Poisson Factorization Models for Spatiotemporal Retrieval

Eliezer de Souza da Silva
NTNU – Norwegian University of Science and Technology
Trondheim, Norway
eliezer.souza.silva@ntnu.no

Dirk Ahlers
NTNU – Norwegian University of Science and Technology
Trondheim, Norway
dirk.ahlers@ntnu.no

ABSTRACT

New retrieval models promise deeper integration of multiple features and sources of information. The inclusion of thematic and location features in a joint factorization model allows location to be modeled as a first-class feature and can improve a range of tasks in geographic information retrieval and recommendation. In this position paper, we describe these factorization models and how they can be useful for corpus and user need understanding and further GIR use cases. We argue that using joint factorization models can be a powerful tool in the integration of complex features and relationships present in many GIR data sources and applications.

CCS CONCEPTS

• **Information systems → Spatial-temporal systems**; *Social recommendation*; *Personalization*; • **Computing methodologies → Factorization methods**;

KEYWORDS

Factorization, Latent Factors, modeling, GIR, retrieval

ACM Reference Format:

Eliezer de Souza da Silva and Dirk Ahlers. 2017. Poisson Factorization Models for Spatiotemporal Retrieval. In *GIR'17: 11th Workshop on Geographic Information Retrieval, November 30-December 1, 2017, Heidelberg, Germany*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3155902.3155912

1 INTRODUCTION

Geographic Information Retrieval aims to include the spatial dimension deep into retrieval systems to use location as a first-class feature and to use a deeper understanding of its specific characteristics to deliver better search experiences. In particular, it aims to understand and query both thematic and geographical scopes of resources and match them to user information needs.

New feature learning models are becoming able to understand implicit structure of multiple sources of semi- and unstructured data sources, for example leveraging prediction of users' information need with understanding of topic structures, community structures in social networks, spatio-temporal clustering and other high-level features. A lot of work on these models comes from recommender systems, dealing with user interactions and user preferences. New work towards queryless predictive assistants starts to blur the boundaries between RecSys and IR. One such joint topic is personalized models (for example personalized learning-to-rank), where we can learn user-specific factors to influence query results. We may further infer the user information need through contextual and historical information without an explicit query, as for example in zero-query systems [3].

We argue that for Poisson factorization models, the joint modeling of location and text can mutually benefit the integrated features, it is easier applied, and can bring additional benefits through the embedding of location and text in the same model at the same abstraction level. For count and implicit user data, using a Poisson likelihood is an improvement over models that rely on Gaussian-based likelihood [6]. We will briefly explain the background of the new framework of models and discuss their applicability to some GIR use cases by including location as a high-level factor.

2 BACKGROUND

New approaches coming from recommender systems and personalized search such as factorization models (matrix factorization, factorization machines), embeddings techniques and personalized learning-to-rank carry a large potential to understand and enrich spatial data and match it to user information needs.

For example, early work has already looked at more generic inclusion of location into recommendation models through additions to models, to recommend indirect locations [7]. Another angle from latent feature analysis is location-aware topic modelling, i.e., finding implicit locations for entity names [4, 8], again through extensions or additions for location to existing topic models. Preprocessing is needed to incorporate the location semantics. [9] presents a cross-model learning to understand urban dynamics and makes the good point that naive implementations will break and we still need an understanding of geospatial semantics to properly build, e.g., similarity measures within the model. With a wider range of more powerful models, we do not need to craft extensions of LDA, but can directly integrate location and semantics into the main model. Embedding models have been proposed for personalized search [2], jointly learning hierarchical user/item past interactions and query embeddings. We envision that the idea of joint learning embeddings can be extended to different domains, including the embedding of geo-spatial information.

Poisson factorization models for recommendation are a family of related models with Poisson distributed likelihood and Gamma distributed latent factors, useful for modeling implicit and count data together with side information from social networks or document topic models into a joint recommendation model [5, 6]. In particular, topic modeling with Poisson factorization can be achieved

GIR'17, November 30-December 1, 2017, Heidelberg, Germany

ACM ISBN 978-1-4503-5338-0/17/11…$15.00
https://doi.org/10.1145/3155902.3155912

by factorizing a word-document count data matrix and achieves results similar to LDA, but using an easier modeling.

In general, the strategy to add contextual information is to link different contexts and location-related information using shared latent variables between the context-specific observations. Some constraints might be included in the model via prior distributions (for example, hierarchy of location factors). One approach to integrate location and textual information is to create a spatial grid over the geographic space and model the document, location, and text (document d at location l using word w) interaction counts as a combination of latent gamma variables and Poisson likelihood, creating a joint factorization of document-word observations $A_{d,w} \sim \text{Poisson}(\sum_k D_{dk} L_{lk})$ and document-location $B_{d,l} \sim \text{Poisson}(\sum_k D_{dk} L_{lk})$ counts, where $D_{dk}, L_{lk}, W_{wk} \sim \text{Gamma}(a, b)$ are document, location and words latent factors. Note that the latent gamma factor D_{dk} is shared, which means that we are creating a particular instance of collective matrix-tensor factorization model. Fig. 1 shows this specific instance of how we can arrange those factors, although many other combinations (e.g. document-word-location-user) can be easily designed using this approach.

It is not always straight-forward to select which factors are important to model as underlying dynamics between people, documents, spatio-temporal activity patterns, and information needs. Instead of defining features a priori as relevant, it is important to create a set of possible tools and models with the inherent ability to include and exclude factors, or weight the distinct relevance of different factors. A general easier applicability also makes it easier to model more implicit features, have the option to explicitly model features (e.g. count data: term frequency, click data, etc.), and provides a principled and more formal way for feature combination.

3 USE CASES

Our application scenarios run along the themes of model, understand, and predict, flowing from user information needs. General applications include the standard recommendation of locations, but also the enrichment of document or item descriptions, the extraction of salient features or locations through latent modeling, query understanding, implicit clustering and semantic annotations, and the ranking of items based on feature combinations, including improved learning to rank methods. This relates strongly to information needs, which can be better understood and extracted from a corpus or user behaviour without prior knowledge.

A major use case is the extraction and understanding of spatio-temporal patterns of implicit feedback features, for example learning the intersection between location, time, and duration or click-through rate. We may want to know when and where a user is more likely to interact with a given type of information. For example, to optimize an indoor product recommendation of items in an indoor space such as shopping malls or museums.

Users can explore the geographical information in different ways, for example by learning topics from textual data in different granularity levels [1], so the user can explore textual data to understand a city, a neighborhood, or a building. It also supports exploration of different locations with a similar topical profile – topic-aware geographical similarity. Path recommendation can integrate specific user interest, either learned by using implicit data or a query; for

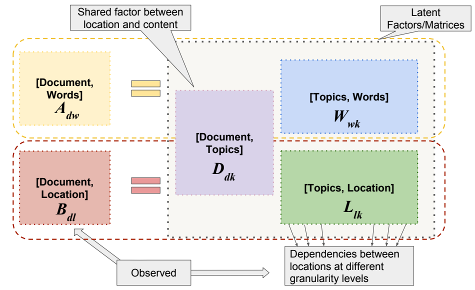

Figure 1: Joint factorization model of document topics and location topics with shared factors and dependencies

example a neighborhood tour that includes nice food and lively music. The result path equally optimizes the distance between topical information of consecutive POIs and the user interest.

The methods allow to use corpus parts to understand other factors, by enabling multi-relational learning. It can be applied on a pure textual corpus with implicit or explicit location, but also on user interactions, social network traces including time, explicit location, text, social relations, and on any combinations of these.

4 CONCLUSION

The models we have described here were originally designed to work on recommender-type data sources, such as users connected to items, locations, and time. However, the integrated nature of the models make it possible to also use them for the common document–topic–location relations and also for other combinations of features. This opens the possibility to work in a framework capable of modeling multi-relational data deeply connected with geo-location data in the same semantic space and to easily include additional factors such as social media, location traces, etc.

We could already show that this model is superior in terms of recall and training time complexity [5] for recommendations and it shows promising improvements for easier modeling as well as count and implicit data (such as user feedback), and sparse observations. We are applying these factorization models in our future work to better understand how they behave on real-world location data and user traces, and which queries will benefit in particular.

REFERENCES

[1] Dirk Ahlers. 2015. Granularity as a Qualitative Concept for GIR *(GIR'15)*. ACM.
[2] Qingyao Ai, Yongfeng Zhang, Keping Bi, Xu Chen, and W. Bruce Croft. 2017. Learning a Hierarchical Embedding Model for Personalized Product Search *(SIGIR '17)*. ACM, 645–654.
[3] Jan R. Benetka, Krisztian Balog, and Kjetil Nørvåg. 2017. Anticipating Information Needs Based on Check-in Activity *(WSDM '17)*. ACM, 41–50.
[4] Da Cao, Liqiang Nie, Xiangnan He, Xiaochi Wei, Shunzhi Zhu, and Tat-Seng Chua. 2017. Embedding Factorization Models for Jointly Recommending Items and User Generated Lists *(SIGIR '17)*. ACM, 585–594.
[5] Eliezer de Souza da Silva, Helge Langseth, and Heri Ramampiaro. 2017. Content-Based Social Recommendation with Poisson Matrix Factorization. In *ECML-PKDD*.
[6] Prem Gopalan, Laurent Charlin, and David M. Blei. 2014. Content-based recommendations with Poisson factorization. In *NIPS 2014*. 3176–3184.
[7] André Sabino and Armanda Rodrigues. 2014. Indirect Location Recommendation *(GIR '14)*. ACM.
[8] Chong Wang, Jinggang Wang, Xing Xie, and Wei-Ying Ma. 2007. Mining Geographic Knowledge Using Location Aware Topic Model *(GIR '07)*. ACM.
[9] Chao Zhang, Keyang Zhang, Quan Yuan, Haoruo Peng, Yu Zheng, Tim Hanratty, Shaowen Wang, and Jiawei Han. 2017. Regions, Periods, Activities: Uncovering Urban Dynamics via Cross-Modal Representation Learning *(WWW '17)*. 361–370.