RECOMMENDER SYSTEM IN A NON-STATIONARY CONTEXT: RECOMMENDING JOB ADS IN PANDEMIC TIMES

G. BIED, S. NATHAN, E. PERENNES, V. ALFONSO, P. CAILLOU, B. CRÉPON, C. GAILLAC, M. SEBAG

MOTIVATIONS

- Goal: improve matching between job seekers and job ads
- Collaboration with the French Public Employment Service
- Formalized as a recommendation problem, learning from employment contracts signed in the past
  - Extreme match matrix sparsity
  - Cold start
  - Low signal to noise ratio (disadvantaged populations)
  - Scalability requirements
- *Contribution*: VADORE, a cascading deep recommender system

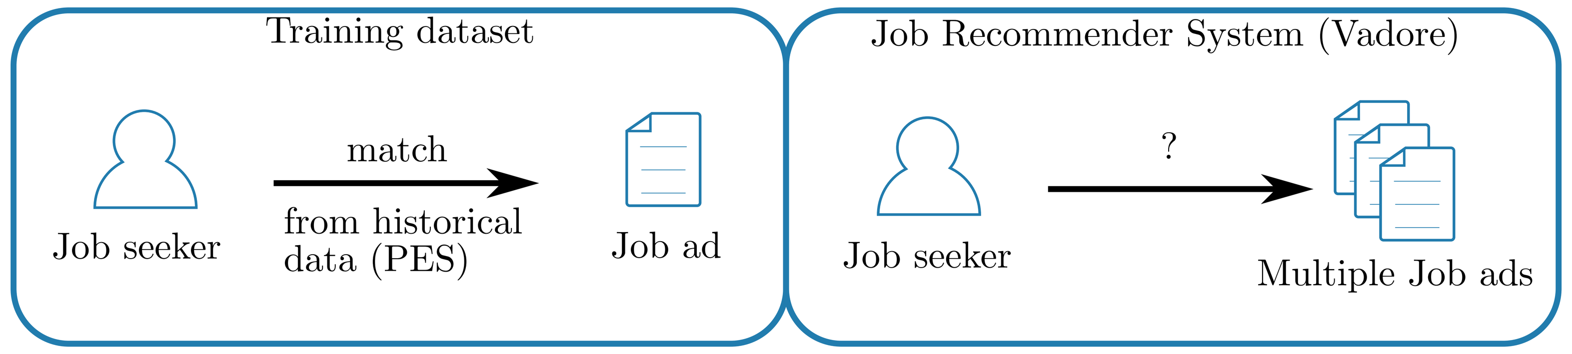

DATA

- Interactions: signed contracts (administrative data)
- Region: Auvergne-Rhône-Alpes
- On average: circa 400k job seekers, 64k job ads and 1.4k matches per week
- Available features:
  - Job seekers: zip code, skills, occupation, experience, education, skills, accepted mobility, text, past experiences, search criteria ...
  - Ads: zip code, skills, occupation, wage, contract, working hours, job & firm descriptions, requirements ...

PROPOSED APPROACH

- Goal: rank job ads j for some job seeker i
- Two-tiered architecture:
  - Tier 1: embedding-based VADORE.0, selects top-1,000 job ads
  - Enables Tier 2, VADORE.1 to re-rank those using a more expressive model

STEP 1: VADORE.0

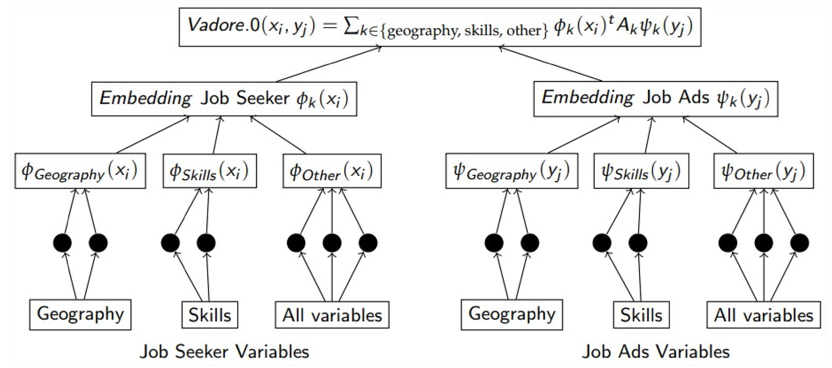

$$\text{VADORE.0}(x_i, y_j) = \sum_{k \in \{\text{general, geography, skills}\}} \phi_k(x_{ik})^T A_k \psi_k(y_{jk})$$

where ϕ_k, ψ_k embeddings for job seekers / ads; A_k matrices

- Three sub-scores, with inputs:
  - General: all relevant information ($d \approx 500$)
  - Geography: tiled representation of job seeker / ad zip-code
  - Skills: multihot of skills in an ontology
- Triplet loss: Weinberger et al., 2009

$$\sum_i \sum_{j' \neq j^*(i)} L(i, j^*(i), j')$$

where $j^*(i)$ is i's job, j' another job ad, and

$$L(i, j^*(i), j') = \max(s(x_i, y_{j^*(i)}) - s(x_i, y_{j'}) + \eta, 0), \quad \eta > 0$$

STEP 2: VADORE.1

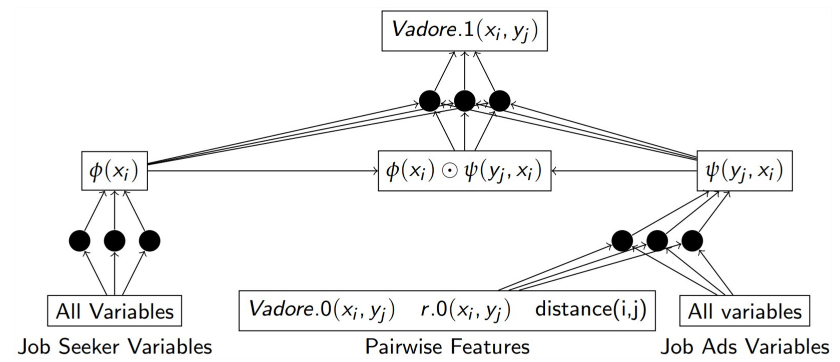

- Input: VADORE.0's top 1,000 job ads for job seeker i
- Job seeker representation as in VADORE.0 "general" submodule; job ads described by

$$\tilde{y}_j = [y_j, \quad \text{VADORE.0}(x_i, y_j), \quad \text{rank}_{\text{VADORE.0}}(i,j), \quad d_{ij}]$$

where d_{ij} distance in kilometers between i et j

- Network structure:

$$s(i,j) = \text{MLP}(z_{ij}), \quad z_{ij} = [\phi(x_i),\ \psi(\tilde{y}_j),\ \phi(x_i) \odot \psi(\tilde{y}_j)]$$

- Logistic loss (predict if pair is a match or not)

METRIC OF INTEREST

- Recall@k: % job seekers i s.t. rank$(j^*(i)) \leq k$

RESULTS

Recall@k	XGBOOST	VADORE.0	VADORE.1
10	**26.83**	22.88	25.96
20	35.59	31.55	**35.65**
50	48.75	44.12	**49.07**
100	**58.88**	53.80	58.67
1,000	**86.47**	82.13	-
Training time (h)	10	7	+ 0.6
Recommendation time per person (s)	7	0.002	+ 0.003

NON-STATIONARITY

- *Problem*: COVID and associated lock-downs affected the labor market. How did VADORE fare?
- COVID changed the job ad distribution; and to a lesser extent the job seeker distribution
- Nominal increase in recall during the pandemic: R@100 of 65 between 1st and 3rd French lock-downs (55 before, 57 after)
- However, after normalization for invariance w.r.t. the number of job ads, the recall remains stable

PERSPECTIVES

- **Randomized control trials** to assess:
  - Job seekers' view of VADORE recommendations
  - Causal impact on labor market outcomes
- Diagnosis of data and model's **biases**
- Mitigating the risk of **congestion**

REFERENCES

- Survey — de Ruijt et al., 2021
- RecSys 16 & 17 challenges — Xiao et al., 2016; Volkovs et al., 2017
- LinkedIn — Kenthapadi et al., 2017
- CareerBuilder — Zhao et al., 2021