

One-class Collaborative Filtering with Random Graphs

Ulrich Paquet
Microsoft Research Cambridge
ulripa@microsoft.com

Noam Koenigstein
Microsoft R&D, Israel
noamko@microsoft.com

ABSTRACT

The bane of one-class collaborative filtering is interpreting and modelling the latent signal from the missing class. In this paper we present a novel Bayesian generative model for implicit collaborative filtering. It forms a core component of the Xbox Live architecture, and unlike previous approaches, delineates the odds of a user disliking an item from simply being unaware of it. The latent signal is treated as an unobserved random graph connecting users with items they might have encountered. We demonstrate how large-scale distributed learning can be achieved through a combination of stochastic gradient descent and mean field variational inference over random graph samples. A fine-grained comparison is done against a state of the art baseline on real world data.

Categories and Subject Descriptors

G.3 [Mathematics of computing]: Probability and statistics

Keywords

One-class collaborative filtering, random graph, variational inference

1. INTRODUCTION

This paper highlights a solution to a very specific problem, the prediction of a “like” or “association” signal from one-class data. One-class or “implicit” data surfaces in many of Xbox’s verticals, for example when users watch movies through Xbox Live. In this vertical, we recommend media items to users, drawing on the correlations of their viewing patterns with those of other users. We assume that users don’t watch movies that they dislike; therefore the negative class is absent. The problem is equivalent to predicting new connections in a network: given a disjoint user and an item vertex, what is the chance that they should be linked?

We introduce a Bayesian generative process for connecting users and items. It models the “like” probability by interpreting the missing signal as a two-stage process: firstly, by modelling the odds of a user considering an item, and secondly, by eliciting a probability that that item will be viewed or liked. This forms a core component of the Xbox Live architecture, serving recommendations to more than

50 million users worldwide, and replaces an earlier version of our system [10]. The two-stage delineation of popularity and personalization allows systems like ours to trade them off in optimizing user-facing utility functions. The model is simple to interpret, allows us to estimate parameter uncertainty, and most importantly, easily lends itself to large-scale inference procedures.

Interaction patterns on live systems typically follow a power-law distribution, where some users or items are exponentially more active or popular than others. We base our inference on a simple assumption, that the missing signal should have the same power-law degree distribution as the observed user-item graph. Under this assumption, we learn latent parametric descriptions for users and items by computing statistical averages over all plausible “negative graphs”.

The challenge for one-class collaborative filtering is to treat the absent signal without incurring a prohibitive algorithmic cost. Unlike its fully binary cousin, which observes “dislike” signals for a selection of user-item pairs, each unobserved user-item pair or edge has a possible negative explanation. For M users and N items, this means that inference algorithms have to consider $\mathcal{O}(MN)$ possible negative observations. In problems considered by Xbox, this amounts to modelling $\mathcal{O}(10^{12})$ latent explanations. The magnitude of real world problems therefore casts a shadow on models that treat each absent observation individually [17].

Thus far, solutions to large-scale one-class problems have been based on one of two main lines of thought. One line formulates the problem as an objective function over all observed and missing data, in which the contribution by the “missing data” drops out in the optimization scheme [7]. It relies on the careful assignment of confidence weights to all edges, but there is no methodical procedure for choosing these confidence weights except an expensive exhaustive search via cross-validation. If a parametric definition of confidence weights is given, a low rank approximation of the weighting scheme can also be included in an objective function [18]. The work presented here differs from these approaches by formulating a probabilistic model rather than an optimization problem, and quantifies our uncertainty about the parameters and predictions.

A second approach is to randomly synthesize negative examples. Our work falls in this camp, for which there already exists a small body of work. The foremost of these is arguably Bayesian Personalized Ranking (BPR), which converts the one-class problem into a ranking problem [21]. In it, it is assumed that the user likes everything that she has seen more than the items that she hasn’t seen. This

assumption implies a constrained ordering of many unobserved variables, one arising from each item. This user-wise ranking of items facilitates the inference of latent features for each user and item vertex. By design, there is no distinction between missing items in BPR; however, popularity sampling of the unobserved items was employed to give more significance to popular missing items [5]. This approach was effectively utilized by many of the leading solutions in the KDD-Cup’11 competition [4]. An alternative, more expensive approach is to construct an ensemble of solutions, each of which is learned using a different sample of synthesized “negative” edges [19].

We motivate our approach by discussing properties of typical bipartite real world graphs in Section 2. A generative model for collaborative filtering when such graphs are observed is given in Section 3. A component of the model is the hidden graph of edges—items that a user considered, but didn’t like. Section 4 addresses the hidden graph as a random graph. Section 5 combines variational inference and stochastic gradient descent to present methods for large scale parallel inference for this probabilistic model. In Section 6, we show state of the art results on two practical problems, a sample of movies viewed by a few million users on Xbox consoles, and a binarized version of the Netflix competition data set.

2. TYPICAL REAL WORLD DATA

The frequency of real-world interactions typically follows some form of power-law. In Xbox Live, we observe a bipartite graph \mathbf{G} of M users and N items, with two (possibly vastly) different degree distributions for the two kinds of vertices. Figure 1 (*top*) illustrates the degree distribution of a sample of $M = 6.2 \times 10^6$ users that watched $N = 1.2 \times 10^4$ different movies on their Xbox consoles, where an edge appears if a user viewed a movie. Throughout the paper the edges in the *observed* graph \mathbf{G} will be denoted with the binary variable $g_{mn} \in \{0, 1\}$ for vertices m and n , with a zero value indicating the absence of an edge. We denote the observed degree distributions as $p_{\text{user}}(d)$ and $p_{\text{item}}(d)$. If a user viewed on average μ items, and an item was viewed on average ν times, then $\mathbb{E}_{p_{\text{user}}}[d] = \mu$ and $\mathbb{E}_{p_{\text{item}}}[d] = \nu$, and the constraint

$$\frac{\mu}{N} = \frac{\nu}{M} \quad (1)$$

should hold [15]. In Figure 1 (*top*), the empirical distributions satisfy $\mu = 7.1$ and $\nu = 3780$, validating the mean constraint $\mu/N = \nu/M = 0.0006$. We overlay a power law degree distribution to items $p_{\text{item}}(d) \propto d^{-0.77}$. The user distribution exhibits an marked exponential cut-off, with $p_{\text{user}}(d) \propto d^{-1.4} e^{-d/70}$, and shares its form with many scientific collaboration networks [16]. The degree distribution of the publicly available Netflix data set is shown in Figure 1 (*bottom*). In it, we have $M = 4.8 \times 10^5$ users and $N = 1.8 \times 10^4$ items. We took a positive edge to be present if a user rated an item with four or five stars.

Given $p_{\text{user}}(d)$ and $p_{\text{item}}(d)$, one can sample i.i.d. graphs with the given degree distribution. Firstly, generate vertex degrees for each user and item at random, and calculate their sum. If the sums are unequal, randomly choose one user and item, discard their degrees, and replace them with new degrees of the relevant distributions. This process is repeated until the total user and item degrees are equal, after which vertex pairs are randomly joined up [15].

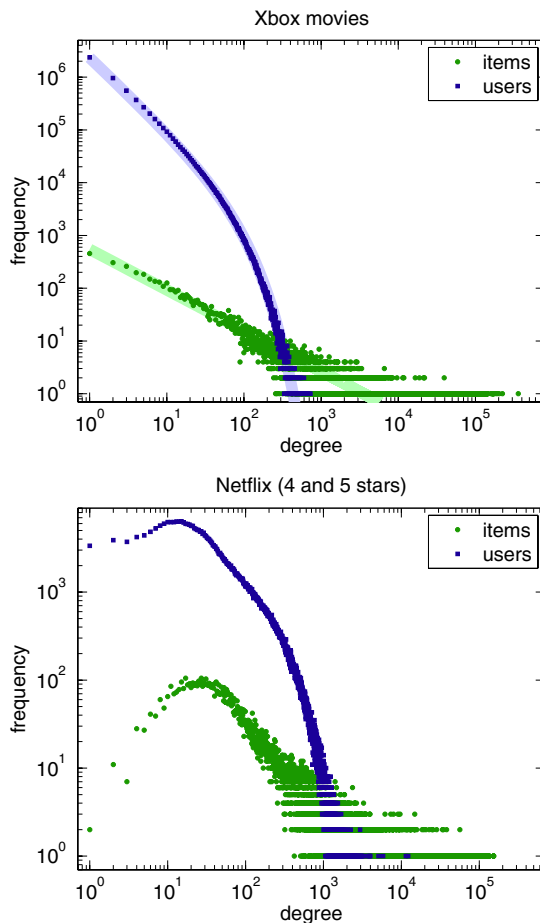


Figure 1: Degree distributions for two bipartite graphs between users and movies: a sample of 4.4×10^7 edges for movies viewed on Xbox (*top*) and the 5.6×10^7 four and five starred edges in the Netflix prize data set (*bottom*).

3. COLLABORATIVE FILTERING

Our collaborative filtering model rests on a basic assumption, that if an edge $g_{mn} = 1$ appears in \mathbf{G} , user m liked item n . However, a user must have considered additional items that she didn’t like, even though the dislike or “negative” signals are not observed. This hidden graph with edges $h_{mn} \in \{0, 1\}$ is denoted by \mathbf{H} . We say that a user considered an item if and only if $h_{mn} = 1$, and the rule $g_{mn} = 1 \Rightarrow h_{mn} = 1$ holds; namely, a user must have considered all the items that she “liked” in \mathbf{G} . The latent signal is necessary in order to avoid trivial solutions, where the interpretation inferred from data tells us that everyone likes everything or that every edge should be present. It strongly depends on our prior beliefs about \mathbf{H} , like its degree distribution or power-law characteristics. \mathbf{G} is observed as a subgraph of \mathbf{H} , while the rest of the edges of the hidden graph \mathbf{H} form the unobserved “negative” signals.

3.1 The likelihood and its properties

On knowing the hidden graph, we define a bilinear or “matrix factorization” collaborative filtering model. We asso-

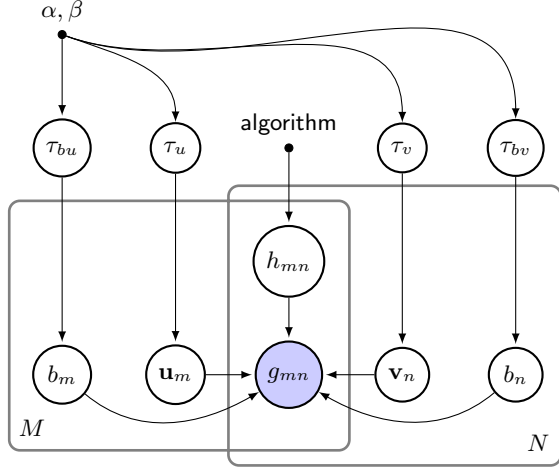


Figure 2: The graphical model for observing graph \mathbf{G} connecting M user with N item vertices. The prior on the hidden graph \mathbf{H} is algorithmically determined to resemble the type of the observed graph.

create a latent feature $\mathbf{u}_m \in \mathbb{R}^K$ with each user vertex m , and $\mathbf{v}_n \in \mathbb{R}^K$ with each item vertex n . Additionally, we add biases $b_m \in \mathbb{R}$ and $b_n \in \mathbb{R}$ to each user and item vertex. The odds of a user liking or disliking an item under consideration ($h = 1$) is modelled with

$$p(g | \mathbf{u}, \mathbf{v}, b, h = 1) = \sigma(\mathbf{u}^T \mathbf{v} + b)^g [1 - \sigma(\mathbf{u}^T \mathbf{v} + b)]^{1-g}, \quad (2)$$

with the logistic or sigmoid function being $\sigma(a) = 1/(1 + e^{-a})$, with $a \stackrel{\text{def}}{=} \mathbf{u}^T \mathbf{v} + b$. Subscripts m and n are dropped in (2) as they are clear from the context; b denotes the sum of the biases $b_m + b_n$. The likelihood of g for any h is given by the expression

$$p(g | a, h) = [\sigma(a)^g (1 - \sigma(a))^{1-g}]^h \cdot (1 - g)^{1-h}. \quad (3)$$

As $g = 1 \Rightarrow h = 1$ by construction, the last factor can be ignored in (3). If the binary ‘‘considered’’ variable h is marginalized out in (3), we find that

$$\begin{aligned} p(g = 1 | a) &= p(h = 1) \sigma(a), \\ p(g = 0 | a) &= p(h = 1)(1 - \sigma(a)) + (1 - p(h = 1)). \end{aligned} \quad (4)$$

In other words, the odds of encountering an edge in \mathbf{G} is the product of two probabilities, separating *popularity* from *personalization*: $p(h = 1)$, the user considering an item, and $\sigma(a)$, the user then liking that item.

3.2 The full model

The probability of \mathbf{G} depends on the prior distributions of the vertices’ hidden features. We choose them to be Gaussian: $p(\mathbf{U}) = \prod_{m=1}^M \mathcal{N}(\mathbf{u}_m; \mathbf{0}, \tau_u^{-1} \mathbf{I})$ for the users, where $\mathbf{U} \stackrel{\text{def}}{=} \{\mathbf{u}_m\}_{m=1}^M$, with similar Gaussian priors on the parameters governing the item vertices. These are shown in the graphical model in Figure 2. To infer the various scale parameters τ , we place a conjugate Gamma hyperprior on

each, for example

$$\mathcal{G}(\tau_u; \alpha, \beta) = \beta^\alpha / \Gamma(\alpha) \cdot \tau_u^{\alpha-1} e^{-\beta \tau_u}.$$

The only prior beliefs in Figure 2 that do not take an explicit form is that of \mathbf{H} . It could be parameterized with a particular degree distribution, be it Poisson, exponential, or a power law with an exponential cut-off. However, we would like this to (approximately) be in the same family as the observed data, and determine an algorithm which can generate such graphs. Section 4 elaborates on this, including the closure of the graph family under random sampling of subnetworks.

We collectively denote our parameters by $\boldsymbol{\theta} \stackrel{\text{def}}{=} \{\mathbf{H}, \mathbf{U}, \mathbf{V}, \mathbf{b}, \boldsymbol{\tau}\}$, with $b_{mn} \stackrel{\text{def}}{=} b_m + b_n$ as shorthand notation. The joint density of all the random variables, given the hyperprior parameters α and β , is

$$\begin{aligned} p(\mathbf{G}, \boldsymbol{\theta}) &= \prod_{m=1}^M \prod_{n=1}^N \overbrace{\sigma(\mathbf{u}_m^T \mathbf{v}_n + b_{mn})^{g_{mn}} \dots}^{\ell_{mn}} \\ &\quad \dots [1 - \sigma(\mathbf{u}_m^T \mathbf{v}_n + b_{mn})]^{h_{mn}(1-g_{mn})} \\ &\quad \cdot \prod_{m=1}^M \mathcal{N}(\mathbf{u}_m; \mathbf{0}, \tau_u^{-1} \mathbf{I}) \mathcal{N}(b_m; 0, \tau_{b_m}^{-1}) \cdot \mathcal{G}(\tau_u; \alpha, \beta) \\ &\quad \cdot \prod_{n=1}^N \mathcal{N}(\mathbf{v}_n; \mathbf{0}, \tau_v^{-1} \mathbf{I}) \mathcal{N}(b_n; 0, \tau_{b_n}^{-1}) \cdot \mathcal{G}(\tau_v; \alpha, \beta) \\ &\quad \cdot \mathcal{G}(\tau_{b_m}; \alpha, \beta) \cdot \mathcal{G}(\tau_{b_n}; \alpha, \beta) \cdot p(\mathbf{H}). \end{aligned} \quad (5)$$

The sigmoid product is denoted with ℓ_{mn} , and will later appear in a variational bound in (11). Obtaining a posterior approximation to (5) would follow known literature [20, 25], were it not for the unknown occurrence of edges h_{mn} in \mathbf{H} . Sections 4 and 5 are devoted to treating \mathbf{H} .

One might also consider placing Normal-Wishart hyperprior on the means and variances of \mathbf{u}_m and \mathbf{v}_n [20, 23]. In practice, we benefit from additionally using meta-data features in the hyperpriors. They allow us to learn how shared features connect the prior distributions of various items, but is beyond the scope of this paper.

3.3 Factorized approximation

It is analytically intractable to compute the Bayesian averages necessary for marginalization in (5). This hurdle is commonly addressed in one of two ways: Samples from the posterior can be drawn by simulating a Markov chain with the posterior as its stationary distribution, and these samples used for prediction [14]. Alternatively, one might substitute the integration problems required for Bayesian marginalization with an optimization problem, that of finding the best deterministic approximation to the posterior density [9].

We *approximate* the posterior from (5), rather than sample from it, as it allows a compact representation to be serialized to disk. The posterior from (5) is approximated with the fully factorized distribution q ,

$$\begin{aligned} p(\boldsymbol{\theta} | \mathbf{G}) \approx q(\boldsymbol{\theta}) &\stackrel{\text{def}}{=} \prod_{m=1}^M q(b_m) \prod_{k=1}^K q(u_{mk}) \cdot \prod_{n=1}^N q(b_n) \prod_{k=1}^K q(v_{nk}) \\ &\quad \cdot q(\tau_u) q(\tau_v) q(\tau_{b_u}) q(\tau_{b_v}) q(\mathbf{H}). \end{aligned} \quad (6)$$

The factors approximating each of the vertex features in \mathbf{U} , \mathbf{V} , and \mathbf{b} are chosen to be a Gaussian, for example

$q(u_{mk}) = \mathcal{N}(u_{mk}; \eta_{mk}, \omega_{mk}^{-1})$. Similarly, the τ 's are approximated by Gamma factors in the conjugate exponential family, for example $q(\tau_u) = \mathcal{G}(\tau_u; \phi_u, \varphi_u)$.

The remaining question is, what to do with $p(\mathbf{H})$, and the posterior marginal approximation $q(\mathbf{H})$?

4. RANDOM GRAPHS

Although an observation $g_{mn} = 1$ implies that $q(h_{mn} = 1) = 1$, we cannot estimate every one of MN $q(h_{mn})$'s, as there are typically $\mathcal{O}(10^{12})$ or more of them. As a recourse, we shall specify q as an algorithm that stochastically generates connections $h_{mn} = 1$, so that $p(\mathbf{H})$ produces (roughly) the same type of graphs as is observed in \mathbf{G} .

The graphical model in Figure 2 specifies that every “considered” edge (m, n) in \mathbf{H} contains a “like” probability σ_{mn} . For each edge in \mathbf{H} , a coin is flipped, and revealed with probability σ_{mn} to give \mathbf{G} . If we assume that the coin is on average unbiased, half the edges will be revealed, and $|\mathbf{H}| \approx 2|\mathbf{G}|$. Alternatively, \mathbf{G} is a subnet of \mathbf{H} , containing half (or some rate of) its connections. Working back, we sample graphs \mathbf{H} at this rate, and the family of graphs \mathbf{H} that can be generated this way constitutes our prior. This places no guarantee that the two graphs will always be of the same type, as not all graph types are closed under random sampling. For example, random subnets drawn from exact scale-free networks are not themselves scale-free [26]. However, the practical benefits of this algorithmic simplification outweigh the cost of more exact procedures.

4.1 Sampling $q(\mathbf{H})$

The factor $q(\mathbf{H})$ is defined stochastically, with the criteria that it should not be too expensive to draw random samples \mathbf{H} . One approach would be to generate samples, similar to Section 2, by specifying a degree distribution conditioned on the number of degrees d that each user and item vertex in \mathbf{G} has. If the mean of each is $2d$, one can show that a version of (1) will also hold for \mathbf{H} . At the cost of many redraws, one can sample half-edges, as in Section 2, and connect them until all half-edges are paired.

We propose a simpler scheme here, which samples \mathbf{H} from \mathbf{G} in $\mathcal{O}(|\mathbf{G}| \log N)$ time. The scheme has the flavour of “sampling by popularity” [5]. We define a multinomial histogram $\mathcal{M}(\boldsymbol{\pi})$ on the N items, where $\pi_n \geq 0$ for $n = 1, \dots, N$. This mimics a pseudo degree distribution for missing degrees. Let user m have degree d_m , or have viewed d_m items. For user m , the subset of d_m edges in \mathbf{H} that corresponds to $g_{mn} = 1$ is marked. We then sample d_m random “negative” edges from $\mathcal{M}(\boldsymbol{\pi})$ *without replacement*—this fills in the remaining values for row m in \mathbf{H} , i.e. h_{mn} for $n = 1, \dots, N$. For user m the sample *without replacement* can be drawn in $\mathcal{O}(\log N)$ time by doing bookkeeping with a weighed binary tree on the items.

There are many ways to define histogram $\boldsymbol{\pi}$, one of which is to simply let $\pi_n = d_n$, the number of degrees (or views) of item n . This is effectively a uniform prior: each item should have the same rate of negatives. If we believe that there is some quality bar that drives popular items to be more generally liked, the histogram can be adjusted with

$$\pi_n = d_n^\gamma \quad (7)$$

so that it obeys a version of the observed power law. A free rate parameter r is introduced, so that the most popular

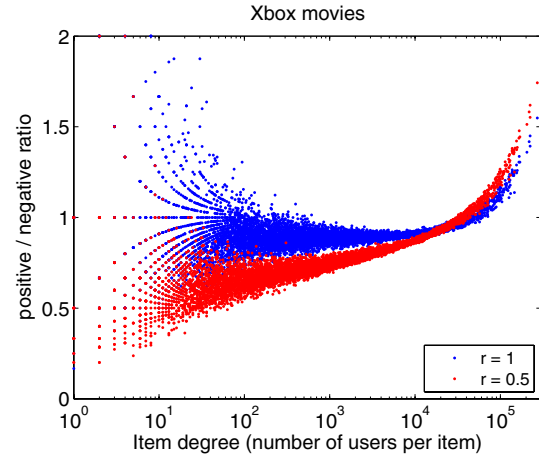


Figure 3: The ratio of positive to negative edges per item, from a single sample from $q(\mathbf{H})$. (The ratio is skewed at the head: sampled edges to more popular items have higher odds to already exist in \mathbf{G} . Discarding and resampling them leaves popular items underrepresented in the “negative” set. This can be overcome with another adjustment of $\boldsymbol{\pi}$ in $\mathcal{M}(\boldsymbol{\pi})$.)

item with degree $d_{\max} = \max\{d_n\}$ has histogram weight

$$\pi_{\max} = r d_{\max} . \quad (8)$$

As an example, $r = \frac{1}{2}$ will add half as many unobserved edges to \mathbf{H} for that item. A substitution gives a power

$$\gamma = 1 - \log d_{\max} / \log r \quad (9)$$

with which the histogram is adjusted in (7).

Figure 3 shows two samples of the edges of \mathbf{H} for two settings of r . For each item, it shows the ratio of “positive” to “negative” edges. A side effect is that at the head, the most popular items are underrepresented in the remainder of \mathbf{H} . This is because the items (or edges) sampled from $\mathcal{M}(\boldsymbol{\pi})$ might already exist in \mathbf{G} , and are discarded and another edge sampled.

5. VARIATIONAL INFERENCE

The approximation $q(\boldsymbol{\theta})$ in (6) is found by maximizing a variational lower bound on the partition function of (5), with

$$\log p(\mathbf{G}) \geq \mathcal{L}[q] = \int q(\boldsymbol{\theta}) \log p(\mathbf{G}, \boldsymbol{\theta}) d\boldsymbol{\theta} + \mathcal{H}[q(\boldsymbol{\theta})] . \quad (10)$$

Here $\mathcal{H}[q]$ is the (continuous) entropy of our choice of q . The expression in (10) is not analytically tractable due to the sigmoids in ℓ_{mn} , which appear in $p(\mathbf{G}, \boldsymbol{\theta})$ in (5), as they are not conjugate with respect to the $q(u_{mk})$'s or any of the other vertex factors. We additionally lower-bound ℓ_{mn} with the logistic or Jaakkola-Jordan bound [8], introducing an additional variational parameter ξ_{mn} on each edge. The logistic bound is

$$\ell \geq e^{g(\mathbf{u}^T \mathbf{v} + b)} \left[\sigma(\xi) e^{-\frac{1}{2}(\mathbf{u}^T \mathbf{v} + b + \xi) - \lambda(\xi)((\mathbf{u}^T \mathbf{v} + b)^2 - \xi^2)} \right]^{g+h(1-g)}, \quad (11)$$

where subscripts m and n that are clear from the context are suppressed. The bound depends on a deterministic function

$\lambda(\xi) \stackrel{\text{def}}{=} \frac{1}{2\xi}[\sigma(\xi) - \frac{1}{2}]$. The substitution of the lower bound in (11) to ℓ_{mn} creates a $p_{\xi}(\mathbf{G}, \boldsymbol{\theta})$ that leaves the bounded likelihood conjugate with respect to its prior. The bound \mathcal{L}_{ξ} ,

$$\mathcal{L}[q] \geq \mathcal{L}_{\xi}[q] = \int q(\boldsymbol{\theta}) \log p_{\xi}(\mathbf{G}, \boldsymbol{\theta}) \, d\boldsymbol{\theta} + \mathcal{H}[q], \quad (12)$$

is therefore explicitly maximized over both the (variational) distribution q and the additional variational parameters $\boldsymbol{\xi} = \{\xi_{mn}\}$.

5.1 Variational updates

The variational updates for the user factors $q(u_{mk})$ are presented in this section. As the model is bilinear, the gradients of \mathcal{L}_{ξ} with respect to the item factors can be set to zero following a similar pattern. To minimize \mathcal{L}_{ξ} with respect to $q(u_{mk})$, one might take functional derivatives $\partial \mathcal{L}_{\xi} / \partial q(u_{mk})$ with respect to each $q(u_{mk})$, and sequentially equate them to zero. This is slow, as each update will require a loop over all the vertex's edges: for the user, K loops over all the items will be required. The vertex factor can alternatively be updated in bulk, by first equating the gradients of \mathcal{L}_{ξ} with respect to a full Gaussian (not factorized) approximation $\tilde{q}(\mathbf{u}_m)$ to zero. The fully factorized $q(u_{mk})$ can then be recovered from the *intermediate* approximation $\tilde{q}(\mathbf{u}_m)$ as those that minimize the Kullback-Leibler divergence $D_{\text{KL}}(\prod_{k=1}^K q(u_{mk}) \| \tilde{q}(\mathbf{u}_m))$: this is achieved when the means of $q(u_{mk})$ match that of $\tilde{q}(\mathbf{u}_m)$, while their *precisions* match the diagonal precision of $\tilde{q}(\mathbf{u}_m)$.

How do we find $\tilde{q}(\mathbf{u}_m)$? The functional derivative $\partial \mathcal{L}_{\xi} / \partial \tilde{q}(\mathbf{u}_m)$ is zero where $\tilde{q}(\mathbf{u}_m)$ has as natural parameters a precision matrix of

$$\mathbf{P}_m = \sum_{n=1}^N \mathbb{E}_q[h_{mn}] \cdot 2\lambda(\xi_{mn}) \cdot \mathbb{E}_q[\mathbf{v}_n \mathbf{v}_n^T] + \mathbb{E}_q[\tau_u] \mathbf{I} \quad (13)$$

and mean-times-precision vector $\boldsymbol{\mu}_m \mathbf{P}_m$, which will be stated in (15). Apart from having to average h_{mn} over $q(\mathbf{H})$, which we cannot do analytically, the update in (13) suffers from having a summation over all N item vertices.

The burden of having to determine a sum over a full item catalogue in (13) can be removed with a clever rearrangement of expectations. As h_{mn} is binary,

$$\begin{aligned} \sum_{n=1}^N \mathbb{E}_q[h_{mn}] f(\mathbf{v}_n) &= \sum_{\mathbf{H}} q(\mathbf{H}) \sum_{n=1}^N h_{mn} f(\mathbf{v}_n) \\ &= \sum_{\mathbf{H}} q(\mathbf{H}) \sum_{n:h_{mn}=1} f(\mathbf{v}_n). \end{aligned} \quad (14)$$

The sum over \mathbf{H} in (14) runs over all 2^{MN} possible instantiations of \mathbf{H} . A rearrangement of (13) therefore allows the updates to appear as a *stochastic average*,

$$\begin{aligned} \mathbf{P}_m &= \mathbb{E}_{q(\mathbf{H})} \left[\sum_{n:h_{mn}=1} 2\lambda(\xi_{mn}) \cdot \mathbb{E}_q[\mathbf{v}_n \mathbf{v}_n^T] + \mathbb{E}_q[\tau_u] \mathbf{I} \right] \\ \boldsymbol{\mu}_m \mathbf{P}_m &= \mathbb{E}_{q(\mathbf{H})} \left[\sum_{n:h_{mn}=1} \left(g_{mn} - \frac{1}{2} \cdots \right. \right. \\ &\quad \left. \left. \cdots - 2\lambda(\xi_{mn}) \cdot \mathbb{E}_q[b_m + b_n] \right) \mathbb{E}_q[\mathbf{v}_n] \right]. \end{aligned} \quad (15)$$

Inside the expectation over $q(\mathbf{H})$, the mean field update in (15) is a quantity specified on the hidden graph \mathbf{H} only, and

not all N plausible edges for the user. We are able to sample graphs from $q(\mathbf{H})$ according to Section 4. Retrospectively, this choice now bears fruit, as the update exists as an average amenable to stochastic gradient descent. We remark, too, that the natural parameters in (15) define the *natural gradients* of the variational objective function [1, 24]. The full natural gradient is periodic in the number of vertices and the updates are component-wise, and convergence with such updates can also be achieved using a stochastic gradient algorithm [12].

There are additional variational parameters at play in (15). For the required edges $h_{mn} = 1$ that connect user m with items n , the values ξ_{mn} that maximize \mathcal{L}_{ξ} or $\mathbb{E}_q[\log p_{\xi}(\mathbf{G}, \boldsymbol{\theta})]$ are each given by

$$\xi_{mn}^2 = \mathbb{E}_q[(\mathbf{u}_m^T \mathbf{v}_n + b_m + b_n)^2], \quad (16)$$

and they are computed and discarded when needed. We take the positive root as ξ_{mn} , and refer the reader to Bishop [2] for a deeper discussion.

Given \mathbf{P}_m and $\boldsymbol{\mu}_m \mathbf{P}_m$ from (15), we have sufficient statistics for $\tilde{q}(\mathbf{u}_m)$, and hence for updating each of the K $q(u_{mk})$'s in bulk. Deriving sufficient statistics for $q(v_{nk})$, $q(b_m)$ and $q(b_n)$ is similar to that presented in (15), and the derivation will not be repeated. Given these, optimization proceeds as follows: At time t , we sample a hidden graph \mathbf{H} , over which the user and item vertex factors are updated. Focussing on user m , let $\mathbf{P}_m^{(t-1)}$ be the (diagonal) precision matrix of the factorized distribution $\prod_{k=1}^K q(u_{mk})$. We then find \mathbf{P}_m in (15), and now the precision matrix of $\tilde{q}(\mathbf{u}_m)$ will be $\mathbf{P}_m^{(t)}$, found through $\mathbf{P}_m^{(t)} = \epsilon_t \mathbf{P}_m + (1 - \epsilon_t) \mathbf{P}_m^{(t-1)}$, where $\epsilon_t \in [0, 1]$. The factors $q(u_{mk})$ are then recovered from the bulk computation of $\tilde{q}(\mathbf{u}_m)$. The mean-times-precision vector of $\tilde{q}(\mathbf{u}_m)$ is given through a similar stochastic update. The series $\{\epsilon_t\}_{t=1}^{\infty}$ should satisfy $\sum_{t=1}^{\infty} \epsilon_t = \infty$ and $\sum_{t=1}^{\infty} \epsilon_t^2 < \infty$, guarding against premature convergence and infinite oscillation around the minimum [22].

Finally, the marginal approximations for the hyperparameters are updated by setting the functional derivatives, say $\partial \mathcal{L}_{\xi} / \partial q(\tau_u)$, to zero. For instance for $q(\tau_u) = \mathcal{G}(\tau_u; \phi_u, \varphi_u)$ the shape ϕ_u and rate φ_u are

$$\begin{aligned} \phi_u &= \alpha + KM/2 \\ \varphi_u &= \beta + \frac{1}{2} \sum_{m=1}^M \mathbb{E}_q[\mathbf{u}_m^T \mathbf{u}_m]. \end{aligned} \quad (17)$$

As $q(u_{mk})$ is dependent on \mathbf{H} , the rate is also stochastically updated as described above.

5.2 Large scale inference

The use of a bipartite graph ensures that variational updates are parallelizable. For instance, by keeping all $q(v_{nk})$, $q(b_n)$ and $q(b_m)$ fixed for the item and user vertices, the gradients $\partial \mathcal{L}_{\xi} / \partial \tilde{q}(\mathbf{u}_m)$, and hence the stochastic updates resulting from (15), have no mutual dependence. Consequently, the loop over user vertex updates $m = 1 \dots M$ is embarrassingly parallel; the same is true for other updates. This will not hold for more general graphs like those of social networks, though, where more involved logic will be required.

Due to the fact that a variational lower bound is optimized for, optimization can also be distributed across multiple machines, *as long as the bound holds*. For example, one might distribute the graph according to item vertices in blocks \mathcal{B}_b , and iteratively optimize one block at a time, or optimize

blocks concurrently (with embarrassingly parallel optimization *inside* the blocks, as discussed earlier). In this example the sparse user-item graph (matrix) \mathbf{G} is distributed such that *all* observations for a set \mathcal{B}_a of items are co-located on the same machine. The natural gradients for the users then distribute across machines, and can be written so that the dependence on the data blocks on various machines separates. When optimizing using the item-wise data block \mathcal{B}_a on one machine, we write \mathbf{P}_m in (15) as

$$\mathbf{P}_m = \mathbb{E}_q(\mathbf{H}) \left[\sum_{\substack{n: h_{mn}=1 \\ n \in \mathcal{B}_a}} 2\lambda(\xi_{mn}) \cdot \mathbb{E}_q[\mathbf{v}_n \mathbf{v}_n^T] \cdots \right. \\ \left. + \sum_{b \neq a} \underbrace{\sum_{\substack{n: h_{mn}=1 \\ n \in \mathcal{B}_b}} 2\lambda(\xi_{mn}) \cdot \mathbb{E}_q[\mathbf{v}_n \mathbf{v}_n^T]}_{\text{block } b\text{'s natural gradient } \mathbf{X}_m^{(b)}; \text{ fixed}} \right] + \mathbb{E}_q[\tau_u] \mathbf{I}. \quad (18)$$

Update (18) defines a thin *message interface* between various machines, where each block has to communicate only its *natural gradients* $\mathbf{X}_m^{(b)}$ —and similar mean-times-precision gradients—to other blocks.¹ In block \mathcal{B}_a we might iterate between updates (18) and full item updates for all $n \in \mathcal{B}_a$, whilst keeping the incoming messages $\mathbf{X}_m^{(b)}$ from other machines fixed. After a few loops over users and items, one can move to the next block. Similarly, different machines can optimize on all the blocks $\{\mathcal{B}_b\}$ in parallel, as long as the natural gradient messages are periodically communicated to other machines. The scheme presented here generalizes to a further subdivision of user vertices into blocks.

6. RESULTS

Given \mathbf{G} , a pivotal task of collaborative filtering is that of accurately predicting the future presence of an edge. This allows online systems to personalize towards a user’s taste by recommending items that the user might like.

The collaborative filtering model in Section 3 explicitly separated the probability of a user considering an item from σ , the probability for the user liking the item. The odds of liking an item depends on our inferred certainty of the user and item parameters,²

$$p(g = 1 | h = 1) \approx \iiint \sigma(\mathbf{u}^T \mathbf{v} + b) q(\mathbf{u}) q(\mathbf{v}) q(b) d\mathbf{u} d\mathbf{v} db \\ \approx \int \sigma(a) \mathcal{N}(a; \mu_a, \sigma_a^2) da \approx \sigma\left(\mu_a / \sqrt{1 + \pi \sigma_a^2 / 8}\right). \quad (19)$$

The random variable a was defined as $a \stackrel{\text{def}}{=} \mathbf{u}^T \mathbf{v} + b$, with its density approximated with its first two moments under q , i.e. $\mu_a \stackrel{\text{def}}{=} \mathbb{E}_q[\mathbf{u}^T \mathbf{v} + b]$ and $\sigma_a^2 \stackrel{\text{def}}{=} \mathbb{E}_q[(\mathbf{u}^T \mathbf{v} + b - \mu_a)^2]$. The final approximation of a logistic Gaussian integral follows from MacKay [13].

¹The division of data to machines will be dictated by the size of M and N ; for $N \ll M$ a user-wise division gives a smaller message interface, as only natural gradients for the items’ updates will be required.

²We suppress subscripts m and n for clarity, and write $q(\mathbf{u})$ for the diagonal Gaussian $\prod_{k=1}^K q(u_{mk})$.




Figure 4: The classification error on \mathbf{G}_{test} , given $h = 1$ (the ground truth is $g = 1$). The full histograms of probabilities $p(g = 1 | h = 1)$ are presented in Figure 5.

6.1 Evaluation

We evaluated our model by removing a test set from the Xbox movies and Netflix (4 and 5 stars) data sets. The degree distributions for these data sets are presented in Figure 1. The training data $\mathbf{G}_{\text{train}}$ was created by randomly removing one edge (or item) for each user from \mathbf{G} ; the removed edges formed the test set.

A core challenge of *real world* collaborative filtering algorithms is to find a balance between popular recommendations and personalized content in a structured form. Based on our experience, a criteria of a good recommender is the ability to suggest non-trivial items that the user will like, and surface less popular items in the tail of the item catalogue. In the evaluations we highlight this by grouping results according to item popularity in Figure 6, for example.

Two evaluations are discussed below. Firstly, given that an item is presented to a user with $h_{mn} = 1$, we are interested in the classifying $g_{mn} \rightarrow \{0, 1\}$. This is one of the key contributions that our model brings to the table. As far as we are aware, there are no other algorithms that isolate $p(\text{like})$ in this way. To be able to draw a comparison with a known state-of-the-art algorithm, we consider various forms of a rank-based metric in a second evaluation.

In the tests below, $K = 20$ latent dimensions were used. The user biases were clamped at zero, as $q(\mathbf{H})$ was defined to give balanced samples for each user. The rate and shape parameters of the hyperprior were set to $\alpha = \beta = 0.01$, giving a hyperprior on the τ ’s with mean 1 and a flexible variance of 100. The means of the hyperparameter posterior estimates were $\mathbb{E}[\tau_{b_v}] = 0.4$, $\mathbb{E}[\tau_v] = 3.5$, and $\mathbb{E}[\tau_u] = 2.0$. When rounded to one decimal place, these values were similar on both the Netflix (4 and 5 stars) dataset and the Xbox Movies dataset.

6.1.1 The “like” probability

The classification error on the held-out data converges to a stable value as users view between ten and twenty items. Its plot is presented in Figure 4, and has a natural interpre-

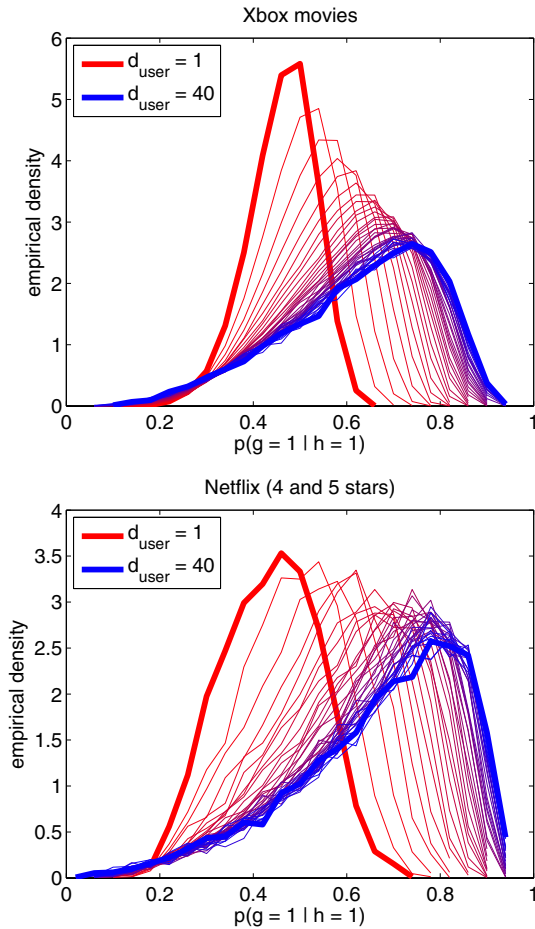


Figure 5: The distribution of $p(g_{mn} = 1 | h_{mn} = 1)$ on the held out items in the evaluation, sliced incrementally according to users connected to $d_{\text{user}} = 1$ to 40 items. The ground truth is $g_{mn} = 1$.

tation. Conventional wisdom dictates that the error rates for explicit ratings-based recommendation systems are typically in the order of 20% of the ratings range. For Netflix’s five-star ratings datasets, this error is around one star [11], while an error of around 20 points in the 0-100 scale of the Yahoo! Music dataset is usual [3]. The 16-19% classification error in Figure 4 is therefore in line with the signal to noise ratio in well known explicit ratings datasets. When users viewed only one item, the bulk of the predictive probability mass $p(g = 1 | h = 1)$ is centered around 50%, slightly skewed to being less certain. This is illustrated in Figure 5. As users view more items, the bulk of the predictive probability skews towards being more certain³.

The probability $p(g = 1 | h = 1)$ is useful in presenting a user with interesting recommendations, as it is agnostic

³A property of a good probabilistic classification system is that it produces an exact *calibration plot*. For example, we expect 10% of edges to be misclassified for the slice of edges that are predicted with $p(g = 1 | h = 1) = 10\%$. The calibration plot requires a ground truth negative class $g = 0$, which is *latent* in our case. Figures 4 and 5 aim to present an equivalent to a calibration plot.

to each item’s popularity. It is therefore possible to define a *utility function* that trades this quantity off with an item’s popularity, effectively giving a knob to emphasize exploration or exploitation. Such a utility can be optimized through A/B tests in a fighting framework, but is beyond the scope of this paper.

6.1.2 Average rank

We turn to a ranking task to draw a comparison against known work, as we are unaware of other algorithms that isolate $p(\text{like})$. On seeing $\mathbf{G}_{\text{train}}$, the absent edges (where $g_{mn} = 0$) are ranked for each user m . The ranking is based on various scores s_{mn} :

like the odds of a user liking an item, namely $s_{mn} = p(g_{mn} = 1 | h_{mn} = 1)$ as approximated in (19);

popularity $s_{mn} = \pi_n$;

popularity \times like the odds of a user considering *and* liking an item, namely $s_{mn} = \pi_n p(g_{mn} = 1 | h_{mn} = 1)$.

We evaluated models for the two settings of r in (9); a sample from \mathbf{H} for each was shown in Figure 3.

Our metric is computed as follows: If item n' was removed, the rank score counts the position of n' in an ordered prediction list

$$S_{\text{rank}}(m, n') \stackrel{\text{def}}{=} \frac{\sum_{n: g_{mn}=0} \mathbb{I}[s_{mn'} > s_{mn}]}{\sum_{n: g_{mn}=0} 1}. \quad (20)$$

Random guessing would give $S = 0.5$, while $S = 1$ places the held-out item at the head of the list.

As a benchmark, we use the Bayesian Personalized Ranking (BPR) model of Rendle et al. [6, 21]. It has shown state of the art performance on ranking metrics against methods ranging from singular value decompositions and nearest neighbours to weighed regularized matrix factorization [19].

BPR was also used as a key component in many of the leading solutions for the second track of the KDD-Cup’11 competition [4]. The competition was designed to capture the ability of models to personalize recommendations that “fit” specific users regardless of an item’s popularity. In that setting, BPR was trained with missing items sampled with probabilities proportional to their popularity as described in [5]. We therefore implemented and trained two BPR models:

BPR-uniform with missing items sampled uniformly;

BPR-popularity with missing items sampled proportional to their popularity.

These two models capture two different aspects of recommender systems. **BPR-uniform** is optimized to learn a user-wise ranking of items, where the objective function specifies that items that are liked (i.e. $g_{mn} = 1$) should be ranked above missing items (i.e. $g_{mn} = 0$).

The metric in (20) follows [21]. Because **BPR-uniform** directly optimizes this metric, it should come as no surprise that it will perform better than methods that do not optimize it directly (see Figure 6). However, meaningful insights can still be gleaned from the comparison. **BPR-popularity** is aimed at ranking observed “liked” items above other *popular* items that are missing from the user’s history. While two BPR models are required to capture these two different aspects of recommendations, our generative model captures both of these aspects in a structured manner.

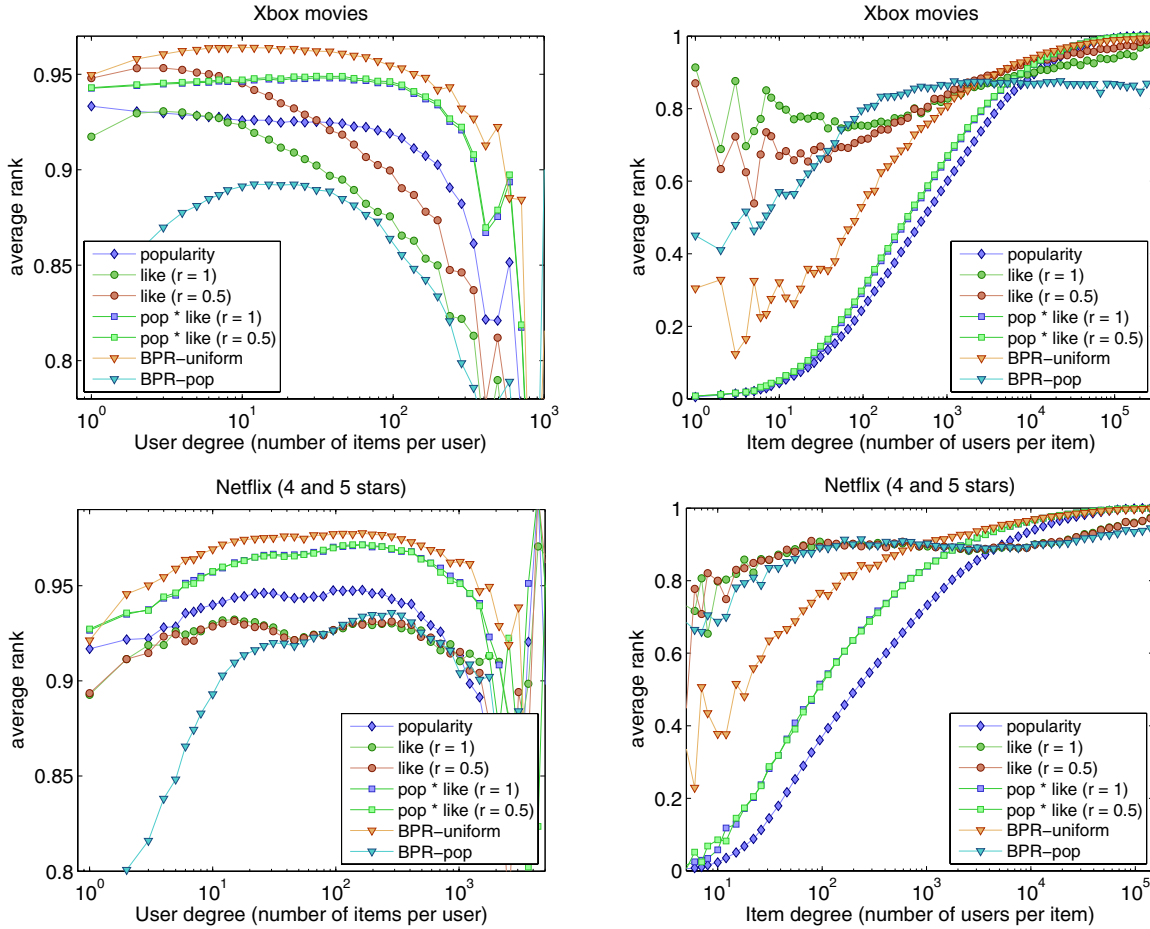


Figure 6: The rank $S_{\text{rank}}(m, n)$ in (20), averaged over users (*left*) and items (*right*), grouped logarithmically by their degrees. The *top* evaluation is on the Xbox movies sample, while the *bottom* evaluations are on the Netflix set, as given in Figure 1.

Figure 6 illustrates the mean rank scores, grouped logarithmically by user and item degrees. In the plots that are grouped by user degrees, we see improved results for algorithms that prefer popularity, i.e. **popularity** \times **like** and **BPR-uniform**. This is explained by the dominance of popularity biases in both datasets. As expected, **BPR-uniform** show best results as it optimizes the task at hand directly. The estimates for users with an order of 10^3 to 10^4 degrees are noisy as the data is very sparse (see Figure 1). However, when looking at the per item breakdown, we learn that **BPR-uniform** and the **popularity** \times **like** models perform poorly on less popular items and their superior results are based on recommendations from the short head of the popular items. When it comes to recommending from the long tail of the less familiar items, the **like** models show best results, with **BPR-popularity** just behind. These trends are consistent on both datasets.

The distribution of the ranks over all users (and items) is heavy-tailed, and whilst the average is often reported, the median is much higher than the average reported in Figure 6. Figure 7 shows the *error bars* using the percentiles of the rank scores for tests **like** and **popularity** \times **like** for $r = \frac{1}{2}$. The rank variance decreases as users view a few movies,

but increases for heavy users which are harder to model. When popularity is included in the ranking, the error bars get tighter for heavy users, which implies that these users' lists are mostly governed by popularity patterns.

7. CONCLUSIONS

Random graphs can be leveraged to predict the presence of edges in a collaborative filtering model. In this paper we showed how to incorporate such graphs in an inference procedure by rewriting a variational Bayes algorithm in terms of random graph samples. As a result, we were able to explicitly extract a “like” probability that is largely agnostic to the popularity of items. The use of a bipartite graph, central to this exposition, is not a hindrance, as user-user interactions in a general network can be similarity modelled with $\sigma(\mathbf{u}_m^T \mathbf{u}_{m'})$. While scalable parallel inference is not immediately obvious, we believe this to be a worthwhile pursuit. By employing the same machinery on general graphs, one should be able to model connections in social or other similar networks.

The use of a Bayesian graphical model makes it easy to adapt the model to incorporate richer feedback signals. Similarly, both structured and unstructured meta-data can be

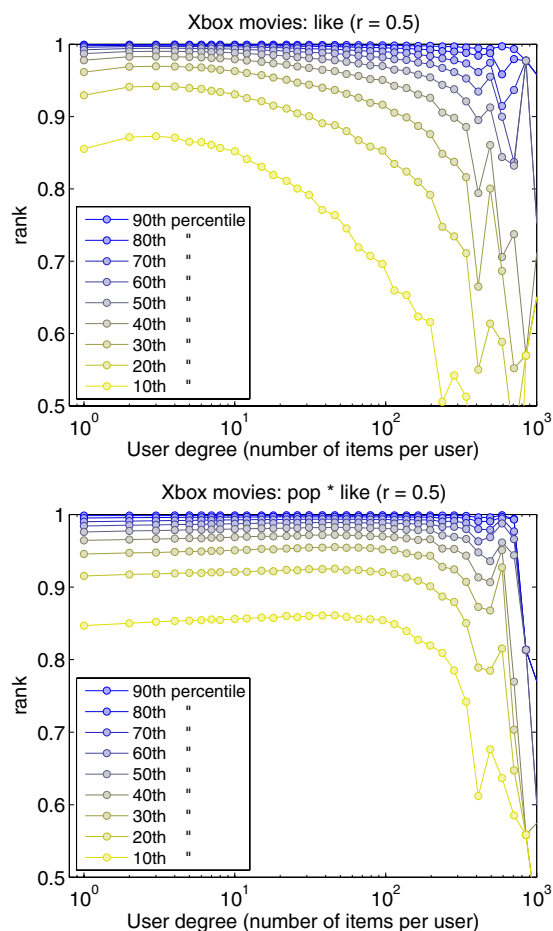


Figure 7: Error bars on the rank tests. The median is much higher than the average rank reported in Figure 6.

plugged into the graphical model. The hidden graph \mathbf{H} may also be partly observed, for example from system logs. In that case some true negatives exist. Alternatively, we may know *a priori* when a user could never have considered an item, fixing some h at zero. In both these scenarios the process of drawing random hidden graphs \mathbf{H} can be adjusted accordingly. For the sake of clarity, none of these enhancements were included in this paper.

8. ACKNOWLEDGMENTS

The authors are indebted to Nir Nice, Shahar Keren, and Shimon Shlevich for their invaluable input, management, and stellar engineering skills.

9. REFERENCES

- [1] S. Amari. Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276, 1998.
- [2] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [3] G. Dror, N. Koenigstein, and Y. Koren. Yahoo! music recommendations: Modeling music ratings with temporal dynamics and item taxonomy. In *Proc. 5th ACM Conference on Recommender Systems*, 2011.
- [4] G. Dror, N. Koenigstein, Y. Koren, and M. Weimer. The Yahoo! music dataset and KDD-Cup’11. *Journal Of Machine Learning Research*, 18:3–18, 2012.
- [5] Z. Gantner, L. Drumond, C. Freudenthaler, and L. Schmidt-Thieme. Personalized ranking for non-uniformly sampled items. *Journal of Machine Learning Research*, 18:231–247, 2011.
- [6] Z. Gantner, S. Rendle, C. Freudenthaler, and L. Schmidt-Thieme. MyMediaLite: A free recommender system library. In *5th ACM International Conference on Recommender Systems*, 2011.
- [7] Y. F. Hu, Y. Koren, and C. Volinsky. Collaborative filtering for implicit feedback datasets. In *IEEE International Conference on Data Mining*, 2008.
- [8] T. Jaakkola and M. Jordan. A variational approach to Bayesian logistic regression problems and their extensions. In *Artificial Intelligence and Statistics*, 1996.
- [9] M. Jordan, Z. Ghahramani, T. Jaakkola, and L. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37:183–233, 1999.
- [10] N. Koenigstein, N. Nice, U. Paquet, and N. Schleyen. The Xbox recommender system. In *Proc. 6th ACM Conference on Recommender Systems*, 2012.
- [11] Y. Koren. The BellKor solution to the Netflix Grand Prize. 2009.
- [12] H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, 2003.
- [13] D. J. C. MacKay. The evidence framework applied to classification networks. *Neural Computation*, 4(5):698–714, 1992.
- [14] R. M. Neal. Probabilistic inference using Markov chain Monte Carlo methods. Technical Report CRG-TR-93-1, Dept. of Computer Science, University of Toronto, 1993.
- [15] M. E. J. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Phys. Rev. E*, 64:026118, 2001.
- [16] M. E. J. Newman, D. J. Watts, and S. H. Strogatz. Random graph models of social networks. *Proc. Natl. Acad. Sci. USA*, 99:2566–2572, 2002.
- [17] K. Palla, D. A. Knowles, and Z. Ghahramani. An infinite latent attribute model for network data. In *29th International Conference on Machine Learning*, 2012.
- [18] R. Pan and M. Scholz. Mind the gaps: Weighting the unknown in large-scale one-class collaborative filtering. In *KDD*, pages 667–675, 2009.
- [19] R. Pan, Y. Zhou, B. Cao, N. Liu, R. Lukose, M. Scholz, and Q. Yang. One-class collaborative filtering. In *IEEE International Conference on Data Mining*, pages 502–511, 2008.
- [20] U. Paquet, B. Thomson, and O. Winther. A hierarchical model for ordinal matrix factorization. *Statistics and Computing*, 22(4):945–957, 2012.
- [21] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. BPR: Bayesian personalized ranking from implicit feedback. In *Uncertainty in Artificial Intelligence*, pages 452–461, 2009.

- [22] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.
- [23] R. Salakhutdinov and A. Mnih. Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In *Proceedings of the 25th International Conference on Machine Learning*, pages 880–887, 2008.
- [24] M. Sato. Online model selection based on the variational Bayes. *Neural Computation*, 13(7):1649–1681, 2001.
- [25] D. Stern, R. Herbrich, and T. Graepel. Matchbox: Large scale Bayesian recommendations. In *International World Wide Web Conference*, 2009.
- [26] M. P. H. Stumpf and C. Wiuf. Sampling properties of random graphs: The degree distribution. *Phys. Rev. E*, 72:036118, 2005.