



**Manchester
Metropolitan
University**

Rainer, A, Williams, A ORCID logoORCID: <https://orcid.org/0000-0002-6888-0521>, Garousi, V and Felderer, M (2020) Retrieving and mining professional experience of software practice from grey literature: An exploratory review. IET Digital Library, 14 (6). pp. 665-676. ISSN 1462-5970

Downloaded from: <https://e-space.mmu.ac.uk/627118/>

Version: Accepted Version

Publisher: Institution of Engineering and Technology

DOI: <https://doi.org/10.1049/iet-sen.2020.0109>

Please cite the published version

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

Retrieving and mining professional experience of software practice from grey literature: an exploratory review

Austen Rainer^a, Ashley Williams^b, Vahid Garousi^a, Michael Felderer^c

^a*School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast, Northern Ireland, U.K.*

^b*Department of Computing and Mathematics, Manchester Metropolitan University, Manchester, U.K.*

^c*Department of Computer Science, University of Innsbruck, Innsbruck, Austria; and Department of Software Engineering, Blekinge Institute of Technology, Karlskrona, Sweden*

Abstract

Background: Retrieving and mining practitioners' self-reports of their professional experience of software practice could provide valuable evidence for research. We are, however, unaware of any existing reviews of research conducted in this area. **Objective:** To review and classify previous research, and to identify insights into the challenges research confronts when retrieving and mining practitioners' self-reports of their experience of software practice. **Method:** We conduct an exploratory review to identify and classify 42 articles. We analyse a selection of those articles for insights on challenges to mining professional experience. **Results:** We identify only one directly relevant article. Even then this article concerns the software professional's emotional experiences rather than the professional's reporting of behaviour and events occurring during software practice. We discuss challenges concerning: the prevalence of professional experience; definitions, models and theories; the sparseness of data; units of discourse analysis; annotator agreement; evaluation of the performance of algorithms; and the lack of replications. **Conclusion:** No directly relevant prior research appears to have been conducted in this area. We discuss the value of reporting negative results in secondary studies. There are a range of research opportunities but also considerable challenges. We formulate a set of guiding questions for further research in this area.

Keywords: Empirical software engineering, Data mining, Natural language processing

1. Introduction

1.1. Motivation

Software practitioners form beliefs about software practice directly from their professional experience of that practice and indirectly from the information shared by other professionals about their respective experiences [68, 18].

Social media provides a platform for the *social programmer* [75] to share their experiences of software practice online. Retrieving and mining practitioners' experience from social media therefore provides the opportunity for research to automatically and non-invasively collect evidence from practitioners about software practice. This evidence can complement evidence collected through the more traditional methods of data collection, such as interviews and surveys. More generally, retrieving and mining professional experience would contribute to the aims of Evidence Based Software Engineering (EBSE), i.e., to integrate best evidence from research with *practical experience* and human values [21].

Retrieving practitioners' professional experience from online, often communicated through *high-quality* grey literature, is challenging [70]. There is a very large, and exponentially increasing, volume of online articles; there is considerable 'noise' in the data; and experience is not articulated in text in any standard way that can be effectively searched by keyword-based search engines. And once retrieved, documents then need to be mined, e.g., to identify and extract the segments of text that represented the experience.

Soldani *et al.* [72] and Garousi and Mäntylä [27] both provide examples of researchers using the professional experience of the author/s of online grey literature as an inclusion/rejection criterion in their systematic grey literature reviews. In both studies, the authors manually retrieved and analysed online articles. Rainer [67] demonstrates a method for extracting arguments, experience and explanation from blog articles. As with Soldani *et al.*, and Garousi and Mäntylä, Rainer's method is currently applied manually.

Given the preceding discussion, we are interested in the (semi) automatic retrieval and mining of self-reports of professional experience of software practice. Ideally, we want to select articles based on the author's self-reporting of professional experience of software practice, and exclude articles

that report an individual’s non-professional experience or an individual’s non-software professional experience. We are unaware of any existing literature reviews of research into the (semi) automatic retrieval and mining of self-reports of professional experience of software practice. We therefore conduct and report what we believe to be the first review in this area. We position our review as an *exploratory review* because, as we will show, there is very little – *if any* – prior work to form the basis for a Systematic Mapping Study or a Systematic Literature Review.

1.2. Aims, objectives and contributions

The overall aim of this paper is to review empirical studies that investigate how to retrieve and mine from online articles the self-reporting of professional experience of software practice. Our aim is translated into two objectives: 1) to classify the articles in our dataset so as to better understand prior work in this area; and 2) to identify challenges and issues with the retrieval and mining of professional experience of software practice.

The paper makes four contributions:

1. We report a *negative* result for previous studies i.e., our review identifies only *one* directly relevant article and this single article concerns the software professional’s emotional experiences rather than the software professional’s reporting of behaviour and events occurring during software practice. The implication of our negative result is that – to the best of our knowledge – no prior research (primary or secondary) has been conducted in this area. A study that reports a negative result may raise concerns about the rigour of the enquiry (e.g., was the review conducted properly?) and about the value of the enquiry (e.g., why conduct a review when there is little to review?). We are not aware of any prior review of the kind we report here, and it is only with the conduct of this review that we discover there is, or was, little prior work to review. Phrased another way, there was no prior indication that there would be no prior research. An alternative strategy might be (have been) to relax the search and selection criteria so as to include more studies for review. But by relaxing our search strategy we change the objectives of the review and, by so doing, we fail to report a negative result and potentially contribute to publication bias. We discuss this issue in more detail in Section 7.

2. A classification of studies related to the retrieval and mining of experience *in general*. Given our first contribution, this second contribution must necessarily be more general.
3. The identification of several challenges and issues that researchers will need to confront in the retrieval and mining of professional experience of software practice.
4. A simple set of guiding questions for primary studies, to support future researchers in this area.

1.3. Structure of the paper

The remainder of the paper is structured as follows. Section 2 explains the scope of our work, presenting and discussing illustrative examples of the problem being tackled, and discussing working definitions. Section 3 discusses prior work. Section 4 explains the method we use to conduct our exploratory review. Section 5 presents and discusses our classification of reviewed articles. Section 6 then analyses several articles in more depth. Finally, Section 7 reviews the objectives of the paper, discusses the value of negative results in secondary studies, offers guiding questions for conducting primary studies in this area, identifies threats to the validity of this study and opportunities for further research, and offers a brief conclusion.

2. Professional experience of software practice

In this section, we present and discuss several contrasting examples of the kinds of experience we are seeking (or not seeking) to retrieve and then mine. We then provide working definitions.

Table 1 and Table 2 present examples to illustrate the problem we are investigating. All of the examples in the two tables are taken from a single blog article [73], *Language Wars*, written by Joel Spolsky. We select these examples because they have been studied more formally in our prior work [67]. The exploratory review we report in the current paper is focused on identifying *primary studies* that have: a) sought to retrieve and mine blog articles for the kinds of example presented in Table 1, whilst b) avoiding the kinds of example presented in Table 2. In other words, our exploratory review is not directly investigating the challenges arising from the examples in the two tables, but instead investigating whether and how *previous work* has investigated these kinds of example, and the challenges that have arisen during those investigations.

The examples in both tables comprise fragments of verbatim text from Spolsky’s blog article.

For the first example in Table 1, the example comprises two fragments of text that are located at different places in the source blog article. For transparency, each text fragment is uniquely identified, e.g., E1.1 and E1.2. For all examples in both tables, URLs are indicated with underlined text together with additional comments in squared parentheses.

Example E1 describes Spolsky’s experience of a group of interns working in his company. As we discuss later in this paper, story is a common approach to articulating experience. Spolsky’s description in example E1 may be understood as a story: there are actors (e.g., the interns), a time and location (e.g., last summer, at the company), decisions (e.g., what language to use), actions (e.g., build Copilot using C# and ASP.NET), outcomes (e.g., we didn’t get into trouble later) and benefits (e.g., not one minute was spent on a fruitless debate). There is even the suggestion of a hero: an intern with “enough experience” to do a “beautiful job architecting the ASP.NET code exactly the right way to begin with so we didn’t get into trouble later”. Examples of this kind might be relevant to software engineering researchers interested in aspects of the software process (or educators interested in internships).

Example E2 describes the technology that Spolsky’s company used or has developed, and how that technology has changed over time. This example might be relevant to researchers interested in the prevalence of technology–use by practitioners, or in reasons for companies to develop their own technologies rather than use existing technologies, or in how technology–use changes over time.

Examples E3 and E4 describe situations where Spolsky recognises external events and outcomes (e.g., events relating to *37 Signals*) and provides his beliefs about those events and outcomes.

Example E5 does not directly describe actual behaviours and experience but instead provides an indication of Spolsky’s accumulated experience. This example is useful to researchers interested in establishing the credibility of a practitioner, e.g., as an expert.

Examples E6 and E7 are both ‘negative examples’ in the sense that they both present arguments about which technologies one should use, but do not describe experience. Example E6 refers to evidence but does not directly present any evidence. Instead, the example implicitly relies on the *reader’s* experience of software practice. Example E7 refers to knowledge, which again

implies expertise and experience of using the technologies.

Example E8 presents another ‘negative example’ concerning an action taken by Spolsky, i.e., Spolsky wrote an article. While the article itself may be about software practice, the *action* Spolsky describes is not itself about software practice.

Notice that in several examples the presence of URL links would, in principle, allow the researcher to corroborate Spolsky’s claims using independent information. As one example, the first URL link in example E1.1 directs to a website (i.e., <https://www.projectaardvark.com/>) that describes the experience of the interns. As a second example, example E8 provides a link to an article that might contain a description of experience.

Precisely defining the relevant concepts for this review is not straightforward, e.g., because the field is still emerging. Consequently, we present and use working definitions.

We recognise the distinction between an experience itself and information shared to convey that experience. Our focus here is on information shared to convey that experience, as we cannot gain direct access to the experience itself. The examples in Table 1 all present information to convey experience. For simplicity in writing, we refer to that information–sharing as the reporting of professional experience. The examples in Table 1 suggest at least the following types of *reporting of professional experience*:

Type I: A person’s description of their own behaviour (e.g., decisions, actions) as they interact with other individuals within a professional situation, e.g., a nurse’s description of her or his behaviour within a particular hospital situation. Example E1 and E2 both contain elements of this type of reporting: E1 refers to “what we ended up doing”, and E2 refers to “what we use”.

Type II: A person’s description of their direct observations of others’ behaviour. Examples E1, E3 and E4 present elements of this type: E1 refers to the behaviour of interns, E3 refers to events relating to the company *37 Signals*, and E4 refers to “Paul’s” behaviour.

Type III: A description of a person’s accumulated experience. Example E5 presents elements of this type. This accumulated experience might then be used to assess the credibility of the practitioner, and from that the credibility of the practitioner’s claims and arguments, such as those presented in examples E6 and E7.

The reporting of the professional experience of a *software practitioner* would describe her or his behaviour in relation to the domain of software engineering: designing, testing, deploying and maintaining software-based systems. It follows that there are other *domains* of experience. Whilst we recognise that there are different domains of non-professional and professional experience, we consider it premature to *a priori* specify or enumerate those domains. We identify some of the domains as part of our review.

We distinguish in principle the reporting of a professional *experience* of actual behaviour and events (e.g., a software practitioner describes in writing her experience working on a software project) from the articulation of *beliefs* based on or inferred from that experience (e.g., that practitioner's opinions about some aspect of that project). Examples E1 and E6 provide illustrative examples of this distinction. One reason we make this distinction is to help separate beliefs based on one's direct experience from beliefs formed from information provided by others. In other words, we seek to confine a witness's statement to a description of what occurred, and distinguish the witness statement from expert testimony. In practice, it may be very challenging to confine such witness statements and to distinguish them from expert testimony.

Just as there are challenges isolating description from opinion, so there will likely be challenges separating factual information from other kinds of information, such as evaluative and emotional information, e.g., examples E3, where Spolsky writes of being "scared". These are some of the challenges that the successful retrieval and mining of professional experience will need to tackle.

Finally, whilst there is no standard way for reporting professional experience, we observed in example E1 — and discuss later in the paper — that *story* appears to be a common approach used by people to report their own and others' behaviour. Again, difficulties arise between the use of a factual or partially factual story to describe an experience from the past, in contrast to an abstract story describing a desired future state (e.g., *User stories*), or in contrast to fictional stories (e.g., [66]) used for entertainment or to convey moral value or other meanings. A further challenge is that, for reasons of confidentiality, commercial sensitivity or diplomacy etc., a contrived story may be used to illustrate a real situation.

3. Related work

3.1. *The value of practitioner experience*

Evidence Based Software Engineering (EBSE) seeks to integrate best evidence from research with practical experience and human values [21]. Devanbu *et al.* [18] and Rainer *et al.* [68], amongst others, have observed that software practitioners value their own experience, and that of their peers, over empirical evidence, placing most value on local expertise.

Storey *et al.* [75] discuss the emergence of the *social programmer*, as social media has dramatically changed the way that software practitioners share information. There are large amounts of online articles, as people share information (though not necessarily about professional issues, or about their experience of practice). In terms of software practice, Choi [17] maintains a reference list of 650 blogs, and Soldani *et al.* [72] observe a “. . . massive proliferation of grey literature [on microservices], with more than 10,000 articles on disparate sub-topics. . .”.

Overall, we conjecture that, in due course, experience mining might help research to achieve EBSE, i.e., by integrating professional experience of software practice with best evidence from software engineering research with human values.

3.2. *Grey literature in software engineering research*

There is emerging interest in incorporating grey literature into software engineering research as such literature provides insight into industry practice. Garousi *et al.* [26] advocate the use of Multi-Vocal Literature Reviews (MLRs) to synthesise peer-reviewed academic literature with grey literature. Such syntheses can help to address the disconnect between the state-of-art and the state-of-practice. (Garousi *et al.* have published guidelines for conducting MLRs [25].) There is the need to automate aspects of the review process to handle the large volume of articles being considered. Rainer [67] presents a methodology for analysing one type of grey literature, the practitioner-written blog article, for reports of the software practitioner’s experience and for the beliefs that practitioners infer from those experiences. Rainer recognised the need to automate the methodology so that it could be scaled to the large volume of online articles produced by software practitioners.

A comparable situation is found in the health sector: Adams *et al.* found that many evaluations of public health interventions “. . . may be predomi-

nantly, or only, held in grey literature and grey information...” [3]. For Adams *et al.* there are at least three reasons why researchers may want to include grey literature in their studies. (We rephrase them here in terms of software engineering research.) First, including grey literature can reduce publication bias. Studies with null findings are less likely to be published in peer-reviewed software engineering journals and conferences, so grey literature can help to compensate for unpublished null findings. Second, grey literature can provide useful contextual information on how, why, and for whom software engineering practices and technologies are effective: context is recognised as a significant challenge in software engineering research. Third, synthesising grey literature can help software engineering researchers and practitioners understand what solutions exist for a particular software engineering problem, the full range of evaluations (if any) that have been conducted, and where further development and evaluation is needed.

Overall, we conjecture that, in due course, experience mining might help research automate MLRs and grey literature reviews (GLRs).

3.3. Mining opinions, arguments and experiences

Cabrio and Villata [11] present a five-year review (c2014 — c2018) of argumentation mining, complementing the previous reviews of Lippi and Torroni [46] and Peldszus and Stede [61]. (Mochales and Moens [54] provide an earlier review restricted to argumentation mining in relation to law.)

Cabrio and Villata [11] distinguish *opinion mining*, which they consider focuses on understanding *what* users think about a topic, from *argument mining*, which focuses on *why* users have a certain opinion. To these we may add *experience mining*, which focuses on a person’s experiences, distinguished from their opinion and argument.

Cabrio and Villata [11] distinguish the identification of an argument or arguments within a text, from predicting the relations between the extracted arguments. They recognise that argumentation mining requires high-quality annotated corpora, that there is the additional challenge in argumentation mining of establishing the factuality of information in an argument; and that there is the significant challenge of aligning the heterogeneity of datasets and approaches used by researchers, partly explained by the immaturity of the field: each dataset has been annotated on the basis of different definitions of argument components and relations, thus preventing the possibility of a straightforward alignment of these datasets.

Lippi and Torroni [46] structure their review around their argumentation pipeline. A version of the pipeline (based on Lippi and Torroni’s [47]) is presented in Figure 1, modified to include our continuing research on the search and selection of grey literature [69]. Lippi and Torroni’s pipeline uses the sentence as the basic unit of discourse. As we show later in this paper, previous research on experience mining has used a variety of units of discourse, for which the sentence is one.



Figure 1: Argumentation pipeline, modified from [47]

Lippi and Torroni [47] use an IBM–developed annotated corpus¹ to evaluate their MARGOT (Mining ARGuments frOm Text) software system, the first publicly accessible web server for argumentation mining. The corpus was developed by IBM as part of its Debater project. The annotation guidelines for the IBM corpus distinguish three types of *evidence* for arguments: *study* (e.g. results of quantitative analyses), *expert* (e.g. testimony by someone/s with known expertise on, or authority in, some topic) and *anecdotal* (a description of specific event instances or concrete examples). It is the third type of evidence that is the particular focus of the current paper (though the second type is relevant too).

Other research on argumentation defines experience in terms of detecting events and stories. Bex [7] defines stories as “...a coherent sequence of events, often involving subjects, objects, outcomes and other attributes”, and Twining [78] defines a story as “...a narrative of particular events arranged in a time sequence and forming a meaningful totality.”

Overall, we conjecture that experience mining might effectively complement the existing sub–fields of text mining.

3.4. Tool support for systematic reviews

Software engineering research has investigated how to reduce the time and effort for systematic reviews. For example, Marshall *et al.* [49, 50] have reviewed tools to support systematic reviews (SR) in software engineering.

¹https://www.research.ibm.com/haifa/dept/vst/mlta_data.shtml

These tools support a number of different steps in the SR process. None of the tools appear to consider the mining of professional experience.

Unsurprisingly, evidence based medicine has considered some of the issues we raise in this paper. Tsafnat *et al.* [77] report a survey of tools for automating parts of (evidence based medicine’s) systemic reviews, structured in terms of the steps of a review (such as search and extraction). Their survey does not appear to address the mining of professional experience. O’Mara-Eves *et al.* [57] report the systematic review of text mining for study identification. They observe an active and diverse body of research, but with almost no replication between studies or collaboration between research teams. They conclude that text mining is ready to be used in ‘live’ systematic reviews (in medicine) for *prioritising* the order in which items are reviewed; and may be used cautiously as a ‘second screener’ to complement the human reviewer. Their focus appears to be on *screening* articles.

Overall, we conjecture that experience mining might inform the development of appropriate software tools.

4. Method

In the following subsections, we first describe our process for searching and selecting articles, then explain our rationale for performing two sets of analyses, and then describe how each set of analyses was conducted.

4.1. Searching & selecting articles

The second author performed 18 queries, using five search engines, to identify the initial set of candidate articles. These searches were conducted in July 2018. We followed the recommendations of Singh and Singh [71] to tailor our search strings to the specific requirements of each search engine. Table 3 summarises the searches. Table 4 summarises the steps involved in the subsequent filtering of articles.

For each search query, the first 50 results were considered for inclusion in the review. (Given that, subsequently, only 15% of the search results were classified, and less than 1% of the search results are directly relevant to our topic, selecting the first 50 search results was an effective, pragmatic inclusion criteria for this study.) Each of those 50 results was selected as a candidate if that result could *not* be immediately rejected based on the respective article’s title, e.g., an article was removed where its title was referring to mineral mining. Results that were written in a non-English language, or

that were patents, were *retained* at this stage, because they could not be immediately rejected based on title. The initial search process yielded 309 candidate articles.

The second author then filtered the 309 articles following the steps summarised in Table 4. Following the filtering, 42 articles remain for classification.

4.2. Rationale for our two sets of analyses

We perform two sets of analyses of the 42 articles: first, a classification of the articles; second, a comparison of a selection of articles. Our classification finds only *one* article that is directly concerned with the retrieval and/or subsequent mining of the reporting of professional experience of software practice. Furthermore, this singular article concerns the software professional's *emotional experiences* rather than the software professional's reporting of behaviour and events occurring during software practice. Our second set of analysis therefore compares a *selection* of articles for their insights into challenges confronted by, and issues recognised by, the researchers in their respective research. We expect these challenges will recur as and when researchers investigate the retrieval and mining of professional experience of software practice.

4.3. Classifying the full set of articles

The first and second authors initially reviewed all 42 articles to better understand the full set of articles, e.g., reviewing a subset of articles and then returning to re-consider the full set of articles. (We actually reviewed a slightly larger set, 48 articles, but excluded 6 of these for reasons indicated in Table 4.) This initial review produced a rudimentary classification scheme and a preliminary classification of the articles.

The first author then developed a simple classification framework to be used to re-classify the full set of 42 articles. All 42 articles were independently classified by the first and the second authors. As a further check, and to enable all authors to contribute to the full paper, the third and fourth authors each independently classified a quarter of the articles, and a colleague also independently classified a quarter, resulting in the majority of articles being independently reviewed. Differences in classifications were resolved by the first author.

4.4. Analysing a subset of articles

During the development of the classification for our first set of analyses, the first and second authors developed a deeper appreciation of the full set of articles. We selected eight articles for further comparison. These articles were selected because they offered insights into the challenges confronted by, and the issues recognised by, the researchers involved in the primary studies. We do not claim that we provide a complete coverage of all challenges. Of these eight articles, seven were selected from the list of 42 articles. We identified an eighth article [14] from our fuller set of articles. We included that eighth article because the article provided a useful contrasting datapoint. We emphasise therefore that the second analyses does not analyse a subset of the 42 articles but instead analyses a selection of articles to identify and illustrate challenges. The second analyses does not claim to be exhaustive in the challenges identified.

5. Classification of the articles

Table 5 presents our classification of the 42 papers. The table suggests that the first papers explicitly investigating experience mining appear to be Kurashima *et al.* [41] in 2006, followed by Inui *et al.* [35] in 2008. Table 5 includes two papers [9, 45] that precede Kurashima *et al.* [41] and Inui *et al.* [35]; a third paper [64] was published in the same year. These three papers are concerned with story detection and event detection rather than directly with experience mining.

Table 5 shows that there is only one article [22] that is directly concerned with the retrieval and subsequent mining of the reporting of professional experience of software practice. Even then, the article focuses more on the developers' emotional responses to situations rather than the developers recounting of their experiences.

Table 5 also suggests that there are few articles on the mining of professional experience in general: there are many more articles that investigate the reporting of non-professional (personal) experience compared to professional experience. The two articles by Gonçalves *et al.* [5, 4] are the only articles in our dataset to examine professional experience. They use a tool to *generate* stories from the professional participants in contrast to, for example, mining reports of professional experience from practitioners' grey literature. Gonçalves *et al.* also intend to use the generated stories to define business rules and processes, rather than as evidence for research.

There are a number of articles that detect experience in terms of story although, by contrast, the one article [22] concerned with detecting the reporting of professional experience of software practice does *not* detect experience in terms of story. The table identifies those articles concerned with story and organises them into those that detect professional experience and those that detect non-professional experience. Though not shown in the table, there are a number of different types of story: story as public news (e.g., a news story), story as a stream of events, folk stories, telling stories. Table 8 presents examples of story based and non-story based experience.

On average, about three papers are published per year in this area, however 75% of the papers have been published in the most recent eight years (the second half of the sixteen year period in which the set of articles have been published). This suggests a slight growth in the area.

Figure 2 visualises the *multi-dimensional* ‘experience space’ characterised by Table 5. The figure illustrates the challenges of conducting retrieval and mining that *include* online articles about the professional experience of software practice whilst *excluding* articles about a number of other closely-related issues.

Our overall conclusion is that whilst there has been more than 15 years’ interest in mining experience, it appears that there has been limited work published in this area. More specifically, there appears to be a very limited number of studies – and arguably no previous study – of the retrieval and mining of self-reports of professional experience of software practice. This suggests a nascent field of research that has yet to establish itself and that remains indistinguished from already-established fields such as opinion mining [34], event detection [16] and argumentation mining [46].

In the following section we consider some challenges likely to arise during the retrieval and mining of self-reports of professional experience of software practice.

6. Analysis of challenges

In this section we discuss several challenges and issues that arise with the retrieval and mining of online articles that contain reports of experience. The eight papers considered in this section are summarised in Table 6.




Figure 2: An illustration of the ‘experience space’

6.1. Characteristics of experience-reporting articles

Jijkoun *et al.* [37] found that online articles reporting experience are different to those that do not report experience: experience-articles have twice as many sentences, twice as many words (these two statistics are clearly related), twice as many references to the pronoun *I*, and twice as many non-modal verbs. There is, therefore, an indication that experience-reporting articles are different to non-experience articles, however we are unaware of any other study that has reported such characteristics and, as indicated by our visualisation in Figure 2, we further distinguish experience-reporting articles from professional-experiencing-reporting articles.

6.2. The prevalence of experience in online articles

There is very limited information reported in the articles on the prevalence of experience in online articles. Swanson *et al.* [76] cite their previous work [28] that found about 5% of blog posts are about personal stories. Ceran *et al.* [14] found 3,301 story-paragraphs in a dataset of 16,930 (19%; though

these are paragraphs rather than documents). On the basis of this limited information, it appears that there would likely be a low prevalence of reporting professional experience in online articles. This degree of prevalence increases the challenges of generating appropriate datasets, of classifier training and evaluation, and then of effective application of trained classifiers e.g. to automate aspects of grey literature reviews (GLRs) and multivocal literature reviews (MLRs).

6.3. Definitions and models of experience, and theories

Table 6 suggests that a limited number of papers base their work on prior theory, though most papers have some kind of model of experience to inform the selection of features. The most frequent model of professional experience appears to be in terms of some kind of story, though there are several different types of story being used by research e.g. news stories, folk stories. Swanson *et al.* [76] asserts that storytelling is considered a fundamental aspect of human social behaviour. Swanson *et al.* [76] also considered that the narrative structures of personal narratives posted in blog posts are likely to be more similar to oral narratives than to classical stories, and consequently chose Labov and Waletzky’s [43] theory of oral narrative to distinguish causal relationships. By contrast, Goncalves *et al.* [4] refer to story mining and story telling, but use a metamodel of scenarios [2] to model experience rather than a model of stories. Table 8 contrasts the conceptual elements of the different models used in our sample, and Table 9 summarises some of the challenges of ‘mapping’ the conceptual elements to linguistic elements of a text. Aside from datasets, the keywords to use for retrieving online articles, and then the features to use (e.g., in machine classifiers) to mine or classify the articles, there is also the conceptual challenge of modelling the reporting of experience in online articles.

6.4. The sparseness of data

Kurashima *et al.* [40] acquired a large dataset of 48M blog posts, but found that 38% of the dataset was not suitable because it did not contain a sufficient co-occurrence of words they needed for their association rules. The authors observed a sparseness in their data e.g. the average frequency of their *activity* words per blog post was 4.8; for *location* words the average frequency was 0.85, and for emotion words the average frequency was 0.26.

6.5. Unit/s of discourse analysis

Table 6 suggests that the majority of studies use the sentence as the basic unit of discourse. This is consistent with the findings of Lippi and Torroni [46] in their review of argumentation mining. Reporting of professional experience will likely fall across sentences and paragraphs (see for example the analysis in Rainer [67]) so there will likely be the challenge of appropriate text segmentation and aggregation.

6.6. Annotating, and annotated datasets

Table 7 shows that only three of our studies report measures of annotation agreement. Swanson *et al.* state, "... the annotation task is highly subjective, requiring interpreting the narrative and the author's intention, which prevents us from obtaining high levels of inter-rater agreement." ([76]; p. 175). They also observe earlier in their paper that a previous study [65] found both a high level of annotator agreement and an extremely high machine learning accuracy for Aesop's Fables. Swanson *et al.* [76] infer that the 'classical', written-down stories are therefore (much) easier to work with than blog-posts. This observation is consistent with the progress made by argumentation mining i.e. to first begin by mining texts that have a consistent structure and content e.g. legal texts. Swanson *et al.* [76] also recognise that a more sophisticated analytical framework, and annotation scheme, such as that of DramaBank has advantages, but developing an annotated corpus using such an analytical framework for blog posts would be (highly) resource intensive.

Kurashima *et al.* [40] focused on two types of association rules. For the first type, they identified 1659 candidate rules, for which their annotators confirmed 129 correct rules (8%); and for the second type, they identified 1330 candidates rules, for which their annotators confirmed 55 correct rules (4%). Again, this suggests a considerable challenge with annotating the data and then mapping the annotated data to the classified data.

6.7. Evaluating the performance of algorithms

Different articles report different approaches to retrieving and mining their data, and this makes it difficult to compare the performance of different approaches. We contrast three examples here.

Swanson *et al.* ([76] investigate performance both in terms of (linguistic) clauses and also stories. They observe that, "... some clauses are ambiguous and difficult [for the annotators] to label, but also that some *stories* are more

difficult to classify.” ([76], p. 177); and, “Performance is different for results by story rather than over all clauses. This indicates that some stories are more difficult to classify than others and that ambiguous clauses are not uniformly distributed but are likely to be correlated with particular authors or writing styles.” ([76], p. 178). Swanson *et al.* [76] is the only article to present a confusion matrix; also they discussed sources of error (which we do not consider here, due to space constraints).

By contrast, Goncalves *et al.* [4] evaluate the overall tool–process, with one case example, rather than the method by which linguistic features are mapped to process attributes. As another contrast, Kurashima *et al.* [40] sought to identify the more ‘interesting’ association rules, and were also seeking an appropriate measure for ‘interestingness’.

6.8. Replications

In our dataset of articles, we found no replication, independent or otherwise, of previous work and no attempts to independently validate the findings, or assumptions, of others’ work.

7. Discussion

7.1. A review of the paper’s objectives

Recall from section 1 that the two objectives for the paper were: 1) to classify the articles in our dataset so as to better understand prior work in this area; and 2) to identify challenges and issues with the retrieval and mining of professional experience of software practice.

For our first objective, we identified and classified 42 articles. In so doing, we showed that there were very few previous studies investigating the retrieval and mining of professional experience of software practice.

For our second objective, we found that it will be very challenging to retrieve and mine *only* those online articles that are about the professional experience of software practice whilst *excluding* online articles about a number of other closely–related issues. Challenges include:

1. There is a low prevalence of experience (not necessarily professional experience) reported in online articles, the implication being that a very large dataset would be needed to, for example, train and test automated tools.

2. There is no consistently used definitions, models or theories of experience. A consistent definition etc. would help with, as examples, replication and the accumulation of evidence. Conversely, contrasting definitions etc. are important for exploration. Overall, the lack of consistent use of definitions etc. suggests the field is still maturing. In our sample, story is the most common approach to articulate experience however story is defined in different ways by different researchers (see Table 9 for examples).
3. Datasets are sparse of features, i.e., had very low frequencies of the presence of features.
4. There is variation in the units of discourse analysis used, with the sentence being the most common unit. It is also not clear how a text would be segmented for experience.
5. Previous work tends not to report annotator agreement and, where annotator agreement is reported, the authors recognise considerable difficulties achieving agreement.
6. There is considerable variation in the use of approaches and algorithms, and difficulty comparing their performance.
7. There has been no replication of previous research.

Overall, our findings are consistent with those of O’Mara-Eves *et al.* [57] and Cabrio and Villata [11] in their respective research domains.

7.2. Guiding questions for primary studies

Given the lack of prior research in this area, and the challenges of even identifying and then comparing studies from prior research, we propose a simple set of guiding questions for the design, conduct and reporting of the retrieval and mining of professional experience. Space prevents a detailed discussion of these questions. We present the questions in Table 10, with brief cross-reference to relevant subsections earlier in paper.

7.3. Reporting negative results in secondary studies

One recurring item of feedback on earlier versions of this paper may be summarised with the following question: What is the contribution of a negative result in a secondary study?

Primary and secondary studies both contribute knowledge to a field or fields of research. A secondary study is often also an empirical study in its own right, albeit that the *unit of data* for the secondary study (e.g.,

information from primary studies) is different to that for the primary study (e.g., information direct from practitioners or observations of the phenomena of interest itself). But both types of study follow methodology to analyse data and to interpret the results of that analysis. And both types of study may report false–negative and false–positive results.

Because primary studies reporting positive results are much more frequently published, secondary studies are often unavoidably ‘loaded’ toward analysing primary studies that report positive results. We suggest this recurring situation leads to a normalisation of the expected results of a secondary study and to a subtle assumption as to what secondary studies *should* report, i.e., that a defining characteristic of a secondary study is that it reports ‘positive’ results about a field of research (and, by implication, that a field of research must have positive results for reporting in a secondary study).

The first and second authors of the current paper were interested in experience mining because of how it relates to research they were conducting on the credibility of grey literature. Our initial motivation, therefore, was not to advance the field of experience mining (either through conducting a primary study or a secondary study) but rather to use existing knowledge from the field. With several related fields of research already relatively well–developed (e.g., opinion mining, event detection and argumentation mining) we hoped to find an existing secondary study on experience mining, and we assumed we would find at least some primary studies on experience mining. (In other words, *our* expectations were for positive results.) When we didn’t find a secondary study, we decided to conduct one. In conducting our secondary study we found a lack of primary studies.

One counter–argument is that we might have quickly detected the lack of primary studies early, or earlier, in the conduct of our secondary study, e.g., some ‘quick’ Google Scholar searches would have shown there was little prior work. As we report in section 4, our initial search returned 300 results and, as we progressed in the study, our preliminary selection and rejection process (see Table 3) continued to suggest we would find at least some relevant primary studies. A risk with making a decision based on the quick search for results is that we terminate the study prematurely, i.e., we make a decision based on false–negative results. One of the principles of systematic reviews is that the reviewer performs a thorough and balanced search for relevant literature. One potential outcome of a balanced secondary study is to report negative results.

As noted in the introduction, an alternative strategy might be (have

been) to relax the search and selection criteria so as to include more studies for review. For example, we might have expanded our secondary study to look at mining in general, drawing on argument mining, opinion mining etc. Whilst relaxing the search strategy would have widened the number of primary studies, the relaxed strategy would also have diverted us from the focus of our research, i.e., to understand how *experience mining* might help us with assessing the credibility of grey literature. Also, relaxing our search strategy would change the objectives of the review and, by so doing, would fail to report a true negative and potentially contribute to publication bias.

7.4. Threats to validity

There are of course threats to the validity of our review. The conduct of the searches and the initial selection of articles was performed by one author, the second author. There may be bias in that search and selection process. We might have applied a more extensive set of search terms and search strategies (including snowballing) to ensure we found all relevant literature to review. Restricting our review to the first 50 results may have limited the articles we found, though we found that a very small percentage of articles were relevant in the first 50 results and we did conduct a series of searches using a number of different academic search engines. There are also threats relating to the application of the inclusion and exclusion criteria, as well as other aspects of the selection process. The classification framework was prepared by one author, the first author. This framework was applied independently by the other authors, with all articles being classified independently by at least two authors, and most articles being classified independently by three.

7.5. Further research

There are several directions in which to extend this research. First, an unresolved and fundamental challenge for mining professional experience for GLRs and MLRs in software practice is establishing the appropriate features to use when classifying experience (cf. [20]). We observe in our review a lack of consensus on appropriate features and, of course, features may be domain specific. Second, there is also the challenge of generating appropriate datasets. Garousi et al. [27] and Soldani et al.[72] provide two datasets as starting points, however both are very small (each approximately 50 articles) for training classifiers. In addition to the size of the dataset, there is the related issue of the sparseness of data. A third challenge is the effective

reporting of the full details of a study, so as to support reproducibility and replication. A fourth direction is to apply lessons learned and technology from the retrieval and mining of other domains to the retrieval and mining of software practice. Finally, there is the need to evaluate the effective contribution that experience mining could make in due course to secondary studies, such as GLR (e.g. [72]) and MLRs (e.g. [24]), and to primary studies (e.g. [60]). Recall that O’Mara-Eves *et al.* [57] found *text mining* to be ready to use for *screening papers* in their domain.

7.6. Conclusions

The overall aim of this paper was to review articles that investigate how to retrieve and mine from online articles the self-reporting of professional experience of software practice. We set two objectives: 1) to classify the articles in our dataset; and 2) to identify challenges.

Our review identified only *one* directly relevant article and this single article concerns the software professional’s emotional experiences rather than the software professional’s reporting of behaviour and events occurring during software practice. It appears that no prior directly-relevant research has been conducted in this area: experience mining is a nascent field of research. We complement our near-null result with a classification of the studies, and with the identification of several challenges. We also propose a simple set of guiding questions to support future researchers. We discuss the value of negative results in secondary studies.

Acknowledgments

We thank Dionysis Athanasopoulos for his independent classification of a subset of the articles presented in Table 5.

References

- [1] Shuya Abe, Kentaro Inui, Kazuo Hara, Hiraku Morita, Chitose Sao, Megumi Eguchi, Asuka Sumita, Koji Murakami, and Suguru Matsuyoshi. “Mining personal experiences and opinions from Web documents”. In: *Web Intelligence and Agent Systems: An International Journal* 9.2 (2011), pp. 109–121.
- [2] Camille Ben Achour. “Guiding scenario authoring”. In: *Information Modelling and Knowledge Bases X* 51 (1999), p. 152.

- [3] Jean Adams, Frances C Hillier-Brown, Helen J Moore, Amelia A Lake, Vera Araujo-Soares, Martin White, and Carolyn Summerbell. “Searching and synthesising ‘grey literature’ and ‘grey information’ in public health: critical reflections on three case studies”. In: *Systematic reviews* 5.1 (2016), p. 164.
- [4] João Carlos de AR Gonçalves, Flávia Maria Santoro, and Fernanda Araujo Baião. “A case study on designing business processes based on collaborative and mining approaches”. In: *Computer Supported Cooperative Work in Design (CSCWD), 2010 14th International Conference on*. IEEE. 2010, pp. 611–616.
- [5] João Carlos de AR Gonçalves, Flávia Maria Santoro, and Fernanda Araujo Baião. “Collaborative narratives for business rule elicitation”. In: *2011 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE. 2011, pp. 1926–1931.
- [6] M. Behrooz, R. Swanson, and A. Jhala. “Remember that time? Telling interesting stories from past interactions”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 9445 (2015). cited By 1, pp. 93–104. DOI: 10.1007/978-3-319-27036-4_9. URL: https://www.scopus.com/inward/record.uri?eid=2-s2.0-84951871183&doi=10.1007%5C%2f978-3-319-27036-4_9&partnerID=40&md5=ce1b5a85586a8a12ee9d78efdd93fa48.
- [7] FJ Bex. “Values as the point of a story”. In: *From knowledge representation to argumentation in AI, law and policy making. A Festschrift in honour of Trevor Bench-Capon* (2013), pp. 63–78.
- [8] Francesco Bonchi, Ilaria Bordino, Francesco Gullo, and Giovanni Stilo. “Identifying buzzing stories via anomalous temporal subgraph discovery”. In: *Web Intelligence (WI), 2016 IEEE/WIC/ACM International Conference on*. IEEE. 2016, pp. 161–168.
- [9] Thorsten Brants, Francine Chen, and Ayman Farahat. “A system for new event detection”. In: *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM. 2003, pp. 330–337.
- [10] Kevin Burton, Akshay Java, Ian Soboroff, et al. “The ICWSM 2009 Spinn3r dataset”. In: *Third Annual Conference on Weblogs and Social Media (ICWSM 2009)*. 2009.

- [11] Elena Cabrio and Serena Villata. “Five Years of Argument Mining: a Data-driven Analysis.” In: *IJCAI*. 2018, pp. 5427–5433.
- [12] Ricardo A Calix, Ravish Gupta, Matrika Gupta, and Keyuan Jiang. “Deep gramulator: Improving precision in the classification of personal health-experience tweets with deep learning”. In: *2017 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE. 2017, pp. 1154–1159.
- [13] B. Ceran, N. Kedia, S. R. Corman, and H. Davulcu. “Story detection using generalized concepts and relations”. In: *2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. Aug. 2015, pp. 942–949. DOI: 10.1145/2808797.2809312.
- [14] Betul Ceran, Ravi Karad, Steven Corman, and Hasan Davulcu. “A hybrid model and memory based story classifier”. In: *Proceedings of the 3rd Workshop on Computational Models of Narrative*. 2012, pp. 58–62.
- [15] Betul Ceran, Ravi Karad, Ajay Mandvekar, Steven R. Corman, and Hasan Davulcu. “A Semantic Triplet Based Story Classifier”. In: *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012)*. Washington, DC, USA: IEEE Computer Society, 2012, pp. 573–580. ISBN: 978-0-7695-4799-2. DOI: 10.1109/ASONAM.2012.97. URL: <http://dx.doi.org/10.1109/ASONAM.2012.97>.
- [16] Xieling Chen, Shan Wang, Yong Tang, and Tianyong Hao. “A bibliometric analysis of event detection in social media”. In: *Online Information Review* (2019).
- [17] Kilim Choi. *Software Engineering Blogs*. URL: <https://github.com/kilimchoi/engineering-blogs#-individuals>. (accessed: 04.08.2018).
- [18] Prem Devanbu, Thomas Zimmermann, and Christian Bird. “Belief & evidence in empirical software engineering”. In: *Proceedings of the 38th International Conference on Software Engineering*. ACM. 2016, pp. 108–119.
- [19] Cristiano Di Crescenzo, Giulia Gavazzi, Giacomo Legnaro, Elena Troccoli, Ilaria Bordino, and Francesco Gullo. “HERMEVENT: A News Collection for Emerging-event Detection”. In: *Proceedings of the 7th International Conference on Web Intelligence, Mining and Semantics*.

- WIMS '17. Amantea, Italy: ACM, 2017, 11:1–11:10. ISBN: 978-1-4503-5225-3. DOI: 10.1145/3102254.3102262. URL: <http://doi.acm.org/10.1145/3102254.3102262>.
- [20] Pedro Domingos. “A few useful things to know about machine learning”. In: *Communications of the ACM* 55.10 (2012), pp. 78–87.
 - [21] Tore Dybå, Barbara A Kitchenham, and Magne Jørgensen. “Evidence-based software engineering for practitioners”. In: *IEEE Software* 22.1 (2005), pp. 58–65.
 - [22] Awdren Fontão, Oswald M Ekwoje, Rodrigo Santos, and Arilo Claudio Dias-Neto. “Facing up the primary emotions in Mobile Software Ecosystems from Developer Experience”. In: *Proceedings of the 2nd Workshop on Social, Human, and Economic Aspects of Software*. ACM. 2017, pp. 5–11.
 - [23] Elena García-Barriocanal, Miguel-Angel Sicilia, and Nikolaos Korfiatis. “Exploring hotel service quality experience indicators in user-generated content: A case using Tripadvisor data”. In: *Proceedings of the 5th Mediterranean Conference on Information Systems (MCIS 2010)*. 2010.
 - [24] Vahid Garousi, Michael Felderer, and Mika V Mäntylä. “Guidelines for including grey literature and conducting multivocal literature reviews in software engineering”. In: *Information and Software Technology* (2018).
 - [25] Vahid Garousi, Michael Felderer, and Mika V Mäntylä. “Guidelines for including the grey literature and conducting multivocal literature reviews in software engineering”. In: *arXiv preprint arXiv:1707.02553* (2017).
 - [26] Vahid Garousi, Michael Felderer, and Mika V Mäntylä. “The need for multivocal literature reviews in software engineering: complementing systematic literature reviews with grey literature”. In: *Proceedings of the 20th International Conference on Evaluation and Assessment in Software Engineering*. ACM. 2016, p. 26.
 - [27] Vahid Garousi and Mika V Mäntylä. “When and what to automate in software testing? A multi-vocal literature review”. In: *Information and Software Technology* 76 (2016), pp. 92–117.

- [28] Andrew Gordon and Reid Swanson. “Identifying personal stories in millions of weblog entries”. In: *Third International Conference on Weblogs and Social Media, Data Challenge Workshop, San Jose, CA*. Vol. 46. 2009.
- [29] Anuradha Goswami and Ajey Kumar. “A survey of event detection techniques in online social networks”. In: *Social Network Analysis and Mining* 6.1 (Nov. 2016), p. 107. DOI: 10.1007/s13278-016-0414-1. URL: <https://doi.org/10.1007/s13278-016-0414-1>.
- [30] Anja Gruenheid, Donald Kossmann, Theodoros Rekatsinas, and Divesh Srivastava. “StoryPivot: comparing and contrasting story evolution”. In: *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. ACM. 2015, pp. 1415–1420.
- [31] Ehab Hassan, Davide Buscaldi, and Aldo Gangemi. “Event-Based Recognition of Lived Experiences in User Reviews”. In: *European Knowledge Acquisition Workshop*. Springer. 2016, pp. 320–336.
- [32] Yuki Hattori and Akiyo Nadamoto. “Extracting tip information from social media”. In: *Proceedings of the 14th International Conference on Information Integration and Web-based Applications & Services*. ACM. 2012, pp. 205–212.
- [33] Yuki Hattori and Akiyo Nadamoto. “Tip information from social media based on topic detection”. In: *International Journal of Web Information Systems* 9.1 (2013), pp. 83–94. DOI: 10.1108/17440081311316406. eprint: <https://doi.org/10.1108/17440081311316406>. URL: <https://doi.org/10.1108/17440081311316406>.
- [34] Fatemeh Hemmatian and Mohammad Karim Sohrabi. “A survey on classification techniques for opinion mining and sentiment analysis”. In: *Artificial Intelligence Review* (2019), pp. 1–51.
- [35] Kentaro Inui, Shuya Abe, Kazuo Hara, Hiraku Morita, Chitose Sao, Megumi Eguchi, Asuka Sumida, Koji Murakami, and Suguru Matsuyoshi. “Experience mining: Building a large-scale database of personal experiences and opinions from web documents”. In: *Proceedings of the 2008 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology-Volume 01*. IEEE Computer Society. 2008, pp. 314–321.

- [36] Keyuan Jiang, Ravish Gupta, Matrika Gupta, Ricardo A Calix, and Gordon R Bernard. “Identifying personal health experience tweets with deep neural networks”. In: *39th IEEE Annual International Conference of Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2017, pp. 1174–1177.
- [37] Valentin Jijkoun, Maarten de Rijke, Wouter Weerkamp, Paul Ackermans, and Gijs Geleijnse. “Mining user experiences from online forums: an exploration”. In: *Proceedings of the NAACL HLT 2010 Workshop on Computational Linguistics in a World of Social Media*. Association for Computational Linguistics. 2010, pp. 17–18.
- [38] Houda Khrouf, Vuk Milicic, and Raphaël Troncy. “Mining events connections on the social web: Real-time instance matching and data analysis in EventMedia”. In: *Web Semantics: Science, Services and Agents on the World Wide Web 24* (2014), pp. 3–10.
- [39] Matthew Krawczyk and Zheng Xiang. “Perceptual mapping of hotel brands using online reviews: a text analytics approach”. In: *Information Technology & Tourism* 16.1 (2016), pp. 23–43.
- [40] Takeshi Kurashima, Ko Fujimura, and Hidenori Okuda. “Discovering association rules on experiences from large-scale blog entries”. In: *European Conference on Information Retrieval*. Springer. 2009, pp. 546–553.
- [41] Takeshi Kurashima, Taro Tezuka, and Katsumi Tanaka. “Mining and visualizing local experiences from blog entries”. In: *DEXA*. Springer. 2006, pp. 213–222.
- [42] Erdal Kuzey, Jilles Vreeken, and Gerhard Weikum. “A fresh look on knowledge bases: Distilling named events from news”. In: *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. ACM. 2014, pp. 1689–1698.
- [43] William Labov and Joshua Waletzky. “Narrative analysis. Essays on the verbal and visual arts”. In: *Essays on the verbal and visual arts*. Ed. by June Helm and June Helm MacNeish. University of Washington Press, Seattle, 1967.
- [44] Kyung-Soon Lee and Kyo Kageura. “Korean–Japanese story link detection based on distributional and contrastive properties of event terms”. In: *Information processing & management* 42.2 (2006), pp. 538–550.

- [45] Kyung-Soon Lee and Kyo Kageura. “Multilingual story link detection based on event term weighting on times and multilingual spaces”. In: *International Conference on Asian Digital Libraries*. Springer. 2004, pp. 398–407.
- [46] Marco Lippi and Paolo Torroni. “Argumentation mining: State of the art and emerging trends”. In: *ACM Transactions on Internet Technology (TOIT)* 16.2 (2016), p. 10.
- [47] Marco Lippi and Paolo Torroni. “MARGOT: A web server for argumentation mining”. In: *Expert Systems with Applications* 65 (2016), pp. 292–303.
- [48] Yunzhong Liu, Yi Chen, Jiliang Tang, and Huan Liu. “Context-aware experience extraction from online health forums”. In: *2015 International Conference on Healthcare Informatics (ICHI)*. IEEE. 2015, pp. 42–47.
- [49] Christopher Marshall and Pearl Brereton. “Tools to support systematic literature reviews in software engineering: A mapping study”. In: *2013 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*. IEEE. 2013, pp. 296–299.
- [50] Christopher Marshall, Pearl Brereton, and Barbara Kitchenham. “Tools to support systematic reviews in software engineering: a feature analysis”. In: *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering*. ACM. 2014, p. 13.
- [51] Dimitriy V Masterov, Uwe F Mayer, and Steven Tadelis. “Canary in the e-commerce coal mine: Detecting and predicting poor experiences using buyer-to-seller messages”. In: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM. 2015, pp. 81–93.
- [52] Béatrice Mazoyer, Julia Cagé, Céline Hudelot, and Marie-Luce Viaud¹. “Real-time collection of reliable and representative tweets datasets related to news events”. In: *BroDyn 2018: 1st Workshop on Analysis of Broad Dynamic Topics over Social Media*. 2018.
- [53] Hye-Jin Min and Jong C Park. “Identifying helpful reviews based on customer’s mentions about experiences”. In: *Expert Systems with Applications* 39.15 (2012), pp. 11830–11838.
- [54] Raquel Mochales and Marie-Francine Moens. “Argumentation mining”. In: *Artificial Intelligence and Law* 19.1 (2011), pp. 1–22.

- [55] Federico Nanni, Simone Paolo Ponzetto, and Laura Dietz. “Building entity–centric event collections”. In: *Proceedings of the 17th ACM/IEEE Joint Conference on Digital Libraries*. IEEE Press. 2017, pp. 199–208.
- [56] Samira Noferesti and Mehrnoush Shamsfard. “Using Linked Data for polarity classification of patients’ experiences”. In: *Journal of Biomedical Informatics* 57 (2015), pp. 6–19. ISSN: 1532-0464. DOI: <https://doi.org/10.1016/j.jbi.2015.06.017>. URL: <http://www.sciencedirect.com/science/article/pii/S1532046415001276>.
- [57] Alison O’Mara–Eves, James Thomas, John McNaught, Makoto Miwa, and Sophia Ananiadou. “Using text mining for study identification in systematic reviews: a systematic review of current approaches”. In: *Systematic Reviews* 4.1 (Jan. 2015), p. 5. ISSN: 2046-4053. DOI: 10.1186/2046-4053-4-5. URL: <https://doi.org/10.1186/2046-4053-4-5>.
- [58] Myle Ott, Yejin Choi, Claire Cardie, and Jeffrey T Hancock. “Finding deceptive opinion spam by any stretch of the imagination”. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies–Volume 1*. Association for Computational Linguistics. 2011, pp. 309–319.
- [59] Keun Chan Park, Yoonjae Jeong, and Sung Hyon Myaeng. “Detecting experiences from weblogs”. In: *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics. 2010, pp. 1464–1472.
- [60] Chris Parnin and Christoph Treude. “Measuring API documentation on the web”. In: *Proceedings of the 2nd international Workshop on Web 2.0 for Software Engineering*. ACM. 2011, pp. 25–30.
- [61] Andreas Peldszus and Manfred Stede. “From argument diagrams to argumentation mining in texts: A survey”. In: *International Journal of Cognitive Informatics and Natural Intelligence (IJCINI)* 7.1 (2013), pp. 1–31.
- [62] Saša Petrović, Miles Osborne, and Victor Lavrenko. “Streaming first story detection with application to twitter”. In: *Human language technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics. 2010, pp. 181–189.

- [63] Saša Petrović, Miles Osborne, and Victor Lavrenko. “Using Paraphrases for Improving First Story Detection in News and Twitter”. In: *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. NAACL HLT '12*. Montreal, Canada: Association for Computational Linguistics, 2012, pp. 338–346. ISBN: 978-1-937284-20-6. URL: <http://dl.acm.org/citation.cfm?id=2382029.2382072>.
- [64] Arun Qamra, Belle Tseng, and Edward Y Chang. “Mining blog stories using community-based and temporal clustering”. In: *Proceedings of the 15th ACM International Conference on Information and Knowledge Management*. ACM. 2006, pp. 58–67.
- [65] Elahe Rahimtoroghi, Reid Swanson, Marilyn A Walker, and Thomas Corcoran. “Evaluation, orientation, and action in interactive storytelling”. In: *Proceedings of Intelligent Narrative Technologies 6* (2013).
- [66] Austen Rainer. “Anaya’s Journey: A Vision for a Future Software Academy”. In: *IEEE Software* 37.2 (2020), pp. 95–96.
- [67] Austen Rainer. “Using argumentation theory to analyse software practitioners’ defeasible evidence, inference and belief”. In: *Information and Software Technology* (2017).
- [68] Austen Rainer, Tracy Hall, and Nathan Baddoo. “Persuading developers to “buy into” software process improvement: a local opinion and empirical evidence”. In: *Proceedings of the 2003 International Symposium on Empirical Software Engineering (ISESE 2003)*. IEEE. 2003, pp. 326–335.
- [69] Austen Rainer and Ashley Williams. “Heuristics for improving the rigour and relevance of grey literature searches for software engineering research”. In: *Information and Software Technology* 106 (2019), pp. 231–233.
- [70] Austen Rainer and Ashley Williams. “Using blog-like documents to investigate software practice: Benefits, challenges, and research directions”. In: *Journal of Software: Evolution and Process* 31.11 (2019), e2197.

- [71] Paramvir Singh and Karanpreet Singh. “Exploring Automatic Search in Digital Libraries: A Caution Guide for Systematic Reviewers”. In: *Proceedings of the 21st International Conference on Evaluation and Assessment in Software Engineering*. ACM. 2017, pp. 236–241.
- [72] Jacopo Soldani, Damian Andrew Tamburri, and Willem–Jan Van Den Heuvel. “The Pains and Gains of Microservices: A Systematic Grey Literature Review”. In: *Journal of Systems and Software* (2018).
- [73] Joel Spolsky. *Language Wars*. URL: <https://www.joelonsoftware.com/2006/09/01/language-wars/>. (accessed: 18.03.2020).
- [74] PK Srijith, Mark Hepple, Kalina Bontcheva, and Daniel Preotiuc-Pietro. “Sub-story detection in Twitter with hierarchical Dirichlet processes”. In: *Information Processing & Management* 53.4 (2017), pp. 989–1003.
- [75] Margaret–Anne Storey, Leif Singer, Brendan Cleary, Fernando Figueira Filho, and Alexey Zagalsky. “The (r) evolution of social media in software engineering”. In: *Proceedings of the Future of Software Engineering*. ACM. 2014, pp. 100–116.
- [76] Reid Swanson, Elahe Rahimtoroghi, Thomas Corcoran, and Marilyn Walker. “Identifying narrative clause types in personal stories”. In: *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 2014, pp. 171–180.
- [77] Guy Tsafnat, Paul Glasziou, Miew Keen Choong, Adam Dunn, Filippo Galgani, and Enrico Coiera. “Systematic review automation technologies”. In: *Systematic Reviews* 3.1 (July 2014), p. 74. ISSN: 2046-4053. DOI: 10.1186/2046-4053-3-74. URL: <https://doi.org/10.1186/2046-4053-3-74>.
- [78] William Twining. *Rethinking evidence: Exploratory essays*. Northwestern University Press, 1994.
- [79] J. Vails-Vargas, J. Zhu, and S. Ontanon. “Towards automatically extracting story graphs from natural language stories”. In: vol. WS-17-01 - WS-17-15. cited By 0. 2017, pp. 1006–1013. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85046093315&partnerID=40&md5=4f5906d4e39d5b384381758d698a589e>.
- [80] Zeno Vendler. *Linguistics in Philosophy. ithaca*. 1967.

- [81] Max L Wilson, Susan Ali, and Michel F Valstar. “Finding information about mental health in microblogging platforms: a case study of depression”. In: *Proceedings of the 5th Information Interaction in Context Symposium*. ACM. 2014, pp. 8–17.
- [82] Jia Yu, Lei Xie, Xiong Xiao, and Eng Siong Chng. “Learning distributed sentence representations for story segmentation”. In: *Signal Processing* 142 (2018), pp. 403–411.

Table 1: Illustrative examples of experience

#	Verbatim quote for example
<i>Example E1: Description of behaviour</i>	
E1.1	“Last summer when we had a <u>group of interns</u> [URL in original] build <u>Copilot</u> [URL in original], we had to decide what language to use for new code. . .
E1.2	. . . What we ended up doing with the interns was just telling them that they had to build it in C# and ASP.NET. In particular, one intern, who wrote the website part of Copilot, had enough experience with ASP.NET to know what things to avoid (like viewstate) and knew to avoid the gotchas that make it impossible to have two <u>forms</u> in one page, etc. etc., so he did a beautiful job architecting the ASP.NET code exactly the right way to begin with so we didn’t get into trouble later. And the benefit was that not one minute was spent debating the merits of programming languages, a fruitless debate if I’ve ever seen one.”
<i>Example E2: Description of technology used and developed</i>	
E2	“Finally – as to what we use – Copilot is C# and ASP.Net, as I mentioned, although the Windows client is written in C++. Our older in-house code is VBScript and our newer in-house code is C#. FogBugz is written in Wasabi, a very advanced, functional-programming dialect of Basic with closures and lambdas and Rails-like active records that can be compiled down to VBScript, JavaScript, PHP4 or PHP5. Wasabi is a private, in-house language written by one of our best developers that is optimized specifically for developing FogBugz; the Wasabi compiler itself is written in C#.”
<i>Example E3: Observation of external event</i>	
E3	“. . . yes I’ve heard of 37 Signals and they’re making lovely Ruby on Rails apps, and making lots of money, but that’s not a <i>safe</i> choice for at least another year or six. I for one am scared of Ruby because (1) it displays a stunning antipathy towards Unicode and (2) it’s known to be slow. . .” (emphasis in original)
<i>Example E4: Observation of external event</i>	
E4	“Oh and I know <u>Paul told you</u> [URL in original] that he made his app in Lisp and then he made millions of dollars because he made his app in Lisp, but honestly <u>only two people</u> [URL in original] ever believed him and, a complete rewrite later, they won’t make <i>that</i> mistake again. . .” (emphasis in original)
<i>Example E5: Indication of author’s accumulated experience</i>	
E5	“ About the Author: I’m your host, Joel Spolsky, a software developer in New York City. Since 2000, I’ve been writing about software development, management, business, and the Internet on this site. For my day job, I run <u>Fog Creek Software</u> [URL in original], makers of <u>FogBugz</u> [URL in original] — the smart bug tracking software with the stupid name, and <u>Fog Creek Copilot</u> [URL in original] — the easiest way to provide remote tech support over the Internet, with nothing to install or configure.” (emphasis in original)

Table 2: Illustrative ‘negative examples’, i.e., of non–experience

#	Verbatim quote for example
<i>Example E6: An argument about evidence</i>	
E6	“The safe answer, for the Big Enterprisy Thing where you have no interest in being on the cutting edge, is C#, Java, PHP, or Python, since there’s so much evidence that when it comes right down to it zillions of people are building huge business–critical things in those languages and while they may have problems, they’re not life–threatening problems.”
<i>Example E7: An argument about experience</i>	
E7	“How do you decide between C#, Java, PHP, and Python? The only real difference is which <i>one</i> you know better. If you have a serious Java guru on your team who has build several large systems successfully with Java, you’re going to be a hell of a lot more successful with Java than with C#, not because Java is a better language (it’s not, but the differences are too minor to matter) but because he knows it better. Etc” (emphasis in original)
<i>Example E8: Something</i>	
E8	“A while ago I wrote an article called <u>Lord Palmerston on Programming</u> [URL in original] in which I claimed that some of these programming worlds, like .NET and Java, were so huge and complicated that you never could really tell if they were going to be good enough for your needs until it was too late.” (emphasis in original)

Table 3: Summary of searches conducted

#	Search Strings	Results per search engine				
		Google Scholar	IEEE Xplore	ACM	Science Direct	Scopus
1	“experience mining”	664	8	9	40	105
2	“experience indicators”	1,850	2	534	160	82
3	indicators AND “personal experience”	132,000	3	25	9,977	4,148
4	indicators AND “professional experience”	43,500	1	6	2,582	549
5	indicators AND experience	3,680,000	1,199	5,176	407,703	291,425
6	allintitle: indicators personal experience	2	81	0	0	1
7	allintitle: indicators personal experience online articles	0	27	0	0	0
8	allintitle: indicators experience web	3	8,675	1	0	0
9	allintitle: “experience indicators”	38	0	0	0	7
10	allintitle: detecting experience	266	3,118	65	87	162
11	allintitle: “detecting experience”	2	0	4	0	2
12	allintitle: “detecting” AND “personal experience”	4	0	0	0	1
13	allintitle: “detecting” AND “professional experience”	0	0	0	0	0
14	“detecting professional experience”	0	0	0	0	0
15	“detecting personal experience”	32	0	0	3	0
16	“detecting experience”	67	0	30	62	26
17	“story mining”	209	2	1	5	12
18	“detecting stories”	63	1	5	19	25

Table 4: Summary of filtering

Filtering step	Remaining articles
Initial results from searches	309
Obvious duplicate articles removed ($n = 71$)	238
Review article’s abstract, introduction and conclusion to remove irrelevant articles, i.e., did the article discuss detecting experience, stories or events? ($n = 154$)	84
Non-accessible articles removed ($n = 21$)	63
No abstract or article ($n = 15$)	
Abstract available, but not article ($n = 6$)	
Apply exclusion criteria, e.g., patents, similar papers, datasets only, theses & duplicate ($n = 21$)	42

Table 5: Classification of the retained articles

Ref	Year	Nw	SP	PE	PI	Bl	Em	Domain	St	Tw
<i>Articles about the experience of software practice</i>										
[22]	2017		•	•			•	Mobile app		
<i>Articles that represent non-software professional experience in terms of story</i>										
[5]	2011			•				Processes	•	
[4]	2010			•				Processes	•	
<i>Articles that represent non-professional experience in terms of story</i>										
[76]	2014				•	•	•	Personal	•	
[15]	2012				•			Extremists	•	
[30]	2015	•			•	•		News	•	
[19]	2017	•			•			News	•	•
[74]	2017	•			•			News	•	•
[63]	2012	•			•			News	•	•
[62]	2010	•			•			News	•	•
[82]	2018	•						News	•	
[44]	2006	•						News	•	
[13]	2015				•	•		Other	•	
[64]	2006				•	•		Other	•	
[8]	2016				•			Other	•	
[6]	2015							Other	•	
<i>Other articles</i>										
[59]	2010				•	•	•	Generic		
[40]	2009				•	•	•	Generic		
[35]	2008				•	•	•	Generic		
[52]	2018	•			•			News		•
[55]	2017	•						News		
[42]	2014	•						News		
[45]	2004	•						News		
[9]	2003	•						News		
[38]	2014				•			Events		
[1]	2011				•	•	•	Other		
[41]	2006				•	•		Visiting		
[12]	2017				•			Health		•
[36]	2017				•		•	Health		•
[31]	2016				•			Product		
[39]	2016				•			Hotel		
[51]	2015				•			Retail		
[48]	2015				•			Health		
[56]	2015				•		•	Patient		
[81]	2014				•		•	Health		•
[32]	2012				•		•	Other		
[53]	2012				•		•	Product		
[23]	2010				•		•	Hotel		
[37]	2010				•		•	Product		
[79]	2017							Other		
[29]	2016							Other		
[33]	2013							Other		

Legend – acronyms:

Nw: News SP: Software practice

PE: Professional experience PI: Personal experience

Bl: Blog posts Em: Emotional experience St: Stories Tw: Twitter

Table 6: Summary classification of the subset of papers

First author	Ref	Date	Prior theory	Model of exp'nce	Unit of discourse	Story Pipeline	
Hassan	[31]	2016	No	No	Sentence	No	Yes
Swanson	[76]	2014	Yes: Labov & Waletzky [43]	Yes	Clause	Yes	No
Gonçalves	[4]	2010	Implied (see [5]): CREWS [2]	Yes	Sentence	Yes	Yes
Park	[59]	2010	Yes: Vendler [80]	Yes	Sentence	No	No
Kurashima	[40]	2009	No	Yes	Document	No	No
Inui	[35]	2008	No	Yes	Sentence	No	No
Kurashima	[41]	2006	No	No	Sentence	No	Yes
Ceran	[14]	2012	No	Yes	Sentence	Yes	Yes

Table 7: Summary of the datasets and their annotation

Author	Studied size, overall size, and source/s	Lang.	#	κ
Hassan [31]	383 true reviews (268 training, 115 test) from 400 true reviews [58]	English	1	–
Swanson [76]	50 personal stories drawn from 5000 posts [28] taken from 44M articles [10]	English	3	0.58
Gonçalves [4]	One case selected from stories generated for the study, relating to one business process	Portug'se	–	–
Park [59]	588 sentences from 6000 blog posts on Wordpress	English	3	–
Kurashima [40]	29M blog posts from 48M blog posts collected using the Blogranger 2.0 API	–	3	–
Inui [35]	50M posts, related to pre-selected topics, from 150M weblog posts [No stated source]	Japanese	2	0.68
Kurashima [41]	62,396 articles from Two Japanese blog hosting sites	Japanese	–	–
Ceran [14]	16,930 paragraphs from 1,256 documents (13629 non-story, 3301 story)[No stated source]	English	–	0.83

An em dash denotes the source paper does not report the information

Table 8: Elements of conceptual models of experience

Elements of story based models of experience

Swanson [76]	<i>Action</i> : clauses describing causal or event relationships. A story must have an <i>actor</i> or actors. <i>Orientation</i> : clauses describing properties about settings (e.g. time and location) and actors. <i>Evaluation</i> : clauses describing emotional reactions, and the reason for telling the story. <i>Abstract</i> : an initial clause summarising the entire sequence of events. <i>Coda</i> : a final clause ending the story and providing a 'moral' to the story.
Gonçalves [4]	Uses a subset of the CREWS scenario meta-model [2], focusing on Activities (<i>Actions</i> and <i>Flows of actions</i> in CREWS) and Actors (<i>Objects</i> in CREWS).
Ceran [14]	A story must have an <i>actor</i> or actors. The actor/s must be performing <i>actions</i> . The actions must result in a <i>resolution</i> .

Elements of non-story based models of experience

Park [59]	The authors state that experiences can be recorded in terms of attributes, such as: <i>location</i> , <i>time</i> and <i>activity</i> and their <i>relations</i> . They focus on two of Vendler's [80] four basic classes of verbs: <i>states</i> and <i>activities</i> (excluding <i>achievements</i> and <i>accomplishments</i>).
Inui [35]	Each experience is represented as a piece of structured information comprising: A person's experience is expressed by five attributes: <i>Topic</i> : what the experience is about. <i>Experiencer</i> : the person experiencing the event. <i>Event expression</i> : The event that is experienced. <i>Event type</i> : The semantic type of the event (not defined in the paper). <i>Factuality</i> : Whether the event indeed took place i.e the temporal and modal status of the event <i>Source pointer</i> : a reference to the source text
Kurashima [40]	A person's experience is expressed by five attributes: <i>Time</i> : when a person (actor) acted. <i>Location</i> : Where the person acted. <i>Activity</i> : An action and its object. <i>Opinion</i> : an evaluation about the object. <i>Emotion</i> : The feeling of the person.

Table 9: Challenges of mapping conceptual elements to linguistic elements (for story-based experience)

Swanson [76] *et al.* identify six primary sources of annotator disagreement:

- 1 Clauses of more than one category are common with types of *action* and *evaluation* and e.g. containing elements of *orientation*, *action* and *evaluation*.
- 2 There are *actions* implied but not explicitly stated in the text.
- 3 Stative descriptions of the world, as a result of an *action*, are not intuitively *orientation*.
- 4 Stative descriptions of the world that are localised to a specific place in the narrative does not easily align with *orientation*
- 5 Disambiguating the functional purpose of clauses that describe *actions*, but may be intended to be *orientation*
- 6 Disambiguating the functional purpose of subjective language in the description of an event or state e.g. the *gigantic* tree, the *radiant* blue of the sky.

Ceran [14] *et al.* make the following inferences:

- 1 Stories will have a lower proportion of stative verbs than non-stories, because stories describe actions.
- 2 Stories will include more named entities, especially person names, than non-stories.
- 3 Stories will use more personal pronouns than non-stories.
- 4 Stories may repeat similar nouns.
- 5 Paragraphs with stories in them will have different sentence lengths than paragraphs without stories [it is not clear whether sentences of story-paragraphs would be longer or shorter.]

Table 10: Guiding questions for primary studies investigating the mining of professional experience

#	Question [with cross-reference to section §, Table or Figure]
<i>Theory</i>	
T1	How is the study informed by an appropriate theory of experience? A theory would support generalisation and/or aggregation of independently conducted studies. [cf. §6.3]
T2	How is the study positioned relative to, and distinct from, fields such as opinion mining, event detection, argument mining etc.? [cf. §3]
<i>Definitions</i>	
D1	What is your definition of professional experience of software practice? [cf. §6.3]
D2	Is your definition derived from theory? [cf. T1 and §6.3]
D3	Do you have appropriate sub-definitions of experience, e.g., different domains of software practice, different domains of experience? [cf. §6.3, Figure 2]
<i>Representation of experience</i>	
R1	What representation, or representations, of experience will you seek to retrieve and mine? [cf. §6.3] Story is a common representation, but even story is defined and measured in different ways.
<i>Natural language</i>	
L1	What natural language or languages will you retrieve and mine, e.g., English, Japanese, Portuguese? [cf. Table 7]. Research appears to have focused more on English and Japanese.
<i>Social media</i>	
S1	From what social media will you retrieve information? [cf. §6.2]
<i>Units of discourse</i>	
U1	What is your basic unit/s of discourse, e.g., sentence, clause, paragraph? [cf. §6.5]
<i>Measures and features</i>	
M1	How will you map your definition of experience, and representation of experience (e.g., story) to measures and features in the dataset? [cf. §6.1]
<i>Annotation</i>	
A1	How will you design, conduct and report the annotation process? [cf. §6.6]
A2	How will you handle disagreements during annotation?
<i>Technologies</i>	
X1	What technologies (e.g., natural language tools) will you use and how will you report those in a way that allows future replication, evaluation and comparison? [cf. §6.7, 6.8]
<i>Open science</i>	
O1	How will you make your study artefacts (e.g., source data) available for reuse by others, including replication? [cf. §6.8]