

Finding expertise using online community dialogue and the Duality of Expertise

Michael Niemann

Faculty of Information Technology,
Monash University

Melbourne, AUSTRALIA

michael.niemann@monash.edu

Abstract

The Duality of Expertise considers the “Expert” to be a social role dependent on an individual’s expertise claims and the opinion of their community towards those claims. These are the internal and external aspects of a person’s expertise. My Expertise Model incorporates this duality in a process designed for expertise finding software in an online community forum. In this model, a posting’s term usage is evidence of expertise claims. The dialogue acts in replies to those postings are evidence of the community’s opinion. The model’s preprocessing element uses a bricolage of linguistic and IR tools and methods in a novel way to construct each author’s expertise profile. For any topic query, the profiles are ranked to determine the Community Topic Expert. A series of experiments demonstrate the advantage of utilising the Duality of Expertise when ranking experts rather than just the internal or external aspects of expertise.

1 Introduction

The Internet provides people and organisations with access to new resources for problem solving and guidance. While there are skillful staff or friends within their own networks, sometimes they need to look outside their own group to find the required knowledge. Specialised online communities are one such source of expertise.

Traditionally, expertise is treated as a combination of knowledge, training and experience (Ericsson, 2000; Gould, 1999). However, expertise is also relative to the context in which it is sought (Mieg, 2006). Who is regarded as a suitable topic expert depends on who is available and the depth and breadth of their expertise. Any local resident

may be able to tell you where the local train station is, but you may need a railway employee if you want know how regular the trains are.

Likewise, it is not enough to simply determine who in an online community knows something about a topic. Many “know-it-all” profess to be experts but do not have much expertise. For guidance, we look to others in the community for advice on who they consider to be Community Topic Experts. This “Expert” label is a social role bestowed by the community (Mieg, 2006). It is relative to the community’s knowledge on the topic and the expertise they have encountered when interacting with members. Who is a suitable expert depends on the community’s opinion of the relative expertise of its members.

My research looks at modelling and identifying someone’s expertise by considering both their expertise claims and their community’s opinion towards those claims. This Duality of Expertise is investigated by examining the linguistic interactions within an online community’s forum. The content of a forum author’s postings is evidence of their expertise claims. The dialogue acts of the community’s replies to those postings is evidence of the community’s opinion.

My Expertise Model utilises the Duality of Expertise in such a way that the model is easily incorporated within expertise finding software for such a forum. This is evaluated through experiments based on the TREC2006 Expert Search task and the related forum postings from the W3C corpus.

This paper is structured as follows. Section 2 discusses previous research relating to expertise and expertise finding technology. The approach and model used in my research is outlined in Section 3. Section 4 explains the experiments run to evaluate the model. Their outcomes are discussed in Section 5. The conclusion in Section 6 summarises my research.

2 Related research

Research has identified various aspects of what makes up an expert. An expert is traditionally seen as someone who is knowledgeable. However, Ericsson (2000) argues that an expert learns how to use that knowledge through deliberate practice. This expertise is only related to one topic or field (Gould, 1999) but every expert is different as their experience and reflections of their actions are unique (Bellot, 2006; Schön, 1983), as is the manner in which they utilise those skills (Rolfe, 1997). For this reason, a person's expertise cannot be simply defined by labels or fields e.g., 'law'.

Simple "yellow pages" software systems try to use such labels to declare the expertise of an organisation's staff (Ackerman and Halverson, 2003). People searching for expertise (expertise seekers) can use simple search engines to search through related databases or staff web-pages for possible experts (Yimam-Seid and Kobsa, 2003), but these systems rely on staff to update them and are often incomplete or inaccurate with little quality control (Lemons and Trees, 2013).

More specialised expert finding systems have been developed to form expert profiles based on evidence of people's expertise. This evidence ranges from authoring academic and company papers (Richards et al., 2008; Yan et al., 2006), being mentioned on web-pages (Balog et al., 2006; Campbell et al., 2003; Liu et al., 2012) or social media and email communications (Guy et al., 2013; Kautz et al., 1997; Kumar and Ahmad, 2012). The survey by Balog et al. (2012) shows that expertise finding is often treated very much as an information retrieval (IR) problem. The relevance of a person's expertise to the specific topic is commonly evaluated based on simple term usage, ranking candidate experts using a probabilistic or vector-based model.

Some early researchers have considered the social network of experts. Schwartz and Wood (1993) formed specialization subgraphs (SSG) based on emails between people. They found each author's Aggregate Specialisation Graph (ASG) was particular to them, supporting the concept of expertise being particular to the individual. The graphs can be used to refer an expertise seeker to someone they know, who may then refer them to someone they consider to be an expert based on previous interactions. Similarly, the Referral Web system described by Kautz et al. (1997) links peo-

ple to experts through the co-occurrence of their names in a document, like an email, research article or web-page. This relies on the social network associations being related to particular areas of expertise, yet the manner in which the associated graph is constructed is very ad-hoc.

More recently, Guy et al. (2013) investigated using social media for expertise finding and found that feedback to social media messages can be a good indication of people's expertise. The ComEx Miner system (Kumar and Ahmad, 2012) attempted to identify the sentiment in feedback to blog entries using lists of adverbs and positive, negative and neutral adjectives. However, a blog provides very one-sided discussions as the opening post is always by the same author and it is not always clear whether comments are replying to the blog or other comments. Therefore a blog may not be a good example of community interaction.

Weigand considers dialogue to be a dialogic act "to negotiate our purposes" (Weigand, 2010, p. 509) through language. Part of that purpose is to be accepted in the community through interaction. Dialogue acts represent the intentions of speakers (Austin, 1975; Searle, 1969; Searle, 1975).¹ Traditionally these are applied to utterances but researchers have attempted to classify the dialogue acts of email sentences (Cohen et al., 2004; Jeong et al., 2009) and entire emails (Carvalho and Cohen, 2005; Feng et al., 2006) and forum messages (Bhatia et al., 2012; Kim et al., 2010). To Weigand, each dialogue act has a corresponding reactive action. Socially, the dialogue act is a claim, such that the claimant desires the suitable reaction to fulfill this claim (Weigand, 2010). Thus dialogue is a negotiation between participants who respond to each other's dialogue acts whilst trying to achieve their own objectives, including social acceptance within the community.

Likewise McDonald and Ackerman (1998) and Hytönen et al. (2014) found that when seeking assistance from experts, people often had to consider various psychological costs like a potential loss of status, expected reciprocity and social equity. Through the sharing of knowledge and the use and recognition of expertise during the interaction, dialogue participants negotiate an outcome that meets their personal objectives (Mieg, 2001). This may include establishing the value of the exchange and nature of the truth. Similarly, in an online fo-

¹ The term 'dialogue act' is used in this paper rather than 'speech act' due to the absence of speech in my data.

rum the dialogue acts of the replies are responses to the content and intent of the previous posting and are indicative of the forum community’s opinion towards the author’s proposals. Therefore, the group’s opinion of an author’s claims and thus their expertise can be judged by examining the dialogue acts in the group’s replies to the author.

3 Model

My Expertise Model is a process that enables the incorporation of the Duality of Expertise in an expertise finding system (Figure 1). The Duality of Expertise is a relationship between an expert and their community, based on their expertise claims and the community’s opinion of those claims.

A person’s expertise claims are their representation of the topics about which they assert to be knowledgeable. An internal aspect of expertise, it relates to how an individual presents themselves to others. The claim demonstrates their topics of interest, but does not judge the accuracy of the claim nor whether they are an expert on the topic.

The external aspect of expertise is the community’s opinion towards the relative expertise of the member claiming expertise. Based on the claims as well as other interactions and expertise within the community, the community judges whether the person is the Community Topic Expert, or whether they are not as knowledgeable on the topic as others in the community.

The Expertise Model is designed to evaluate and find expertise within online communities by examining forum postings. Each discussion thread is a linguistic interaction between community members with each posting being an author’s contribution to the community’s knowledge. Term usage in postings is evidence of the author’s expertise claims. It is assumed authors claim expertise on what they write about. The community’s opinion is recognised through the responses to a member’s postings and expertise claims. The dialogue acts in reply postings are evidence of this opinion.

The Expertise Model uses a combination of these internal and external aspects of expertise to construct expertise profiles for each author then evaluate the relevance of their profile to the topic of expertise being sought. The outcome is a list of authors ranked according to whether they are the Community Topic Expert.

There are four main stages in the model: pre-processing, profiling, topic querying and ranking.

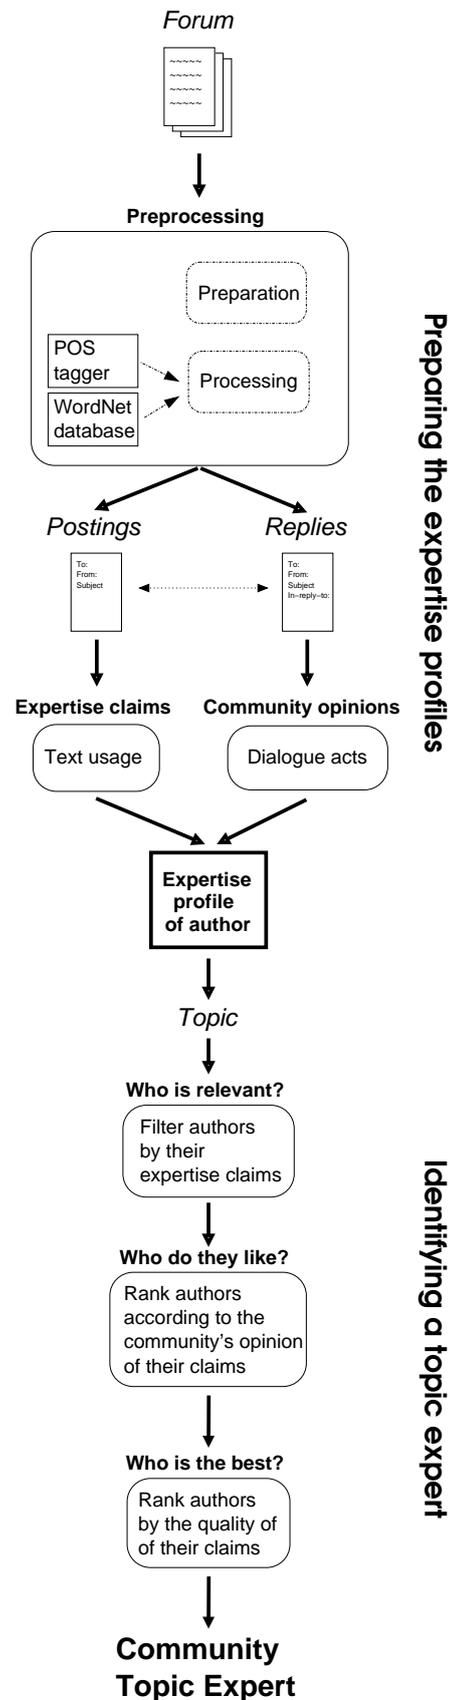


Figure 1: Expertise Model

3.1 Preprocessing

For a given forum, the preprocessing prepares the data to a standardised format and processes it according to various linguistic criteria. The preparation includes extracting postings from their source files (e.g., web-pages or digests), standardising the form of the metadata, and identifying quotations and non-dialogue lines in the postings. The linguistic processing includes sentence segmentation, term tokenisation, part-of-speech tagging, lemmatisation, and identification of the semantic concepts associated with the terms. These tasks can be completed using third-party software like a part-of-speech tagger and the WORDNET lexical database, but many expertise finding systems I have reviewed did not include such preprocessing.

The preprocessing also identifies the dialogue acts in the reply postings. The dialogue acts used (Table 1) were based on those used in the Verbomobil, TRAINS and DAMSL research (Alexandersson et al., 1995; Core, 1998; Ferguson et al., 1996). The decision to simplify the number of acts to six was made after reviewing the dialogue in the 20 NEWSGROUPS corpus (Rennie, 2008) and the CORVUS corpus which I collected from five professional and semi-professional mailing lists. These acts were broadly related to the historical attributes of experts, e.g., supplying and seeking information and reflection, the community's attitude when responding, e.g., support, rejection or enquiry, and other acts not related to expertise. The dialogue studied was found to be more a discussion than questions and answers so less focus was given to related acts (e.g., *Answer*, *Clarification*). The six acts used subsumed these acts.

3.2 Profiling

For the profiling, the preprocessing established metadata for forum discussions, including lists of each author's postings and the reply postings. This enables the profiling to be divided into the data relating to the internal aspect of expertise, being the term and semantic concept usage of each author, and the external aspect of expertise, namely the dialogue acts in the replies to each author's postings. This data is indexed per author and can be updated whenever new postings appear in a forum.

3.3 Topic querying

Any expertise finding system needs to interface with an expertise seeker to determine what exper-

Dialogue act	Example sub-categories subsumed
Inform	Inform, Answer, Clarification, Suggestion, Explanation, Order, Instruction, Statement, Opinion, Signal Not Understanding,
Positive	Agreement, Acceptance, Acknowledgement, Support, Thanks
Negative	Disagreement, Rejection, Criticism
Question	Yes/No, Rhetorical Query, Request Query
Reflection	Reflection, Correction, Experience
Other	Greeting, Bye, Coding, Graphic, Numeric, Quotation, Signature

Table 1: The dialogue acts

tise is sought. For my Expertise Model, the expertise topic is indicated by one or more query terms. There are no restrictions on how many terms or which particular terms can be given as the topic. Just as the profiling does not represent expertise topics by a finite set of labels, neither is there any restriction on the topic query terms. While the interface is presumed to be part of any expertise finding system that may incorporate my Expertise Model, any topic terms still undergo the same linguistic processing as the posting content.

3.4 Ranking

Various ranking methods are used to identify the Community Topic Expert that best meets the user's needs. This ranking utilises existing IR methods in novel ways. The relevance of the term usage is ranked through a combination of the vector space model with the term frequency-inverse document frequency (tf-idf) measure to identify the specialised usage of terms. This allows the filtering of authors not relevant to the topic as well as ranking their topic relevance when there is ambiguity as to who claims to have greater expertise.

For the community opinion, each author has an opinion vector based on the dialogue acts used in their postings' replies. For each act, two dimensions were added to the vector:

1. Number of replies with at least one instance of the dialogue act
2. Average number of instances of the dialogue act per reply

	Gold Replies	Bad Replies
Reply postings	378	531
Reply authors	169	241
Average postings per author	2.2	2.2
Postings replied to	304	386
Authors replied to	92	199
Average replies per author	4.1	2.7
Annotated Postings	378	531
Annotated Sentences	6731	11908
Postings with at least one		
- Positive DA	121 (32%)	178 (33%)
- Negative DA	20 (5%)	79 (15%)
- Question DA	126 (33%)	179 (34%)
- Reflection DA	6 (2%)	10 (2%)
- Inform DA	338 (89%)	498 (94%)
- Other DA	361 (96%)	499 (94%)
- Positive & Negative DAs	13 (3%)	39 (7%)
- Positive & Question DAs	46 (12%)	83 (16%)
- Negative & Question DAs	12 (3%)	45 (8%)
- Negative, Positive & Question DAs	11 (3%)	30 (6%)
Average quantity in a posting		
- Positive DA	1.9	2.1
- Negative DA	1.6	1.6
- Question DA	2.1	2.3
- Reflection DA	1.1	1.6
- Inform DA	9.3	10.9
- Other DA	8.4	11.0

Table 2: Statistics of the annotated reply postings

The community opinion of an author’s expertise on a topic was scored using the Expertise Opinion Measure (EOM, Equation 1), where $v(q, a)$ is an opinion vector for author a , topic query q and constant α . This formula is similar to Rocchio’s Algorithm in the use of weighted comparisons of an author’s vector to the centroid vectors of relevant and irrelevant results (Manning et al., 2008). The GOLD centroid is formed from the opinion vectors of known experts on the topic q . The BAD centroid is formed from the vectors of non-experts on the topic. The similarity measure (sim) uses a method like cosine comparison.

Expertise Opinion Measure (EOM)

$$EOM(a, q) = \alpha \times sim(v(q, a), centroid(q, GOLD)) - ((1 - \alpha) \times sim(v(q, a), centroid(q, BAD))) \quad (1)$$

Through these measures, the Expertise Model

1. filters out authors whose expertise claims are

not relevant to the topic,

2. determines which authors the community thinks highly of, and
3. determines the best of the topic experts.

Thus the linguistic interaction is used to determine the community’s expert on a required topic on the basis of the author’s claims and the forum community’s opinion towards those claims.

4 Experiments

The Expertise Model was evaluated through a series of experiments, each examining an aspect of how the Duality of Expertise is represented. This evaluation was conducted using a preprocessor and the Lemur INDRI IR system,² modified to act as an expertise finding system. The preprocessor utilised a novel combination of scripts written by me based on established linguistic and IR

² <http://www.lemurproject.org/indri.php>

technologies and methodologies. The preprocessor also made use of the C&C Tools POSTAGGER³ and WORDNET 3.0 for the lemmatisation and semantic relationships.

The test-set was based on the TREC2006 Expertise Search task in the Enterprise track (Soboroff et al., 2006). This task had participants identify experts from a corpus of W3C website files. My test-set only included the pages containing forum postings from W3C mailing lists, about 60% of the original corpus documents (Craswell et al., 2005). Personal homepages and other documents were ignored as they do not relate to the linguistic interactions within the W3C community.

The 49 TREC2006 queries were used, each set of terms referring to a topic of expertise, e.g., ‘SOAP security considerations’. TREC supplied a list of candidate experts that linked identification numbers to names and email addresses, e.g., *candidate-0025 Judy Brewer jbrewer@w3.org*. This was based on a list of people involved in the W3C (Soboroff et al., 2006). TREC2006 gave a “goldlist” for each topic query of who were judged to be relevant experts and who were not experts. This judgement did not consider all candidate experts but was based on the top 20 responses from TREC2006 participants and human judgement, given a review of documents related to the candidates. For my evaluation I increased the list from 1092 to 1844 candidate experts by including any unlisted forum authors and included heuristics in the preprocessor to recognise when an author used a nickname or alternate email address. This allowed each author’s postings to be better identified.

For each topic, the top 50 ranked authors were evaluated using the *trec_eval* software⁴ from TREC. This tool evaluates TREC results in various ways but I focused on the Mean Average Precision (MAP)⁵ and the Interpolated Precision at Recall 0.0. MAP is commonly used as the main measure for TREC participants. The interpolated precision represents the highest precision for any recall level above 0.0 (Buckley and Voorhees, 2005; Manning et al., 2008). If the rank 1 author for a topic is a known expert, the recall

³ <http://svn.ask.it.usyd.edu.au/trac/candc/wiki/POSTagger>

⁴ http://trec.nist.gov/trec_eval/trec_eval_latest.tar.gz

⁵ precision = number of relevant experts returned ÷ number of candidate experts returned

recall = number of relevant experts returned ÷ total number of relevant experts

may be low but the precision will be 1.0. If no known expert for the topic is rank 1, then the interpolated precision at recall 0.0 will be lower, due to the lower rank of the known experts. Therefore, the interpolated precision at recall 0.0 can be considered a measure of the degree to which known experts are given the highest ranks in the results.

The evaluation was divided into three stages (Table 3). First the internal aspect of the Expertise Model (the Knowledge Model or KM) was examined, using only it to determine the community’s experts. Then only the external aspect of the Expertise Model (the Community Model or CM) was utilised for the expertise finding process. Finally, all aspects of the Expertise Model were evaluated in combination. This allowed comparisons to be made between when the aspects are used individually, as is commonly done by other researchers, and when the expertise ranking is conducted based on the Duality of Expertise.

The vector space model IR method provided the baseline for these evaluations. This used the raw, unaltered forms of the topic query terms and sought perfect matches in the postings, the contents of which were indexed according to the posting, not the author. The topic expert was the author of the most relevant posting. No consideration was made of what else each author had posted about outside the posting being ranked.

For the Knowledge Model, there were three main experiments:

- **KM1: Raw terms by author** – The baseline method was modified with all of an author’s postings being indexed together, so their contributions to the forum were examined as a set when considering their expertise claims.
- **KM2: Lemmatised terms by author** – Before indexing, terms were tagged as a noun, verb, adjective, adverb or other, then lemmatised, e.g., ‘antennas’ becomes ‘antenna#n’. This utilises the linguistic processing from the preprocessing and considers the linguistic context of the term usage.
- **KM3: Semantic data by author** – Each lemmatised term was indexed as their corresponding WORDNET synsets, e.g., ‘antenna#n’ is indexed as synsets 02715229, 04843270 and 02584915 because it has three senses in WORDNET. The hyponyms and hypernyms of these synsets were also indexed, e.g., ‘03204955 directional antenna’

Experiment	MAP	Interpolated Precision at Recall 0.0
<i>Baseline: Raw terms by Post</i>	0.0785	0.5539
KM1: Raw Terms by Author	0.1348	0.7065
KM2: Lemmatised Terms by Author	0.1344	0.7355
KM3: Semantic Data by Author	0.1302	0.7113
CM1: Generalised EOM ($\alpha = 0.6$)	0.1156	0.6061
CM2: Non-author EOM ($\alpha = 1.0$)	0.0990	0.6061
CM3: Topic-specific EOM ($\alpha = 0.75$)	0.1250	0.7834
Expertise Model ($\alpha = 0.5$)	0.1461	0.8682

Table 3: Experimental results

and ‘03269401 electrical device’. This extends the idea of terms being evidence of an author’s expertise claims by making each term represent the broader semantic concepts that the claims relate to, not just words.

The experiments for the Community Model utilised the dialogue acts in the replies to the top 50 postings from the baseline experiment. I hand-annotated the dialogue acts using the six acts in Table 1. While other research has attempted to automate similar annotation using a classifier (Cohen et al., 2004; Jeong et al., 2009), my research focused on the effectiveness of the data and the model, not designing a classifier. The Expertise Opinion Measure used the dialogue acts to evaluate community’s opinion of each author’s expertise. After reviewing the spread of annotated dialogue acts (Table 2), I omitted the *Inform*, *Reflection* and *Other* acts from the EOM as they were either too rare or too general. In contrast, the use of *Positive*, *Negative* and *Question* acts seemed to differ depending on whether they were in response to an expert or not. It was hoped that this would aid the Community Model.

Three main experiments were conducted, each training the centroids on different sets of postings.

- **CM1: Generalised EOM** – These centroids used the opinion vectors for experts and non-experts of any query other than the current topic query and excluding the opinion vectors of the expert being ranked.
- **CM2: Non-author EOM** – The opinion vectors of any author relevant to any topic contributed to the centroids, still excluding the current author’s vector.
- **CM3: Topic-specific EOM** – The centroids were formed only from the replies to authors

relevant to the current topic query, including the author being ranked.

The opinion vector of the author being ranked only used dialogue acts from the replies to their topic-related postings. These experiments examined whether the community opinion should be considered as dependent on the topic or whether there can be a single opinion of each author.

Finally, the relevance scores from the Knowledge Model and the Expertise Opinion Measure from the Community Model were used together as shown in Figure 1. The relevance score is used to filter out non-relevant authors, based on each author’s lemmatised term usage and treating their postings as a collection. Authors are then ranked based on the community’s responses to their postings, using the EOM with topic-based centroids. Finally, any tied rankings are resolved using the relevance scores. The top ranked author is judged to be the Community Topic Expert, according to the linguistic evidence of the internal and external aspects of their expertise.

5 Discussion

As shown on Table 3, the Expertise Model achieved the best results but lessons can be learnt from the experiments with the Knowledge and Community Models.

The experiments with the Knowledge Model made it clear that the contents of single documents in isolation cannot be considered good evidence of someone’s expertise. Their expertise claims are better recognised through an examination of all their postings, treating them as a body of alleged knowledge. Further processing like lemmatisation allows the lexical evidence to be better associated with the context in which it was used. While the MAP value for the lemmatisation experiment

(KM2) was similar to that without lemmatisation (KM1), there was a marked improvement in the interpolated precision. This indicates that lemmatisation gave higher ranks to the top experts. Conceptualising the term usage further through the use of WORDNET synsets, hyponyms and hypernyms was not as successful. This was mainly due to ambiguity caused by the multiple senses in WORDNET for each term, as no sense disambiguation was included in the preprocessing because it was not a focus area of my research. However, early experiments tested using synsets alone. The results improved when only each term's first WORDNET sense was used. When the hypernyms and hyponyms were also considered, the results improved further but the best results (as shown for KM3 on Table 3) occurred when hypernyms and hyponyms for only a single sense were considered per term. This suggests that with improved sense handling, the internal aspect of the Expertise Model is best represented by considering all of an expert's lexical contributions to the community and how they are associated with each other through hypernyms and hyponyms.

The experiments with the Community Model were not as successful as those for the Knowledge Model. Various α values were tried for each run with the most successful indicated on Table 3. While no single α value was best for all experiments, there was a general preference for $\alpha > 0.5$. This indicates the importance of an author having a similar opinion vector to that of known experts. However, the results also indicate that with the EOM, the community opinion is best represented when centroids are related to opinion vectors for authors of topic-relevant postings. This was supported by earlier experiments that ranked authors using opinion vectors based on responses to any of their postings, not just relevant ones. These experiments were far less successful with MAP values below 0.1. Therefore, the community's opinion of an author's expertise is topic-specific. The community does not simply consider an expert at one topic to be an expert at all fields, regardless of the specialised nature of the community. This supports the concept of expertise being particular to each individual and dependent on the context in which expertise is sought.

The best results were achieved when the internal and external aspects of expertise are combined in the Expertise Model. The MAP and the in-

terpolated precision for the Expertise Model were clearly better than those for any of the previous experiments. Furthermore, the best results occurred when $\alpha = 0.5$. This differed from the results for the Community Model in the equal weight given to what the EOM considers to be standard community responses to experts and those of non-experts. Therefore, knowing information about non-experts is just as vital as knowing experts.

This demonstrates how the Duality of Expertise can be incorporated in the Expertise Model and automatically identify Community Topic Experts in an online forum. Using a bricolage of freely available linguistic and IR resources and methods, the model processes forum postings in a novel way, enabling the ranking of authors' expertise through the community's linguistic interaction.

In future research, the definitions and choice of the dialogue acts will be reviewed before further annotations. The automated classification of the acts and sense disambiguation will be trialled. Further experiments will use deeper hypernyms and hyponyms to increase the number of synsets associated with each author's expertise. The expertise and dialogue in social media like LinkedIn and Facebook groups will be examined.

6 Conclusion

This research examined the presence of the Duality of Expertise in online community forums. This duality is used within my Expertise Model to determine the Community Topic Expert, considering the expertise claims in their postings and the community's opinion towards these postings and claims. This is achieved through use of the term usage in the postings and the dialogue acts in the replies. Experiments showed that the best representation for the expertise claims is achieved when all of an expert's contributions are considered together and the terms are lemmatised. Results when linking terms to semantic concepts are encouraging. The Expertise Opinion Measure is used to score the community's opinion of each expert based on the similarity of their opinion vector to those of other experts and non-experts. The experiments also showed that the community has a different opinion about every forum author for each topic of expertise sought. This and the Duality of Expertise supports the concept of expertise being relative to the context in which it is found, such that it has internal and external aspects.

Acknowledgments

The usage of data from public online forums was approved by the Monash University Human Research Ethics Committee (MUHREC Project Number CF10/2626 - 2010001463). This research was partly supported by funding from an Australian Postgraduate Award and the Monash University Faculty of Information Technology. Thank you to the reviewers of this paper for their comments.

References

- Mark S. Ackerman and Christine Halverson. 2003. Sharing expertise: The next step for knowledge management. *Social Capital and Information Technology*, pages 273–299.
- Jan Alexandersson, Elisabeth Maier, and Norbert Reithinger. 1995. A robust and efficient three-layered dialogue component for a speech-to-speech translation system. In *Proceedings of the seventh conference on European chapter of the Association for Computational Linguistics*, EACL '95, pages 188–193, San Francisco, U.S.A. Morgan Kaufmann Publishers Inc.
- John L. Austin. 1975. *How to do things with words*. Clarendon Press, Oxford, England, 2nd edition.
- Krisztian Balog, Leif Azzopardi, and Maarten de Rijke. 2006. Formal models for expert finding in enterprise corpora. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR, pages 43–50, New York, U.S.A. ACM.
- Krisztian Balog, Yi Fang, Maarten de Rijke, Pavel Serdyukov, and Luo Si. 2012. Expertise retrieval. *Foundations and Trends in Information Retrieval*, 6(2-3):127–256, February.
- Andrea Bellot. 2006. Advanced practioners' use of reflexivity in decision-making. *Nursing Times*, 102(45):33, December.
- Sumit Bhatia, Prakhar Biyani, and Prasenjit Mitra. 2012. Classifying user messages for managing web forum data. In *Fifteenth International Workshop on the Web and Databases (WebDB 2012)*, WebDB, pages 13–18, Scottsdale, U.S.A.,.
- Chris Buckley and Ellen Voorhees, 2005. *TREC: Experiment and Evaluation in Information Retrieval*, chapter 3. MIT Press.
- Christopher S. Campbell, Paul P. Maglio, Alex Cozzi, and Byron Dom. 2003. Expertise identification using email communications. In *Proceedings of the twelfth international conference on Information and knowledge management*, CIKM, pages 528–531, New York, U.S.A. ACM.
- Vitor R. Carvalho and William W. Cohen. 2005. On the collective classification of email speech acts. In *SIGIR '05: Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR, pages 345–352.
- William W. Cohen, Vitor R. Carvalho, and Tom M. Mitchell. 2004. Learning to classify email into “speech acts”. In Dekang Lin and Dekai Wu, editors, *Proceedings of EMNLP 2004*, EMNLP, pages 309–316, Barcelona, Spain, July. Association for Computational Linguistics.
- Mark G. Core. 1998. Analyzing and predicting patterns of DAMSL utterance tags. Technical report, AAI Technical Report SS-98-01.
- Nick Craswell, Arjen P. de Vries, and Ian Soboroff. 2005. Overview of the TREC-2005 enterprise track. In *The Fourteenth Text REtrieval Conference (TREC 2005) Proceedings*.
- K. Anders Ericsson. 2000. How experts attain and maintain superior performance: Implications for the enhancement of skilled performance in older individuals. *Journal of Aging and Physical Activity*, 8(4):346–352.
- Donghui Feng, Erin Shaw, Jihie Kim, and Eduard Hovy. 2006. Learning to detect conversation focus of threaded discussions. In *Proceedings of HLT-NAACL 2006*, HLTNAACL, pages 208–215.
- George Ferguson, James Allen, and Brad Miller. 1996. TRAINS-95: Towards a mixed-initiative planning assistant. In *Proceedings of the 3rd Conference on AI Planning Systems*, AIPS-95, pages 70–77.
- Nick Gould. 1999. Qualitive practice evaluation. *Evaluation and Social Work Practice*, pages 63–80.
- Ido Guy, Uri Avraham, David Carmel, Sigalit Ur, Michal Jacovi, and Inbal Ronen. 2013. Mining expertise and interests from social media. In *Proceedings of the 22nd International Conference on World Wide Web*, WWW '13, pages 515–526, Republic and Canton of Geneva, Switzerland. International World Wide Web Conferences Steering Committee.
- Kaisa Hytönen, Tuire Palonen, and Kai Hakkarainen. 2014. Cognitively central actors and their personal networks in an energy efficiency training program. *Frontline Learning Research*, 2(2):15–37.
- Minwoo Jeong, Chin-Yew Lin, and Gary Geunbae Lee. 2009. Semi-supervised speech act recognition in emails and forums. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, EMNLP, pages 1250–1259, Singapore, August. Association for Computational Linguistics.
- Henry Kautz, Bart Selman, and Mehul Shah. 1997. ReferralWeb: Combining social networks and collaborative filtering. *Communications of the ACM*, 40(3):63–65, March.

- Su Nam Kim, Li Wang, and Timothy Baldwin. 2010. Tagging and linking web forum posts. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning*, pages 192–202, Uppsala, Sweden, 15–16 July.
- Akshi Kumar and Nazia Ahmad. 2012. Comex miner: Expert mining in virtual communities. *International Journal of Advanced Computer Science & Applications*, 3(6):54–65.
- Darcy Lemons and Lauren Trees. 2013. Expertise location in the social media age. In *Presentations from APQC's 2013 Knowledge Management Conference*, Houston, U.S.A. APQC. <http://www.apqc.org/knowledge-base/documents/expertise-location-social-media-age>. Accessed on 29/07/2014.
- Xiaomo Liu, G. Alan Wang, Aditya Johri, Mi Zhou, and Weiguo Fan. 2012. Harnessing global expertise: A comparative study of expertise profiling methods for online communities. *Information Systems Frontiers*, pages 1–13, September.
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press, New York, U.S.A.
- David W. McDonald and Mark S. Ackerman. 1998. Just talk to me: A field study of expertise location. In *Proceedings of the 1998 ACM conference on Computer supported cooperative work, CSCW*, pages 315–324, Seattle, U.S.A.
- Harald A. Mieg. 2001. *The Social Psychology of Expertise: Case Studies in Research, Professional Domains, and Expert Roles*. Expertise Series. Lawrence Erlbaum Associates.
- Harald A. Mieg. 2006. Social and sociological factors in the development of expertise. *The Cambridge Handbook of Expertise and Expert Performance*, (Chapter 41):743–760.
- Jason Rennie. 2008. 20 Newsgroups. <http://people.csail.mit.edu/jrennie/20Newsgroups/>. Accessed 26/06/2009.
- Debbie Richards, Meredith Taylor, and Peter Busch. 2008. Expertise recommendation: A two-way knowledge communication channel. In *Fourth International Conference on Autonomic and Autonomous Systems, 2008, ICAS*, pages 35–40.
- Gary Rolfe. 1997. Beyond expertise: theory, practice and the reflexive practitioner. *Journal of Clinical Nursing*, 6:93–97, March.
- Donald A. Schön. 1983. *The reflective practitioner: how professionals think in action*. Basic Books, New York, U.S.A.
- Michael F. Schwartz and David C. M. Wood. 1993. Discovering shared interests using graph analysis. *Communications of the ACM*, 36(8):78–89.
- John Searle. 1969. *Speech Acts*. Cambridge University Press.
- John Searle. 1975. A taxonomy of illocutionary acts. In K. Günderson, editor, *Language, Mind and Knowledge*, pages 344–369. University of Minnesota Press.
- Ian Soboroff, Arjen P. de Vries, and Nick Craswell. 2006. Overview of the TREC 2006 enterprise track. In E. M. Voorhees and Lori P. Buckland, editors, *The Fifteenth Text REtrieval Conference (TREC 2006) Proceedings*, pages 32–51. National Institute of Standards and Technology (NIST).
- Edda Weigand. 2010. Language as dialogue. *Intercultural Pragmatics*, 7(3):505–515, August.
- Xin Yan, Dawei Song, and Xue Li. 2006. Concept-based document readability in domain specific information retrieval. In *Proceedings of the 15th ACM international conference on Information and knowledge management, CIKM*, pages 540 – 549, Arlington, U.S.A.
- Dawit Yimam-Seid and Alfred Kobsa. 2003. Expert-finding systems for organizations: Problem and domain analysis and the DEMOIR approach. *Journal of Organizational Computing and Electronic Commerce*, 13(1):1–24.