

# Uncovering the Dynamics of Crowdlearning and the Value of Knowledge

Utkarsh Upadhyay, Isabel Valera, and Manuel Gomez-Rodriguez

Max Plank Institute for Software Systems,  
{utkarshu, ivalera, manuelgr}@mpi-sws.org

## Abstract

Learning from the crowd has become increasingly popular in the Web and social media. There is a wide variety of *crowdlearning* sites in which, on the one hand, users learn from the knowledge that other users contribute to the site, and, on the other hand, knowledge is reviewed and curated by the same users using assessment measures such as upvotes or likes. In this paper, we present a probabilistic modeling framework of crowdlearning, which uncovers the evolution of a user’s expertise over time by leveraging other users’ assessments of her contributions. The model allows for both off-site and on-site learning and captures forgetting of knowledge. We then develop a scalable estimation method to fit the model parameters from millions of recorded learning and contributing events. We show the effectiveness of our model by tracing activity of  $\sim 25$  thousand users in Stack Overflow over a 4.5 year period. We find that answers with high knowledge value are rare. Newbies and experts tend to acquire less knowledge than users in the middle range. Prolific learners tend to be also proficient contributors that post answers with high knowledge value.

**Keywords:** User modelling; Education; Markets and crowds; Social and information networks.

## 1 Introduction

*“By learning you will teach; by teaching you will learn.”*  
—Latin proverb

Question answering (Q&A) sites, online communities, wikis and microblogs offer unprecedented opportunities for people to learn about a wide variety of topics, acquire specialized knowledge or be up to date with latest breaking news. A common feature shared by most of these platforms is that knowledge is contributed by the crowd – it is *crowdsourced* – and it is also the crowd who reviews and curates the contributed knowledge. For example, in Q&A sites, users can learn by reading answers others post to their own or similar questions; in wikis, a set of editors write and review the content of pages in a collaborative fashion, and this content is then publicly accessible to others; in microblogs, users post small pieces of information, which can then be assessed by other users by means of likes, shares or replies. All of the above are examples of *crowdlearning*, in which users can play the role of a *learner*, a *contributor*, or switch between both. There have been many recent works on identifying experts (or potential experts) in Q&A websites [15, 36] and microblogs [13, 25], as well as modeling learning in controlled settings [6, 12]. However, general models of crowdlearning are largely inexistent to date. Such models are of outstanding interest since they would allow us:

- (i) to better understand how people learn over time and become experts;
- (ii) to identify questions with high knowledge value, which systematically help users increase their expertise;

(iii) to investigate the interplay between learners and contributors.

In this paper, we propose a probabilistic generative model of crowdlearning, especially designed to fit fine-grained crowdlearning event data [1]. The key idea behind our modeling framework is simple: every time a user learns from a *knowledge item* contributed by other users, she may increase her expertise and, as a consequence, her subsequent contributions be more knowledgeable and assessed more highly by others in terms of, *e.g.*, upvotes, likes or shares. Thus, by jointly modeling *learning events*, in which users acquire *effective knowledge*, and *contributing events* (in short, *contributions*), in which users contribute with their expertise to a knowledge item, our framework will reach the above mentioned goals. In this work, we aim to measure those aspects of the learning process for which we have evidence in the observed data, *i.e.*, a measure of *effective knowledge* that leads to measurable increase in users' *effective learning*. A general model of abstract knowledge and learning remains a challenging endeavor.

In more detail, we model each user's expertise as a latent stochastic process that evolves over time and think of the other users' assessment of her contributions as noisy samples from this stochastic process localized in time. Moreover, this stochastic process is driven by two types of learning: off-site learning and on-site learning. The proposed formulation also captures characteristic properties of the learning process, previously studied in the literature, such as forgetting [21] and initial expertise [30]. We then develop an efficient parameter estimation method that finds the model parameters that maximize the likelihood of an observed set of learning and contributing events via convex optimization. Finally, we show the effectiveness of our model by tracing learning and contributing events in data gathered from Stack Overflow over a 4.5 year period. Our experiments reveal several interesting insights:

- I. The knowledge value of items follow a log-normal distribution.
- II. Users with very low or very high initial expertise, *i.e.*, newbies and experts, tend to increase their knowledge the least and, in contrast, users in the middle range tend to increase it the most. This suggests that the learning curve may be sigmoidal, in agreement with existing literature [20].
- III. Although there are fewer contributors than learners in absolute numbers, the distribution of knowledge in the contributions is fat tailed while the distribution of knowledge learned is heavy tailed.
- IV. Users who learn from high knowledge items are also more proficient at providing answers with high knowledge value.

**Related work.** Our work lies in the intersection between expert identification, learning and knowledge tracing and student modeling.

Identifying topical authorities or experts, *i.e.*, users who provide high quality contributions, on Q&A [15, 17, 26, 36] and microblogging sites [13, 25], has received a lot of attention recently. The problem of expert finding in Q&A sites was first studied by Zhang et al. [36], who formulated it as a ranking problem and developed several PageRank based methods. Shortly after, Jurczyk and Agichtein [17] tackled the problem using link analysis techniques. More recently, Pal and Konstan [26] approached the problem from the perspective of supervised learning and developed Gaussian classification models to distinguish between ordinary and (potential) experts users, and Hanrahan et al. [15] described a method to find experts given a specific target question. In the context of microblogging, the problem of expert finding was first studied by Pal and Counts [25], who proposed a set of features for characterizing contributors and then formulated the problem using unsupervised learning in this feature space. Since then, Ghosh et al. [13] mined Twitter users' lists to find topical authorities and Kao et al. [19] and Paulina et al. [27] leveraged temporal statistics on the users' activity to identify experts. Finally, in the context of web search, White et al. [34] studied how expertise influence search and Eickhoff et al. [9] investigated how a user can increase her expertise as she looks for *procedural* and *declarative* knowledge using a search engine. However, in contrast to our work, previous work on expert identification did not capture the evolution of users' expertise over extended periods of time nor accounted for the knowledge value of their contributions.

The interest in the field of modeling and measuring learning is very old and several paradigms have been developed over the last century in the experimental psychology literature [10, 22, 24]. Most of the research has, however, either happened in strictly controlled environments (*i.e.*, schools or study groups) or used centralized assessments (*e.g.*, SAT or local testing). The work most closely related to ours is knowledge tracing and student modeling, which has been carried on by researchers from the learning analytics, educational data mining and intelligent tutoring communities. In this line of work, several probabilistic models have been proposed: Bayesian knowledge tracing [8, 14, 31, 35], performance factor analysis [28] and ensembles [5]. However, previous work typically relies on controlled assessment and manually annotated knowledge items, even if allowing for different knowledge values per item [6]. Only very recently, Piech et al. [29] solved this limitation by resorting to recurrent neural networks to model the learning of students, unfortunately, they use metrics that are not suitable for crowdlearning.

In summary, our goal here is a general understanding of crowdlearning dynamics, from uncovering the evolution of users' expertise over time and understanding the interplay between learning and contributing, to identifying questions with a high knowledge value, which systematically help users to increase their expertise. In contrast, previous work has focused either on identifying experts, not their expertise evolution, or modeling learning of students in controlled environment with manually annotated knowledge items.

## 2 A Crowdlearning Model

In a crowdlearning site, users often play two different functional roles: *contributors* and *learners*. In the former role, they share their knowledge on a topic (or topics) with other users within the site; and, in the latter role, they acquire knowledge by reading what other users contributed to the site. Then, we can think of users' expertise as latent stochastic processes that evolve over time, and think of the assessments of their contributions to the site as noisy samples from these stochastic processes localized in time. Here, we propose a modeling framework that uncovers the evolution of these processes by modeling two types of learning:

- I. *Off-site learning*, which accounts for the knowledge that the user accumulates outside the site; and,
- II. *On-site learning*, which accounts for the knowledge that the user gains by reading other users' contributions within the site.

Next, we formulate our generative model, starting from the data it is design for.

**Crowdlearning data.** Given a crowdlearning site with a set of users  $\mathcal{U}$  and a set of learning areas (or topics)  $\mathcal{A}$ , we first define a knowledge item  $q$  as the smallest quantum of knowledge a user can learn from within the site. For example, in a Q&A site, a knowledge item corresponds to a question and its answer(s); in Twitter, it corresponds to a tweet; and in a wiki site, it corresponds to a wiki page. Intuitively, each knowledge item  $q$  provides certain (latent) knowledge value,  $k_q \in \mathbb{R}^+$ , and contains knowledge about a subset of topics  $\mathcal{A}_q \in \mathcal{A}$ . Here, we assume that knowledge is additive, *i.e.*,  $k_q = \sum_{a \in \mathcal{A}_q} k_{qa} = \mathbf{w}_a^T \mathbf{k}_q$ , where  $k_{qa} \in \mathbb{R}^+$  is the knowledge value contained in item  $q$  about topic  $a$ ,  $\mathbf{k}_q = [k_{qa}]_{a \in \mathcal{A}}$ , and  $w_{qa} = 1$  if  $a \in \mathcal{A}_q$  and  $w_{qa} = 0$ , otherwise. The model can be extended to non-binary weights to represent fractional presence of topics in a knowledge item [7].

Then, we define two types of events: *learning events*, in which users acquire knowledge by reading contributions by other users, and *contributing events* (or *contributions*), in which users contribute to the crowd by sharing their knowledge. Formally, we represent each learning event as a triplet

$$l := \left( \underset{\substack{\uparrow \\ \text{user}}}{u}, \underset{\substack{\uparrow \\ \text{time}}}{t}, \overset{\substack{\downarrow \\ \text{knowledge item}}}{q} \right), \quad (1)$$

which means that a user  $u \in \mathcal{U}$  learned from knowledge item  $q$  at time  $t$ . Here, a knowledge item  $q$  may contain one or more contributions by other users. For example, in a Q&A site, a knowledge item corresponds to a question and its answers, typically contributed by different users. In a learning event, we do not distinguish the knowledge provided by individual contributions, but instead, consider the knowledge of the item as a whole. Moreover, note that, if the knowledge value of an item is zero, the learning event will not increase the expertise of the learner. Then, we denote the history of learning events associated to user  $u$  up to time  $t$  by  $\mathcal{H}_u^l(t) = \bigcup_{i:t_i < t} \{l_i \mid u_i = u\}$ , and the history of learning events in the whole crowdlearning site up to time  $t$  by  $\mathcal{H}^l(t) = \bigcup_{i:t_i < t} \{l_i\}$ .

Similarly, we represent each contribution as a quadruplet

$$c := \begin{array}{cccc} & & \text{knowledge item} & \\ & & \downarrow & \\ (u, & t, & q, & s) \\ \uparrow & \uparrow & & \uparrow \\ \text{user} & \text{time} & & \text{score} \end{array}, \quad (2)$$

which means that a user  $u \in \mathcal{U}$  contributed to a knowledge item  $q$  at time  $t$ , and other users assigned a score  $s$  to her contribution. For example, in a Q&A site, this may be the number of upvotes an answer receives. We gather the history of contributions in the whole crowdlearning site up to time  $t$  by  $\mathcal{H}^c(t) = \bigcup_{i:t_i < t} \{c_i\}$ , and the history of contributions and learning events up to time  $t$  by  $\mathcal{H}(t) = \mathcal{H}^c(t) \cup \mathcal{H}^l(t)$ .

**Crowdlearning generative process.** We represent each user's expertise as a multidimensional (latent) stochastic process  $e_u^*(t)$ , in which the  $a$ -th entry,  $e_{ua}^*(t) \in \mathbb{R}^+$ , represents the user  $u$ 's expertise on topic  $a$  at time  $t$ . Here, the sign  $*$  means that the expertise  $e_{ua}^*(t)$  depends on her learning history  $\mathcal{H}_u^l(t)$ . Then, every time a user  $u$  contributes to a knowledge item  $q$  at time  $t$ , we draw the contribution's score from a distribution  $p(s|\mathcal{A}_q, e_u^*(t))$ . Further, we represent the times of the learning and contributing events within the site by two sets of counting processes, denoted by two vectors  $\mathbf{N}^l(t)$  and  $\mathbf{N}^c(t)$ , in which the  $u$ -th entries,  $N_u^l(t)$  and  $N_u^c(t)$ , count the number of times user  $u$  learned from and contributed to the crowdlearning site up to time  $t$ . Then, we can characterize these counting processes using their corresponding intensities as

$$\mathbb{E}[d\mathbf{N}^l(t) \mid \mathcal{H}(t)] = \boldsymbol{\lambda}^l(t) dt \quad \text{and} \quad \mathbb{E}[d\mathbf{N}^c(t) \mid \mathcal{H}(t)] = \boldsymbol{\lambda}^c(t) dt$$

where  $d\mathbf{N}^l(t) := [dN_u^l(t)]_{u \in \mathcal{U}}$  and  $d\mathbf{N}^c(t) := [dN_u^c(t)]_{u \in \mathcal{U}}$  denote the number of learning and contributing events in the window  $[t, t + dt)$  and  $\boldsymbol{\lambda}^l(t) := [\lambda_u^l(t)]_{u \in \mathcal{U}}$  and  $\boldsymbol{\lambda}^c(t) := [\lambda_u^c(t)]_{u \in \mathcal{U}}$  denote the vector of intensities associated to all the users. Here, there is a wide variety of intensity functions one can choose from [1]. However, modeling the times of learning and contributing events is not the main focus of this work – we refer the reader to the growing literature on social activity modeling using point processes [11, 33, 37]. Next, we specify the functional form of each user's expertise  $e_u^*(t)$  and the score distribution  $p(s|\mathcal{A}_q, e_u^*(t))$ .

**Stochastic process for expertise.** The expertise  $e_{ua}^*(t)$  of a user  $u$  on a topic  $a$  at time  $t$  takes the following form:

$$e_{ua}^*(t) := \underbrace{\alpha_{ua}}_{\text{initial expertise}} + \underbrace{\mu_{ua} \cdot t}_{\text{off-site learning}} + \overbrace{\sum_{i:q_i \in \mathcal{H}_u^l(t)} k_{q_i a} \cdot \kappa_\omega(t - t_i)}^{\text{on-site learning}}$$

where the first term,  $\alpha_{ua} \in \mathbb{R}^+$ , models the initial expertise of user  $u$  on a topic  $a$  when she joined the crowdlearning site; the second term,  $\mu_{ua} \in \mathbb{R}^+$ , assumes a linear trend for the off-site learning process as a first order approximation<sup>1</sup>; and, the third term models the knowledge a user acquires by means of learning events within the crowdlearning site. Here,  $\kappa_\omega(t)$  is a nonnegative kernel function that models the rate at which users forget the knowledge they learn from knowledge items. Following previous work on the psychology literature [4, 21], which argues that people forget at an exponential rate, we opt for

<sup>1</sup>Several other shapes for the learning curve have been proposed in Heathcote, et al. [16]. However, we chose the linear form for its simplicity and ease in model estimation, as suggested by Skinner [32].

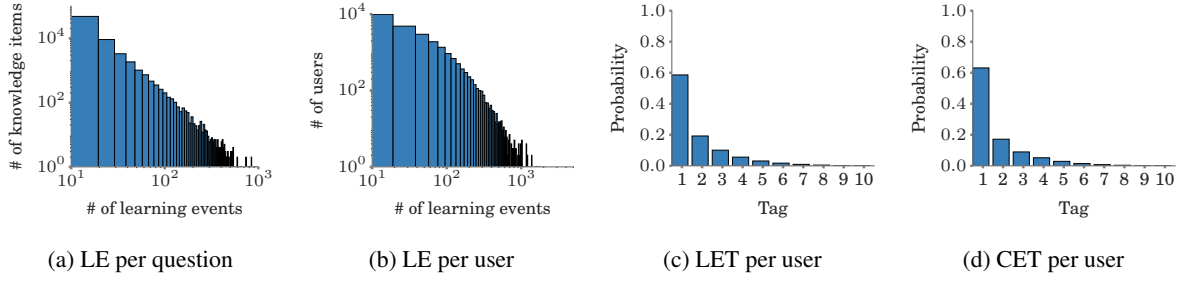


Figure 1: Statistics of learning events (LE), tags of learning events (LET) and tags of contributing events (CET) in the Stack Overflow dataset. In Panels (c) and (d), the  $x$ -axis denotes the tag index in order of popularity for each user.

an exponential kernel  $\kappa_\omega(t) := \exp(-\omega t)\mathbb{I}(t \geq 0)$ . However, our model estimation method does not depend on this particular choice.

For compactness, we write each user’s expertise as a row vector of length  $|\mathcal{A}|$ , *i.e.*,

$$\mathbf{e}_u^*(t) = \boldsymbol{\alpha}_u + \boldsymbol{\mu}_u \cdot t + \sum_{i:q_i \in \mathcal{H}_u^l(t)} \mathbf{k}_{q_i} \cdot \kappa_\omega(t - t_i) \quad (3)$$

where  $\boldsymbol{\alpha}_u = [\alpha_{ua}]_{a \in \mathcal{A}}$ ,  $\boldsymbol{\mu}_u = [\mu_{ua}]_{a \in \mathcal{A}}$  and  $\mathbf{k}_{q_i} = [k_{q_i a}]_{a \in \mathcal{A}}$ . Here, by definition,  $k_{q_i a} = 0$  if  $a \notin \mathcal{A}_{q_i}$ . Then, we can gather the model parameters for all users in three matrices  $\boldsymbol{\alpha}$ ,  $\boldsymbol{\mu}$  and  $\mathbf{k}$  with sizes  $|\mathcal{U}| \times |\mathcal{A}|$ ,  $|\mathcal{U}| \times |\mathcal{A}|$  and  $|\mathcal{Q}| \times |\mathcal{A}|$ .

**Score distribution.** Given a contribution  $c = (u, t, q, s)$ , the particular choice of score distribution  $p(s|\mathcal{A}_q, \mathbf{e}_u^*(t))$  depends on the observed data. In this work, we consider discrete non-negative scores,  $s \in \mathbb{N}$ , which fit well several scenarios of interest. For example, in Stack Overflow, scores may correspond to the number of upvotes that answers receive; in Twitter, to the number of likes or retweets that tweets receive; and, in Pinterest, to the repins that a pin receives. A natural choice in such cases is the Poisson distribution:

$$p(s|\mathcal{A}_q, \mathbf{e}_u^*(t)) \sim \text{Poisson} \left( \frac{\mathbf{w}_q^T \mathbf{e}_u^*(t)}{\mathbf{w}_q^T \mathbf{1}} \right), \quad (4)$$

Here,  $\mathbf{1}$  is a column vector of ones with length  $|\mathcal{A}|$ . With this choice, the average of the score distribution is simply the average expertise of user  $u$  at time  $t$  across the topics  $\mathcal{A}_q$  the knowledge item  $q$  is about. Moreover, the greater the expertise of a user, the greater the scores given by other users to her contributions, as one may expect in real-world data.

Note that depending on the recorded data, we could choose a different score distribution, *e.g.*, for continuous assessments like time elapsed between the question and the answer, one may choose a continuous distribution. Our model estimation method in Section 2 can be easily adapted to any distribution that is jointly log-concave with respect to the model parameters  $\boldsymbol{\alpha}$ ,  $\boldsymbol{\mu}$  and  $\mathbf{k}$ .

**Efficient Parameter Estimation.** Given a collection of learning and contributing events,  $\mathcal{H}^l(T)$  and  $\mathcal{H}^c(T)$ , recorded during a time period  $[0, T)$ , we find the optimal model parameters  $\boldsymbol{\alpha}$ ,  $\boldsymbol{\mu}$  and  $\mathbf{k}$  by solving the following maximum likelihood estimation problem:

$$\underset{\boldsymbol{\alpha} \geq 0, \boldsymbol{\mu} \geq 0, \mathbf{k} \geq 0}{\text{maximize}} \quad \mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\mu}, \mathbf{k}), \quad (5)$$

where we compute the log-likelihood  $\mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\mu}, \mathbf{k})$  using Eq. 3 and Eq. 4, *i.e.*,

$$\mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\mu}, \mathbf{k}) = \sum_{\substack{(u,t,q,s) \\ \in \mathcal{H}^c(T)}} s \cdot \log \left( \frac{\mathbf{w}_q^T \mathbf{e}_u^*(t)}{\mathbf{w}_q^T \mathbf{1}} \right) - \frac{\mathbf{w}_q^T \mathbf{e}_u^*(t)}{\mathbf{w}_q^T \mathbf{1}}. \quad (6)$$

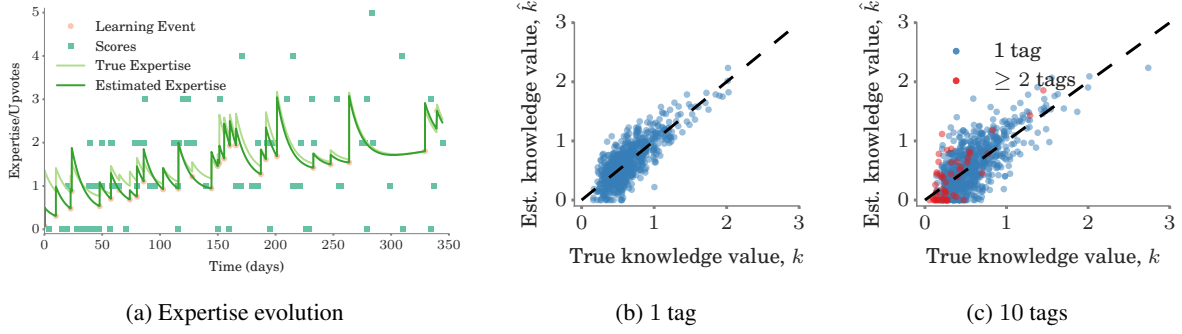


Figure 2: In (a), estimated and true expertise evolution for a user, picked at random, in the 1-tag synthetic dataset. In (b) & (c), estimated (y-axis) against true (x-axis) knowledge item values. Each point corresponds to a knowledge item variable, and the line defined by  $x = y$  corresponds to zero estimation error. Our estimation method achieves Spearman’s correlations  $\rho_{1\text{-tag}} = 0.74$  and  $\rho_{10\text{-tag}} = 0.64$ .

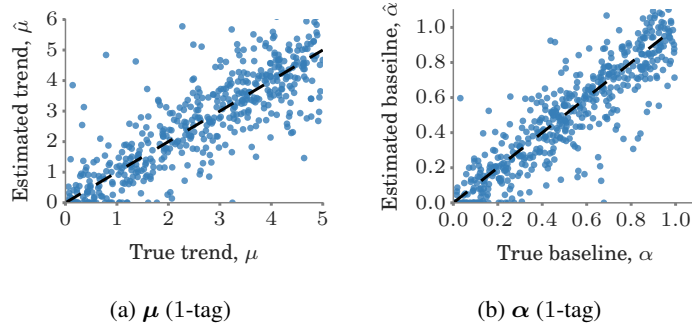


Figure 3: Estimated (y-axis) against true (x-axis) model parameters for the 1-tag synthetic datasets. Each point corresponds to a user’s (a) trend  $\mu_u$  or (b) baseline  $\alpha_u$  variable, and the line defined by  $x = y$  corresponds to zero estimation error. Our estimation method achieves a Spearman’s correlation  $\rho_\mu = 0.82$  and  $\rho_\alpha = 0.89$ . The results for 10-tag synthetic datasets are qualitatively similar.

Since  $e_u^*(t)$  is linear in the model parameters  $\alpha$ ,  $\mu$  and  $k$ , the function  $\log x - x$  is concave, and a composition of a linear function with a linear combination of concave functions is concave, the optimization problem above is jointly convex in  $\alpha$ ,  $\mu$ , and  $k$ . As a consequence, the global optimum can be efficiently found by many algorithms. In practice, the limited memory BFGS with bounded variables (L-BFGS-B) algorithm [38] worked best for our problem.

**Remarks.** In this work, we are measuring *effective learning*, which accounts for the ability of a user to get better assessment of her posts, and *effective knowledge*, which accounts for the gain in this ability that learning from a knowledge item causes. Making these quantities correspond to real-life expertise and knowledge value on a crowdlearning website requires careful mapping from the features on that website to learning events and scores.

Moreover, using our model, one can only measure learning and knowledge if there is overlap between the topics of a user’s learning and contributing events. Therefore, there is a trade-off between the granularity of the topics chosen and the amount of data available for inference: increasing the granularity ensures accurate mappings between learning and contributing events, but reduces the amount of data available to learn the model parameters. We discuss this further in Section 5.

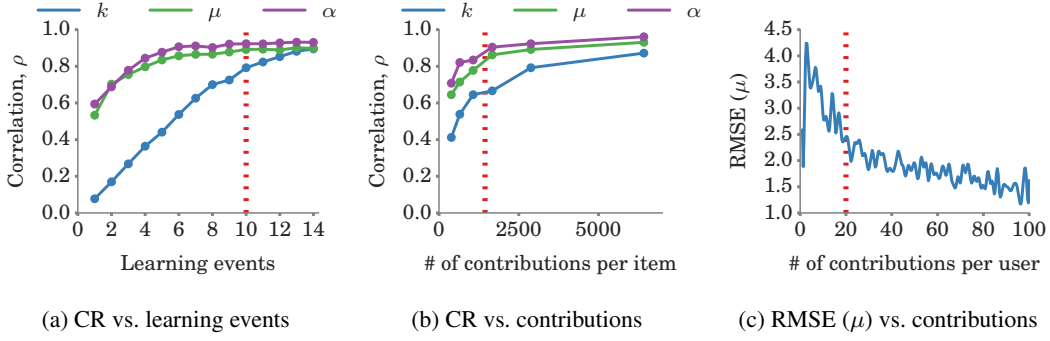


Figure 4: Estimated against true model parameters for the 1-tag synthetic dataset. Panels (a) and (b) show the correlation (CR) between the estimated and true model parameters against number of learning events and median number of contributions per knowledge item, respectively. Panel (c) shows the RMSE for the estimated trends,  $\mu$ , against number of contributed events per user. In Panel (a), the number of contributing events is 255,000 and the red dotted line shows the threshold (10) we chose for the learning events per knowledge item in our dataset (see Section 4). In Panels (b) and (c), the number of learning events is always 13,000 and the red dotted lines show the median number of contributions per knowledge item and the minimum number of contributions per user in the experiments on our dataset, respectively (see Section 4).

### 3 Experiments on synthetic data

In this section, we first show that our model estimation method can accurately recover the true model parameters from learning and contributing events generated under *realistic* conditions. We then show that, as long as there are a sufficient number of contributions *per* learning event, the estimation becomes more accurate as we feed more events into the estimation procedure. Finally, we show that our estimation method can easily scale up to millions of users, knowledge items, and learning and contributing events.

**Experimental setup.** We carefully craft an experimental setup to closely mimic some of the empirical patterns observed in real crowdlearning data, as given by Figure 1. Here, for simplicity, we assume the topics associated to each knowledge item are specified by means of tags.

Given a set of users and knowledge items, we draw the users’ offsite learning rates  $\{\mu_{ua}\}$  and initial expertise  $\{\alpha_{ua}\}$  from  $U(0, 5)$  and  $U(0, 1)$ , and the knowledge value of the items from the rescaled log-normal distribution  $0.05 \times \ln \mathcal{N}(0, 1)$ . These choices ensure that the distribution of scores which users receive resembles the distribution in real data. We set the users’ forgetting decay rate to  $\omega = (11.6 \text{ days})^{-1}$ , such that 50% of the knowledge is forgotten roughly after the first week, and assume that the intensities of both users’ learning and contributing events are (homogeneous) Poisson processes. We denote the total simulation time by  $T$ . We set each user’s learning event rate to  $T/n$ , where  $n$  is drawn from a log-normal distribution, so that the number of events per user fits well the empirical distribution (see Figure 1b), and each user’s contributing rate to  $T/m$ , where  $m$  is drawn from a uniform distribution for easy control. Finally, for each user, we shuffle the tag labels and set her tag learning propensity, defined as the probability that she up-votes a knowledge item with a given tag, and her tag contributing propensity, defined as the probability that she contributes to a knowledge item with a given tag, using the empirical distributions (see Figures 1c and 1d).

Then, we generate learning and contributing events as follows. First, we generate the timings of each user’s learning events by drawing samples from the corresponding Poisson process, and assign each learning event to a knowledge item such that the user’s tag learning propensity is satisfied. Then, we generate the timings of each user’s contributions by drawing samples from the corresponding Poisson process, assign each contributing event to a knowledge item such that the user’s tag contributing

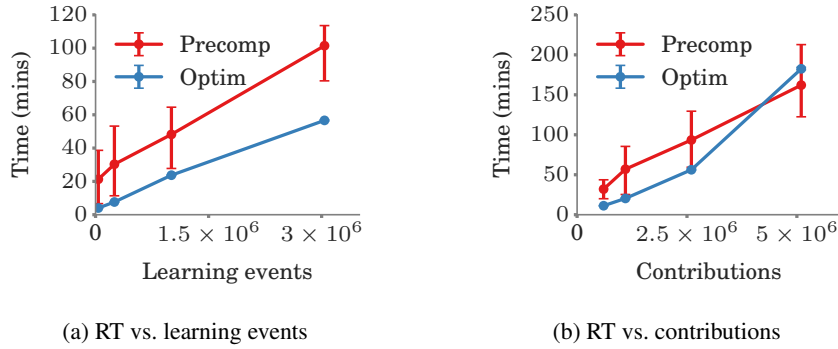


Figure 5: Running time (RT) of our model estimation method. In Panel (a), we consider  $\sim 2$  million contributing events while varying number of learning events (and knowledge items); in Panel (b), we consider  $\sim 1.8$  million learning events while varying the number of contributions (per learner). For pre-processing, we used ten machines with 48 cores and, for the optimization itself, we used a single machine with 48 cores. The memory requirements were below 16 GB at all points of the pre-processing and optimization.

propensity is satisfied, and draw the quality score from a Poisson distribution that depends on the user’s expertise on the item tags at the time of the event, as given by Eq. (4). Unless explicitly stated, we only consider knowledge items with at least 10 associated learning events. Given this data, our goal is to find the knowledge value of the items users learned from, as well as the users’ offsite learning rates and initial expertise by solving the maximum likelihood estimation problem defined in Eq. (5).

**Parameter estimation accuracy.** We evaluate the accuracy of our model estimation procedure across all users and knowledge items for a 1- and 10-tag dataset with  $\sim 800$  knowledge items. Figure 2 summarizes the results for the estimation of the knowledge item values by means of two scatter plots. In all cases, we find that points lie close to the line  $x = y$ , *i.e.*, their estimation error is close to zero. We also observe that the estimation of knowledge items in the 10-tag dataset is more challenging than in the 1-tag dataset. Additionally, Figure 3 summarizes the results for the estimation of the user’s expertise baseline and trend variable for the 1-tag dataset using scatter plots. Results for the 10-tag dataset are qualitatively similar although the estimation is more challenging. In particular, if we look at the estimation of the knowledge values, trends and baseline for the 1-tag dataset, our estimation method achieves a Spearman’s correlations  $\rho_k = 0.74$ ,  $\rho_\mu = 0.82$  and  $\rho_\alpha = 0.89$  while, for the 10-tag dataset, it achieves  $\rho_k = 0.64$ ,  $\rho_\mu = 0.76$  and  $\rho_\alpha = 0.81$ . This is most likely due to the mixing of tag knowledge variables within the same knowledge item, *i.e.*, the linear combination of knowledge variables in Eq. (6).

**Parameter estimation accuracy vs. number of learning events.** In our model, we can think of learning events as measurements of the amount of knowledge in a knowledge item, which are accumulated over time in the users’ expertise, and of contributing events as noisy samples of the users’ expertise at particular points in time. Therefore, intuitively, the more users learn from a knowledge item the easier it should become to accurately estimate their associated knowledge value, as long as these users also contribute to other knowledge items with overlapping topics. Figure 4a confirms this intuition by showing the Spearman’s correlation against minimum number of learning events per knowledge item in a 1-tag dataset with 255,000 contributions.

**Parameter estimation accuracy vs. number of contributing events.** As pointed out above, we can think of the score of contributing events as noisy samples of users’ expertise at particular points in time. Therefore, one may expect the accuracy of our model parameter estimation to improve as the number of contributions increases, due to a more fine-grained sampling of each user’s expertise. Fig-



ure 4b gives empirical evidence that this indeed happens, by showing the Spearman’s correlation against average number of answers *per* learning event in a 1-tag dataset with 13,000 learning events. Figure 4c shows how the RMSE of the estimation of  $\mu$  decreases as the number of contributions made by the user increases.

**Scalability of parameter estimation.** Crowd-learning sites such as Stack Overflow or AskReddit are rapidly increasing their number of active users, questions and answers. For example, Stack Overflow recently crossed the  $\sim 10$  million questions mark<sup>2</sup>. The pre-computation of all coefficients in Eq. (6), which is the running time bottleneck, can be readily parallelized. Figure 5 shows that our model estimation method easily handles up to millions of learning events and contributing events, and scales almost linearly with the number of learning events and contributions.

Thus, it should be possible to scale up our estimation method even further.

## 4 Experiments on Real data

In this section, we apply our model estimation method to a large-scale crowdlearning dataset from Stack Overflow. First, we evaluate our model quantitatively by means of a prediction task: given two different answers to a question, predict which one will receive a higher score. Then, we discuss the distribution of the knowledge values and the effect of the kernel parameter on the estimation, identify different types of learners and derive insights into their main characteristics. Finally, we study the interplay between learners and contributors in crowdlearning sites and investigate to which extent users switch between learning and contributing over time.

**Data description.** Our Stack Overflow dataset comprises  $\sim 8$  million questions,  $\sim 13.7$  million answers, and  $\sim 47.2$  million upvotes. These questions and answers were posted by  $\sim 1.9$  million users during a six year period, from the site’s inception on July 31, 2008 to September 14, 2014. Importantly, for each upvote, our dataset contains its associated user identity, question or answer identity and timestamp<sup>3</sup>. We discard the events which happened before 2010-01-01 (before the site had fully matured) and after 2014-06-01 (the extent of the data-dump we had access to). Whenever in the data a user *upvotes* (*writes*) an answer, we record it as a learning (contributing) event involving the user and the knowledge item containing the answer. Moreover, we select the number of upvotes a user’s answer received in the first week after posting it as the score of the contribution; downvotes were discarded because they constitute less than 3% of total votes cast. Here, we consider only the first week of voting to prevent old contributions from gaining an unfair advantage as they have more time to accumulate upvotes. Figure 1 provides general statistics on learning and contributing events and tags usage. We find that the learning events per user (per question) follow a log-normal (power-law) distribution. As shown, the tag usage is highly skewed towards few tags; most users contribute and learn only from their favorite tags.

**Data preprocessing.** In Section 3, we have shown that the accuracy of our estimation method depends dramatically on the number of learning and contributing events per question and user (refer to Figure 4). As a consequence, we can only expect our model estimation method to provide reliable and accurate results in real data if the data we start with contains enough learning and contributing events per question and user. To this aim, we carefully pre-process our large-scale dataset of learning and contributing events. We only consider:

- (i) Knowledge items with more than 10 associated learning events, which corresponds to a correlation value  $\geq 0.8$  between true and estimated knowledge parameters in synthetic data, as shown in Figure 4a.

---

<sup>2</sup><http://meta.stackoverflow.com/questions/303045/10-million-questions>

<sup>3</sup>Stack Overflow generously gave us access to these additional metadata, which allows us to readily fit our model.

Score difference	# of pairs	Off-site only	Our model
$\geq 1.0$	31,639	52.5%	<b>61.9%</b>
$\geq 2.0$	19,253	52.9%	<b>64.8%</b>
$\geq 3.0$	10,804	53.2%	<b>67.0%</b>
$\geq 4.0$	5,910	53.7%	<b>70.7%</b>
$\geq 5.0$	3,250	55.0%	<b>71.6%</b>
$\geq 6.0$	1,935	56.0%	<b>73.3%</b>
$\geq 7.0$	1,159	56.8%	<b>73.8%</b>

Table 1: Performance of our model against a linear baseline model at predicting which one of two answers to a question will receive a higher score. As the difference in score between the answers (and, hence, the users’ expertise) increases, the competitive advantage of our model becomes more pronounced.

- (ii) Users that contribute (answer) more than 20 times in at least 10 unique months, which corresponds to a RMSE value  $\leq 2$  for the estimated users’ baseline and trend parameters in synthetic data, as shown in Figure 4c; and,
- (iii) Top 10 tags in terms of number of learning events in the recorded data (*i.e.*, java, c#, javascript, php, android, jquery, python, html, c++, and mysql).

After these preprocessing steps, our dataset consists of  $\sim 25$  thousand users who learn from  $\sim 66$  thousand knowledge items by means of  $\sim 1.4$  million learning events, and contribute to  $\sim 2.5$  million knowledge items, by means of  $\sim 3.8$  million contributing events. Then, we correct for the overall decreasing trend on number of upvotes per answer over time<sup>4</sup> and since, for each knowledge item, most learning events occur after all the contributions (answers) to the knowledge item took place, we assume its knowledge value to be constant. We use the first event of each user in our dataset as an estimate of her joining time.

Finally, we would like to highlight that the preprocessing steps above do not aim to reduce the size of the original dataset but to increase the accuracy of our estimated model parameters and the reliability of our derived qualitative insights — our model estimation method does easily scale to millions of learning and contributing events. In this case, the pre-processing of the raw data using five machines with 48 cores each took  $\sim 30$  minutes and our estimation method, implemented using the Intel MKL libraries, took  $\sim 11.5$  hours on a single machine with 48 cores. The memory requirements were below 16 GB at all points of the pre-processing and optimization.

**Quantitative evaluation.** We evaluate our model quantitatively by means of the following prediction task: given two different answers to a question, predict which one will receive a higher score, *i.e.*, more number of upvotes in the week after posting it.

— *Experimental setup:* We train our model using the first 80% of the answers provided by each learner, as well as the learning events which occurred before them. Then, we match pairs of answers to the same questions from the remaining 20% and predict which one will receive a higher score. Here, we only consider questions with pairs of answers provided by users from our dataset such that their scores differ by at least one upvote. There are  $\sim 32$  thousand such pairs in our dataset. Finally, we compare our model against a baseline linear model which only accounts for off-site learning to show the benefits of including knowledge item variables.

— *Results:* Table 1 compares the performance of our model against the baseline model as the difference in score between the answers (and, hence, the users’ expertise) increases. Our model consistently

<sup>4</sup>The number of upvotes per answer decreases over time because the number of answers grows at a faster rate than the number of learners.

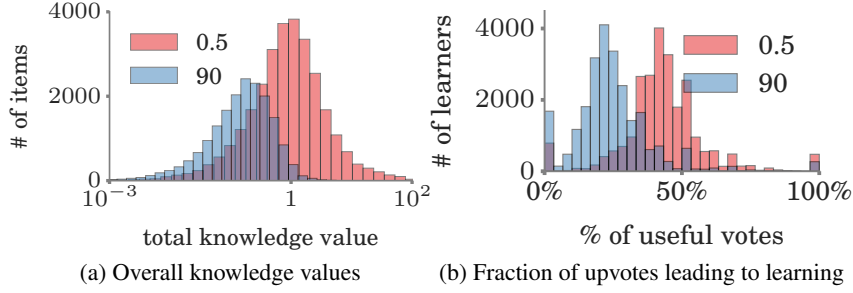


Figure 6: Estimated knowledge values for knowledge items and *useful upvotes* for two different kernel parameters, with *half-life* 0.5 days (12 hours) and 90 days. Panel (a) shows the distribution of knowledge value per knowledge item follows a log-normal distribution with longer half-life leading to smaller knowledge values and higher sparsity. Panel (b) shows what fraction of upvotes were useful (*i.e.*, led to learning) per learner: higher half-life leads to higher sparsity, which leads to fewer fraction of upvotes causing effective learning.

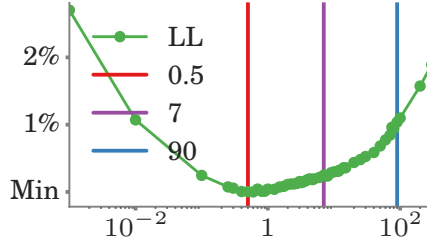


Figure 7: Negative log-likelihood plotted for different values of kernel parameter (expressed as half-life in days). The  $y$ -axis shows the relative difference with respect to the minimum value. The likelihood nearly plateaus ( $\sim 1\%$ ) for half-life between 0.5 and 90 days. The results we present are robust to parameter changing within the range and we chose 7 days as a representative value.

outperforms the baseline for any score difference and the competitive advantage becomes more pronounced as the score difference increases, reaching  $>73\%$  accuracy when the score difference is  $\geq 6$ .

**Knowledge value and forgetting rate.** In this section, we leverage our model to give insights on the knowledge values across items in Stack Overflow for different forgetting rates, *i.e.*, the kernel decay parameter  $\omega$ . We express the kernel decay parameter  $\omega$  in units of *half-life* in days, *i.e.*, the time to forget 50% of the knowledge in an item. Figure 6a shows the distribution of estimated knowledge value across knowledge items for two kernel parameters with half-life 0.5 days and 90 days. We find several interesting patterns. Knowledge values in both settings follow a log-normal distribution, in which  $\sim 10\%$  of the items account for  $\sim 75\%$  of the overall knowledge. However, while for a half-life of 0.5 days,  $\sim 53\%$  of the knowledge items do not contribute knowledge, this fraction increases to  $\sim 70\%$  for 90 days. A potential explanation for this difference is that, by increasing the half-life, a knowledge item must show evidence of *effective learning* over longer stretches of time to contribute knowledge and this happens more rarely. As a consequence, a smaller fraction of upvotes lead to effective learning (*i.e.*, being *useful*) when the half-life is high, as shown in Figure 6b – when the half-life is 0.5 days (90 days), 42% (24%) of upvotes lead to learning.

In the remaining sections, for ease of exposition, we set the kernel parameter such that the half-life of knowledge is 7 days (refer Figure 7), however, the insights obtained in the following sections are robust to changes in the kernel parameter.

**Types of learners.** Here, our goal is to better understand the type of learners that use crowd-learning sites as well as their characteristic properties. To this aim, we start by visualizing the estimated learning

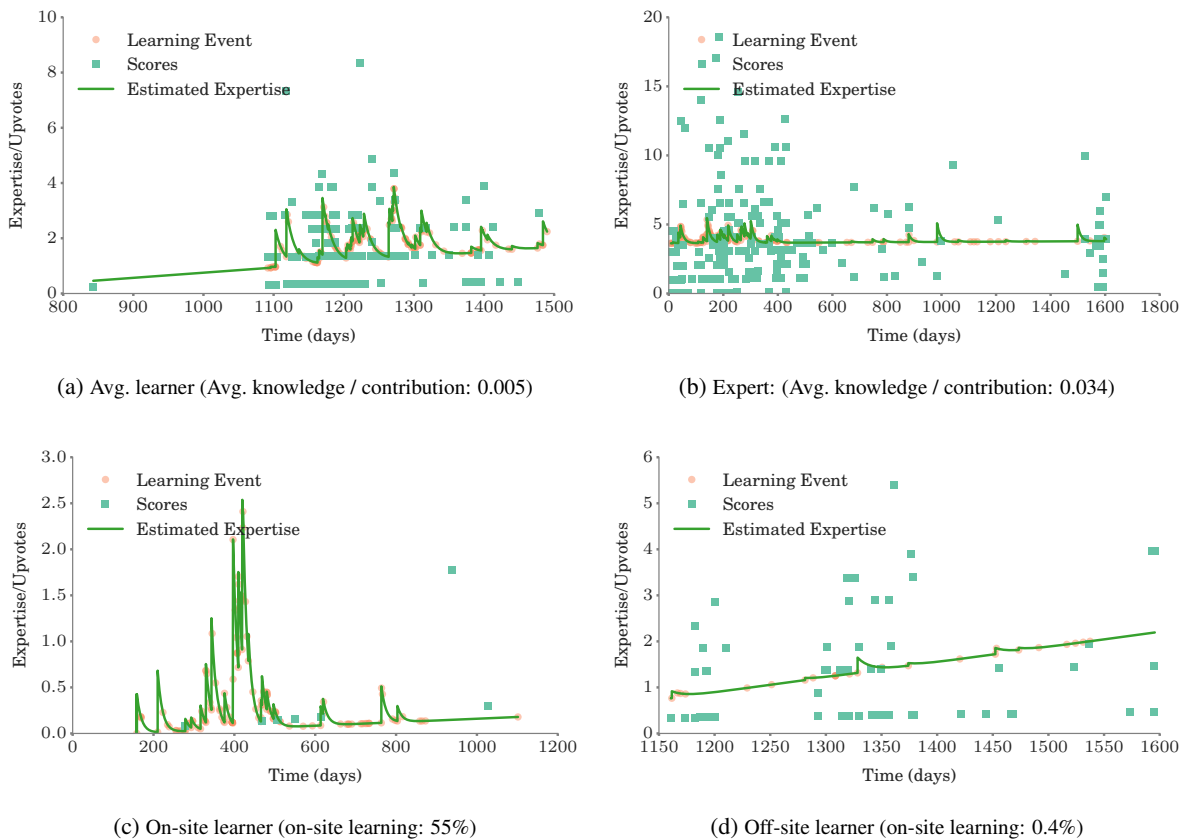


Figure 8: Estimated learning trajectory for four characteristic Stack Overflow users. The (average) learner (a) contributes answers with much less knowledge value than the expert (b), *i.e.*, 0.005 vs 0.034. The on-site learner (c) acquires 55% of her knowledge by learning from items in Stack Overflow in contrast with the off-site learner (d), who only learns 0.4% of her knowledge by those means. Day 0 in the plots is the date 2010-01-01.

trajectory for four different users — an average learner, an on-site learner, an off-site learner and an expert — in Figure 8. Each of the users exhibits different characteristic properties. For example, the average learner contributes answers with much less knowledge value (0.005) than the expert (0.034), and the on-site learner acquires 55% of the knowledge by learning from items in Stack Overflow in contrast with the off-site learner, who only learns 0.4% of the knowledge by those means.

Next, we investigate the interplay between on-site and off-site learning across all users. Here, given user  $u$ , we define on-site learning as the total expertise gathered by reading the knowledge items,  $\sum_{a \in \mathcal{A}} \sum_{q \in \mathcal{H}_u^l(T)} \int k_{qa} \kappa_w(t) dt$ , off-site learning as the expertise gathered outside Stack Overflow,  $\sum_{a \in \mathcal{A}} \int \mu_{ua} t dt$ , and overall learning as the sum of both. One can think of these quantities as the aggregate number of upvotes (*i.e.*, score) users would have received on the site through either their on or off-site learning if they were posting answers at the same rate. Note that unlike the reputation on Stack Overflow, which is a measure of how much a user has effected others on the site, the on-site and off-site learning reflects how much a user has learned. Figure 9a compares users' on-site and off-site learning by means of a box plot. For  $x \leq 2000$ , users achieving higher on-site learning also achieve higher off-site learning, but over  $x > 2000$ , off-site learning becomes more dominant. Our results seem to indicate that quick learners rely less on on-site learning, in relative terms.

Finally, we investigate the role that a user's starting expertise plays on her overall learning over time

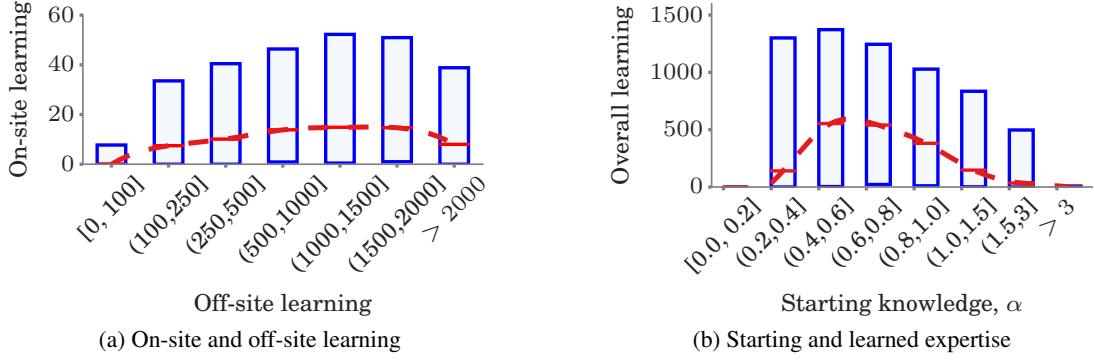


Figure 9: Behavior of learners in for tag  $c\#$ . Panel (a) compares users’ on-site and off-site learning in a box plot. For  $x \leq 2000$ , users achieving higher on-site learning also achieve higher off-site learning, however, over  $x > 2000$ , users off-site learning becomes more dominant. Panel (b) shows users’ overall learning against starting expertise in a box plot. Users with very low or very high initial expertise, *i.e.*, newbies and experts, tend to increase their knowledge the least, in contrast, users in the middle of the range tend to increase it the most. In both panels, the limits of the boxes are the 25%–75% percentiles and the red dashed lines shows the median value.

by means of a box plot, shown in Figure 9b. Here, the  $x$ -axis corresponds to a user’s starting expertise,  $\alpha_u$ , and the  $y$ -axis to her overall learning. Interestingly, we find that users with very low or very high initial expertise, *i.e.*, newbies and experts, tend to increase their knowledge the least, in contrast, users in the middle of the range tend to increase it the most. This is in agreement with previous research, which indicated that in presence of only positive reinforcement, the gain in expertise has a sigmoidal shape for learners, *i.e.*, the newbies and experts increase their expertise at lower rates than learners with medium levels of expertise [20].

**Learners vs contributors.** A crowd-learning site is only useful if it has both learners and contributors. Here, we investigate two natural questions that emerge in such context:

- I. Are learners and contributors equally common?
- II. Are more prolific learners better contributors?

To answer the first question, we compute the distribution of learned and contributed knowledge per user. Here, we estimate the knowledge value of each contribution (*i.e.*, answer) in a knowledge item by dividing the total knowledge item value across contributions proportionally to their quality scores (upvotes). Figures 10a and 10b summarize the results. Although, in absolute numbers, there are more learners than contributors in our dataset, the amount of knowledge fed into the site by the contributors shows higher variability than the knowledge learned by users – the distribution of contributed knowledge is fat tailed ( $\alpha \approx 2.26$ ).

Next, we investigate the second question and assess whether more prolific learners are better contributors. To do so, we calculate the average knowledge value per contribution across users that have learned similar amount of knowledge over time, *i.e.*, sum of the knowledge value of all the knowledge items the user learned from,  $\sum_{a \in \mathcal{A}} \sum_{q \in \mathcal{H}_u^i(T)} k_{qa}$ . Figure 10c shows that the users that learn more knowledge are also more proficient at producing high knowledge contributions. In other words, our results suggest that “*by learning you will teach; by teaching you will learn.*”

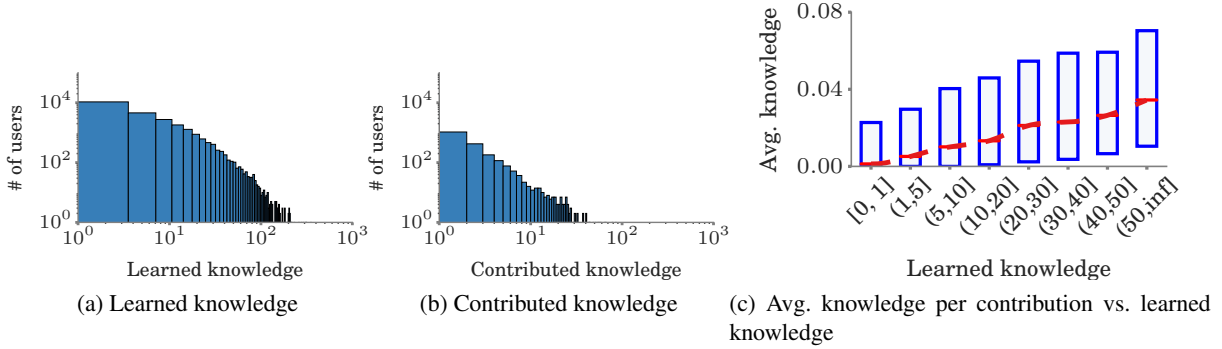


Figure 10: Learners vs contributors in Stack Overflow. Panels (a) and (b) show the distribution of overall learned and contributed knowledge per user. The former follows a log-normal distribution ( $\mu \approx 1.39, \sigma \approx 1.09$ ), while the latter follows a power-law ( $\alpha \approx 2.26, x_{min} \approx 1.84$ ). This shows that through the contributors are fewer in number than learners in absolute terms, they show much higher variability. Panel (c) shows a user’s average knowledge value per contribution against overall learned knowledge in a box plot. The red dotted line shows the median values and the box limits are the 25%–75% percentiles. Interestingly, the users that learn more knowledge are also more proficient at producing high knowledge contributions.

## 5 Discussion

In this section, we take a step back and discuss the limitations of our model. First, we remark that, due to the large number of parameters in the model, it is necessary to have access to large amount of data for our model estimation method to be accurate. However, this limitation can be overcome, to some extent, by linking expertise of a user across different platforms or sites (e.g., MOOCs), *i.e.*, our model can easily assimilate traces available for the same user from those sites.

In our model, it is also crucial that the score reflects the true assessment of the knowledge content of the item and not of, say, the popularity of the contributor. In the case of Stack Overflow, which is a strict and self-regulated community, upvotes are seldom granted to answers which do not address the question — cases of serial upvoting are caught and remedied quickly which (mostly) prevents users from voting as a *thank you* gesture. As a consequence, on Stack Overflow, upvotes on answers are a good assessment of the quality of the posts. However, a sensible choice for scores in platforms or sites with milder self-regulation may be challenging.

Finally, the learning events also need to be chosen such that they are not conflated with other objectives the user may have on the website. On Stack Overflow, if a user only upvotes a question, it indicates that she relates with the problem but none of the answers (if any) provide a solution. However, upvoting an answer is evidence that the Q&A pair *taught* the user something.

These unique features and mechanisms afforded by Stack Overflow allow us to easily identify learning events and assessments. Finding similar features in a different social network would require careful reasoning and justification.

## 6 Conclusions

In this paper, we proposed a probabilistic model of crowdlearning, naturally designed to fit fine-grained learning and contributing event data. The key innovation of our model is modeling the evolution of users’ expertise over time as a latent stochastic process, driven by both off-site and on-site learning. Then, we developed a scalable estimation method to fit the model parameters from millions of recorded learning

events and contributions. Finally, we applied our model to a large set of learning and contributing events from Stack Overflow and found several interesting insights. For example, items with high knowledge value are rare. Newbies and experts acquire less knowledge than users in the middle range. Prolific learners tend to be also proficient contributors that share knowledge with high knowledge value.

Our work also opens many interesting venues for future work. For example, natural follow-ups to potentially improve the expressiveness of our modeling framework include:

1. Consider more complex off-site learning trends, *e.g.*, isotonic regression [18] or exponential/power-law [16].
2. Allow for a knowledge item to have different knowledge values per user by considering a knowledge distribution per item, and use Bayesian inference [23] to learn the model parameters.
3. Perform a non-parametric estimation of the kernels that model the users' forgetting process. We expect this to allow clustering of knowledge items which provide short-lived and long-lasting knowledge.
4. Incorporate incentives mechanisms such as badges, which are often used in crowdlearning sites [2] and MOOCs [3].

One of the key modeling ideas behind our framework is realizing that users' contributions can be viewed as noisy discrete samples of the users' expertise at points localized (non uniformly) in time. We could generalize this idea to any type of event data and derive sampling theorems and conditions under which an underlying general continuous signal of interest (be it user's expertise, opinion, or wealth) can be recovered from event data with provable guarantees. Finally, we experimented with data gathered exclusively from Stack Overflow. It would be interesting to apply our model to Stack Exchange at large, to other questions and answers websites (*e.g.*, AskReddit), microblogging platforms (*e.g.*, Twitter), social networking sites (*e.g.*, Pinterest), or even offline crowdlearning networks (*e.g.*, citation network).

**Acknowledgements.** The authors would like to thank Sam Brand from Stack Overflow for providing data to make this work possible.

## References

- [1] O. Aalen, O. Borgan, and H. Gjessing. *Survival and event history analysis: a process point of view*. Springer Science & Business Media, 2008.
- [2] A. Anderson, D. Huttenlocher, J. Kleinberg, and J. Leskovec. Steering user behavior with badges. In *22nd International Conference on World Wide Web*, 2013.
- [3] A. Anderson, D. Huttenlocher, J. Kleinberg, and J. Leskovec. Engaging with massive online courses. In *23rd International Conference on World Wide Web*, 2014.
- [4] L. Averell and A. Heathcote. The form of the forgetting curve and the fate of memories. *Journal of Mathematical Psychology*, 55(1), 2011.
- [5] R. S. Baker, Z. A. Pardos, S. M. Gowda, B. B. Nooraei, and N. T. Heffernan. Ensembling predictions of student knowledge within intelligent tutoring systems. In *User Modeling, Adaption and Personalization*. Springer, 2011.
- [6] J. E. Beck and J. Mostow. How who should practice: Using learning decomposition to evaluate the efficacy of different types of practice for different types of students. In *Intelligent Tutoring Systems*, 2008.
- [7] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *The Journal of Machine Learning Research*, 2003.
- [8] A. T. Corbett and J. R. Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction*, 1994.
- [9] C. Eickhoff, J. Teevan, R. White, and S. Dumais. Lessons from the journey: A query log analysis of within-session learning. In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining*. ACM, 2014.
- [10] S. E. Embretson and S. P. Reise. *Item response theory*. Psychology Press, 2013.
- [11] M. Farajtabar, N. Du, M. Gomez-Rodriguez, I. Valera, L. Song, and H. Zha. Shaping social activity by incentivizing users. In *Advances in Neural Information Processing Systems*, 2014.
- [12] M. Feng, N. T. Heffernan, and K. R. Koedinger. Looking for sources of error in predicting student’s knowledge. In *Educational Data Mining: AAAI Workshop*, 2005.
- [13] S. Ghosh, N. Sharma, F. Benevenuto, N. Ganguly, and K. Gummadi. Cognos: crowdsourcing search for topic experts in microblogs. In *35th international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2012.
- [14] J. Gonzalez-Brenes and J. Mostow. What and when do students learn? fully data-driven joint estimation of cognitive and student models. In *6th International Conference on Educational Data Mining*, 2013.
- [15] B. V. Hanrahan, G. Convertino, and L. Nelson. Modeling problem difficulty and expertise in Stack Overflow. In *ACM 2012 Conference on Computer Supported Cooperative Work Companion*. ACM, 2012.
- [16] A. Heathcote, S. Brown, and D. Mewhort. The power law repealed: The case for an exponential law of practice. *Psychonomic bulletin & review*, 7(2), 2000.
- [17] P. Jurczyk and E. Agichtein. Discovering authorities in question answer communities by using link analysis. In *16th ACM Conference on Conference on Information and Knowledge Management*. ACM, 2007.
- [18] S. M. Kakade, V. Kanade, O. Shamir, and A. Kalai. Efficient learning of generalized linear and single index models with isotonic regression. In *Advances in Neural Information Processing Systems*, 2011.
- [19] W.-C. Kao, D.-R. Liu, and S.-W. Wang. Expert finding in question-answering websites. In *ACM Symposium on Applied Computing*. ACM, 2010.
- [20] N. Leibowitz, B. Baum, G. Enden, and A. Karniel. The exponential learning equation as a function of successful trials results in sigmoid performance. *Journal of Mathematical Psychology*, 54(3), 2010.
- [21] G. R. Loftus. Evaluating forgetting curves. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(2), 1985.



- [22] B. Means, Y. Toyama, R. Murphy, M. Bakia, and K. Jones. Evaluation of evidence-based practices in online learning: A meta-analysis and review of online learning studies. *US Department of Education*, 2009.
- [23] K. P. Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [24] G. R. Norman and H. G. Schmidt. The psychological basis of problem-based learning: A review of the evidence. *Academic Medicine*, 67(9), 1992.
- [25] A. Pal and S. Counts. Identifying topical authorities in microblogs. In *Fourth ACM International Conference on Web Search and Data Mining*. ACM, 2011.
- [26] A. Pal and J. A. Konstan. Expert identification in community question answering: Exploring question selection bias. In *19th ACM International Conference on Information and Knowledge Management*, 2010.
- [27] A. Paulina and J. Marta. Study of the Temporal-Statistics-Based Reputation Models for Q&A Systems. *Computer Science*, 16(3), 2015.
- [28] P. I. Pavlik, H. Cen, and K. R. Koedinger. Performance factors analysis—a new alternative to knowledge tracing. In *2009 conference on Artificial Intelligence in Education: Building Learning Systems that Care: From Knowledge Representation to Affective Modelling*. IOS Press, 2009.
- [29] C. Piech, J. Bassen, J. Huang, S. Ganguli, M. Sahami, L. J. Guibas, and J. Sohl-Dickstein. Deep knowledge tracing. In *Advances in Neural Information Processing Systems*, 2015.
- [30] D. Posnett, E. Warburg, P. T. Devanbu, and V. Filkov. Mining stack exchange: Expertise is evident from initial contributions. In *International Conference on Social Informatics*, 2012.
- [31] J. Qiu, J. Tang, T. X. Liu, J. Gong, C. Zhang, Q. Zhang, and Y. Xue. Modeling and predicting learning behavior in moocs. In *Fourth ACM International Conference on Web Search and Data Mining*. ACM, 2016.
- [32] C. H. Skinner. Applied comparative effectiveness researchers must measure learning rates: A commentary on efficiency articles. *Psychology in the Schools*, 47(2), 2010.
- [33] I. Valera and M. Gomez-Rodriguez. Modeling adoption and usage of competing products. In *2015 IEEE International Conference on Data Mining*, 2015.
- [34] R. W. White, S. T. Dumais, and J. Teevan. Characterizing the influence of domain expertise on web search behavior. In *Proceedings of the Second ACM International Conference on Web Search and Data Mining*. ACM, 2009.
- [35] M. V. Yudelson, K. R. Koedinger, and G. J. Gordon. Individualized bayesian knowledge tracing models. In *Artificial Intelligence in Education*. Springer, 2013.
- [36] J. Zhang, M. S. Ackerman, and L. Adamic. Expertise networks in online communities: Structure and algorithms. In *16th International Conference on World Wide Web*, 2007.
- [37] K. Zhou, H. Zha, and L. Song. Learning triggering kernels for multi-dimensional hawkes processes. In *30th International Conference on Machine Learning*, 2013.
- [38] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal. Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software*, 23(4), 1997.