush, and those that *do* factor in crush employ *computationally expensive* Newtonian force calculations.

The two major problems we address are as follows: firstly, the consideration of crush within existing simulation frameworks requires the use of computationally expensive Newtonian force calculations. These drastically slow down simulations, restricting their applicability in the *rapid prototyping* of building designs and crowd control procedures. The second problem is that the monitoring of crush within *real crowds* is rudimentary, at best, and relies largely on personal observation and interpretation of crowd patterns by security professionals [76]. This method of crush detection is inherently problematic.

We therefore seek a method for the detection of crush conditions that is relatively “cheap” in terms of computational effort, and which can, in future, be easily integrated into existing software for crowd monitoring. Such a method will have a significant impact on both simulation-based evacuation studies and real-time analysis of video images (facilitating, for example, the future development of automated crush alarms based on CCTV images).

In essence, we propose that the breakdown of order, i.e. smooth or laminar flow, within an evacuating crowd may be used to metricise the amount of crush danger that individuals within that crowd may face. As order breaks down, the predominant behaviour within the crowd will transition to a disordered phase, in which individuals exhibit competitive behaviours such as pushing or overtaking. We suggest that the identification of the level of order may therefore be used as a de-facto measure of the amount of force that will build within the evacuating crowd.

In this thesis we describe our proposed method, an information theory based technique which treats the onset of dangerous behaviours during an

evacuation as a form of *phase transition* within the evacuating crowd, i.e. an observable, and measurable, change from one set of exhibited behaviours to another. We show how our method can be easily integrated into an existing simulation framework, and test it using details of a historical event. Our results show that we can provide a robust warning indicator of the emergence of crush conditions.

1.1 Scope of Study

This study is concerned with the identification of the formation of crush conditions within *in silico* evacuations. The analytic technique developed will therefore focus on the *prediction* of physical force, or the *identification* of known dangers (during our historical recreation), in two-dimensional pedestrian evacuation simulations. The precise physical force that an *individual* pedestrian is subject to at any one time is not considered.

1.2 Identification of Crush Conditions in *in silico* Simulations

To prove the usefulness of our detection methodology we implement our technique as part of an *in silico* simulation environment, and test the technique on a recreation of an historical evacuation where high levels of force are known to be present. The advantage of using an *in silico* evacuation is twofold; firstly, there are a number of simulation environments which offer the ability to measure the physical forces building within crowds of people, which allows for the confirmation of the presence of dangerous levels of force using our technique. The second advantage of choosing an *in silico* environment is that obtaining empirical data relating to high density crowd situations is notoriously difficult, with good quality video footage being extremely rare. To mitigate the disadvantages inherent in the use of *in silico* data rather than *real-life* data, we only consider items of performance data from the *in silico* evacuations that could feasibly be obtained from the video feed of a “real-world” crowd event (i.e. variables describing an individual’s motion, such as density, velocity, or direction). In this way it is anticipated that the technique will eventually be shown to be applicable to the *real time* detection of crush conditions.

1.3 Contributions

1. **To identify factors that contribute to the development of crush conditions.**

We identify multiple contributing factors that lead to the initial formation of crush conditions, and show how they can be used for the analysis of crush formation in historical events.

2. **To review historical incidences of crush.**

We review numerous historical incidents in which crush conditions have been found to have caused serious injury or loss of life. The aim of this review is the better understanding of the types of situation which lead to crush, and consequently the difficulties in predicting the occurrence of crush conditions in some situations. Further analysis involves the application of the contributing factors to the investigation of these events.

3. **To identify the most suitable methodology for simulating crush conditions during an evacuation.**

The accuracy and integrity of crush simulation will depend greatly on the specific methodological choices made when designing a simulation environment. To ensure the correct choices are made during this project certain different methodologies are investigated by means of a literature review, the results of which will inform our choice of testing environment in subsequent work.

4. **To develop a technique to identify the presence of crush conditions.**

The initial aim of this project is to define a technique which can be used to signal the presence of crush during an *in silico* evacuation by means of *passive* analysis, i.e. without the explicit calculation of the level of force present within an evacuating population.

5. **To test our candidate technique on a simple simulation of crowd behaviour.**

As a *proof of concept* measure, the analytical technique is tested on a simplified evacuation, so that its basic operation can be confirmed, that is, to detect changes in crowd behaviour which can be used as an identifier of the presence of crush conditions during an evacuation.

6. **To validate the working of our technique by using it to identify dangerous conditions within the recreation of an historical emergency.**

To show the usefulness of our technique for the measurement of crush risk within a simulation, an analysis is carried out on a well-known, and previously investigated, disaster in which crush conditions played a role.

1.4 Thesis Outline

The intention of this thesis is to lead someone with a basic knowledge of either evacuation or agent based systems through the process and application of our analytic technique to *in silico* simulations. This thesis is also intended to be a self-contained document that requires no outside information to enable the understanding of its central concepts.

The structure of this thesis is as follows.

Chapter 2 introduces the field of evacuation studies, and defines the *real-world* problems with which this work is concerned. Evacuation and crush conditions are defined, and we identify and define five key factors which can be shown to contribute to the formation of crush conditions. A number of historical examples where the presence of crush conditions lead to serious injury or loss of life are investigated, and analysed using the five factors identified previously. We conclude with an interesting issue relating to the post-mortem diagnosis of crush deaths, which suggests that the number of deaths attributed to crush conditions may often be under-estimated. In Chapter 3 we establish the *state-of-the-art* with respect to computational evacuation models, and investigate the three main modelling techniques relevant to this work. The difference between strict *movement models* and the more comprehensive *behaviour models* are also discussed. We carry out an investigation into current trends in both evacuation modelling as a field, and the popularity of specific computational models over the past twenty years. Previous methodologies for detecting crush are investigated, and the need for a new method of crush detection is identified. The chapter concludes with a brief overview of our proposed technique for identifying the onset of crush conditions within evacuation scenarios. Chapter 4 defines the model which we have chosen as the initial *test bed* for our chosen technique, the

social forces model, and covers the underlying concepts and mathematical equations by which the model operates. A critical review of the strengths and weakness of the model is also carried out. In Chapter 5 we investigate the mathematical techniques which we employ in our crush detection technique, which originated in the field of *Information Theory*, and define the techniques and methods which were employed in later work. A review of some of the more relevant applications of these techniques is carried out, and the chapter concludes with a more in-depth discussion on the way that we apply these techniques to the task of crush detection. The initial *proof of concept* is detailed in Chapter 7, where we show that the application of information theory techniques to an *in silico* evacuation simulation *can detect* changing levels of force within the evacuating crowd. A statistical analysis is carried out on the data obtained from these experiments and the results show that there is strong evidence to support our hypothesis. In Chapter 8 we recreate the evacuation of the Station Nightclub, a well-documented example of an evacuation in which crush conditions caused serious injury and loss of life. Our results confirm the findings of the official investigation, and our technique is shown to adequately distinguish between a *safe* evacuation and an evacuation in which the population is put at considerable risk. The results of these experiments are shown to be highly statistically significant, and we conclude that our technique does have the ability to metricise the relative level of threat present in an *in silico* evacuation. Tests are also run to ensure that the technique can distinguish between normal pedestrian movement and evacuation behaviour. Chapter 9 gives a summary of our findings, and discusses the possibilities for this work going forward.

Chapter 2

Evacuation and Crush

2.1 Introduction

This Chapter introduces the field of evacuation studies, and defines the real-world problems with which this work is concerned. We begin by defining what evacuation is, and also what evacuation is not; the latter includes the identification of three well-established fallacies of evacuation behaviour which are so commonly found within the literature that they require negation. We define and investigate crush conditions, and identify factors that lead to the formation of these dangerous conditions. Historical examples of crush disasters are reviewed. We conclude with an interesting issue relating to the post-mortem diagnosis of crush deaths, which suggests that the number of deaths attributed to crush conditions may be continually under-estimated.

2.2 What is Evacuation?

Evacuation, with respect to the movement of a person or persons, is the act of evacuating a person or a place. Evacuate is defined as:

Evacuate, *verb, trans.*

[To] Remove (inhabitants, inmates, troops), esp. to a place of safety from a place that has become dangerous. [131]

Dangers which precipitate an evacuation are often caused by fire or toxic materials, but could equally be related to a natural disaster or an impending

conflict. This definition of an evacuation is far too broad for our purposes, encompassing both the evacuation of towns and cities, and the evacuation of smaller structures with limited spatial constraints.

For the purposes of this work we therefore define evacuation as;

Evacuation

A process caused by the requirement, or perceived requirement, of a large number of people to effect egress from, or ingress to, a structure, under strict temporal conditions.

There are numerous distinctions between these two definitions of evacuation, the first being the presence of a *hazardous event or threat of such an event*, which we do not consider a requirement. In many cases, large scale crowd disasters can be found to have no precipitating factor that may be described as “hazardous” (at least in any traditional sense), and are often found to be the result of various other factors. Examples of crowd disasters such as “The Who Concert Stampede” [64], the Hillsborough disaster (see Section 2.7.1, or the Hajj pilgrimage [52], involved no actual or perceived dangers that were not a result of human factors. Factors contributing to the breakdown of an evacuation, such as spatial constraints and perceptual issues, are discussed later, alongside a description of the Hillsborough Stadium disaster (see Section 2.7.1).

Secondly, the requirement of *strict temporal conditions* is present in our redefinition of evacuation. For the purposes of this work we do not consider the exiting of a building as an evacuation behaviour, and therefore require there to be either a time limit to be placed upon the population, or for the population to believe that a time limit has been placed upon them.

Thirdly, we have the caveat that evacuation must consider *a large number of people*. This part of the definition is specific to this project, as whilst a single individual leaving a house due to a fire is technically an evacuation, it is not in any way useful to this work. The term *large* is necessarily relative, a large number of people in a residential house would be different to a large number of people in a sports stadium. As a guideline, we do not consider any evacuation during this thesis which involves any less than 200 people.

Finally, we have made the addition of *egress from, or ingress to, a structure* to our definition, to ensure that it is both distinct from the case of the evacuation of a town or city, and also to include both egress and ingress

within our definition. As is shown later (see Section 2.7), some examples of large scale crowd disasters exist in which the population concerned were attempting to enter a structure, rather than to exit, and we will consider these to be valid cases of evacuation for the purposes of this work.

This is this definition of evacuation that we will be using for the entirety of this thesis. It should be noted that we are not suggesting that this definition of evacuation is suitable to describe *all* situations in which it is necessary to effect egress from or ingress to an area, or that other definitions of evacuation are in some way lacking, this is simply the most suitable definition for the situations in which we are interested.

2.3 The Behaviour of Evacuating Crowds

The behaviour of evacuating crowds has been of interest to researchers for many years, with the first theories on crowd behaviours appearing in print during the late nineteenth century [83, 121, 134]. This area of research has been active ever since, with sociologists, psychologists, and mathematicians all adding to knowledge of the behaviour of crowds. The history and intricacies of crowd or group psychology are outside the boundaries of this work, but we briefly address some basic misconceptions of crowd psychology that may be relevant to the understanding of this thesis.

It is commonly thought that the overriding behaviour exhibited during an emergency evacuation is that of *panic* [2], i.e. illogical, irrational, or *crazed* behaviour, brought on by a rise in adrenaline caused by the precipitating disaster, but this is far from the case. The proliferation of this point of view can be seen in the language used to describe many crowd disasters, such as *stampede*, *frenzy*, or simply *panic* [114], and the point of view itself can often be seen repeated in sociology or psychology texts. In fact, the majority of crowd behaviours observable during evacuation scenarios have been shown to not only be entirely rational, but are the most logical survival strategy given the circumstances.

2.3.1 Fallacies of Crowd Behaviour

There are many commonly-held *fallacies* surrounding the behaviour of crowds or large groups of people. As these fallacies might influence the development and interpretation of the work reported here, we now briefly address them.

Panic or Irrationality

The most widely-held, and incorrect, belief concerning the behaviour of crowds is the phenomenon commonly known as *panic* [2], which for our purposes is defined as either a state of mind which may lead people to (inadvertently) cause injury to themselves or others whilst effecting their egress, or a highly emotional and irrational state, in which an individual's actions are detrimental to themselves or their position. Indeed this second definition is the more popularly held, and is regularly used by both the media and academics to describe mass evacuations [114].

All current research suggests that the idea of panic or irrationality in evacuating crowds is actually exceedingly *rare*, despite the widespread use of the term. If irrationality were the norm in a crowd, the entire idea of evacuation modelling would become defunct, as each member of an evacuating crowd would always exhibit entirely random and unstructured behaviour. In reality, panic type behaviours have not been shown to be the overriding behaviour, even in the most extreme of emergency situations. An investigation into the World Trade Centre disaster [100] discovered that panic behaviour extremely rare

... classic panic action or people behaving in an irrational manner was noted in $\frac{1}{124}^{th}$ (0.8%) [of] cases.

Blake *et al* (2004) [10]

We might ask, "If panic does not exist, then how do people come to harm during evacuations?". Research shows that under most circumstances the crowd is unaware that their actions are causing harm, or the scarcity (or *perceived* scarcity) of resources is such that individuals begin to compete with one another for access to the resources. This behaviour is known as *competitive egress* and, unlike panic, has been shown to be both rational and beneficial to the individual.

An example of this is found, as an analogy, in the Prisoners' Dilemma as stated by Brown [13]. The Prisoners' Dilemma frames a well-known problem of game theory relating to the reward structure of cooperation and competition in limited pay-off games. The game begins with two people being arrested for the same crime, and interrogated separately. Each prisoner has

		A	
		Cooperate	Compete
B	Cooperate	A -2, B -2	A -10, B 0
	Compete	A 0, B -10	A -5, B -5

Table 2.1: Prisoners’ Dilemma payoff matrix

the option of either refusing to talk, or blaming his adversary in the cell next door, and neither prisoner has any information about the decision that the other may make. If both prisoners refused to talk, they will both be sentenced to a short time in prison, but if both blamed the other then this sentence will be much greater. In the eventuality that one prisoner were to refuse to talk and the other was willing to lay blame upon him, then the sentence for the silent man will be great, whereas the man who cooperated with the police will be set free. These options, with associated sentences, are detailed in the pay-off matrix shown in Table 2.1.

Brown theorised that this situation is wholly analogous to the problem of evacuation where, in a cooperative crowd, any individual can increase their chances of escape, and minimise their own evacuation time, by competing. No matter how large an evacuation, if people are cooperating it is *always* possible to improve one’s current situation by making the decision to compete. This analysis of the evacuation problem is widely accepted, and there are many studies into *in silico* evacuation which treat the evacuation scenario as a game theoretic problem [7, 13, 31].

This simple analogy shows that, under any circumstances, an individual taking part in an evacuation can increase their chances of escape by competing with those around them. In a situation where ample time is given for evacuation, and no additional stresses are placed upon the population (i.e. no visible fire or smoke, little perceived threat of structural collapse) perceived benefit of competition can appear quite low, but as time frames are shortened or the perceived level of threat rises, the benefits of competition can begin to outweigh that of cooperation.

We see then an example of what is termed “non-adaptive” group behaviour (or panic behaviour), that may initially seem irrational, but which can be shown to be the best option available to the individual.

Homogeneity

This belief can be traced back throughout the field of crowd psychology, right back to the earliest works of Freud [36], McDougall [93], Sighele [121] and Le Bon [83]. The idea of the crowd becoming *homogeneous*, either in thought or in action, is less often perpetuated than the myth of panic, yet it is still present in many modern sociology texts [118].

The idea that a crowd may be thought of as *unanimous* in thought and action has been refuted by many, including Turner and Killian, who debunk the idea as “the illusion of unanimity” [144], and offer the theory that crowd thought and action is governed by *differential expression* in place of a supposed state of homogeneity brought about by the mere fact that one has become part of a crowd.

This is not to say that members of a crowd do not take *cues* from those around them, or are unlikely to imitate other members of the crowd, but it is generally believed that individuals within a crowd continue to exhibit *individual* thought and action.

Anonymity

The idea of the anonymity in a crowd is a long held belief of crowd psychology, and has been used to explain the “tendency” of crowds toward violence [36, 83, 93, 103, 135]. The theory is that the anonymity felt by members of the crowd, due to their number, allows the crowd to act without the usual fear of accountability, and possible retribution, that they may experience if acting as individuals.

This sense of anonymity has often been associated with the formation of dangerously competitive evacuations, the reasoning being that the removal of any consequences for one's actions allow people to compete for evacuation capacity in a more aggressive manner, which can lead directly to the formation of crush conditions. This has rarely been found to be the case, and in fact in most situations in which serious injury has been caused in dense crowd situations, the devastation caused by people's actions are not known to them. A good example of this is the Hillsborough disaster (see Section 2.7.1), where the people entering were completely unaware of the crush forming within the stands.

Recent research into communication within crowds also supports this

argument, showing that under high-density situations the communication between individuals at the front of the crowd and those at the rear (known as “front to back communication” [104]) is limited to very short distances [54]. This evidence shows that in many situations where individuals are compounding the problems suffered by a different part of the crowd, it is highly likely that the consequences of their actions are completely unknown to them at the time.

2.4 What are Crush Conditions?

An informal definition of *crush conditions*, with reference to evacuation scenarios, can be considered as the point at which the build up of physical force within a crowd of people reaches levels likely to cause serious injuries or death. Yet the simplicity of this definition can serve to obscure the true complexity of the phenomenon.

The majority of deaths or injuries caused by crush conditions are attributed to *compressive asphyxia* (or simply asphyxia [50]), commonly known as chest compression, which is the posterior compression of the torso to the point at which it effects breathing, leading to hypoxia [50] (partial deprivation of oxygen to the body), anoxia [50] (the extremity of hypoxia), and hypercapnia [50] (raised levels of carbon dioxide in the blood). The term *traumatic asphyxia* [2] is also regularly used to describe this phenomenon, but in some areas is reserved for specific uses, e.g. much of the evacuation literature uses these terms interchangeably, but it can often be found that the term *traumatic asphyxia* is reserved for the use of describing cases of asphyxia resulting solely from sudden and severe trauma to the upper body. For the purposes of this work we make no distinction between compressive and traumatic asphyxia caused by crush conditions.

The term *crush conditions* should not be confused with the medical conditions known as *crush syndrome* (also known as Bywater’s syndrome or Rhabdomyolysis) [2, 8, 50] which is a condition caused by an extreme pressure being placed and held on human tissue, and then subsequently released. The condition is common in situations such as earthquakes and structural collapses, which is why the appearance of crush syndrome is common across much of the evacuation literature, and therefore the reason that this important distinction has to be made.

2.5 What Causes Crush?

Having discussed the medical factors relating to crush injuries, it is useful to define the *human factors* which lead to the build up of sufficient levels of force within an evacuation to cause injury or death. Whilst no one suggests that the formation of crush conditions can be reliably defined in all circumstances, there are certain factors that can be shown to contribute to the likelihood of the formation of crush in emergency situations. We classify the main factors that lead to the initial formation of crush conditions under the broad headings of spatial, temporal, perceptual, procedural, and cognitive components.

2.5.1 Spatial

The spatial components of crush conditions are the simplest to quantify. They relate to the ratio of space available for egress to the number of persons that are expected to use the escape routes. It is obvious that if the density of a crowd does not reach a critical level, then the formation of crush is an impossibility, but density measured at two different points within the same crowd can vary greatly (i.e. the density distribution in crowds is rarely uniform).

Fruin devised a general metric with which to classify different degrees of crowd density. He termed these the *levels of service* [37], and highlighted the level at which the population density has the potential to facilitate the formation of crush as “Level of Service F” (see Table 2.2), which is the density at which a single individual has, on average, less than $0.46m^2$ of space available to them. It has never been suggested that the level of service F will *definitely* lead to the formation of crush, but Fruin suggests that if an emergency situation were to occur and an evacuation were to reach this density, then it is likely that injury will be caused to the evacuating population.

It should also be noted that whilst these levels of service are widely regarded as a reliable metric to describe immediate spatial concerns of an evacuating population, there are other criteria which are used in different circumstances. An example of this is the guidelines of the International Maritime Organisation (IMO), who consider an evacuation to be unsafe if, for 10% of the overall evacuation time, the density of the evacuating popu-

Level of Service	Walkway (m^2p^{-1})	Stairs (m^2p^{-1})	Queue (m^2p^{-1})
A	$s > 3.42$	$s > 1.85$	$s > 1.21$
B	$3.24 \leq s < 2.32$	$1.85 \leq s < 1.39$	$1.21 \leq s < 0.93$
C	$2.32 \leq s < 1.39$	$1.85 \leq s < 0.93$	$0.93 \leq s < 0.65$
D	$1.39 \leq s < 0.93$	$0.93 \leq s < 0.65$	$0.65 \leq s < 0.28$
E	$0.93 \leq s < 0.46$	$0.65 \leq s < 0.37$	$0.28 \leq s < 0.19$
F	$s \leq 0.46$	$s \leq 0.37$	$s \leq 0.19$

Table 2.2: Pedestrian Levels of Service (LoS). Available space (s) is measured in m^2 per person [37].

lation reaches 4 persons per square metre [58]. This is due to the fact that, even at relatively low levels of force, prolonged exposure to “light” crush conditions can still cause serious injury or death, and when dealing with an environment in which physical space is already highly constrained (e.g. a seagoing vessel) additional precautions must be taken to avoid overcrowding during evacuation scenarios.

It may seem logical that the main factors determining the probability of crush conditions forming are spatial, but *this is not strictly the case*. There are numerous counter examples, situations in which there are known to be very high crowd densities, yet have regularly shown to present relatively low risk to the population. These situations will be familiar to those who used to attend sporting events in the U.K., previous to the introduction of laws which prohibit the provision of standing tickets, where the crowd densities regularly reached such levels that members of the crowd were physically lifted off their feet by the force of the crowd. Crowd density in these situations, or in other large scale events such as concerts, rallies and religious festivals can often reach levels that are considered to be highly dangerous, yet they are held on a regular basis without incident. This suggests that spatial considerations *alone* cannot be used as a measure of danger present in all situations, there must be further factors which can cause a situation to transition from *high density* conditions to *dangerous* conditions.

2.5.2 Temporal

Temporal factors of egress vary, and depend heavily upon the rate at which conditions change. The traditional metrics used to evaluate the time that an evacuation will take are the RSET (Required Safe Egress Time); defined

as the elapsed time between the initialisation of an evacuation and the final evacuee reaching safety [122], i.e. the time required for a complete evacuation under ideal circumstances, and the ASET (Available Safe Egress Time), defined as the total time *available* for evacuation [122].

ASET and RSET calculations are strictly defined measures. The ASET metric is defined as the amount of time between ignition to the moment at which the conditions within the structure become so severe that further evacuation becomes an impossibility. The RSET metric is defined as the amount of time between ignition and the time at which the last occupant has exited the structure and is in no further danger. In addition to this amount of time the RSET metric must also include a “safety margin”, an additional amount (or proportion) of time included in the metric to ensure the safety of the population. The RSET metric can be further subdivided into more fine-grained time measurements. As an example of how this calculation could be further fine-grained, we could divide the RSET into the amount of time between ignition to detection (how long it takes to discover the fire), the amount of time between detection to raising the alarm, the amount of time between raising the alarm and the beginning of the evacuation (known as “pre-movement time”) and the amount of time between the end of pre-movement and the last evacuee reaching safety. As can be seen from this subdivision of the RSET calculation there may be specific time delays in each one of these measurements, and engineers err on the side of caution when working with these calculations, always considering the worst-case scenario.

Traditionally, the RSET and ASET metrics are used to determine whether or not the occupants of a building are able to evacuate safely under specific conditions. Generally, a structure can be considered “safe” if the ASET value exceeds that of the RSET by an *acceptable* margin, i.e. there is sufficiently more time available for an evacuation than would be required. The margin by which the ASET value should exceed the RSET value is usually decided on a *structure-by-structure* basis, and will vary greatly depending on factors such as building size, capacity, occupant familiarity, etc.

The rate at which conditions change can compound time constraints. As events unfold during an evacuation, the perceived time-scales within which occupants believe that they must escape will change dynamically, i.e. conditions which change the available escape time, such as exits becoming unusable, visibility becoming reduced, and the evacuation capacity

exit structures becoming overwhelmed, all lead to immediate changes in the evacuees “perceived available escape time”. The Rhode Island Nightclub fire (see Section 2.7.2), is a good example of this, and shows how the rapidity with which an incident escalates can place severe additional and novel time constraints on the evacuating population. Whilst modern evacuation analysis does allow for the simulation of an unlimited number of different evacuation scenarios from a single structure, real world constraints make the testing of all of these a physical impossibility, a prioritisation process which selects an adequate subset of these events is therefore employed.

2.5.3 Perceptual and Cognitive Factors

Perceptual and cognitive factors that lead to the formation of crush conditions are intrinsically linked, as an individual must rely on their perception of events in order to decide upon a course of action. The individuals’ perceived level of threat plays a large part in this, as it has the most direct effect on the decision making process [27], as this can lead people to instigate sudden movement from within a crowded situation, producing rapid changes in local densities. An additional side-effect caused by perceptual factors is the frustration caused to pedestrians when their objectives are continually not met [142], this can increase the probability of competitive egress behaviour within a crowd. Whilst the perception of threat plays a great part in the decision making process, one might assume it provides a good indicator of outcome, i.e. the greater the perception of threat the more likely that an individual would exhibit non-adaptive evacuation behaviour. However in reality the relationship between *perception* and *cognition* is highly complex, and can result in individuals displaying a wide range of behaviour, from the altruistic at one end of the scale, through to highly competitive egress behaviour.

The perception of information also plays a key part in the formation of crush. During emergency situations, it is often found that information relating to the current conditions is slow to propagate throughout a crowd of people, for example evacuees that are placed further back in a crowd may not necessarily be aware of the conditions further ahead [105]. This has been found in many situations, such as the Hillsborough disaster (see Section 2.7.1), where the people attempting to enter a structure were unaware of the already dangerously overcrowded conditions that existed inside, and

were effectively exacerbating the situation. In these cases the persons at the rear of a crowd can compound the situation by producing additional force that will propagate forward through the crowd, and also by limiting the extent to which the pressure might be alleviated, by inadvertently and unknowingly blocking the most immediate exit routes.

Work has been done on the propagation of information throughout a crowd of people, and has found that the process of passing information from the front of a crowd to the back is extremely ineffective [54].

2.5.4 Procedural

The procedural components of crush centre around the procedural failings of crowd management. As is shown during the investigations into historical events (see Section 2.7), many crowd disasters have serious procedural factors which contribute to their severity. Failures to plan effective evacuation strategies, to provide enough persons in a crowd management capacity, and to meet the local requirements for building and event management can all be seen to play a role in either causing or compounding crowd disasters.

Another procedural factor found to be commonplace is the inability of evacuees to follow strict evacuation plans in emergency situations, either due to the perceived level of threat being too great, or due to the conditions during the evacuation making following evacuation procedures impossible. This is not to say that the individuals involved are at fault, but that in extreme situations decision making processes are shortened, and therefore may not represent an optimal evacuation strategy.

This type of problem is extremely common in public buildings, where a great number of the occupants will be unfamiliar with the structure and have little, or no, knowledge of the evacuation plans, e.g. hospitals, town halls, museums, stadia, etc. When an evacuation takes place under these circumstances, the crowd often leaves by the most familiar route, generally the route by which they entered, even though there may be exits in closer proximity. An example of this type of behaviour can be found in the Rhode Island Nightclub incident (see Section 2.7.2), where the majority of the crowd converged at just one point of escape, the main entrance to the building, even though there were numerous other exits available.

2.5.5 Structural

Structural factors effecting crush may initially seem to be an extension of the spatial constraints, but there are distinct differences between the two categories. Structural design decisions can have a great effect on the safe evacuation of a structure; simple design decisions such as spacing emergency exits a sufficient distance apart can avoid the build up of a large crowd at a single exit, and aid the safe evacuation of the entire population. Examples of good design practice can be found in much of the engineering and design literature, and a graphic example of a small structural change which can make a large difference in the evacuation of a structure may be seen in Section 2.7.2.

2.6 Types of Force in Evacuation Scenarios

With respect to evacuation and crowd dynamics, there are three distinct types of force, or more properly *force propagation*, which are commonly seen, we will term these as *pushing*, *leaning*, and *stacking*. These types of force can, and do, lead to the formation of dangerous levels of force, but they are distinct in their physical mechanics and therefore in the modelling methods used to recreate these dynamics computationally. There follows a discussion on the specifics of these forces, with examples to models, or methods, by which they are investigated.

2.6.1 Pushing

Pushing force is the most commonly modelled of forces, and is the force with which we primarily address in this thesis. The term pushing force, as used here, defines the specific situation in which force is exerted by one individual upon the body of another individual whilst both individuals are standing. The criteria that both individuals are standing is most important in this situation, as it allows the modelling of this force in two-dimensions, e.g. on a Cartesian plane.

In large-scale crowd scenarios, the propagation of pushing force throughout the crowd has a summative effect, which can lead to members of a crowd being subject to high levels of physical force. In the most extreme of cases the propagation of pushing force alone can cause serious physical injury due

to compressive asphyxia, but it is often the side-effects of the pushing force building up with a crowd which causes the most serious issues.

There are issues surrounding the classification of injuries caused by pushing force (see Section 2.8), and commonly a large number of crowd tragedies are incorrectly attributed to trample injuries, especially by the media. In the words of John Fruin “Virtually all crowd deaths are due to compressive asphyxia and not the “trampling” reported by the news media. Evidence of bent steel railings after several fatal crowd incidents show that forces of more than 4500N (1,000lbs) occurred” [38].

The presence of high levels of pushing force has, in itself, shown to be a serious danger to a crowd, but there are additional *side-effects* of this type of force which present additional dangers. The presence of pushing force within an evacuating crowd has been repeatedly shown to increase the evacuation time of the crowd [65], a phenomenon known as the *faster-is-slower* effect [51]. The cause of this increase in evacuation time is due to the friction which builds within a crowd when physical force is present. As friction builds between individuals within a crowd their overall movement speed decreases, which will effectively increase the RSET of that evacuation. It may be argued that the friction effects caused by the presence of pushing forces within a crowd are far more dangerous than the level of physical force itself, as an increase in RSET in situations where smoke or toxins are present in the environment leads to increased exposure, but this risk is difficult to quantify.

The modelling of pushing force, and its frictional component, is not too uncommon in evacuation simulations, with both the Helbing model (see Chapter 4) and the Fire Dynamics Simulator (see Chapter 8) incorporating the modelling of pushing force and friction into their evacuation simulations. A notable set of evacuation models which partially include these factors are CAFE models (Cellular Automata with Force Essentials) [65, 126], which ignore the propagation of force within a crowd, but include methods by which the friction between evacuees may be modelled (see Section 3.2.1).

2.6.2 Stacking

Stacking force may be defined as the force produced by one body on another body when the bodies are vertically stacked. Stacking force is common in many evacuation scenarios, particularly where the evacuation route requires

the traversal of stairs (due to the increased risk of trips and falls).

We have seen from Fruin's Level of Service metrics (see Section 2.5.1) that the evacuation capacity of a stairway is significantly lower than that of a walkway of comparable width, which can lead to bottlenecks forming at the top of flights of stairs. It has also been found that a natural hesitation as evacuees transfer from walkway to stairway can lead to these bottlenecks in relatively low-density evacuations [117].

The case of stacking is not to be confused with that of trampling, which is also often recorded in the literature. Trampling injuries occur when an individual loses their footing and falls to the ground whilst within a crowd of people, at high densities it can be difficult, or impossible, to get back to their feet. This will often lead to the individual being unintentionally or unavoidably stepped upon by other persons in the crowd. Figure 2.1 shows examples of both stacking (left) and a trampling (right). We see that in the event of stacking the distribution of force is more even across the body, whilst during a trample incident the pressure is focussed on one specific area. There are also commonly different injuries arising from these situations, with stacking regularly leading to compressive asphyxia as, though spread out over a wider area, the force applied is constant, and often prolonged, whereas in the case of trample injuries it is common to see sharp blunt injuries as force is concentrated on specific areas of the body.

In the most extreme of cases the stacking of bodies can cause mass injury, with both the level of force and the prolonged duration of application leading to large numbers of cases of compressive asphyxia. An example of this effect can be found in the Ibrox stadium disaster of 1971, in which the stacking of bodies occurred during the exit from a football stadium. The incident is thought to have been triggered by a person falling on a large open stairway [149], which caused a chain reaction of trips and falls which led to a large-scale stacking of bodies upon the stairs. In the words of John Fruin;

In the Ibrox Park soccer stadium incident, police reported that the pile of bodies was 3m (10ft) high. At this height, people on the bottom would experience chest pressures of 3600-4000N (800-900lbs), assuming half the weight of those above was concentrated in the upper body area.

Fruin (2004) [38]

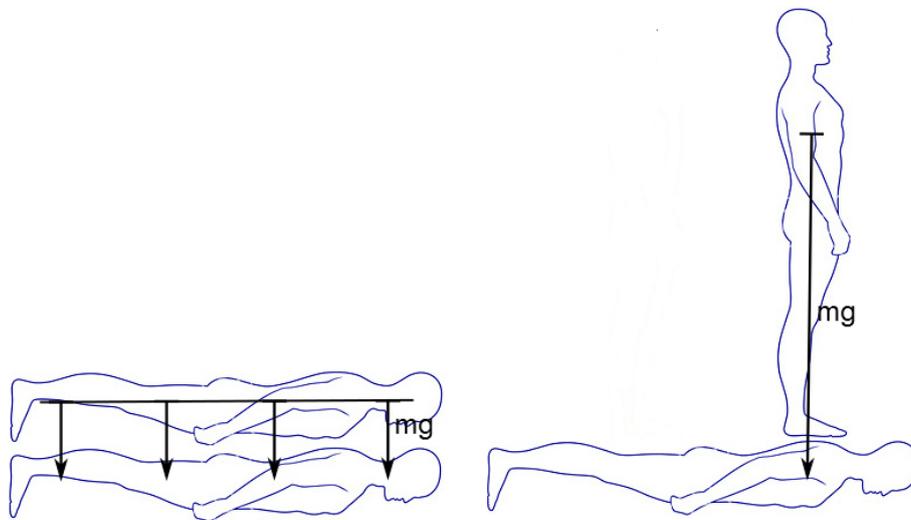


Figure 2.1: Left: A visual example of stacking, in which the force of the person above is distributed across the body of those below. This situation may well lead to compressive asphyxia, but is unlikely to result in diagnosable trample injuries as the weight is more evenly distributed about the body. Right: A visual example of trampling, in which the entire weight of the person on top is channelled down through the feet to small portion of the body underneath. This situation may well lead both trample injuries and compressive asphyxia, as the person above transmits force to a highly localised area.

The dual danger in this situation was not only the level of force, but also the prolonged period for which this force was applied, and is common in situations involving crowd stacking. The Ibrox stadium disaster of 1971 lead to the deaths of 66 people and the further serious injury of at least 145.

Many other highly publicised evacuations have lead to injuries and deaths due to the stacking of bodies after disastrous evacuations through stairwells, such as the Gothenberg Dancehall evacuation (see Section 2.7.3) , the Bethnall Green tube station disaster [28], and the e2 nightclub incident [154].

2.6.3 Leaning

The final type of force propagation which we will address is *leaning force*, which is the force propagating through a crowd which are not standing vertically, i.e. they are leaning either forwards or backward. This type of force is common in situations in which a high density crowd is standing on steps, stairs, or sloping ground, as the angle of the incline raises the probability of the crowd leaning or even falling forward.

This type of force is most commonly recorded in stadia and concert venues, therefore these are the areas in which it is most often investigated, but can be found in any situation in which densely packed crowds stand on uneven ground.

An important example of the disastrous effects of leaning forces can be found in the Hillsborough disaster (see Section 2.7.1), which saw an over-crowded terrace in a football stadium lead to the deaths of some 96 spectators. This example is particularly pertinent, and the investigation into the event saw one of the first instances of a mathematically defined *leaning crowd* model being applied to study the effects of these forces. During the aftermath of the Hillsborough disaster the British government launched an inquiry into the events, lead by Lord Justice Taylor [138]. The report included an investigation by Smith and Games [125], into the leaning forces which may have been present during the disaster itself, and is the first example of a leaning crowd model being applied in this manner.

The work by Smith and Games considered a single line of individuals spaced one step (i.e. stair) apart, leaning forward at an angle of θ , and aimed to calculate the supporting force which would be required to hold each one of these individuals in place. The supporting force required for the front most individual in the crowd was then, effectively, used as a metric by

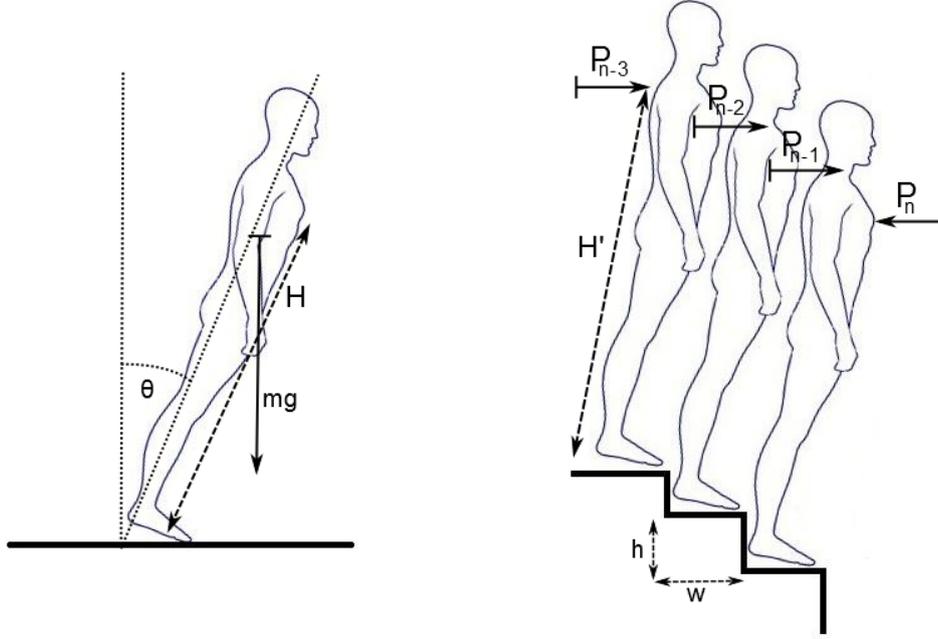


Figure 2.2: A visual example of the values used in Smith and Games[125] leaning crowd model. The figure on the right displays the lean angle, θ , the centre of mass height, H , and the gravitational force acting on the individual, mg . The figure on the right shows the push height, H' , and how the push force, $P_{0...N}$ propagates through a line of people.

which the supporting strength of crowd barriers could be calculated.

In the Smith and Games model, a single line of persons, each with mass m , are leaning forward at the angle θ on a set of steps of width w and height h , and they are assumed to be touching. The values of lean and mass are combined with both the *push height*, H' , (i.e. the height at which a person will push the individual directly in front of them) and the centre of mass height, H , to provide a geometric progression that describes the force required to hold the person at the front of the line in place, P_n . A visual example of the measures and values used in this equation may be seen in Figure 2.2.

$$P_n = \frac{mgH \sin \theta}{h} \left[\left(1 + \frac{h}{H' \cos \theta} \right)^n - 1 \right] \quad (2.1)$$

The result of Smith and Games work, with respect to the Hillsborough disaster, was the calculation of the force that was applied to a specific crowd barrier at the Hillsborough stadium, barrier 124a (see Section 2.7.1). Their calculations showed that the barrier would have been subjected to in excess of $8000Nm^{-1}$, that is $8000N$ per metre of the barrier length. Assuming a spectator width of $0.355m$, as in Dickie and Wanless [29], this would equate to an equal and opposite force of approximately $2840N$ acting on each individual pressed against the barrier.

Whilst commonly seen in stadia, due to the necessity of having inclined standing/seating areas, leaning force is found to occur during many other situations. The Mihong bridge disaster (see Section 2.7.5), during which the leaning force which built up on a bridge lead to high force propagation through a densely packed crowd, is another example the devastating effect of leaning forces.

2.7 Historical Examples of Crush

Here we present case studies of situations where the formation of crush conditions led to both serious injuries and fatalities. Each case study also represents some failure within a system (e.g. failure to limit the capacity of a structure to safe levels, failure to adhere to official guidelines or fire laws, failure to follow crowd control policies, etc). These types of failure are often observed in cases where the evacuation of a building leads to the death or injury of many people. Failures of this kind are common, and we believe that they should not only be expected, but also be considered during the design of buildings, the creation of evacuation plans, and especially during simulated evacuation exercises.

2.7.1 Hillsborough

The Hillsborough disaster [136] (Sheffield, UK), claimed the lives of 96 people and caused the hospitalisation of a further 300. The disaster happen at a football match between Liverpool and Nottingham forest, taking place at Hillsborough stadium, the home of Sheffield Wednesday, on the 15th April 1989. Due to the heightened public interest in the incident (the match was scheduled to be transmitted on English television), and also because of the multiple perceived failures on the part of the authorities, the Hillsborough

disaster has become one of the most thoroughly investigated crowd disasters in living memory.

The tragedy at Hillsborough stadium occurred when police stewarding the match made the decision to open an extra set of gates, intended as an exit, in order to relieve the extreme levels of congestion that were forming as the crowds tried to enter the stadium through the turnstiles at the Lepping's Lane end of the ground. These gates did not have turnstiles, and the result was an influx of up to 5,000 fans through the narrow corridor that lead into the standing terrace, see Figure 2.4. The sudden arrival of so many additional fans pushed the capacity of the central pens far above their legal maximum, and soon a dangerous crush formed at the front of the stands. Those fans still entering the stadium were unaware of this, and continued to attempt to enter the stand as the people inside were slowly crushed against the crowd barriers and fences at the front of the stands. The conditions at the front of the terrace became so bad that most of the 96 victims died from asphyxiation, or other crush related injuries, within five minutes of the game starting.

The maximum capacity of the stands at Hillsborough stadium were a source of great debate during the aftermath of this tragedy, and factored heavily in the technical investigation into the disaster. The initial calculation, made previous to the disaster, had suggested that the two central pens, pens 3 and 4, had a maximum capacity of 1,200 and 1,000 persons respectively, but investigations after the tragedy occurred resulted in much lower figures. If we centre on the capacity of pen 3 we see that the post-event investigation, in which calculations were made according to existing official guidelines, estimated the maximum capacity of pen 3 to be just 822 persons, a reduction of 378 people from the original capacity. According to the report, the recalculation was carried out to

...compensate for some departures from the recommendations of Chapter 9 of the [Guide to Safety at Sports Grounds]. We therefore calculated capacities for Pens 3 and 4 on the basis of the areas behind crush barriers and perimeter fences in which a crowd packing density of 5.4 persons/ m^2 would be permissible ... [20]

A separate strand of the investigation, carried out at the same time,

aimed to estimate the actual occupancy of these areas during the disaster. Using photographs taken during the event, estimates for the occupancy of pen 3 were calculated to be up to 1,576 [101], which would make the actual occupancy 180% of the recommended maximum. At this level the average crowd density across the entire pen would have been approximately $9.8pm^{-2}$. In actuality the empirical observations made from analysis of the video footage taken during the build-up of the crush suggested that the crowd density was far from uniform.

Figure 2.3 shows the suggested capacity of pen 3 (bottom), and the observed crowd density across the rows of pen 3 during the disaster (top). The suggested capacity was calculated as the number of people that could stand behind each crowd barrier, assuming a maximum density of 5.4 persons per square metre [20]. The actual occupancy, and therefore crowd density, was estimated from photographs taken during the disaster, and was calculated on a row-by-row basis.

The pathologists' reports into the deaths at Hillsborough [138] found all but 9 of the victims to have died from the primary cause of compressive asphyxia. In six of these nine cases the victims were found to have injuries to the back or chest, one had suffered a ruptured aorta, and the remaining two had existing medical conditions which were considered to be major contributory factors. None of the pathologists' reports returned stated trample injuries as the primary cause of death, although there were 18 cases which had presented with injuries that may have been indicative of trample injury [138].

An investigation into the events at the Hillsborough stadium showed that the level of force present could have reached over $8000Nm^{-1}$, which was the force required to cause the damage observed to crash barriers at the ground. Figure 2.5 shows the damage caused to barrier 124A at the Lepping's Lane terrace of the Hillsborough stadium.

It should be stated that the security and crowd control arrangements during the Hillsborough stadium tragedy have been a source of controversy from the date of the event, and this controversy continues to the present day. Disregarding the political, societal, and security considerations surrounding the event, it is widely regarded that the security arrangements of the day compounded the conditions within the stadium.

Later investigations into the disaster, most notably the Taylor report

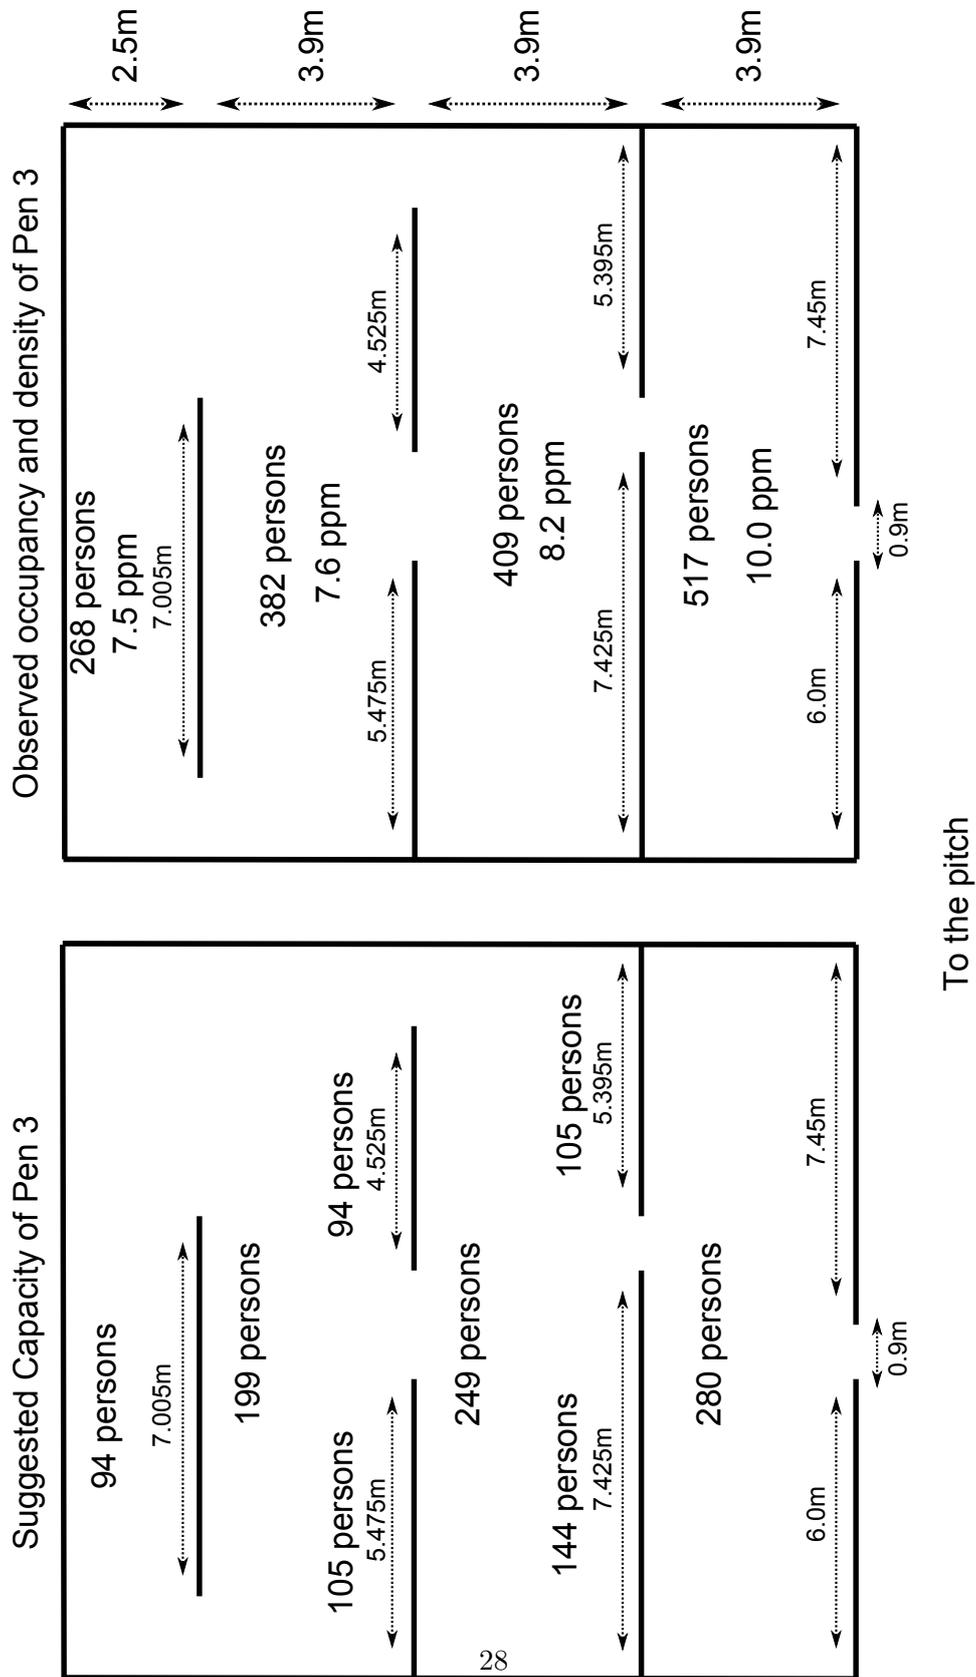


Figure 2.3: Bottom: Safe occupancy levels of Pen 3 at Hillsborough stadium. Top: Estimated actual occupancy and density of Pen 3 at the time of the disaster. All internal sections represent the crowd barriers in place at the time.

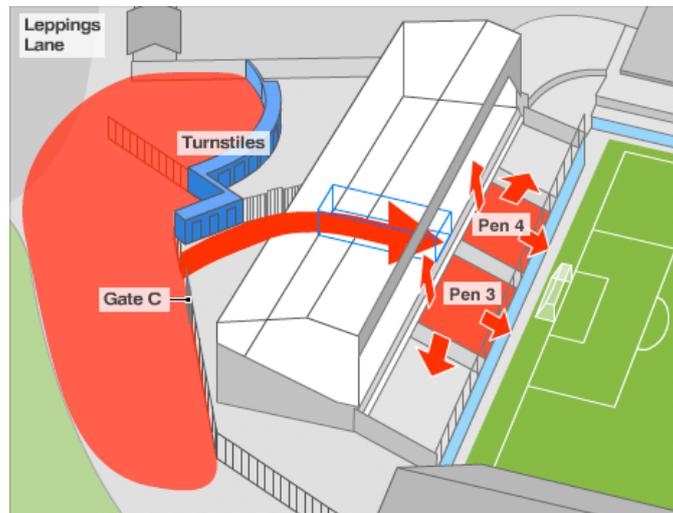


Figure 2.4: Diagram showing the route that the influx of supporters took into the ground, and the routes at which fans already in pens 3 and 4 tried to take to alleviate the crush inside the pens [5].



Figure 2.5: Barrier 124A at the Leppings Lane terrace of Hillsborough stadium after the crowd crush. Figures estimate the force required to cause this crush barrier to collapse was in excess of $8kNm^{-1}$, or $8000Nm^{-1}$ [99].

[138], have molded the security and policing of all subsequent sporting events within the U.K. and further afield.

Types of Force Present

During the Hillsborough disaster, the primary type of force present was a leaning force, a combination of pushing force with the additional effect of gravity pulling the crowd forward from their pivot points (see Section 2.6.3). All investigations into the Hillsborough disaster thus far have focussed on the propagation of force throughout a leaning crowd, with the initial investigation presenting the first documented example of a mathematical leaning force model.

The fact that leaning forces were present in situations such as football stadia was accepted long before the Hillsborough disaster, but true investigations of its effect had not previously been completed. The presence of crowd barriers, placed periodically throughout standing terraces, were meant to reduce the effects of force propagation by absorbing force propagating forward through the crowd at strategic points. The failure of this system, in this specific case, came from the overcrowding of terraces, as the crowd barriers in place were designed to support a crowd of a density of $5.4pm^{-2}$, which was breached quite severely in this case. This additional force was compounded by the failure of certain crowd barriers, which would have allowed propagation of more force throughout any person who occupied the space in front of the collapsed barrier.

Five Factor Analysis

Spatial

The overloading of pens three and four were obviously a direct cause of the Hillsborough disaster, as it was this extreme density that caused the deaths and injuries suffered. As has been stated previously, the occupancy of pen 3 at the time of the disaster was greater than 180% of the suggested safe maximum, and density levels were breaching the $10pm^{-2}$ mark. That said, the overloading of parts of the stadium was the effect of numerous other failings, rather than a failing in itself. The other factors that lead to the Hillsborough disaster are discussed below.

Temporal

The rapidly changing conditions during the incident were a major factor in the ensuing crowd crush. Late arrivals at the stadium caused high density conditions leading into the stadium, and the turnstiles at the entrance were overwhelmed. To ease the crowds conditions outside the stadium, gates that were intended as an exit were ordered to be opened, and a large number of additional fans entered the stadium through this route [137](p. 11). The late influx of fans (some estimate up to 5000) into the central pens of the Lepping's Lane Stand caused a sudden increase in the crowd density within these areas, which lead to massive compressive forces building up at the front of these stands.

Perceptual/Cognitive

The perceptual factors in this disaster are twofold. Firstly, the fans outside the stadium were not fully aware of the state of the game, and when cheers were heard from inside the stadium, the perceived need to effect ingress was raised significantly. These cheers were not caused by the start of the game, but instead were for the initial entrance of the teams onto the pitch; the game itself would have not started for at least 10 minutes so time was still available for supporters outside to enter through the designated turnstiles into the correct pens. Secondly, the tunnel through which the supporters eventually entered was entirely enclosed from the actual stands, so the terrible events unfolding on the terraces would not have been obvious to those still entering. It was found that once these fans had gained access to the stands many attempted to go back, but were not able to do so [137] (p. 13).

Procedural

There are many procedural failings that lead to the Hillsborough disaster, and they have been extensively documented as a result of numerous inquests into different aspects of the disaster. The police have been criticised for the opening of an exit which was not fitted with turnstiles, to alleviate the pressure building up outside the ground [137] (p. 40). The official report into the disaster, known as the Taylor Report [138], stated that the cause of the disaster was a failure of Police control. The start of the match should have been postponed to allow the influx of fans to be correctly distributed about the stadium, as it was the majority of fans were directed by existing signage into pens 3 and 4 only. Other failures have been noted in the handling of the aftermath of the event [137] (p. 44), such as access to the ground by the emergency services, but these are outside the scope of this work.

Structural

Previous to this event the general level of safety at Hillsborough Stadium was considered good, but criticisms were made of the relatively low number of turnstiles at the Lepping's Lane Stand, and of the poor condition of the crush barriers in place to stop the propagation of force [137] (p. 21-23). Failures of crush barriers allow the propagation of force throughout a much larger section of the population, and can compound already hazardous conditions. Figure 2.5 shows crush barrier 124A of the Lepping's Lane stand after the incident; it is estimated that the force required to cause this damage was up to $8kNm^{-1}$, which equates to approximately $3kN$ of force acting on each individual in direct contact with the barrier (see Section 2.6.3).

2.7.2 Rhode Island Nightclub

At the Station Nightclub, Rhode Island, on February 20th 2003, a fire during a rock concert caused 96 fatalities, alongside numerous other serious injuries [45]. The fire was started when the band's manager discharged pyrotechnics which ignited a large section of flammable polyurethane foam, which had been used to soundproof the drummers' alcove. The resulting dense, choking smoke quickly filled the club. The fire spread from the stage, igniting other portions of the ceiling and wooden structure of the building, and within five minutes of the initial ignition those outside the club observed flames breaking through a portion of the roof. Figure 2.6 shows the two ignition points from which the fire started, and contains clearly marked positions of all the exits to the club.

Official estimates of the occupancy for the club range between 400-450 patrons on the night of the fire, which is approximately the maximum capacity for a building of its dimensions. It was the sudden ignition and fast spread of the fire and smoke which caused such a high density evacuation with very short time-scales.

As is often found in cases in which the evacuating population is unfamiliar with the emergency exits of a structure, the distribution of evacuees across the possible exits was non-uniform. Despite the existence of four possible exits at the start of the fire, the majority of the crowd headed for the most familiar exit: the entrance to the club (it is estimated that up to two thirds of the patrons attempted to exit via the main door [45]). Investigations suggest the kitchen door was used by just 12 people throughout

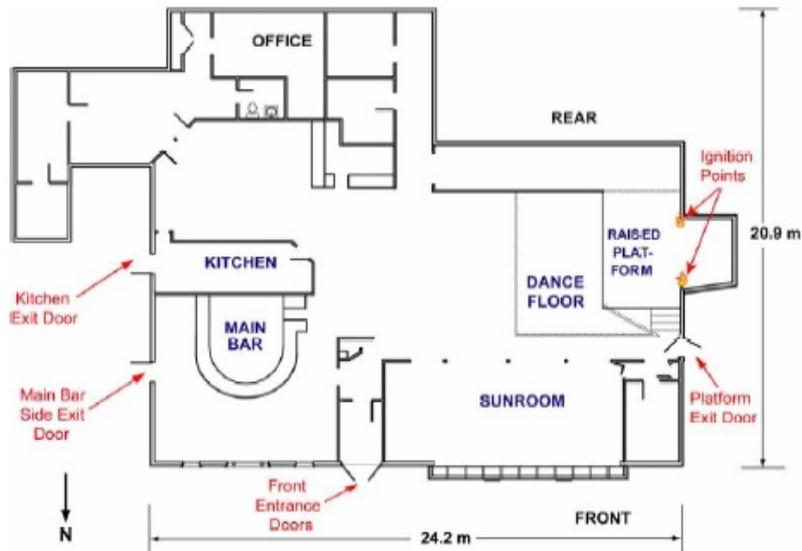


Figure 2.6: Floor plan of the Station Nightclub. Ignition points mark the area in which the pyrotechnics were discharged [45].

the evacuation, and that the stage door became impassable (due to the high levels of heat and smoke) after roughly 30 seconds from the ignition point. These dynamically changing conditions served to compound the already overcrowded situation at the main entrance to the club, which was soon overwhelmed, and people began to trip or fall during their escape.

Figure 2.7, compiled by the West Warwick Rhode Island Police Department, shows the location of the deceased victims in the aftermath of the Station Nightclub fire. We see that the majority of victims were found around the main entrance, with most of the rest of the victims being found at the rear of the building, where it was assumed they had retreated to in an attempt to find another exit. The seven victims marked as being found outside the club were deaths due to smoke inhalation after leaving the building.

The official time-line of the fire (compiled by NIST [45]), states that just 1 minute and 42 seconds after the start of the fire, there existed a “pile” of people in the entrance corridor, blocking the main escape route and making further egress via this route impossible.

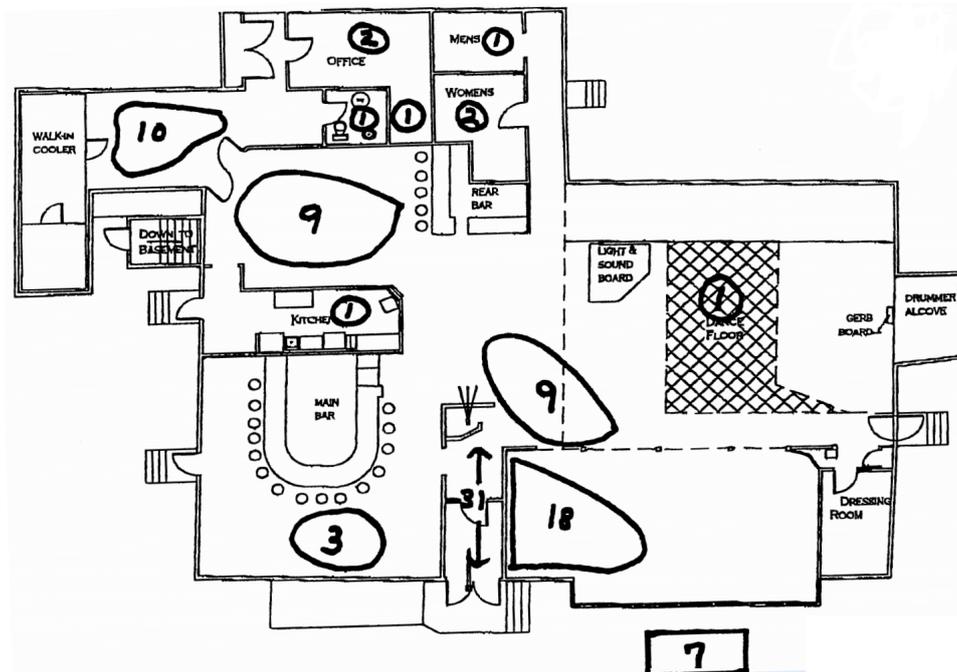


Figure 2.7: Diagram showing the location of the deceased in the aftermath of the Station Nightclub fire. This diagram was compiled by the West Warwick, Police Department (Rhode Island), and was released during the investigation into the incident.

There were many factors which contributed to the scale of the tragedy at the Station nightclub, but the two primary factors were the speed at which the fire spread (mainly due to the highly flammable nature of the ignition material) and the subsequent overloading of exit routes. The occupancy of structures of this type are calculated according to floor space and exit capacity, yet in the case of the Station nightclub the latter was severely affected during this incident. The main entrance of the Station nightclub was approximately two metres wide, yet within the entrance corridor was a opening of just one metre, which is an effective halving of the exit capacity of the venue. When this is considered alongside the under-utilisation of the kitchen exit, and the stage door becoming impassable within thirty seconds of ignition, we can see that the effective exit capacity of the venue during the fire was less than half that which would have been considered during the occupancy calculations.

The structural failure point of the Station Nightclub was arguably the configuration of the main entranceway. As we can see from Figures 2.6 and 2.7, the route from the inside of the club into the main entrance lobby had two openings, of approximately 1m in width each, with an exit at the end of the lobby of 2m in width. Between these two points there was an opening of just 1m in width, which acted as a bottleneck to those attempting to effect egress via this route. It is easy to imagine that were an evacuee to be attempting to exit via the door from the bar, that upon finding the main entrance inaccessible they could easily exit through the bar door at the other side of the room. Were an evacuee to be attempting to exit from the dancefloor or sunroom via the main entrance there would be no other visible exits available to them, so they would have no choice but to continue to effect egress via the main entrance, no matter how congested this route had become. In the case of the Station Nightclub numerous evacuees managed to escape by breaking the large front windows of the sunroom and using these as exits, reports suggest that up to 79 people may have escaped via the windows in the main bar and sunroom. Many of the people escaping via the windows did so aided by the emergency services [45].

Simulations of the fire that consumed The Station Nightclub showed that, at approximately 90 seconds after the initial ignition point, the temperature in the club during the evacuation reached levels of up to $1,000^{\circ}\text{C}$. To put this figure in perspective, the melting point of iron is roughly $1,375^{\circ}\text{C}$.

The Station Nightclub disaster contains archetypal examples of evacuees displaying both adaptive and non-adaptive evacuation behaviours in their attempts to effect egress. The crowding at the main exit, and the resulting crush, is an example of non-adaptive behaviour, as at finding a large mass of persons at the main exit and adaptive choice would have been to search for another exit route. This type of behaviour is both common and understandable in high stress situations, as a people under stress often exhibit highly shortened decision making processes, making them less likely to consider the possibility of other options. That said, there were individuals whom showed adaptive behaviours, by effecting egress through the large windows present in both the sun room and the main bar, some later aided by the police and fire services.

The Station Nightclub evacuation is investigated further in Chapter 8, where we also detail the results of our own simulations of the event.

Types of Force Present

The primary type of force present during the Station Nightclub disaster was pushing force, as the crowd overwhelmed the main entrance and a crowd crush occurred. As with many situations such as this, there is the inherent presence of stacking force also, as the severe levels of crowd density occurring at and about the main entrance increases the probability of trips and falls. In these situations it is not uncommon for the presence of high levels of pushing force to cause localised stacking behaviour within and around exit structures, as the presence of fallen evacuees further increases the probability of trips and falls until a large number of individuals to “stack”, and in some cases this will entirely block an exit route.

There is definite evidence of this type of *pushing-to-stacking* phenomenon occurring during the Station Nightclub evacuation, as eye witness reports had stated that within approximately 90 seconds of ignition there existed a pile of people within the entrance corridor [45].

Five Factor Analysis

Spatial

Accounts suggest that the structural capacity of the Station Nightclub, at the time of the evacuation, had *not* been exceeded (according to the Rhode

Island fire laws in place at the time) [45] (p. 6-25).

Temporal

An extremely fast moving fire caused severe temporal constraints to be placed on the evacuating population. Investigation suggested that, even under ideal circumstances (i.e. no blocking of the stage door, and an even distribution of evacuees across the other exits), the fire would still have caused significant loss of life and injury to the evacuating crowd [45].

Perceptual/Cognitive

We would not consider that the events at the Station Nightclub were compounded by any perceptual or cognitive factors concerning the evacuating crowd. The speed at which the fire spread necessitated an extremely fast evacuation, and the perception of risk was fully justified in this situation.

Procedural

Evacuation plans for buildings such as the Station Nightclub are notoriously difficult to implement, due to naïve populations (i.e. patrons who would not necessarily be familiar with the building, or the position of emergency exits). This said, better usage of the available exits *could* have distributed the evacuation more evenly across the structure and reduced egress time.

Structural

Disregarding the installation of sprinklers throughout the building, which was not a legal requirement for buildings of this type, it has been found that the minor structural change of removing the smaller interior opening present in the front entrance/exit corridor (i.e. the “lips” visible in the entrance corridor on Figure 2.6) [45] (p. 6-26), would have increased the evacuation capacity of that specific exit, from approximately one metre of evacuation space to two metres.

2.7.3 Gothenburg Dancehall

When fire broke out in a crowded dance hall in Gothenburg, Sweden, on October 28th 1998, it claimed the lives of 63 people and injured over 200 others [21]. The first floor venue in question was packed to nearly triple its 150 capacity, with officials estimating that there were over 400 people in attendance. Eye-witness accounts of the incident state that population density prior to the start of the fire was already at dangerously high levels, with a number of the occupants observing that there were so many people present that they were barely able to move [21].

Figure 2.8 shows the floor plan of the building (second floor only). The ignition point, in the East staircase, is clearly marked.

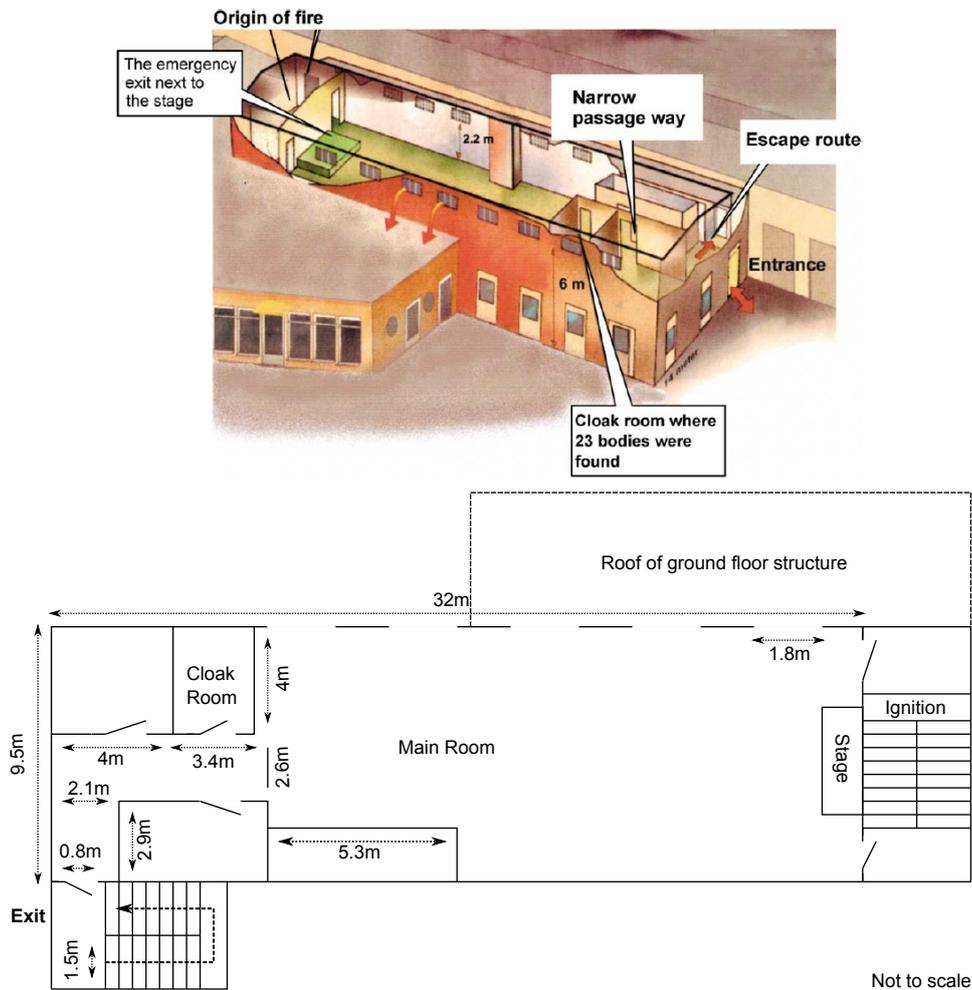


Figure 2.8: Top: Representation of the dancehall in Gothenburg [15]. Bottom: More detailed floorplan showing the first floor of the building, the stairwell at the south-west of the structure lead down to the only available exit. The similar stairwell at the east of the structure marks the ignition point of the fire, and was also blocked as it had been used to store furniture not in use during that night's festivities.

Shortly before midnight, the DJ playing at the event discovered that a stack of chairs stored in one of the only two stairways leading out of the first floor dance hall had been ignited. The fire had already reached an advanced stage, so he proceeded to inform the authorities and to make his

escape rather than to attempt to extinguish the blaze. Finding that the crowd density had reached such levels that crossing the dancefloor to make his escape was an impossibility, he decided to make his escape via a window and across the adjoining roof.

Announcements had been made to the occupants, but it is suggested that as little as 50% of the occupants made attempts to exit the building at this early stage [15]. Some survivors who had been at the far end of the hall when the fire was initially discovered stated that they smelled smoke but had believed it to be cigarette smoke and therefore felt no need to evacuate. Other eye-witnesses stated that they had gone towards the emergency exit to investigate the fire, and that the emergency exit door may have opened “multiple times” by persons wishing to take a look at the fire. On the last occasion the exit door was left open, as it was reported that the handle had become too hot to be touched, which allowed additional oxygen to the fire. A report into the incident states;

Shortly after this, lamps located close to the emergency exit began to explode and thick black smoke entered the dancing floor. People began running towards the main entrance door ... [15]

At this point the major portion of the evacuation began, and the main exit was quickly overwhelmed. Evacuees took to using the windows, situated approximately $2m$ above the floor, as an evacuation route, despite there being a drop of approximately $6m$ to the ground below.

Fire-fighters attending the scene attempted to enter the structure through the front stairwell, yet their progress was hindered by large numbers of severely injured people whom they had to remove from the stairwell before continuing upwards. The fire-fighters reported that upon reaching the top of the stairs they discovered a wall of bodies blocking the entire of the upper doorway. These were removed and shortly the evacuation continued, simultaneously other fire-fighters were trying to facilitate egress of evacuees through the first floor windows.

There follows a time-line of the events at the Gothenburg dancehall, any estimated times will be prefixed with “ \approx ”. The horizontal line through the event time-line marks the point at which conditions within the structure would have become untenable.

	Time	Since Ignition	Events
≈	11:30	0:00	Estimated time of ignition
≈	11:30 - 11:40	0:00 - 0:10	Fire discovered
	11:42	0:12	Emergency services informed
≈	11:42	0:12	Door to the stairway in which the fire started was left open
	11:44	0:14	Smoke spread throughout the main room, and evacuation began
	11:45	0:15	Emergency services dispatched from station just 2km away
	11:47	0:17	Additional units requested due to severity of fire
	11:49	0:19	Fire fighters report people jumping from the windows of the building
≈	11:57	0:27	Fire is fully developed and flames are visible in all windows of the hall

We can see from this time-line that the the elapsed time between ignition and the point at which conditions within the structure became untenable could have been as low as 17 minutes. This however would not represent the time available for evacuation, as it had been stated that the evacuation did not begin “in earnest” until approximately 14 minutes after the discovery of the fire. This leads to the estimation of the ASET for this evacuation as between 3 to 13 minutes. If we err on the side of caution, and suggest that the ASET for this evacuation was at the lower bounds of this figure, we have an ASET time of just 3 minutes. This represents the time between the start of actual evacuation (11:44) and the time at which the fire brigade dispatched additional units due to the severity of the situation (11:17). The RSET for this event can be estimated using the “Maximum specific flow rate” [98], a rate of flow used to calculate the egress capacity of exit structures. The flow rate of the 80mm door leading to the stairs is approximately 1.04 persons per second, whilst the stairway itself (150mm) has a flow rate of 1.41 person per second. Using the value of the smaller bottleneck, we can estimate the RSET for 400 people from this building would be approximately 6m 25s ($RSET = \frac{Occupancy}{Exit\ Capacity} = \frac{400}{1.04} \approx 385s$), almost twice the lower boundary of

the ASET calculation.

We can see from the ASET and RSET figures, calculated above, that the reduced exit capacity of the structure meant that the single exit had less than half the required evacuation capacity that would be required for the safe and timely evacuation of this structure within the given time-frame. These estimates do not take into account the extreme situations within the hall, which was reportedly filling with thick smoke at an alarming rate, so it could be argued that even these figures do not represent the true extent by which the RSET for this evacuation exceeded the ASET.

Types of Force Present

In the case of the Gothenberg Dancehall fire it is very difficult to ascertain which was the primary source of force. There was certainly a large element of stacking force present during this evacuation, as rescue crews arriving on the scene had to extricate fallen evacuees from a pile of bodies present in the only accessible stairwell to the structure. There must, however, also have been a element of pushing force present within the earlier stages of this evacuation, due to both the extreme number of evacuees present (compared to the maximum allowable capacity) and the size of the doorway which lead into the stairwell. As stated previously, the safe exit capacity of this upper doorway was just $1.04ps^{-1}$, which was less than half the capacity which would have been required to evacuate this number of people in less than the ASET time for this incident. It would seem that, considering these factors, the presence of a high level of pushing force through and around this doorway could not be ruled out.

Five Factor Analysis

Spatial

On the night of the fire the building was severely over capacity, with a total of approximately 400 persons occupying a venue which was legally allowed to be occupied by no more than 150 persons at any one time. Crowd density is difficult to calculate for this event, as the dancehall was separated into numerous rooms, but an estimate based on the size of the dancefloor suggests that an occupant may have had as little as $0.37m^2$ space in which to move *before* the evacuation had started. This is comparable to Fruin's "Level of

Service” F, and would not be considered a safe density at which to occupy a structure under even the most ideal of circumstances.

Temporal

In this case the temporal factors were, in our opinion, secondary. There was obviously insufficient time for the large number of patrons to safely evacuate the building, but this is due to other factors detailed during this analysis.

Perceptual/Cognitive

Knowledge of the fire was slow to propagate throughout the crowd, which slowed the initial evacuation down greatly. Interestingly in this case though is the apparent perception of danger amongst many members of the crowd, which was far less than was merited by the situation. As mentioned previously, survivors stated that many people had been unworried by the knowledge that there may have been a fire in the building, and some actively sought out the blaze rather than effecting egress.

Procedural

The procedural factors facilitating this incident are twofold. Firstly, one entire exit to the dancehall had been rendered unusable by the placing of a large number of chairs on the stairwell leading to it. In actuality the stairwell would not have been of use during the evacuation, as it was the origin of the fire, but it is a serious procedural failure that would effectively halve the evacuation capacity of this structure. Secondly, and more importantly, the extreme breach of the building’s maximum legal capacity (by a factor of over 260%) was a major factor in the formation of the crush conditions.

Structural

In this case, the structural factors appear to have had little effect on the overall incident. This can be partially verified by calculating the RSET for this event were the occupancy to within the suggested maximum for this venue, which was 150 persons. Using this figure we can estimate that the time taken to evacuate the venue via the one usable exit (exit capacity of 1.04 persons per second) would have been approximately 2 minutes 45 seconds ($RSET = \frac{Occupancy}{Exit\ Capacity} = \frac{150}{1.04} \approx 144s$), which falls 15 seconds under the lower bounds of the ASET for this event.

2.7.4 E2 Nightclub Incident

In Chicago’s E2 Nightclub on Feb 17th 2003, the security guards’ use of pepper spray, to intervene during an altercation, became the catalyst for an

evacuation that claimed the lives of 21 patrons [127, 154].

The evacuation began after a serious altercation between patrons caused the security staff to attempt to diffuse the situation by deploying pepper spray. Unfortunately, the use of pepper spray in such an enclosed space caused a sense of fear throughout the population as, to the patrons, this was interpreted as an unknown chemical irritant present in the atmosphere. The physical effects of the pepper spray on the surrounding crowd being significant meant that those close to the altercation began to rush toward the exit to escape the pepper spray, which by this point was already spreading around the club. As the initial wave of evacuees made their way through the club, those who had not witnessed the incident began to fear for their safety, especially as it became obvious that some form of chemical agent was present.

Within seconds the entire crowd, estimated at over 1,100 people (the club's capacity was only 240), rushed towards the main exit. The door to the street opened inwards, whilst the door leading to the dance floor opened outwards. As people rushed from the club, the upper door flew outwards, pushing those on the upper landing down the steep flight of stairs. As more people exited, they were forced on top of the fallen evacuees, and the bodies began to "stack up" and block the exit. It was the tremendous pressure placed upon the fallen evacuees that caused the 21 deaths during this incident.

Figure 2.9 shows an example of the type of stairway configuration which was present in the E2 Nightclub. The actual building plans for this structure are not publicly available, so exact measurements are not shown on this diagram. We can see that the stairway entrance on the first floor of the building has an outward opening door, which opens onto a landing leading to the stairs, whilst the exit to South Michigan Avenue has an inward opening door, which also opens into a landing at the bottom of the stairway. Eye witness reports state that as the crowd rushed to exit the first floor of the building the doors to the upper landing were thrown open, pushing those who were standing on them down into the stairwell [151]. Reports suggest that as the evacuation progressed this stacking behaviour on the upper landing and at the top of the stairs continued until any egress was impossible. During the court case testimony from the security personnel working that night suggested that the pile of bodies on the stairwell reached

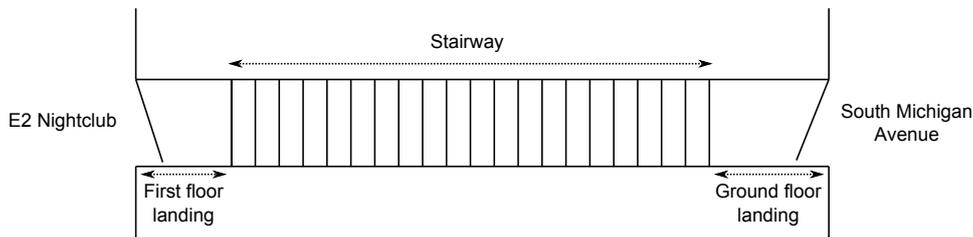


Figure 2.9: A representation of the stairway leading from the E2 Nightclub down to South Michigan Avenue. The first floor doorway (left) opens into a landing at the top of the stairway, whilst the ground floor door (right) opens into a similar landing at the bottom of the stairway.

over six feet in height [95], and that attempts to extricate people from this stack of bodies fast became unsuccessful [55].

This event has been included here for two very important reasons, the first being that this event did not involve what are traditionally thought of as an evacuation catalysts (e.g. fire, flood, earthquake, etc), which makes it an interesting case to investigate. The second reason is that, as there were no hazardous toxins in the environment (the pepper spray being classed as non-lethal, and unlikely to cause lasting injury), all deaths in this disaster were caused by the presence of force, either in the form of crush conditions or via trampling deaths. Of the 21 fatalities during this disaster, nine of these were confirmed as crush deaths (compressive or traumatic asphyxia) at autopsy.

It should be noted that the E2 nightclub had, previous to this event, been cited numerous times for building code violations, but the club had remained open regardless. In the aftermath of the event, the club’s owners were convicted of indirect criminal contempt over repeated failures to close the club despite specific court orders requiring them to do so [1].

Types of Force Present

The primary type of force present within the E2 Nightclub evacuation was a stacking force, caused by the large number of people attempting to effect egress through the sole exit to the structure. On the night of this incident the occupancy of the E2 Nightclub was at over four times it’s suggested capacity. The two bottlenecks present in the exit structure were the doors leading into and out of the stairway, which were approximately 150cm wide,

and the width of the stairway itself, which was approximately 140cm wide. If we assume the lower bounds of the maximum specific flow rate of these structures [98], we can estimate the outflow through the doorway and stairs of the E2 Nightclub. The flow rate for a door of 1.5m is approximately $1.97ps^{-1}$, whilst a stairway of 1.4 metres is approximately $1.3ps^{-1}$. We see from these figures that the stairway becomes the bottleneck in this situation, which is not too unusual. What is unusual in this situation is the presence of a door which opens outwards into the upper landing. As stated by witnesses to this disaster, the doorway to the upper landing was repeatedly thrown open as people were attempting to escape, this action exerted an accidental toppling force on those individuals already on the upper landing, which was the primary cause of the severe levels of stacking which built up during this evacuation.

It should be noted that the author accepts that there would also have likely been both pushing force, and leaning force present in this situation. Pushing force caused by the large number of people attempting to exit through the upper landing door at the same time, and leaning force caused by the toppling of those on the upper landing itself towards the stairway. Although in this case it is not thought that these forces were a primary cause of the disaster.

Five Factor Analysis

Spatial

The number of patrons of E2 nightclub, at the time of the incident, was more than four times its suggested capacity of 240 people, with estimated occupancy figures ranging between 1100 and 1500 patrons. It has been claimed after the event that the City of Chicago failed to set a maximum occupancy for the venue, but it is widely held that the occupancy levels of that night could not have been deemed to be reasonable for a building of its size [24].

Temporal

The speed at which the events unfolded played a large part in the formation of crush conditions during the E2 Nightclub incident. The release of the pepper spray caused fear in a large portion of the crowd, and created a high desire to leave. The ensuing rush toward the main exit overwhelmed the exit capacity, which would have been insufficient for the population even under

non-emergency conditions.

Perceptual/Cognitive

The perception of threat in the E2 incident was far greater than the threat actually present, and this caused an extremely high desire to leave in many of the club's occupants. Additional factors may have also played a part in raising the perception of threat. An example of this can be found in the reports that one patron was heard to shout "I'll bet it's [Osama] Bin Laden!" during the evacuation [127]. Combined with the unknown (to a large number of the patrons) chemical substance present in the environment, it is understandable that the level of fear experienced by some patrons would have been heightened.

Procedural

The procedural failings found in this case are numerous. Repeated failures to meet building and fire codes by the owners are the main procedural failures, which were compounded by the alleged failure of the City of Chicago to set an adequate maximum capacity to the venue. Also, the release of a chemical toxin in an enclosed space by the security staff is a major failing, as this acted as the precipitating factor in this case, leading to the evacuation taking place.

Structural

According to the City of Chicago and evidence given at trial [24], the E2 Nightclub venue was not fit for purpose, and it had been shown on multiple occasions that this was the case. Ignoring these past code violations, the main evacuation issues on this occasion were caused by the inward opening door at the main exit. Once a number of patrons had come up against this inward opening door, and others had followed behind them, the exit became completely impassable, and the build up of forces that resulted were the sole cause of death in this event which took the lives of 21 people in total.

2.7.5 Mihong Bridge Spring Festival Disaster

The second annual Lantern Exhibition of Miyun county was due to take place from the 31st January until the 10th February 2004, in Miyun county China [155]. Organisers expected less than four thousand people per day to attend, based on the previous year's figures, and for the first five days of the exhibition this was shown to be the case, with attendance figures falling between two and three thousand per day. The sixth day of the festival

however (5th February), the traditional day on which the Chinese Lantern Festival itself would have been held, saw attendance levels increase to as many as forty thousand people, more than ten times that of previous day's figures and of official estimates. This increase in attendees alone did not cause any serious issues, as the festival site was spread out across a large area, but when fireworks were set off near the river, many within the crowd headed towards the Rainbow Bridge, which was thought to be the best position to view the firework display.

It has been estimated that, previous to the fireworks being released, the number of people on the bridge was less than three hundred, which is well within the structure's capacity. The sudden influx of pedestrians caused by the desire to get a better view of the firework display, caused an overloading of the bridge's capacity, and the crush disaster began. Eyewitness accounts suggest that the crush disaster only lasted between seven and eight minutes, at which point stewards began to alleviate the situation by guiding people away from the bridge, but during the short crowd crush thirty-seven people lost their lives.

Estimates suggest that the number of people on the bridge could have reached as high as one thousand three hundred people during the period immediately after the fireworks had been released (these estimates are taken from both survivor accounts and the results of simulations [155]). At this point the crowd density at some points on the bridge would have reached over eight persons per square metre ($8pm^{-2}$), and possibly as high as ten persons per metre, allowing personal space of just $0.13mp^{-2}$, which is lower than Fruin's "Level of Service F" for all types of pedestrian situation (i.e. walkways, stairways, or queues) [37], but less than one-third of LoS F for a stairway, which is the most comparable pedestrian situation to the Miyun bridge. The physical situation on the bridge was compounded by the steepness of the sides of the bridge, which were over 31° from the horizontal. Simulations suggest that the high crowd density combined with the steep angle of the walkway caused additional injury as such conditions reduce pedestrians' ability to remain upright under high density conditions. This led to increased "leaning force" and also contributed to the presence of "trample deaths" in which pedestrians lose their footing and fall to the floor. Surrounding pedestrians, unable to readily effect their own movement have difficulty in avoiding placing their weight on the fallen pedestrian, and

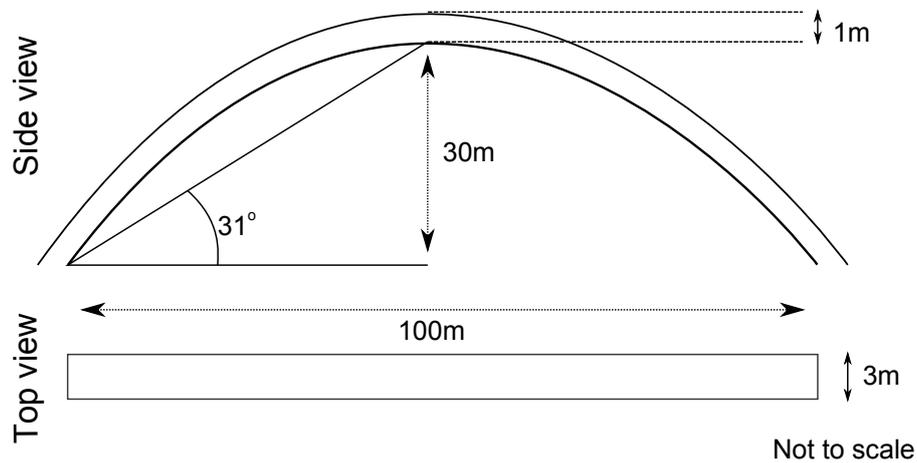


Figure 2.10: Top: The Miyun county “Rainbow Bridge”, the scene of the Miyun bridge disaster of 5th February 2004. Bottom: Diagram showing dimensions of Miyun bridge (not to scale).

in more extreme densities may not even be aware of their presence. Calculations ascertained that the leaning force produced by this crowd could have exceeded $6kNm^{-1}$.

The reason for these high levels of leaning force are partially down to the overcrowding of the bridge, and partially down to the structure of the steps on the bridge itself. Two of the defining characteristics of steps, mathematically at least, are *rise height*, i.e. the difference in height between the top

of one step and the top of another, and *tread depth*, i.e. the actual depth of each step. The function of the steps should be taken into account during the design of these two facets, e.g. a structure used for general standing and milling will have lower rise height and greater tread depth than a set of stairs designed to transport pedestrians from one level of a structure to another. The Rainbow bridge had an incline of 31° , a riser height of 25cm, and a tread depth of 40cm. Compare this to an investigation of football terracing in the UK, carried out by Dickie and Wanless [29], which considering the bridge was being used for a viewing area for a fireworks display is a roughly analogous purpose, and we see that for a milling or viewing area a reasonably safe configuration could be considered to be a 15° incline, a riser of 9.5cm, and a tread of 35.5cm. Comparing these two situations it becomes readily apparent that the steps on the Rainbow bridge were drastically unfit for purpose.

Simulations of the over-crowding of the bridge were run as part of the investigation into the disaster. The simulation environment buildingEXODUS was employed for this, and were initially configured so that they matched the conditions on the bridge before the influx of persons began, this equates to approximately 19:20 in real time. The influx of pedestrians was simulated from this point, and density at different points on the bridge was measured until conditions became untenable. Snapshots of these simulations may be seen in Figure 2.11.

The simulation's start, 0:00, corresponds to the conditions on the bridge at 19:20, and ends at approximately 15:00, which was when emergency measures began to be taken to alleviate the crush. We can see from the simulation results that the conditions on the bridge degraded at an alarming rate, with each snapshot (approx 3 mins apart) showing marked increase in higher density sections of the bridge.

Types of Force Present

The primary type of force present during the Mihong Bridge disaster was that of leaning force, caused by a high density crowd standing on a stepped area. Investigations into this disaster have suggested that, in the case of the Rainbow Bridge, the step height of the bridge, which was 25cm, was not suitable for a venue of public gathering. This step height, whilst considered within reasonable boundaries for an evacuation route, is not considered suit-

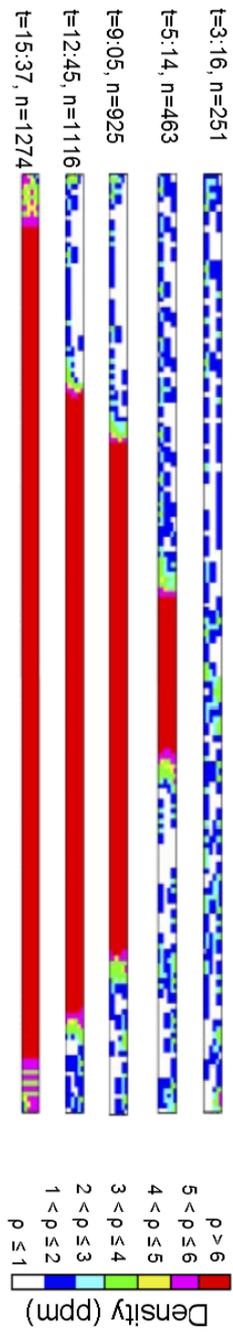


Figure 2.11: Top down density maps of Mihong bridge at different times from the start of the incident [155]

able for an area where pedestrian may gather and “mill”, to put this figure in perspective the standard step height, also known as *rise*, of a domestic stairway would be 15-22cm, whilst the steps at a public venue, e.g. standing areas at football stadia (a reasonable approximation to the purpose of the Rainbow Bridge), are found to be closer to 10cm. The steep incline (31° from the base of the bridge to its summit) and large rise of the steps on the Mihong bridge was not fit for purpose, and would have increased the amount of leaning force generated by the densely packed crowd.

Five Factor Analysis

Spatial

The first five days of Festival had seen roughly three thousand tourists per day visiting the park, but on the day of the disaster this rose dramatically to approximately forty thousand tourists. This in itself would not have been a problem, as the festival was held in a large park. The problems occurred when the occupancy of the bridge showed a sharp rise. Preceding the disaster, the occupancy of the bridge was estimated at three hundred people, but the influx of pedestrians caused this figure to rise sharply, and within minutes the occupancy of the bridge may have reached as high as one thousand three hundred people. At this level of occupancy the density of the crowd could have reached as high as $8pm^{-2}$ [155] (p. 8).

Temporal

The rapidly changing conditions on the bridge were certainly a factor in the onset of crush, as little time was available for officials to react. This said, there were security personnel stationed on or near the bridge at the time, see *Procedural factors*, below.

Perceptual/Cognitive

The desire to get a good view of the fireworks display drew masses of spectators toward the bridge, although accounts do not suggest a high level of *competition* (i.e. pushing, shoving, etc) for places on the bridge. Once again the inability to perceive the severity of conditions played a large part in this disaster, as those entering the area would have been unaware of the conditions ahead, so made no attempt to turn back or to alleviate the pressure.

Procedural

The procedural factors in this incident were mainly failures of planning.

The under-estimation of the the crowd levels played a large part but, as the festival ground was spread out over a large park, was not the main failure. During the previous year's event, there had been security staff stationed by the bridge to limit the flow of pedestrians to a reasonable number. Unfortunately no one had been assigned this duty in 2004 [155] (p. 6). This measure could have averted the disaster. This was also the finding of the People's Procurator, who in 2005 sentenced three officials to prison sentences for their part in this disaster.

Structural

An important factor in this disaster was the design of the bridge itself, which was found to have an unacceptable gradient and large step height [155] (p. 7), see Section 2.6.3 for discussion on leaning forces. Investigations have shown that this step height (25cm) would have made it exceedingly hard to keep one's balance if being pushed from behind. This is thought to have contributed to the high levels of force present during this disaster due to "leaning force", which is the cumulative force created when numerous pedestrians lean forward at unnatural angles, and often found in stadium or arena settings.

2.8 A Diagnosis Issue in Crowd Crush Situations

It is known that a problem exists with the classification of *crush related* injuries and deaths, namely in the process of posthumous diagnoses (ascertaining a *cause of death*). Firstly, there is the problem with the attribution of injury to trampling. In the words of Gill and Landi

These deaths often are attributed mistakenly to blunt impacts from trampling. The autopsy, however, typically finds inconsequential blunt injury but does find signs of traumatic asphyxia.

Gill *et al* (2004) [42]

This problem arises from the level of force present during crush disasters, in short, the high crowd densities found during crush conditions are known to prevent the victim from falling to the floor, as the compression that they have been placed under has a tendency to hold them in an upright position due to friction effects between members of the crowd. It is often

only after the pressure has subsided slightly that the victim will collapse to the floor, at which point the fatal injury would have already occurred but subsequent crowd movement can still cause *superficial* injuries (termed by Gill and Landi as “inconsequential blunt injury”) that are normally consistent with a trampling death [42]. In these cases, the attribution of mortality to compressive or traumatic asphyxia is only possible after a full autopsy has been carried out.

A similar problem exists when the presence of fire, smoke or other toxins in the environment has compounded the effects of the compressive asphyxia.

In modern building fires the two main asphyxiant gases present are usually carbon monoxide (CO) and hydrogen cyanide (HCN). Carbon monoxide is produced when organic matter undergoes an incomplete combustion due to a shortage of oxygen, a situation common in building fires where enclosed spaces cause oxygen supplies to quickly deplete. Hydrogen cyanide is a common bi-product of the combustion of many synthetically produced items found in the modern environment. Considered more toxic than CO, HCN can be found in large quantities wherever man-made synthetics are burned (especially in low oxygen environments). Whilst it is true that natural materials such as cotton or paper produce small amounts of HCN the primary sources in most building fires would be from nylon, polyurethane, and acrylonitrile, which are commonly found in carpets, foam insulation, clothing, and plastic products. CO and HCN within the environment further increase the risk of asphyxia. If we take the example of a fire, which is a common cause of evacuations, the combustion process requires both a fuel (e.g. wood, cloth) and an oxidant (e.g. oxygen, fluorine) to enable continuous burning. This means that any oxidants in the vicinity are likely to be used up by the fire, and without ventilation this will result in an extremely low-oxygen environment which is also rich in asphyxiant gases.

As we can see from this example, the major signs of a crush related death can easily be mimicked by a death in close proximity to fire, smoke, or other toxins. Byard *et al* state, in a study spanning 25 years of traumatic asphyxial mortalities, that

[A diagnosis of] fatal crush asphyxia may have to be a diagnosis of exclusion, made only when there are characteristic death scene findings . . .

Byard *et al* (2006) [14].

It is also common that the presence of crush, whilst not attributed as the cause of death at autopsy, has been seen to play a significant part in the individual's demise. This is common when in situations where crush conditions are found to have occurred at the same time as significant smoke or toxins are prevalent in the environment. In situations such as these, this forms a dual threat to the individuals' ability to breathe, with the asphyxia causing little breath to be taken in whilst any breaths that are taken would contain an extremely low level of oxygen, a high concentration of toxins, or both. It seems obvious then that, even when a victim's cause of death is not *directly* attributable to the presence of force within an evacuation, the presence of force in any evacuation must be considered an exacerbating factor, and therefore be planned for, and designed against at any opportunity.

2.9 Detecting Crush Conditions via Phase Transitions - An Initial Idea

During the Hajj pilgrimage, Johansson *et al* observed unusual behaviour in the crowds of people [63]. The behaviour observed appeared to show that immediately preceding times of high turbulence within the crowd, a behavioural *phase transition* could be seen, which marked the transition between the smooth laminar flow of an ordered crowd, to a turbulent state in which the onset of crush conditions *could* begin.

This type of phase transition can be seen in many kinetic systems, from Ising-spin systems to theoretical particle systems (see Section 5.10). It has been found that the point at which the system shifts *phase*, that is the point at which the behaviour of the system changes state, can be reliably identified by mathematical means.

The five factors contributing to the formation of crush conditions (Section 2.5) all have the potential to effect the behaviour of pedestrians. Within a kinetic system, a change in behaviour can often be identified by the subsequent change in the *movement patterns* within that system. Poorly designed structures, fast changing temporal conditions, procedural issues, tight spatial constraints and perceptual and cognitive factors will all effect the behaviour of a pedestrian, and also all contribute to the formation of crush.

What we suggest then, is that the analysis of *movement patterns* can be used to predict the probability of crush formation.

In evacuation terms, and specifically to the evacuations we will consider in this thesis, the build-up of crush can be considered in a sequence of *steps*. At the beginning of an evacuation there will most likely be a number of people dispersed throughout a structure, who will begin to move towards the exit. After this stage, when the evacuees get close to the exit the dispersal of these individuals will reduce, and the crowd density around the exit(s) will increase, as per any evacuation. It is at this point that crush may occur, if the evacuees sense of urgency causes competition for the available exit, and leads to the sort of non-adaptive crowd behaviour that we have discussed previously. At this point the interpersonal forces (such as friction) increase, the net speed of the evacuation is reduced, and physical force begins to build up within the evacuating crowd. The point which we would like to identify is that where the usual evacuation pattern of densely packed but non-competitive evacuees transitions to non-adaptive, competitive behaviour. It is this point where the situation has the potential to lead to crush conditions.

We propose that by tracking and identifying changes in pedestrian movement patterns, we can identify the probability of crush conditions forming at any one time. Over the next Chapters, we define the technique, and apply it to evacuation systems, showing that the level of *risk* to pedestrians within an evacuation can be measured in this way.

2.10 Scope of this Study

As we have seen from the material presented in this Chapter there are a wide range of situations in which the build up of force within a crowd of people can lead to serious injury or death. The types of force detailed have all shown the possibility of leading to large scale disaster if present during an evacuation scenario, but the focus of this thesis going forward will be to specifically look at the build-up of *pushing force* within crowds of people as they evacuate. The reason for this is twofold. Firstly, the examples previously detailed all contained elements of pushing force within the evacuation, and by extension the evacuations were hindered by the friction generated by these forces (see Section 2.6.1). Secondly, the modelling of pushing force

has been more thoroughly investigated (certainly in the field of agent-based modelling) than either stacking or leaning forces, although more simulation environments are starting to model the effects of leaning forces as they become ever increasingly more complex.

Therefore, for the rest of this thesis, when the term *force* is used with respect to evacuation it will be used to define the pushing forces within an evacuating crowd, unless explicitly stated otherwise.

2.11 Summary

We have defined the term evacuation, and how it will be applied during this research, and briefly summarised the current knowledge of the behaviour of individuals and crowds during an evacuation. We have addressed popular myths regarding non-adaptive crowd behaviours, and shown that the decisions and behaviour of evacuees, far from being irrational, is most often found to be both rational and logical, being based on decision making rather than blind panic.

Crush conditions have been shown to be possible in any situation in which large numbers of people gather. We have detailed examples in which crush conditions presented during emergency evacuations, sporting events, and religious festivals, but many further examples can be found in the literature. We have shown that the formation of crush is a complex emergent phenomenon, which can be difficult to predict and therefore is hard to protect against.

Considering this, methods by which the evacuation of a structure, and the danger that the formation of crush presents during this evacuation must be investigated. In the next Chapter, the use of computational evacuation models is discussed, and the current state-of-the-art regarding the inclusion of crush detection measures within these models is investigated.

Chapter 3

Computational Studies of Evacuation

3.1 Introduction

To establish the state of computational crush modelling today, we identify three main methodologies for evacuation modelling found in current circulation, and describe the difference between *movement models* and *behavioural models*. The current trends in evacuation models and evacuation modelling are investigated, and we track the trends of evacuation models over the past twenty years.

3.2 Computational Evacuation Modelling

The field of evacuation modelling grew from advances in the field of fire safety that occurred during the mid stages of the 20th century, and has evolved from the earliest *hand calculations* and general design rules into a field spanning many forms of model, environment, and techniques. The current state of the field will be discussed in the following section, with specific reference to the factors and techniques relevant to this work going forward. We will concern ourselves primarily with computational models, rather than the theoretic or mathematical models underlying their operation, unless these factors are of direct importance to later work.

The current state of the field of evacuation modelling contains many varied and diverse modelling techniques, but for the purposes of this work

we will divide them into three general categories of approach, these being network node, cellular automata [148], and continuous field. These three categories were created for their use in later stages of this project, and concern the movement of persons through their environment.

3.2.1 Representations of Physical Environment

One of the defining factors of an evacuation model is the way in which it represents both the agents, individual pedestrians, and the environment in which they exist, such as the buildings, ships, or stadia. Initial computational models of evacuation were an extension of the queueing models (which had been used previous to simulation), systems of connected areas through which pedestrians may pass to effect egress. As models progressed, further options became available for these representations, options which define far more than a simple representation of structure or interactivity. In order to accurately simulate pedestrian movement the physical environment must be described. The next sections cover the three main types of physical representation in use today, and discuss the strengths and weakness of each approach.

Network Node Models

The network node models are some of the oldest in the field, yet are still extensively used for purposes such as shortest/fastest path finding and minimum cost network flow. These models originated from research into the movement of pedestrians in public spaces by academics such as Fruin [37], Pauls [105, 106], and Predtetchenski and Milinski [109]. These models operate by reducing architectural structures to their base components, which can then be visualised as vertices within a graph, which each have an associated weighting which represent the time that it would take the average evacuees to traverse it. This time varies with current capacity, or by means of an additional traversal weighting equation representing differing conditions during an evacuation. These timings and weighting are mostly based on traditional *hand calculations* that were used for evaluation of evacuation times, and have been shown to remain reasonably accurate in their predictions [87, 126].

The main advantage of these models is that they are extremely fast run-

ning, relative to those implemented using either cellular automata or continuous field approaches, as the minimum amount of computing power is used to calculate the actual traversal of the structure. The higher computational overhead is found in the optimal route findings algorithms implemented in these models, which can enable evacuees to calculate their escape path. However, even when this is considered, these type of models still simulate evacuations at much reduced computational cost relative to models based on other approaches.

The principle of these models is mathematically sound, with many of them employing techniques from the fields of *queueing theory* and *graph theory* [17, 123], and having been tested over many years by validation against empirical data. Yet there are many situations in which the network node approach cannot be employed, for example an engineer modelling the effects that interactions between individuals during an evacuation has on the overall egress time, or the modelling of crush conditions such as we deal with during this project. A problem is also presented when modelling non-standard evacuation topologies that do not lend themselves to being rendered in the traditional room-corridor-stairwell-corridor fashion that is required for the network node approach to be effective.

In short, the network-node modelling approach has a tendency to manipulate the complex, *real-world* design of buildings into structures that are well represented by a network graph. This is not a problem with smaller, more traditional structures, but larger and more complex structures present significant difficulties.

Many examples of this type of simulation environment remain in distribution and are still utilised to great effect within the field, models of particular note are; EVACNET4 [67] and EXIT89 [33].

Strengths

- Lower run-times per simulation than either cellular automata or continuum models.
- Allows the simple inclusion of advanced route finding algorithms.
- Ideally suited to the calculation of ASET and RSET times across a large number of eventualities.

- The simplicity of the models require less variables to be supported with empirical data.

Weaknesses

- Poor physical representations of pedestrians.
- Heavy reliance on statistical data, which can make them unsuitable for previously untested topologies.
- Limited ability to model inter-personal communication.

Cellular Automata

The Cellular Automata (CA) [148] approach (also referred to as *fine grid*, *course grained*, or *discrete floor field*) subdivides the evacuation grid into a finite number of discrete cells which the occupants transition to and from directly as they effect egress. These models allow the simulation of individuals in a more realistic manner than is possible using the network node modelling approach, and have a number of advantages. Firstly, as an individual is represented in a more accurate manner than in the network node models (i.e. having physical dimensions, able to effect movement in two dimensions rather than just one, etc.) the reliance on statistical data regarding pedestrian flow is partially removed.

An example of this can be found by comparing the modelling of an open space in hypothetical network node and cellular automata models. Consider a square atrium, such as that found in structures like hotels, apartment buildings, or offices, with four possible entrances/exits. Using a network node model, the atrium will be filled according to the capacity defined during its specification, and pedestrians are allowed to pass through this space in a predefined time. The approach is not in any way incorrect, as in low density situations the amount of time taken to cross the atrium will remain reasonably stable, but if we consider the case of a higher density situation in which the population are attempting to evacuate, we must consider the possibility that pedestrian flow may reduce the time required to pass through this space in unforeseen ways. Consider Figure 3.1, which shows the influx

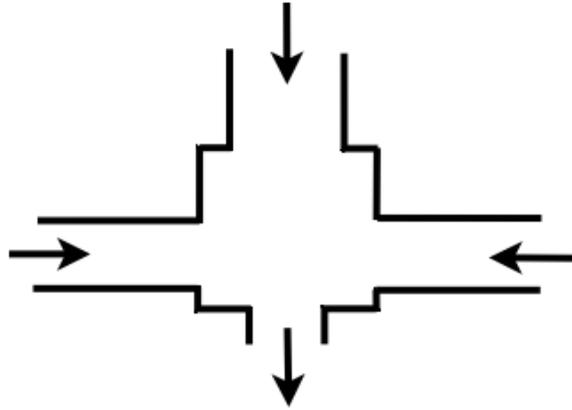


Figure 3.1: An atrium as may be found in large buildings, during evacuation inflow will occur at North, East, and West intakes, whilst exit will occur at the South.

of pedestrians from the North, West, and East intakes, who are attempting to exit the atrium via the South exit. At low densities the time taken to cross this atrium will be roughly the same as the time taken by someone to cover the same distance at average walking speed. At higher densities however, the interactions between evacuees can cause unpredicted conflicts that will increase the evacuation time per person, as different streams of people must make their way through the structure.

A cellular automata model allows the more *fine grained* simulation of pedestrian movement, and can also model complex behaviours such as person to person interactions and obstacle avoidance. They are considered by many to offer more accurate simulation of occupant movement than is possible with the network node technique [126].

The computation speed of these models is far slower than that of the network node models as this is an entirely *agent based* approach, which means each pedestrian is modelled as an individual discrete unit. Applying a CA approach to modelling allows significant optimisations to be made to many parts of a model, most notably by applying the *floor field* technique [66] which associates certain evacuee variable values with the grid structure, rather than requiring the calculation of each value for each evacuee at every time-step. Direction finding for example may be calculated once per each discrete element of the floor field, the variable value being assigned as the vector pointing directly towards the nearest exit, or away from the nearest

obstacle. The savings afforded by this type of modelling are many, but the grid structure presents problems when modelling crush conditions within an evacuation. The two most pertinent problems are the modelling of force and friction between evacuees, which arise from the reduction of the evacuation grid limiting the level of interaction and movement of evacuees. Commonly, evacuees are able to traverse the grid in just 4, 6, or 8 possible directions, as opposed to the (effectively) infinite number of possible headings allowable in a continuum model. This reduces the effects that interpersonal friction, a force that inhibits free movement in other models, has within a simulation and has traditionally restricted some of the behaviours that can be recreated using this approach. A solution to this problem is found in certain CAFE models (Cellular Automata with Force Essentials) which apply conflict resolution algorithms to evacuees attempting to move into the same grid cell. These algorithms can be as simple as a *dice rolling* solution in which two or more agents who are attempting to move into the same cell are required to resolve the conflict via random number generation, the conflict is resolved only when one agent “rolls a six”. This approach is taken in [66], and has been shown to partially mimic the effects of inter-personal friction seen in continuum models [126, 150]. These type of models are used extensively within the field, and have been shown to produce accurate results across many scenarios.

Popular examples of environments which utilise the cellular automata approach are SGEM [86] and buildingEXODUS [102].

Strengths

- Lower run-times per simulation than continuum models.
- Able to model non-adaptive group behaviours and competitive egress.
- Low reliance on statistical data during operation.
- Allows the inclusions of more accurate environmental conditions (fire, smoke, toxins, etc).

Weaknesses

- Higher run-times per simulation than network-node models.
- Unable to reliably model physical inter-personal contact.
- Larger *cells* (a more coarse-grained environment) can effect physical realism of movement.
- Cannot model inter-personal friction, although there are ways to simulate this factor.

Continuum

The final form of grid structure in use today is the Continuous floor (or *Continuum* model), in which the evacuation grid is represented as a continuous plane, separated only by obstacles and the form of the architectural structure. Continuum models allow movement throughout the evacuation grid at a highly accurate level, with evacuees able to assume any heading and to take steps in any valid direction. Due to this continuous nature, these models must include collision detection algorithms to ensure that the evacuees physical form is accurately represented at all times, which adds another significant computational overhead. The continuous evacuation grid also hinders the use of floor fields within most models, forcing the calculation of many variables at each time-step, although for certain evacuation parameters floor fields can still be implemented, FDS [74] implements their path finding algorithm using a CA style floor field approach (see Chapter 8).

The continuum models have been shown to offer a great level of realism when modelling the movements of pedestrians throughout a structure, but the additional computational expense of continuous floors and physical contact often means that complex behavioural models are not included within these simulations.

A popular example of this type of model is the Fire Dynamics Simulator, which we discuss further in Chapter 8.

Strengths

- Able to model physical movements to a near *real-life* accuracy
- Can accurately model physical inter-personal contact (force and friction)
- Allows greater physical diversities within pedestrian population (height, weight, etc)
- Most accurate representation of human form of all methodologies

Weaknesses

- Higher run-times per simulation than either network-node or cellular automata models.

Critical Analysis

The different approaches to floor field representation each have their own strengths and weaknesses, and examples of each type of model are still in use to this day. This can be seen in Section 3.3.1, where out of the eight most popular evacuations models in use during the past ten years, there can be found at least one example of each of these methodological approaches.

The choice of differing floor field representations with evacuation models depends on the desired use of the model, as the floor field representation can have a great effect on the strengths and weaknesses that a model may exhibit. We have seen, for example, that the network-node approach to modelling allows for extremely fast computation times, which means that large numbers of distinct evacuation scenarios could be tested in very little time. This could allow for Monte-Carlo modelling of all predictable eventualities, an option that would be cumbersome to achieve using a model which incurred a higher computational cost. Equally, incorporating advanced behavioural models within a network-node simulation presents myriad problems, as inter-personal interaction is difficult to define, so designers considering this type

of additional functionality would be inclined to choose a cellular automata or continuum approach.

The modelling of crush, covered in more detail in Section 3.4, presents a more serious decision to modellers, this being that the inclusion of crush measurement techniques have traditionally required the use of a continuous floor model, which necessitates a large computational overhead. The requirement of the continuous approach for crush modelling is due to the inter-personal friction required for the modelling of high density pedestrian situations, which is not easily achieved using cellular automata or network-node approaches.

We propose applying a technique which can assign a metric to the formation of dangerous crush conditions, which would not require the calculation of physical force levels. If such a technique were to be successful, it would allow the inclusion of a crush detection metric within a cellular automata model, and negate the need for computational expensive physical force calculations. Over the next chapters we introduce further elements leading to the application of this technique, and show that it is possible to identify behaviours which are likely to lead to the formation of crush using movement variables alone.

3.2.2 Behavioural/Movement Modelling

There is a distinction between the modelling of *behaviour* and the modelling of *movement*, which can often become blurred in the literature. There are a large number of *particle* or *fluid* models which ostensibly model the behaviour of evacuees, but on inspection they are more strictly modelling their movement. We next define and contrast the two techniques of modelling human movement generally taken when designing/specifying an evacuation simulation.

Behaviour Models

Comprehensive behavioural models are now included in many evacuation simulations [35, 84, 102, 140], and are being used to model a far greater number of behaviours than have previously been possible. An example of a behavioural model that may be used in an evacuation simulation is shown in Figure 3.2, which shows a subset of the interactions that can occur as

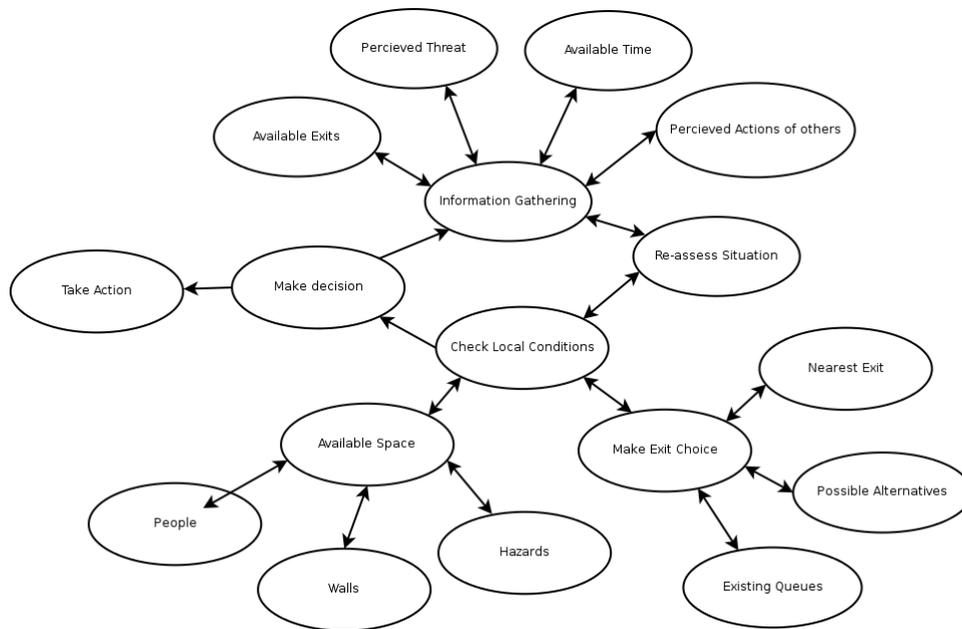


Figure 3.2: Example of a cognitive behavioural model

part of the decision making process.

The underlying computation in these systems can work in many ways, and common implementations range from probabilistic rule-based models through to the more complex systems implementing Artificial Intelligence (AI) techniques to create realistic human decision making processes. Well known examples of such models are VEgAS [62] and Legion [22]. For an overview of models in circulation, including a breakdown of types of behavioural models, see Kuligowski and Peacock [79].

Movement Models

A classic example of a movement model is the Social Forces model (SFM) [51], examined more thoroughly in Chapter 4, in which evacuees are represented as particles subject to forces around them, and able to exert force to drive them toward the exit. The movement model present in the SFM is visualised in Figure 3.3. We see that the model present in the SFM is highly simplified when compared to the example behavioural model, taking into account just the immediate surroundings of the evacuee in question. This models only the movement of the evacuee, and cannot take into account factors such as inter-personal relationships, knowledge of the structure to be

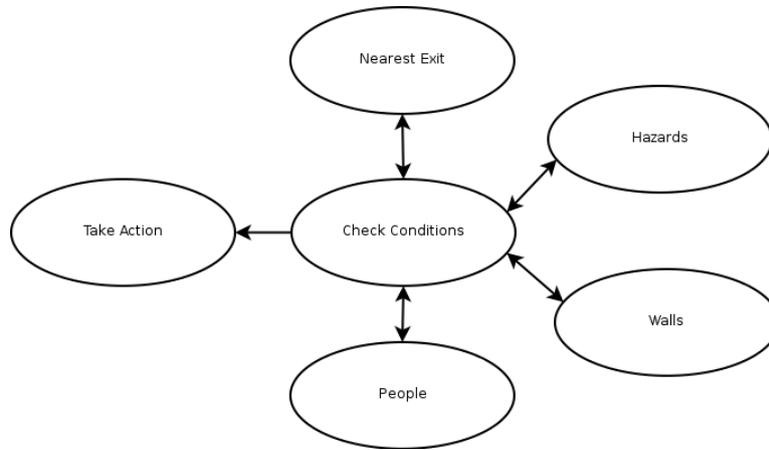


Figure 3.3: Example Movement Model. Based on Helbing’s pedestrian model

evacuated, availability of additional information in the form of exits signs, perception of risks, etc.

As can be seen when comparing the movement and behavioural processes, a movement model utilises just a fraction of the decision making parameters present in the behavioural model, yet can still accurately represent many situations to a great degree of accuracy.

3.3 Recent Trends in Research Activity

To investigate the popularity of computational evacuation modelling, a publication review was carried out to assess the varying level of interest these models have been achieving in the field over the past twenty years (1991 to 2010 inclusive). Three journal indexing sites were identified for data collection, these being Science Direct [32], IEEE Xplore [57] and Web of Knowledge [141]. These sites were chosen both for the known inclusion of evacuation materials (from personal experience) but also for their accuracy of their search functionality (i.e. the correctness of the results returned compared with the exact terms searched). The terms chosen for these searches, and the way in which they were searched for, is shown below:

$$(evacuation \vee egress \vee ingress) \wedge (model \vee models \vee modelling \vee modeling)$$

The data obtained may be seen in Figure 3.4.

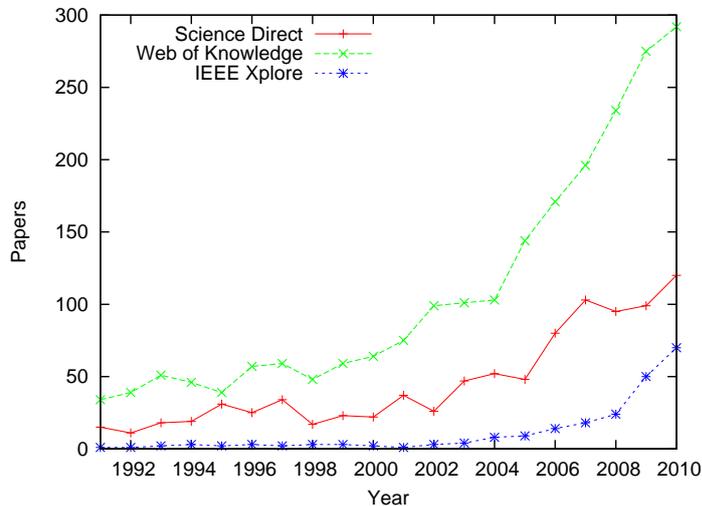


Figure 3.4: Publications on evacuation modelling over the last twenty year period.

As we observe from the data, there has been a continual growth in interest in computational evacuation over the twenty year period in question. The years 1991 to 2000 show a steady, yet slow, growth in the number of articles across all three of the data sources, but the rate of growth increases noticeably in the period 2000 to 2010. This increased rate of growth could be attributed to many factors, such as the increasing power of personal computers allowing for more complex functionality to be added to models, or increased interest in the field after numerous high-profile events which involved the fast evacuation of people (i.e. the World Trade Centres attack [100] and the London Underground bombings [145]). It could be said that the trends are indicative of the rise in output seen across all sciences over the past twenty to forty years, but estimates suggest that this rise is likely to be in the region of 2.5% per annum [97]. This is significantly less than the trends seen here between 2001 and 2010, which was an approximate 24% increase per annum (data taken from “Science Direct” and averaged across the 10 year period from 2001). This does show that the increase in output across the field of evacuation modelling cannot solely be ascribed to an increase in output across academia in general.

Whatever the cause of this increased growth, the rising popularity of

computational evacuation has caused the field to become a more “mainstream” discipline, and the computational modelling of evacuations is now often found to be carried out routinely during the design and planning stages of large building projects.

3.3.1 Trends in Model Usage

We have seen that the research area has been in ascendancy for the last two decades, but this doesn’t tell us much about the popularity of the specific models that people have been using across this time. We can apply a similar method to that used previously, to assess the popularity of specific models. For this we will focus on the last decade (2000 to 2010), as this is where the most pronounced growth in the field was identified, and collect data on the number of journal papers which directly employ specific models.

The specific models included in this research have been compiled during the course of the work, and are not meant to represent a comprehensive list of all evacuation models available. The list of models which have been encountered during this work are listed in Table 3.1. The models listed in Table 3.1 do not represent a comprehensive list of all available evacuation models. The list was compiled throughout this research, and includes any model that has been included in the available literature, alluded through external sources, or found during the review process. This does not imply that a model listed is in current use or is being actively developed at the time of writing. A community updated list of known models can be found at [evacmod.net](http://www.evacmod.net) (<http://www.evacmod.net>), and lists of models in active use can be obtained from the NIST website (<http://www.nist.gov>).

The data for the IEEE Xplore engine was not included in this section of the research as there were very few publications on the subject of evacuation included in their journals between the years 2000 to 2010 (compared with the other two sources). We have therefore analysed the activity using Science Direct and Web of Knowledge only, results are shown in Figure 3.5.

Discussion

As we can see from the data obtained from both sources, the models which currently dominate the field, as far as research is concerned, have been FDS+Evac, buildingEXODUS, Simulex, and the Social Forces model. Of

Model Name		
AENEAS	Allsafe	ASERI
ASET / ASET-B	BFIRES-2	BGRAF
BuildGEM	BUMMPEE	Cube Avenue
buildingEXODUS	CRISP	DBES
EARM	EESCAPE	EGRESS
Egress Complexity Model	EgressPro	ELVAC
ENTROPY	EPT	ERM
E-SCAPE	ESM	EVACNET4/EVACNET+
EVACS	EVACSIM	EvacuationNZ
Evi	EXIT89	EXITT
FAST	FDS+Evac	Firescap
FlowTech	FPETool	GridFlow
Helios	Legion Studio	MA&D
MAGNETIC Model	maritimeEXODUS	MASCM
MASSEgress	MASSIVE Software	MASSMotion
Myriad II	Nomad	PathFinder
PEDFLOW	PedGo	PEDROUTE / PAXPORT
PedSim	S-Cape	SEVE_P
SGEM	SimPed	Simulex
SimWalk	Social Forces Model	SMART Move
SpaceSensor	STEPS	Takahashi's Fluid Model
TIMTEX	TSEA	VEgAS
VISSIM	WAYOUT	ZET

Table 3.1: List of Evacuation models considered during literature review. Note: A model's inclusion in this table does not imply that it is widely used or actively developed, just that research papers exist which have explicitly referenced the model in question.

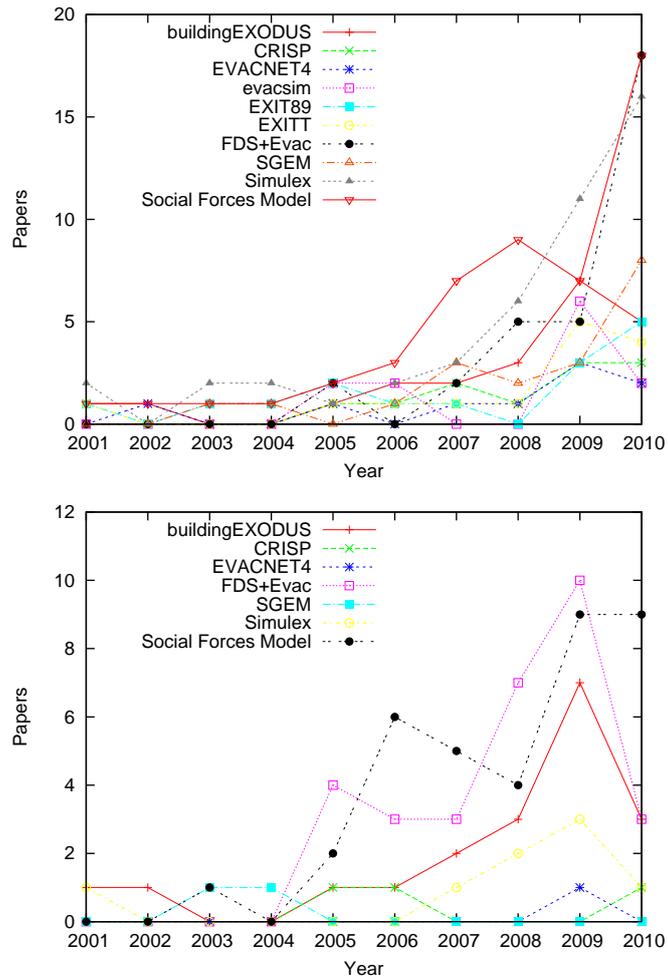


Figure 3.5: Number of papers published each year in which the evacuation model has been directly referenced. Data taken from *Science Direct* [32] (top), and from *Thompson Reuters Web of Knowledge* [141] (bottom).

these four models, the cellular automata approach is employed in building EXODUS and Simulex, whilst the continuum approach has been employed in FDS+Evac and the Social Forces model.

Obviously the age of these evacuation models may have an impact on the amount of interest shown in them. Figure 3.6 shows the time-line of model release of all the models investigated previously. Model appearance is defined by the first instance of their reference within a published research paper.

We can see from the time-line that the older models, pre-1990, have remained at a reasonably static rate of reference over the past ten years, albeit a closer look at the manner of their reference shows that they are most regularly used as benchmarking tools for the newer generation of model. If we take the case of EXIT89 as an example, this was directly referenced in six papers published through journals indexed by “Science Direct” in the five years from 2003, but in each of these papers they were used only as comparative models, and none contained new results from *in silico* evacuations modelled using the simulation environment.

These trends do show that the field in general is moving toward a more “agent based” approach to the modelling of evacuation, which naturally favours the techniques of the CA and continuum model over that of the older network node approach. This said, the network-node models included in this survey (EVACNET+, EXIT89, etc) are still actively referenced by many researchers within the community, and may be prevalent within certain sectors of industry.

3.4 Modelling of Crush

In general, each crush detection method that has been used to date can be classified into one of two generic groups; explicit methods and implicit methods. These two generic methodologies are outlined below, along with a brief discussion of their relative strengths and weaknesses.

3.4.1 Implicit

The implicit methodology is the original crush detection approach, and is still highly popular, being used in a large number of simulation models [79]. This methodology relies on the expert analysis of factors such as population

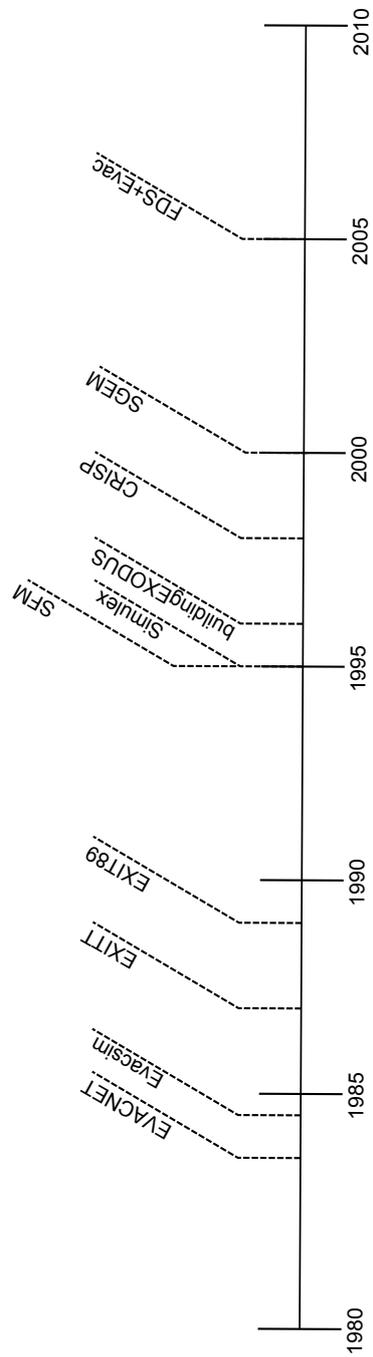


Figure 3.6: Time-line showing the release of evacuation models included in our survey.

density (see Section 2.5.1), behavioural analysis, and environmental considerations. The analysis of conditions within these models, therefore, is left to the engineer, who interprets the output of the simulation to determine whether crush conditions have occurred.

Implicit modelling does not take into account the possibility that evacuees exhibit any competitive egress behaviours, such as pushing, as there is no accurate method for simulating these behaviours without the inclusion of force calculations. This makes it best suited for general evacuation simulations; i.e. timely evacuations under “ideal” conditions.

As the exact force being exerted upon individuals is never calculated, the precise physical danger that exists in the evacuation can never be quantified. The only assertion that can be made, based on an implicit analysis, is that crush conditions are likely to form during the evacuation in question. The benefit of this approach is that, as the physical force calculation are not performed, it requires far less processing power than other methods.

Among many factor’s that may be considered during an implicit analysis are crowd density and flow rate. Measurement of crowd density can provide a good indicator of the danger present from crush during a simulation but it can only be confidently employed as an indicator that crush conditions are not present. If we take sports stadia as an example, the density of persons at post-war football matches in the UK has been empirically observed to have reached at least 10 persons per square metre under relatively standard conditions [29], at which point there were still no injuries recorded to the persons in question. A similar example can be found during the Saudi Arabian Hajj, where it is not uncommon to find similar densities of 10 persons per square metre. This is not to say that these situations should be considered “safe”, far from it, but the mere presence of high density crowds does not in itself signify the presence of dangerous crush conditions within a crowd.

There are too many implementations of the implicit methodology to list here but a popular, well documented example is Simulex [139], from Crowd Dynamics Ltd.

3.4.2 Explicit

The explicit modelling of crush conditions incorporates an assessment of crush into the model itself, and therefore requires less user analysis than the

implicit approach. Often based on the calculation of Newtonian force values, and generally operating in two-dimensional space, explicit methodologies can be used to detect the presence of crush conditions much more precisely than is possible with implicit modelling techniques. By simulating the exact forces being exerted by each individual, and enabling the propagation of forces throughout a crowd, the explicit methodology can be used to measure the exact amount of force that any individual is subject to. This, therefore, offers the possibility of *quantifying* the dangers that individuals face, which is not possible using the implicit modelling techniques.

Whilst the explicit methodologies offer an accurate measure of the forces acting within a crowd, the calculations needed to measure force require much more processing power than an implicit implementation, so there exists a definite trade-off between the two techniques.

The most well-known implementation of this methodology is the Social Forces Model [51] (see Chapter 4), which combines the force equations mentioned above with the modelling of the social forces acting within crowds. Although the original Social Forces Model was created as a research exercise, rather than a full-featured simulation environment, the model has appeared in many variations since its first appearance [110, 113], and has recently been incorporated into the FDS environment [73], a popular model of fire, smoke, and pedestrian flow.

3.5 Difficulties in Modelling Injuries Caused by Force

The final factor relating to the modelling of crush conditions is the difficulties that modellers and engineers have found in establishing exactly how much physical force is required before it becomes apparent that crush has occurred. There are many different metrics that are used to quantify the force affecting a single evacuee, or a cluster of evacuees, none of which have proven faultless. The evidence of the amount of force suffered by some individuals during the Hillsborough stadium disaster for example, suggests that forces within the crowd had breached levels of $4400N$ [29] but it is unknown what, if any, direct injuries (e.g. the breaking or fracturing of bones) can be caused by such force. Moreover, it is not known exactly how prolonged an exposure to this level of force would need to be to result in the presence of suffocation

injuries (e.g. compressive or traumatic asphyxia).

Different environments, models and researchers state their assertions on the levels of force required for injuries of varying levels. Fruin [37] suggests that physical injuries begin to occur at levels approaching $1500N$ suffered by an individual, whereas Helbing *et al* use the more malleable metric of a radial force acting about the circumference of an individual exceeding $400Nm^{-1}$. Conversely, the more implicit approach taken by the International Maritime Organisation (IMO) [58] is to measure the danger present in any situation using a relative metric associating the time-span of the evacuation to the crowd density at all times, and therefore states that an evacuation must be considered unsafe if the crowd density exceeds $4pm^{-1}$ for 10% of the overall evacuation time.

As we can see from these three examples, there still exists a great deal of uncertainty within the community on exactly how to metricise the dangers to individual crowd members that crush conditions cause. Notably, the Fire Dynamics Simulator (FDS), which includes injury behaviour from fire, smoke, and toxin exposure, does not include any criteria by which an individual may succumb to crush related injuries, and instead the force that individuals are subject to is presented to the user for later analysis.

Considering the difficulties in establishing the presence of injuries due to the formation of crush, it has been decided that during this work we will not use injury itself as a primary identifying factor of the relative danger of an evacuation. We instead opt for the measurement of either the maximum amount of force that any individual in an evacuation is placed under at any one time, or an average (arithmetic mean) measure of the force across the population at any one time. It is hoped that circumventing the actual classification of crush injury during this work will lead to a more robust and straightforward analytical methodology, that can be applied regardless of the exact manner in which injury metrics are carried out in the future.

3.6 Our Proposed Approach

We propose a new approach to the identification of crush conditions within evacuation scenarios. By identifying the underlying behavioural patterns that lead an evacuation towards the formation of crush, we aim to ascertain the relative likelihood of crush forming in a given evacuation.

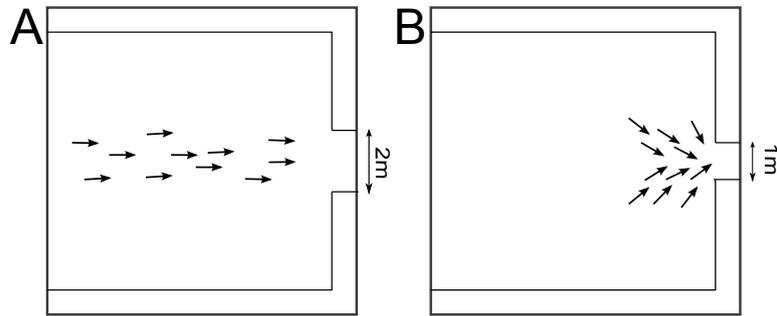


Figure 3.7: Figure A shows *laminar flow*, where the exit capacity appears sufficient for egress. Figure B shows the reduction of exit capacity inducing a transition into *turbulent flow*.

The behaviours within an evacuation that we would like to identify mark the transition between laminar flow (smooth and uninterrupted) and turbulent flow (haphazard and disjointed movement of evacuees), examples of which may be seen in Figure 3.7.

At its simplest, this technique will require the identification of large-scale, non-adaptive evacuation behaviours from the output of a simulation environment, which would suggest presence of high levels of physical force. The suggested approach applies mutual information (MI), a statistical measure of the interdependence of two signals (see Chapter 7), to the movement data taken from an evacuation simulation. We suggest that the MI of the system can then be used as a quantification of the amount of order within that evacuation, i.e. high MI suggests a well ordered laminar flow of pedestrians, whereas low MI suggests the breakdown of order and the onset of crush formation.

3.7 Summary

In this Chapter we described and summarised the details of two relevant methodological choices made in evacuation models, these being the representation of the physical environment and the choice between behavioural modelling and movement modelling. It should be noted that these classifications were designed for their relevance to the formation of crush conditions. In general, the most popular CA models include a greater depth of behavioural model than that found in continuum models, which tend to favour the movement modelling approach which allows very accurate contin-

uous movement. Our proposed technique enables the measurement of crush conditions via the analysis of movement data, available from both CA and continuum models. In this respect the technique can allow the inclusion of a crush measurement system within a CA model, allowing a model that runs at speed to have an indicator of crush that was previously only possible in the more computationally expensive continuum models.

We have seen that there has been a growing interest in the field of evacuation modelling over the past twenty years, and it is currently a highly active field of research.

During the next two chapters we outline both the test bed for our chosen technique, the social forces model (Chapter 4), and the mathematical basis of the proposed technique, MI (Chapter 5).

Chapter 4

Social Forces Model

4.1 Introduction

In this section we describe the history and working of our chosen test model for the crush detection technique, the Social Forces Model [51, 53], and examine its strengths and weaknesses in relation to both the modelling of evacuation behaviour and the calculation of physical force. The model is examined in depth as it not only forms the *test bed* for our technique, where we show that MI is suitable for the measurement of force within a simulation, but it was also used as the mathematical basis of the pedestrian movement within the evacuation simulation which we employ later in the project (see Chapter 8).

During this section we refer to two different versions of the Social Forces Model. The first was presented by Helbing *et al* in 1995, and did not model the physical force between pedestrians during an evacuation [53]. This will always be referred to as the *original social forces model*, or *original SFM*, in this Chapter. The second model was presented in 2000 [51] and extended the original SFM to include the calculation of physical forces, and it is this model which will be referring to as the *social forces model* or *SFM*.

4.2 Background

The original social forces model was introduced by Helbing *et al* in 1995 [53], and initially did not include the ability to calculate the physical forces between pedestrians which this project is most concerned with. The *forces*

referred to in the title of the model were therefore strictly *social forces*, which were defined as the sociological forces which drive pedestrians to both remain at a *reasonable* distance away from others (retaining personal space), and also to remain within a reasonable proximity of them (increasing the chances of successful exit/evacuation from a building). The other forces included within the original SFM are the psychological forces which repel pedestrians to a *comfortable* distance away from walls and boundaries (again, to retain personal space), and the driving force which attracts pedestrians toward their desired goal (usually the exit to the structure). The use of these forces within the original SFM enabled the model to recreate many empirically observed evacuation phenomena that are known to exist in real-life evacuations, such as natural pedestrian lane formation and the *faster-is-slower-effect*.

The updated SFM, hereby referred to as simply the SFM, was presented by Helbing *et al* in 2000 [51], included the addition of force and friction effects to the original model. This updated version was able to recreate additional phenomena observed in the field, including *arching* behaviour at exits (see Figure 4.5) in which the friction between pedestrians causes a solid arch-like structure about an exit, and *exiting bursts* in which the outflow of pedestrians turns from a uniform flow into a more turbulent and uneven exiting pattern under higher densities.

4.3 Description

The SFM operates under simple to understand principles. A pedestrian’s movement at every time step is an accumulation of simple decisions regarding their current situation and their desired goals. In short, a pedestrian will move away from other pedestrians or obstacles if they are “too close”, and will move towards the exit. In this way a pedestrian is able to avoid walls and obstacles whilst moving closer to their desired goal at every time-step, thus effecting egress.

The “urgency level” in the social forces model, i.e. the level of the perceived threat, is controlled by a parameter which defines the pedestrian’s desired escape velocity. In simple terms this parameter defines the walking/running speed that a pedestrian would attempt to reach were there no obstacles, such as walls, pillars, or other pedestrians blocking their path. As

will be shown by the mathematical definition of the model, the “desire to leave” is integral to the working of the SFM, as it will effect not just the speed at which a pedestrian will travel, but also effects the distance that they will keep from other pedestrians. In this way, a high desire to leave, combined with an insufficient exit capacity, will lead to high competition for the exit, increased crowd densities, and the build-up of physical force within the crowd. It is by increasing the value of a pedestrian’s desire to leave that crush conditions within the social forces model may be simulated.

An example of how changing the pedestrians’ desire to leave within a simulation will effect the crowding behaviour exhibited by the pedestrian may be seen in Figure 4.7.

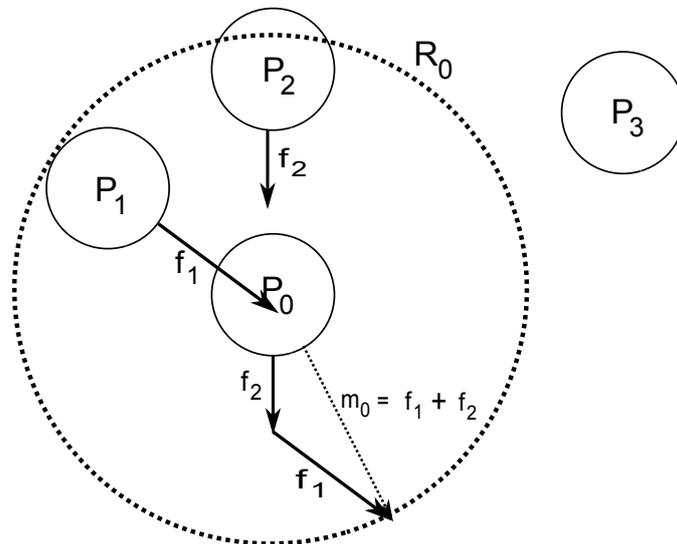


Figure 4.1: Visualisation of interaction between four pedestrians. Forces f_1 and f_2 show the social force exerted on p_0 by p_1 and p_2 respectively. Pedestrian p_3 does not contribute to movement, as they lie outside the interaction area R_0 of p_0 . Vector m_0 shows the final movement of p_0 , away from p_1 and p_2 and a rate inversely proportional to their proximity.

4.4 Mathematical Definition

The Social Force Model consists of five interacting forces, these being; repulsive social force, repulsive boundary force, repulsive psychological force, body force, and sliding friction force. The equations governing the forces within the model are defined below.

4.4.1 Movement

$$m \frac{dv_i(t)}{dt} = m \frac{v_0 e_i - v_i(t)}{\tau} + \sum_{j \neq i} f_{ij} + \sum_W f_{iW} \quad (4.1)$$

Where:

- \mathbf{v} , and \mathbf{m} represent the velocity, and mass of pedestrian \mathbf{i}
- \mathbf{v}_0 is desired velocity of \mathbf{i} , also used here to assign the level of desire, or *urgency*, the pedestrian has to exit the structure.
- \mathbf{e}_i is the direction in which the pedestrian would, if unimpeded, desire to travel. In the case of an evacuation this is likely to be the direction to the closest exit.
- τ is the *relaxation* parameter. This governs the acceleration and deceleration of pedestrians, and acts as a form of damping force to prevent extraordinary movements.
- \mathbf{f}_{ij} is the repulsive social force acting between pedestrian \mathbf{i} and \mathbf{j} , or between pedestrian \mathbf{i} and obstacle \mathbf{j} . In the original experiments the obstacles were circular pillars, analogous to the representation of the evacuees, hence including the repulsive obstacle force in the person-person interaction force equation.
- \mathbf{f}_{iW} is the repulsive boundary force of boundary \mathbf{W} acting on pedestrian \mathbf{i} .

4.4.2 Social Force

The movement that the SFM produces in an evacuee is therefore a summation of the vectors returned by the functions e_i , f_{ij} , and f_{iW} , with the value of e_i being a vector pointing toward the nearest exit.

$$f_{ij} = [Ae^{(r_{ij}-d_{ij})^{-B}+kg(r_{ij}-d_{ij})} + \kappa g(r_{ij} - d_{ij})]n_{ij}\Delta v_{ji}t_{ij} \quad (4.2)$$

Where:

- \mathbf{A} , \mathbf{B} , \mathbf{k} , and κ are mathematical constants.

- \mathbf{d}_{ij} is distance between pedestrian \mathbf{i} and \mathbf{j} .
- \mathbf{r}_{ij} is the sum of the radii of pedestrian \mathbf{i} and pedestrian \mathbf{j} $\therefore \mathbf{r}_{ij} = \mathbf{r}_i + \mathbf{r}_j$
- $\mathbf{g}(\mathbf{x})$ is a function which returns zero if $\mathbf{r}_{ij} < \mathbf{d}_{ij}$ (i.e. pedestrians have no physical contact) returns \mathbf{x} otherwise.
- \mathbf{n}_{ji} is the normalised 2-dimensional vector pointing from pedestrian \mathbf{j} to pedestrian \mathbf{i} .
- \mathbf{t}_{ij} is the tangential direction vector between pedestrian \mathbf{i} and pedestrian \mathbf{j}
- $\Delta \mathbf{v}_{ji}$ is the tangential velocity difference between the two pedestrians
- $\mathbf{k}\mathbf{g}(\mathbf{r}_{ij} - \mathbf{d}_{ij})$ represents the body force, which counteracts the pressure placed on a pedestrian's body. Scaling the constant \mathbf{k} effects the exact level of body force present in specific simulations.
- $\kappa\mathbf{g}(\mathbf{r}_{ij} - \mathbf{d}_{ij})\Delta \mathbf{v}_{ji}\mathbf{t}_{ij}$ represents the friction force between \mathbf{i} and \mathbf{j} , which impedes the tangential motion of pedestrian \mathbf{i} according to the magnitude of κ .

The social force equation governs the *desire* which an agent has to create space between themselves and any other agents within their interaction radius R . The output from f_{ij} , across the range $0 \geq R \leq 2$ (the default interaction radius within the SFM), is displayed in Figure 4.2.

At values of $r_{ij} - d_{ij} < 0$, the friction co-efficient κ begins to effect agent movement. When the function $g(r_{ij} - d_{ij})$ returns a figure lower than zero, the agent's movement begins to be hampered by the tangential friction acting between pedestrians i and j , which restricts movement in the x and y plane. The default value for the friction co-efficient from the original model is set at $\kappa = 3000$. Figure 4.3 shows the effect of the friction parameter κ when a pedestrian's proximity to another pedestrian is less than zero, i.e. the formation of crush occurs.

The effective *overlapping* of agents' physical representations that is required for the friction co-efficient to begin to hamper movement has been questioned by some (see Section 4.6.1), as the amount of additional *compression* suffered by the evacuees has not been either strictly defined nor empirically tested. There are extensions to the original model which effec-

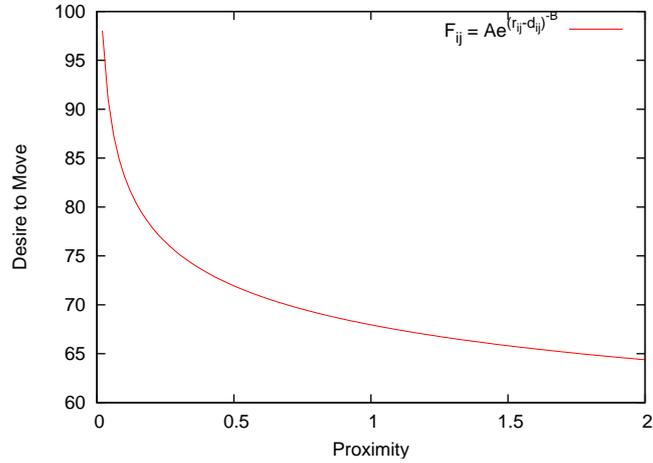


Figure 4.2: The changing desire of agent i to decrease their proximity to agent j , shown across the entire interaction radius ($0m \geq R \leq 2m$).

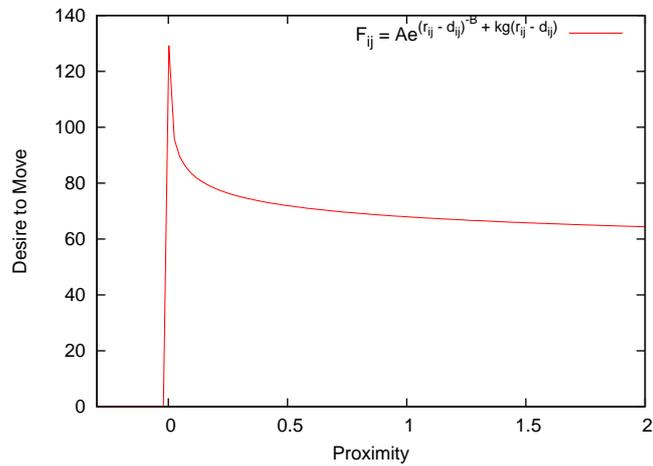


Figure 4.3: The effect of the force co-efficient κ on agent movement at very high density. At proximities less than zero, the agents body representations are effectively overlapping.

tively remove this behaviour and replace it with a more traditional collision detection method which prevents the overlapping of agent forms [81].

4.4.3 Boundary Force

The equation modelling the interactions between pedestrians and boundaries is similar to the equation governing the social force (Equation 4.2), and contains many of the same component variables.

$$f_{iW} = [Ae^{(r_i - d_{iW})^{-B} + kg(r_i - d_{iW})} + \kappa g(r_i - d_{iW})]n_{iW} - \kappa g(r_i - d_{iW})(v_i \cdot t_{iW})t_{iW} \quad (4.3)$$

Where:

- \mathbf{n}_{iW} is the perpendicular vector from pedestrian \mathbf{i} and boundary \mathbf{W} .
- \mathbf{t}_{iW} is the tangential direction vector between pedestrian i and boundary W .

4.4.4 Goal Finding

The goal finding behaviour within the SFM is governed by an equation returning the unit vector pointing directly at the agents' desired goal.

$$e_i = \frac{x_i^0 - x_i(t)}{\|x_i^0 - x_i(t)\|} \quad (4.4)$$

Where:

- \mathbf{x}_i^0 is the position of the desired goal.
- $\mathbf{x}_i(\mathbf{t})$ is the position of agent \mathbf{i} at time \mathbf{t} .

4.5 Visual Example

To understand the operation of the SFM, it is often best to view the interacting forces graphically. The equations that model movement in the SFM

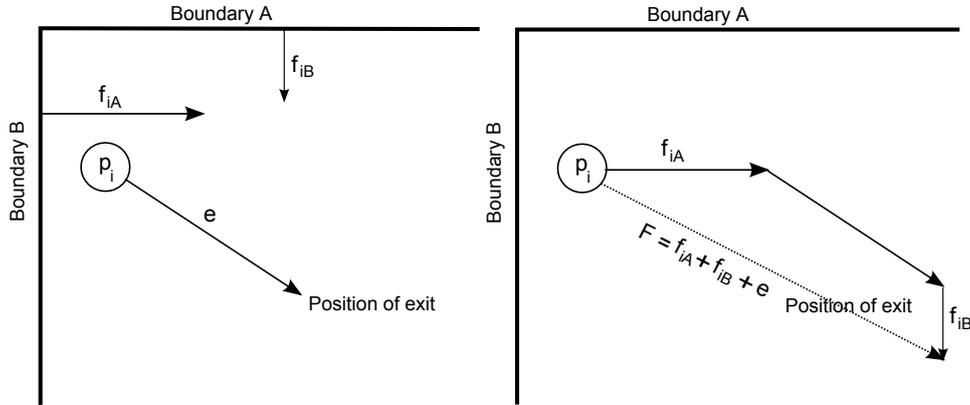


Figure 4.4: Graphical example of movement vectors within the SFM. Figure One shows the component vector leading an agent away from walls and towards the goal, Figure Two shows the resultant movement vector.

return vectors which either repulse (guide away from) or attract (pull toward) certain points of the evacuation space. It is the summation of these vectors which determines the direction in which the agent will travel at each time-step, and the acceleration equation which determines at exactly what speed.

The simplest visual example of the operation of the SFM's movement model can be seen during the unimpeded travel of a pedestrian. At low density the main forces which will act on the pedestrian are the repellent boundary forces (Equation 4.3) which act to steer the pedestrian a comfortable distance away from walls, and the attractive force of the pedestrian's desired goal (Equation 4.4). Under these conditions a pedestrian will *steer* away from walls, whilst making steady progress toward their goal.

We can see from Figure 4.4 that the result of force additions from F_iW and e give the direction of travel of the agent at the next time-step, t_{n+1} . The strength of repulsion is inversely proportional to the distance between the agent and the boundary in question, e.g. repulsion from boundary A is significantly greater than that of boundary B, which roughly models the priorities of the agent with consideration of personal space. The vector e draws the agent toward the exit, ensuring that movement is never severely detrimental to the progress of the agent toward their desired goal.

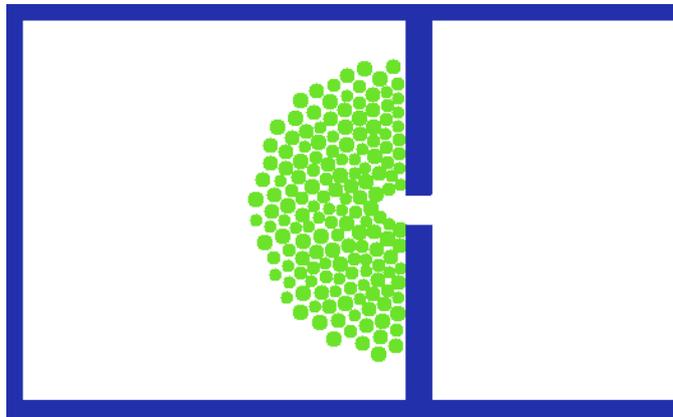


Figure 4.5: The arching behaviour about the exit can clearly be seen during this simulation.

4.6 Critical Analysis

The Helbing model is highly regarded as a microscopic model of pedestrian movement, and has been shown to recreate many of the behaviours present during evacuations, such as *exiting bursts* and the *faster-is-slower* effect [51, 53]. The SFM, or variations of the classic model, are currently being incorporated into many environments, and it has been shown to provide accurate results across a number of evacuation scenarios.

Most of the strengths of the SFM stem from the inclusion of force and friction within the simulation. That said, the original social forces model [53] was notable for its recreation of numerous evacuation behaviours despite the fact that it included no *physical* force calculations. The version of the SFM which we utilise as our test bed has numerous strengths above and beyond the ability to quantify the level of physical force, a number of which are defined below.

The inclusion of inter-personal friction within the SFM allowed the model to recreate the observed phenomena of *arching*, in which the evacuating crowd cause an arch-like structure to form in front of the main evacuation route (Figure 4.5), which has only been accurately recreated, in an *in silico* environment, by including inter-personal friction forces.

We can see that the pedestrians have formed a semi-circle around the exit, which is held tight by the friction acting between their bodies, in the same way that stone arches can be held upright due to the friction acting be-

tween blocks. Some cellular automata (CA) models recreate this behaviour by employing measures such as *friction probability*, which disables a pedestrian from moving to a new position if another pedestrian is trying to move into that position at the same time by implementing a probabilistic conflict determination [150]. This technique has shown results which are qualitatively similar to those achieved by the SFM, but requires correct parametrisation to operate, unlike the SFM which relies on strict motion equations to calculate the resolution of conflicts.

The counterpoint to the arching phenomenon is that of exit bursts, which are often observable immediately after the natural breakdown of an “arch”. At the point at which a single evacuee manages to break free of an arch, the pressure placed on those immediately surrounding them is released, and a number of other evacuees are often able to clear the congested area before the pressure becomes too great and a second arch is formed in place of the first.

This said, the usefulness of particle simulations in general is being questioned by some, and there are specific areas in which the SFM has been found to be lacking. A critical analysis of the SFM follows, in which a number of the limitations of the SFM are addressed.

4.6.1 Limitations

Despite the strengths of the SFM, it was originally built as a *toy model* (i.e. as a tool for investigation rather than for commercial use), and was designed as a means to investigate high density pedestrian movement, interaction and the formation of physical force, rather than as a means to simulate the evacuation of a building. As such, there are many areas in which the model could be improved. Below we list the areas identified by ourselves and others which are considered weaknesses of the SFM, and look at the improvements currently being made by researchers in the field to update the model and to expand its use.

Representation of Agents

The traditional SFM represented agents with a perfectly spherical two-dimensional shape, accounting for the body mass of the pedestrian. This reduces the maximum observable density within the simulation, by effectively

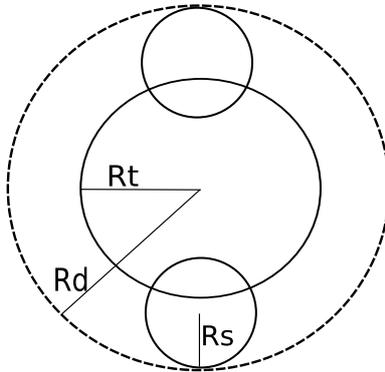


Figure 4.6: Body representation similar to that suggested by Thompson and Marchant [140], where R_t and R_s represent the radius of the torso and shoulders respectively, whilst R_d represents the radius of the body at its greatest point.

reducing the number of agents which could conceivably be compressed into a space of $1m^2$. A more accurate representation of agents was suggested by Fruin [37], later expanded by Thompson and Marchant [140], and utilises three overlapping circles to better represent the torso and shoulders of a pedestrian, see Figure 4.6.

At each shoulder, the radius of the agent is R_d , but at the torso section the body radius is R_t , significantly smaller. This elliptical footprint allows simulations to reach much greater densities than is possible using the circular representation. Our simulations with the Fire Dynamics Simulator (see Chapter 8) have shown that the elliptical shape allows crowd density to reach a maximum of $10pm^{-2}$, whereas the maximum density measured in the SFM was only $6pm^{-2}$.

This configuration gives the overall representation of an agent a roughly elliptical footprint, which it is argued can offer both more accurate visualisations of crowd movement and more accurate simulation results when compared to empirical data [81].

Lack of a Behavioural Model

As has been stated previously, the SFM acts as an excellent descriptor of human *movement* under certain environmental conditions. However it cannot be said that the SFM accurately models human behaviour or decision making in the same manner as other simulation environments (see Section 3.2.2).

Movement in the SFM is governed by attractive and repulsive forces acting against each other, and in this way the model considers only the minimum number of factors with which an evacuation can be represented. The complexity of human interactions within their environment, and with other human beings, makes the evacuation process itself a highly complex and intricate task. A full behavioural model would include tasks undertaken during the preparation to evacuate e.g. gathering knowledge from the immediate environment (some of this knowledge may be correct, but this is not necessarily the case), assessing the most immediate factors relevant to successful egress, judging the expected time-scales necessary to complete certain tasks, the sharing of information with other evacuees, and the evaluation of past experiences of such events, before making decisions on their best course of action. The social forces model only accounts for the most immediate threats with the environment, so cannot be considered a behavioural model in the true sense of the term.

Required Damping Force

The presence of a damping parameter is standard across many physical simulations. A damping force allows a body to settle into a state of rest, and reduces the tendency of a system to oscillate. Its use in evacuation simulations allows better modelling of the *absorption* of force at the point of contact between two pedestrians, and has been shown to reduce the effects of modelling humans as solid structures rather than deformable bodies. The work of Langston [81] introduced the use of a damping parameter to models similar to the SFM, and experimental results show that the addition has a positive effect on the modelling of evacuation.

The FDS simulation environment, which we will use for later testing, does include a damping parameter when implementing the SFM's physical force equation, as can be seen below

$$f_{ij}^c = (k(d_{ij} - r_{ij}) + C_d \delta v_{ij}^n) n_{ij} + \kappa(d_{ij} - r_{ij}) \delta v_{ij}^t t_{ij} \quad (4.5)$$

The size of the damping parameter C_d is implementation specific, but the FDS default value is set at 500Kgs^{-1} . Note that damping is usually measured in Newton-seconds per meter (Nsm^{-1}), which sets the FDS pa-

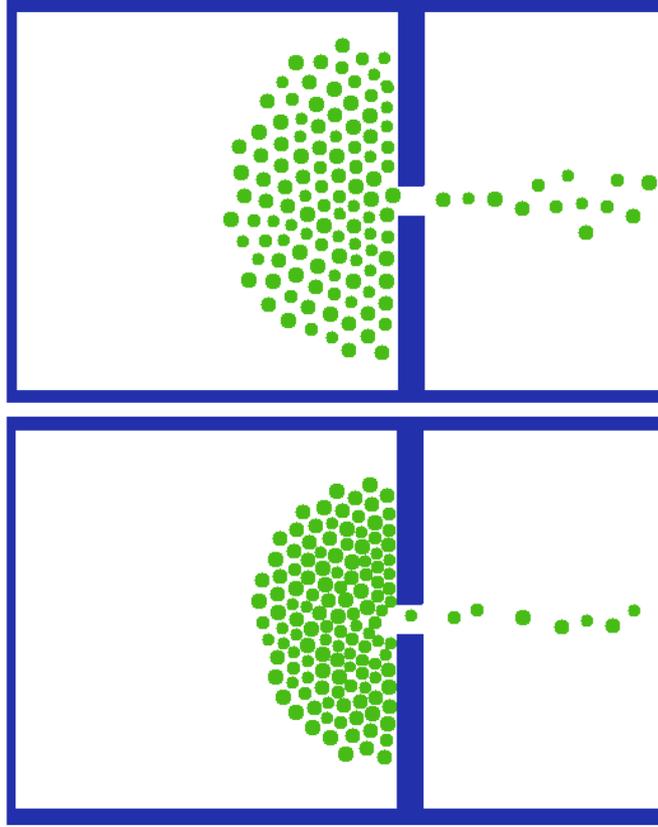


Figure 4.7: SFM agent distributions at differing values of desired velocity (V_0). Top: $V_0 = 0.5ms^{-1}$. Bottom: $V_0 = 5.0ms^{-1}$

parameter at a value of $C_d \approx 4905Nms^{-1}$.

Operation Under Non-emergency Conditions

The SFM contains no defined *queueing behaviour*, which means that an agent cannot *wait* or allow someone ahead of them. This lack of queueing behaviour creates a slightly unrealistic behaviour when an evacuation is attempted in which there is a relatively low desire to leave. Rather than a queue forming for the exit, or even agents making the decision not to progress toward the exit, the evacuation will take place in exactly the same manner as an emergency evacuation. This means that the behaviour in the SFM at emergency and non-emergency conditions are qualitatively similar, as seen in Figure 4.7.

As we can see from these figures, the distribution of agents at low veloci-

ties in the SFM is qualitatively similar to the distribution at higher velocities. This is to be expected, as the SFM implements exactly the same movement equations for any desired velocity, but it does mean that non-emergency usage of a structure cannot be modelled using the traditional SFM. An example of this is queueing behaviour, in which a system of pedestrians will organise themselves into informal queues during their egress from a building (or other analogous situation). During queueing behaviour, people will often exhibit behaviour such as low-level altruism (i.e. common courtesy) in allowing others to pass, and regularly use verbal and non-verbal cues to enforce this societal norm.

The SFM takes an approach more akin to *game theory* in its prescription of evacuation behaviour, in that at any one point a pedestrian will make the movement that minimises their net distance to their goal and maximises their distance to another pedestrian, wall, etc. In this respect, the SFM fails to recreate many basic behaviours that can be seen during non-emergency usage of structures, and we conclude that it is unfit for purpose were it to be applied to any non-emergency evacuation.

Physical Compression

The SFM allows a manner of body compression, but this compression is not strictly defined within the parameters of the model. If we look at Figure 4.3, we can see that when $d_{ij} - r_{ij} < 0$ the friction parameter begins to affect the movement of the pedestrians, but at these distances the compression of evacuees is more akin to overlap rather than the physical force of compressing a body. Whilst there has been some work into the modelling of compressive forces acting within crowds [18], which suggested that the addition of compressive force has beneficial results in respect of the modelling of evacuation, the strictly defined compression criteria detailed therein has not yet been included in a mature simulation environment. Some argue that if the overlapping or compression of pedestrians has not been strictly defined then it is wise to remove the ability for bodies to overlap from simulations entirely, or for it to be replaced with a more *rigid* collision detection method that will enforce a strict adherence to the agents' physical proportion [80].

Force Feedback

Under high density crowd conditions, it is known that individuals are often found to lose the ability to control their own actions, this can be found in the reports into the Hillsborough disaster [136], in which it was known that the pressure placed upon people in the crowd was enough to lift them off their feet. At this point, their movement was wholly controlled by the pressures around them, rather than by their own desired direction of movement. This is not modelled, or accounted for, in the SFM, as the friction effects of this type of situation are extremely high, and difficult to model in two-dimensions.

Criteria for Injury

The criteria by which an evacuee is considered to have succumbed to injury within the original work is very strict, and follows engineering guidelines which suggest that a pedestrian may become injured if the sum of the magnitudes of the radial forces acting on them, divided by their circumference, exceeds a pressure of $1600Nm^{-1}$ [124].

There is still a great deal of debate over the levels of force that are likely to cause crush related injuries (see Section 3.5). High levels of force are known to cause serious physical trauma and serve to incapacitate individuals unfortunate enough to be subject to them. The medical term for this is *traumatic asphyxia*. Likewise, lower levels of force that are applied for a prolonged period of time can lead to similar incapacitation due to suffocation. Reflecting this, the International Maritime Organisation (IMO) have, when setting their standard for a safe evacuation, included the dangers of prolonged exposure to lower amounts of force; stating that an evacuation may be considered unsafe if the local population density of the structure (ship) exceeds 4 persons per metre for over 10% of the overall evacuation time [58].

It can be said that the criteria for injury used in the SFM models the effects of sudden *traumatic asphyxia*, as the pedestrian is considered injured if the level of force they are subject to breaches a maximum limit for any amount of time but were a pedestrian subject to fractionally less force for a prolonged periods they will not suffer any injury, which is counter-intuitive.

This is not a specific failing of the SFM, as there is much still to learn

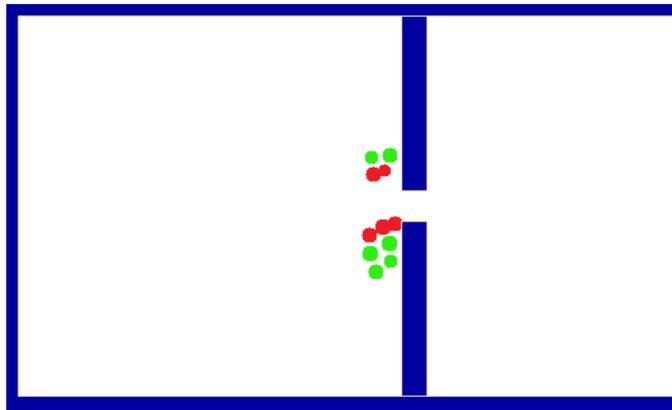


Figure 4.8: Barricading behaviour of the SFM. Injured pedestrians (Red) become immovable obstacles, which the active pedestrians (Green) are unable to circumvent. The system will remain in this state permanently once this behaviour has been established.

about the levels of force that cause death or serious injury, and these are further complicated by the uniqueness of individuals, i.e. height, weight, age, existing medical conditions. It is clear that there is still much to learn about the complex relationships between the conditions placed on the human body and the likelihood of these conditions leading to serious injury or death.

Behaviour of Injured Pedestrians

The behaviour of an injured pedestrian within the traditional SFM is another point of interest. At the point at which a pedestrian becomes injured, the pedestrian immediately becomes unable to affect movement, but also becomes a solid unmovable mass, remaining in the same position for the remainder of the simulation.

The effects of this injury behaviour on the remaining evacuees can be seen in Figure 4.8, which shows the injured pedestrians (red) blocking the route to the exit for the remaining pedestrians in the room (green). The exact simulation from which this graphic was taken was carried out using the SFM default parameters for all variables except the *desired velocity*, which was set to a value of $5ms^{-1}$.

4.7 Time-lapsed Visual Example

A visual example of a single run of the social forces model may be found in Figure 4.9. Images were recorded at intervals of $0.1s$ throughout an entire simulation, and relevant frames were then selected by hand. Experimental parameters used were those presented in the original literature [51] and the desired escape velocity, V_0 , was set to $V_0 = 5ms^{-1}$, as this is the first point at which injuries are found to occur. In addition to this the injury criteria of the SFM has been switched off during this simulation, so agents can continue to function regardless of the amount of force that they are subject to.

4.8 Differences in the FDS+Evac Implementation

Later in this thesis the Fire Dynamics Simulator is used to recreate a historic example of a crush disaster (see Chapter 8), more specifically it is the FDS+Evac module which is used to model human evacuation from a structure. The FDS+Evac module uses an implementation of the social forces model as a basis for evacuee movement, and whilst the implementation is similar to the original model there are a number of small differences. The most important deviations from the original model are discussed below, with reference to the material presented previously.

1. Agent Representation

As stated previously (see Section 4.6.1), the traditional SFM represented agents as a perfect circle, which sufficed as an initial approximation but is not a reasonable fit to the actual frame of a pedestrian. The FDS+Evac module models pedestrians in a slightly more accurate manner, by representing the body and shoulders as three connected circles. The FDS representation is more akin to that suggested by Thompson and Marchant (see Figure 4.6) than that employed by the original SFM, and the additional granularity allows more complex behaviours. An example of this is the pivoting behaviour that may happen if a person is being subject to a force on just one side of their body, forcing them turn about their centre and change the direction they face. The modelling of this behaviour was not possible in the original SFM, but can be seen in the FDS+Evac implementation.

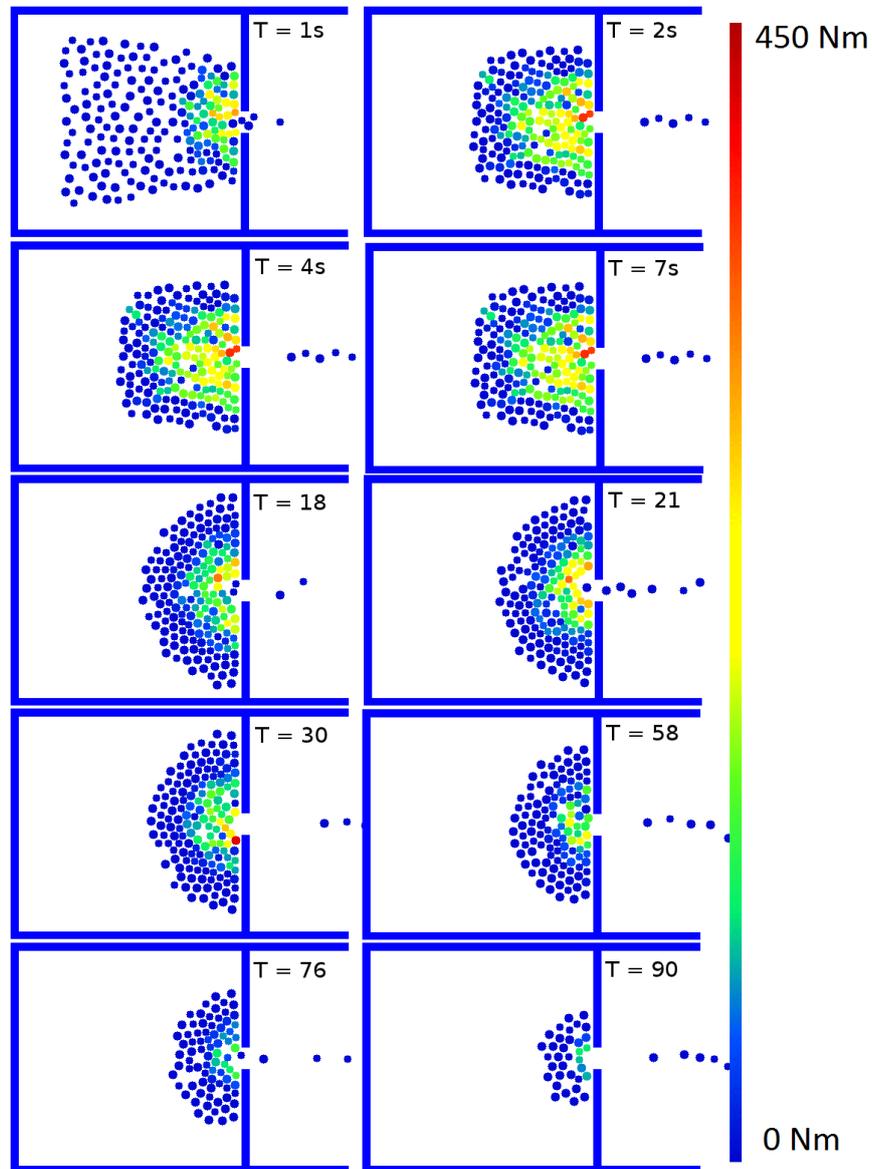


Figure 4.9: Example visual output at different times during a single evacuation using the social forces model, $V_0 = 5\text{m s}^{-1}$, $N = 200$. Force listed as per the traditional Helbing model, i.e. radial force about the specific circumference of each evacuee.

2. Inclusion of Collision Detection

The possibility of pedestrian overlap behaviour described previously (see Section 4.6.1) has been removed from the FDS+Evac movement model. In the traditional SFM it was possible, under some circumstances, for pedestrians bodies to “overlap”, that is two or more pedestrians could partially inhabit the exact same physical space at the same time. This behaviour is a result of the movement equations used to define the friction force between two pedestrians, but the behaviour was not based on empirical observations, i.e. it was not known that this type of compression was either possible or realistic in an evacuation scenario. The FDS+Evac model retains the calculation of friction force but does not allow the overlapping of a pedestrians physical representation, which is achieved through the inclusion of a simple collision detection model.

3. Removal of Injury Behaviour

As discussed in Section 4.6.1, the criteria for defining an pedestrian as having become “injured” due to their being subject to a physical force is still a question which remains to be answered. Rather than include an approximate definition of this, the FDS+Evac model relies on the output of the physical forces calculated during an evacuation simulation to metricise the possible danger that an evacuee may be subject to by the build-up of force.

4. Flow-based Directional Algorithm

An algorithmic change within the FDS+Evac environment occurs in the goal finding equation, which defines the direction that the pedestrian would travel in if entirely unimpeded (i.e. the direction of the exit). The traditional SFM makes use of a simple equation which will point an agent directly toward the nearest exit, which is calculated for each agent at each time-step. The FDS+Evac environment takes a flow field approach to this problem, which reduces overall computation time. As the FDS contains a comprehensive fluid dynamics engine, the

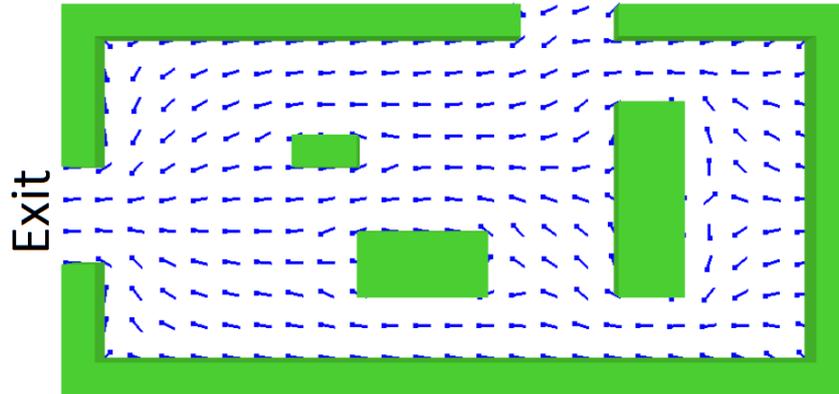


Figure 4.10: A graphical representation of an FDS+Evac “flow field” [72]. In this case all evacuees would head towards the left exit, using the flow field (blue arrows) to find their way towards this exit.

modelling of the underlying evacuation path is achieved by treating all exits as “vents”, and calculating the path to these exits from any point in the environment as the direction of flow at that one point.

Figure 4.10 shows an example of the flow field for a simple evacuation topology. The exit, marked on the left, is treated as a vent, and the blue arrows shown at all points within the room described the direction of travel toward that vent. The approach is common in cellular automata models (see Section 3.2.1), but not as often employed in continuous models. An additional advantage of this technique is that it is required to run just one time at the start of a simulation, rather than for every agent at every time-step.

4.9 Summary

The original implementation of the Social Forces Model was chosen as a test bed for this project for two main reasons. Firstly, the SFM has the option to calculate force, which is required as a measure of the physical danger that evacuees face in certain situations, and can be used to ensure the correct working of the analytical technique. Secondly, many variations of the SFM have been implemented, and versions of the model are incorporated into large scale evacuation simulations, which will aid the integration of our analytic technique in further models.

Despite the limitations listed previously (see Section 4.6.1) the model is seen as an ideal test bed for our initial experiment, but it is accepted that any further experiments must be carried out using a fully fledged simulation environment, as the SFM does not contain much of the functionality required to test our technique conclusively.

Chapter 5

Mutual Information

5.1 Introduction

In this Chapter we investigate the information theory techniques which we will be applying to the identification of crush conditions later in this work. The techniques are summarised and mathematically defined, and previous uses of these techniques across different areas of the sciences are investigated. We show that the techniques produced from initial research into information theory have applications far beyond this field, and provide robust and malleable statistical measures of order, interdependence and data complexity that can be applied to many and varied data-sets.

The Chapter concludes with a more in-depth discussion of the way in which we intend to apply these measures to the field of crush detection.

5.2 Background

The need to measure the specific amount of information held within a signal became imperative during the early days of telegraph and radio communication uptake. During the early 1900s, many researchers and engineers were attempting to quantify the information content of specific communications, and much of the initial literature is of this era. In 1928, Hartley published his paper *Transmission of Information* [49], in which he defined a general measure of the information content of a variable length message formed from a known alphabet of symbols. Taking a message of n symbols in length, consisting of symbols taken from an alphabet of size s , he sought

to define a measure of information content that increases linearly with n , it therefore follows that if a message were doubled in length it can be said to contain double the amount of information. Hartley arrived at the measure H , where

$$H = \log s^n = n \log s \quad (5.1)$$

This measure of information complies with Hartley's restriction that the information content of a message *must* scale linearly with the length of the message n . As the alphabet of possible symbols remains static, the logarithm of s can be considered a constant, with a value dependent on the number of symbols contained in the specific alphabet used. Therefore the metric that Hartley created may be shown to be equal to nK , where K is the alphabet specific information constant, and it can be seen that the measure H is linearly dependent on the message length n , for any fixed alphabet.

Hartley's measure of information content assumes no syntactic rules are present in the language, i.e. that the language possesses no grammar, and each symbol in the alphabet is assumed to contain inherent information when viewed abstractly. This type of communication can be quite common in some circumstances, such as in electronic communication systems, but the grammatical complexity of human language makes this technique of limited use in the measurement of information content of the written word. If we look at the English language, the frequency with which certain characters appear in text is far from uniform, with many letters appearing constantly in generic sections of text whilst others appear with relative infrequency (which is the reason that the points on Scrabble tiles are not uniformly distributed throughout the alphabet). If we consider this on top of the use of separate words in the language, and the syntactic rules that must be followed to allow the connection of these words correctly, we see that the complexities of information content in natural language cannot be accurately modelled by Hartley's formula. A measure of information content that relied upon more than the size of the alphabet was needed.

In 1948, Shannon introduced a measure which relied on the probabilistic appearance of each letter contained within the possible alphabet [119], and which accounted for (at least partially) the nuances of natural language.

Given the observed probabilities of the letters of our alphabet appearing are $p_1, p_2 \dots p_n$, the information content (H) of a given message may be described by

$$H = - \sum_i^n p_i \log p_i \quad (5.2)$$

This value gives us the *average* amount of information that may be taken from message of n symbols in length using the alphabet in question. It has become known as the *Shannon Entropy*, as there have since been numerous methods of calculation of Entropy defined (e.g. Gibbs Entropy [115], Boltzmann Entropy [60], and Renyi entropy [112]). All references to Entropy in this paper will refer to the Shannon Entropy described previously, unless explicitly stated as being otherwise.

Entropy can have multiple definitions but the method defined in Equation 5.2, being the quantification of the amount of information a message could contain when it is received, is that pertaining to the original design criteria for the calculation.

5.3 Joint Entropy

An extension of the Entropy measure is found in *Joint Entropy* [119], which is the measure of the joint information content of two related variables, signals, or events. If we take two discrete variables X and Y , we can calculate their Joint Entropy by considering the probabilities of each pair of possible outcomes. If the two variables have an observed probability distribution $P_{X,Y}$, then their Joint Entropy, $H(X, Y)$, is defined as:

$$H(X, Y) = - \sum_{x,y} p(x, y) \log p(x, y) \quad (5.3)$$

Where $p(x, y)$ is defined as the observed frequency of signal X taking the value x whilst Y takes the value y within the paired variable set.

5.3.1 Conditional Entropy

A related metric to the Joint Entropy is the Conditional Entropy of two signals, which is defined as the amount of “uncertainty” that remains about

variable two, once the value of variable one has been ascertained [119]. If we take two discrete variables X and Y , the conditional entropy is defined as:

$$H(Y | X) = - \sum_{x,y} p(x,y) \log \frac{p(x,y)}{p(x)} \quad (5.4)$$

This is simply the measure of the uncertainty that remains in Y after the value of X is known. The equivalence therefore exists that:

$$H(Y | X) \equiv H(X, Y) - H(X) \quad (5.5)$$

5.4 Mutual Information

Mutual Information (MI), is the final information metric that will be discussed here. It is a metric that quantifies the amount of information that two signals or variables share. Again defined by Shannon [119], it was presented in the same publication as his entropy findings, although named at a later date. It is a measure of the mutual dependence of the two variables, and can be thought of as a general measure of the reduction in uncertainty that is gained about one variable by knowing the value of the other.

If we take (again) our two discrete variables X and Y , the Mutual Information, $I(X, Y)$ is defined as:

$$I(X, Y) = \sum_{x,y} p(x,y) \log_n \frac{p(x,y)}{p(x)p(y)} \quad (5.6)$$

Where $p(x, y)$ is the joint probability distribution of X and Y , and $p(x)$ and $p(y)$ are the probability distributions of X and Y respectively.

The Mutual Information of two variables is *always* non-negative ($I \geq 0$), and is *only* zero when the two variables are entirely independent. Like joint entropy, MI is also commutative $\therefore I(X, Y) \equiv I(Y, X)$.

As Mutual Information is a measure of the amount of information shared by two signal or variables, the following equivalences with other entropy measures exist:

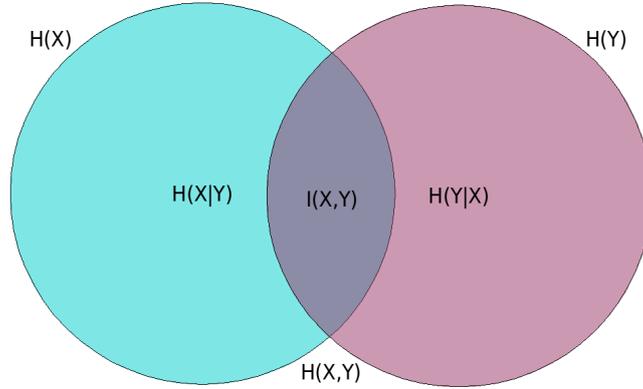


Figure 5.1: Venn diagram displaying a quantitative representation of the Shannon entropies, joint entropy, Conditional entropies, and the mutual information of a pair of variables X and Y .

$$\begin{aligned}
 I(X, Y) &\equiv H(X) - H(X | Y) \\
 &\equiv H(Y) - H(Y | X) \\
 &\equiv H(X, Y) - H(X | Y) - H(Y | X) \\
 &\equiv H(X) + H(Y) - H(X, Y)
 \end{aligned} \tag{5.7}$$

These equivalences can be better viewed, graphically, in Figure 5.1.

5.4.1 Conditional Mutual Information

A useful construct within information theory is the ability to express the mutual information of two random variables conditional on a third [129]. For three discrete variables, X , Y , and Z , where the value of variable Z is known, the conditional mutual information is defined as

$$I(X, Y | Z) = \sum_{i,j,k} P(z_k) P(x_i, y_j | z_k) \log \frac{P(x_i, y_j | z_k)}{P(x_i | z_k) P(y_j | z_k)} \tag{5.8}$$

$$= \sum_{i,j,k} P(x_i, y_j, z_k) \log \frac{P(z_k) P(x_i, y_j, z_k)}{P(x_i, z_k) P(y_j, z_k)} \tag{5.9}$$

5.4.2 Multi-variate Mutual Information

The logical extension of the bi-variate (two dimensional) MI, defined previously, is to extend the calculation to enable the measurement of information content shared by *more than two* signals simultaneously. This calculation of MI involving more than two variables has turned out to be far less trivial than it initially sounds, and subsequently there have been many different definitions of *Multi-variate Mutual Information* (MMI) over the past 50 years. The first definition, and most widely accepted, was that of McGill [94], who defined the MMI of three variables (A , B , and C) as:

$$I(A, B, C) = I(A, B) - I(A, B | C) \quad (5.10)$$

This equation also extends to higher dimensions, so for a total of n input variables $X_1 \dots X_n$

$$I(X_n) = I(X_1, \dots, X_{n-1}) - I(X_1, \dots, X_{n-1} | X_n) \quad (5.11)$$

The use of *Multi-variate Mutual Information* (MMI) as an analytical tool is not as widespread as the traditional MI calculation, despite being well defined and heavily documented in theoretical literature [23, 91, 94, 129, 143]. One hurdle to the application of MMI in real-world applications is found in its counter-intuitive tendency to offer a MI metric that can be *negative* ($I < 0$), which implies that the *uncertainty* of one input is actually *increased* when we have knowledge of the value of another related variable.

A good example of this type of negative results can be found in the instance of an *XOR* gate, with binary inputs X and Y , and a binary output Z , see Figure 5.2. The inputs and corresponding output from this gate may be seen in Table 5.1

In this example application, we see that the three variables must have dependence of some form, as they are logically related. If we analyse the information content shared between these variables, we find that both $I(Y, Z)$ and $I(X, Z)$ are zero as knowing either X or Y will tell us nothing about the corresponding value of Z , and $I(X, Y)$ is also zero as the two inputs are wholly independent. This tells us that our first term, according to Equation 5.10 must be zero, whatever the arrangement of the input variables,

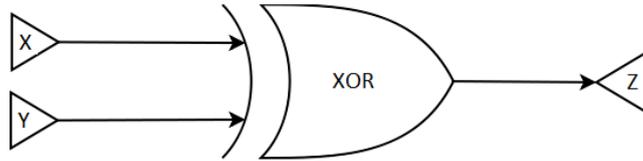


Figure 5.2: A dual input, single output XOR gate. X and Y are binary inputs, whilst Z forms the binary output. Table 5.1 contains the input and output data

X	Y	Z
0	0	0
0	1	1
1	0	1
1	1	0

Table 5.1: The inputs, X and Y , and the corresponding output, Z , of the dual input binary XOR gate, as shown in Figure 5.2

therefore $I(X, Y, Z) = 0 - I(X, Y|Z)$. The value of $I(X, Y|Z)$ will be positive, and in this case will have the value of *1bit*, as once Z is known, the value of Y is wholly dependent on the value of X , and vice versa. This leaves us with:

$$\begin{aligned}
 I(X, Y, Z) &= I(X, Y) - I(X, Y | Z) \\
 I(X, Y, Z) &= 0 - 1 \\
 I(X, Y, Z) &= -1bit
 \end{aligned}$$

It is this counter-intuitive inference, i.e. that the knowledge we have of one variable actually increases the uncertainty of the values of the other two, that makes MMI a difficult concept to grasp, and may have hindered its wider adoption as an analytical technique. The MMI technique is, in various forms, still employed in certain fields, most prominently in medical imaging (see Section 5.9).

5.4.3 The Logarithmic Base

In all formulae presented in this Chapter, the base of any logarithms involved is not explicitly stated. This is due to the fact that the logarithmic

x	$P(X = x)$	y	$P(Y = y)$
A	$\frac{9}{20}$	A	$\frac{1}{4}$
B	$\frac{3}{10}$	B	$\frac{1}{4}$
C	$\frac{1}{5}$	C	$\frac{1}{4}$
D	$\frac{1}{20}$	D	$\frac{1}{4}$

Table 5.2: Probability distributions for variables X and Y

		$P(Y)$			
		A	B	C	D
$P(X)$	A	$\frac{1}{20}$	$\frac{3}{20}$	0	$\frac{1}{4}$
	B	$\frac{1}{10}$	$\frac{1}{20}$	$\frac{3}{20}$	0
	C	$\frac{1}{10}$	0	$\frac{1}{10}$	0
	D	0	$\frac{1}{20}$	0	0

Table 5.3: Joint probability distributions for variables X and Y

base dictates only the units in which the mutual information or entropy is calculated. As these techniques originate in Information Theory the standard unit is the *bit* (base 2) but *nats* (base e), or *Hartleys* (base 10) are also widely used.

5.5 MI Example

The Mutual Information of two discrete, and short, signals is simple to calculate. There follows an example of how the mutual information of two discrete variables may be calculated. Given the discrete variables X and Y, where both X and Y are comprised of 20 values taken from an alphabet of no more than 4 possible values.

$X =$ A B B A C A D A C A B A A C A B C B A B
 $Y =$ D C A D A A B B C D A B B C D C A B D C

We can calculate the probability distributions of X and Y by hand, see Table 5.2. The joint probability distribution of X and Y is also listed in Table 5.3.

Using values from Table 5.2 we can calculate the Entropy of signals X and Y:

$$\begin{aligned}
H(X) &= - \sum_i P(X = x_i) \log_2(P(X = x_i)) \\
&= -\left(\frac{9}{20} \log_2\left(\frac{9}{20}\right) + \frac{3}{10} \log_2\left(\frac{3}{10}\right) + \frac{1}{5} \log_2\left(\frac{1}{5}\right) + \frac{1}{20} \log_2\left(\frac{1}{20}\right)\right) \\
&= 1.72 \text{ bits}
\end{aligned}$$

Calculating the entropy of signal Y, using the same method, gives us:

$$H(Y) = 2.00 \text{ bits}$$

The joint entropy of X and Y can now be calculated, using the probabilities given in Table 5.2, giving us:

$$\begin{aligned}
H(X, Y) &= - \sum_{i,j} P(x_i, y_j) \log_2 P(x_i, y_j) \\
&= 2.97 \text{ bits}
\end{aligned}$$

As we now have the values for $H(X)$, $H(Y)$, and $H(X, Y)$, we may use these to acquire the MI of the two signals:

$$\begin{aligned}
I(X, Y) &= H(X) + H(Y) - H(X, Y) \\
&= 1.72 + 2.00 - 2.97 \\
&= 0.75 \text{ bits}
\end{aligned}$$

As we can see from the result of the calculation, the Mutual Information of X and Y, in this case, is measured as 0.7540 *bits*. The same result can be reached by solving Equation 5.6 using the data provided in Table 5.2.

5.6 Information Theory

Initial uses of Mutual Information were mainly Information Theory based [92, 119, 120], and centred around its use as a measure of *real-world* communication channel transmission rate, i.e. the rate at which data is actually

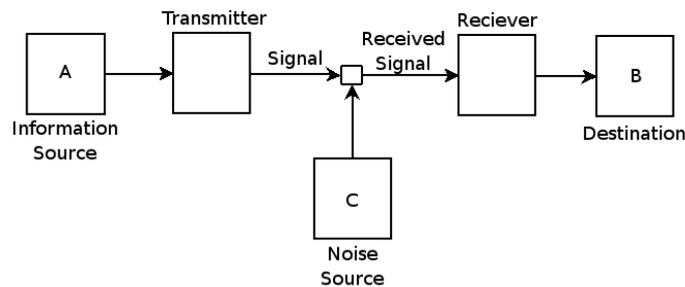


Figure 5.3: Schematic diagram of general communication system. Information A is transmitted across the channel, where it is modified by with the addition of noise (source C), the resulting information content received at point B will therefore only contain a subset of the original content ($B \subset A$).

transmitted across a channel, including error correction rates. This measurement is the intuitive use of MI, as it quantifies the information common to both signal A and signal B . If we were to consider that A were a transmitted signal and B a received signal, then the MI of the two signals, $I(A, B)$, must be a quantification of the amount of signal A that was transmitted without errors *and* the amount of signal B that was received without error. For a generic communication system, as depicted in Figure 5.3, the MI between the signal transmitted from A and received at B tends towards infinity for low noise, but tends towards zero for extremely high levels of noise. Note that the MI of A and B will equal zero ($I(A, B) = 0$) only if the noise introduced to the system was of such intensity that there was no part of A still present in B , i.e. A and B were statistically independent.

The use of MI within Information Theory was extended further, and is routinely used as a measure of the security offered by a particular cipher (encryption method). Shannon defined the perfect cipher as one in which the text A and the cipher-text B shared no information whatsoever, i.e. $I(A; B) = 0$ [92, 120]. This may seem counter intuitive, as a coded message *must* contain the original message, in at least some form, but nevertheless this is strictly the definition of the perfect cipher. If we analyse this claim from an information theory perspective, we see that it is far less bizarre. If the communication system in Figure 5.3 were to have the noise source C replaced with a cipher, and following this $I(A; B) = 0$, then the cipher applied at C is said to be perfect. In this way it can be said that the goal of any complex cipher is to reduce the MI of the source message and cipher

text to the lowest possible value.

Recent research has applied this analytical approach to the wholly opposite purpose, to break encryption systems used on embedded devices. The work of Kocher *et al* [71, 70] identified, and subsequently exploited, a new form of attack on embedded hardware known as “side-channel leakage analysis”. This form of attack revolves around the analysis of implementation specific, hardware based “leaked information”, information gained by analysis of secondary processes such as computation timings [70], physical power consumption [71] or electromagnetic emanation [111]. Analyses of the data obtained by observing these secondary sources under known conditions is analogous to the transceiver application described previously; the statistical analysis of measurements gained from one or more observed “noisy” physical variables. Gierlichs *et al* successfully applied both MI [40] and subsequently MMI [41] to the analysis of these readings, and to great effect. As the MI approach to statistical analysis requires no knowledge of the likely dependencies which exist between input and output variables, the application proved to be a “universal tool” able to be applied to any system in which the observations of secondary information sources can be reliably made.

5.7 Biological Stimuli-response Systems

The use of MI as a measure of the information transferred across a communication channel has been extended to other fields, most notable of which is the application to biological stimuli-response systems [30, 69, 89, 107, 130]. The application to these types of systems has been aided by the overall similarity in approach to the original use of MI, with areas such as neural coding being wholly analogous to communication channels described by Shannon *et al* in the original literature on the field.

The MI of the stimulus (S) and the response (R) in many biological models involves specific timing issues which can often be disregarded in other applications. The response to certain stimuli has been found to exhibit at a later time than the stimuli is applied. This fact may often appear trivial, with the “lag” between stimuli and response being measured in milliseconds, but when dealing with stimuli which change dynamically, and with a high temporal resolution, this can cause serious issues in data analysis. To combat this, the use of a *time delay* during MI calculations is introduced to “align”

the stimuli and response signals, thus reducing the effect of a constant or measurable time delay within the system. There are many ways in which the exact time delay required for each application can be calculated, which is out of the scope of this review, but were we given a system of stimuli S and response R with a time delay of δt we can calculate the MI of the system as follows:

$$I(S, R) = I(S^t, R^{\delta t}) = \sum_{i,j} p(s_i^t, r_j^{\delta t}) \log \frac{p(s_i^t, r_j^{\delta t})}{p(s_i^t)p(r_j^{\delta t})} \quad (5.12)$$

This simple shifting (effectively a backwards shifting of the response variables) cancels any standardised time-lag behaviour that inhibits the operation of the MI analysis.

The MI calculation itself can be used to calculate the optimum time lag, in a similar manner to that of a cross-correlation analysis. One signal can be set as a *base signal*, and set for reference, whilst the other is progressively incremented forward or backwards along the time-line by changing the value of δt , until $I(S; R)$ has reached a maxima. At this point it can be assumed that the signals are aligned, yet this does not necessarily yield optimum alignment. This alignment technique is more suited for complex data than some traditional approaches, such as cross-correlation, as when dealing with a response that may only have been hypothetically defined, the MI alignment will optimise regardless of the type of correlation present in the data.

5.8 Chaotic Systems

The use of MI for the analysis of chaotic systems is an area which has garnered much interest in recent years and has, in reciprocity, added greatly to the advance of the information theory techniques which inspired their analysis. A classic paper in the field is by Fraser and Swinney [34] which detailed the application of MI to the analysis of Belousov-Zhabotinskii reaction, the application of MI to the measurement of a system's time delay (or lag), and an algorithm to maximise analytical accuracy by the dynamic partitioning of the MI calculation space. This latter algorithm is known as the *Fraser-Swinney algorithm* [16, 34], and has been shown to produce good results across the systems tested.

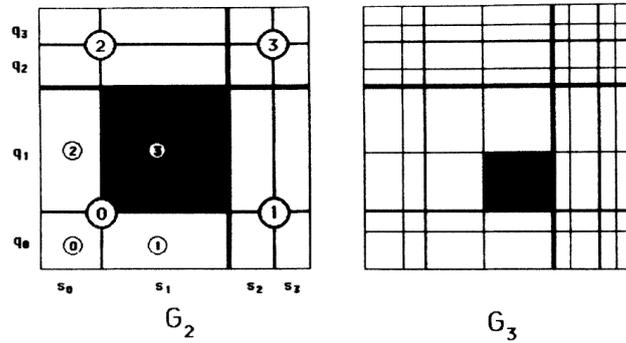


Figure 5.4: Visual example of Fraser-Swinney recursive subdivision. Using this method the calculation space may be discretised according to the amount of information they contain. Referencing is achieved by *tree notation*, therefore R_x is referenced as $R(0, 3, 1)$.

The Fraser-Swinney algorithm effectively operates by recursively subdividing the calculation space (Cartesian space, in the case of this work) into progressively finer subsections, dependent on the amount of information that it contains. This process of recursive division continues until it is found that no calculation *cell* contains any additional substructure. The probability of further subdivision in this manner is measured via a χ -square test.

This approach allows areas of far higher interest to be more accurately accounted for, whilst avoiding the error introduced into an MI calculation by dividing the entire calculation space into a fine-grained grid. To allow this technique to work, the calculation space was arranged into *tree notation* rather than the more common grid notation that is often used in mutual information calculations, See Figure 5.4.

The application of this algorithm results in a mixed-size partition being used to calculate the MI, with the recursive subdivisions going deeper in areas of the calculation space which the data-points are more numerous, and has been shown to provide accurate results in the analysis of many chaotic systems.

5.9 Medical Image Processing

The arena of medical image processing presents numerous unique challenges that had not become apparent in other areas, one of which is the issue of *image registration* (also known as *image alignment*). Image registration is the

process of combining multiple images of the same subject, taken at different times, distances, or orientations, into one image. The act of registering and *summing* multiple copies of the same subject can decrease the signal to noise ratio (i.e. increase clarity) or improve image resolution [59] when compared to any of the initial input images. It is also highly useful for combining the output from numerous imaging or data sources (e.g. CT, MRI, MEG/EEG, etc), to allow the more accurate visualisations of scans or readings obtained in isolation.

There are many methods by which to register images, for an overview of traditional techniques see [12]. The automated registration of multiple images has been aided greatly by the application of MI analysis to identify the correct scale, orientation and perspective of images [4, 61, 90, 147, 152], most often in relation to a *base image*. In the simplest form, the MI technique is repeatedly applied to a pair of two-dimensional images of the same object, with some form of translation being applied to one image at each iteration. The MI of the images in question will peak at the point at which the two images are maximally aligned, i.e. the point at which object one shares the most amount of information with object two. This is a necessary oversimplification of the technique, as the decision on the *type* of translation and the specific *degree* of translation required at each step are far from trivial.

There are two forms of MI which are, to our knowledge, applied uniquely to the field of image registration. These measures are derived from the same information theoretic calculations as the traditional MI, but have yet to make great impact outside of the field. These two measures are *Normalised Mutual Information* [132] and *Multivariate Mutual Information* [94] (also known as *Generalised Mutual Information*). A brief description of these two measures follows, with an example of the use of each measure.

5.9.1 Normalised Mutual Information

An interesting advance in the MI technique can be found in the work of Studholme *et al* [132], in which the use of *Normalised Mutual Information* (NMI) is defined. The NMI is defined as:

$$I_{nmi}(X, Y) = \frac{H(X) + H(Y)}{H(X, Y)} \quad (5.13)$$

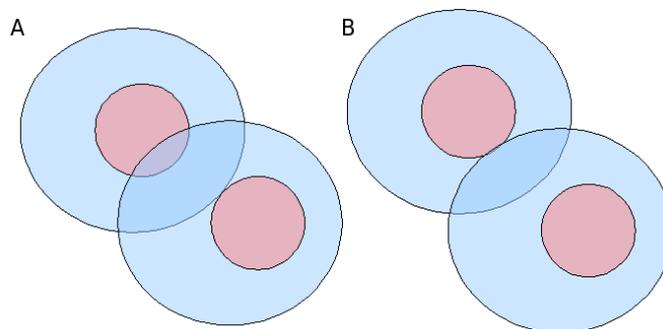


Figure 5.5: An example of the type of error that can cause a rise in MI even though the images are becoming increasingly misregistered.

This measure of MI was suggested to counter the problem of repeated *misregistration* (incorrect alignment, or orientation) of images of differing size. The problem tackled was that of image overlap, which was found to alter the MI readings to the point where the MI of the images can actually *rise* as the registration of the images *decreases*. The problem is twofold; firstly, it was found that a decrease in image overlap decreases the number of working samples within the calculation, as there can be no MI reading from sectors in which the images do not overlap. This decreases the relevancy of the MI analysis, as the probability distributions on which it relies is calculated from a much smaller sample set, and therefore can be considered less representative. Secondly, as the MI value of two misregistered images has been shown to be flawed in some cases, the probability of the MI rising as the images become even further misregistered is also increased. An example of this kind of misregistration error can be seen graphically in Figure 5.5. In this figure, A shows the partial registration of two images, the MI for these two images will be within a middle range, as at the places where the images intersect there is reasonable similarity with the colour schemes used. B shows a further misregistration of the two images, yet the MI of the intersecting segments of image will yield a higher results than in A, this is due to the identical colour (or intensity) of the overlapping image segments.

5.9.2 Higher Dimensional Mutual Information

Whilst Multivariate Mutual Information has not been widely adopted as an analytical technique across multiple disciplines, one area in which it has been

successfully employed is in the field of image registration [11, 88, 108, 132], although often not via the original technique (See Equation 5.11). In the literature relating to medical imaging alternative definitions for MMI have been proposed, most interestingly by Studholme *et al* [133], who define their calculation of MMI as:

$$I(A, B, C) = \sum_{i,j,k} p(a_i, b_j, c_k) \log \frac{p(a_i, b_j, c_k)}{p(a_i)p(b_j)p(c_k)} \quad (5.14)$$

This definition varies greatly from the original work of McGill [94], see Section 5.4.2. The level of this variance can be seen when comparing the entropy methods of calculating the two metrics. If we call McGill's MMI metric I_M and Studholme I_S the two entropy methods are:

$$\begin{aligned} I_M(A, B, C) = & H(A) + H(B) + H(C) - H(A, B) \\ & - H(A, C) - H(B, C) - H(A, B, C) \end{aligned} \quad (5.15)$$

$$I_S(A, B, C) = H(A) + H(B) + H(C) - H(A, B, C) \quad (5.16)$$

The difference in the two proposed methods is easily seen from Equations 5.15 and 5.16. The calculation of Studholme's metric includes far more of the joint entropies of the three images in question than that of McGill. If we return to the Venn diagrams used previously, we can see graphically the regions of shared information which will be negated or included when applying these two techniques. Figure 5.6 shows the areas of joint or mutual information content analysed during the two calculations.

This method of calculation has become widely adopted within the field of medical imaging, yet has not been adopted by other information theoretic research areas with such success.

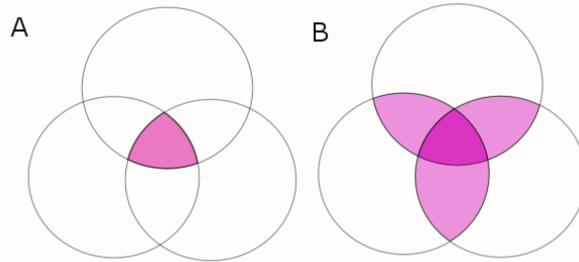


Figure 5.6: Comparative depictions of the level of shared information measured during a) McGill's [94] traditional MMI technique, and b) the technique used by Studholme [133]. In diagram B, the darker shaded area in the centre of the Venn diagram denotes that this central union is counted twice when calculating the final MI value

5.10 Complex Systems

The application of information theoretic techniques to complex systems is still in its infancy, but significant advances have been made in certain areas. The use of MI in complex systems has focused on the application of MI to identify *phase-transitions* in the complex systems in question, i.e. the point at which the behaviour of these systems change from one qualitatively distinct behaviour to another.

The work of Langton [82] centred about the use of physical systems for viable computational tasks. The use of Cellular Automata was employed as a *proof of concept* for this idea, and was used to display the possibility that a physical or biological system could provide support for the primitive functions required for computation, i.e. transmission, storage, and modification of information. To enable computation using these CA models, it had to be shown that two cells must be able to show a degree of *cooperation*, i.e. the behaviour of one cells must be able to directly impact the state of another, and *vice versa*. If this were the case, then it would be possible to find some form of correlation between the events that take place at the two cells. Traditional correlation methods require some form of order, or linearity, to be apparent in the mutual behaviour of the two cells in question, whereas Mutual Information can be applied to measure the correlation between two *unordered* variables. Langston applied the MI analysis to cells and compared the results to values of his behaviour parameter (the λ -parameter) to

investigate the point at which the system moved away from a static state and begins to show signs of complexity. It was found that the MI of the system could clearly show the transition from static to dynamic behaviour, but also showed the further transition from dynamic (or complex) behaviour to chaotic (or random) behaviour.

An interesting application, which we will cover extensively as specific features are directly employed later, is the work by Wicks *et al*, who applied MI to the identification of a kinetic phase transition in the Scalar Noise Model (SNM) [146, 43, 96, 25]. The SNM is a model of self-propelled particles that is known to exhibit self-organisational behaviour under certain parameter values. The SNM model is described by four equations, describing coordinate, heading, and velocity.

$$x_{n+1}^i = x_n^i + v_n^i \delta t \quad (5.17)$$

$$y_{n+1}^i = y_n^i + v_n^i \delta t \quad (5.18)$$

$$\theta_{n+1}^i = \langle \theta_n^{NR} \rangle + \delta \theta \quad (5.19)$$

$$v_{n+1}^i = v_0 (\cos \theta_n^i + \sin \theta_n^i) \quad (5.20)$$

At each time-step t , each particle i within the system performs a *parallel* update of its heading, which is taken as an average of the headings of all particles within distance R of particle i , including i itself, with a random angle $\delta \theta$ added to their heading which causes fluctuation within the system. The fluctuation term $\delta \theta$ is a randomly chosen *independently identically distributed* angle in the range $-\eta < \delta \theta < \eta$. In this respect the term η can be viewed as the *noise parameter* which controls the degree of stochastic, or random, behaviour within the system. Figure 5.7 gives a visual representation of the heading update for particle i .

By changing values of η , the behaviour of the systems can be manipulated. Figure 5.8 shows the effects on the flocking behaviour of the SNM for different values of η . At low noise levels (e.g. $\eta \rightarrow 0$), we see the highly ordered behaviour expected from a flocking model with the particles exhibiting a high degree of clustering. As η approaches $\frac{\pi}{2}$ the behaviour has changed, and whilst there are still clusters of particles present, the distribution appears far more stochastic than at low noise. The higher η is raised above

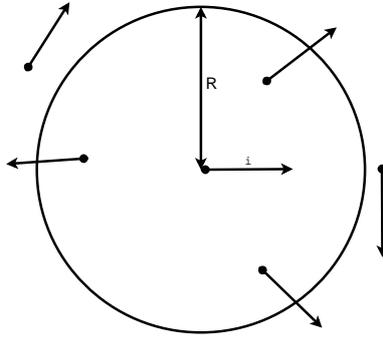


Figure 5.7: Visualisation of the heading update for particle i , utilising all particles within radius R , corresponds to term $\langle \theta_n^{NR} \rangle$ in Equation 6.3.

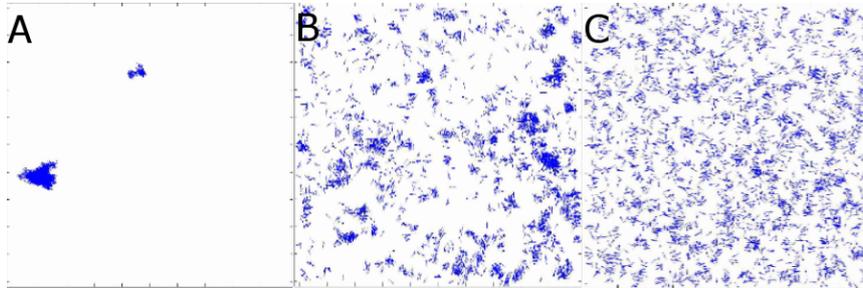


Figure 5.8: Three snapshots of the behaviour of the SNM under varying values of η . A) $\eta = 0$, B) $\eta = \frac{\pi}{2}$, and C) $\eta = \pi$

the critical threshold, the lower the degree of order in the system appears to be, with all values of $\eta \geq \pi$ showing few signs of clustering or order. The work of Wicks *et al* sought to identify the critical phase transition in the SNM, η_c , i.e. the point at which the system begins to tend away from order, and starts to exhibit disordered and noisy behaviour.

By applying a mutual information calculation to the agent's position and direction coordinates, as seen in Equation 5.21, the position of the phase transition could be identified.

$$\begin{aligned}
I(X, \Theta) &= \sum_{i,j} P(x_i, \theta_j) \log_2 \frac{P(x_i, \theta_j)}{P(x_i)P(\theta_j)} \\
I(Y, \Theta) &= \sum_{i,j} P(y_i, \theta_j) \log_2 \frac{P(y_i, \theta_j)}{P(y_i)P(\theta_j)} \\
I &= I(X, \Theta) + I(Y, \Theta)
\end{aligned} \tag{5.21}$$

It is known that the scalar noise model exhibits a kinetic phase transition at $\eta \approx \frac{\pi}{2}$. This can be confirmed by calculating the value of the Binder cumulant (β) which will take a value of $\beta \approx \frac{2}{3}$ at periods of order, and $\beta \approx \frac{1}{3}$ at periods of low order.

$$\beta = 1 - \frac{\langle \phi^4 \rangle}{3\langle \phi^2 \rangle^2} \tag{5.22}$$

Where ϕ is an *order parameter* of the system, in this case the *speed of the net movement* of the system

$$\phi = \frac{1}{Nv_0} \left| \sum_{i=0}^N v_i \right| \tag{5.23}$$

We can see then that the order parameter of the system is basically the *average velocity* of the system at any one time-step, whilst the Binder cumulant acts as a general measure of the system state ($\beta \approx \frac{2}{3}$ implying order, whilst $\beta \approx \frac{1}{3}$ implies disorder). The final measure of the system state is the *susceptibility*, χ , which is the variance of the order parameter

$$\chi = \sigma^2(\phi) = \frac{1}{N} (\langle \phi^2 \rangle - \langle \phi \rangle^2) \tag{5.24}$$

Figure 5.9 shows these measures alongside the MI calculated according to Equation 5.21.

We see that the Binder cumulant signifies the phase transition point at $\eta \approx \frac{\pi}{2}$, which is confirmed by a change in the behaviour of the order parameter at this value. We also see that the MI of the system shows a definite peak at this point, alongside a reduction in the size of the error

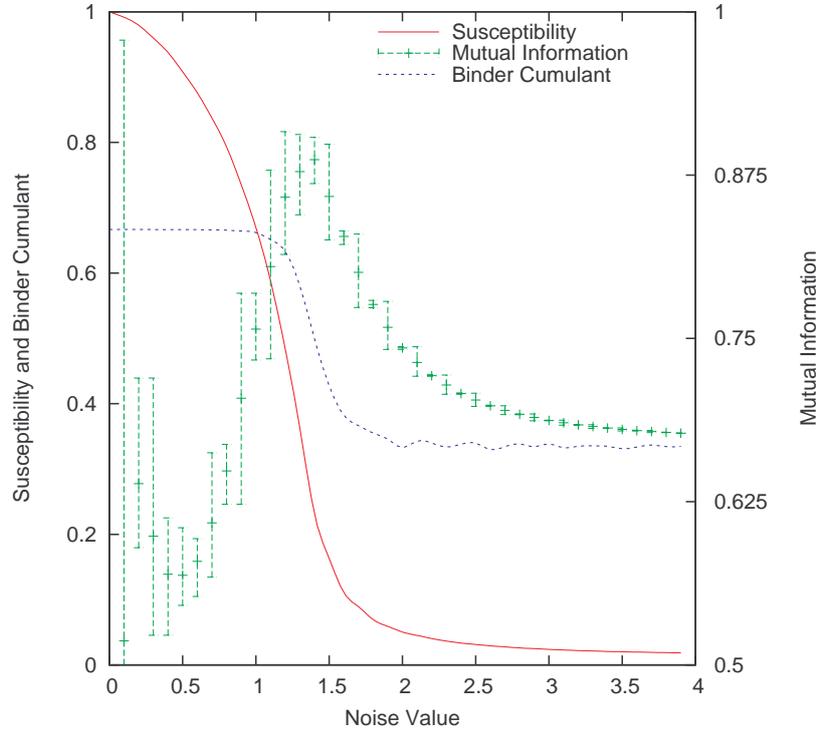


Figure 5.9: Output from the SNM across different values of the Noise value (η), showing the susceptibility, the MI, and the Binder cumulant. Error bars represent minimum and maximum values, and are shown for MI only.

bars. These results show that the MI of a system can be used to identify the point of a kinetic phase transition within a system of inter-connected particles.

Figure 5.9 also shows that there are very large error bars for the MI at low levels of η , this is due to the more varied states of the system for low noise level experiments. At low noise levels the system can settle into a state in which N distinct clusters of agents exist, the lower the value of N the higher the MI of the system will rise, whereas for larger values of N the MI of the system could be extremely low. At values of η greater than the transition point of the system, $\eta > \frac{\pi}{2}$, the errors bars reduce considerably, as the system settles into a far more predictable state of disorder.

5.11 Summary

We have seen that the use of mutual information is common across many computer science related disciplines, from image processing to chaotic systems. The usefulness of MI as a statistical measure of the interdependence of signals or data-sets is clear, but it has never previously been applied to an evacuation system.

The results obtained when using MI to identify the point of the kinetic phase transition within the scale noise model clearly shows that the MI of a system can be used as *some* measure of order, but we suggest that within an evacuation the MI could be used to measure the order of the evacuation. We propose that the MI of an evacuating population may be used as a measure of the *order* or *disorder* of the evacuation itself.

In Chapter 7 we detail the application of the mutual information technique to measure order in an *in silico* evacuation environment.

Chapter 6

Detection of a Phase Transition

6.1 Introduction

During this chapter we will cover initial work into the application of MI techniques, described previously (see Chapter 5), to the detection of a kinetic phase transition within a previously analysed agent-based particle model. The model chosen for this investigation is the *Scalar Noise Model* [146], an agent-based particle swarm model with a defined phase transition point. We start with a definition of the model, including a technique which may be used to identify the phase transition in the model. We then investigate the manner by which different MI measures may be used to measure order within the system, and conclude with an analysis of their efficacy.

6.2 Scalar Noise Model

The Scalar Noise Model (SNM) is a mathematically defined model of interacting particles [146, 96, 43, 25]. At each time-step within the SNM, all particles perform an asynchronous update of their direction of movement, during which each particle takes the average heading of all other particles, including itself, that lie within its interaction radius, R , plus an additional random noise measure, $\delta\theta$. A visual example of this heading update may be seen in Figure 6.1.

The Equations governing the operation of the SNM, on a per-particle

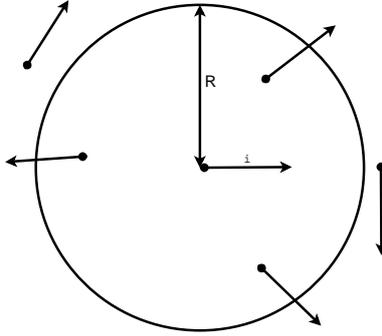


Figure 6.1: Visualisation of the heading update for particle i , utilising all particles within radius R , corresponds to term $\langle \theta_n^{NR} \rangle$ in Equation 6.3.

basis, can be seen below.

$$x_{n+1}^i = x_n^i + v_n^i \delta t \quad (6.1)$$

$$y_{n+1}^i = y_n^i + v_n^i \delta t \quad (6.2)$$

$$\theta_{n+1}^i = \langle \theta_n^{NR} \rangle + \delta \theta \quad (6.3)$$

$$v_{n+1}^i = v_0 (\cos \theta_n^i + \sin \theta_n^i) \quad (6.4)$$

The term $\delta \theta$ is an independent identically distributed variable (IID) chosen from within the range $-\eta \leq \delta \theta \leq \eta$, and it is the value of η which defines the level of noise within the system. With very low values of η the SNM produces a very strong flocking behaviour, as the particles slowly begin to take a similar heading, reducing the overall level of noise with the system. At higher levels of η the particles have so much additional noise added to their heading at each time-step that a coherent behaviour within the system is not able to emerge, so the overall behaviour is that of random movement of particles. It is the point within this range that the phase transition lies, and the point which we wish to identify be means of Mutual Information analysis.

6.2.1 Phase Transition Point

The phase transition within the scalar noise model may be identified in a number of ways, but the simplest of these means is the *Binder cumulant*. If we may define the *order* of the system to be measured by the *net transport* of the system at one time-step, as for highly ordered simulations in which all particles have assumed a similar heading the net transport will be high, whereas as the level of disorder in the system increases the net transport tends towards zero. We will call this net transport, or order parameter, ϕ , and it is defined mathematically in Equation 6.5.

$$\phi = \frac{1}{Nv_0} \left| \sum_{i=0}^N v_i \right| \quad (6.5)$$

The order parameter, measured across changing values of the Noise parameter, η , can be seen in Figure 6.2, in which the error bars represent the standard deviation of the values recorded across 64 simulation runs. We see that as the noise within the system is raised the order shows a steady drop, which then plateaus at the point when the system has passed into a more disordered phase. A rough estimation of the phase transition point, from this output, can be made at approximately 1.6.

We see from Figure 6.2 that during the order to disorder transition there is a far greater level of variation in the values recorded for ϕ . The Binder Cumulant, which we will denote β , uses this variation to more accurately measure the range in which the transition happens, the mathematical form of β can be seen in Equation 6.6.

$$\beta = 1 - \frac{\langle \phi^4 \rangle}{3\langle \phi^2 \rangle^2} \quad (6.6)$$

The value of β will assume a value of approximately $\frac{2}{3}$ when the system is in a ordered state, whilst during periods of disorder will assume the value of approximately $\frac{1}{3}$. During the phase transition itself it will take values within this range. Figure 6.3 shows the calculation of β for changing values of η . We see from the Binder cumulant output that the system begins to degrade from its state of order at approximately $\eta = 1.1$, entering into a regular state of disorder at a value of approximately $\eta = 1.6$.

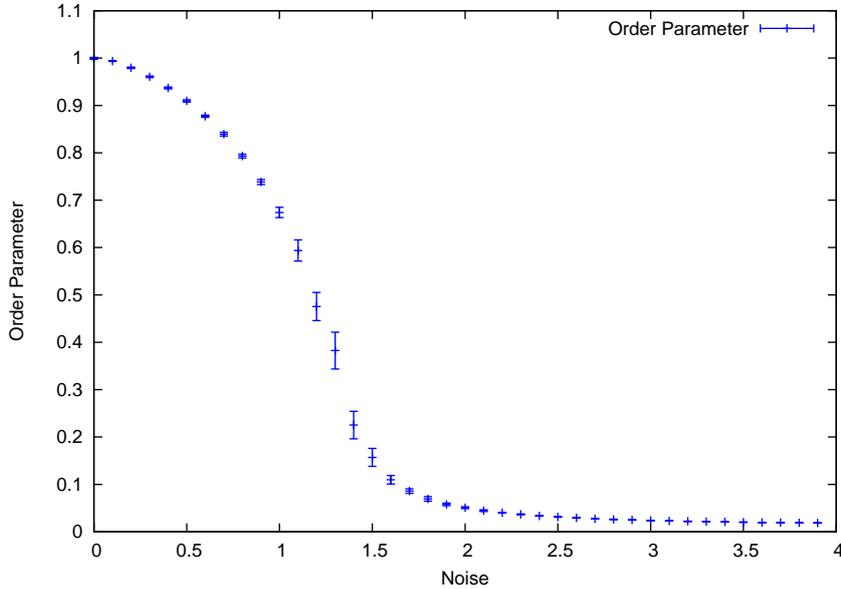


Figure 6.2: The changing values of the order parameter, ϕ , for different values of the noise parameter, η . The error bars represent the standard deviation of the values across the 64 simulation runs.

From these results, we see that to properly detect the order to disorder transition we must find a metric which exhibits quantitatively different behaviour about this range of values $1.1 < \eta < 1.6$, or quantitatively different behaviours at $\eta < \frac{\pi}{2}$ and $\eta > \frac{\pi}{2}$, before we can safely assume that the state of order within the system may be measured. The exact mathematical value of the transition point in this system is at a value of $\eta = \frac{\pi}{2}$ [146].

6.2.2 Susceptibility

One metric by which to identify the first order phase transition within the SNM is by calculating the susceptibility of the data [3] which is, in plain english, the variance present in the order parameter. We saw previously that during the phase-transition of the SNM the variance of the order parameter increases, as can be seen from the increased size of the error bars in Figure 6.2, if we let the variance of this order parameter be the susceptibility, χ .

$$\chi = \sigma^2(\phi) = \frac{1}{N} (\langle \phi^2 \rangle - \langle \phi \rangle^2) \quad (6.7)$$

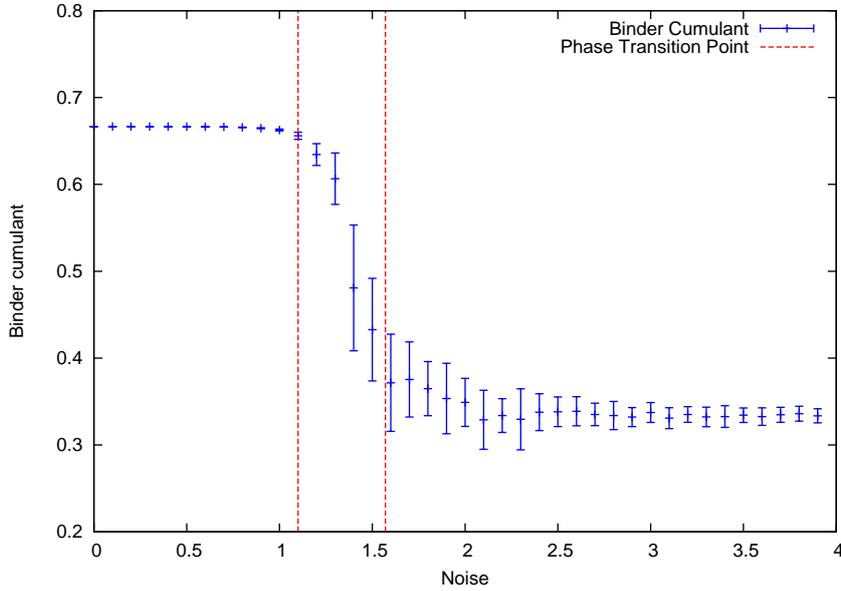


Figure 6.3: The value of the Binder cumulant, β , in the scalar noise model, calculated for different values of the noise parameter, η .

This measure yields the susceptibility of the SNM at different values of η (noise), and can be seen in Figure 6.4, where the error bars represent the standard deviation of the data.

We see that the susceptibility of the system shows a distinct peak during the phase transition, but that the values recorded also show far greater variance across this range.

The susceptibility, i.e. the variance of the order parameter within a system, has been extensively used for the identification of phase transitions within complex systems [3, 9, 26, 75] and has been found to be a powerful and malleable technique by which to perform this analysis. It is, then, the susceptibility of the SNM that we will be using as a general measure of the success of the following techniques to adequately identify the phase-transition point within the scalar noise model.

6.2.3 Methodology

The aim of this Chapter is to investigate possible methods of employing our chosen form of analysis, Mutual Information, for the identification of a

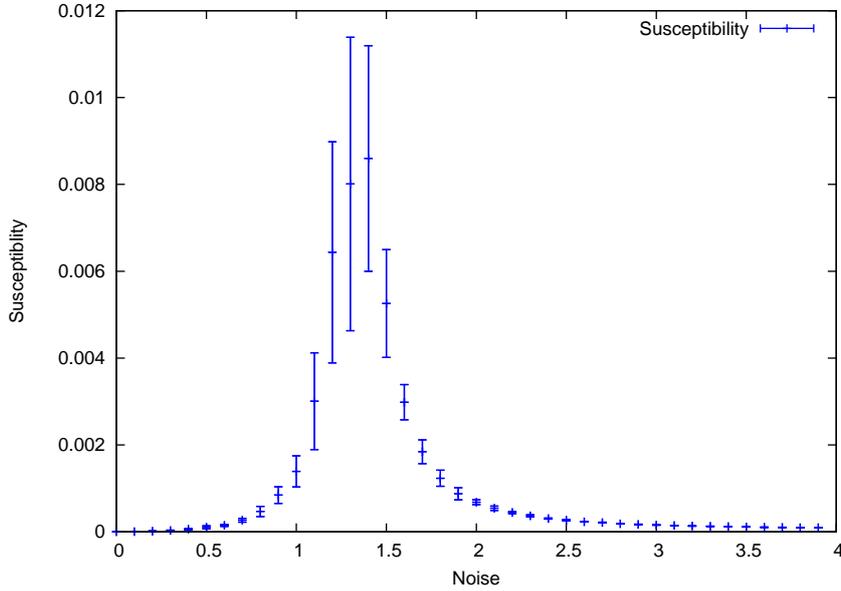


Figure 6.4: The susceptibility of the scalar noise model, χ , calculated for different values of the noise parameter, η .

phase transition within the SNM. Numerous forms of Mutual Information, covered in detail during Chapter 5 are applied to the three agent variables of the SNM model (position, X and Y , and heading, Θ) and their identification of the phase transition is classified in two ways. Firstly, as our ideal method would give a simple, human readable metric by which to measure order, the output is analysed by eye to confirm that the point of the phase transition can be identified from visual output only. Secondly a correlation analysis will be carried out to measure the mathematical correlation between the numerical output of our analytical methods and the level of order within the system measured via more traditional means.

Data Collection

At the start of each experiment, N particles are randomly distributed throughout the game space, L . As time progresses the behaviour of the particles falls into a steady state, defined by the noise parameter, η . As we are looking to analyse the behaviour of the system during the steady state, data is not collected until 50,000 iterations have been completed. At this stage, the various MI metrics, defined later in the Chapter, are applied at every

Variable	Designation	Value
Number of Particles	N	3000
Game Area	L	50
Velocity	v_0	0.15
Interaction Radius	R	0.5
Total number of iteration	t	60,000
Iteration at which data collection begins	τ	50,001

Table 6.1: Variable values used during the experimentation with the SNM.

time-step for the next 10,000 steps. The values of the metrics collected during these 10,000 steps are then averaged (arithmetic mean) and it is these values which are presented in all the figures and analysis in this Chapter.

Parameters

The parameters used whilst collecting the data for the experiments in this Chapter are defined in Table 6.1.

These values are identical to those presented in [153], in which similar experiments have been run. This was done to ensure that the data collected was a true representation of the system whilst in a well mixed state.

6.3 Development

As existing implementations of the SNM were not freely available, an implementation of the SNM was developed specifically for these experiments. It was decided that as an agent based model the system was best developed using an object oriented programming language, as each agent and it's functionality was inherently suited to being encapsulated within a single class. The chosen programming language was C++, due mainly to developer experience. All MI and entropy libraries were also developed specifically for these experiments, and also coded in C++.

The algorithm underlying the SNM is of $O(N^2)$ complexity, in that at each time-step every agent (or particle) within the model must cross reference with every other agent within the system to ascertain whether or not these agent's are within their interaction radius, R . This computational complexity results in a function whose runtime grows non-linearly according to the the number of agents, N . All other equations in this model, and the

associated code for experimentation, are of $O(kN)$ complexity, including all MI and entropy libraries, and grow linearly with N .

The initial test system was an Intel Core i7 2.0GHz 4-core processor with 8GB RAM, and compiled with Visual Studio 2010. The application was found to run at sufficient speed for small population sizes. For a population of 300 agents, and all other variables set as detailed in Section 6.2.3, the simulations took an average of 0.0012 seconds per time-step, but when attempting to operate at the system size stated for these experiments ($N = 3000$) the computation time per time-step was found to be in the region of 0.119 seconds. Due to the data collection only being initialised after the system has settled into a steady state, i.e. after 50,000 time-steps of the simulation, this lead to run-times of up to 7140 seconds (or approximately 2 hours) , which was considered unacceptable.

The initial avenue considered for optimisation was a course graining of the particles into computation grids, which would allow each particle to poll only a subset of the entire population to update it's own heading value. This technique would have increased computation speed during the initial time-steps of the algorithm, but the behaviour would be undefined as the time-steps grew. In the worst case scenario, e.g. for low values of η , the computation speed would be reduced to such a point that this optimisation method would be no faster per time-step than the un-optimised code. Instead it was decided to opt for a parallelisation approach to the optimisation, splitting computation of multiple cores, threads, or physical systems to reduce overall runtime.

The first attempt at parallelisation was an openMP parallelisation, a software parallelisation method that exploits multicore/multithread processors to reduce the runtime of applications. By splitting the execution of the heading update functionality across multiple computation streams it was possible to reduce the running time of the experiments from 2 hours to just 39 minutes, a 66% reduction in overall computation time.

Figure 6.5 shows the computation time against the increase in population size for both the Serial C++ and openMP implementations of the SNM. Further, smaller speed increases are available within this code, but the reduction of computation time to under the 1 hour mark was considered acceptable for experimentation, so no further optimisations were pursued.

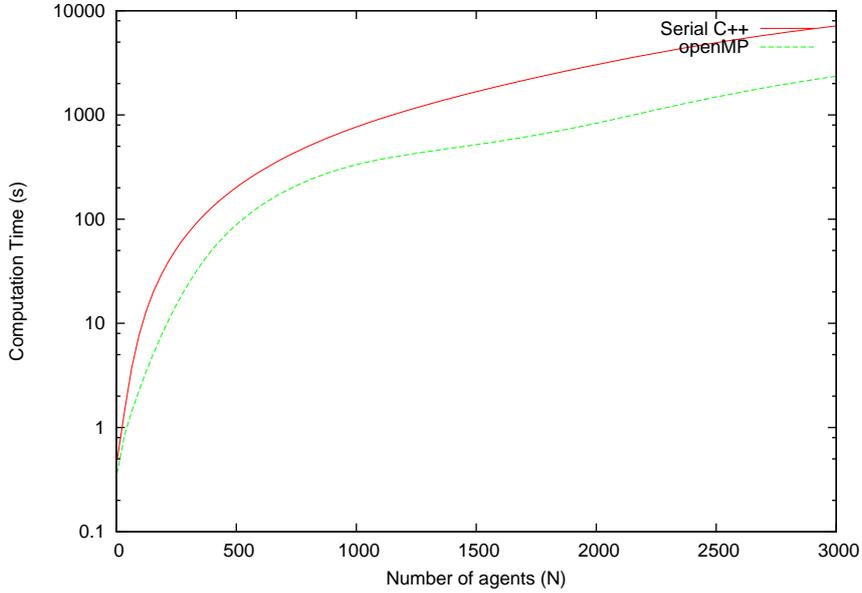


Figure 6.5: Computation times, in seconds, of 60,000 time-steps of the SNM for varying values of the population size, N . The timings of both the Serial C++ implementation and the openMP implementation are shown. The y-axis in this figure is logarithmically scaled.

6.4 Experimentation

6.4.1 Multivariate Mutual Information

The first metric which will be tested is the traditional multivariate mutual information of three variables. The metric is simply defined as the MI of two of the variables less the conditional MI of the those variables considering the value of the third variable is known, see Equation 6.8 [94, 128].

$$\begin{aligned}
 I(X, Y, \Theta) &= I(X, Y) - I(X, Y | \Theta) \\
 I(X, Y, \Theta) &= \sum_{i,j} P(x_i, y_j) \log \frac{P(x_i, y_j)}{P(x_i)P(y_j)} \\
 &\quad - \sum_{i,j,k} P(x_i, y_j, \theta_k) \log \frac{P(\theta_k)P(x_i, y_j, \theta_k)}{P(x_i, \theta_k)P(y_j, \theta_k)} \quad (6.8)
 \end{aligned}$$

The results from this analysis can be seen in Figure 6.6, where the error bars represent the standard deviation of the metric across all test runs.

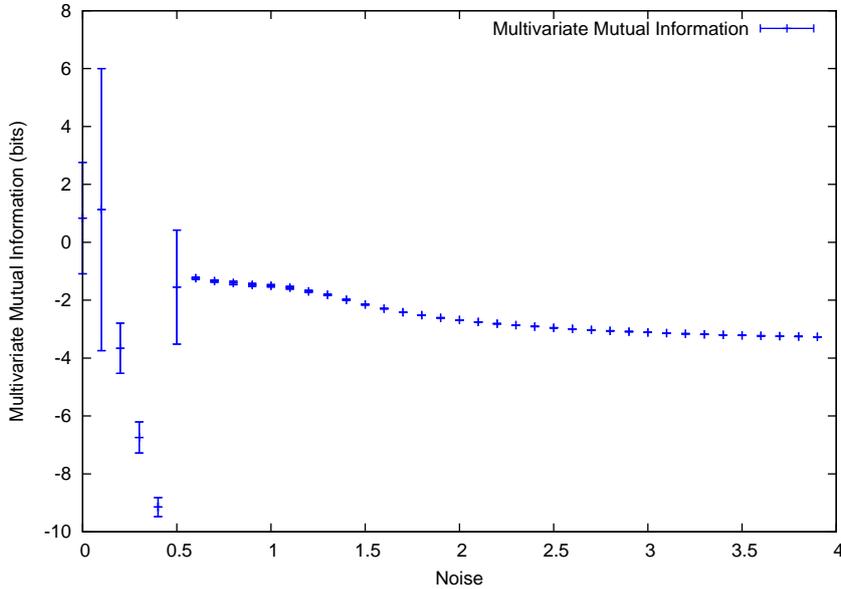


Figure 6.6: The multivariate mutual information, $I(X, Y, \Theta)$, measured across 40 runs of the scalar noise model across different values of the noise parameter, η . The error bars represent the standard deviation of the data.

We see from this figure that the MMI shows a high level of variance during the period of order within the system, which drops drastically when the system begins to shift into a state of disorder. The mean value of the MMI though is also highly variable at low levels of noise, plateauing as the order within the system breaks down. If we re-plot the output to investigate the variance of the MMI, we can see clearly the trend as the system falls into disorder, see Figure 6.7.

The figure shows the trend of the standard deviation of $I(X, Y, \Theta)$ to drop significantly after the phase transition point. Viewing the data shows that for values of η below the phase transition point $0.015 < I(X, Y, \Theta) < 4.87$, whereas after the phase transition the range of values recorded was $0.0009 < I(X, Y, \Theta) < 0.008$. We could, then, infer from this result that the MMI may, in effect, be used to identify the presence of order within a system, but that this could only be achieved with a significant number of test runs.

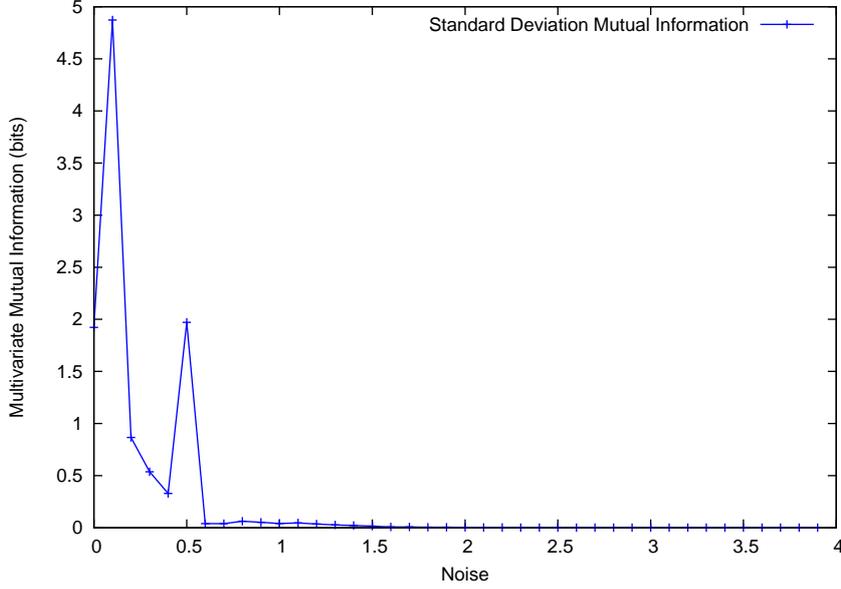


Figure 6.7: The standard deviation of $I(X, Y, \Theta)$ measures across all experiment runs.

6.4.2 Conditional Mutual Information

Conditional mutual information (CMI) is a further information theoretic measure that will be tested. The CMI of variables A and B conditional on C, denoted $I(A, B|C)$, can be thought of as the amount of uncertainty remaining in variables A and B when the value of C is known. Like MMI described previously, the CMI may be applied directly to the three agent variables in the SNM.

$$\begin{aligned}
 I(X, Y | \Theta) &= \sum_{i,j,k} P(\theta_k) P(x_i, y_j | \theta_k) \log \frac{P(x_i, y_j | \theta_k)}{P(x_i | \theta_k) P(y_j | \theta_k)} \\
 &= \sum_{i,j,k} P(x_i, y_j, \theta_k) \log \frac{P(\theta_k) P(x_i, y_j, \theta_k)}{P(x_i, \theta_k) P(y_j, \theta_k)} \quad (6.9)
 \end{aligned}$$

Being applied in this manner, i.e. with Θ as the conditional variable, will ensure that for greater overall order within the system we will have a lower value of CMI. The results of the CMI analysis of the SNM may be seen in Figure 6.8.

We see from the output that visualising the CMI in this way shows a

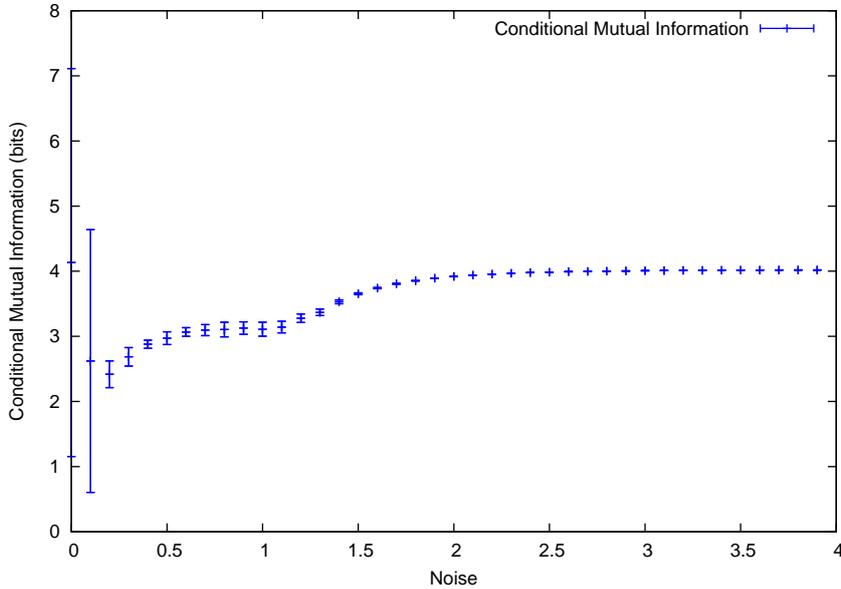


Figure 6.8: The conditional mutual information of X and Y conditional on Θ , $I(X, Y|\Theta)$, across changing values of the noise parameter, η . The error bars represent the standard deviation of values across all test runs.

smoother transition as the system passes from a state of order into that of disorder, will a reasonably steady rise in line with η . As with the MMI the variance of the values recorder is far greater during periods of order, dropping to values of $0.0003 < I(X, Y|\Theta) < 0.0098$ after the phase transition point.

6.4.3 Normalised Mutual Information

The Normalised Mutual Information (NMI) is a measure used extensively in image processing, and particularly in the field of medical image registration [68, 85, 116]. The NMI measure used here is a variant of that presented in [132], and discussed in Section 5.9.1, which extends the technique to a third dimension.

$$I_{norm}(X, Y, \Theta) = \frac{H(X) + H(Y) + H(\Theta)}{H(X, Y, \Theta)} \quad (6.10)$$

Output from the NMI analysis of the SNM can be found in Figure 6.9, with the error bars representing the standard deviation of the recorded data

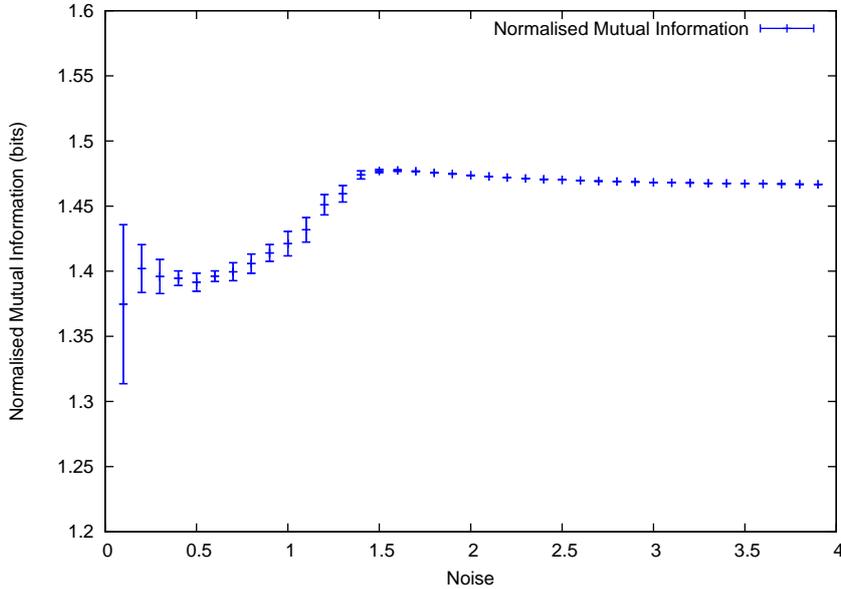


Figure 6.9: The Trivariate Normalised Mutual Information of the SNM measured for changing values of η . Error bars represent the absolute standard deviation from the mean.

items across the entirety of the test runs.

We see from the graphical output that the NMI metric provides results which are qualitatively similar to that of the CMI (see Figure 6.8), with the same trend of decreased variance in values of η greater than the phase transition point. There is an obvious upward trend in the NMI as the value of η is increased, with a very gradual drop as η progresses past the phase transition point.

6.4.4 Studholme’s Mutual Information

This second multivariate MI measure presented is also extensively used in medical imaging, and is a variant of traditional two variable MI which has been extended to account for a third variable. As discussed previously, see Section 5.9.2, the multivariate MI measure presented here differs from traditional multivariate MI primarily in the amount of conditional MI between each possible pair of input variables, see Figure 5.6. First presented by Studholme *et al* [132], we will refer to this metric as Studholme’s Mutual Information. It is defined by Equation 6.11.

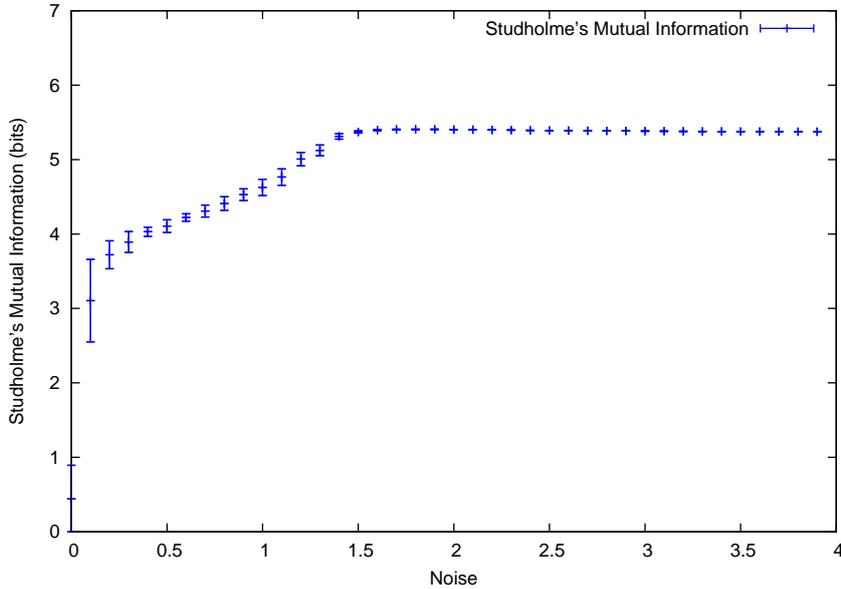


Figure 6.10: Studholme’s Multivariate Mutual Information of the SNM measured across changing values of the noise parameter, η . Error bars represent the absolute standard deviation from the mean.

$$I(X, Y, \Theta) = \sum_{i,j,k} p(x_i, y_j, \theta_k) \log \frac{p(x_i, y_j, \theta_k)}{p(x_i)p(y_j)p(\theta_k)} \quad (6.11)$$

Output from Studholme’s Mutual Information can be seen in Figure 6.10, with the error bars representing the standard deviation of the recorded data across the entirety of the test runs.

We see from Figure 6.10 that Studholme’s Mutual Information (SMI) gives a visually meaningful output, with the SMI rising consistently up to the point at which the system falls into a state of disorder, at which point it plateaus and remains at a constant value with a low deviation across test runs.

6.4.5 Wicks’ Mutual Information

The final information measure we investigate here is an MI measure suggested by Wicks *et al* [153], and originally applied to identify the phase transition in the model in question. The MI measure, we will refer to it as Wicks’ MI (WMI), is a combinatorial measure which averages the MI of the

positional input variables, X and Y, with the directional input variable, Θ , according to Equation 6.12.

$$\begin{aligned}
 I(X, \Theta) &= \sum_{i,j} P(x_i, \theta_j) \log_2 \frac{P(x_i, \theta_j)}{P(x_i)P(\theta_j)} \\
 I(Y, \Theta) &= \sum_{i,j} P(y_i, \theta_j) \log_2 \frac{P(y_i, \theta_j)}{P(y_i)P(\theta_j)} \\
 I(X, Y, \Theta) &= I(X, \Theta) + I(Y, \Theta)
 \end{aligned} \tag{6.12}$$

The combination of the variables in this way enables the calculation of the amount of order in the particles of the simulation to be measured across both planes. The output of the WMI analysis of the SNM is shown in Figure 6.11, where the error bars represent the absolute standard deviation of the values across the entirety of the test runs.

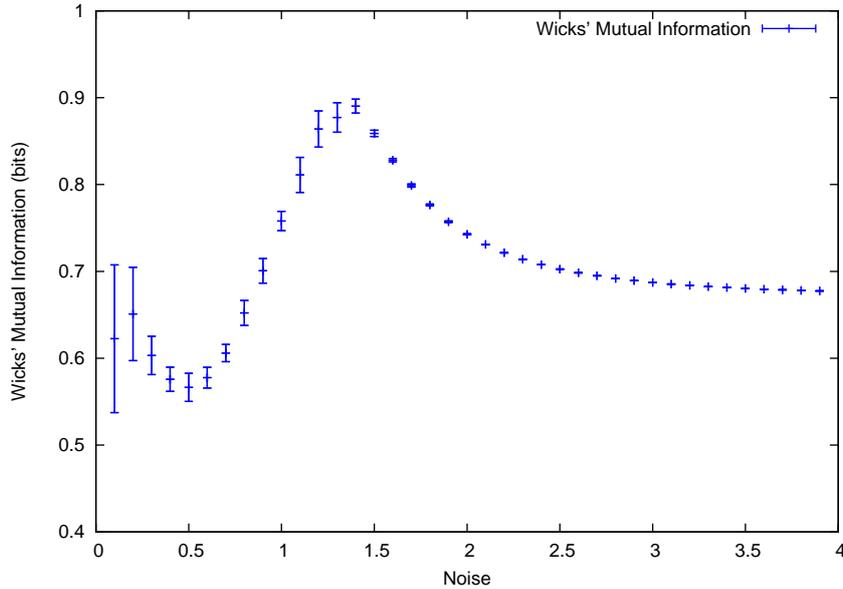


Figure 6.11: The Wicks Mutual Information, $I(X, Y, \Theta)$, measured across varying values of the noise parameter, θ . The error bars represent the standard deviation across all test runs.

We see from the graphical output that the WMI shows a distinct peak during the phase transition point and, as been seen with earlier metrics, the

variability of the metric is also reduced at levels of noise greater than the phase transition point, i.e. $\eta > 1.6$. Figure 6.12 shows the variability of the WMI metric, taken from the same data set as the previous figure.

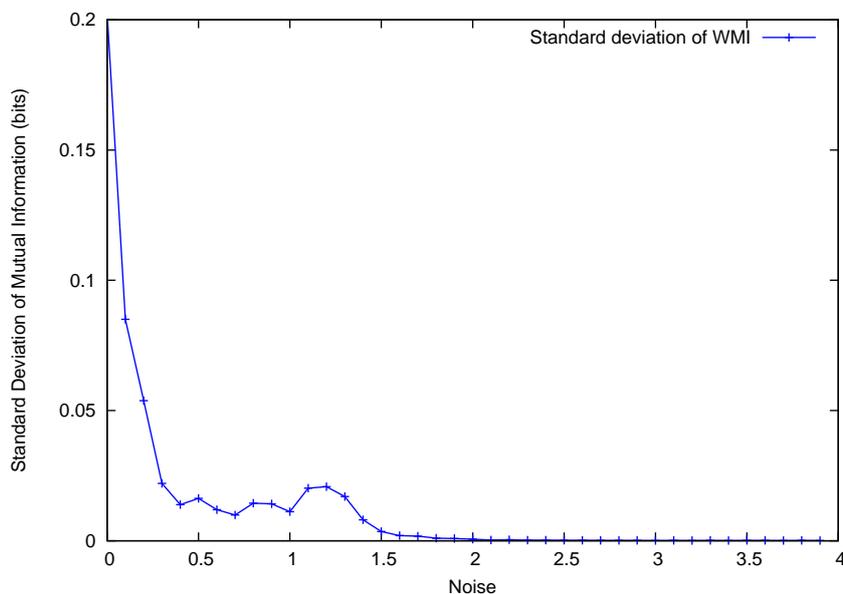


Figure 6.12: The variance of the WMI, as calculated from the entirety of the test runs.

6.5 Correlation Analysis

The correlation analysis will measure the amount by which the MI metrics, presented earlier in this Chapter, relate to the susceptibility of the system, χ . The results are presented in Table 6.2, the correlation measure used is the Spearman’s rank correlation coefficient, and all values presented are absolute.

We see from the data that, considering a standard two-tailed significance test ($\alpha = 0.05$), all results apart from the conditional mutual information (CMI) are statistically significant.

The strongest correlation seen is the Wicks’ MI, showing an R value of $R_s = 0.91$, which can be said with a high level of confidence (actual p-value was $p = 9.9e^{-16}$).

Metric	Absolute Correlation	p-value
MMI	0.51	0.001
CMI	0.27	0.091
NMI	0.57	0.000
Studholme's MMI	0.41	0.009
Wicks MI	0.91	0.000

Table 6.2: The absolute values of the Spearman's rank correlation coefficient, and the associated p-values, of the MI metrics and χ , the susceptibility of the system, across all test runs of the SNM.

6.6 Conclusions

We have seen throughout this Chapter the application of numerous analytical techniques for the detection of the phase transition within the scalar noise model. Some of these techniques have been used previously (e.g. susceptibility or the Binder cumulant), whereas others are newly applied during this work (e.g. conditional mutual information and Studholme's multivariate mutual information).

We have shown that the newly applied measures of order, apart from the MMI (see Figure 6.6), give visually meaningful output which may be used to identify the different levels of order within the SNM. The mathematical analysis though showed that whilst visually meaningful, the output from some of the MI-based metrics showed either low correlation or a lack of statistical significance.

The highest significant correlation found was shown to be the Wicks' MI measure, which showed a strong correlation with the susceptibility of the SNM. Figure 6.13 shows the Wicks' MI, the Binder cumulant, and the susceptibility of the SNM, plotted on the same figure, error bars represent the standard deviation of the data.

As we can see from this graphic, whilst both susceptibility and WMI show a definite peak about the phase transition point of this system the large error bars present on the susceptibility make it a far less usable metric. Additional to this, it can be seen that the WMI shows measurably different trends below $\eta = \frac{\pi}{2}$, i.e. the ordered phase of the system, than above $\eta = \frac{\pi}{2}$, i.e. the disordered phase of the system. We believe that these factors will allow the WMI metric to be applied as an order metric to evacuation systems.

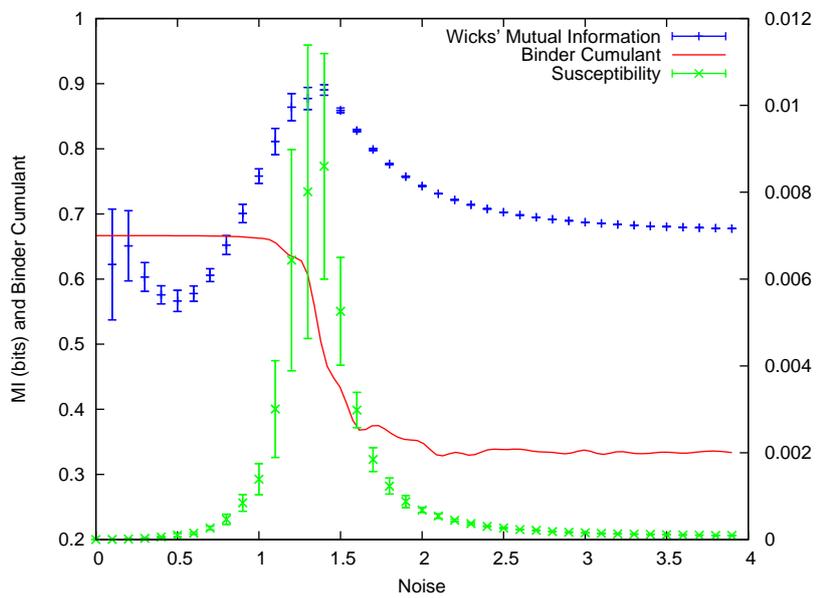


Figure 6.13: The values of Wicks' MI (blue), susceptibility (green), and Binder cumulant (red) calculated from the scalar noise model for different values of the noise parameter, η . Error bars present represent the standard deviation of the data.

Chapter 7

Initial Evacuation Simulations

7.1 Introduction

This chapter details the methodology which we use to identify the onset of crush conditions within an *in silico* evacuation simulation. The work detailed in this chapter is a *proof of concept* (i.e. initial evidence that the proposed technique is fit for purpose) rather than an extensive testing of this technique, further testing and large scale experimentation is summarised in Chapter 8.

Whilst initial experimentation with the mutual information (MI) analysis is carried out on a trivial evacuation scenario modelled using the original implementation of the Social forces model (SFM), it is accepted that this simulation environment cannot be used for the extensive testing of the MI technique, as the modelling of complex evacuation topologies are not possible. However, this simulation has been chosen to confirm the general suitability of MI to the purpose of crush detection. Following successful experimentation, advanced testing and validation will continue with a more fully-featured simulation package.

7.2 Hypotheses

We argue that the transition of evacuations from *laminar* to *turbulent* states can be used to identify the build up of crush conditions during an evacua-

tion. The MI analysis is used here to measure the *order* within an evacuation, therefore we expect the MI of the system to drop as competition for exit capacity increases. It is also expected that as the competition for exit capacity increases, the amount of force that builds up within the evacuating population will increase [51]; we therefore have our hypothesis for the acceptance of the MI technique as a plausible method for measuring crush.

- **Null hypothesis**

$$H_0 : R = 0$$

There **is no** correlation between mutual information and force.

- **Alternative hypothesis**

$$H_1 : R \neq 0$$

There **is** a correlation between mutual information and force.

7.3 Experimental Aims

There are two main aims of the experimentation during this Chapter. Firstly, we will show that the MI analysis which we apply can offer a dynamic and visually meaningful representation of the amount of order present within an *in silico* evacuation. This will be achieved by examination of the changing output of the MI of the system, alongside a manual examination of visual output of the SFM during the evacuation run. Secondly, we will show that the level of order, as measured by the MI of the system, shows a strong correlation with the level of force present at any one point in the evacuation. This will be achieved by means of a correlation analysis, showing that our alternative hypothesis (presented in Section 7.2) is correct.

In short, as we have proposed that a build up of high levels of force will follow a breakdown of order within an evacuation, we expected to see that as the level of force rises within the evacuating crowd the MI will show a marked drop, and vice-versa.

7.4 Expected Outcomes

We propose that the MI technique can be used to measure the *order* within a crowd of people, i.e. the higher the MI measure, the more ordered the evacuation. We believe that the breakdown of order within an evacuating

crowd is a major contributor to the build up of dangerous levels of physical force, and therefore that as MI drops the level of force measured will necessarily rise.

We therefore expect a significant negative correlation between the level of force present at any point in the evacuation and the corresponding MI measure at that point.

7.5 Methodology

Initial simulations were carried out on a modified version of the Social Forces Model¹. The model was written in serial C, and is identical to that which formed the basis of the simulations detailed in the original literature [51].

Additional libraries for calculation and analysis were based on those developed for use during the Scalar Noise Model analysis detailed previously, converted from parallelised C++ to serial C to enable simpler integration into the existing code-base. The MI analysis, and required libraries, were integrated with the existing code in an entirely passive manner, i.e. their inclusion cannot effect the result of simulations. As during the analysis of the SNM, this analysis is performed once at the end of each time-step during the simulations, with data being copied directly from the data structures present in the simulation into custom structures that form no part of existing computation to ensure integrity.

All simulations during this Chapter of the thesis were performed on the same evacuation topology, this being a single room evacuation shown in Figure 7.1. The reasons for carrying out experimentation on just one topology are twofold. Firstly, as this original model was designed for experimental purposes, this implementation does not include the means by which an evacuation topology can be loaded from file, therefore any changes to the environment must be hard-coded into the model itself. Secondly, the purpose of these initial experiments is to provide an indicator that the MI technique can be used to track changes in evacuation behaviour, rather than to show that it may be used to analyse realistic evacuation scenarios. More complex scenarios, as to be found in Chapter 8, will be analysed using a more mature and fully-featured simulation environment.

¹Model selectively available for non-commercial purposes from <http://angel.elte.hu/panic/>

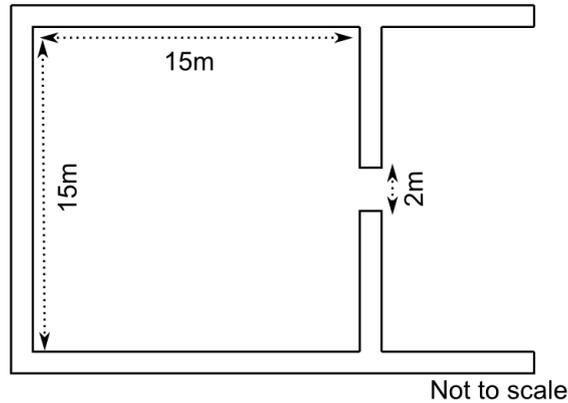


Figure 7.1: Representation of the evacuation topology used during the simulations in this Chapter. All evacuees are distributed evenly across the central 15m-by-15m room, and effect egress through the 2m wide door.

There were two changes made to the working of the original SFM model during this work. Firstly, in the original model an injured agent forms an immovable obstacle, incapable of movement yet still able to exert force (physical and social) onto any agent that comes into contact with them. This behaviour caused problems with simulations, as it was possible for multiple agents to become injured and form a “barricade” between the agents remaining in the room and the only available exit. This causes the simulation to end with a number of evacuees still remaining in the structure. Simulations in which this occurred were necessarily declared void, their results unusable, and experiments had to be restarted with the minimum amount of changes made to the initial conditions to avoid this situation. To counteract this issue the rule was added that were an agent to succumb to injury, the agent in question is removed from the simulation after an arbitrary amount of time. This allows the increase in force that an injured agent causes to be fully taken into account within the simulations, but negates the “barricade” behaviour mentioned previously.

Secondly, to obtain a baseline for the MI of the system (i.e. a null value), a period of *milling* was introduced, this took to the form of a 10 second addition to the start of each experiment run, during which the agents had no clearly defined goal. The purpose of this addition was to provide a baseline value for the MI in each simulation, i.e. the value that the MI takes during the random movement of agents. Hence, the data collected from

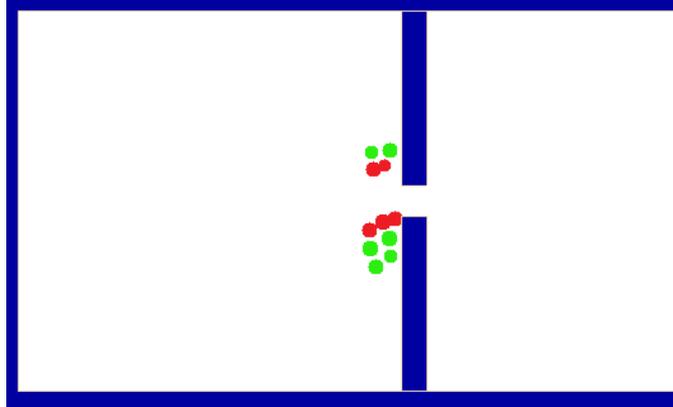


Figure 7.2: Image showing the barricading behaviour of the SFM. Red agents represent those incapacitated due to injury, whilst green agents are free to move, but unable to escape past the injured parties. In these situations the simulation will continue indefinitely with no resolution.

experiments we have carried out all show the start of the experiment at $t = -10$, with an actual evacuation beginning at $t = 0$ as per the standard model.

The measure of force stated at all points within this Chapter is identical to that used in the original literature. This force metric is a summation of the radial forces acting about the circumference of an individual, divided by that individual's circumference. This measure was initially used to replicate the ability of larger individuals (e.g. adult males), to suffer a higher level of force than than smaller individuals (children or young adults), with fewer ill effects.

The point at which an individual is considered to become “injured” is when the sum of the radial forces acting about their circumference exceeds $1600Nm^{-1}$.

All simulations carried out within this Chapter use the default parameters for the original implementation of the social forces model [51], with the exception of the random seed which is changed before the start of each simulation run. The value of the “driving force”, or the evacuees desire to effect egress, for these simulations is set to a constant value of $V_0 = 5ms^{-1}$. This value was chosen as it marks the first point at which injuries were found during the original work.

At the start of each of these simulations the room is populated with

exactly 200 individuals, who are equally spaced apart, once the simulation has begun Force and MI data are collected at the end of each and every time-step. As stated previously, each simulation begins with exactly 10s of “milling behaviour” in which agents move randomly, the evacuation stage begins immediately after this 10 seconds had passed. All times are stated with the start of evacuation as the reference point, therefore the start of each experiment is marked as $t = -10s$.

7.6 Experimentation

Within the Social Forces Model, there are many variables that are suitable for inclusion in the Mutual Information analysis, but initially solely geographic data (i.e. the 2-dimensional Cartesian coordinates of each agent) were analysed. Using the default room size from the original implementation [51] ($15m \times 15m$) as a guide to discretisation, the signals are sorted into *bins* measuring $1m$, assigns each agent to a space in a 15×15 grid, which is used to calculate the probabilistic distribution of the agents for use in the MI analysis. Mutual information for this experiment is calculated according to Equation 7.1, and the results of initial experiments can be seen in Figure 7.3.

$$I = I(X, Y) = \sum_{i,j} P(x_i, y_j) \log_2 \frac{P(x_i, y_j)}{P(x_i)P(y_j)} \quad (7.1)$$

As the results show, the trough that is apparent in the measurement of the MI in the system is not an accurate representation of the moment at which the largest amount of force is recorded, and is poorly defined when compared to the initial stage of measurement ($-10 > t < 0$). Also, as the simulation continues, it becomes impossible to ascertain the levels of force that are present solely by observing the MI of the system.

This is an unusual result, as both published works [153] and our initial investigations into the application of the MI method to complex systems, suggested that an analysis of the interdependence of the geographical variables should offer a reasonable measure of the spatial clustering within the system. The evidence for this came from investigations into the scalar noise model (SNM) [146], that suggested as the degree of noise (η) in the system decreases, the MI of the system will rise sharply. Whilst Figure 7.3 does

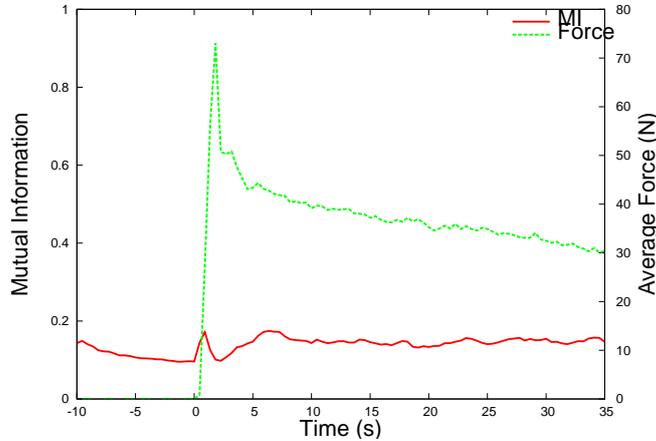


Figure 7.3: Results from the MI analysis of the SFM, MI was calculated according to Equation 7.1

show a peak in the value of the MI, at roughly the same point as the rise in average force, the change in value is very minor and the changes in crowd density are poorly defined and could not be reliably identified. There is, however, an explanation for this result.

The tendency of the MI ($I \rightarrow \infty$) in a highly ordered instance of the SNM is caused by the dimensionless nature of the agents, i.e. their lack of mass. Under exceedingly low noise conditions ($\eta \ll 1$) the particles within the SNM will, as $t \rightarrow \infty$, exhibit an extremely high level of clustering. These conditions can lead to *all* agents occupying the same *bins* (physical space), which causes the MI of the system to peak in this way. The agents in the SFM contain sufficient mass to negate this specific problem, as there is an absolute maximum number of agents that can occupy one area, but a by-product of this is the poorly defined clustering metric that the MI provides in this particular model. Therefore, to analyse the SFM correctly, a technique must be found which abstracts the spatial clustering of the particles from the calculation, and centres further on the analysis of the behaviour of the velocity vectors of the agents, a more accurate indicator of behavioural change than spatial clustering.

To rectify the reliance on spatial clustering present in our initial experiments, test runs were carried out using a combination of the coordinate and heading information from each agent, as detailed in Equation 7.2.

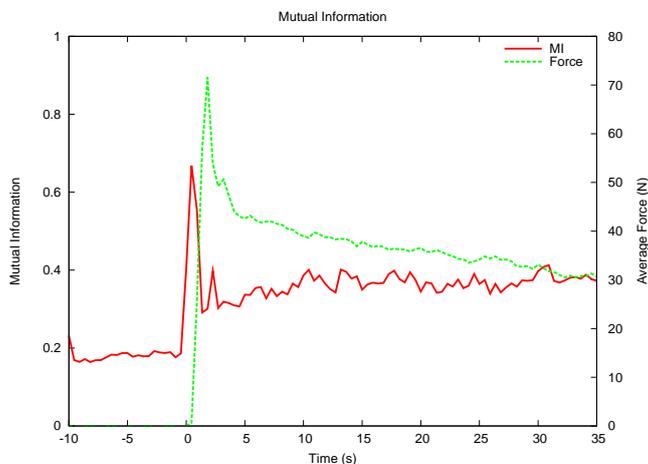


Figure 7.4: MI for this result was calculated according to Equation 7.2. Significant improvements can be seen, most importantly the improved definition of the phase-transition displayed at $0 > t < 4$. This data represents the average of 64 identical experimental runs.

$$\begin{aligned}
 I(X, \Theta) &= \sum_{i,j} P(x_i, \theta_j) \log_2 \frac{P(x_i, \theta_j)}{P(x_i)P(\theta_j)} \\
 I(Y, \Theta) &= \sum_{i,j} P(y_i, \theta_j) \log_2 \frac{P(y_i, \theta_j)}{P(y_i)P(\theta_j)} \\
 I &= I(X, \Theta) + I(Y, \Theta)
 \end{aligned} \tag{7.2}$$

With this approach, in which the coordinate and directional data on each agent is analysed in such a way that the spatial clustering is removed from the analysis, the analysis relies more heavily on the changing behaviours of the agents (more precisely, the changing velocity vectors), than previously. The results gained from this approach can be seen in Figure 7.4.

As we can see, the peak in MI using this technique is much more pronounced, with a large increase in the MI as the agents' vectors become ordered (at $t > 0$), displaying the characteristic rise in MI that is expected as a system attains order. The more relevant characteristic of the mutual information, the severe drop in MI that identifies the deterioration of the system into a state of disorder, is also more pronounced in using this technique. It can be seen that, in addition to identifying the initial rise in force,

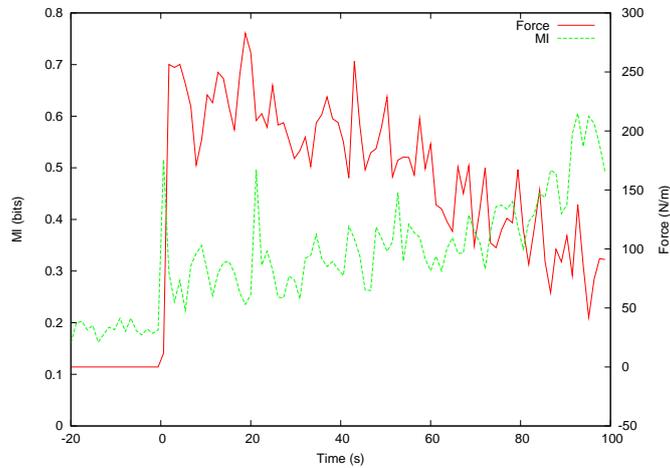


Figure 7.5: The *smoothed* output from a single run of the SFM.

the geo-directional analysis also makes obvious further peaks and troughs in the average force measure during simulations.

7.6.1 Single Run

The data shown in Figure 7.4 represents the average of 64 experiment runs using identical parameter values. It was considered, and had been confirmed by previous experiments with the Scalar Noise Model (see Chapter 6), that for a model that has been previously tested for numerical stability as few as 10 experiments runs should suffice to ensure numerical integrity. The figure of 64 runs used here was decided upon due to hardware constraints, namely that access to a 64 node HPC (High Power Computing) cluster enabled the execution of up to 64 experiments in approximately the same time as a lower number of runs. We see from this figure that MI can be used as a measure of the amount of force present within an *in silico* evacuation. To be of real use to evacuation modellers though, the technique must be shown to be applicable in *real-time*, i.e. as a simulation is running, or at least immediately after a single simulation has finished.

It has been found that the raw data output has a high level of noise, but the application of a trivial running average of the data is sufficient to *smooth* the MI output into a more readable form. The MI and force readings from a run of the SFM were therefore *coarse grained*, using the running arithmetic mean of the previous 10 MI readings. The results are shown in Figure 7.5.

As we can see, the application of a simple smoothing technique leaves the data in a readable format, and the same trends that could be seen in the aggregate data are clearly visible. Statistical analysis of these results shows that, in this case, the correlation (Pearson's) between force and MI produces a value of $R = -0.8132$ with high significance ($p = 2.2e^{-16}$), but this will obviously vary depending on the specific simulation.

7.6.2 Partial Data Analysis

Previous work [153] suggested that, in certain applications, the measurement of the MI of a proper subset of the particles within a simulation could offer an accurate indication of the point at which the phase transition occurs (relative to existing methods). To test if this applies to the Social Forces model, a technique was used which extrapolates an agent set from the time-series MI readings of a small subset of the total number of agents.

This time-series subset analysis relies on the fact that MI is entirely insensitive to the scaling of signals to which it pertains, see Equation 7.3. We can, therefore, calculate the MI of a subset of particles using the entirety of the time-series values that were collected for each agent. Using this method the exact values recorded for each of the n agents in the sample set can be used for analysis.

$$I(A, B) \equiv I(nA, nB), n \in \mathbb{R} \quad (7.3)$$

For testing, we recorded the MI values across a set of 10 agents for 100 internal time-steps during a simulation run of the SFM. The agents' positions and headings were recorded at each time-step, and the entirety of this data was used for the MI analysis.

The results gained from this analysis can be seen in Figure 7.6. The results show a peak in Mutual Information at the point at which the order in the system increases, however as order decays the MI readings fall back to a small, unchanging value.

We can see from Figure 7.6 that the partial data analysis can be used to detect the point at which evacuation behaviour changes, at $t \approx 0$, but the accuracy with which the analysis of the full system identified the changing levels of order (and therefore force) is lost at all higher values of t .

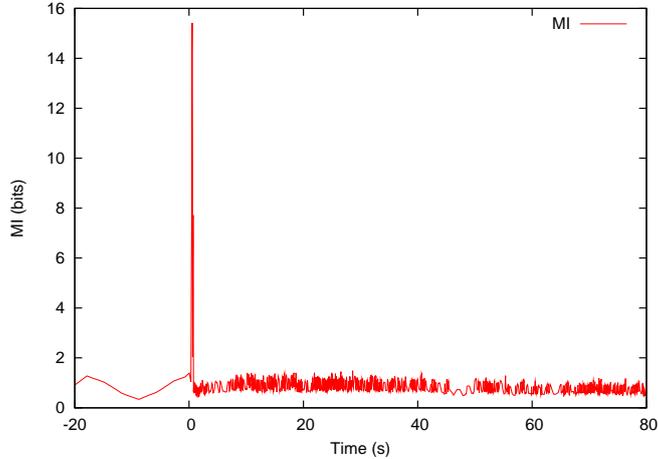


Figure 7.6: Results of the partial data MI analysis of the SFM.

It is suggested that the poor performance of the MI analysis on the SFM, relative to that presented in the literature, is partly due to the physical mass of the agents. The mass of the agents prevents large numbers of pedestrians from occupying the same physical space, which reduces the MI of the entire system by effectively removing the ability of the agents to cluster. Therefore, when taking a subset of the agents for analysis the probability of taking more than one agent from the same coarse-grained grid sections is further reduced.

7.7 Calibration

There are just three variables specific to the MI calculations, which require parametrisation, these being the size of the *bins*, or discretisation value, for the signals X , Y , and Θ . As the MI technique accepts only discrete variables for analysis, these three continuous variables are course grained previous to the application of the MI technique. We will assume that the bin size for x and y will always be equivalent, as they measure the same fundamental unit (i.e. length).

The initial values used to discretise these variables were bins x and y , b_x, b_y , were equal to exactly $1m$, giving us 15 possible values for both x and y , and b_θ was set to a value of $\frac{2\pi}{6}$, giving 6 possible values for variable θ . These values were arbitrarily chosen, and therefore we must investigate the possibility that other values for b_x, b_y , and b_θ could yield a better measure of order. To enable the best possible result from the MI analysis, calibration

experiments were carried out on our three analytic variables, these being the size of the bins used to discretise X and Y , and the size of the bins used to discretise θ . To calibrate the MI analysis, a range of values were chosen for the binning size of the variables, and both Pearson's and Spearman's correlation tests were run on each variable. For the purposes of these tests we were only concerned with the *absolute* value of the correlation co-efficient, $\therefore R = |R|$, as we do not want to assume that the correlation between the force and MI will take any specific form, and we are solely interested in the general level of correlation between the two values. The correlation analyses are carried out using both Pearson's and Spearman's coefficient calculations, as again we do not want to make assumptions about the nature of the relationship between the MI and the force during this specific analysis.

7.7.1 X and Y Discretisation

The first variables to be calibrated are the binning values for X and Y , which we will refer to as b_x and b_y . The value of b_θ for these calibration will remain at the original value chosen ($b_\theta = \frac{2\pi}{6}$). As these variables represent values measured in the same unit (metres), it was decided that for the purposes of experimentation the binning values will remain equivalent, i.e. $b_x \equiv b_y$. The range of values tested were $0.25m \leq b_x, b_y \leq 7.5m$, at intervals of $0.25m$. The correlation results are shown in Figure 7.7. All R values shown in this figure were found to be significant, i.e. $p \ll 0.05$.

We can see from the graph that there is a clear peak in the correlation for binning values in the range of $\frac{1}{2}m < b_x, b_y < 1m$, where the absolute correlation achieves a value of $R \approx 0.9$. There is a drop in correlation after this, but the general correlation rises again at a value of $b_x, b_y \approx 5$.

The results show that the size of the binning values b_x and b_y have a definite effect on the extent to which the MI can be used as a signifier of force, but that the initial (arbitrarily chosen) binning values that were used for the previous analysis were within the optimum values ranges. Initial b_x and b_y values were set at exactly $1m$, whilst b_θ had a value of 6.

7.7.2 Orientation Discretisation

Using the optimum binning values found in Section 7.7.1 ($b_x, b_y = 1m$), we ran equivalent tests to ascertain the effects that changing b_θ has on the

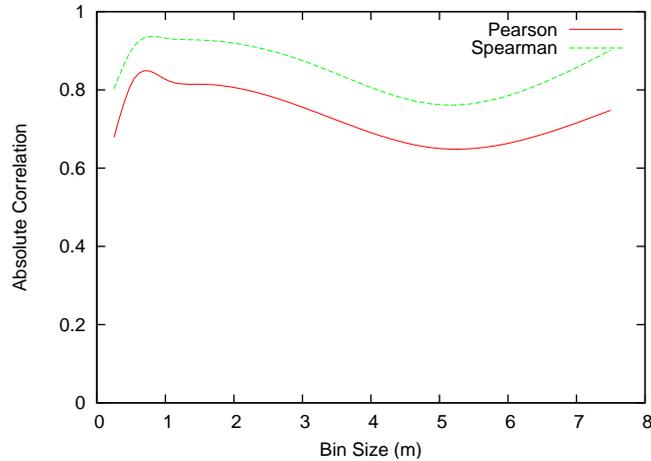


Figure 7.7: Absolute correlation (both Pearson’s and Spearman’s) between force and MI, across different binning values for variables b_x and b_y .

system as a whole. We again ran both Pearson’s and Spearman’s correlation tests on the force and MI values across a range of values for b_θ (measured in radians), the results are shown in Figure 7.8. All R values shown in this figure were found to be significant, i.e. $p \ll 0.05$.

We see a severe drop for low values of b_θ ($b_\theta \ll 1$), where the correlation values drop to as little $R \approx 0.1$. A plateau exists with values at $b_\theta > \frac{\pi}{4}$, where we see little improvement past this point, which gives us correlation values of up to $R \approx 0.95$. This bin size equates to discretising the value of θ into more than 8 bins of uniform size.

We see from Figure 7.8 that a peak of correlation occurs at $b_\theta = 1^c$, and for all values greater than this the amount of correlation between the MI and force shows a gradual degradation. Initial experiments were carried out discretising the value of θ into 6 bins of equal size, which was again arbitrarily chosen but according to this data the value does in fact represent a reasonable approximation to the optimum value for this parameter.

7.8 Negation of False Positives

As the SFM offers an accurate representation of pedestrian behaviour under emergency conditions, it can be said that the Mutual Information analysis can accurately identify the presence of order under these conditions. The question still remains about the accuracy of the MI technique under non-

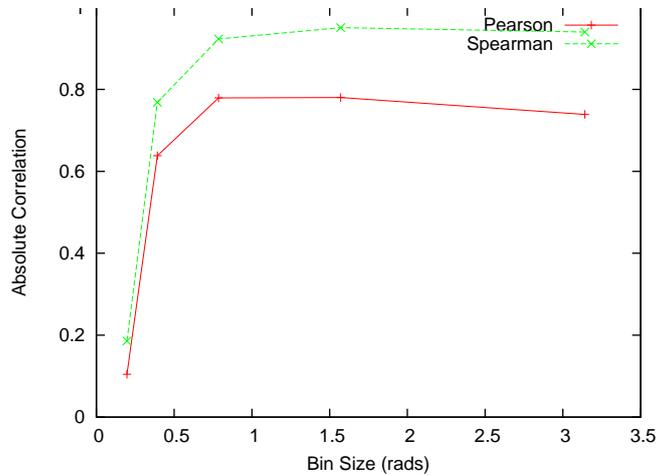


Figure 7.8: Absolute correlation (both Pearson’s and Spearman’s) between force and MI, across different binning values for variable b_θ .

emergency conditions, which are far more complex to simulate. As a means of negating the possibility of the MI technique flagging non-emergency conditions as highly dangerous, the metric was extended to classify the inherent danger of the specific evacuation (or, more accurately, a specific time in a single evacuation) as a function of both the MI of the system and the average crowd density ($\bar{\rho}$) at that point in time. This metric, which we will term $f(I, \rho_{max})$ is defined as:

$$f(I, \rho_{max}) = \frac{\rho_{max}}{I} \quad (7.4)$$

This function of the MI and average density of a specific time during an evacuation will return very high values at low I and high ρ_{max} , a highly turbulent and densely distributed evacuation, and low values when presented with very high I values and low values of ρ_{max} , a very ordered and sparsely distributed evacuation. The general distribution of $f(I, \rho_{max})$ is seen in Figure 7.9

The tracking changes between force and MI are highly pronounced, and show a correlation (Pearson’s) value of $R = 0.96$ with a high significance ($p = 2.2e^{-16}$).

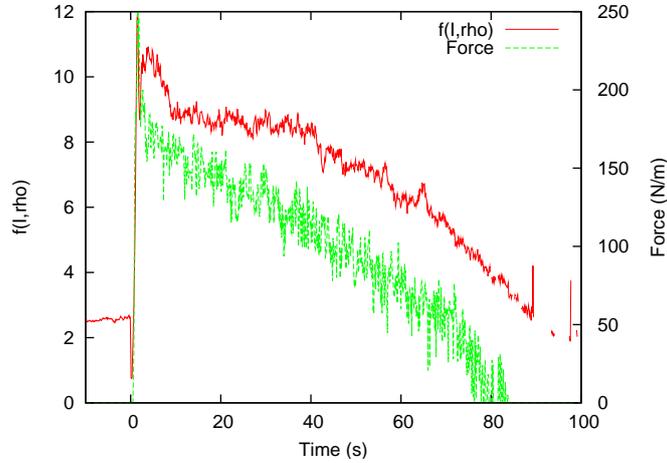


Figure 7.9: Simulation showing the changing values of Force and $f(I, \rho_{max})$

7.9 Results

As can be seen from the results in Figures 7.4 and 7.5, the potential for Mutual Information to be used as an analytical measure of the general presence of force within this simulation is clear. The experimental output shows that upon a rise in the *average force* within a simulation, there is a corresponding drop in the value of the MI. Statistical analyses have shown that the R value for this data falls at $R = 0.8132$ with a high degree of certainty ($p = 2.2e^{-16}$).

Calibration of the two discretisation values showed that the initial values chosen have given a reasonable correlation between force and MI, although these values were not optimal. Across the calibration experiments the performance of the analysis was improved by approximately 17%, with the absolute correlation between force and MI rising from a value of $R = 0.81$ to $R = 0.95$.

We have also shown that the combination of MI with the average crowd density could slightly improve the operation of the analysis, with a marginal increase in the correlation of approximately 1%. Whilst the calculation of the metric in this way shows only a marginal improvement to the analysis, it does serve to accentuate the increase and decrease in force levels in a far more *human readable* manner, i.e. as the force increases, so does $f(I, \rho)$, and vice versa. Were the MI technique to be incorporated into a simulation environment in place of explicit force calculation it is far more agreeable to

employ the metric $f(I, \rho)$ rather than I , as this can be directly substituted for force and could allow a more reliable estimate of the true force levels to be presented to the user.

7.10 Hypothesis Testing

Previous to the calibration of the system, the correlation calculated between the MI of the system and the level of force present reached a value of $R \approx -0.81$, and the statistical significance of the result was calculated at $P = 2.2e^{-16}$, which is far lower than the result required to reject the null hypothesis in either a two tailed test ($P < 0.05$) or one tailed ($P < 0.01$) test. Hence, according to these results we reject H_0 , our null hypothesis, and accept H_1 , the alternative hypothesis:

There **is** a correlation between mutual information and force.

7.11 Distribution of Force Across Agents

The force measured during these experiments, and therefore the metric which we compare MI against, is not evenly distributed across the population, with the majority of the force measured being found in the agents closest to the exit. Figure 7.10 shows the distribution of force across agents during different points during a simulation identical to those presented throughout this Chapter.

We can see from the figure that the highest level of force recorded during these simulations occur within the first 10 seconds of the evacuation, during which there are the maximum number of agents within the structure. The formation of arching can clearly be seen at $t = 7$, where the agents which form the arch are subject to far greater levels of force than others in the simulation. An arch breaking can be seen at $t = 18$, where the two agents directly inside of the exit have broken free from an arch, which has already begun to close behind them.

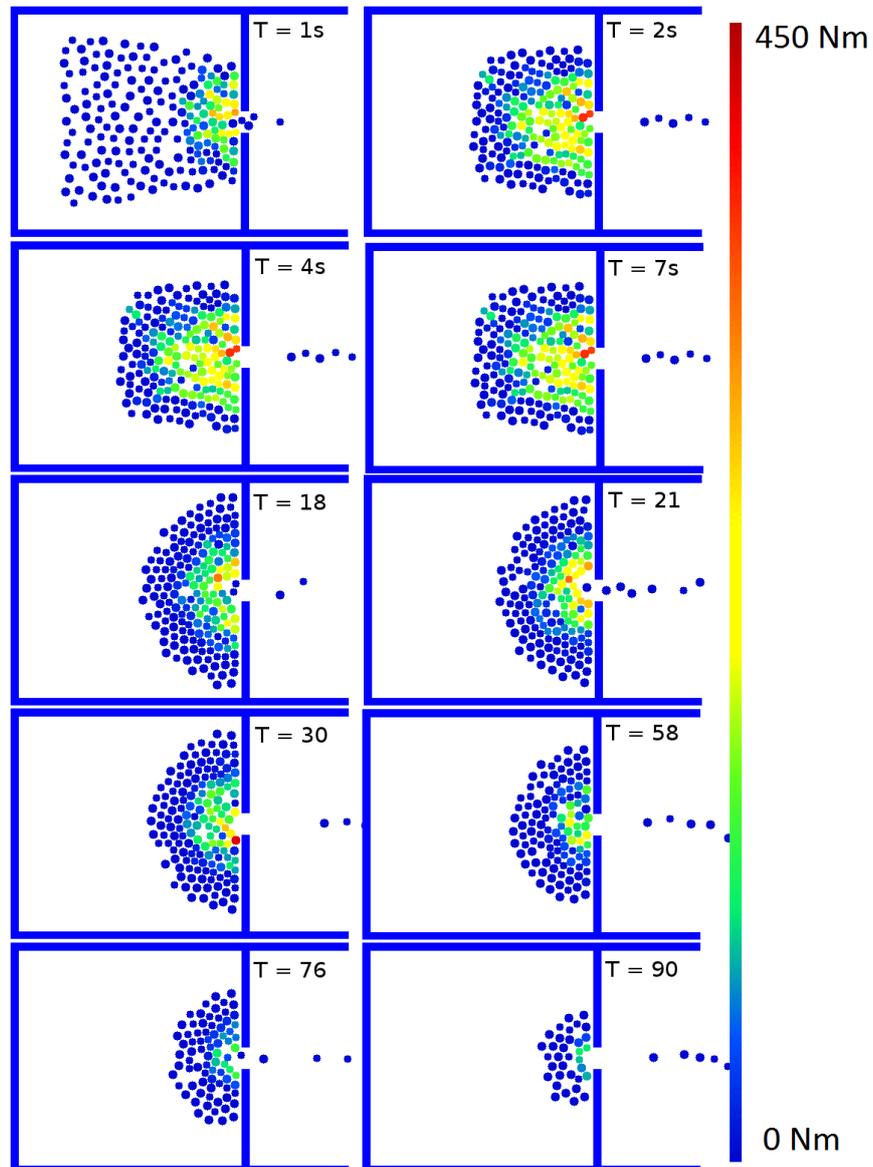


Figure 7.10: An example of the forces measured from each agent at multiple points during a simulation identical to that presented in this Chapter.

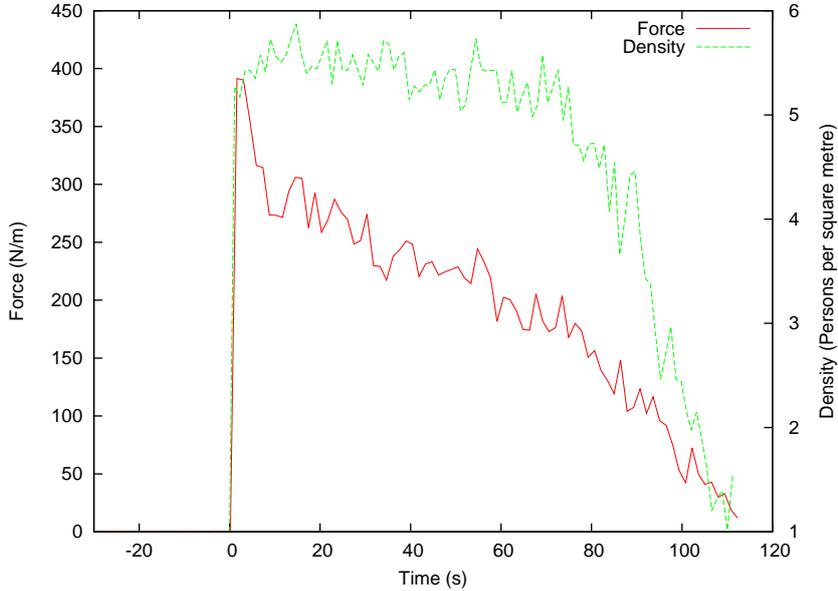


Figure 7.11: Average (arithmetic mean) values found during of 64 simulations in which the maximum crowd density (pm^{-2}) and force (Nm^{-1}) were recorded. The crowd density and the force recorded during the entire milling period ($t < 0$) were $1pm^{-2}$ and $0Nm^{-1}$ respectively.

7.12 Comparison with the Crowd Density Measure

Another technique by which force may be approximated is via a crowd density measure, a measure of the maximum crowd density across the entire evacuation. Identical simulations to those previously described were run, and at each time-step the crowd density (pm^{-2}) and the force measure was calculated for each agent within the simulation, and the maximum value found for each measure at each time-step was recorded. The experimental parameters for these simulations were identical to those defined in the original paper, and as with all previous simulations the agents' "desire to leave" was set at $V_0 = 5ms^{-1}$. The results of these experiments can be seen in Figure 7.11.

We can see from Figure 7.11 that both the force and the crowd density rise sharply at the start of the evacuation, but that the maximum force produced quickly begins to reduce as evacuees leave the room. The density metric remains high for a prolonged period, and maintains a value of between

5 and 6 pm^{-1} until approximately 80s into the evacuation. This is due to high crowd densities being achievable by a relatively small number of people, whereas high levels of force are not.

If we measure the correlation of these two metrics in the same manner as the MI and force (Pearson's Product Moment Correlation with a two-tailed significance level, $\rho = 0.01$) we find that the correlation between force and density in the SFM has an R value of $R = 0.885$, with $P = 2.2e^{-16}$. This result, as with the MI, is statistically significant, but the correlation is lower than that of the MI, which after calibration achieved ($R = 0.95$). As the number of samples used for both of these correlation calculations are extremely large (greater than 1000 samples), any difference in R value would be considered statistically significant.

A method by which the significance of this results may be formally confirmed is the Fisher's z-test, used to test that two correlation co-efficients taken from independent samples are equal [19]. The fisher's z-test returns a z-score which may be compared to the Gaussian distributed to ascertain statistical significance. The z-score obtained when comparing the Density-Force coefficient to the MI-Force coefficient is $z = -14.938$. In the Fisher's z-test the significance is confirmed when the absolute value of the z-score is greater than the calculated p-value. The p-value for this, two tailed, test is $p \leq 0.000001$, which is not unusual when dealing with such a large number of samples. From this we can establish that the difference in the two correlation coefficients is statistically significant ($|z| > p$), and that the MI-Force correlation is significantly higher than that of the Density-Force correlation ($z < 0$).

7.13 Summary

During this chapter we have defined and tested the MI analysis on a simplified evacuation consisting of a single room with 200 evacuees, and shown that the MI of the system can be used to measure the amount of force present. We have shown that the MI technique can be improved by the consideration of global densities during the analysis and that in these experiments a strong and highly significant correlation with the amount of force present has been shown.

We have also shown that the MI metric offers a better indicator of the

presence of significant levels of force than measuring the density of the evacuating crowd. Whilst this improvement is relatively small, it has been shown to be statistically significant.

These preliminary results are not considered sufficient to prove that the MI analysis can be used on a large scale evacuation, but rather represents a *proof of concept* that shows that the technique has promise and merits further testing and experimentation.

The following chapter covers the application of the MI analysis to an evacuation using a fully featured simulation environment.

Chapter 8

Analysis of a Historic Event

8.1 Introduction

In this Section we describe the results of experiments to investigate the applicability of MI as a plausible tool for crush detection. In order to ensure its broad applicability, we first show how MI is easily integrated into an existing, industry-standard simulation framework. We then validate the technique, by using it to analyse a historical event. By demonstrating that the MI technique correctly detects known incidences of crush within this scenario, we provide support for its adoption as a standard tool.

8.2 Fire Dynamics Simulator

The base simulation environment which we will employ for these extended tests is the Fire Dynamics Simulator (FDS), a fluid dynamics-based model of fire and smoke flow. Originally designed to simulate the spread of fire, heat and smoke throughout structures, the FDS environment has recently been updated to include the ability to model evacuation from structures.

The FDS+Evac module [74, 73] is the evacuation simulation extension to FDS, and is based on the *social forces model* (see Chapter 4) [53, 51] of pedestrian movement. Unlike the original SFM, presented by Helbing *et al*, the FDS+Evac environment enables engineers to accurately model the intricacies of a building's design, which allows the simulation of large-scale evacuations from complex structures.

8.3 The Station Nightclub Disaster

As mentioned previously (see Section 2.7), the Station Nightclub fire is a well-known example of the type of hazards that an emergency evacuation presents. In 2003 the Station Nightclub (Rhode Island, U.S.A.) was the scene of one the worst nightclub fires in modern American history, when a pyrotechnic device ignited a flammable polyurethane foam used for sound insulation. According to the official report into the incident [44] a crush had formed at the main escape route within 90 seconds of the start of the fire, trapping patrons inside the club as it filled with smoke. Estimates for the occupancy of the nightclub on the night of the fire vary, with figures stated by the media of between 420 and 458 people [44] this cannot be confirmed. A total of 96 of person died during the evacuation, with more than thirty bodies were recovered from the crush that formed at the main exit.

This particular event was selected on the basis of (a) the existence of a significant amount of professional film footage taken inside the nightclub during the incident¹, (b) availability of supporting witness evidence and other associated documentation, and (c) results from substantial simulation tests using FDS for fire simulation as part of the formal investigation. We therefore have information on the initial distribution of individuals at the *beginning* of the incident, visual evidence of crush *during* the incident, and the *final locations* of each of the victims, as well as a set of validated simulations with which to compare our own results.

8.4 Methodology

To confirm the ability of the mutual information technique to metricise force during a complex *in silico* evacuation two distinct evacuations of the Station Nightclub were simulated.

Firstly, an idealised simulation was designed, which represented the minimum time in which a building such as the Station Nightclub could be evacuated under emergency conditions. During this simulation it was assumed that each evacuee had “perfect” knowledge of their environment, i.e. each person within the building had knowledge of all possible escape routes. The simulation would, therefore, represent an even use of available exit capacity

¹Ironically, the film crew was present to record a documentary on nightclub safety, after a fatal incident elsewhere four days previously.

throughout the building. Also, unlike during the actual event, it is assumed that no exit becomes unusable due to fire during this simulation.

Secondly, a more realistic recreation of the events which occurred during the Station Nightclub disaster was simulated. In this scenario the evacuees did not have perfect knowledge of their environment, and each evacuee was assigned knowledge of each exit, by means of probability (see Section 8.4.1). In effect this means that evacuees may not be in a position to make optimal exit choices, as their limited knowledge of the building layout may not allow for this. To recreate the blocking of the Stage door due to fire (see Section 2.7.2) the Stage exit of the club is “closed” exactly 30 seconds after ignition, and is not usable for the rest of the simulation.

The force and MI outputs of these two evacuation simulations will then be compared, to ascertain whether the results are both visually meaningful (i.e. the MI provides a reasonable visual indicator of force present) and mathematically accurate (measured via correlation analysis).

8.4.1 Experimental Setup

The floor plan of the Station Nightclub was rendered in FDS, using official architectural plans taken from [44] (Figure 8.1). The figure of 420 is used for the number of pedestrians to be simulated, as this was used during the NIST investigations into the disaster, the true number of patrons is unknown but estimates range between 420 and 460 occupants. The initial distribution of occupants is specified so as to create high crowd densities in the Dancefloor and Sunroom areas, and lower densities in other areas. For both simulations the distribution of evacuees throughout the club was as follows; dance floor - 225 persons, main bar - 35 persons, sunroom - 110 persons, and the kitchen and the rear of the club - 50 persons. Evacuees are distributed evenly across the respective areas of the building at the start of each simulation.

The kitchen, bar and stage exits were all set to $0.9m$, and the main exit had a width of $1.8m$. The smaller opening inside the corridor leading to the main exit had a width of $0.9m$. There is no pre-evacuation time included for either of the scenarios simulated, so that all time measurements are stated from the actual start of evacuation. There is no fire present in either of the scenarios which are tested here. The imperative to evacuate is controlled via the pedestrian’s “desire to leave” (variable $V0$ in the original equations), which was set with a mean of $V0 = 5ms^{-1}$, and a standard deviation of

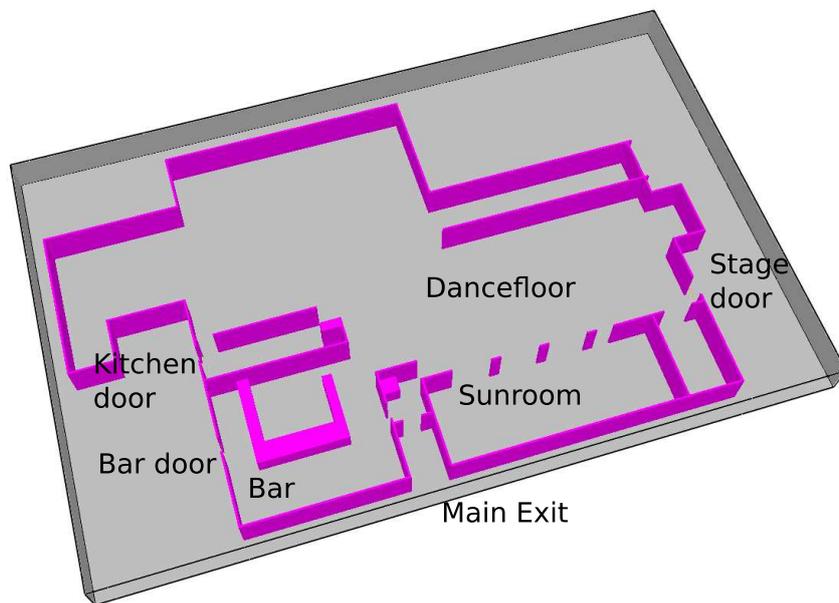
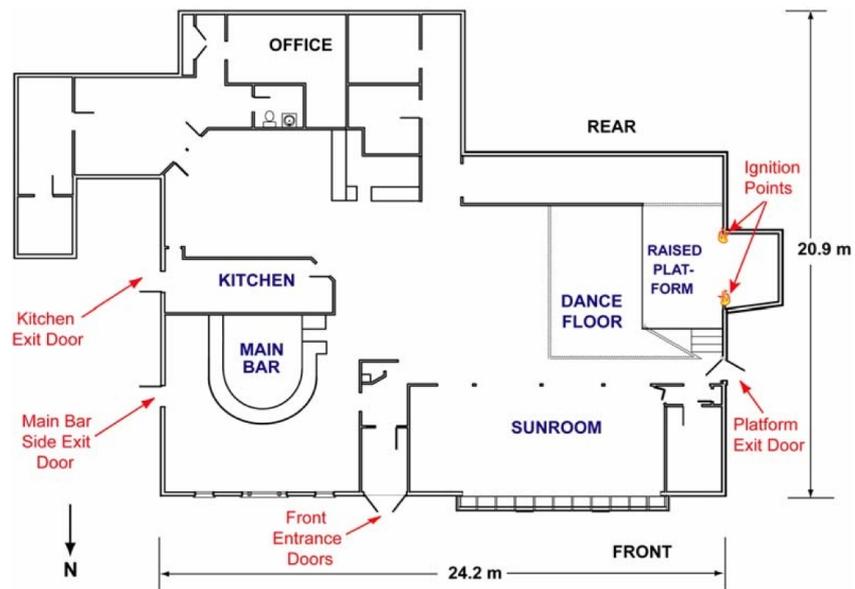


Figure 8.1: (Top) Floorplan of Station nightclub, taken from official report. (Bottom) Rendering in FDS+Evac. [44]

Area	Exit Probabilites			
	Main exit	Main bar exit	Stage exit	Kitchen exit
Dancefloor	1.0	0.5	0.4	0.0
Sunroom	1.0	0.5	0.2	0.0
Rear	1.0	0.5	0.0	0.1
Main bar	1.0	0.9	0.0	0.0

Table 8.1: The probabilities of agents starting at each area of the building knowing of the existence of of each of the possible exit routes.

$0.5ms^{-1}$.

Idealised

During the idealised simulations the assumption is made that each evacuee has a perfect knowledge of their environment, and therefore all exits from the building are known to each evacuee from the start of the evacuation. The simulations begin with the evacuees distributed as defined previously, and with the defined “desire to leave”.

Actual

At the start of these simulations, each pedestrian is assigned knowledge of exit routes via a probability. The exact probabilities given for each exit was estimated by examining firstly the placement of the agents, i.e. which part of the Nightclub they occupy at the start of the simulation, and secondly by the suggested exit use found in the literature. As an example of this, we assume that most of the people occupying the main bar room will be aware of the exit within that room, but also that the entire population would be aware of the main entrance, as this would be the most likely entry route for all occupants of the building. Therefore the 35 persons occupying the main bar would have a 90% probability of knowing of the existence of the main bar exit, and a 100% probability of knowing about the main exit. Persons in this instance who had knowledge of both exits would be able to choose the closest, or least congested, according the FDS exit choice algorithm [72].

There follows a breakdown of the persons placed in each area of the building at the start of the simulation, which is identical to that in the idealised simulation, and their respective known door probabilities, see Table 8.1.

As mentioned previously, the probabilities stated in Table 8.1 are esti-

mates based upon actual exit usage during the event and the area of the building in which agents' begin the simulation. In the case of the kitchen exit, it was assumed that only people in that specific area would know of the exit and, as this exit was only used by approximately 12 person during the event, that only a very small number of individuals near the kitchen exit knew of it's existence.

During the actual evacuation of the Station Nightclub, it was found that the Stage door (see Figure 8.1) was rendering unusable due to the proximity of the fire at approximately 30 seconds after the start of the evacuation. Whilst during these simulations we will not be modelling the fire itself, the urgency to evacuate being controlled by the agent's $V0$ parameter, the closing of this exit route is still modelled. Therefore in these simulations this exit route will be removed from use at exactly $t = 30$, and any evacuees which were planning to use this exit will be forced to find another escape route.

8.4.2 Validation

We compare the leaving profiles gained from our simulation with those obtained by similar simulations by the National Institute of Standards and Technology (NIST), and detailed in the official investigation report [44]. In these experiments, NIST investigators used both Simulex [139] and buildingEXODUS [46] to evaluate both idealised and actual evacuation scenarios. The results obtained were very similar for both packages, so we concentrate on the buildingEXODUS output.

It is considered that the ideal scenario does not require further validation, as no *complex* configuration has been required for the FDS simulation, i.e. it was an entirely standard evacuation. This cannot be said for the more realistic simulation, as assumptions have been made about exit knowledge, this scenario will therefore be compared to previous simulations in the literature as an additional validation step. The NIST simulations (detailed in [44]) which we will compare against were not identical to those run in this instance, but certain published statistics may be used to verify the simulation of our more realistic scenario. The two relevant metrics which are available are the number of evacuees remaining in the building at 90 seconds, and the number of occupants which effected egress through each exit from the structure. The 90 second marker was used in this case as this was the value

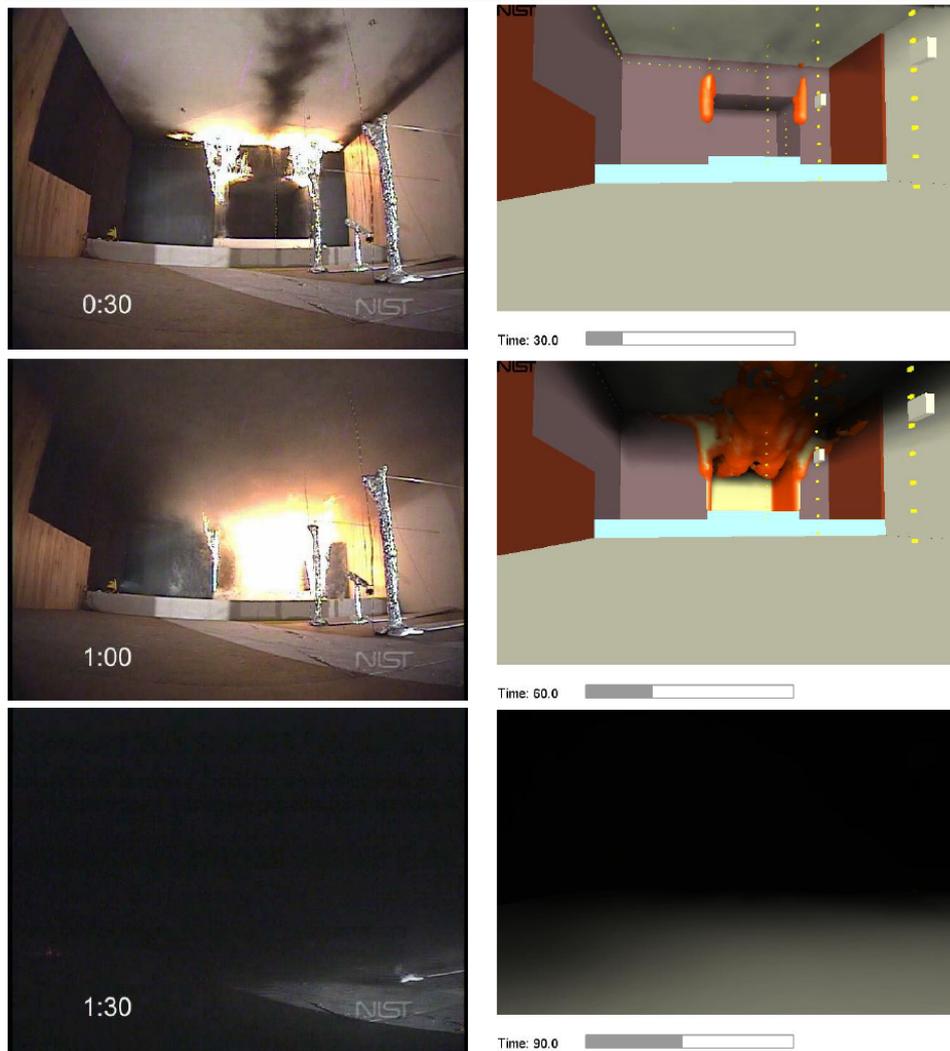


Figure 8.2: Results of NIST recreation (left) and simulation (right) of the fire and smoke spread during the initial 90 seconds of the Station nightclub fire. [44]

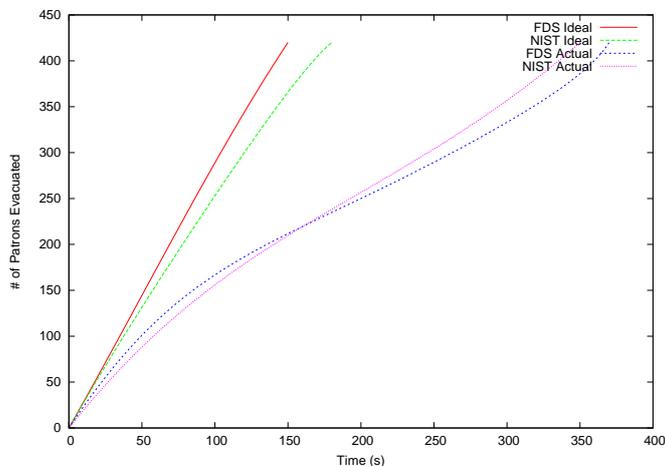


Figure 8.3: Comparison of leaving profiles between our simulation (FDS) and official NIST results.

Environment	Main door	Bar door	Stage door	Kitchen door	Total remaining at 90 seconds
Simulex	356	22	39	3	256
EXODUS	364	20	32	4	274
FDS	352	26	26	14	278

Table 8.2: Comparison of valid exit metrics for the two NIST simulations [44], using Simulex [139] and EXODUS [46], and our simulations using FDS.

chosen by NIST, due to the suggestion that the main exit became blocked at approximately this time.

We see from Table 8.2 that the number of occupants exiting through the main door and bar door during our simulations were within reasonable range of that recorded during the NIST simulations with both EXODUS and Simulex. The stage door however shows an average of 6 and 13 less evacuees when compared to EXODUS and Simulex respectively. We believe this to be due to the friction force between evacuees (not modelled by either other environment) slowing the flow rate through this door. There is also a discrepancy when examining the usage of the kitchen door, which is used by just 3-4 evacuees in the NIST simulations. The usage of this door in the FDS environment could be reduced to this level by approximately halving the probability of an evacuee knowing of this door’s existence. This would

bring the exit usage in line with NIST's figures, but it was decided that as the known exit usage during the disaster was recorded at 12 persons this makes our simulation a better representation of the actual usage during the event. The total persons remaining after 90 seconds for all three simulations are also considered to be within reasonable bounds.

We therefore conclude that, in this respect, the official NIST simulations provide a sound basis for validating our own simulations. The results of the comparison are depicted in Figure 8.3. We note that the results obtained (in terms of leaving profiles over time) are very similar to those reported by NIST, and also that the available metrics from the NIST simulations compare favourably with those from our simulation (see Table 8.2), which supports the argument in favour of the soundness of our model.

8.4.3 Detection of Crush

Having established the validity of our simulation in terms of broad outcomes, the next stage was to investigate the *emergence of crush*, and to assess if this is easily detectable using Mutual Information. To achieve this we measured the average force and the level of MI within our simulated population of 420 individuals, for both “actual” and “idealised” evacuations.

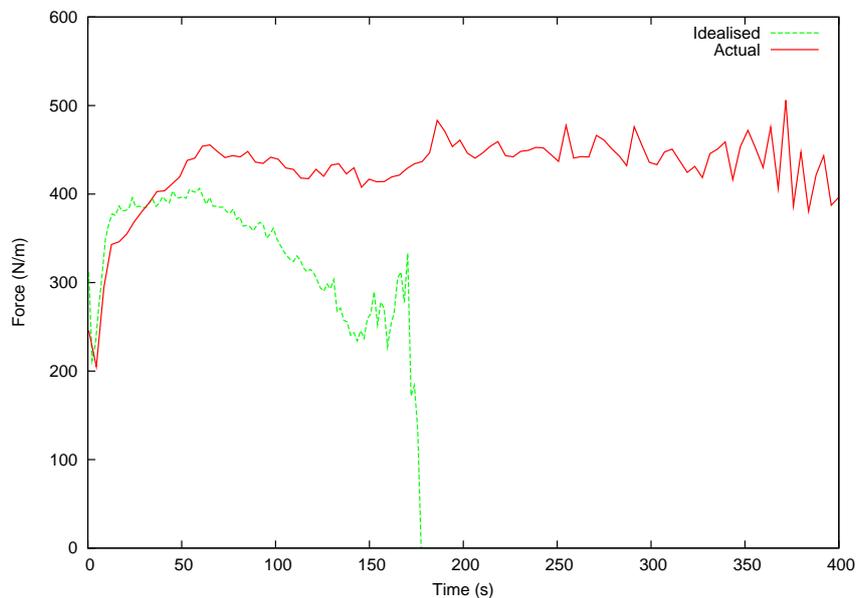


Figure 8.4: Comparison of average force between real and idealised scenarios.

We first consider the results of the force measurements, comparing them with evidence from the investigation. The force measurements for both scenarios are depicted in Figure 8.4. It should be noted that these graphics represent the average force at each point during the simulations. As simulations finish at different times, the force readings presented represent the average of all simulations where force could be measured, i.e. the average of those simulations which are *still running*.

The force measure used here is identical to that presented in the original SFM [51], that is the sum of the forces acting on any individual divided by the circumference of that individual. The measure of force in FDS has not been calibrated against empirical data, and is therefore meant primarily as a guide to the amount of force that any individual may be subject to at any one time.

Across both scenarios the levels of force initially increased as the evacuation commenced, but rapidly decayed during the idealised version of events. Force levels dropped to zero at around 175s, when everyone has left the building, which is broadly in line with the findings of the NIST idealised situation simulation ($195s \pm 7s$).

In the “actual” scenario, sharp initial rises in force were observed, which peak after around 65 seconds, or 95 seconds after ignition. This is directly in line with the findings of the official investigation, which states that a significant crowd crush occurred by the main entrance (where around a third of the fatalities occurred) beginning during the time period 71-102 seconds after ignition. Assuming an approximate 30 seconds between ignition and the start of evacuation, i.e. a 30 second pre-evacuation time, this puts the peak in force directly within NIST’s 71-102 second range at 95 seconds.

Prior to 1-1/2 minutes into the fire, a crowd-crush occurred in the front vestibule which almost entirely disrupted the flow through the main exit. Many people became stuck in the prone position in the exterior double doors . . .

The camera angle shifts away from this door after 0:07:33 (0:01:11 fire time) and does not return to the front door until 0:08:04 (0:01:42 fire time). When the camera returns at 0:08:04 (0:01:42 fire time) a pile-up of occupants is visible. Details regarding how the pile-up occurred are not available from the WPRI-TV

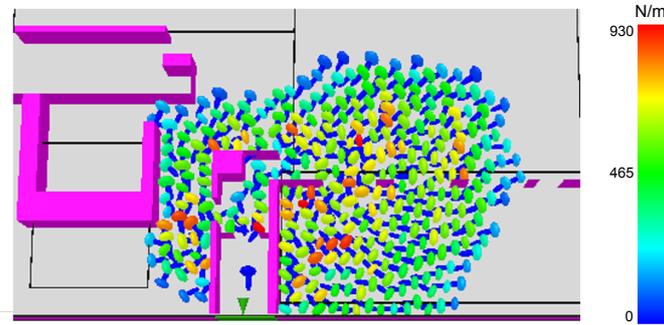


Figure 8.5: Screenshot of our fire scenario simulation after 65 elapsed seconds.

video; however, the interruption in flow of evacuating occupants apparent [in Figure shown in document] supports the contention that the disruption may have initiated early during the 31 second period when the camera was pointed elsewhere.

Grosshandler *et al* [44]

Figure 8.5, shows a screenshot of the simulation after 65 seconds, which graphically illustrates the significant crush around the main entrance and sunroom area (high levels of force are shown in red).

Again, the analysis of MI during evacuation was performed using only *observable* variables, i.e. those with values that can be obtained via direct observation of the evacuation. This is to ensure that the results were not *implementation specific*, and to maximise the possibility of applying the technique in future to other environments or video-captured data from *real-life* evacuations. Therefore, the three variables considered for analysis were the 2-dimensional *Cartesian coordinates* (x_i and y_i) of each individual, i , together with their *heading* (Θ_i). The use of *speed* within our analysis was again avoided, as during *in silico* evacuations there may often be little variation in speed during high population density simulations.

The MI was measured using Equation 8.1, taken from [153];

$$\begin{aligned}
I(X, \Theta) &= \sum_{i,j} p(x_i, \theta_j) \log_2 \frac{p(x_i, \theta_j)}{p(x_i)p(\theta_j)} \\
I(Y, \Theta) &= \sum_{i,j} p(y_i, \theta_j) \log_2 \frac{p(y_i, \theta_j)}{p(y_i)p(\theta_j)} \\
I &= \frac{I(X, \Theta) + I(Y, \Theta)}{2}
\end{aligned} \tag{8.1}$$

The MI measurements are depicted in Figure 8.6. We would expect to see, as the simulations begin, an initial rise in the MI of the system. As evacuees prepare to exit the structure they tend towards *alignment*, exhibiting similar escape trajectories to other evacuees in their locale. In a maximally efficient evacuation this period of *high order* (and high MI) would be sustained throughout, as evacuees would not alter their course in order to increase their chances of effective egress. However, in an evacuation with a great deal of competition, the order in the system quickly breaks down, as the evacuees reposition themselves in order to increase their probability of escape. MI can therefore be used as an *order parameter*, where falling values of MI signify the breakdown of order within a specific evacuation. We observe marked quantitative differences in the MI readings between the two simulations. During periods of disorder, MI will tend towards zero, whereas, during ordered segments of the evacuation, MI will rise significantly.

8.4.4 Idealised Scenario

In the idealised simulation, we see a sharp initial peak, as individuals all make for the exits at the same time. We then observe a drop, as the evacuees begin to compete for the available exit capacity. An increase in order is seen as one exit route begins to clear, creating the rise in MI at $50 < t < 75$, falling back into a state of disorder as the final evacuees clear this (main bar) exit. The MI reading then shows a progressive rise as the final evacuees exit the structure. The sharp drop in MI at the end of the simulation occurs when the number of remaining evacuees falls below the threshold at which calculating MI is viable (approximately 5 - 10 evacuees).

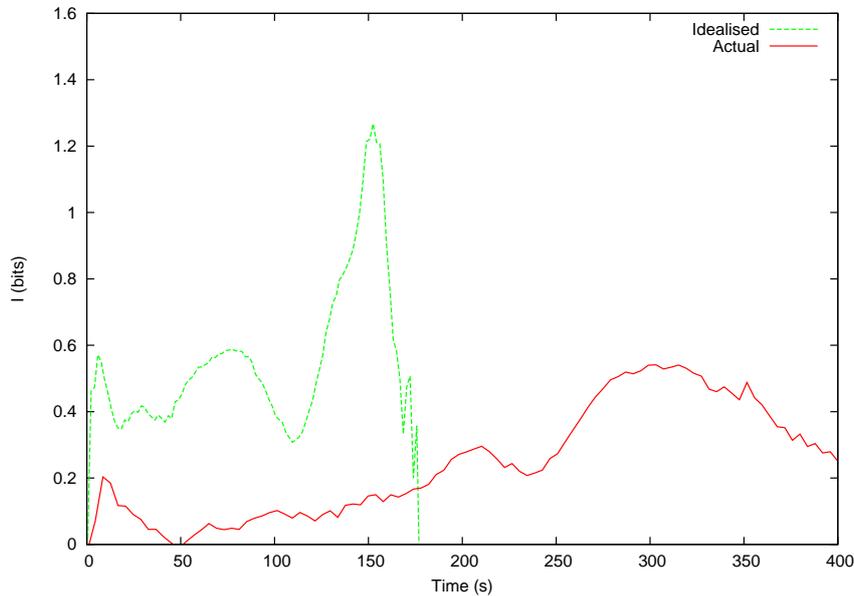


Figure 8.6: Comparison of Mutual Information between idealised and actual scenarios.

8.4.5 Realistic Scenario

The MI readings obtained from the simulation of *actual* events show a far more disordered evacuation, with an initial rise in MI (signifying order) quickly disintegrating into *disorder*. The MI reading at $t \approx 50s$ approaches zero; this period of highly disordered evacuation remains as the exits to the structure are overwhelmed (see Figure 8.5). The exit rate of evacuees during this period is also low, which is confirmed by the exit profiles (see Figure 8.3). The MI level slowly rises towards the end of the evacuation, but, notably, the higher levels of order seen in the idealised evacuation are not reached until $t \approx 300s$, 5 minutes after the start of the evacuation.

8.4.6 Correlation Analysis

A correlation analysis was performed in order to establish the relationship (if any) between force and Mutual Information. A scatterplot of force versus MI suggests the existence of a statistical association (Figure 8.7), so a Pearson’s correlation test was applied. The results of this are as follows;

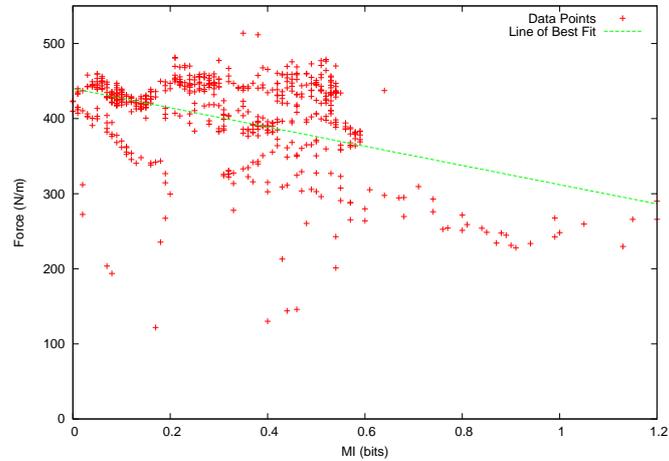


Figure 8.7: Scatterplot of Force versus Mutual Information.

$$P = 2.2e^{-16}$$

$$R_p = -0.571$$

The P-value obtained is much lower than the standard significance level for a two tailed test ($\alpha = 0.01$), ($P \ll \alpha$), which confirms the significance of the result. The correlation coefficient, $R_p = -0.571$, confirms that there exists a negative correlation between MI and force within an evacuation scenario.

8.5 False Positives

The ability of our technique to detect crush during an emergency evacuation has been demonstrated but the possibility that normal crowd movement would cause false positives still remained, i.e. what is the possibility that our analysis could flag *normal* crowd movement as having the potential to cause crush?

As crowd behaviour is an inherently complex, emergent phenomena, and relies upon a myriad of factors, it can be difficult to mathematically prove a technique to be fail-safe. What can be shown is when presented with *normal* (non-emergency) crowd movement throughout a structure, the MI technique shows sufficiently different result to that of an emergency evacuation. To

this end, the analysis of a trivial evacuation topology under non-emergency and emergency conditions was carried out. The aim of these simulations was to test the capacity of the MI technique to distinguish between complex (yet laminar) flow, and the presence of turbulence and disorder within the system.

8.5.1 Specification

The topology chosen was a single room, measuring $25m \times 50m$, with an exit placed at the east wall, and an identical entrance occupying the same position on the west. The room contains a single, large obstacle, see Figure 8.8, placed in such a way that it disrupts the flow of evacuees. The test will take two parts; firstly, the usage of the structure under normal conditions was analysed, this provided data on the ordinary usage of the structure. Secondly, the structure's evacuation capacity was overloaded to mimic an evacuation, which gave a comparative measure showing the MI readings under abnormal conditions.

The MI and physical force were recorded once for every second of evacuation time, using the same method described previously. The results of the simulations were as follows.

8.5.2 Baseline

The simulation began with 20 evacuees at the west of the structure, and additional evacuees were added through the west entrance at a rate of 10 evacuees per second of simulation time. The simulation continued for 1000 seconds, in which time the structure did not exceed capacity, and outflow continued at a steady rate. The desired leaving speed for these experiments was set to the FDS default value, $1.25ms^{-1}$, which whilst expedient, is far lower than is expected in an emergency evacuation. All other variables were also set to the FDS default values during these experiments.

This first test provided a baseline for the usage of the structure, which was used to classify the MI readings taken under abnormal crowd conditions.

Evacuation

The second test overwhelmed the capacity of the structure under evacuation conditions, for this the input rate was increased to 30 evacuees per second,

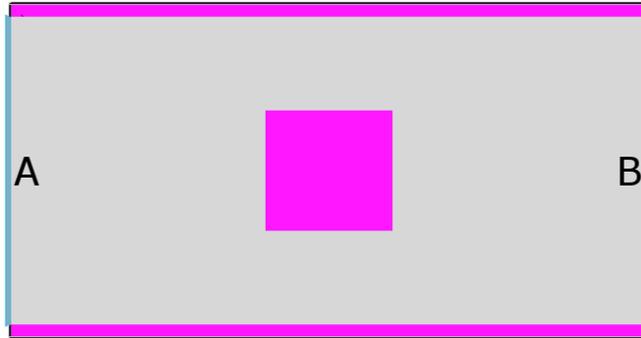


Figure 8.8: Topography of test configuration. Position A marks the centre of the entry point for pedestrians, position B marks the centre of the exit.

and the desired escape velocity was increased to $3.5ms^{-1}$, which is more in line with that of an emergency evacuation. Figure 8.10 shows the MI and force recorded during this simulation.

This second simulation was to be compared to the baseline results to see if the different MI values between simulations can be used to identify the changing levels of order present in the two instances.

False Positives Results

Simulations were run 64 times, and the results averaged to obtain the data shown here. The MI of the system under normal usage, Figure 8.10, reaches a stable level of $I \approx 0.6bits$ after roughly 50 seconds of simulation, and remains at this level for the duration of the simulation. If we compare this MI reading to that obtained from the simulations modelling the *actual* events during the Station nightclub evacuation, we see that the MI in this system is considerably higher than that recorded at even the most ordered section of that evacuation (maximum recorded during station simulation was $I \approx 0.2bits$). The force figures recorded during this test run were negligible, with the average force reading being $F \approx 30Nm^{-1}$ across the population.

The results from the simulations in which the structure is overwhelmed, Figure 8.11, show a far lower base MI reading, $I \approx 0.2bits$, after approximately 50 seconds of simulation time. The force readings, again averaged across all agents, show a drastic increase, with an average value of $F > 100Nm^{-1}$ for the majority of the simulation.

These results show that the MI analysis, in this case, is relatively insen-

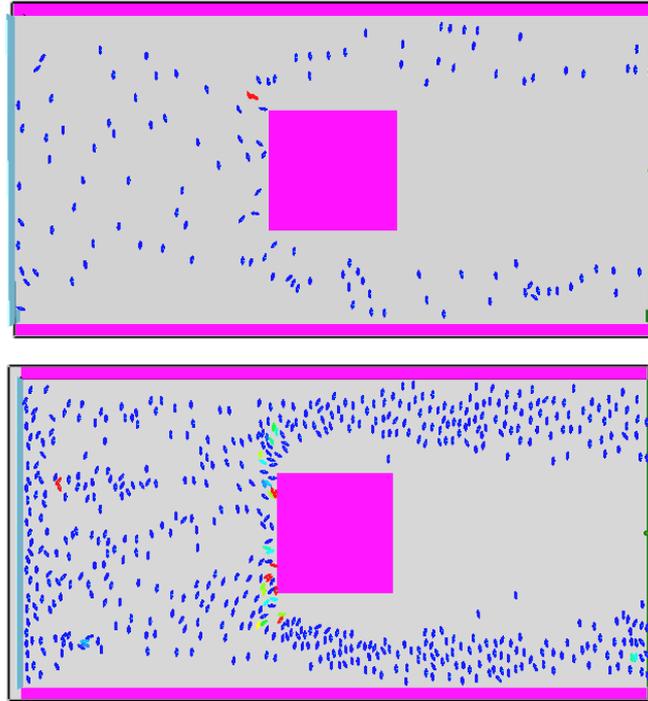


Figure 8.9: Images showing the false positives tests after a sufficient amount of time for the system to settle into a representative state. Top: First test, showing low usage. Bottom: Higher usage test.

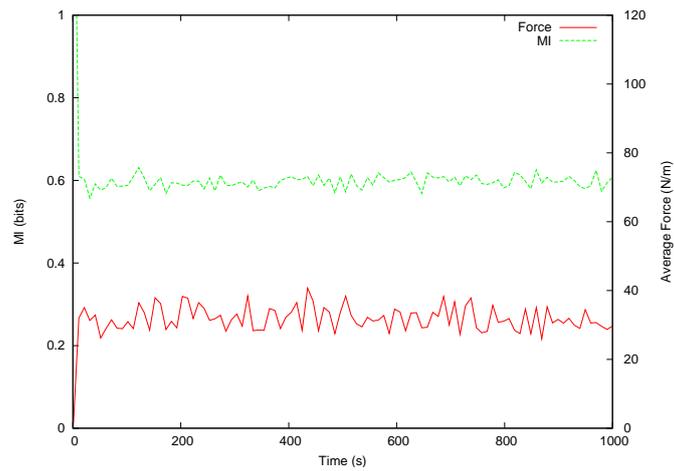


Figure 8.10: MI (green) and Average Force (red) against time for the simulation described previously.

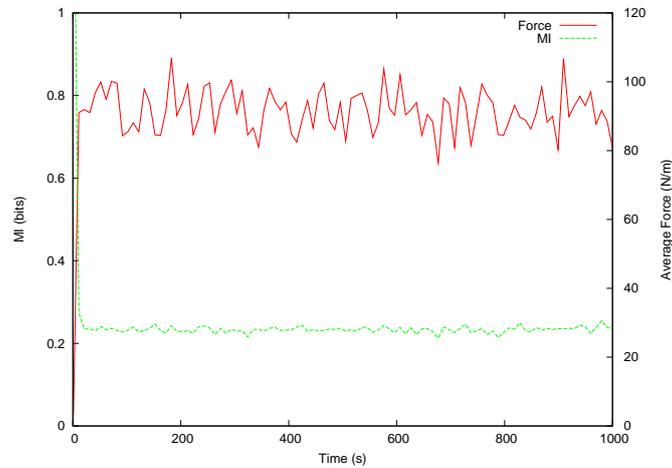


Figure 8.11: MI (green) and Average Force (red) against time for the simulation in which the evacuation capacity of the structure is overwhelmed.

sitive to minor local disorder caused by the specific geometry of the topology in question, yet is robust enough to register a lower MI level as the disorder in the system increases.

If we contrast these two sets of results, we can see the difference in the magnitude of the MI between normal and evacuation conditions is pronounced.

False Positives Conclusions

As we can see, the simulations shown here suggest that the MI technique can accurately distinguish between normal crowd movement and the disorder caused by evacuation conditions. Despite the fluctuations in crowd movement, caused by the obstacle placed in the agents' path, the MI remains at a relatively high level during the baseline simulation. The same structure, when overwhelmed, shows much lower MI readings, which is in general accord with the changes in the amount of force present during each simulation.

This is not to say that our technique could not cause false positives, as the analysis cannot be mathematically proven and it is therefore impossible to state that the technique is infallible. Were this technique to be applied to the real-time analysis of CCTV footage, the analysis under non-emergency conditions will be used to form a type of baseline for the MI of the sys-

tem under *normal* operating parameters, which will allow for more accurate identification of *abnormal* behaviours or usage patterns which will signify that a problem exists.

8.6 Summary

During this Chapter we have described the application of the MI technique to analyse a historical example of a crowd disaster . By calculating the Mutual Information of a system of interacting individuals, we are able to determine the level of disorder present within a crowd, which correlates strongly with the amount of force present. We have shown that consistently low levels of Mutual Information are correlated with high levels of force within a crowd. This method removes the need for computationally expensive physical force calculations, and allows planners to quickly and easily incorporate an explicit measure of crowd disorder and crush (see Section 3.4) into their simulation scenarios.

Chapter 9

Conclusions and Recommendations

9.1 Conclusions

This study defined a technique which offers a metric that can be used to ascertain the threat of crush during an evacuation without requiring computationally expensive physical force calculations.

1. Demonstration of the Need for a New Methodology

During this work we defined the two existing methodologies for the identification of crush, these being the explicit and implicit methodologies. We have discovered that the two methods of detection offer distinctly different benefits. The explicit methodology is the less implemented method of measurement, and relies on computationally expensive force calculations to be carried out to calculate the amount of force that each pedestrian is subject to. The implicit methodology has no such overhead, as it relies on the analysis of simulation output and the experience of the engineer to allow a subjective classification of the crush danger present in a specific simulation. The trade-off between these two techniques was identified as that of the accuracy and objectivity of the explicit method for the running speed and malleability of the implicit.

The result of this finding was the identification of the desirability of

an analytical technique which offers the objectivity, or automation, of the explicit methodology (thus going part of the way to removing the need for highly trained engineers to facilitate the simulation of large-scale evacuation scenarios), but also negates the need for physical force calculation, allowing simulations to run a much greater speed.

2. Identification of MI

The demonstration of this *gap in the market* for a low computational cost technique that metricises crush danger, lead us to investigate statistical techniques that can be applied during a simulation that allow the measurement of order and turbulence, thus allowing a comparison to physical force.

In mutual information (MI) we discovered a malleable technique which has been widely implemented for numerous classification, statistical, and measurement tasks across multiple disciplines. The MI metric had never previously been employed for the task of crush detection or evacuation analysis, and has never before been applied to the general analysis of a human or social system.

3. Proof of the Technique

The MI technique was tested on a simple evacuation topology, in which crush conditions could be ensured. This scenario was a single $225m^2$ room with one exit of just $2m$, and a population of 200 persons (see Section 7.5). In this scenario the MI technique performed excellently, offering results which correlated with the measurement of force to a degree of 0.98.

4. Analysis of a Historical Event

The Station Nightclub disaster, a well-known example of an evacuation in which the presence of crush conditions was known to have lead to serious injury and loss of life, was recreated, and analysed using the MI technique. Two sets of experiments were undertaken (see Section 8.4.1). Firstly, the evacuation was set up analyse the affects of an “ideal” evacuation of the station nightclub, i.e. an emergency evacuation in which the optimum use of the entire exit capacity was ensured. Secondly, a recreation of the evacuation conditions during the disaster itself, in which exit capacity was both reduced (as happened during

the event due to the stage door becoming inaccessible due to fire), and an uneven utilisation of the remaining exit capacity was introduced. These experiments demonstrate the suitability of the MI analysis to measure the force within a full scale evacuation, with the differences in MI between the idealised evacuation and the recreation of the conditions of the disaster itself showing highly noticeable differences in output that can be used to measure the safety of both the *idealised* scenario and the historical recreation of the disaster.

9.2 Recommendations for Further Research

The application of the MI technique to a cellular automata (CA) model will allow, for the first time, a model which can measure the danger of crush conditions forming whilst operating at a drastically reduced computation time compared to other methods. The inclusion of the MI metric in a CA model will allow myriad further possibilities for research, most notably the use of a force measurement model for experimentation with genetic algorithms, which may be required to run simulations millions of times before they satisfy the termination condition, e.g. a predetermined minimum evacuation time or acceptable levels of force throughout the simulation. This level of computation is inhibitive for current force measurement models due to their long run-times, but is possible using a modified CA that implements MI as their primary force metric.

The MI technique could be improved after an investigation into the analysis of multiple subsystems of an evacuating population. As we have seen from the work with the SFM, the changing patterns of evacuation can be seen by viewing the MI of the system as it changes over time. It has been shown that in a single room evacuation, the MI technique can identify a breakdown in ordered flow into turbulence and disorder. The logical extension of this is to discover a way by which the entire population of a large scale evacuation can be reliably segmented into different sub-populations, to allow for the calculation of both highly local and global order simultaneously. A method such as the k-means clustering algorithm could be employed for this purpose, but the limits at which this type of analysis may operate will have to be thoroughly tested.

Methods such as the Fraser-Swinney algorithm (see Section 5.8) provide

means by which a system of interconnected particles can be subdivided into uneven regions of interest, which has been shown to offer more accurate analytical results than a strictly defined analysis such as that presented here. The problem encountered with applying the Fraser-Swinney algorithm to the field of evacuation is that of agent mass, i.e. the physical space that may be occupied by a single agent at any one time. Many works that deal with dynamically subdividing the sphere of operation (such as [34]), or extrapolating results from a subset of data (such as [153]) generally deal with mass-less particles which, theoretically, can allow the entire population to occupy the exact same physical space at one time. When dealing with these systems, the dynamic subdivision of the game-space has shown reasonable results, yet the same technique offers multiple problems when applied to a system that contains, and also relies on, mass. The usefulness of such techniques, and their applicability to evacuation systems must be investigated, and possible modifications to existing algorithms which will allow their application to such systems researched.

This thesis provides the ground work from which this, and other, research can be investigated. The aim being the application of the techniques contained within this thesis in a *real world* context.

Bibliography

- [1] ABC NEWS. Verdict in e2 housing court case. Published online, September 2009. Last accessed: 14.03.11.
- [2] ANDERSON, D. *Dorland's illustrated medical dictionary*. Saunders, 2000.
- [3] BAGLIETTO, G., ALBANO, E., ET AL. Finite-size scaling analysis and dynamic study of the critical behavior of a model for the collective displacement of self-driven individuals. *Physical review. E, Statistical, nonlinear, and soft matter physics* 78, 2 (2008), 021125.
- [4] BAILLET, S., GARNERO, L., MARIN, G., AND HUGONIN, J. Combined meg and eeg source imaging by minimization of mutual information. *IEEE Transactions On Biomedical Engineering* 46 (1999), 522–534.
- [5] BBC NEWS. How the hillsborough disaster happened. Published online, April 2009. Last accessed: 15.03.11.
- [6] BBC NEWS. Stampede at german love parade festival kills 19. Published online, July 2010. Last accessed: 14.03.11.
- [7] BERK, R. A gaming approach to crowd behavior. *American Sociological Review* 39, 3 (1974), 355–373.
- [8] BETTER, O. The crush syndrome revisited (1940–1990). *Nephron* 55, 2 (1990), 97–103.
- [9] BINDER, K., AND LANDAU, D. Finite-size scaling at first-order phase transitions. *Physical Review B* 30, 3 (1984), 1477.

- [10] BLAKE, S., GALEA, E., WESTENG, H., AND DIXON, A. An analysis of human behaviour during the world trade center disaster of 11 september 2001 based on published survivor accounts. In *Proceedings of Third International Symposium on Human Behaviour in Fire* (2004).
- [11] BOES, J. L., AND MEYER, C. R. Multivariate mutual information for registration. In *Medical Image Computing and Computer Assisted Intervention*, Lecture Notes in Computer Science. Springer-Verlag, 1999.
- [12] BROWN, P., DE SOUZA, P., MERCER, R., PIETRA, V., AND LAI, J. Class-based n-gram models of natural language. *Computational Linguistics* 18 (1992), 467–479.
- [13] BROWN, R. *Social Psychology*, 2nd ed. Free Press, New York, USA, 1965.
- [14] BYARD, R., WICK, R., SIMPSON, E., AND GILBERT, J. The pathological features and circumstances of death of lethal crush/traumatic asphyxia in adults—A 25-year study. *Forensic science international* 159, 2-3 (2006), 200–205.
- [15] CASSUTO, J., AND TARNOW, P. The discotheque fire in gothenburg 1998:: A tragedy among teenagers. *Burns* 29, 5 (2003), 405–416.
- [16] CELLUCCI, C. J., ALBANO, A. M., AND RAPP, P. E. Statistical validation of mutual information calculations: Comparison of alternative numerical algorithms. *Physical Review E* 71, 6 (2005), 066208.
- [17] CHEAH, J., AND SMITH, J. Generalized m/g/c/c state dependent queueing models and pedestrian traffic flows. *Queueing Systems* 15 (1994).
- [18] CHUNXIA, L. Analysis of compressed forces in crowds. *Journal of Transportation Systems Engineering and Information Technology* 7, 2 (2007).
- [19] COHEN, J. *Applied multiple regression/correlation analysis for the behavioral sciences*, vol. 1. Lawrence Erlbaum, 2003.

- [20] COLLINS, A., AND WATERHOUSE, D. An estimation of the maximum allowable capacities of pens 3 and 4, 1990.
- [21] COMAEAU, E., AND DUVAL, R. F. Dance hall fire, gothenburg, sweden, oct 28, 1998. Published online at <http://www.nfpa.org>.
- [22] CORPORATION, L. Online, 2011. Last accessed: 9.02.12.
- [23] COVER, T. M., AND THOMAS, J. A. *Element of Information Theory*. Wiley, 1991.
- [24] CROWD MANAGEMENT STRATEGIES. Fire safety expert says e2 night-club defendants violated capacity limits; but, wait... Published online, September 2011. Last accessed: 14.03.11.
- [25] CZIROK, A., AND VICSEK, T. Collective behavior of interacting self-propelled particles. *Physica A: Statistical Mechanics and its Applications* 281, 1-4 (June 2000), 17–29.
- [26] DASGUPTA, C., AND HALPERIN, B. Phase transition in a lattice model of superconductivity. *Physical Review Letters* 47, 21 (1981), 1556–1560.
- [27] DASH, N., AND GLADWIN, H. Evacuation decision making and behavioral responses: Individual and household. *Natural Hazards Review* 8 (2007), 69.
- [28] DICKIE, J. Major crowd catastrophes. *Safety science* 18, 4 (1995), 309–320.
- [29] DICKIE, J., AND WANLESS, G. Spectator terrace barriers. *The Structural Engineer* 71, 12.
- [30] ECKHORN, R., AND POPEL, B. Rigorous and extended application of information theory to the afferent visual system of the cat. *Biology and Cybernetics* 16 (1974), 191–200.
- [31] EHTAMO, H., HELIÖVAARA, S., HOSTIKKA, S., AND KORHONEN, T. Modeling evacueeS exit selection with best response dynamics. *Pedestrian and Evacuation Dynamics 2008* (2010), 309–319.
- [32] ELSEVIER. Science direct. Online, 2011. Last accessed: 23.09.11.

- [33] FAHY, R. F. Exit 89 - an evacuation model for high-rise buildings - model description and example applications. In *Proceedings of the Fourth International Symposium on Fire Safety Science* (1994).
- [34] FRASER, A. M., AND SWINNEY, H. L. Independent coordinates for strange attractors from mutual information. *Physical Review A* 33 (1986), 1134–1140.
- [35] FRASER-MITCHELL, J. Modelling human behaviour within the fire risk assessment tool CRISP. *Fire and materials* 23, 6 (1999), 349–355.
- [36] FREUD, S. Group psychology. *The standard edition of the complete psychological works of Sigmund Freud* (1921), 1953–1974.
- [37] FRUIN, J. J. *Pedestrian Planning and Design*. Metropolitan Association of Urban Designers and Environmental Planners, Inc, NY, USA, 1971.
- [38] FRUIN, J. J. The causes and prevention of crowd disasters. In *First International Conference on Engineering for Crowd Safety* (1993), Elsevier.
- [39] GALEA, E., BLAKE, S., GWYNNE, S., AND LAWRENCE, P. The use of evacuation modelling techniques in the design of very large transport aircraft and blended wing body aircraft. *The Aeronautical Journal of the Royal Aeronautical Society* (April 2003), 207–218.
- [40] GIERLICH, B., BATINA, L., TUYLS, P., AND PRENEEL, B. Mutual information analysis - a generic side-channel distinguisher. In *Cryptographic Hardware and Embedded Systems - CHES 2008*. Springer-verlag, 2008.
- [41] GIERLICH, B., BATINA, L., AND VERBAUWHEDE, I. Revisiting higher-order dpa attacks: Multivariate mutual information analysis. *Lecture Notes in Computer Science* 5985 (2010), 221–234.
- [42] GILL, J., AND LANDI, K. Traumatic asphyxial deaths due to an uncontrolled crowd. *The American journal of forensic medicine and pathology* 25, 4 (2004), 358.

- [43] GÖNÇI, B., NAGY, M., AND VICSEK, T. Phase transition in the scalar noise model of collective motion in three dimensions. *The European Physical Journal* 157, 1 (2008), 53–59.
- [44] GROSSHANDLER, W., BRYNER, N., MADRZYKOWSKI, D., AND KUNTZ, K. Report of the Technical Investigation of the Station Nightclub fire. Tech. rep., National Institute of Standards and Technology, USA, 2005.
- [45] GROSSHANDLER, W., BRYNER, N., MADRZYKOWSKI, D., AND KUNTZ, K. Report of the technical investigation of the station nightclub fire. Published online at <http://fire.nist.gov/bfrlpubs/fire05/PDF/f05032.pdf>, Nov 2007.
- [46] GWYNNE, S., GALEA, E., LAWRENCE, P., AND FILIPPIDIS, L. Modelling occupant interaction with fire conditions using the buildingEXODUS evacuation model. *Fire Safety Journal* 36, 4 (2001), 327–357.
- [47] GWYNNE, S., GALEA, E., LYSTER, C., AND GLEN, I. Analysing the evacuation procedures employed on a thames passenger boat using the maritimeEXODUS evacuation model. *Fire Technology* 39, 3 (2003), 225–246.
- [48] GWYNNE, S., GALEA, E. R., OWEN, M., LAWRENCE, P. J., AND FILIPPIDIS, L. A review of the methodologies used in the computer simulation of evacuation from the built environment. *Building and Environment* 34 (1999), 741–749.
- [49] HARTLEY, R. V. L. Transmission of information. *Bell System Technical Journal* 7, 3 (1928), 535–563.
- [50] HAVARD, C. *Black’s medical dictionary*. Rowman & Littlefield Publishers, Inc., 1990.
- [51] HELBING, D., FARKAS, I., AND VICSEK, T. Simulating dynamical features of escape panic. *Nature* 407 (2000), 487–490.
- [52] HELBING, D., JOHANSSON, A., AND AL-ABIDEEN, H. Dynamics of crowd disasters: An empirical study. *Physical Review E* 75, 4 (2007), 46109.

- [53] HELBING, D., AND MOLNÁR, P. Social force model for pedestrian dynamics. *Physical Review E* 51, 5 (1995), 4282–4286.
- [54] HENEIN, C., AND WHITE, T. Front-to-back communication in a microscopic crowd model. *Pedestrian and evacuation dynamics 2008* (2010), 321–334.
- [55] HERMAN, E. Some laughed in e2 stampede, January, 19 2007.
- [56] HUGHES, R. The flow of human crowds. *Annual review of fluid mechanics* 35, 1 (2003), 169–182.
- [57] IEEE. Ieeexplore. Online, 2011. Last accessed: 23.09.11.
- [58] INTERNATIONAL MARITIME ORGANISATION. Ineterim guidelines for evacuation analyses for new and existing passenger ships. MSC/Circ. 1033, IMO, 2002.
- [59] IRANI, M., AND PELEG, S. Improving resolution by image registration. *Graphical Models and Image Processing* 53 (1991), 231–239.
- [60] JAYNES, E. T. Gibbs vs boltzmann entropies. *American Journal of Physics* 33 (1965), 391–398.
- [61] JEONG, J., GORE, J., AND PETERSON, B. Mutual information analysis of the EEG in patients with alzhiemer’s disease. *Clinical Neurophysiology* 112 (2001), 827–835.
- [62] JIN, Y. Studies on human behavior and tenability in fire smoke.
- [63] JOHANSSON, A., HELBING, D., AL-ABIDEEN, H. Z., AND AL-BOSTA, S. From crowd dynamics to crowd safety: A video-based analysis. *Advances in Complex Systems (ACS)* 11, 04 (2008), 497–527.
- [64] JOHNSON, N. Panic at the who concert stampede: An empirical assessment. *Social Problems* 34 (1987), 362.
- [65] KIRCHNER, A., NISHINARI, K., AND SCHADSCHNEIDER, A. Friction effects and clogging in a cellular automaton model for pedestrian dynamics. *Physical review. E, Statistical, nonlinear, and soft matter physics* 67, 5 Pt 2 (2003), 056122.

- [66] KIRCHNER, A., AND SCHADSCHNEIDER, A. Simulation of evacuation processes using a bionics-inspired cellular automaton model for pedestrian dynamics. *Physica A: Statistical Mechanics and its Applications* 312, 1-2 (2002), 260 – 276.
- [67] KISKO, T., AND FRANCIS, R. EVACNET+: A computer program to determine optimal building evacuation plans. *Fire Safety Journal* 9, 2 (1985), 211–220.
- [68] KNOPS, Z., MAINTZ, J., VIERGEVER, M., AND PLUIM, J. Normalized mutual information based registration using $i_j, k_j/i_j$ -means clustering and shading correction. *Medical image analysis* 10, 3 (2006), 432–439.
- [69] KNUDSEN, E., AND KONISHI, M. Mechanisms of sound localization in barn owl (*tyto alba*). *Journal of Computational Physiology* 133 (1979), 13–21.
- [70] KOCHER, P. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In *Advances in Cryptology - CRYPTO '96* (1996), pp. 104–113.
- [71] KOCHER, P., JAFFE, J., AND JUN, B. Differential power analysis. In *Annual international cryptology conference: CRYPTO '99* (1999), vol. 1666, Springer Berlin, pp. 388–397.
- [72] KORHONEN, T., AND HOSTIKKA, S. Fire dynamics simulator with evacuation: Fds+ evac, technical reference and users guide, 2009.
- [73] KORHONEN, T., HOSTIKKA, S., HELIOVAARA, S., AND EHTAMO, H. Fds+evac: An agent-based fire evacuation model. *Proceedings of the 4th Intl. Conference on Pedestrian and Evacuation Dynamics* (2008).
- [74] KORHONEN, T., HOSTIKKA, S., HELIÖVAARA, S., EHTAMO, H., AND MATIKAINEN, K. FDS+ Evac: Evacuation module for fire dynamics simulator. In *Proceedings of the Interflam2007: 11th International Conference on Fire Science and Engineering* (2007), pp. 1443–1448.
- [75] KOSTERLITZ, J., AND THOULESS, D. Ordering, metastability and phase transitions in two-dimensional systems. *Journal of Physics C: Solid State Physics* 6 (1973), 1181.

- [76] KRAUSZ, B., AND BAUCKHAGE, C. Automatic detection of dangerous motion behavior in human crowds. In *Advanced Video and Signal-Based Surveillance (AVSS), 2011 8th IEEE International Conference on* (2011), IEEE, pp. 224–229.
- [77] KUKLA, R., KERRIDGE, J., WILLIS, A., AND HINE, J. PEDFLOW: Development of an Autonomous Agent Model of Pedestrian Flow. *Transportation Research Record 1774*, -1 (2001), 11–17.
- [78] KULIGOWSKI, E. Review of 28 egress models. In *Workshop on Building Occupant Movement During Fire Emergencies* (2005), R. Peacock and E. Kuligowski, Eds., National Institute of Standards and Technology, Gaithersburg, MD, pp. 68–90. <http://fire.nist.gov/bfrlpubs/fire05/PDF/f05008.pdf>.
- [79] KULIGOWSKI, E. D., AND PEACOCK, R. D. A review of building evacuation models. Tech. rep., National Institute of Standards and Technology, NIST, Technology Administration, US Department of Commerce, 2005.
- [80] LAKOBA, T., KAUP, D., AND FINKELSTEIN, N. Modifications of the Helbing-Molnar-Farkas-Vicsek social force model for pedestrian evolution. *Simulation* 81, 5 (2005), 339.
- [81] LANGSTON, P. A., MASLING, R., AND ASMAR, B. N. Crowd dynamics discrete element multi-circle model. *Safety Science* 44 (2006), 395–417.
- [82] LANGTON, C. Computation at the edge of chaos: Phase transitions and emergent computation. *Physica D: Nonlinear Phenomena* 42, 1-3 (1990), 12–37.
- [83] LE BON, G. *The Crowd: A Study of the Popular Mind*. Viking, NY, USA, 1960, 1896.
- [84] LEVIN, B. Exitt—a simulation model of occupant decisions and actions in residential fires. In *Fire Safety Science-Proceedings of the Second International Symposium* (1988), pp. 561–570.

- [85] LIKAR, B., AND PERNUS, F. A hierarchical approach to elastic registration based on mutual information. *Image and Vision Computing* 19, 1 (2001), 33–44.
- [86] LO, S. M., FANG, Z., LIN, P., AND ZHI, G. S. An evacuation model: the sgem package. *Fire Safety Journal* 39, 3 (2004), 169–190.
- [87] LU, Q., GEORGE, B., AND SHEKHAR, S. Capacity constrained routing algorithms for evacuation planning: A summary of results. In *Proceeding of 9th International Symposium on Spatial and Temporal Databases (SSTD'05)* (2005).
- [88] LYNCH, J. A., PETERFY, C. G., WHITE, D. L., HAWKINS, R. A., AND GENANT, H. K. MRI-SPECT image registration using multiple MR pulse sequences to examine osteoarthritis of the knee. In *Medical Imaging: Image Processing* (1999), K. M. Hanson, Ed.
- [89] MACKAY, D., AND W, M. The limiting information capacity of a neuronal link. *Bulletin of Mathematical Biophysics* 14 (1952), 127–135.
- [90] MAES, F., COLLIGNON, A., VANDERMEULEN, D., MARCHAL, G., AND SUETENS, P. Multi-modality image registration maximization of mutual information. In *Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA '96)* (1996).
- [91] MATSUDA, H. Physical nature of higher-order mutual information: Intrinsic correlations and frustration. *Physical Review E* 62 (2000), 3096–3102.
- [92] MAURER, U., AND STEFAN, W. Unconditionally secure key agreement and the intrinsic conditional information. *IEEE Transactions on Information Theory* 45 (1999), 499–515.
- [93] MCDUGALL, W. *The Group Mind - A sketch of the principles of collective psychology, with some attempt to apply them to the interpretation of national life and character*. Cambridge University Press, Cambridge, UK, 1920.
- [94] MCGILL, W. J. Multivariate information transfer. *IEEE Transactions on Information Theory* 4 (1954), 93–111.

- [95] NADILE, L. Looking back at e2. *NFPA Journal* (March/April 2008).
- [96] NAGY, M., DARUKA, I., AND VICSEK, T. New aspects of the continuous phase transition in the scalar noise model SNM of collective motion. *Physica A: Statistical and Theoretical Physics* 373 (2007), 445 – 454.
- [97] NATIONAL SCIENCE FOUNDATION. Science and engineering indicators. Online, 2010. Last accessed: 04.06.11.
- [98] NELSON, H., AND MOWRER, F. Emergency movement. Society of Fire Protection Engineers, Bethesda, MD, 2002, pp. 367–380.
- [99] NICHOLSON, C. E., AND ROEBUCK, B. The investigation of the hillsborough disaster by the health and safety executive. *Safety Science* (1995), 249–259.
- [100] NIST. Final report of the National Construction Safety Team on the collapse of the World Trade Center tower. Tech. rep., NIST, 2005. url: <http://wtc.nist.gov/NISTNCSTAR1CollapseofTowers.pdf>.
- [101] NORTON, G. An investigation into the number of people in pens 3 and 4, 1989.
- [102] OWEN, M., GALEA, E., AND LAWRENCE, P. The EXODUS evacuation model applied to building evacuation scenarios. *Journal of Fire Protection Engineering* 8, 2 (1996), 65.
- [103] PARK, R., AND BURGESS, E. *Introduction to the Science of Sociology*. University of Chicago Press Chicago, 1921.
- [104] PAULS, J. The movement of people in buildings and design solutions for means of egress. *Fire Technology* 20, 1 (1984), 27–47.
- [105] PAULS, J. The movement of people in buildings and design solutions for means of egress. *Fire Technology* 20, 1 (1984), 27–47.
- [106] PAULS, J. Calculating evacuation times for tall buildings. *Fire Safety Journal* 12, 3 (1987), 213–236.
- [107] PAULUS, M. P., FEINSTEIN, J. S., SIMMONS, A., AND STEIN, M. B. Anterior cingulate activation in high trait anxious subjects is related to

- altered error processing during decision making. *Biological Psychiatry* 55, 12 (2004), 1179 – 1187.
- [108] POMPE, B., BLIDHE, P., HOYER, D., AND EISELT, M. Using mutual information to measure coupling in the cardiorespiratory system. *IEEE engineering in medicine and biology magazine* 17 (1998), 32–39.
- [109] PREDTECHENSKY, V., AND MILINSKY, A. *Planning for foot traffic flow in buildings*. Amerind Publishing, 1978.
- [110] QUINN, M., METOYER, R., AND HUNTER-ZAWORSKI, K. Parallel implementation of the social forces model. In *Proceedings of the Second International Conference in Pedestrian and Evacuation Dynamics* (2003), pp. 63–74.
- [111] QUISQUATER, J.-J., AND SAMYDE, D. Electromagnetic analysis (ema): measures and counter-measures for smart cards. In *E-SMART '01: Proceedings of the International Conference on Research in Smart Cards* (London, UK, 2001), Springer-Verlag, pp. 200–210.
- [112] RÉNYI, A. On measure of information theory. In *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics and Probability* (1960), pp. 547–561.
- [113] RICHMOND, P., AND ROMANO, D. Agent based gpu, a real-time 3d simulation and interactive visualisation framework for massive agent based modelling on the gpu. In *Proceedings International Workshop on Supervisualisation* (2008).
- [114] ROGSCH, C., SCHRECKENBERG, M., TRIBBLE, E., KLINGSCH, W., AND KRETZ, T. Was it panic? an overview ab out mass-emergencies and their origins all over the world for recent years. In *Proceedings of the 4th Intl Conference on Pedestrian and Evacuation Dynamics* (Berlin, Germany, 2008), Springer.
- [115] RONDONI, L., AND COHEN, E. Gibbs entropy and irreversible thermodynamics. *Nonlinearity* 13 (2000), 1905.
- [116] RUECKERT, D., SONODA, L., HAYES, C., HILL, D., LEACH, M., AND HAWKES, D. Nonrigid registration using free-form deformations:

- application to breast mr images. *Medical Imaging, IEEE Transactions on* 18, 8 (1999), 712–721.
- [117] SCHADSCHNEIDER, A., KLINGSCH, W., KLÜPFEL, H., KRETZ, T., ROGSCH, C., AND SEYFRIED, A. Evacuation dynamics: empirical results. *Modeling and Applications (in Encyclopaedia of Complexity and System Science, Ed: Meyers, B., Springer, Berlin* (2008).
- [118] SCHWEINGRUBER, D., AND WOHLSTEIN, R. T. The madding crowd goes to school: Myths about crowds in introductory sociology textbooks. *Teaching Sociology* 33, 2 (2005), 136–153.
- [119] SHANNON, C. A mathematical thoery of communication. *The Bell System Technical Journal* 27 (1948), 379–423,623–656.
- [120] SHANNON, C. E. Communication theory of secrecy systems. *Bell System Technical Journal* 28 (1949), 656–715.
- [121] SIGHELE, S. *La Foule criminelle: Essai de Psychologie Collective*. Paris: Felix Alcan., 1892.
- [122] SIME, J. An occupant resp onses escap e time (oret) model. *Proceeding of the First International Symposium* (1998).
- [123] SMITH, J., AND TOWSLEY, D. The use of queuing networks in the evaluation of egress from buildings. *Environment and Planning B: Planning and Design* 8, 2 (1981), 125–139.
- [124] SMITH, R., AND DICKIE, J. *Engineering for crowd safety*. Elsevier, 1993.
- [125] SMITH, R., AND GAMES, G. Collapse load calculations for barrier 124a. Tech. rep., HSE Report No. IR/L/ME/MM/89, 1989.
- [126] SONG, W., YU, Y., XU, X., AND ZHANG, H. Evacuation analysis of a commercial plaza with cafe model. *International Journal on Engineering Performance-Based Fire Codes* 7, 4 (2005), 182–191.
- [127] SOUCEK, G. *Chicago Calamities: Disaster in the Windy City*. History Press, 2010.

- [128] SRINIVASA, S. A review on multivariate mutual information. Available online at <http://www.nd.edu/~jnl/ee80653/Fall12005/tutorials/sunil.pdf>.
- [129] SRINIVASA, S. A review on multivariate mutual information. University of Notre Dame, Notre Dame, Indiana, USA, 1995.
- [130] STEIN, R., FRENCH, A., AND HOLDEN, A. The frequency response, coherence, and information capacity of two neuronal models. *Journal of Biophysics* 12 (1972), 295–322.
- [131] STEVENSON, A., AND WAITE, M. *Concise Oxford English Dictionary*. Oxford University Press, 2011.
- [132] STUDHOLME, C., HILL, D., AND HAWKES, D. Overlap invariant entropy measure of 3d medical image alignment. *Pattern Recognition* 32 (1999), 71–86.
- [133] STUDHOLME, C., HILL, D. L. G., AND HAWKES, D. J. Incorporating connected region labelling into automated image registration using mutual information. In *Mathematical Methods in Biomedical Image Analysis*, A. A. Amini, F. L. Bookstein, and D. C. Wilson, Eds. IEEE Computer Society Press, 1996, pp. 23–31.
- [134] TARDE, G. Les crimes des foules. *Archives de l'Anthropologie Criminelle* 7 (1892), 353–386.
- [135] TARDE, G., AND CLARK, T. *On communication and social influence: Selected papers*. University of Chicago Press, Chicago, 1969.
- [136] TAYLOR, L. The Hillsborough Stadium Disaster. *Final Report on Enquiry by Rt. Hon. Justice Taylor, London, UK: Home Office, HMSO* (1989).
- [137] TAYLOR, P. The Hillsborough Stadium Disaster: Interim Report. *London: HMSO* (1989).
- [138] TAYLOR, S., AND BRITAIN, G. *The Hillsborough Stadium Disaster: 15 April 1989: Inquiry by the Rt Hon Lord Justice Taylor: Final Report: Presented to Parliament by the Secretary of State for the Home Department by Command of Her Majesty January 1990*. HMSO, 1990.

- [139] THOMPSON, P., AND MARCHANT, E. A computer model for the evacuation of large building populations. *Fire Safety Journal* 24, 2 (1995), 131–148.
- [140] THOMPSON, P., AND MARCHANT, E. Testing and application of the computer model SIMULEX. *Fire Safety Journal* 24, 2 (1995), 149–166.
- [141] THOMPSON REUTERS. Web of knowledge. Online, 2011. Last accessed: 23.09.11.
- [142] TONG, D., AND CANTER, D. The decision to evacuate: a study of the motivations which contribute to evacuation in the event of fire. *Fire Safety Journal* 9, 3 (1985), 257–265.
- [143] TSUJISHITA, T. On triple mutual information. *Advances in Applied Mathematics* 16 (1995), 269–274.
- [144] TURNER, R. H., AND KILLIAN, L. *Collective behavior*. Prentice-Hall, 1957.
- [145] UK GOVERNMENT. Report of the official account of the bombings in london on 7th july 2005. *The Stationary Office* (2006).
- [146] VICSEK, T., CZIROK, A., BEN-JACOB, E., COHEN, I., AND SCHOCHECHET, O. Novel type of phase transition within a system of self-driven particles. *Phys. Rev. Lett* 75 (1995), 1226–1229.
- [147] VIOLA, P., AND WELL, W. Alignment by maximization of mutual information. *International Journal of Computer Vision* 24 (1997), 137–154.
- [148] VON NEUMANN, J., AND BURKS, A. *Theory of self-reproducing automata*. University of Illinois Press, 1966.
- [149] WALKER, G. 3 the ibrox stadium disaster of 1971. *Soccer & Society* 5, 2 (2004), 169–182.
- [150] WEI-GUO, S., YAN-FEI, Y., BING-HONG, W., AND WEI-CHENG, F. Evacuation behaviors at exit in ca model with force essentials: A comparison with social force model. *Physica A* 371 (2006), 658–666.

- [151] WEINSTEIN, J. Video of e2 stampede shown on first day of testimony, 2007.
- [152] WELLS, M., VIOLA, P., ATSUMI, H., NAKAJIMA, S., AND KIKINIS, R. Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis* 1, 1 (1996), 35–51.
- [153] WICKS, R. T., CHAPMAN, S. C., AND DENDY, R. O. Mutual information as a tool for identifying phase transitions in dynamical complex systems with limited data. *Physical Review E* 75, 5 (2007).
- [154] WILGOREN, J. 21 die in stampede of 1,500 at Chicago nightclub. *New York Times* (2003). February 18th.
- [155] ZHENA, W., MAOA, L., AND YUANC, Z. Analysis of trample disaster and a case study mihong bridge fatality in china in 2004. *Safety Science* 46, 8 (2008), 1255–1270.

Appendix A

Publications

Mutual Information for the Detection of Crush Conditions - Journal article

This article contains much of the material covered in Chapter 8 of this thesis. It was published in PLoSone in December 2011.

Details

Author(s) : Harding, P. J., Gwynne, S., and Amos, M.

Title: Mutual Information for the Detection of Crush

Journal: PLoS ONE (Public Library of Science)

Year: 2011

Issue: 6

Number: 12

DOI: [10.1371/journal.pone.0028747](https://doi.org/10.1371/journal.pone.0028747)

Mutual Information for the Detection of Crush

Peter Harding¹, Steve Gwynne², Martyn Amos^{1,*}

1 School of Computing, Mathematics and Digital Technology, Manchester Metropolitan University, Manchester, UK.

2 Hughes Associates, London, UK.

* E-mail: M.Amos@mmu.ac.uk

Abstract

Fatal crush conditions occur in crowds with tragic frequency. Event organizers and architects are often criticised for failing to consider the causes and implications of crush, but the reality is that both the prediction and prevention of such conditions offer a significant technical challenge. Full treatment of physical force within crowd simulations is precise but often computationally expensive; the more common method of human interpretation of results is computationally “cheap” but subjective and time-consuming. This paper describes an alternative method for the analysis of crowd behaviour, which uses information theory to measure crowd disorder. We show how this technique may be easily incorporated into an existing simulation framework, and validate it against an historical event. Our results show that this method offers an effective and efficient route towards automatic detection of the onset of crush.

Introduction

Overloading pedestrian routes can quickly lead to the development of *crush conditions* (should the necessary conditions be evident), as observed in the Hillsborough [1], Station nightclub [2] and Saudi Arabian Hajj [3] incidents. A more sophisticated understanding of how crush conditions form is therefore critical for the architectural design of highly-populated, contained regions (such as ships, nightclubs and stadia), as well as for the planning of events and formulation of incident management procedures. Using this insight, we can begin to understand *how* and *why* crush forms as a result of poor design or lack of strategic planning. A first step towards this is a method for *detecting* the early-stage formation of crush, which is the problem we address here.

Computer-based simulation studies are often used to analyse the movement of individuals in various scenarios, often as part of a performance-based design. Such work encompasses the study of historical events [3], the examination of evacuation procedures [4], and the design of aircraft [5]. Existing simulation frameworks include EXODUS [6], PEDFLOW [7] and EVACNET [8], and these offer a range of “real world” features, including exit blockage/obstacles, occupant impatience and route choice [9]. However, the phenomenon of *crush* is one that has received relatively little attention so far from the designers of evacuation simulations, and any simulations do not *explicitly* consider the effects of crush.

We therefore seek a method for the detection of crush conditions that may be easily integrated into existing software for crowd simulation. Such a method will have a significant impact on both computer-based evacuation studies and real-time analysis of video images (facilitating, for example, the development of automated crush alarms based on CCTV images). In this paper we give a description of our proposed method, which is based on applying *information theory* to a system of interacting particles. We show how our method may be easily integrated into an existing simulation framework, and test it using details of an historical event. Simulation results show that our method provides an excellent “early warning” indicator of the emergence of crush conditions.

Methods

Within an evacuation simulation, the two distinct states of a crowd are characterised by the behaviour of individuals. Under “normal” conditions, crowd flow is highly *ordered*, with the orientation and speed of a specific individual being similar to that of those in their immediate locality. The onset of more *turbulent* flow sees individuals exhibit a marked change in behaviour, as they change speed and alter course in order to avoid others [3]. We therefore wish to identify these distinct states, and achieve this by applying statistical analysis techniques to the movement of individuals within crowds.

In the general case, the Mutual Information (MI) of two *discrete* time-series variables, A and B, is defined as:

$$I(A, B) = \sum_{i,j} p(a_i, b_j) \log_n \frac{p(a_i, b_j)}{p(a_i)p(b_j)} \quad (1)$$

where $p(a_i)$, $p(b_j)$, and $p(a_i, b_j)$ are the individual probability and joint probability distributions of A and B. In general terms, MI quantifies the interdependence of two variables; therefore if A and B are entirely *independent*, then $I(A, B) = 0$, but in *all* other cases $I(A, B) > 0$. In the context of crowd behaviour, we measure the interdependence of both *location* and *heading* over a population of individuals, in order to establish the degree of order within the crowd. An *ordered* crowd (e.g., one exhibiting stable laminar flow) will have relatively high MI, since individuals are moving in a synchronised fashion. An entirely *disordered* (i.e. turbulent) crowd will exhibit an MI value of zero, since individuals are acting completely independently of one another. We seek to detect the onset of such turbulence, as an early indicator of crush.

The three variables considered for analysis are the 2-dimensional *Cartesian coordinates* (x_i and y_i) of each individual, i , together with their *heading* (Θ_i). We forego the use of *speed* within our analysis, as there is often little variation in speed during incidents with high population density. We measure MI using Equation 2, taken from [10]:

$$\begin{aligned} I(X, \Theta) &= \sum_{i,j} p(x_i, \theta_j) \log_2 \frac{p(x_i, \theta_j)}{p(x_i)p(\theta_j)} \\ I(Y, \Theta) &= \sum_{i,j} p(y_i, \theta_j) \log_2 \frac{p(y_i, \theta_j)}{p(y_i)p(\theta_j)} \\ I &= \frac{I(X, \Theta) + I(Y, \Theta)}{2} \end{aligned} \quad (2)$$

The base simulation environment used is the Fire Dynamics Simulator (FDS) [11], a fluid dynamics-based model of fire and smoke flow. The FDS+Evac module [12] is an agent-based evacuation simulation extension for FDS, and is based on the established *social forces model* [13, 14] (SFM) of pedestrian movement. An important feature offered by FDS+Evac is that of route selection, which allows the user to embed “knowledge” about available exits into each individual.

Importantly, the evacuation module for FDS includes the calculation of *physical forces*, which we will need in order to assess the correlation between crush conditions and mutual information. We integrate the MI analysis code into the FDS environment as a set of natively coded (FORTRAN 90) libraries. As the technique is entirely passive, i.e. it will not affect the results of the evacuation, there are no concerns regarding the effect this may have on the *outcome* of the simulations (although there is clearly a small overhead incurred by the MI calculations). The MI of the system is calculated at every *simulation* time step, and the results averaged over 100 time steps before being recorded. This equates to one MI reading

per second of *real-life* evacuation time, which gives sufficient granularity. We record the average physical force within a simulation in the same way. In what follows, we use the default FDS+Evac parameter values, as described in [15]. All simulation code is available at <http://code.google.com/p/mi-crush/>

Results

In order to validate the technique, we choose a well-documented incident that illustrates the significant hazards that an emergency evacuation may present. In 2003, the Station Nightclub (Rhode Island, USA) was the scene of one the worst nightclub fires in recent history, when a pyrotechnic device, used by the rock band Great White, ignited sound insulation foam in the walls and ceiling of the venue. According to the official report into the incident [2], a crush formed at the main escape route within 90 seconds of the start of the fire, trapping patrons inside the club as it filled with smoke. Estimates of the nightclub occupancy vary between 440 and 460; a total of 96 people died during the incident.

We select this particular event on the basis of (a) the existence of a significant amount of professional film footage taken inside the nightclub during the incident - ironically, the film crew was present to record a documentary on nightclub safety, after a fatal incident elsewhere four days previously, (b) availability of supporting witness evidence and other associated documentation, and (c) results from substantial simulation tests using FDS as part of the subsequent (extensively documented) formal investigation. We therefore have information on the initial distribution of individuals at the *beginning* of the incident, visual evidence of crush *during* the incident, and the *final locations* of each of the victims, as well as an additional set of validated simulations with which to compare our own results. We first ensure that our simulation produces *valid outcomes* in terms of evacuation profiles (by testing it against the historical event), and then specifically test the MI technique.

Exit profile validation

Here, we first ensure that our own simulation produces *general evacuation outcomes* that are in line with reality (as well as previously validated simulations). We begin by rendering the floor plan of the Station in FDS, using official architectural plans taken from [2] (Figures 1- 2). We use a figure of 450 for the number of agents to be simulated, and their initial distribution is specified according to [2] (i.e., with high crowd densities in the Dancefloor and Sunroom areas, and lower densities in other areas).

We run two sets of experiments; the first, *idealised* set is designed to provide baseline evacuation data, and the second set replicates, as closely as possible, the conditions and events in the nightclub during the event. Investigation findings into the spread of the fire suggest that the Stage door became impassable 30 seconds from the start of the incident, so we reflect this fact in our simulation by closing that exit after that period has elapsed. The official investigation was able to identify the exit paths for 248 of the 350 people who escaped from the building. The distribution of evacuees through the three other available exit routes was found to be non-uniform, with estimates of between one-half and two-thirds of patrons attempting to leave via the familiar main exit, rather than the under-utilised (and less familiar) Main Bar and Kitchen doors. Reports suggest that only 12 people left via the Kitchen door during the evacuation. In order to simulate this distribution of path choices, patrons are assigned a *probability of knowledge* for each exit route. Exactly 12 evacuees are made aware of the existence of the Kitchen exit, and of the remaining patrons, 100% are given knowledge of the main door, 50% are given knowledge of the main bar door, and 25% are given knowledge of the stage door. On the other hand, the idealised evacuation was structured as follows: there was no blocking of the Stage door, and agents in the simulation had full knowledge of all exit routes. This scenario represents the minimum time it would take to evacuate 450 people from the Station Nightclub, with optimum use made of available exit structures and no hindrance from fire, smoke, or unfavourable environmental conditions.

We compare our simulation results with those obtained by the National Institute of Standards and Technology (NIST), and detailed in the official investigation report [2]. In these experiments, NIST investigators used both Simulex [16] and buildingEXODUS [17] to evaluate both idealised and realistic evacuation scenarios. The results obtained were very similar for both packages, so we concentrate on the buildingEXODUS output. Within the “realistic” simulation, occupants were instructed to always select the nearest exit, and the Stage door was also closed after 30 seconds. In the NIST simulation, 91 simulated occupants left via the building front door, which is precisely the number reported in the official investigation. Thirty-five simulated occupants used either the platform door or the kitchen door, which, again, is consistent with the evidence.

We therefore conclude that the official NIST simulations provide a sound basis for validating our own simulations. The results of the comparison are depicted in Figure 3. We note only that the results obtained (in terms of leaving profiles over time) are very similar to those reported by NIST, which supports the argument in favour of the soundness of our model.

MI technique validation

Having validated the model in terms of broad outcomes, we now consider the problem of Mutual Information “false positives” (that is, a situation in which “normal” pedestrian flow is incorrectly flagged, via MI measurement, as potentially leading to crush). In order to mitigate against this, we first benchmark the method using a trivial evacuation topology under both emergency and non-emergency conditions. This structure is designed to test the capacity of the MI technique to distinguish between laminar flow and turbulence within the system.

The topology chosen is a single room, measuring $25m \times 50m$, with an exit placed at the east wall, and an identical entrance occupying the same position on the west (Figure 4). The room contains a single, large obstacle, placed in such a way that it disrupts the flow of evacuees. We then perform two sets of runs; the first set tests usage of the structure under “normal” conditions, and the second set tests it during an evacuation situation.

For the normal situation, we begin with 20 evacuees at the west of the structure, with additional evacuees added through the west entrance at a rate of 10 evacuees per second of simulation time. The desired leaving speed for is initially the FDS default value of $1.25ms^{-1}$. All other parameters are set at the FDS default values. For the simulated evacuation, we aim to overwhelm the capacity of the structure by increasing the input rate to 30 evacuees per second, and increasing the desired escape velocity to $3.5ms^{-1}$.

We now compare the results of both sets of runs to see if the values for MI differ between them (and thus may be used to identify the different levels of order observed in each situation). Each situation is simulated 50 times, and the results averaged. The MI of the system under *normal* usage (Figure 7 reaches a stable level of $I \approx 0.6$ bits after roughly 50 seconds of simulation (after which point there are sufficient individuals in the system to render the results meaningful), and remains at this level for the duration of the simulation. The force figures recorded during this test run are negligible, with the average force reading being $F \approx 30Nm^{-1}$ across the population.

The results from the simulations in which the structure is overwhelmed (Figure 8) show a far lower basal MI reading, $I \approx 0.2$ bits, after approximately 50 seconds of simulation time. The force readings, again averaged across all agents, show a significant increase, with an average value of $F > 100Nm^{-1}$ for the majority of the simulation.

These results confirm that MI analysis is relatively insensitive to minor local disorder, but is robust enough to register a lower MI level as disorder in the system increases. We observe a significant difference in MI between normal and evacuation conditions, leading us to conclude that our method is unlikely to generate false positive results, and is capable of detecting the disorder present at the onset of crush.

Crush detection

The next stage is to specifically investigate the emergence of crush in our “real-world” scenario, and to see if crush is easily and reliably detectable using Mutual Information. We repeat the validation experiments described above, but this time we measure the average force and the level of MI within a simulated population of 450 individuals (again, for both idealised and representative evacuation scenarios). For each scenario, the simulation was run 64 times (across a cluster computer), and the results averaged.

We first consider the results of the force measurements, comparing them with evidence from the investigation. The force measurements for both scenarios are depicted in Figure 9. Across both scenarios the levels of force initially increase as the evacuation commences, but it rapidly decays during the idealised version of events, since evacuees are more uniformly distributed. Force levels drop to zero at around 175s, when everyone has left the building, which is broadly in line with the findings of the NIST idealised situation simulation ($195s \pm 7s$).

In the representative scenario, we observe a sharp initial rise in average force, which initially peaks after around 65 seconds. This is directly in line with the findings of the official investigation, which states that a significant crowd crush occurred by the main entrance (where around a third of the fatalities occurred) at the beginning of the time period 71-102 seconds into the fire.

“Prior to 1-1/2 minutes into the fire, a crowd-crush occurred in the front vestibule which almost entirely disrupted the flow through the main exit. Many people became stuck in the prone position in the exterior double doors [2, p. 232].

The camera angle shifts away from this door after 0:07:33 (0:01:11 fire time) and does not return to the front door until 0:08:04 (0:01:42 fire time). When the camera returns at 0:08:04 (0:01:42 fire time) a pile-up of occupants is visible. Details regarding how the pile-up occurred are not available from the WPRI-TV video; however, the interruption in flow of evacuating occupants apparent [in Figure 6-3] supports the contention that the disruption may have initiated early during the 31 second period when the camera was pointed elsewhere.” [2, p. 182]

In Figure 10, we show a screenshot of the simulation after 65 seconds. The MI measurements are depicted in Figure 11. We expect to see, as the simulations begin, an initial rise in the MI of the system. As evacuees prepare to exit the structure they tend towards *alignment*, exhibiting similar escape trajectories to other evacuees in their locale. In a maximally efficient evacuation this period of *high order* (and high MI) would be sustained throughout, as evacuees would not alter their course in order to increase their chances of effective egress. However, in an evacuation with a great deal of competition, the order in the system quickly breaks down, as the evacuees reposition themselves in order to increase their probability of escape. MI may therefore be used as an *order parameter*, where falling values of MI signify the breakdown of order within a specific evacuation. We observe marked quantitative differences in the MI readings between the two simulations. During periods of disorder, MI should tend towards zero, whereas, during ordered segments of the evacuation, MI will rise significantly.

In the idealised simulation, we see a sharp initial peak, as individuals all make for the exits at the same time. We then observe a drop, as the evacuees begin to compete for the available exit capacity. An increase in order is seen as one exit route begins to clear, creating the rise in MI at $50 < t < 75$, falling back into a state of disorder as the final evacuees clear this (main bar) exit. The MI reading then shows a progressive rise as the final evacuees exit the structure. The sharp drop in MI at the end of the simulation occurs when the number of remaining evacuees falls below some (very low) threshold.

The MI readings obtained from the simulation of *actual* events show a far more disordered evacuation, with an initial rise in MI (signifying order) quickly disintegrating into *disorder*. The MI reading at $t \approx 50s$ approaches zero; this period of highly disordered evacuation remains as the exits to the structure are overwhelmed (see Figure 10). The exit rate of evacuees during this period is extremely low, which is

confirmed by the exit profiles (see Figure 3). The MI level slowly rises towards the end of the evacuation, but, notably, the higher levels of order seen in the idealised evacuation are not reached until $t \approx 300s$, 5 minutes after the start of the evacuation.

We then perform a correlation analysis in order to establish the relationship (if any) between force and Mutual Information. A scatterplot of force versus MI suggests the existence of a statistical association (Figure 12), so we perform a simple linear correlation test. The results of this are as follows:

$$\begin{aligned} P &= 2.2e^{-16} \\ R_p &= -0.571 \end{aligned}$$

The P-value obtained is much lower than the standard significance level for a two tailed test ($\alpha = 0.01$), ($P \ll \alpha$), which confirms the significance of the result. The correlation coefficient, $R_p = -0.571$, confirms that there exists a negative correlation between MI and force within an evacuation scenario.

Discussion

Fatal levels of force can emerge within a crowd as a result of pushing, leaning or (less commonly) vertical stacking of bodies. Images of steel barriers bent out of shape (for example, in the aftermath of the Hillsborough disaster [1]) graphically illustrate the extent to which force levels can grow. Fruin reports the results of several studies (either after-the-event forensic tests, or controlled experiments) which suggest that forces exceeding around 1500N could prove fatal [18]. Crush is therefore an important factor to be considered in simulation studies aimed at improving structural designs or evacuation/control procedures, along with other aspects such as panic or physical obstacles.

Crush detection methods used to date in simulation studies may be classified into two generic groups; *explicit methods* and *implicit methods* [19]. The *implicit methodology* is the traditional approach, and is still highly popular, being the preferred technique in a large number of simulation models (see [20] for an extensive review). It relies on the *expert analysis* of factors such as population density and environmental considerations, yielding a *human interpretation* of the output of the simulation to help determine whether or not crush might have occurred. Although subjective, this method is still popular, because it does not require the use of computationally expensive force calculations, relying instead on human expertise and intuition.

The *explicit modelling* of crush conditions incorporates an assessment of crush into the model itself, and therefore requires less human analysis than the implicit approach. Usually based on the calculation of Newtonian force values, and operating in 2-dimensional space, explicit methodologies are used to detect the presence of crush conditions in a much more objective fashion. By simulating the physical force exerted by each individual, they calculate the precise amount of force present within a crowd. While the explicit methodologies offer a measure of the forces acting within a crowd, the calculations needed to assess levels of force require much more computer processing power than an implicit method. Experiments show that the computation time required by a model that *explicitly* quantifies force can be up to 100 times greater than that required by an *implicit* model [21]. We therefore require a computationally “cheap” alternative if large-scale, iterative studies are to be effective.

In this paper we have described a novel technique for detecting the onset of crush in crowd evacuation scenarios. By calculating the Mutual Information of a system of interacting individuals, we are able to determine the level of order within a crowd. We have shown that consistently low levels of Mutual Information are correlated with high levels of force within a crowd. This method allows planners to quickly and easily incorporate objective measures of crowd disorder and crush into their simulation scenarios. Future work will focus on refinements of the technique, as well as investigation of its “real-world” applicability. A key extension of the method will incorporate partitioning of the simulated space

in order to detect the *location* (as well as the existence) of crush. Another possible addition would be the consideration of social and psychological factors within our simulation. We are also particularly interested in the potential for using our technique to analyse real-time video images, with the eventual aim of developing an on-site automatic early warning system for crush and disorder at large-scale events.

Acknowledgments

References

1. Taylor L (1989) The Hillsborough Stadium Disaster. Final Report on Enquiry by Rt Hon Justice Taylor, London, UK: Home Office, HMSO .
2. Grosshandler W, Bryner N, Madrzykowski D, Kuntz K (2005) Report of the Technical Investigation of the Station Nightclub fire. Technical report, National Institute of Standards and Technology, USA. Available at <http://www.nist.gov/ncst/>.
3. Helbing D, Johansson A, Al-Abideen H (2007) Dynamics of crowd disasters: An empirical study. *Physical Review E* 75: 46109.
4. Gwynne S, Galea E, Lyster C, Glen I (2003) Analysing the evacuation procedures employed on a thames passenger boat using the maritimeEXODUS evacuation model. *Fire Technology* 39: 225–246.
5. Galea E, Blake S, Gwynne S, Lawrence P (April 2003) The use of evacuation modelling techniques in the design of very large transport aircraft and blended wing body aircraft. *The Aeronautical Journal of the Royal Aeronautical Society* : 207–218.
6. Owen M, Galea E, Lawrence P (1996) The EXODUS evacuation model applied to building evacuation scenarios. *Journal of Fire Protection Engineering* 8: 65.
7. Kukla R, Kerridge J, Willis A, Hine J (2001) PEDFLOW: Development of an Autonomous Agent Model of Pedestrian Flow. *Transportation Research Record* 1774: 11–17.
8. Kisko T, Francis R (1985) EVACNET+: A computer program to determine optimal building evacuation plans. *Fire Safety Journal* 9: 211–220.
9. Gwynne S, Galea ER, Owen M, Lawrence PJ, Filippidis L (1999) A review of the methodologies used in the computer simulation of evacuation from the built environment. *Building and Environment* 34: 741–749.
10. Wicks RT, Chapman SC, Dendy RO (2007) Mutual information as a tool for identifying phase transitions in dynamical complex systems with limited data. *Physical Review E* 75.
11. Ryder N, Sutula J, Schemel C, Hamer A, Brunt V (2004) Consequence modeling using the Fire Dynamics Simulator. *Journal of Hazardous Materials* 115: 149–154.
12. Korhonen T, Hostikka S, Heliövaara S, Ehtamo H, Matikainen K (2007) FDS+ Evac: Evacuation module for fire dynamics simulator. In: *Proceedings of the Interflam2007: 11th International Conference on Fire Science and Engineering*. pp. 1443–1448.
13. Helbing D, Molnár P (1995) Social force model for pedestrian dynamics. *Physical Review E* 51: 4282–4286.

14. Helbing D, Farkas I, Vicsek T (2000) Simulating dynamical features of escape panic. *Nature* 407: 487–490.
15. Korhonen T, Hostikka S (2009) Fire dynamics simulator with evacuation: FDS+Evac. Technical Report VTT Working paper 119. Available from <http://www.vtt.fi/proj/fdsevac>.
16. Thompson P, Marchant E (1995) Testing and application of the computer model SIMULEX. *Fire Safety Journal* 24: 149–166.
17. Gwynne S, Galea E, Lawrence P, Filippidis L (2001) Modelling occupant interaction with fire conditions using the buildingEXODUS evacuation model. *Fire Safety Journal* 36: 327–357.
18. Fruin JJ (1993) The causes and prevention of crowd disasters. In: *First International Conference on Engineering for Crowd Safety*. Elsevier.
19. Harding PJ, Amos M, Gwynne S (2010) Prediction and mitigation of crush conditions in emergency evacuations. In: Klingsch WWF, Rogsch C, Schadschneider A, Schreckenberg M, editors, *Pedestrian and Evacuation Dynamics 2008*, Springer Berlin, Heidelberg. pp. 233-246.
20. Kuligowski ED, Peacock RD (2005) A review of building evacuation models. Technical report, National Institute of Standards and Technology, USA.
21. Quinn MJ, Metoyer RA, Hunter-Zaworski K (2003) Parallel implementation of the social forces model. In: *Proceedings of the Second International Conference in Pedestrian and Evacuation Dynamics*. pp. 63–74.

Figure Legends

Figure 1. Environment to be simulated. Floorplan of Station nightclub, taken from official report.

Figure 2. Station nightclub. Rendering in FDS+Evac.

Figure 3. Initial validation results. Comparison of leaving profiles between our simulation (FDS) and official NIST findings.

Figure 4. Layout of benchmarking environment. Position A marks the centre of the entry point for pedestrians, and position B marks the centre of the exit.

Figure 5. Screenshots of benchmarking simulations. Normal scenario.

Figure 6. Benchmarking simulations. Evacuation scenario.

Figure 7. Results of benchmarking simulations. MI (green) and Average Force (red) plotted against time for normal scenario.

Figure 8. Results of benchmarking simulations. MI (green) and Average Force (red) plotted against time for evacuation scenario.

Figure 9. Average force comparison for real and idealised scenarios. Across both scenarios the levels of force initially increase as the evacuation commences, but it rapidly decays during the idealised version of events, since evacuees are more uniformly distributed.

Figure 10. Typical screenshot of our fire scenario simulation after 65 elapsed seconds. This illustrates the significant crush around the main entrance and sunroom area (high levels of force are shown in red).

Figure 11. Mutual Information comparison for idealised and representative scenarios. This illustrates the difference between ordered and disordered evacuations in terms of MI.

Figure 12. Scatterplot of force versus Mutual Information. This suggests the existence of a statistical association.

Mutual Information for the Detection of Crush Conditions - Conference

This conference publication contains much of the material covered in Chapters 7 of this thesis. It was published in Pedestrian and Evacuation Dynamics 2010.

Details

Author(s): Harding, P.J., Amos, M., and Gwynne, S.

Title: Mutual Information for the Detection of Crush

Book: Pedestrian and Evacuation Dynamics 2010

Editor(s): Peacock, R.D., Kuligowski, E.D., and Averill, J.D.

Pages: 778-784

Year: 2010

Publisher: Springer

DOI: 10.1007/978-1-4419-9725-8_73

Mutual Information for the Detection of Crush

P. J. Harding¹, M. Amos¹, and S. Gwynne²

Novel Computation Group¹
Manchester Metropolitan University
Manchester, U.K.

Hughes Associates, Inc²
London, UK

Corresponding Author: p.harding@mmu.ac.uk

Abstract: This paper describes the application of Mutual Information to the detection of crush in a well-established model of pedestrian evacuation. We show that Mutual Information offers a computationally low-cost alternative to "expensive" physical force calculations for the detection of crush in evacuation simulations.

Introduction

A number of software environments [1] exist for the simulation of large-scale egress situations, such as the evacuation of buildings, stadia and other enclosed spaces. These environments offer sophisticated tools for the analysis of human behaviour under evacuation conditions, and can recreate many of the social, environmental, structural, and psychological factors that may affect egress. Although such simulation environments can accurately model many aspects of crowd behaviour, they generally lack the capability to analyse the effects of the *physical forces* that build up within crowds. These forces can give rise to *crush conditions* (or simply *crush*), and are commonly cited as a major cause of injuries and fatalities during emergency evacuations [2-6]. The inclusion of crush analysis in simulations has traditionally been achieved by one of two methods;

1. Implicit

The implicit approach is the traditional method of qualifying the presence of crush within a simulation. It requires experienced engineers and technicians to analyse simulation output, such as population densities and environmental considerations, to ascertain the likelihood of crush becoming a danger during an evacuation. This is, however, a fundamental weakness of this technique - since it requires the knowledge of an expert analyst, the identification of crush is inherently subjective and difficult to automate.

2. Explicit

This approach requires the deployment of physical force calculations to quantify the level of crush that arises within a crowd. Most often based on traditional physical force equations, the explicit analysis of crush conditions offers a highly accurate measure of the force that exists within a simulation, but incurs a significant computational overhead.

We propose the use of Mutual Information (MI) [7] as a new approach to the analysis of crush conditions within a simulation environment. MI is a probabilistic method of analysing order within variable sets, and offers the possibility of automated qualification of the presence of crush within a simulation, whilst requiring a fraction of the computational overhead required by physical force calculations. In this paper we first define the notion of crush, before discussing previous work on the analysis of crowd movements. We then give a formal definition of Mutual Information, before describing its application to an established model of pedestrian movement. We conclude with a discussion of possible future work.

Crush

The danger presented by crush conditions has been recognised for some time as a major cause of injury and death during emergency situations [8,9]. The build-up of force within groups of people is known to be a major cause of *compressive asphyxia* (or *traumatic asphyxia*), which is the application of pressure on or about the

chest or ribcage which leads to shortness of breath and, eventually, suffocation. These types of injury are characteristic of situations in which crush conditions are present.

There have been a number of situations where crush has caused a great number of injuries or deaths. Some of the most notable include the Hillsborough disaster [2], the Gothenburg dancehall fire [3], the E2 nightclub incident [4], the Station Nightclub fire [5], and the Mihong bridge disaster [6]. The precipitating factors for the formation of crush are many and varied, e.g. emergency evacuation due to fire (Gothenburg and the Station Nightclub) or poor event management (Hillsborough, Mihong bridge). The numerous causes of crush, and the dynamically changing nature of crowd behaviour, can therefore make it difficult to precisely define the parameters under which a situation may lead to crush conditions forming.

Previous Work

After analysing video recordings of the Hajj pilgrimage in Saudi Arabia (2006), it was noted by Johansson *et al* [10] that the crowd exhibited a marked change in behaviour under certain conditions. This change in behaviour appeared to mark a *transition* between laminar ("smooth") flow of individuals, and a more *turbulent* flow. We therefore suggest that this latter state of crowd behaviour immediately precedes the formation of crush, and that its detection can therefore act as an indicator of imminent crush conditions. It has already been shown that Mutual Information may be used to identify phase transitions within a system of interacting, self-propelled particles [11]. This work focused on identifying kinetic phase transitions in the Scalar Noise Model [12] (SNM), a system of dimensionless particles that exhibits flocking behaviour under correct parameterisation. By measuring the MI of the system, it is possible to detect the point (the kinetic phase transition) at which the system moves from *chaotic* or stochastic behavioural characteristics to exhibiting signs of *order*.

We suggest that these transitions within evacuating crowds, from one state of collective behaviour to a qualitatively different behavioural state, may be considered analogous to the kinetic phase transition identified in the SNM. It is by treating the formation of crush as a phase transition (which can be identified within an evacuation) that we form the basis of applying the MI technique for crush analysis and detection.

Mutual Information

Mutual Information (MI) is a statistical measure of the mutual dependence of two variables, and has been used extensively as an analytic technique [13,14]. Equation 1 expresses the mutual information (I) of two discrete signals (A and B).

$$I(A, B) = \sum_{i,j} P(a_i, b_j) \log_n \frac{P(a_i, b_j)}{P(a_i)P(b_j)}$$

$P(a_i)$ is the probability of A having the value a_i ; $P(a_i, b_j)$ is the probability of A having the value a_i and B having the value b_j . The base of the logarithm (n) defines the units in which the MI will be measured; this is commonly base 2, giving the MI in *bits*. In general terms, MI quantifies the measure of *interdependence* between two signals or variables; therefore, if A and B are entirely independent then $I(A,B) = 0$, but in all other cases MI is non-zero.

Experimentation

The Social Forces Model [15] (SFM) is a well-established framework for the simulation of pedestrian movement. A particle-based model, the SFM can accurately recreate many of the social, psychological and physical forces present within evacuating crowds. We use the Mutual Information technique to analyse a version of the SFM identical to that presented in [15], but with two important additions. Firstly, in the original model an injured agent forms an immovable obstacle, still able to exert force (both physical and *social*) on agents within their interaction radius. This behaviour causes problems during simulations, as it makes possible the creation of

a *barricade* of injured agents between the evacuating mass and the only available exit. This causes the simulation to end with evacuees remaining in the structure. Simulations in which this occurs are declared void, the results unusable, and the experiments must be restarted. To counteract this issue, we add a rule dictating that when agents succumb to injury they are removed from the simulation after an arbitrary amount of time. This allows the increase in *force* that an injured agent may incur to be fully taken into account within the simulation, but prevents the *barricading* behaviour mentioned previously. Secondly, in order to obtain a baseline for the MI of the system (i.e. a null value), a period of "milling" is introduced. This takes the form of a 10 second "pre-evacuation" period inserted at the start of each experiment, during which agents have no clearly defined goal. This addition yields a baseline value for the MI in each simulation, i.e. the value of the MI for a random geospatial distribution of agents. Therefore all experiments show the start of the experiment, at $t = -10s$, with evacuation beginning at $t = 0s$ as per the standard model.

Tests are run using a combination of the spatial and directional data taken from the agents during multiple simulations of the Social Forces model. The MI is calculated as shown below:

$$I(X, \Theta) = \sum_{i,j} P(x_i, \theta_j) \log_2 \frac{P(x_i, \theta_j)}{P(x_i)P(\theta_j)}$$

$$I(Y, \Theta) = \sum_{i,j} P(y_i, \theta_j) \log_2 \frac{P(y_i, \theta_j)}{P(y_i)P(\theta_j)}$$

$$I = \frac{I(X, \Theta) + I(Y, \Theta)}{2}$$

With this approach the coordinate and directional data on each agent is analysed in such a way that the spatial clustering is abstracted from the analysis, i.e. the two positional variables are analysed separately. This counteracts problems associated with the measurement of the Euclidean distance between particles, in which the MI acts as a better metric of *clustering* than of the *alignment* of behavioural characteristics. This analysis relies solely on the changing behaviours of the agents (more precisely, the changing velocity vectors), rather than their spatial clustering. The results are depicted in Figure 1.

We observe that the peak in MI is pronounced, with a large increase in the MI as the agent vectors become ordered (at $t > 0s$), displaying the characteristic rise in MI that would be expected as a system attains order. The more relevant characteristic of the mutual information - the severe drop in MI that identifies the deterioration of the system into a state of disorder ($1s < t < 3s$) - is also highly pronounced.

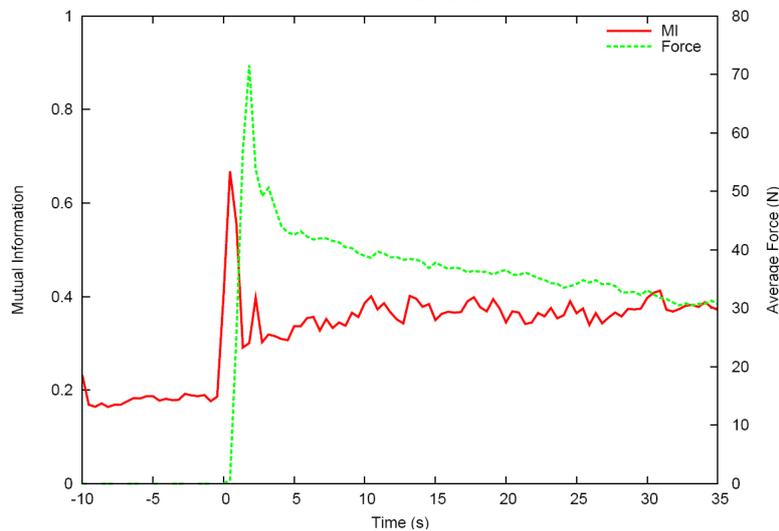


Figure 1 - MI and Average Force, as recorded against simulation time.

Conclusions

Despite this work still being in at an early stage, our preliminary results show that by employing the Mutual Information technique it is possible to detect the point at which the formation of dangerously high levels of force become likely, without the need to calculate the *precise levels of force* present within a specific simulation. As this method of crush analysis negates the need for the actual calculation of physical forces acting between agents, it is thought that it provides a far less computationally expensive method of analysis, although the exact *cost saving* has yet to be calculated.

Future Work

The investigation into the use of Mutual Information as a detector for crush is still in its infancy, but research into this area will continue, with a particular focus on the use of MI as an automated indicator of crush. Exact computational cost savings must also be calculated. These steps will make the case for the MI analysis to be regarded as a valid and effective alternative to traditional crush detection methods

Acknowledgments

The authors would like to thank the Dalton Research Institute for supporting this work, and also Hughes Associates, Inc for allowing Dr Steve Gwynne the time to take part in this research.

References

1. Kuligowski E: A comprehensive review of 28 evacuation models. <http://fire.nist.gov/>, 2004
2. Taylor LJ: The Hillsborough Stadium Disaster. Final Report on Enquiry. UK Home Office, HMSO, 1989.
3. Comaeau E and Duval R: Dance hall fire, Gothenburg, Sweden, 1998. <http://www.nfpa.org>.
4. Wilgoren J: 21 die in stampede of 1,500 at Chicago nightclub. New York Times, February 18th 2003.
5. Grosshandler W *et al*: Report of the technical investigation of the station nightclub fire. <http://fire.nist.gov/>, 2007.
6. Zhena W, Maoa L, and Yuanc Z: Analysis of trample disaster and a case study Mihong bridge fatality in China in 2004. Safety Science, 46(8):1255–1270, 2008.
7. Kraskov A, Stoegbauer H, and Grassberger P: Estimating Mutual Information. Phys. Rev. E. **24**, 2004
8. Predtechenski V and Milinski, A: Planning for Foot Traffic Flow in Buildings. , Amerino Publishing Co., New Delhi, 1978.
9. Fruin, J., Pedestrian Planning and Design, Metropolitan Association of Urban Designers and Environmental Planners, New York, 1971.
10. Johansson A, Helbing D, Al-Abideen H, and Al-Bosta S: From crowd dynamics to crowd safety: A video-based analysis. Advances in Complex Systems (ACS), 11(04):497–527, 2008.
11. Wicks R, Chapman S, and Dendy R: Mutual information as a tool for identifying phase transitions in dynamical complex systems with limited data. Physical Review E, 75(5), 2007.
12. Vicsek T, Czirok A, Ben-Jacob E, Cohen I, and Schochet O: Novel type of phase transition within a system of self-driven particles. Phys. Rev. Lett, 75:1226–1229, 1995.

13. Jeong J, Gore J, and Peterson B: Mutual information analysis of the EEG in patients with Alzheimer's disease. *Neurophysiology*, 112:827–835, 2001.
14. Battiti R: Using mutual information for selecting features in supervised neural net learning. *IEEE Transactions on Neural Networks*, 5, 1994.
15. Helbing D, Farkas I, and Vicsek T: Simulating dynamical features of escape panic. *Nature*, 407:487–490, 2000.

Prediction and Mitigation of Crush Conditions

This conference publication contains much of the material covered in Chapters 2 - 4 of this thesis. It was published in *Pedestrian and Evacuation Dynamics* 2008.

Details

Author(s): Harding, P.J., Amos, M., and Gwynne, S.

Title: Prediction and Mitigation of Crush Conditions in Emergency Evacuations

Book: *Pedestrian and Evacuation Dynamics* 2008

Editor(s): Klingsch, W. W. F, Rogsch, C., Schadschneider, A., and Schreckenberg, M.

Pages: 233-246

Year: 2010

Publisher: Springer

DOI: 10.1007/978-3-642-04504-2_18

Prediction and Mitigation of Crush Conditions in Emergency Evacuations

Peter J. Harding¹, Martyn Amos^{1,2}, and Steve Gwynne³

¹ Manchester Metropolitan University
Manchester, UK

² E-mail: m.amos@mmu.ac.uk

³ Hughes Associates, Inc
Boulder, Colorado, USA
E-mail: sgwynne@haifire.com

Summary. Several simulation environments exist for the simulation of large-scale evacuations of buildings, ships, or other enclosed spaces. These offer sophisticated tools for the study of human behaviour, the recreation of environmental factors such as fire or smoke, and the inclusion of architectural or structural features, such as elevators, pillars and exits. Although such simulation environments can provide insights into crowd behaviour, they lack the ability to examine potentially dangerous forces building up *within* a crowd. These are commonly referred to as *crush conditions*, and are a common cause of death in emergency evacuations.

In this paper, we describe a methodology for the prediction and mitigation of crush conditions. The paper is organised as follows. We first establish the need for such a model, defining the main factors that lead to crush conditions, and describing several exemplar case studies. We then examine current methods for studying crush, and describe their limitations. From this, we develop a three-stage hybrid approach, using a combination of techniques. We conclude with a brief discussion of the potential benefits of our approach.

1 Introduction

The events of 9/11 were widely seen and examined in the safety community and beyond. The catastrophic outcome and the minutiae of the evacuation have been examined by numerous official agencies, research organizations, media outlets, as well as Hollywood. Given this, the events of the day are incredibly well known; possibly more so than any other recent event.

Tall buildings are designed based on the assumption that an evacuation is managed, i.e. that the evacuation will take place in stages, if required, with only certain sections of the population evacuating at any one time. The evacuation will usually take place from those floors closest to the incident, then occur from more distant floors. This assumption is key to the successful evacuation of these tall structures; the stair capacity is calculated based on the assumption that the majority of the population follow the evacuation procedure. This means that the stair capacity within the structure will not be sufficient for the simultaneous evacuation of the entire population.

After 9/11 the assumption that tall buildings can be evacuated in a phased and controlled manner is being questioned. Instead, it is often suggested that evacuees will be reluctant to remain behind in a structure, fearful of a failure in structural integrity similar to that experienced in the twin towers. Given the nature of the incident on 9/11 and the possible consequences of remaining within the building (either by choice or through compulsion), it is now suggested that residents may choose to ignore the instructions of a staged procedure and instead move to the stairwells. This may then overload the available staircase capacity.

Given this is the case, the consequences of failure should be examined. If there is a failure in the acceptance of procedure then either the failure should be made as graceful as possible, or measures should be taken to resolve the issue; in either case, an understanding of the consequences of failure is vital.

It should be noted that during these scenarios it is not assumed that the conditions are dependent upon the existence of panic, which is difficult to predict and rarely the dominant evacuee behaviour[1]. In reality, it has been found that panic and irrational behaviour are a direct effect of the deteriorating conditions, rather than the cause of the deterioration itself. Here we are assuming that crush conditions may develop simply because of the overloading of a route and may therefore be influenced by architectural, procedural, or behavioural factors.

One of the consequences of a full evacuation from a tall structure, that was originally designed for phased evacuation, is the overloading of an escape route in a relatively short period of time. One of the most dangerous consequences of such an incident is that the exits, such as those at the base of stairways, would become overloaded, leading to many evacuees arriving at a bottleneck; i.e. the exit component is used above and beyond its design capacity. This may then lead to conditions similar to those observed at the Rhode Island[2] and Gothenburg[3] incidents, where crush incidents and falls were evident and lead to blocked egress routes and injuries. It is therefore critical for the safety of tall structures to develop an understanding of: (1) Exactly when these conditions may develop? (2) What factors need to be present in order for crush conditions to occur? (3) When do these conditions become critical? (4) How can we establish the possible consequences of this type of incident and design against them?

Here, we outline a program of work that will enable the assessment of architectural and procedural designs in order to establish whether they are prone to crush conditions developing in certain scenarios, what the consequences of this might be, and how we might best mitigate against this event. The development of a similar tool is mentioned in the recommendations within the 9/11 report[4]:

NIST recommends that tall buildings be designed to accommodate timely full building evacuation of occupants when required in building-specific or large-scale emergencies such as widespread power

outages, major earthquakes, tornadoes, hurricanes without sufficient advanced warning, fires, explosions, and terrorist attacks. Building size, population, function, and iconic status should be taken into account in designing the egress system. Stairwell capacity and stair discharge door width should be adequate to accommodate counter-flow due to emergency access by responders.

Improved egress analysis models, design methodologies, and supporting data should be developed to achieve target evacuation performance for the building population by considering the building and egress system designs and human factors such as occupant size, mobility status, stairwell tenability conditions, visibility, and congestion.

Although numerous egress models exist that are able to simulate general movement, none are able to simulate all of the conditions highlighted in NIST recommendations, along with a comprehensive crush model. Developing such a model, that is publicly available and that can be embedded into existing egress tools, meets an identified need and will allow for a broad and vital examination of these situations.

2 Definition of Crush Conditions

There are many factors that play a part in the initial formation of crush conditions during an evacuation, these can be classified under the broad headings of spatial, temporal, perceptual, procedural, and cognitive components.

2.1 Spatial

The spatial components of crush conditions are the simplest to quantify. They relate to the ratio of space available for egress to the number of persons that are expected to use the escape routes. Fruin defined this metric as the “level of service” [5], and highlighted the level at which the population density has the potential to facilitate the formation of crush as “Level of Service F”, which is the density at which a single individual would have, on average, less than $0.46m^2$ of space available to them. It should also be noted that the International Maritime Organisation (IMO) considers an evacuation to be unsafe if, for 10% of the overall evacuation time, the density of the evacuees reaches 4 persons per square metre [6]. This is due to the fact that, even at relatively low levels of force, prolonged exposure to “light” crush conditions may still cause serious injury or death.

2.2 Temporal

Temporal factors of egress vary, and depend heavily upon the rate at which conditions change. The RSET (Required Safe Egress Time), defined as the

elapsed time between the initialisation of an evacuation and the final evacuee reaching safety[7], i.e. the time required for a complete evacuation under ideal circumstances. The ASET (Available Safe Egress Time), defined as the total time *available* for evacuation[7], is a far more specific metric, as it will vary depending on the catalyst for evacuation (i.e. the nature of the emergency). Traditionally, the RSET and ASET metrics have been used to determine whether or not the occupants of a building could evacuate under specific conditions. Generally, a structure could be considered safe if the ASET value exceeds that of the RSET, i.e. there is more time available for an evacuation than would be required. The rate at which conditions change can compound time constraints, as the ASET calculation will change dynamically with the changing conditions. The Rhode Island nightclub fire (see Section 3.1), is a good example of this, and shows how the rapidity with which an incident escalates can place severe time constraints on the evacuating population.

2.3 Perceptual and Cognitive Factors

Perceptual and cognitive factors that lead to the formation of crush conditions are intrinsically linked, as an individual must rely on their perception of events in order to decide upon a course of action. The individuals' perceived level of threat plays a large part in this, as it has the most direct effect on the decision making process. Whilst the perception of threat plays a great part in the decision making process, the relationship between these two factors is highly complex, and can result in individuals displaying a wide range of behaviour, from the altruistic at one end of the scale, right through to highly competitive egress behaviour, e.g. running, pushing, etc.

The perception of information also plays a key part in the formation of crush. During emergency situations, it is often found that information relating to the current conditions is slow to propagate throughout a crowd of people, e.g. the evacuees that are placed further back in a crowd may not be aware of the conditions further ahead. This has been found in many situations, such as the Hillsborough disaster (see Section 3.4), where the people attempting to enter a structure were unaware of the already dangerously overcrowded conditions that existed inside. In these cases the persons at the rear of a crowd can compound the situation by producing additional force that will propagate forward through the crowd, and also by limiting the extent to which the pressure could be alleviated by inadvertently blocking the most immediate exit routes.

2.4 Procedural

The procedural components of crush were already alluded to (see Section 1), and centre around the inability, or unwillingness, of evacuees to follow strict evacuation plans in emergency situations. This type of problem is extremely

common in public buildings, where a great number of the occupants will be unfamiliar with the structure and have little, or no, knowledge of the evacuation plans, e.g. hospitals, town halls, museums, stadiums, etc. When an evacuation takes place under these circumstances the crowd will often attempt to leave by the most familiar route, generally the route by which they entered, even though there may be exits in closer proximity. An example of this type of behaviour can be found in the Rhode Island nightclub incident (see Section 3.1), where the majority of the crowd converged at just one point of escape, even though there were numerous other exits available.

2.5 Summary

The formation of crush conditions within crowds is a highly complex, emergent phenomena, and the causes of this cannot be explained by simply attributing it to the presence of panic within the crowd, which is widely regarded as being somewhat of a fallacy. We suggest that crush conditions can only be reliably defined as a combination of all the factors mentioned above, which culminate in the individuals' inability to fully control their direction and speed of movement, thus leading to an increase in the physical forces that they are subject to.

3 Case Studies

Here we present case studies representing situations where the formation of crush conditions have led to both serious injuries and fatalities. Each case study also represents some failure within a system (e.g. failure to limit the capacity of a structure to safe levels, failure to adhere to official guidelines or fire laws, failure to follow crowd control policies, etc). These types of failure are often observed in cases where the evacuation of a building leads to the death or injury of many people. Failures of this kind are common, and we believe that they should be expected, and be considered during the design of buildings, the creation of evacuation plans, and especially in simulated evacuation exercises.

3.1 Rhode Island Nightclub

The Station Nightclub, Rhode Island, was the scene of a tragedy when, on February 20th 2003, a fire during a rock concert caused 100 fatalities and significant injuries[2]. The fire started when the band's pyrotechnics ignited the flammable soundproofing foam that surrounded the stage, and quickly filled the club with dense, choking smoke. The fire spread from the stage, igniting a large portion of the ceiling, and within five minutes of the initial ignition those outside the club observed flames breaking through a portion of the roof.

Despite the existence of four possible exits, the majority of the crowd headed for the most familiar exit; the entrance to the club. This exit point was soon overwhelmed, and people began to trip or fall during their escape. The official time-line of the fire (compiled by NIST[2]) states that just 1 minute and 42 seconds after the start of the fire, there existed a “pile” of people, blocking the main exit and making further egress through that route impossible.

3.2 Gothenburg Dancehall

When fire broke out in a dancehall in Gothenburg, Sweden, on October 28th 1998, it claimed the lives of 63 people and injured over 180 others. The first floor venue in question was packed to over double its 150 capacity, with officials estimating that there may have been over 400 people in attendance. Eye-witness accounts of the incident state that population density prior to the start of the fire was already at dangerously high levels, with a number of the occupants observing that there were so many people present that they were unable to dance[3]. Shortly before midnight, a fire was discovered in one of the two stairways leading out of the first floor dancehall, and those near to the affected area began to evacuate. No announcement was made to the remaining occupants, and some survivors who had been at the far end of the hall when the fire was initially discovered stated that they smelled smoke but had initially believed it to be cigarette smoke and felt no need to evacuate. As the full evacuation began, the one remaining exit to the building quickly became overwhelmed, and the mass of evacuees began to trip or fall over others, further diminishing the capacity of the exit.

3.3 E2 Nightclub Incident

In Chicago’s E2 nightclub on Feb 17th 2003, the security guards’ use of pepper spray, to intervene during an altercation, became the catalyst for an evacuation that claimed the lives of 21 patrons[8]. When the security guards released the pepper spray in an enclosed space, the effects of the chemical compound on the surrounding crowd were significant. Those close to the attack began to rush toward the exit in an attempt to escape the pepper spray, which by this point was already spreading around the club. As the initial wave of evacuees made their way through the club, those who had not witnessed the incident began to fear for their safety, especially as it became obvious that some form of chemical agent was present.

Within seconds the entire crowd, consisting of over 1500 people, rushed towards the main exit. The door to the street opened inwards, whilst the door leading to the dance floor opened outwards. As people rushed from the club, the upper door flew outwards, pushing those on the upper landing down the steep flight of stairs. As more people attempted to exit, they were forced

on top of the fallen evacuees, and the bodies began to “stack up” and block the exit. It was the tremendous pressure placed upon the fallen evacuees that caused the 21 deaths during this incident. The most common cause of death was asphyxiation.

3.4 Hillsborough

The Hillsborough disaster[9] (Sheffield, UK), claimed the lives of 96 people and caused the hospitalisation of a further 300. Due to the heightened public interest in the incident (the match had been transmitted live on English television), and also because of the multiple perceived failures on the part of the authorities, the Hillsborough disaster has become one of the most thoroughly investigated crowd disasters in living memory.

The tragedy at Hillsborough stadium occurred when police stewarding the match made the decision to open an extra set of gates, intended as an exit, in order to relieve the extreme levels of congestion that were forming as the crowds tried to enter the stadium through the turnstiles at the Lepping’s Lane end of the ground. These gates did not have turnstiles, and the result was an influx of up to 5,000 fans through the narrow corridor that lead into the standing terrace. The sudden arrival of so many additional fans pushed the capacity of the central pens far above their legal maximum, and soon a dangerous crush formed at the front of the stands. Those fans still entering the stadium were unaware of this, and continued to attempt to enter the stand as the people inside were slowly crushed against the crowd barriers and fences at the front of the stands. The conditions at the front of the terrace became so bad that most of the 96 victims died from asphyxiation, or other crush related injuries, within five minutes of the game starting.

4 Previous Work in the Field

In general, each crush detection method that has been used to date can be classified into one of two generic groups; explicit methods and implicit methods. These two generic methodologies are outlined below, along with a brief discussion of their relative strengths and weaknesses.

4.1 Implicit

The implicit methodology is the original crush detection approach, and is still highly popular, being used in a large number of simulation models[10]. This methodology relies on the expert analysis of factors such as population density (see Section 2.1), behavioural analysis, and environmental considerations. The analysis of conditions within these models, therefore, is left to the engineer, who interprets the output of the simulation to determine whether crush conditions have occurred.

Implicit modelling does not take into account the possibility that evacuees may exhibit any competitive egress behaviours (e.g. pushing), as there is no accurate method for simulating these behaviours without the inclusion of force calculations. This makes it ideally suited for general evacuation simulations; i.e. timely evacuations under “ideal” conditions.

As the exact force being exerted upon individuals is never calculated, the precise physical danger that may exist in the evacuation can never be quantified. The only assertion that can be made, based on an implicit analysis, is that crush conditions *may* form during the evacuation in question. The benefit of this approach is that, as the physical force calculation are not performed, it requires far less processing power than other methods.

There are too many implementations of the implicit methodology to list here but a popular, well documented example is Simulex[11], from Crowd Dynamics Ltd.

4.2 Explicit

The explicit modelling of crush conditions incorporates an assessment of crush into the model itself, and therefore requires less user analysis than the implicit approach. Often based on the calculation of Newtonian force values, and generally operating in 2-dimensional space, explicit methodologies may be used to detect the presence of crush conditions much more precisely than would be possible with implicit modelling techniques. By simulating the exact forces being exerted by each individual, and enabling the propagation of forces throughout a crowd, the explicit methodology can be used to measure the exact amount of force that any individual is subject to. This, therefore, offers the possibility of *quantifying* the dangers that individuals may face, which is not possible using the implicit modelling techniques.

Whilst the explicit methodologies offer an accurate measure of the forces acting within a crowd, the calculations needed to measure force require much more processing power than an implicit implementation, so there exists a definite trade-off between the two techniques.

The most well-known implementation of this methodology is the Social Forces Model[12], which combines the force equations mentioned above with the modelling of the social forces acting within crowds. Although the original Social Forces Model was created as a learning tool, rather than a full-featured simulation environment, the model has recently been incorporated into the FDS+Evac Simulation environment[13].

5 Our Proposed Approach

We propose a three stage approach to this problem, consisting of separate processes for the **identification**, **qualification**, and **quantification** of crush

conditions. By employing different methods for all three stages of the analysis, we believe that the entire process may be completed at relatively low computational expense. We hope to implement these techniques as part of a suite of applications, that would offer existing egress simulations the possibility of including either full or partial crush analyses, depending on the level of accuracy required.

Two of the three techniques that we propose are still relatively novel and untested, so will require validation before they would be suitable for integration into existing environments. Each methodology will be fully tested as stand-alone applications, but a full validation will be required before the concepts are proven. At present, the team intends to attempt to integrate the applications into the open source simulation environment FDS+Evac, to enable full validation of the models, including historical data validation and peer validation[14].

5.1 Identification

In order to first identify crush conditions, we propose treating their formation as a simple phase transition, similar to those found in many social and biological systems[15]. In many of these systems a point is reached at which a change (often an abrupt change) can be observed, this change is characterised as a movement away from one general rule of system behaviour to another, different set of observable behaviours that dictate the state of the system as a whole.

In egress situations, a crowd will usually head towards the most familiar exit, often forming groups either before or during this action. The evacuees that make up these groups will have similar trajectories to their closest neighbours and will be travelling at a similar speed (i.e the flow, within each group, can be considered laminar). This would form the general rule for the ordered state of this system (see Fig 1 - A). If the evacuees are impeded in any way during their exit (e.g. they come across an obstacle in their path, or reach a congested area), they will reduce their speed and be forced to change their directions of movement, or forced to remain stationary (i.e. the flow becomes non-laminar, or turbulent). This would form the general rule for the disordered state of this system (see Fig 1 - B).

Buckingham's *II* Theorem[16] is a key theorem in dimensional analysis, and can be used to create a set of dimensionless variables that allow the analysis of an unfamiliar system, i.e. a system for which the equations governing its behaviour are either partially or wholly unknown. We will apply this theorem to the agent data within an egress model, to reduce the system to a number of dimensionless quantities, which can then be analysed to ascertain the state of the system at any one time. The advantage of this approach is that both the agent's physical variables (e.g. speed, direction, mass) and their decision making variables (e.g. perceived level of threat, tendency to-

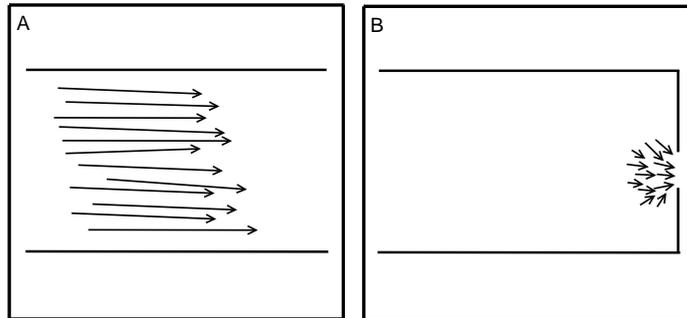


Fig. 1. Slide A shows an example of the movement vectors of evacuees during the ordered state of the system, with all vectors showing a good deal of similarity. Slide B shows example vectors during the disordered state, with the vectors varying a great deal more in both direction and magnitude

ward competition) are considered, which will provide a more comprehensive analysis of crush than could be achieved by movement variables alone.

Further analysis is achieved by the use of Mutual Information (MI)[17], a technique that has been used to quantify the similarity of two signals. This methodology was first used by Wicks *et al*[18] to detect phase transitions within a well-known flocking model[19], and was found to accurately identify the point of phase transition even when only a subset of the agents' data were analysed. We will employ a similar methodology to detect the formation of crush conditions within localised groups of agents, using the MI method to quantify the extent to which our “idealised” (ordered) agent-state (see fig 1), differs to that of the current state. We will dynamically restructure agents into groups, based on their current locale, and treat each group as a system within its own right, tracking a subset of each “sub-system” to identify the earliest stages of crush formation without the need to track *every* agent throughout the *entire* evacuation.

5.2 Qualification

To qualify the presence of crush conditions within the crowd, we intend to use a time-series, neural network classifier[20] to analyse the agent variables and movement patterns. This will give an indication of the amount of pressure that is likely being exerted on the individual in the crowd. The classifier acts as a statistical data analysis tool, and is configured to recognise the conditional similarities shared by individuals affected by the onset of crush conditions.

The neural network approach has been selected for two main reasons. Firstly, after the initial training program, the neural network approach requires far less computational power to make its classification than other statistical analysis techniques, reducing the classification during normal running

conditions to little more than matrix arithmetic. The reduction in computation, in relation to other techniques, will free up system resources for utilisation by other tasks. Secondly, the method of classification used in a neural network is highly robust, as it does not rely on any “system specific” variables, which makes the deployment of this technique possible across a wide range of existing egress simulations, without the need for extensive configuration.

By employing a time-series, neural network[21] (i.e. a neural network that accepts input in the form of sequential data representing changes over time), we also hope to identify the qualitative similarities of individuals exhibiting competitive egress behaviour. It will enable us to analyse growing behavioural trends, rather than just classify an agent’s behaviour at one precise moment in time.

To train the network, we will collect time-series agent data from a “full-force” simulation, i.e. a simulation in which a physical force model is running, which should enable the network to recognise the qualitative similarities that individuals affected by crush share. We hope that training the network using this type of data will allow the network to associate the existence of a variety of conditions to the presence of crush, therefore negating the need to engage a physics engine for all subsequent simulation runs.

5.3 Quantification

To fully quantify the effects of force propagating through a crowd, a physical force model is employed, based on the explicit crush detection method mentioned previously(see Section 4.2). We currently plan to implement this physical force model as a rigid body dynamics engine[22], with representations of such variables as mass, velocity, friction, and force propagation, modelled according to the laws of Newtonian mechanics. The engine will solve simplified physical equations in two dimensional space, resulting in good approximations[23] of force calculations that can be completed in as little time as possible.

The possibility of modelling this phenomena as a soft body dynamical system will be investigated, as recent research has highlighted the need to incorporate calculations for the compression forces acting within crowds[24], but our initial research into the feasibility of this approach leads us to believe that the calculations involved would be prohibitively computationally expensive at this time.

5.4 Hybrid Approach

The methodologies outlined above may each be employed individually, to add differing degrees of crush analysis to a simulation, but we also propose a conceptual framework, within which all three methodologies could be combined to create an analytical tool that applies crush calculations intelligently.

This approach will allow us to retain the accuracy of force calculations whilst reducing the computational expense associated with it.

The proposed approach requires the analysis of conditions based on locale, i.e. analysing conditions in different locations as if they were separate systems, and the escalation of analytical accuracy upon confirmation of crush. Figure 2 shows the flow of control across the three applications.

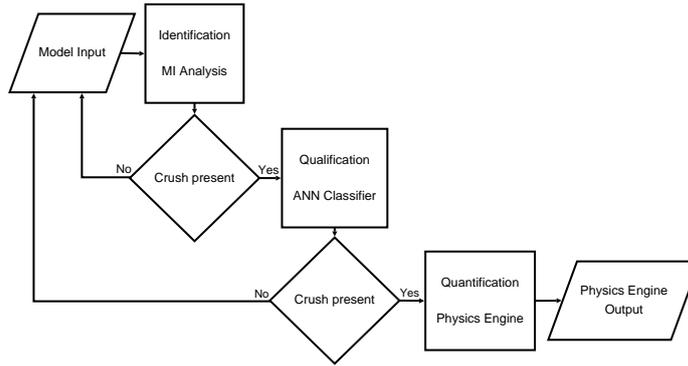


Fig. 2. Process flow diagram depicting the interactions between the three applications, according to the suggested framework.

By applying the more accurate analyses *only* once crush has been confirmed by the previous method, the most computationally expensive techniques will only be applied to affected areas, rather than across the entire behaviour space. This leaves us with the possibility of having different analyses being applied simultaneously, within the same simulation, but in different geographical locales, e.g. the identification method is running on a corridor where the flow of pedestrians is laminar, whilst at the exit of a stairwell, where a crowd has formed, the analysis would be carried out by the quantification method. The advantage of engaging each application in this way is that it will ensure that the most serious effects of crush, the build up of forces within a crowd, are measured precisely, without calculating force for all agents within the simulation.

6 Benefits of our Approach

This approach to crush analysis will provide a new tool, suitable for integration into existing simulation environments, that will allow engineers the ability to incorporate different levels of analysis for each specific simulation. The inclusion of such analytical methods will add a further dimension to traditional models, and further the realism of current simulation tools.

The addition of crush analysis techniques into models will allow engineers to better test the robustness of evacuation procedures, carry out more realistic recreations of historical incidents, and more comprehensively investigate the safety of architectural designs. It is the aim of this project to supply further tools to the evacuation sciences community that will allow this to happen, and act as a further weapon in the armoury of the engineers, technicians, and analysts that operate in this field.

7 Conclusion

The need for further crush analysis techniques has been clearly stated, and the phenomena that we wish to simulate precisely defined. We have presented three methodologies for the detection, confirmation, and measurement of crush conditions within a simulation environment, and a theoretical framework within which they could operate in unison, reducing computational expense without a reduction in accuracy.

The short-term goal of this research is simply to prove the suitability of these concepts for use in the analysis of crush, by the creation of a prototype implementation that may be used for experimentation. In the long-term we are looking to integrate this prototype into a larger simulation environment, to prove its feasibility as an “off the shelf” component to an evacuation model.

Acknowledgements

This work is partially supported by the Dalton Research Institute, Manchester Metropolitan University.

References

1. J Sime. Escape behaviour in fires and evacuations. *Design Against Fire*, 1994.
2. W Grosshandler, N Bryner, D Madrzykowski, and K Kuntz. Report of the technical investigation of the station nightclub fire. Technical report, NIST, 2005.
3. E Comeau and R F Duval. Dance hall fire gothenburg, sweden, october 28, 1998. Technical report, National Fire Protection Association, 2000.
4. Final report on the collapse of the world trade center towers. Technical report, NIST, 2005.
5. J Fruin. *Pedestrian Planning and Design*. Metropolitan Association of Urban Designers and Environmental Planners, 1971.
6. IMO. Interim guidelines for evacuation analyses for new and existing passenger ships. Technical report, International Maritime Organisation, 2002.
7. J Sime. An occupant responses escape time (oret) model. In *Proceeding of the First International Symposium*, 1998.

8. J Wilgoren. 21 die in stampede of 1,500 at chicago nightclub. *New York Times*, 18.02.2003.
9. The hillsborough stadium disaster: final report. Technical report, U.K. Home Office, 1989.
10. E Kuligowski. Review of 28 egress models. In *Workshop on Building Occupant Movement During Fire Emergencies*, 2004.
11. P A Thompson and E W Marchant. A computer model for the evacuation of large building populations. *Fire Safety Journal*, 24:131–148, 1995.
12. Dirk Helbing, Illes Farkas, and Tamas Vicsek. Simulating dynamical features of escape panic. *Nature*, 407:487–490, 2000.
13. Juha-Matti Kuusinen. Group behavior in FDS+evac evacuation simulations. Published online, August 2007.
14. J Kleijnen and R Sargent. A methodology for fitting and validating metamodels in simulation. *European Journal of Operational Research*, 120:14–29, 2000.
15. 2001 — vol. 98 — no. 17 — 9703-9706 PNAS — August 14. M beekman and d j t sumpter and f l w ratnieks. *PNAS*, 98(17):9703–9706, 2001.
16. Malcolm Longair. . *Theoretical Concepts in Physics: An alternative view of theoretical reasoning in physics*. Cambridge Univ. Press, 2 edition, 2003.
17. A Kraskov, H Stogbauer, and P Grassberger. Estimating mutual information. *Phys. Rev. E*, 69, 2004.
18. R Wicks, S Chapman, and R Dendy. Mutual information as a tool for identifying phase transitions in dynamical complex systems with limited data. *Physical Review E*, 75(5), 2007.
19. T Vicsek, E Ben-Jacob A Czirok, I Cohen, and O Shochet. Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.*, 75:1226, 1995.
20. Mohamad H. Hassoun. *Fundamentals of Artificial Neural Networks*. MIT Press, 1995.
21. A Kehagias and V Petridis. Predictive modular neural networks for time series classification. *Neural Networks*, 10(1):31–49, 1997.
22. J Wittenburg. *Dynamics of Multibody Systems: Dynamics of Systems of Rigid Bodies*. Springer-Verlag, 2007.
23. R van Zon and J Scho?eld. Numerical implementation of the exact dynamics of free rigid bodies. *J. Comput. Phys*, 225:145?–164, 2007.
24. L U Chunxia. Analysis of compressed force in crowds. *J Transpn Sys Eng & IT*, 7(2):98–103, 2007.