Distributed Representations of Expertise

Fangqiu Han* Shulong Tan* Huan Sun* Mudhakar Srivatsa† Deng Cai‡
Xifeng Yan*

Abstract

Collaborative networks are common in real life, where domain experts work together to solve tasks issued by customers. How to model the proficiency of experts is critical for us to understand and optimize collaborative networks. Traditional expertise models, such as topic model based methods, cannot capture two aspects of human expertise simultaneously: Specialization (what area an expert is good at?) and Proficiency Level (to what degree?). In this paper, we propose new models to overcome this problem. We embed all historical task data in a lower dimension space and learn vector representations of expertise based on both solved and unsolved tasks.

Specifically, in our first model, we assume that each expert will only handle tasks whose difficulty level just matches his/her proficiency level, while experts in the second model accept tasks whose levels are equal to or lower than his/her proficiency level. Experiments on real world datasets show that both models outperform topic model based approaches and standard classifiers such as logistic regression and support vector machine in terms of prediction accuracy. The learnt vector representations can be used to compare expertise in a large organization and optimize expert allocation.

1 Introduction

Collaborative platforms, such as crowdsourcing service providers, community question answering forums, and customer service centers, are becoming more and more prevalent. Once managed effectively, the rich online human resources have shown great potential to solve problems more economically, efficiently, and reliably [1–3]. In order to effectively manage and utilize expert resources, an essential problem is how to correctly understand/represent human expertise and identify right experts for a certain task [4, 5]. In this paper, we take collaborative networks as an example to derive expertise representation so that multiple experts can be compared in the same framework.

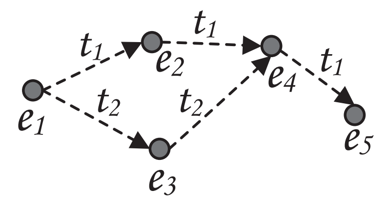

Figure 1: A sample collaborative network. Tasks are routed among experts in a collaborative network until they are resolved.

In collaborative networks, tasks are routed among a network of experts until they are resolved. Fig. 1 shows a sample collaborative network. Task t_1 starts at expert e_1 and is resolved by expert e_5; task t_2 starts at expert e_1 and is resolved by expert e_4. The sequences $e_1 \to e_2 \to e_4 \to e_5$ and $e_1 \to e_3 \to e_4$ are called *routing sequences* of task t_1 and t_2 respectively. One fundamental problem in ticket routing is how to represent experts' knowledge and employ it to estimate the probability of solving a task. Once this problem is solved, the final resolver to a given task can be found quickly.

An expert has to meet two constraints in order to solve a task: (1) Topic Match: the specialized areas of the expert shall match the topic of the task. For example, a programmer can possibly solve a programming problem while he is less likely to solve a physics problem; (2) Difficulty Level Match: the difficulty level of the task should match the proficiency of the expert. Previous studies [6, 7] assumed that a list of possible specialized areas for each expert is given. In real collaborative networks, manually creating these specialized areas is laborious and hardly accurate. An intuitive solution is to use topic modeling to automatically learn human expertise from the previously solved tasks. This solution has two main problems: (1) Tasks that an expert has failed to solve cannot be properly modeled to specify what the expert cannot do. (2) Topic models

*University of California, Santa Barbara. {fhan, shulongtan, huansun, xyan@cs.ucsb.edu.}
†IBM T.J. Watson Research Center. msrivats@us.ibm.com.
‡Zhejiang University. dengcai@gmail.com.

essentially capture the topic distribution of historical tasks. They do not directly measure proficiency level and its difference among experts.

To overcome the aforementioned issues and learn better expertise representations, we propose two expertise models in below:

(Model A) It assumes each expert has one or several specialized functional areas in a collaborative network. A task falling to one of the functional areas will be solved by the expert; otherwise it will be transferred to another expert. Based on this assumption, we define an expertise space in which all experts' expertise and all tasks will be embedded as numerical vectors. Tasks close to one of the expertise of an expert will be resolved by the expert whereas those far from his/her expertise will not. In this model, we combine the two aspects, specialized area and proficiency level of human expertise. The specialized area of an expert is characterized as a ball centered at his expertise vector and the radius of the ball signifies the range of the expert's duty. The ball is named *functional area* of the expert. This model is referred to as *Functional Area Expertise* (FAE).

(Model B) In some collaborative networks, there is no clear division of experts' responsibility. Experts solve tasks just based on their true capability. In this case, experts could deal with tasks in all difficulty levels below his capacity of solving tasks. Therefore, instead of unifying specialized areas and proficiency levels as in the FAE, our second model learns a vector representation of expertise and characterizes the two aspects separately: dimensions of the expertise vector encode specialized areas and the value in each dimension signifies the proficiency level of the expert on the corresponding area. Our intuitions in this formulation are as follows: (1) If an expert can solve a task, his proficiency level should be greater than or equal to the task difficulty; (2) If an expert cannot solve a task, there must be some dimensions in his expertise where their values are smaller than those required by the task. In this way, the specialized areas together with their proficiency levels can be modeled naturally. We refer to this model as *All-Round Expertise* (ARE).

FAE and ARE represent two different strategies of assigning task to experts. FAE is going to reserve the capacity of highly skilled experts for difficult tasks, while ARE tries to shorten task processing time as much as possible. Experimental results on real collaborative networks show that expertise learnt from our models better predict ticket solving than topic model based approaches and other methods in expertise modeling.

2 Conclusion

Expertise modeling is considered as the core content in improving the efficiency of collaborative networks. The goal is to represent experts' abilities in terms of specialization and proficiency level. In this paper, we developed two models to learn distributed representations of expertise which can convey both aspects. The first model assumes that an expert only solves tasks matching his/her specialization and proficiency level. Alternatively, in the second model, the expert can solve all the tasks whose difficulty level is equal to or lower than his/her proficiency level in his/her specialized area. Experimental results on real datasets demonstrated that the proposed models can learn meaningful expertise representations and are effective in predicting task resolution.

3 Acknowledgement

This research was sponsored in part by the Army Research Laboratory under cooperative agreements $W911NF-09-2-0053$. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notice herein. Deng Cai is supported by the National Basic Research Program of China under Grant $2013CB336500$.

References

[1] P. Welinder, S. Branson, P. Perona, and S. J. Belongie, "The multidimensional wisdom of crowds," in *NIPS*, 2010, pp. 2424–2432.

[2] D. R. Karger, S. Oh, and D. Shah, "Efficient crowdsourcing for multi-class labeling," in *ACM SIGMETRICS*, vol. 41, no. 1, 2013, pp. 81–92.

[3] S. Jagabathula, L. Subramanian, and A. Venkataraman, "Reputation-based worker filtering in crowdsourcing," in *NIPS*, 2014, pp. 2492–2500.

[4] I. Guy, U. Avraham, D. Carmel, S. Ur, M. Jacovi, and I. Ronen, "Mining expertise and interests from social media," in *WWW*, 2013, pp. 515–526.

[5] Y. Zhou, G. Cong, B. Cui, C. S. Jensen, and J. Yao, "Routing questions to the right users in online communities," in *ICDE*, 2009, pp. 700–711.

[6] K. Balog and M. De Rijke, "Determining expert profiles (with an application to expert finding)," in *Proc. of IJCAI-07*, 2007, pp. 2657–2662.

[7] K. Balog, Y. Fang, M. de Rijke, P. Serdyukov, and L. Si, "Expertise retrieval," *Foundations and Trends in Information Retrieval*, vol. 6, no. 2–3, pp. 127–256, 2012.