

Please cite the Published Version

Coates, Adam, Han, Liangxiu  and Kleerekoper, Anthony  (2018) A Unified Framework for Opinion Dynamics. In: 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), 10 July 2018 - 15 July 2018, Stockholm, Sweden.

Publisher: International Foundation for Autonomous Agents and Multiagent Systems

Version: Published Version

Downloaded from: <https://e-space.mmu.ac.uk/620221/>

Usage rights:  In Copyright

Additional Information: This is a conference paper which was originally presented at 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)

Enquiries:

If you have questions about this document, contact openresearch@mmu.ac.uk. Please include the URL of the record in e-space. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from <https://www.mmu.ac.uk/library/using-the-library/policies-and-guidelines>)

A Unified Framework for Opinion Dynamics

Adam Coates
Manchester Metropolitan University
Manchester, UK
adam.coates@stu.mmu.ac.uk

Liangxiu Han
Manchester Metropolitan University
Manchester, UK
l.han@mmu.ac.uk

Anthony Kleerekoper
Manchester Metropolitan University
Manchester, UK
a.kleerekoper@mmu.ac.uk

ABSTRACT

Opinion dynamics is the study of how large groups interact with one another and reach consensus, with applications to various areas such as computer networks, politics, and sociology. It is typically explored using agent-based modelling, with a wide variety of available models.

Numerous opinion dynamics models have been proposed, but it has been pointed out that there is a lack of a shared framework. We extend earlier attempts and provide a unified framework. The advantages of such a framework include the reduction of duplication and the identification of unexplored parameter space.

Our framework is implemented in a modular simulator which is then used to verify the validity of the framework. We show that the modular approach we propose is able to perfectly replicate results from purpose-built, stand-alone simulators for two widely used models, namely Relative Agreement and CODA.

KEYWORDS

Opinion Dynamics, Social simulation, Unifying Framework

1 INTRODUCTION

Opinion dynamics is the study of how groups of individuals adjust their beliefs and opinions as a result of interactions with one another, and exposure to additional information, media, or propaganda. It employs agent-based modelling to investigate topics such as consensus forming [11], and the role of argumentation in society [16].

The opinions held by individuals are expressed as variables - either real numbers within set bounds for continuous models, or a selection from a finite set in the case of discrete models. Relationships between individuals are modelled as edges on a graph, and rules to govern their interactions are introduced to represent real life. In each time step of the simulation, a group of agents interact with one another, and as a result shift their opinions towards or apart from one another.

The need for a unifying framework has been identified by a number of researchers. As Castellano notes, “the development of opinion dynamics so far has been uncoordinated and based on individual attempts... without a general shared framework” [1]. Xia concurs, stating that “the related endeavours are largely uncoordinated and presently it may be difficult to construct an integral and coherent framework” [24].

In the literature, we see references to a potential shared framework. Urbig et al. mention a “communications regime” and an “updating mechanism” [23]. Xie et al. display a large table of models broken down into “environmental structure”, “interaction rules”, and two components that Urbig would describe as comprising a communications regime [25]. Xia et al. similarly describes models as composed of “local rules of interaction” and “environmental structure” [24]. This pervading idea of structural, communication, and updating rules forms the cornerstone of our framework. It is possible to transplant a given update rule across various structures [2] and communication rules [3, 23], this indicates that it is the update rule that forms the central and key part of any given model. Finding patterns pervading entire sections of the literature is a strong indicator of a unifying structural framework beneath the surface of these models, as is first noted by Urbig et al. [23]. In section 3 we present our view of this framework, and demonstrate its applicability to an even wider range of models than was previously considered in Urbig et al’s work. We then implement this framework and validate it through replication of prior work in section 5.

There are multiple advantages in constructing this shared framework. Coordinating efforts between researchers minimizes the risk of duplicating work, while at the same time revealing previously unexplored parameter space. As Urbig et al. demonstrated, two models previously thought separate are in fact special cases of one unifying model [23]. This allowed investigation of a spectrum, of which those two previous models are the endpoints. Unifying additional models could reveal more spectra, and allow investigation of further parameters.

Models such as Relative Agreement [5] are described as being created by making modifications to the earlier Deffuant model [6]. This method of swapping and altering rules individually shows a modular construction rather than a fully-connected design. The ease with which structural rules can be replaced in certain models also demonstrates that said structural rules are independent of the communication or update rules [2].

However, despite the shared ideas present across a large span of time, we were unable to find a unifying framework that encompassed a large fraction of the models currently in use. Efforts such as the q-voter model by Castellano and Muñoz [3] and the random-m model by Urbig et al. [23] do make a significant amount of progress towards such a framework, and it is this work we aim to extend.

2 BACKGROUND

Opinion dynamics models can be broadly categorized into discrete models, that permit an agent to hold one of a finite set of opinions, and continuous models that allow for a real-valued number. Below, we briefly summarize the models we have paid particular attention

to, both discrete and continuous. These models were selected due to their prevalence within the literature.

Within these descriptions, we use the following symbols.

- s Opinion
- i The influenced agent
- j The influencing agent
- G A group of agents, for instance G_j for a group of influencers
- u Uncertainty

2.1 Voter Model

Perhaps the earliest discrete opinion dynamics model is the one proposed by Clifford and Sudbury to model conflict between species in an area [4]. In each time step, a single random agent i is considered along with a random neighbour j , and s_i is set to s_j . This model tends towards a stabilizing effect, though with "very rough interfaces" between areas of consensus [1]. In addition to considering single pairs, a variant of this model considers a node and its neighbourhood, N , and a binary choice of opinion [1]. Should the number of nodes in the neighbourhood holding a different opinion exceed a threshold T , the selected node changes opinion.

There are three potential outcomes: fixation, coexistence, and clustering. Fixation describes a scenario where agents converge to a static state of one opinion. Coexistence is a similar outcome, being a static state in which both opinions exist. Clustering is a state that does not stabilize, and instead has many changing clusters of nodes. Equation 1 shows how the outcomes vary with respect to T .

$$T \begin{cases} > \frac{|N|-1}{2} & \text{fixation} \\ < \frac{|N|}{4} & \text{coexisting} \\ = \frac{|N|-1}{2} & \text{clustering} \end{cases} \quad (1)$$

The authors found that these trends are always the case on 1-dimensional structures, but that the predominance towards these results still exists on lattices with more dimensions. When a community fixates, the probability of a given opinion dominating is equal to the initial density of that opinion [1].

Variants of the voter model have explored non-binary discrete opinions [19], various different initial network structures [2], the effects of "zealots" who are unchanging in their opinions [18], and explored various co-evolutionary rules [14].

2.2 Majority Rule

Similar to the voter model, this considers a group of agents with discrete opinions [13]. At each time step, a group of agents G is randomly selected, and $s_i = \text{mode}(G)$ for each agent i in G .

The authors found that for constant G , the mean consensus time, scales with $\ln N$, where N is the number of agents in the network. Furthermore, for two-state systems, the initial majority opinion will always dominate and lead to total consensus in lattices of any dimension higher than 1 [13].

2.3 Sznajd

This model is an extension of the Ising model of ferromagnetism [22]. A pair of neighbours i and j are considered and if $s_i = s_j$, the pair are in agreement and all neighbours of i or j (including non-shared neighbours) are updated with this shared opinion. This

represents the greatly magnified persuasive power of groups, compared to individuals holding an opinion. If the pair do not agree, no opinions are altered.

Original models had disagreeing pairs also influence their neighbours, but this always leads either to a fully agreeing or a chessboard-like state where every agent disagrees with all neighbours. In the updated models, a phase transition can be observed around $P(s_i = +1) = 0.5$, as the likelihood of adopting a consensus of $+1$ drastically increases [21].

2.4 Social Impact Theory

This theory considers agents with three variables - a level of persuasiveness, a level of supportiveness, and a binary opinion [20]. Persuasiveness and supportiveness respectively represent an agent's ability to change the opinion of others, and their ability to influence others to resist having their opinion changed. Edges are also given a single property, immediacy, that describes the ease or probability of communication between any two given nodes. On a lattice this is the Euclidean distance between those nodes, while in other graph topologies it can be handled with weighted edges.

If the total persuasive effect on an agent is greater than the total supportive effect, the agent is persuaded and changes its opinion. In addition an agent changing its mind, sets its persuasiveness and supportiveness to random values. This was designed as a neutral ground to observations of real life - a new convert could be impassioned with their new cause and thus more persuasive, or they could be dismissed as unstable or lacking conviction. As in models of physical forces, the effect of an interaction diminishes with the square of the immediacy.

2.5 Bounded Confidence

The bounded confidence model was developed by Deffaut et al. [6]. Agents discuss with one another and adjust their opinions only if the two agents share common ground. This is modelled by the use of a confidence threshold surrounding an agent's opinion. An agent will only adjust its opinion if the opinion of the other agent is within its confidence threshold.

This model frequently leads to groups of similar beliefs forming, with isolated extremists around the borders of groups. The final number of groups increases as the confidence of the population increases. This demonstrates that more confident or well-informed agents are more inclined to create smaller clusters around opinions they agree with, rather than drastically alter opinion. End results are typically central clustering, convergence to each extreme, or convergence to a single extreme. Increasing the threshold - equivalent to increasing uncertainty within the population - tends towards convergence to extremes.

2.6 Relative Agreement

This modification to the bounded confidence model features individual uncertainty as the threshold when determining if two agents may interact, rather than agents each having the same uncertainty [5]. Interactions are also scaled by the amount of overlap between the two agents' opinions. This is asymmetric, and a highly confident agent will not be persuaded by the efforts of a less confident agent. Furthermore, the uncertainty of agents is also modified as a

result of interactions, such that interacting with a highly confident individual that you agree with increases your own confidence in that shared belief.

2.7 CODA

The Continuous Opinions and Discrete Actions model describes a situation in which agents hold a real-valued opinion s , yet may only express themselves in discrete terms [15]. Again, at each time step a randomly-selected pair of neighbours i and j are evaluated, and their internal opinions updated. If s_j is positive, s_i is incremented by a step size α . Otherwise, s_i is decremented by this step size. As the step size is the same for all interactions, it serves only to control the speed of stabilization, not the end result. A heterogeneity factor h allows for contrarians and inflexibles to be introduced - those who change their opinion in the opposite direction to usual, or not at all, respectively.

This relatively simple update model can be applied to the various discrete models, and has the interesting effect of tending towards extremism. This is because interactions do not depend on whether or not one's partner has a greater or lesser opinion than oneself, but only if it is greater or lesser than the midpoint. Consequently, interactions between two moderate agents that agree reinforce their belief in that shared opinion without limit. Contrarians up to a concentration of 50% reinforce the central or "middle-ground" opinion, while inflexibles have no practical effect on the system.

2.8 Social Judgement Based Opinions

SJBO, or Social Judgement Based Opinions was introduced recently and has a number of features in common with the CODA model from which it draws inspiration [8, 15, 16]. It has two potential scenarios - one in which agents can express their opinion as a real number, and another where they are limited to one of a set of discrete options. Agents are expressed with three properties - an opinion from -1 to +1, an assimilation threshold, and a repulsion threshold.

In the scenario with discrete choices, another property, hesitation h , is added to each agent. Should the inner opinion of an agent lie within the range of h they will not express an opinion. Otherwise, they will express -1 or +1 accordingly. It follows that other agents are unable to discern the true opinion of an agent, only that the magnitude of their support exceeds that of their hesitation. While highly rational agents will believe that agents may not fully support their expressed opinion, to simplify the model agents are believed to fully support their expressed choice. Two global properties are also included. ρ indicates the decay threshold, below which an agent has not yet become fully committed to a cause, and $\lambda < 1$, the decay coefficient. At each time step, a random pair of agents i and j are selected. If j has expressed an opinion, i updates its own opinion. Otherwise, if $s_i < \rho$, then $s_i := \lambda s_i$. This reflects an undecided agent who loses confidence after seeing a lack of conviction among his peers.

This model displays many of the same characteristics as other continuous models, with the addition of hesitation. In a hesitating state the community displays a general preference for one or more options, but with very low consistency. This state either does not stabilize or takes an extremely large number of steps to do so, with

agents constantly alternating between expressing either no opinion, or one for which they display a slight preference.

3 UNIFYING FRAMEWORK

After reviewing the literature, we noted characteristics common to every model, and developed a framework encompassing these. We refer to the modules of this framework as rules. In this section, we list and describe each of the rule modules comprising our unified framework - structural, communication, update, and coevolutionary.

Structural Rules. Structural rules describe the initial population before the simulation begins. They encompass not only the initial distribution and configuration of attributes such as opinion, but also the edges in their network. While early models were frequently limited to complete graphs, lines, or finite lattices, recent modelling techniques allow for scale-free and small-world networks to be investigated as a more realistic model of human relationships [4, 22]. Agent attribute proportions such as contrarians and extremists are also seen as structural rules, as they only directly affect the initial composition of the population [5, 16].

Recent research uses uniformly distributed initial opinions over a complete graph or a Barabasi-Albert scale-free network. Older research - particularly discrete models - was run on lattices and lines. Several models hold properties like bounded confidence interval as global properties. We model these as a property homogeneous for every agent.

Communication Rules. Communication rules handle who interacts with who. Mathematically, they produce a subgraph. Constructing a directed subgraph of $A \rightarrow B$ means that B is influenced by A . The update rule is later applied to each node in the subgraph with incoming edges. Reciprocal edges mean two agents influencing each other, as in the Relative Agreement model.

Often, groups of agents are condensed into their means - for instance, social impact theory essentially considers an agent interacting with two others; the combined force of all those supporting it, and the combined force of those opposing. This can be modelled as a temporary agent.

Update Rules. Once agents i and j are chosen to interact by the communication rules, the update rules determine their resultant changes in opinion. Certain update rules (for instance, Relative Agreement) also alter other parameters about the agents such as their uncertainty. The update rule allows changes such as $s_i := s_j$, to set agent i 's opinion to that of agent j , or more nuanced changes such as the gradual shift exhibited in the Relative Agreement model.

The equation used within update rules is consistently some variation of $s_i := s_i + \alpha(s_j - s_i)$, where α is some scaling factor.

Coevolutionary Rules. Coevolutionary rules affect the structure of the graph itself, rather than individual opinions. These changes can include adding or removing nodes or agents, and changing, adding, and deleting edges or relationships.

The name is borrowed from biology, where changes to one thing triggers and are triggered by another. In this case, the outcome of interactions alters the structure of the graph, which in turn alters the potential interactions. Coevolutionary changes can be any of

Work	Structural				Communication		Update			Coevolutionary	
	Lattice	Complete	Random	Scale Free	Pair	Group	Majority Rule	Bounded Agreement	SJBO	Null	Edge Deletion
Castellano [2]		✓	✓	✓	✓		✓			✓	
Clifford [4]	✓				✓		✓			✓	
Deffaunt et al. [5]		✓		✓				✓		✓	
Deffuant et al. [6]		✓			✓			✓		✓	
Fan & Pedryez [8]	✓	✓			✓				✓	✓	
Fu & Wang [9]	✓				✓		✓				✓
Gil & Zanette [10]		✓			✓		✓				✓
Hegselmann & Krause [11]		✓				✓		✓		✓	
Krapivsky & Redner [13]	✓					✓	✓			✓	
Urbig et al. [23]		✓			✓	✓		✓		✓	

Table 1: The independent modules in our framework can be combined to generate opinion dynamics models. Here we present a selection of papers and their constituent modules.

random, targeted, or reactive. In the random case, elements are altered according to a random selection. In the targeted case, they are altered by some selection algorithm such as age-based removal of agents, or based on their opinions. Reactive changes occur as a result of a failure state in the communication rule. For example, if two agents are unable to reach common ground through bounded confidence mechanisms, they may sever that relationship and seek a new agent with which to interact [12].

3.1 Fitting the Framework

In this section, we describe how the models previously identified fit within our framework, and briefly summarize the mathematical transformations undertaken to link very similar models together. Table 1 displays the framework components comprising the models used in a selection of papers.

Voter Model and Majority Rule. Each of these models operates in the same manner - a point in opinion-space is determined through taking the modal average of the group of influencing agents, and the influenced agent moves to that point. The only difference between these two models lies within the number of agents that are influencing at one time. We express this with an update module that performs $s_i := \text{mode}(G_j)$ for each agent s in G_i . In the voter model, G_i and G_j each consist of a single, different, agent, whereas in the majority rule model $G_i = G_j$, and the group size is larger. Using $|G| = N$ ensures each agent interacts with all its neighbours.

Social Impact Theory. In social impact theory, an agent is considered along with its neighbours, and the persuasive and supportive influence totalled. In the equations below, o and a denote opposing and agreeing opinions, respectively.

$$I_o = |G_o|^g \frac{\sum_j^{G_o} p_j / d_j^2}{|G_o|} \quad (2)$$

$$I_c = |G_a|^g \frac{\sum_j^{G_a} \sigma_j / d_j^2}{|G_a|} \quad (3)$$

Where σ denotes supportiveness, p denotes persuasiveness, and g is a factor denoting the relative persuasive power of groups.

If $I_o > I_c$ then the agent is persuaded, and we perform the following update rule:

$$s_i := -s_i \quad (4)$$

Bounded Confidence, Hegselmann-Krause, and Relative Agreement. Together, we refer to these three similar models as “Bounded Agreement”. Bounded confidence considers N agents and selects a random agent i and one of its neighbours, j each time step. If the difference in opinions between these agents is within the threshold d , the opinions of each are adjusted. As relative agreement is an expansion to bounded confidence, the equations for both are similar. Below, the update equation for bounded confidence is shown in equation 5, and relative agreement in equations 6 and 7.

$$s_i := \begin{cases} s_i + \mu(s_j - s_i) & |s_i - s_j| < u \\ s_i & \text{otherwise} \end{cases} \quad (5)$$

$$h_{ij} = \min\{s_i + u_i, s_j + u_j\} - \max\{s_i - u_i, s_j - u_j\} \quad (6)$$

$$s_i := s_i + \mu \left(\frac{h_{ij}}{u_i} - 1 \right) (s_j - s_i) \quad (7)$$

These can be unified with:

$$s_i := \begin{cases} s_i + \mu\alpha(s_j - s_i) & v \\ s_i & \text{otherwise} \end{cases} \quad (8)$$

where α is some scaling factor and v is a function returning whether the interaction can legally occur. In order to recreate the bounded confidence model, let $\alpha = 1$ and $v(s_i, s_j) = |s_i - s_j| < u$. To instead reproduce the relative agreement model, let $\alpha = \left(\frac{h_{ij}}{u_i} - 1 \right)$ and $v(s_i, s_j) = \text{true}$.

Creating the Hegselmann-Krause model from this is done by using a group communication module and taking the weighted average of that group. This results in the random- m model described by Urbig et al. [23].

CODA. The CODA model is explicitly acknowledged as an update rule - Martins describes his work as combining the CODA model with the Sznajd model [15]. By decomposing the resultant work into distinct communication and update mechanisms, we are left with rule modules for Sznajd and CODA.




Figure 1: Program Control Flow
Control flows from left to right before entering the main loop.
Models are assembled by selecting one module from each column.

SjBO. At each time step, two agents i and j are selected using a pairwise communication rule, and their opinions updated according to the following update rule.

$$s_i := \begin{cases} s_i + a_{i,j}(s_j - s_i) & |s_i - s_j| < \epsilon \\ s_i - r_{i,j}(s_j - s_i)^{\frac{1-|s_i|}{2}} & |s_i - s_j| > \tau \\ s_i & \text{otherwise} \end{cases} \quad (9)$$

Where ϵ is the agent's assimilation threshold, τ is their repulsion threshold, a is the assimilation coefficient, and r is the repulsion coefficient.

4 IMPLEMENTING THE FRAMEWORK

In this section, we describe our implementation of the framework in a simulator. Using this simulator, we attempt to replicate prior work which we discuss in the next Section. Should our simulator produce the same results as earlier work using a different approach, we thereby demonstrate the validity of the framework.

The simulator is a Python program for constructing and evaluating opinion dynamics models according to the unified framework discussed in section 3. It is designed to be easily extensible, system-independent, and simple to use.

4.1 Provided Modules

Several modules are provided for the user, allowing them to begin experimentation without needing to write any code. These modules are all fully independent of each other, for easy exploration of new combinations of modules. The main modules are briefly described here.

Figure 1 shows the control flow of the simulator and how the rule modules map to the rule components of the unified framework.

Graph Generators. The complete graph generator produces a graph in which every agent is connected to every other agent. This is the environment used in models where every agent is able to interact with every other agent, such as random pair interactions.

Erdős-Rényi graphs are a type of graph in which each potential edge has an equal chance to exist. This generator takes a probability and returns a graph with the requested number of nodes, and every possible edge having that probability of existing. This is also frequently referred to as simply a "random" graph.

The 2D Lattice generator produces a square graph of n agents. It is a convenient shorthand for the more general lattice graph.

Barabási-Albert graphs construct a scale-free network using a preferential attachment model, in which edges are added from new nodes to nodes that already have many edges. Scale-free networks obey a power law in their degree distributions - a few agents have orders of magnitude more relationships than others.

The small-world generator produces a graph in which the number of hops between any agent grows proportionally to the logarithm of the number of agents.

Initial Values. The Uniform Distribution generator returns a number uniformly selected from the interval between -1 and $+1$, inclusively.

When provided with a mean and standard deviation, the Normal Distribution generator returns a value selected from that normal distribution, capped to within -1 and $+1$. This generator defaults to $\sigma = 1$ and $\mu = 0$, the standard normal distribution.

Group Selectors. The Random- n group selection module uses the selection process of Urbig et al. [23]. Given an agent, it returns n neighbours of that agent. In addition, the Pairwise module is provided as a convenient shorthand for Random- n with $n = 2$. If




Figure 2: A heatmap of the y metric from the Relative Agreement model using the corrected, purpose-built simulator from Meadows and Cliff[7].




Figure 4: Using a simulator with two minor discrepancies from the original model, produces a significantly different heatmap of the final metric.

n is greater than the number of neighbours, all neighbours are returned.

Updaters. The two main update modules are capable of emulating many of the existing models. In particular, the q-Voter model within the framework emulates discrete models using a variant of the voter model discussed in section 2.1, and the Unified Continuous model (see section 3.1) produces the continuous models using a variant of bounded confidence, discussed in section 2.5.

For convenience and ease of replication of earlier work, modules are provided for those models covered in section 2. For instance, the relative agreement model (see section 2.6) provides a function for calculating the overlap of two agents' opinions to the Unified Continuous module, and then returns the result of that module.




Figure 3: Using our framework to implement the Relative Agreement model by linking independent modules reproduces the heatmap almost exactly.

y	Convergence Type
0.0	Central Convergence
0.5	Dual Convergence
1.0	Single Extreme Convergence

Table 2: Convergence Types Indicated By y

Co-Evolver. The Random Rewire co-evolutionary module takes a pair of agents, severs the edge between them, and then establishes a new edge to a randomly chosen node in the graph.

5 VALIDATING THE FRAMEWORK

Validation against existing work offers a reliable and simple way to ensure that our framework is correct and our implementation is accurate. We reason that should our decomposition of these models into modules fitting our framework and subsequent reassembly produce the same results as the original model, then our framework produces a faithful recreation of that model. This demonstrates that our assertion is correct, that this model is composed of fully separable modules, and that each of those modules lie within the framework.

5.1 Relative Agreement

We selected two pieces of work from Deffaut et al. to compare against [5, 7]. In this second piece of work, the authors contest findings by Meadows and Cliff [17] that differed from their own work. Meadows and Cliff had published different findings using the same parameters as Deffaut et al., and used these findings to propose alternative conclusions to the question of how extremism spreads through networks.

In the initial study by Meadows and Cliff [17], a minor programming discrepancy that partially arose from the model not being fully specified, led to vastly different results than were expected, and they were unable to replicate the findings of Deffaut et al. Firstly,




Figure 5: The distribution of final opinions in terms of the number of steps away from 0, after 800 runs or CODA using Martins' purpose-built simulator.

the simulation was not carried out for enough time steps, and so metrics were calculated before the model had fully converged. Secondly, the threshold to be considered extreme was intended to be lower at the end of the simulation than at the beginning¹, leading to a far smaller number of extremists being reported under certain circumstances. Thankfully, Deffaunt et al. were available to consult, and after resolving these issues both groups arrived at very similar results.

This offers a unique opportunity to compare our work with multiple authors performing the same simulation, each using different purpose-built, one-off programs. These experiments made use of the relative agreement model (see section 2.6) in random pair interactions i.e. a complete graph. In each simulation, the proportion of extremists p_e and the global uncertainty of the non-extremist agents U were varied, and the resultant trends plotted in a heatmap (see² Figure 2). To analyse this trend, Deffaunt introduced a metric, $y = p_+^2 + p_-^2$. These represent, respectively, the number of initially moderate agents who became positive extremists, and those who became negative extremists. The forms of convergence indicated by different levels of y are summarized in table 2.

Our primary validation method is through visual comparison of a graph produced by our study with the graphs produced by the two aforementioned studies, aiming to identify common features in each set of results. Visual inspection of the graphs reveals a number of artefacts:

- (1) A white section on the far left, of central or no convergence.
- (2) A large red area in the bottom right, of extreme single convergence.
- (3) An orange “wedge”, from the top right towards the centre, of extreme dual convergence.
- (4) An unusually low-valued line separating the red and orange areas.

¹Though this was not directly stated in the description of the model.

²Best viewed in colour




Figure 6: Reproducing CODA using our framework produces the same distribution of final opinions.

This test showcases the importance of small factors in this form of simulation. As two relatively minor changes were able to introduce drastically different results to those seen in the original work and the corrected work by Meadows et al., we see that the model is highly susceptible to minor changes (compare the erroneous heatmap in Figure 4 to the correct one in Figure 2). Thus, our simulator producing a graph very similar to the valid one (see figure 3) demonstrates that the framework is valid and that it has been implemented correctly.

5.2 CODA

For a further test, we have analysed the results of our simulator after swapping two modules. The structural rule is changed from complete to lattice, and the update rule is changed from Relative Agreement to CODA. We then use the simulator to replicate the results of Martins [15], seen in figures 5 and 6. 2500 agents in a square 2D lattice are seeded with initial opinions between -1 and $+1$, inclusive, and a step size of 0.2 , and then left to interact 800 times each according to the CODA update rule (see section 2.7).

We measure the distance from 0 of each agent's opinion after $k = 800$ iterations per agent, in terms of multiples of the step size. The histogram displays a penta-modal distribution with primary peaks at $x = \pm k$, secondary peaks at $x = \pm k/2$, and a tertiary peak at 0. The strong similarity demonstrates that this model also fits within the framework.

6 CONCLUSION AND FURTHER WORK

Opinion dynamics is a large and complex field, with a large selection of competing models and no clear unifying framework underlying their development. In this work, we have briefly summarized a number of opinion dynamics models in use within the field, each of which seems to approach from a different angle. These varying angles led to a chaotic and disorganized body of research, potentially leaving large amounts of parameter space unexplored.

We have identified and discussed a common framework underlying all of these models, by expanding upon earlier work by Urbig

et al. unifying two models [23]. This framework uses independent rule modules to control different stages of the simulation, from determining the initial layout of the network, to handling the evolution of opinions and the network itself through coevolutionary changes. The previously discussed models were decomposed into modules fitting within this framework.

The framework was then implemented in order to test its validity. The strong similarity between the output generated by our framework-based simulator and two other purpose-built, one-off simulators is a strong indicator that the framework is a valid system for describing and creating models. The use of this framework could pinpoint possible areas for future exploration, and also highlight common assumptions made across a large variety of models.

In future work, we will add additional rules to encompass a wider range of models, and attempt to further unify existing models. We also aim to use the framework to identify new and unexplored parameter space and investigate those areas.

REFERENCES

- [1] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. 2009. Statistical physics of social dynamics. *Rev. Mod. Phys.* 81, 2 (2009), 591–646. DOI: <http://dx.doi.org/10.1103/RevModPhys.81.591>
- [2] Claudio Castellano, Vittorio Loreto, Alain Barrat, Federico Cecconi, and Domenico Parisi. 2005. Comparison of voter and Glauber ordering dynamics on networks. *Phys. Rev. E* 71, 6 (2005), 066107. <https://arxiv.org/pdf/cond-mat/0501599.pdf>
- [3] Claudio Castellano and Miguel A Muñoz. 2009. The non-linear q-voter model. *Phys. Rev. E* 80, 4 (2009), 041129.
- [4] Peter Clifford and Aidan Sudbury. 1973. A model for spatial conflict. *Biometrika* 60, 3 (1973), 581–588. DOI: <http://dx.doi.org/10.1093/biomet/60.3.581>
- [5] Frédéric; Weisbuch Gérard; Faure Thierry Deffuant, Guillaume; Amblard. 2002. How can extremism prevail? A study based on the relative agreement interaction model. *Journal of Artificial Societies and Social Simulation* 5, 4 (2002).
- [6] Guillaume Deffuant, David Neau, Frederic Amblard, and Gérard Weisbuch. 2000. Mixing beliefs among interacting agents. *Advances in Complex Systems* 03, 01n04 (1 2000), 87–98. DOI: <http://dx.doi.org/10.1142/S0219525900000078>
- [7] Guillaume Deffuant, Gérard Weisbuch, Frederic Amblard, and Thierry Faure. 2013. The Results of Meadows and Cliff Are Wrong Because They Compute Indicator y Before Model Convergence. *Journal of Artificial Societies and Social Simulation* 16, 1 (2013), 11. DOI: <http://dx.doi.org/10.18564/jasss.2211>
- [8] Kangqi Fan and Witold Pedrycz. 2015. Emergence and spread of extremist opinions. *Physica A: Statistical Mechanics and its Applications* 436 (2015), 87–97. DOI: <http://dx.doi.org/10.1016/j.physa.2015.05.056>
- [9] Feng Fu and Long Wang. 2008. Coevolutionary dynamics of opinions and networks : From diversity to uniformity. *Phys. Rev. E* 78, 1 (2008), 016104. DOI: <http://dx.doi.org/10.1103/PhysRevE.78.016104>
- [10] Santiago Gil and Damián H Zanette. 2006. Coevolution of agents and networks: Opinion spreading and community disconnection. (2006).
- [11] Rainer Hegselmann and Ulrich Krause. 2002. Opinion Dynamics and Bounded Confidence Models, Analysis, and Simulation. *Journal of Artificial Societies and Social Simulation* 5, 3 (2002). http://www.math.fsu.edu/~dgalvis/journalclub/papers/02_05_2017.pdf
- [12] Balazs Kozma and Alain Barrat. 2008. Consensus formation on adaptive networks. *Physical Review E* 77, 1 (1 2008), 016102. DOI: <http://dx.doi.org/10.1103/PhysRevE.77.016102>
- [13] P. L. Krapivsky and S. Redner. 2003. Dynamics of Majority Rule in Two-State Interacting Spin Systems. *Physical Review Letters* 90, 23 (6 2003), 238701. DOI: <http://dx.doi.org/10.1103/PhysRevLett.90.238701>
- [14] Nishant Malik and Peter J. Mucha. 2013. Role of social environment and social clustering in spread of opinions in coevolving networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 23, 4 (2013), 043123. DOI: <http://dx.doi.org/10.1063/1.4833995>
- [15] Andre C. R. Martins. 2007. Continuous Opinions and Discrete Actions in Opinion Dynamics Problems. *International Journal of Modern Physics C* 19, 4 (11 2007), 617–624. DOI: <http://dx.doi.org/10.1142/S0129183108012339>
- [16] André C. R. Martins and Cleber D. Kuba. 2009. The Importance of Disagreeing: Contrarians and Extremism in the CODA model. *Advances in Complex Systems* 13, 5 (1 2009), 621–634. DOI: <http://dx.doi.org/10.1142/S0219525910002773>
- [17] Michael Meadows and Dave Cliff. 2012. Reexamining the Relative Agreement Model of Opinion Dynamics. *Journal of Artificial Societies and Social Simulation* 15, 4 (2012), 4. DOI: <http://dx.doi.org/10.18564/jasss.2083>
- [18] Mauro Mobilia. 2003. Does a single zealot affect an infinite group of voters ? *Phys. Rev. Lett.* 91, 2 (2003), 028701. <https://arxiv.org/pdf/cond-mat/0304670.pdf>
- [19] F. Slanina and H. Lavicka. 2003. Analytical results for the Sznajd model of opinion formation. *The European Physical Journal B - Condensed Matter* 35, 2 (9 2003), 279–288. DOI: <http://dx.doi.org/10.1140/epjb/e2003-00278-0>
- [20] Jacek Szamrej, Robert Abelson, Michael Berbaum, Richard Moreland, Serge Moscovici, Brian Mullen, David Myers, Thomas Ostrom, Richard Palmer, Tom Wallsten, and Harry Upshaw. 1990. Andrzej From Private Attitude to Public Opinion: A Dynamic Theory of Social Impact. *Psychological Review* 97, 3 (1990), 362–376. <http://cognitrn.psych.indiana.edu/rgoldsto/complex/nowak90.pdf>
- [21] Katarzyna Sznajd-Weron. 2005. Sznajd model and its applications. (3 2005). <http://arxiv.org/abs/physics/0503239>
- [22] Katarzyna Sznajd-Weron and Józef Sznajd. 2001. Opinion evolution in closed community. *International Journal of Modern Physics C* 11, 6 (2001), 1157–1165.
- [23] Diemo Urbig, Jan Lorenz, and Herzberg. 2008. Opinion dynamics: The effect of the number of peers met at once. *Journal of Artificial Societies and Social Simulation* 11, 4 (2008).
- [24] H Xia, H Wang, and Z Xuan. 2011. Opinion dynamics: A multidisciplinary review and perspective on future research. *Journal of Knowledge and Systems Science* 2, 4 (2011), 72–91. <http://www.igi-global.com/article/international-journal-knowledge-systems-science/61135>
- [25] Ziping Xie, Xiao Song, and Qiyuan Li. 2017. A Review of Opinion Dynamics. In *Theory, Methodology, Tools and Applications for Modeling and Simulation of Complex Systems. AsiaSim 2016, SCS AutumnSim 2016. Communications in Computer and Information Science*, vol 646, L Zhang, X Song, and Y Wu (Eds.). Springer, Singapore, 349–357. DOI: http://dx.doi.org/10.1007/978-981-10-2672-0_36