Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

7-2020

Keen2Act: Activity recommendation in online social collaborative platforms

Roy Ka-Wei LEE
University of Saskatchewan

Thong HOANG
Singapore Management University, vdthoang.2016@phdis.smu.edu.sg

Richard J. OENTARYO
McLaren Applied

David LO
Singapore Management University, davidlo@smu.edu.sg

Citation

Keen2Act: Activity Recommendation in Online Social Collaborative Platforms

Roy Ka-Wei Lee
University of
Saskatchewan, Canada
roylee@cs.usask.ca

Thong Hoang
Singapore Management
University, Singapore
vdthoang.2016@smu.edu.sg

Richard J. Oentaryo
McLaren Applied,
Singapore
richard.oentaryo@mclaren.com

David Lo
Singapore Management
University, Singapore
davidlo@smu.edu.sg

Abstract

Social collaborative platforms such as GitHub and Stack Overflow have been increasingly used to improve work productivity via collaborative efforts. To improve user experiences in these platforms, it is desirable to have a recommender system that can suggest not only items (e.g., a GitHub repository) to a user, but also activities to be performed on the suggested items (e.g., forking a repository). To this end, we propose a new approach dubbed Keen2Act, which decomposes the recommendation problem into two stages: the *Keen* and *Act* steps. The *Keen* step identifies, for a given user, a (sub)set of items in which he/she is likely to be interested. The *Act* step then recommends to the user which activities to perform on the identified set of items. This decomposition provides a practical approach to tackling complex activity recommendation tasks while producing higher recommendation quality. We evaluate our proposed approach using two real-world datasets and obtain promising results whereby Keen2Act outperforms several baseline models.

CCS Concepts

• **Information systems → Learning to rank**; **Rank aggregation**; • **Computing methodologies → Factorization methods**.

Keywords

Activity Recommendation, Factorization Machine, Social Collaborative Platform, Stack Overflow, GitHub

1 Introduction

Users are increasingly adopting social collaborative platforms for collaborative activities. GitHub and Stack Overflow are two such popular platforms; GitHub is a collaborative software development platform that allows code sharing and version control, while Stack Overflow is a technical question-and-answer community-based website. As these social collaborative platforms gain popularity,

many research studies have proposed recommender systems to improve the usability of these platforms. For example, there are works that recommend Stack Overflow questions for users to answer [16, 17]. Similarly, in GitHub, researchers have proposed methods to recommend relevant repositories to a user [5, 7].

Many of the existing works, however, focus largely on recommending either items or a single type of activity to users, ignoring the fact that the users can perform multiple types of activity on these platforms. For example, GitHub users may *fork* or *watch* repositories and Stack Overflow users may *answer* or *favorite* questions. Recommending multiple types of activities to a user is a challenging task. A naïve solution would be to recommend individual activities separately, treating them as independent tasks. However, there might be insufficient observations for learning the user activities at such granularity. Another possible solution is to simply learn the different types of activities together, but such solution creates sparsity issue, having to learn for all possible item-activity pairs.

To tackle these challenges, we propose a new recommendation approach called Keen2Act[1], which learns the users' item-level and activity-level interests in a step-wise manner to achieve (joint) item-activity recommendation. In particular, the main contribution of Keen2Act is that it features a novel two-stage process that decomposes the activity recommendation problem into predicting a user's interests at item level and subsequently predicting at activity level. To the best of our knowledge, this is also the first work that achieves multi-typed activity recommendation in social collaborative platforms. Finally, empirical studies on real-world GitHub and Stack Overflow datasets have shown promising results whereby the proposed approach outperforms several baseline methods.

2 Related Work

Research studies on prediction and recommendation in social collaborative platforms broadly fall into two categories: (i) finding experts to perform certain platform tasks [1, 2, 6, 10, 13, 15, 19–21] and (ii) recommending items to users in a platform [3, 5, 7, 16, 17, 22]. Under category (i), the works on Stack Overflow mainly involve devising methods to find experts to answer questions [2, 13, 15, 19, 20], while for GitHub a user is identified as an expert if (s)he reviews pull-requests and code for repositories [10, 21]. The works under category (ii) largely focus on recommending items to users, without specifying activities to be performed on the items. For example, several works aim to recommend relevant Stack Overflow posts [3, 17] and Github repositories [7, 22] to users.

Our work deviates from the existing works under category (ii) in several ways. Firstly, to the best of our knowledge, Keen2Act constitutes the first work that focuses on recommending not only

---

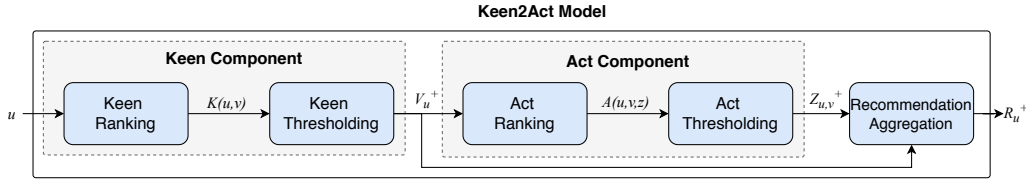[1]Source code: https://gitlab.com/bottle_shop/scp/keen2act

Figure 1: Keen2Act Framework

items but also specific activities under each recommended item. Our approach has also been applied to more than one platform (i.e., GitHub and Stack Overflow).

3 Proposed Approach

3.1 Problem Formulation

We define the activity recommendation problem as follows: *For a given user, which items should (s)he choose and what activities should (s)he perform on those items?* The proposed Keen2Act approach addresses this joint item-activity recommendation problem by breaking it down into two sub-problems: the *Keen* and *Act* tasks. The former aims to identify the set of items that a user is potentially interested in, while the latter aims to subsequently determine the set of activities to perform on the items of interest. An outline of our proposed Keen2Act model is given in Figure 1.

We first denote a particular user, item, and activity using the notation u, v and z, respectively. We also let U, V and Z denote the set of all users, all items, and all activities, respectively. For a given user u, the Keen component determines whether an item v should be included in the set V_u^+ of items selected by that user, as follows:

$$V_u^+ = \{v \in V : \mathcal{K}(u,v) \geq \delta_{\mathcal{K}}(v)\} \tag{1}$$

where $\mathcal{K}(u,v)$ is the Keen ranking score for user-item pair (u,v), and $\delta_{\mathcal{K}}(v)$ is the Keen decision threshold for v. Subsequently, the Act component determines whether an activity z should be part of the set $Z_{u,v}^+$ of activities performed by user u on a selected item v:

$$Z_{u,v}^+ = \left\{z \in Z : \mathcal{A}(u,v,z) \geq \delta_{\mathcal{A}}(z) \land v \in V_u^+\right\} \tag{2}$$

where $\mathcal{A}(u,v,z)$ is the Act ranking score for the tuple (u,v,z), and $\delta_{\mathcal{A}}(z)$ is the Act decision threshold for z.

The main intuition behind the Keen followed by Act steps is that, prior to determining an activity z on item v, user u must be sufficiently keen in item v in the first place. When there is a lack of keenness (i.e., $\mathcal{K}(u,v) < \delta_{\mathcal{K}}(v)$), the user should not perform any activity at all on item v. Conversely, only when the user shows a sufficient level of keenness (i.e., $\mathcal{K}(u,v) \geq \delta_{\mathcal{K}}(v)$), he/she can proceed with selecting which activities to be performed on item v. This two-stage decision process helps not only reduce the search space (by filtering out less relevant items) but also improve the quality of the final item-activity recommendation.

The ranking scores $\mathcal{K}(u,v)$ and $\mathcal{A}(u,v,z)$ can each be realized using any machine learning model. In this work, we choose to use a multilinear model called *Factorization Machine* (FM) [11], which has shown good performance in a variety of recommendation tasks based on sparse data. It is also worth noting that the decision thresholds $\delta_{\mathcal{K}}(v)$ and $\delta_{\mathcal{A}}(z)$ are parameters that are learnable from data and are specific to each item v and activity z, respectively.

Finally, a *Recommendation Aggregation* process takes place to combine V_u^+ and $Z_{u,v}^+$ in order to arrive at the final list R_u^+ of recommended item-activity pairs. More specifically, for a given user-item-activity triplet (u,v,z), the Recommendation Aggregation process corresponds to a decision function $\mathcal{D}(u,v,z)$:

$$\mathcal{D}(u,v,z) = \mathbb{I}\left[v \in V_u^+ \land z \in Z_{u,v}^+\right] \tag{3}$$

where $\mathbb{I}[.]$ is an indicator function and $\land$ is an AND logical operator.

We further elaborate the formulation of learning mechanisms behind the Keen and Act components in Sections 3.2 and 3.3, respectively. We then describe the Recommendation Aggregation process in greater detail in Section 3.4. Finally, we recap the overall Keen2Act learning procedure in Section 3.5.

3.2 Keen Component

Keen ranking. We consider the problem of identifying the parameters $\theta_{\mathcal{K}}$ of the Keen model $\mathcal{K}(u,v)$ as a learning-to-rank task, and it is thus sensible to use a loss function optimized for ranking. Among the myriad of ranking loss functions, the *Weighted Approximately Ranked Pairwise* (WARP) loss [18] in particular has been shown as a good criterion for recommendation tasks. An appealing trait is that WARP works for data that have only implicit feedback, making it well-suited to our problem. The key idea of WARP is that, for a given positive (observed) example, the remaining negative (unobserved) examples are randomly sampled until we find a pair of positive and negative examples for which the current model incorrectly ranks the negative example higher than the positive one. We can then perform a model update only based on these violating examples.

Adopting WARP to the context of our Keen model, we can write the loss function with respect to the Keen model parameters $\theta_{\mathcal{K}}$ as:

$$\mathcal{L}_{\text{WARP}}(\theta_{\mathcal{K}}) = \sum_{u \in U} \sum_{v \in V_u^+} \Phi\left(rank_v(\mathcal{K}(u,v))\right) \tag{4}$$

where $\Phi(.)$ transforms the rank of a positive item $v \in V_u^+$ into a weight. Here the $rank_v$ function can be defined as a margin-penalized rank of the following form:

$$rank_v(\mathcal{K}(u,v)) = \sum_{v' \notin V_u^+} \mathbb{I}\left(\mathcal{K}(u,v) \geq 1 + \mathcal{K}(u,v')\right) \tag{5}$$

Choosing a transformation function such as $\Phi(n) = \sum_{i=1}^{n} \frac{1}{i}$ would allow the model to optimize for a better Precision@k. However, directly minimizing such a loss function via gradient-based algorithms would be computationally intractable, as equation (5) sums over all items, resulting in a slow per-gradient update. We can circumvent this issue by replacing $rank_v$ with the following sampled approximation [18]: we sample N items v' until we find a violation, i.e, $\mathcal{K}(u,v) < 1 + \mathcal{K}(u,v')$, and subsequently estimate the rank, i.e., $rank_v(\mathcal{K}(u,v))$, by $\frac{|V \setminus V_u^+| - 1}{N}$. In addition to the WARP loss,

we need an appropriate regularization term to control our model complexity. In this work, we employ L2 regularization, which is differentiable and suitable for gradient-based methods as well. This leads to an overall, regularized ranking loss function $\mathcal{L}_{\text{rank}}$:

$$\mathcal{L}_{\text{rank}}(\theta_{\mathcal{K}}) = \mathcal{L}_{\text{WARP}}(\theta_{\mathcal{K}}) + \frac{\lambda_{\mathcal{K}}}{2}||\theta_{\mathcal{K}}||^2 \qquad (6)$$

where $\lambda_{\mathcal{K}} > 0$ is a user-specified l_2-regularization parameter. The term $||\theta_{\mathcal{K}}||^2$ is used to mitigate data overfitting by penalizing large parameter values, thus reducing the model complexity.

Keen thresholding. Once the Keen ranking step is done, we need to determine the appropriate decision thresholds $\delta_{\mathcal{K}}$ in order to decide whether to include item v into V_u^+. Ideally, we wish to identify a threshold such that the set of selected items matches the set of ground-truth, observed items as close as possible. However, it is difficult to learn the thresholds using such a set matching objective. We therefore relax this via the *cross-entropy* loss function, and for the Keen model this would be:

$$\mathcal{L}_{\text{thres}}(\delta_{\mathcal{K}}) = \sum_{u \in U} \sum_{v \in V} CE\left(\mathcal{K}(u,v) - \delta_{\mathcal{K}}(v), I(v \in V_u^+)\right) \quad (7)$$

where $CE(x,y) = -\left[y\ln(\sigma(x)) + (1-y)\ln(1-\sigma(x))\right]$ is the cross entropy function, $\sigma(x) = \frac{1}{1+\exp(-x)}$ is the sigmoid function, and $I(v \in V_u^+)$ is an indicator function reflecting the ground truth for whether item v belongs to the set of items selected by user u.

3.3 Act Component

Act ranking. The ranking loss formulation for the Act model is similar to that of the Keen model except that the former deals with ranking of activities at the level of user-item pair. In particular, the WARP loss function associated with the Act model is given by:

$$\mathcal{L}_{\text{WARP}}(\theta_{\mathcal{A}}) = \sum_{u \in U} \sum_{v \in V_u^+} \sum_{z \in Z_{u,v}^+} \Phi\left(rank_z(\mathcal{A}(u,v,z))\right) \qquad (8)$$

where the rank function is likewise estimated by a sampled approximation $rank_z(\mathcal{A}(u,v,z)) \approx \frac{|Z \setminus Z_{u,v}^+|-1}{N}$. Accordingly, adding the L2 penalty to control the model complexity, the overall regularized WARP loss for the Act model $\mathcal{A}$ is given by:

$$\mathcal{L}_{\text{rank}}(\theta_{\mathcal{A}}) = \mathcal{L}_{\text{WARP}}(\theta_{\mathcal{A}}) + \frac{\lambda_{\mathcal{A}}}{2}||\theta_{\mathcal{A}}||^2 \qquad (9)$$

where $\lambda_{\mathcal{A}} > 0$ is the L2-regularization parameter.

Act thresholding. Similar to the thresholding in the Keen model, we can estimate the decision threshold of the Act model using the following cross-entropy loss:

$$\mathcal{L}_{\text{thres}}(\delta_{\mathcal{A}}) = \sum_{u \in U} \sum_{v \in V_u^+} \sum_{z \in Z} CE\left(\mathcal{A}(u,v,z) - \delta_{\mathcal{A}}(z), I(z \in Z_{u,v}^+)\right) \quad (10)$$

where $I(z \in Z_{u,v}^+)$ is an indicator function reflecting the ground truth for whether activity z belongs to the set of activities carried out by user u on a selected item v (i.e., $v \in V_u^+$).

3.4 Recommendation Aggregation

The final step in the Keen2Act model is to generate a sorted list R_u^+ of recommended item-activity pairs for a given user u, by aggregating the selected item set V_u^+ and chosen activity set $Z_{u,v}^+$ as computed by the Keen and Act models respectively. To achieve this, we generate

---

Algorithm 1 Keen2Act Learning Procedure

Inputs:
    Keen interactions $\mathcal{I}_{\mathcal{K}}$ and Act interactions $\mathcal{I}_{\mathcal{A}}$
    Maximum number of epochs T and maximum negative samples N
Outputs:
    Model parameters $\theta_{\mathcal{K}}$ and $\theta_{\mathcal{A}}$
    Decision thresholds $\delta_{\mathcal{K}}$ and $\delta_{\mathcal{A}}$

```
 1: repeat                                          ▷ Phase 1: Rank learning
 2:     for each (u, v) ∈ I_K do                         ▷ Keen ranking loop
 3:         repeat
 4:             Randomly sample a negative item v' by from V \ V_u^+
 5:             if K(u, v) < 1 + K(u, v') then      ▷ Keen rank violation found
 6:                 Perform Adam update on θ_K to minimise (6)
 7:                 break
 8:             end if
 9:         until maximum sampling N
10:     end for
11:     for each (u, v, z) ∈ I_A do                     ▷ Act ranking loop
12:         repeat
13:             Randomly sample a negative activity z' by from Z \ Z_{u,v}^+
14:             if A(u, v, z) < 1 + A(u, v, z') then  ▷ Act rank violation found
15:                 Perform Adam update on θ_A to minimise (9)
16:                 break
17:             end if
18:         until maximum sampling N
19:     end for
20: until maximum epochs T
21: repeat                                        ▷ Phase 2: Threshold learning
22:     for each user u ∈ U do
23:         for all items v ∈ V do                   ▷ Keen thresholding loop
24:             Perform Adam update on δ_K to minimise (7)
25:         end for
26:         for all positive items v ∈ V_u^+ do       ▷ Act thresholding loop
27:             for all activities z ∈ Z do
28:                 Perform Adam update on δ_A to minimise (10)
29:             end for
30:         end for
31:     end for
32: until maximum epochs T
```

R_u^+ by enumerating the selected items $v \in V_u^+$ starting from the one with the highest $\mathcal{K}(u,v)$, and then for each selected item, we enumerate the selected activities $z \in Z_{u,v}^+$ starting from the highest $\mathcal{A}(u,v,z)$. This results in a recommendation list whereby the item ranking takes precedence over the activity ranking.

3.5 Learning Procedure

To minimize the ranking losses $\mathcal{L}_{\text{rank}}(\theta_{\mathcal{K}})$ and $\mathcal{L}_{\text{rank}}(\theta_{\mathcal{A}})$ as well as threshold losses $\mathcal{L}_{\text{thres}}(\delta_{\mathcal{K}})$ and $\mathcal{L}_{\text{thres}}(\delta_{\mathcal{A}})$, we devise an incremental learning procedure based on a variant of stochastic gradient descent called *Adaptive Moment Estimation* (Adam) [8]. Algorithm 1 summarizes the procedure, which takes in the Keen interactions $\mathcal{I}_{\mathcal{K}} = \{(u,v) : u \in U \wedge v \in V_u^+\}$ and Act interactions $\mathcal{I}_{\mathcal{A}} = \{(u,v,z) : u \in U \wedge v \in V_u^+ \wedge z \in Z_{u,v}^+\}$ as inputs, and outputs the model parameters $\theta_{\mathcal{K}}$ and $\theta_{\mathcal{A}}$ as well as thresholds $\delta_{\mathcal{K}}$ and $\delta_{\mathcal{A}}$.

The overall algorithm consists of two phases: *rank learning* and *threshold learning*. In the first phase, we carry out Adam to update the model parameters $\theta_{\mathcal{K}}$ and $\theta_{\mathcal{A}}$ by finding a pair of positive and negative examples that violate the desired (Keen or Act) ranking. In the second phase, we also perform Adam update on $\delta_{\mathcal{K}}$ and $\delta_{\mathcal{A}}$ by enumerating at the item and activity levels respectively.

4 Experiment

Datasets. We experiment on two public datasets: GitHub [4] and Stack Overflow[2]. For both datasets, we retrieve active users (i.e.,

---

[2]https://archive.org/details/stackexchange

Table 1: Dataset summary

GitHub		Stack Overflow	
#Users	33,453	#Users	23,612
#Repositories	461,931	#Questions	1,020,809
#Fork Activities	445,084	#Answer Activities	860,302
#Watch Activities	1,730,181	#Favorite Activities	544,617

users who have performed at least 10 activities) and their activities performed between October 2013 to March 2015. For GitHub dataset, we obtain the *fork* and *watch* activities performed by 33,453 users on over 400k repositories. For Stack Overflow dataset, we retrieve the *answer* and *favorite* activities performed by 23,612 users on over 1 million questions. Table 1 summarizes the two datasets.

Features. From the datasets, we construct two sets of features:

- *User features*: We adapt the *co-participation similarity scores* introduced in [9] to measure the level of co-participation between a given user and other users in the platform. For example, we compute the number of times two users *co-fork* and *co-watch* a GitHub repository. As such, each user is represented with a count vector with d_U dimensions, where d_U is equal to the number of users in the platform.
- *Item features*: For each item (i.e., Stack Overflow question and GitHub repository), we compute its TF-IDF vector based on its description tags. Hence, each item is represented with a TD-IDF vector with d_T dimensions, where d_T is equal to the total number of tags used to describe the items.

Baselines. As a few related works perform activity recommendations, we adapt and apply some of the commonly used item recommendation methods to our scenario. Specifically, we compare our model to two variants of Factorization Machine (FM):

- **FM_BPR** [11]. This method uses *Bayesian Personalized Ranking* (BPR) loss [12] to maximize the rank difference between a positive example and a randomly chosen negative example.
- **FM_WARP** [18]. This method uses WARP loss to maximize the rank of positive examples by repeatedly sampling negative examples until a rank violation is found.

The user by item-activity interaction matrix, user, and item features are used as input for both baseline methods. All the parameters of baseline methods are empirically set to the optimal values. Besides the baselines, we also test several variants of our model:

- **Keen Model**. Using only the Keen model, we retrieve a set of items that the user is interested in and recommend the user to perform all activities on the retrieved items.
- **Act Model**. Using only the Act model, we consider all possible activities for all possible user-item pair and recommend activities which meet the Act threshold.
- **Keen2Act**. Our full model with both Keen and Act modules.

Training and testing splits. In all our experiments, we randomly select 80% activities of each user to form the training set and use the remaining activities as testing set. As such, all users are observed in the training set. However, some items might be new in the test set (corresponding to a cold start problem). Note that the user and item features are computed based on the observations in the training set. We repeat this process 5 times, resulting in 5 training-testing splits based on which we evaluate our model.

Evaluation metric. We use the Mean Average Precision at top k (MAP@k) as a primary metric in our experiments, which is popularly used to evaluate recommendation models [14]. We vary k

Table 2: Experiment results of various methods

Dataset	Model	MAP@5	MAP@10	MAP@20	MAP@50	MAP
GitHub	FM_BPR	0.120	0.128	0.127	0.113	0.036
	FM_WARP	0.300	0.299	0.276	0.233	0.077
	Keen Model	0.298	0.297	0.269	0.224	0.074
	Act Model	0.250	0.243	0.232	0.219	0.058
	Keen2Act	**0.348**	**0.347**	**0.325**	**0.284**	**0.099**
Stack Overflow	FM_BPR	0.103	0.109	0.108	0.100	0.033
	FM_WARP	0.180	0.183	0.177	0.160	0.050
	Keen Model	0.238	0.234	0.230	0.209	0.054
	Act Model	0.191	0.191	0.188	0.186	0.053
	Keen2Act	**0.259**	**0.249**	**0.227**	**0.210**	**0.064**

from 5 to ∞ in order to examine the sensitivity of our model. Note that setting $k = \infty$ is equivalent to computing MAP.

4.1 Results and Discussion

Table 2 shows the MAP@k results averaged over the 5 runs of our activity recommendation experiments. We observe that Keen2Act consistently outperforms all the other methods. Specifically, Keen2Act outperforms FM_BPR and FM_WARP by 175% and 29% respectively in Stack Overflow, and 94% and 28% respectively in GitHub. As we increase k, we also notice deterioration in MAP@k results. This can be attributed to the Precision@k metric favoring a model that outputs a ranked list with the relevant (i.e., observed) item-activity pairs leading the list. That is, as k increases, it is likely that more and more irrelevant item-activity pairs would appear in between the relevant item-activity pairs, pushing the Precision@k lower.

Additionally, we can see that the improvement of Keen2Act over the baselines is greater in Stack Overflow than in GitHub. A possible reason is due to the sparsity of user activities. More specifically, GitHub users are observed to perform more activities concentrated on a small set of items (i.e., denser user by item-activity interaction matrix), whereas Stack Overflow users tend to perform fewer activities spread across a large set of items. The denser interaction matrix for GitHub allows the baseline methods to have sufficient observations to achieve competitive results.

Comparing the different variants of our proposed model, we can see that Keen2Act outperforms the Keen model, suggesting that it is inadequate to learn only the item-level interests of the users when recommending activities. Keen2Act also outperforms the Act model, which demonstrates the importance of the Keen step in learning the item-level interests before activity-level interests. It is also interesting to see that the Keen model performs fairly well in comparison to the other methods. We can attribute this to reduced sparsity in the problem space it is operating at, i.e., the Keen model only makes item-level recommendation while the other methods recommend at the activity level. Finally, it is worth noting that the MAP@k scores of various models are generally low, showcasing the complexity of the activity recommendation problem.

5 Conclusion

In this paper, we put forward a new Keen2Act modeling approach to recommending items and multiple type of activity to a user in a step-wise manner. The efficacy of the approach has been demonstrated in experiments using real-world GitHub and Stack Overflow datasets. In the future, we would like to extend Keen2Act using deep representation learning and conduct more comprehensive experiments to benchmark it against more state-of-the-art methods. We also plan to apply Keen2Act to other social platforms as well as consider a larger number of activity types.

References

[1] Mohammad Y Allaho and Wang-Chien Lee. 2014. Increasing the responsiveness of recommended expert collaborators for online open projects. In *23rd ACM International Conference on Conference on Information and Knowledge Management*. 749–758.

[2] Morakot Choetkiertikul, Daniel Avery, Hoa Khanh Dam, Truyen Tran, and Aditya Ghose. 2015. Who will answer my question on Stack Overflow?. In *24th Australasian Software Engineering Conference*. 155–164.

[3] Lucas BL de Souza, Eduardo C Campos, and Marcelo de A Maia. 2014. Ranking crowd knowledge to assist software development. In *22nd International Conference on Program Comprehension*. 72–82.

[4] Georgios Gousios. 2013. The GHTorent dataset and tool suite. In *10th Working Conference on Mining Software Repositories*. 233–236.

[5] Mohamed Guendouz, Abdelmalek Amine, and Reda Mohamed Hamou. 2015. Recommending relevant GitHub repositories: A collaborative-filtering approach. *International Conference on Networking and Advanced Systems*.

[6] Chaoran Huang, Lina Yao, Xianzhi Wang, Boualem Benatallah, and Quan Z Sheng. 2017. Expert as a service: Software expert recommendation via knowledge domain embeddings in stack overflow. In *IEEE International Conference on Web Services*. 317–324.

[7] Jyun-Yu Jiang, Pu-Jen Cheng, and Wei Wang. 2017. Open source repository recommendation in social coding. In *40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1173–1176.

[8] Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.

[9] Roy Ka-Wei Lee and David Lo. 2017. GitHub and Stack Overflow: Analyzing developer interests across multiple social collaborative platforms. In *International Conference on Social Informatics*. 245–256.

[10] Mohammad Masudur Rahman, Chanchal K Roy, and Jason A Collins. 2016. COR-RECT: Code reviewer recommendation in GitHub based on cross-project and technology experience. In *38th International Conference on Software Engineering Companion*. 222–231.

[11] Steffen Rendle. 2012. Factorization machines with libFM. *ACM Transactions on Intelligent Systems and Technology* 3, 3 (2012), 1–22.

[12] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *25th Conference on Uncertainty in Artificial Intelligence*. 452–461.

[13] Fatemeh Riahi, Zainab Zolaktaf, Mahdi Shafiei, and Evangelos Milios. 2012. Finding expert users in community question answering. In *21st International Conference on World Wide Web*. 791–798.

[14] Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. 2009. Mining multi-label data. In *Data Mining and Knowledge Discovery Handbook*. Springer, 667–685.

[15] Jian Wang, Jiqing Sun, Hongfei Lin, Hualei Dong, and Shaowu Zhang. 2017. Convolutional neural networks for expert recommendation in community question answering. *Science China Information Sciences* 60, 11 (2017), 110102.

[16] Lin Wang, Bin Wu, Juan Yang, and Shuang Peng. 2016. Personalized recommendation for new questions in community question answering. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. 901–908.

[17] Wei Wang, Haroon Malik, and Michael W Godfrey. 2015. Recommending posts concerning API issues in developer Q&A sites. In *IEEE/ACM 12th Working Conference on Mining Software Repositories*. 224–234.

[18] Jason Weston, Samy Bengio, and Nicolas Usunier. 2011. WSABIE: Scaling up to large vocabulary image annotation. In *22nd International Joint Conference on Artificial Intelligence*.

[19] Congfu Xu, Xin Wang, and Yunhui Guo. 2016. Collaborative expert recommendation for community-based question answering. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. 378–393.

[20] Liu Yang, Minghui Qiu, Swapna Gottipati, Feida Zhu, Jing Jiang, Huiping Sun, and Zhong Chen. 2013. CQARank: Jointly model topics and expertise in community question answering. In *22nd ACM international conference on Information & Knowledge Management*. 99–108.

[21] Yue Yu, Huaimin Wang, Gang Yin, and Tao Wang. 2016. Reviewer recommendation for pull-requests in GitHub: What can we learn from code review and bug assignment? *Information and Software Technology* 74 (2016), 204–218.

[22] Lingxiao Zhang, Yanzhen Zou, Bing Xie, and Zixiao Zhu. 2014. Recommending relevant projects via user behaviour: an exploratory study on GitHub. In *1st International Workshop on Crowd-based Software Development Methods and Technologies*. 25–30.