A Deep Latent Recommender System based on User Ratings and Reviews

Dingge LIANG
(Joint work with M. Corneli, C. Bouveyron and P. Latouche)

Université Côte d'Azur, INRIA MaasAI

Introduction: towards a topic-aware recommender system

Consider a dataset with M users scoring/reviewing P products, encoded by two matrices:

- Y is an ordinal matrix in $\mathbb{N}^{M \times P}$, with Y_{ij} the rating of the object j by user i.
- W is a document-term matrix (DTM), with $W^{(i,j)}$ encoding a review about object j by user i.

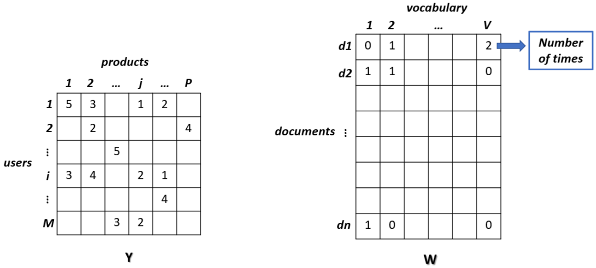

(a) Ordinal matrix (b) Document-term matrix

Figure: Two encoded matrices.

The generative process and VAE inference

We assume the following generative process for rating and review:

- The rating Y_{ij} is:

$$Y_{ij} = \langle R_i, C_j \rangle + \epsilon_{ij}, \tag{1}$$

  where $R_i \sim \mathcal{N}(0, I_D)$, $C_j \sim \mathcal{N}(0, I_D)$, $\varepsilon_{ij} \sim \mathcal{N}(0, \eta^2)$.

- For the reviews, we assume that:

$$p(W^{(i,j)}|\theta_{ij}) \sim \text{Multinomial}(L_{ij}, \beta\theta_{ij}), \tag{2}$$

  where

  - L_{ij} is the number of words in $W^{(i,j)}$,
  - $\beta \in \mathbb{R}^{V\times K}$ is the matrix whose entry β_{vk} is the probability that word v occurs in topic k,
  - $\theta_{ij} = \sigma(f_\gamma(R_i, C_j))$ is the topic proportion, where $f_\gamma : \mathbb{R}^{2D} \to \mathbb{R}^K$ is a continuous function approximated by a neural network parametrized by γ, $\sigma(\cdot)$ denotes the Softmax function.

A Variational auto-encoder is used for the inference:

$$\log p(Y, W|\eta^2, \gamma, \beta) \geq \mathbb{E}_{q(R,C)}\left[\log \frac{p(Y,W,R,C|\eta^2,\gamma,\beta)}{q(R,C)}\right],$$

where

$$q(R_i) = g(\mu_i^R := h_{1,\phi}(Y_i, W^{(i,\cdot)}), S_i^R := h_{2,\phi}(Y_i, W^{(i,\cdot)})),$$
$$q(C_j) = g(\mu_j^C := l_{1,\iota}(Y^j, W^{(\cdot,j)}), S_j^C := l_{2,\iota}(Y^j, W^{(\cdot,j)})).$$

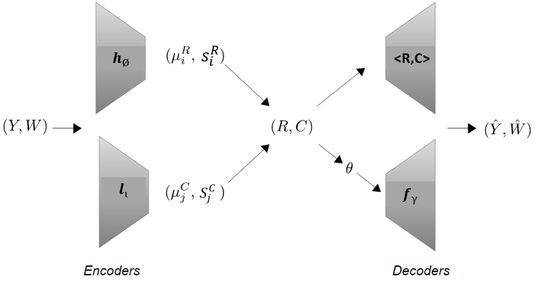

Figure: A deep learning view of deepLTRS.

Numerical experiments

- Benchmark on simulated data:

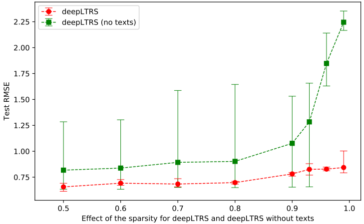

Figure: Comparison of deepLTRS with and without texts.

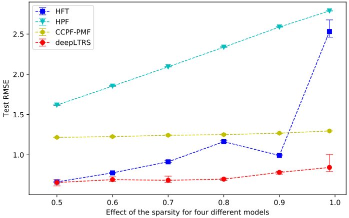

Figure: Test RMSE of models: HFT, HPF, CCPF and deepLTRS with different sparsity level on simulated data.

- Amazon Fine Food data:

Table: Test RMSE on Amazon Fine Food data.

Model	Run 1	Run 2	Run 3	Runt 4	Run 5	Average
HFT	1.4241	1.5327	1.4737	1.4228	1.3850	1.4477 ($\pm$0.0510)
HPF	2.9486	2.9682	2.9311	2.9428	2.9734	2.9528 ($\pm$0.0158)
CCPF-PMF	1.2695	1.2964	1.3035	1.2923	1.2950	1.2913 ($\pm$0.0115)
deepLTRS	**1.1364**	**1.2595**	**1.2445**	**1.1710**	**1.2475**	**1.2518** ($\pm$0.0489)