Master Internship 2022

Combatting Unemployment: Data, Recommendations and Fairness

SUPERVISORS: Michèle Sebag, Philippe Caillou, Armand Lacombe, Guillaume Bied
{sebag,caillou,bied} @lri.fr, Armand.Lacombe@polytechnique.edu
LAB: TAU – CNRS – INRIA – LISN, U. Paris-Saclay

ABSTRACT :
One of the causes for unemployment is the so-called "frictional information", the fact that retrieving the job ad or the CV most appropriate to your current needs requires (lots of) time and knowledge.

Recommender systems (RS), at the core of the e-commerce, handle a similar retrieval task: they allow the customer to navigate in a virtually infinite catalog and propose expectedly appropriate items, that is, items similar to what the customer (or someone similar to the customer) bought in the past [1].

Classical collaborative filtering proceeds by exploiting the description of users $X = \{x_i, i = 1 \ldots m\}$, of items $Y = \{y_j, j = 1 \ldots n\}$ and the (m, n) scoring matrix M, with $M(x_i, y_j) = 1$ iff x_i selected y_j. This matrix is embarrassingly sparse (one selects a very tiny fraction of the whole set of items) and the goal consists in estimating all missing entries (would x_i select y_k ?)

In many domains, it does not harm to recommend the same item to many people (think of block-busters). In the domain of human resources however, e.g. job market, it isn't a good idea to recommend all job seekers the same job ad as only one will get the job; it's a waste of time and energy.

This issue, referred to as *congestion*, is is even more problematic as the number of job seekers is about 8x the number of job ads. In practice, it turns out that 75% of the job ads do not appear in the top recommendations to the job seekers.

The goal of the internship is to investigate how to handle this issue.
A first direction is to consider Optimal transport (OT) [2,3], and to try to transport the whole population of job seekers to the whole population of job ads. Formally, assuming some cost matrix C (with $C(x, y)$ denoting the opportunity of transporting x onto y), OT aims to build a (probabilistic) mapping from a set of x_i onto a set of y_j (with $\gamma(x_i, y_j)$ denoting the probability of mapping x_i onto y_j) in such a way that it yields maximal expected opportunities (i.e. it maximizes $\sum_{i,j} \gamma(x_i, y_j) C(x_i, y_j)$).

Along this line, two options will be investigated:
• How to define the opportunity $C_{i,j}$ based on the recommendation score ? Can we propose an end-to-end learning of the recommendation score and opportunity ? (very daring!)
• How to define a new recommendation policy based on the recommendation scores and the estimated congestion (presenting a mix of popular and non-popular job offers).

THE INTERNSHIP :
The goal is to allocate the X population to the Y population in a "globally optimal" manner, in the sense that each user will be recommended relevant items while ensuring that the same item is recommended to few users.

The student will i) propose performance indicators relevant to congestion-avoiding recommendation; ii) design and implement an appropriate recommendation algorithm (enforcing some trade-off between high accuracy and low congestion, in particular, ensuring that the recommended people do not form a sub-distribution of the whole people distribution [4]); iii) if time permits, conduct experiments on real-world data; the internship will be conducted in the context of a national hiring agency and data are available.

The internship requires excellent theoretical and algorithmic skills (programming environment in Python).

Bibliography

[1] *Recommender Systems*, Aggarwal, Springer 2016.
[2] *Computational Optimal Transport*, Peyré & Cuturi, Foundations and Trends in Machine Learning 11(5-6), 2019
[3] *CO-Optimal Transport*, Redko et al., ArXiv 2020
[4] *A Kernel Two-Sample Test*, Arthur Gretton et al., JMLR 2012.