Online Learning for Distributed and Personal Recommendations A Fair Approach

Martin Tegnér[1 2] <Martin.tegner@ingka.ikea.com>

[1]IKEA Retail Business [2]Oxford-Man Institute, University of Oxford

1 Recommendations

A **recommendation system** presents items (`80257360`) to a user. The aim is to make *relevant* suggestions, for a product purchase, friend request or movie watch. **Supervised learning** trains a map that recommends an item y from user inputs x.This can be seen as a **graph**, where the model $\mathcal{M}$ connects all items based on x. The **shortest edge** from the current item yields a recommendation.

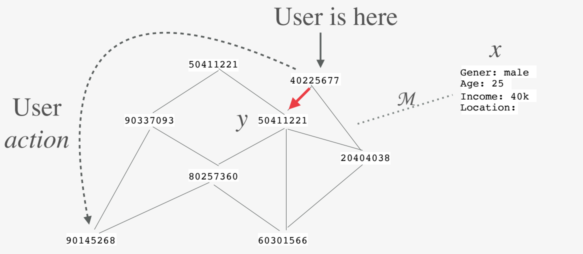

Personalised models are expressive but **data hungry** to be able to generalise: informative features encoded in x typically rely on *personal* and *private* information, and model training on centralised storage of user data. On the other hand, **contextual** graphs can be constructed from item features, or inferred from sequences of item-views alone.

This work presents a method **not** relying on private user information and data collection, still providing **personalised** recommendations.

2 Probabilistic Mixture Model

At the base, we assume a **collection of graphs**, representing different (contextual/inferred) *activities*, i.e. relevant ways of relating items for recommendations. We model each with a **Markov chain**. A high (low) probability $p(y_j|y_i) = \mathbf{A}_{ij}$ yields a strong (weak) relation between item y_i and y_j, where $\mathbf{A}$ is the chain's transition matrix.

For a model both expressive of different activities and adaptive to individual behaviours, we use a latent *categorical* variable $C\sim\mathrm{Cat}(\pi)$ to **mix** the Markov chains at user level: $p(C=k) = \pi_k$ is the probability of k^{th} activity with items distributed by Markov chain $\mathbf{A}^{(k)}$.

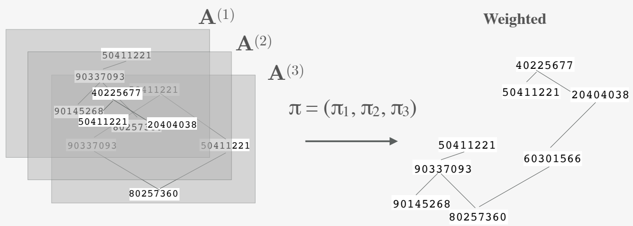

Let $\mathbf{y}_t = (\ldots, y_{t\text{-}2}, y_{t\text{-}1}, y_t)$ be user-viewed items at iteration t. If $\mathcal{L}(\mathbf{y}_t)_k$ is likelihood under chain k, we have that the posterior $C|\,\mathbf{y}_t$ remains a $\mathrm{Cat}(\pi(t))$ with probs $\pi(t)_k = Z^{-1}\mathcal{L}(\mathbf{y}_t)_k\,\pi_k$. For new data $\mathbf{y}_{\text{new}}$, Algorithm 1 updates the posterior in an *online* and *distributed* way:

Algorithm 1 Online posterior update

Input: matrices $\{A^{(k)}\}_{k=1}^{L}$, current mixing parameter π and data $\mathbf{y}_{\text{new}}$ observed *after* previous update
for $k = 1$ **to** L **do**
  compute $\mathcal{L}(\mathbf{y}_{\text{new}})_k$ from $A^{(k)}$
  update $\pi_k \leftarrow \mathcal{L}(\mathbf{y}_{\text{new}})_k \pi_k$
end for
normalise:
$Z = \sum_k \pi_k$
$\pi_k \leftarrow \pi_k / Z$
Output: updated π (last value of $\mathbf{y}_{\text{new}}$, to include in data of next update)

(1) only $\mathbf{y}_{\text{new}}$ is needed for computing π of the new posterior, and (2) inference is *distributed* using user data locally, at the user's device. Further, the *posterior predictive* $p(y_{\text{rec}}|\mathbf{y}_{\text{new}})$ for recommendations is distributed.

3 Experiments

We construct three activities for 300 home-furnishing products:

- $\mathbf{A}^{(1)}$ a *functional* context, according to (i) kitchen (100 items), (ii) living-room (100), and (iii) bathroom (100). See figure below.
- $\mathbf{A}^{(2)}$ an *inspirational* context, extracted from product images designed by home-furnishing specialists.
- $\mathbf{A}^{(3)}$ an *add-to-cart* activity, with $\mathbf{A}^{(3)}$ estimated from item-sequences of order data.

Visualisation (right): A blue dot is a 2-dim t-SNE representation of an item according to $\mathbf{A}^{(1)}$.

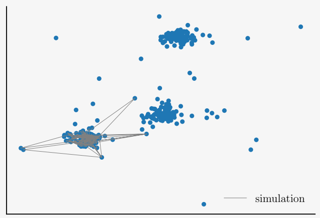

The functional shopper. We simulate 100 items from $\mathbf{A}^{(1)}$ and feed the algorithm with two items for each posterior update.

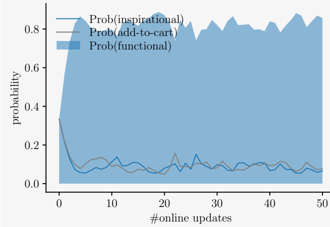

Resulting $\pi(t)$ (left). Weights stabilise after a few iterations, with a high probability for the generating functional activity. t-SNE visualisation of transitions generated by $\pi(t)$ (below) along with posterior recommendations

Similarly, posterior (below) from a **random shopper**.