Analysis of co-movement pattern mining methods for recommendation
(extended abstract)

Sergio Torrijos, **Alejandro Bellogín**
Universidad Autónoma de Madrid
Spain

CIRCLE, July 2020

Recommender Systems

- You already know about them (from previous talks)

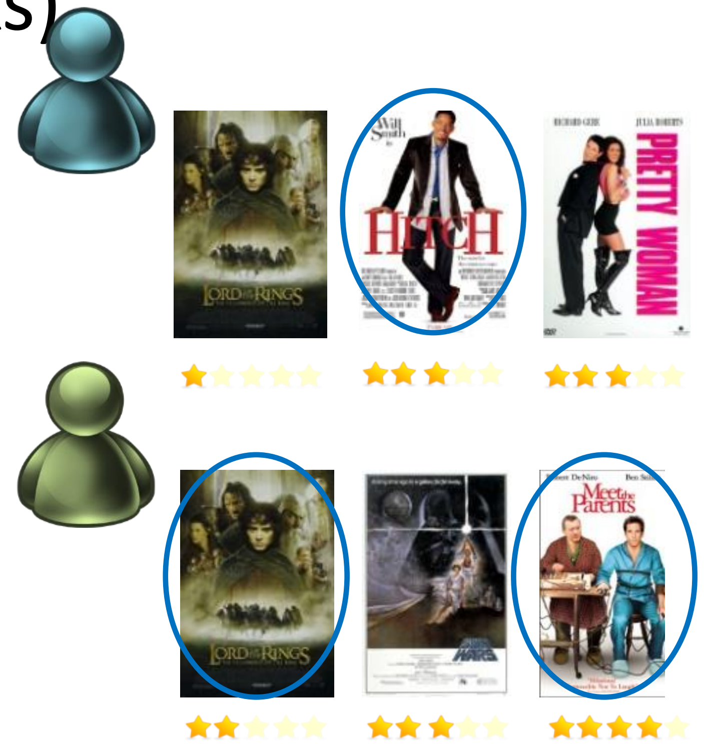

- Just in case:
  - Users interact (rate, purchase, click) with items
  - Which items will (another) user **like**?
    - Based on content of items, experiences of other users, etc.

Recommender Systems for tourism

- We need to consider some particularities of this domain
  - Types of venues
  - Proximity
  - Order

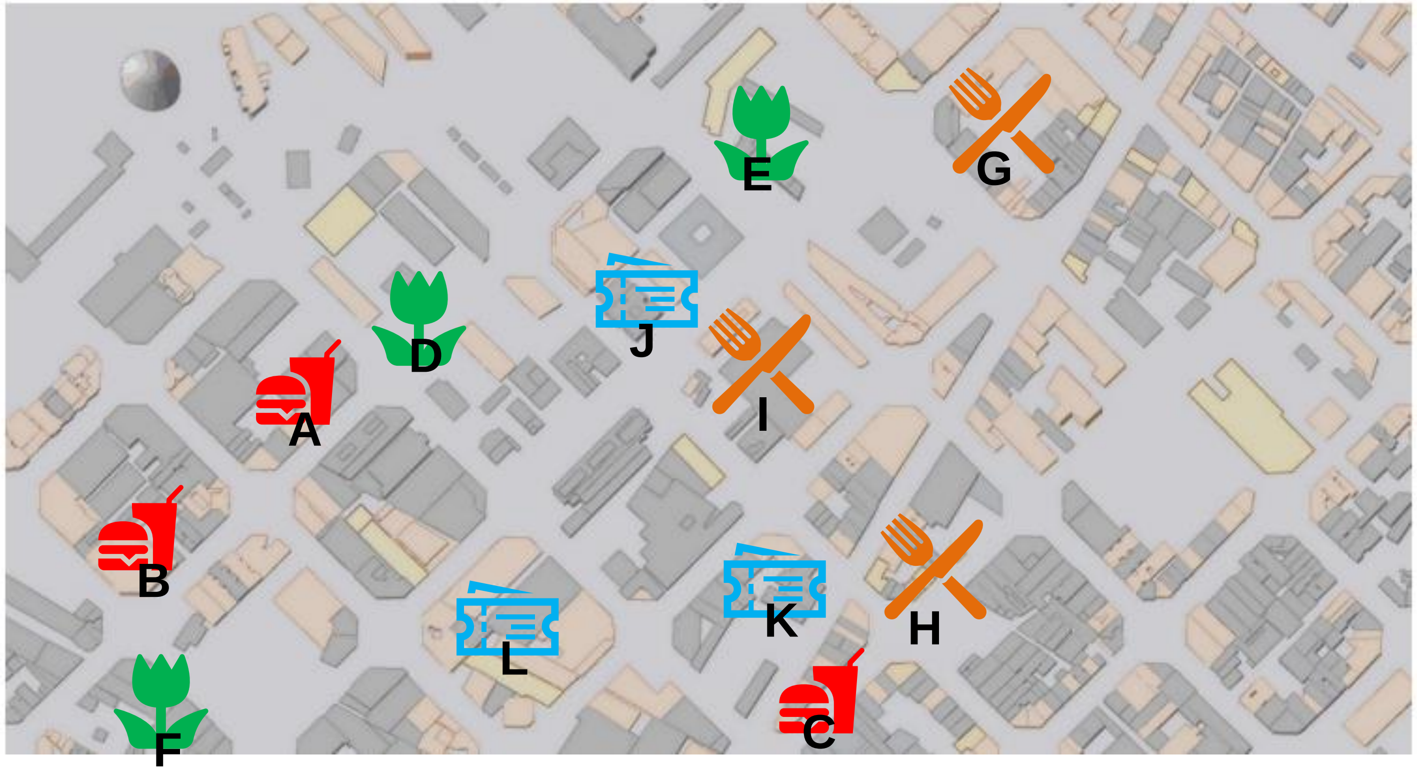

Recommender Systems for tourism

- We need to consider some particularities of this domain
  - Types of venues

    categories
  - Proximity

    geographical coordinates
  - Order

    sequences → trajectories

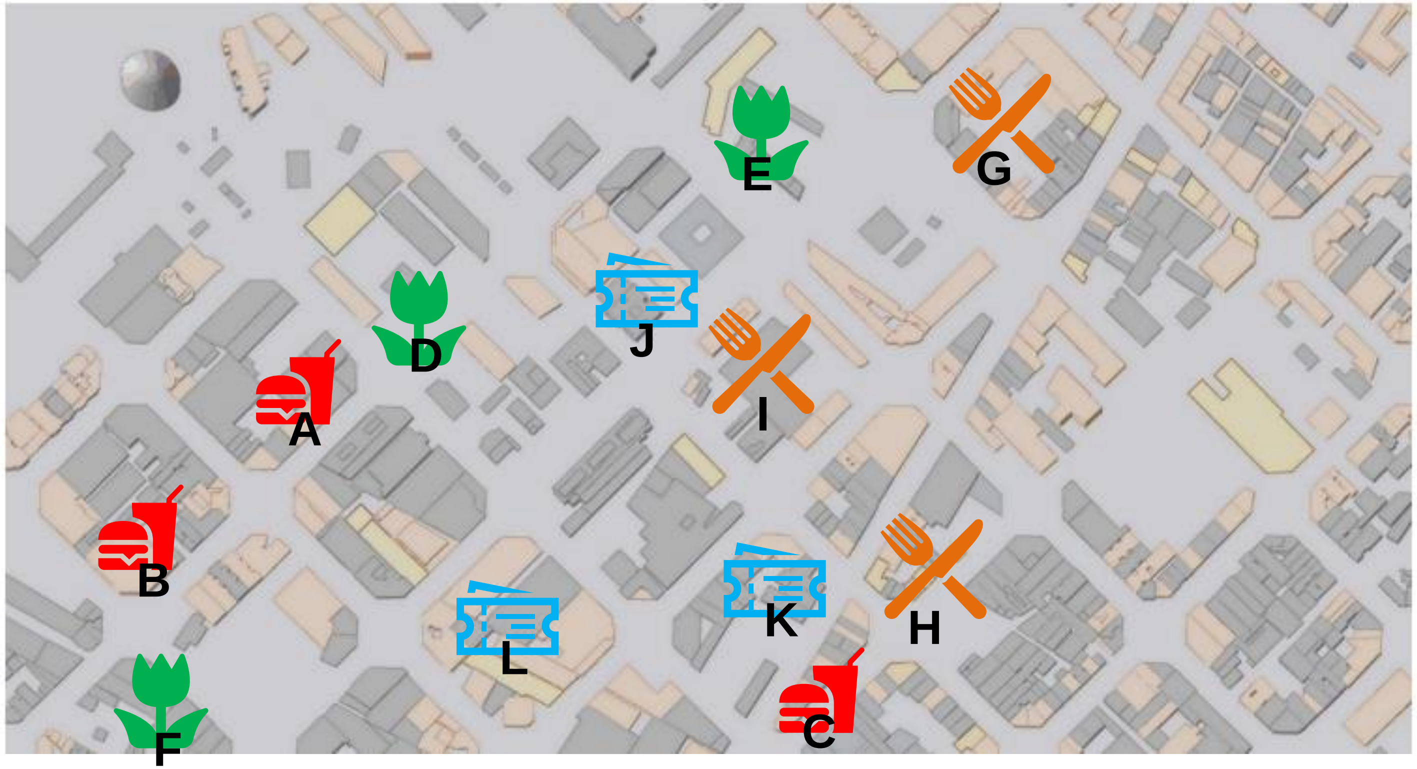

Trajectory pattern mining

- A very broad area

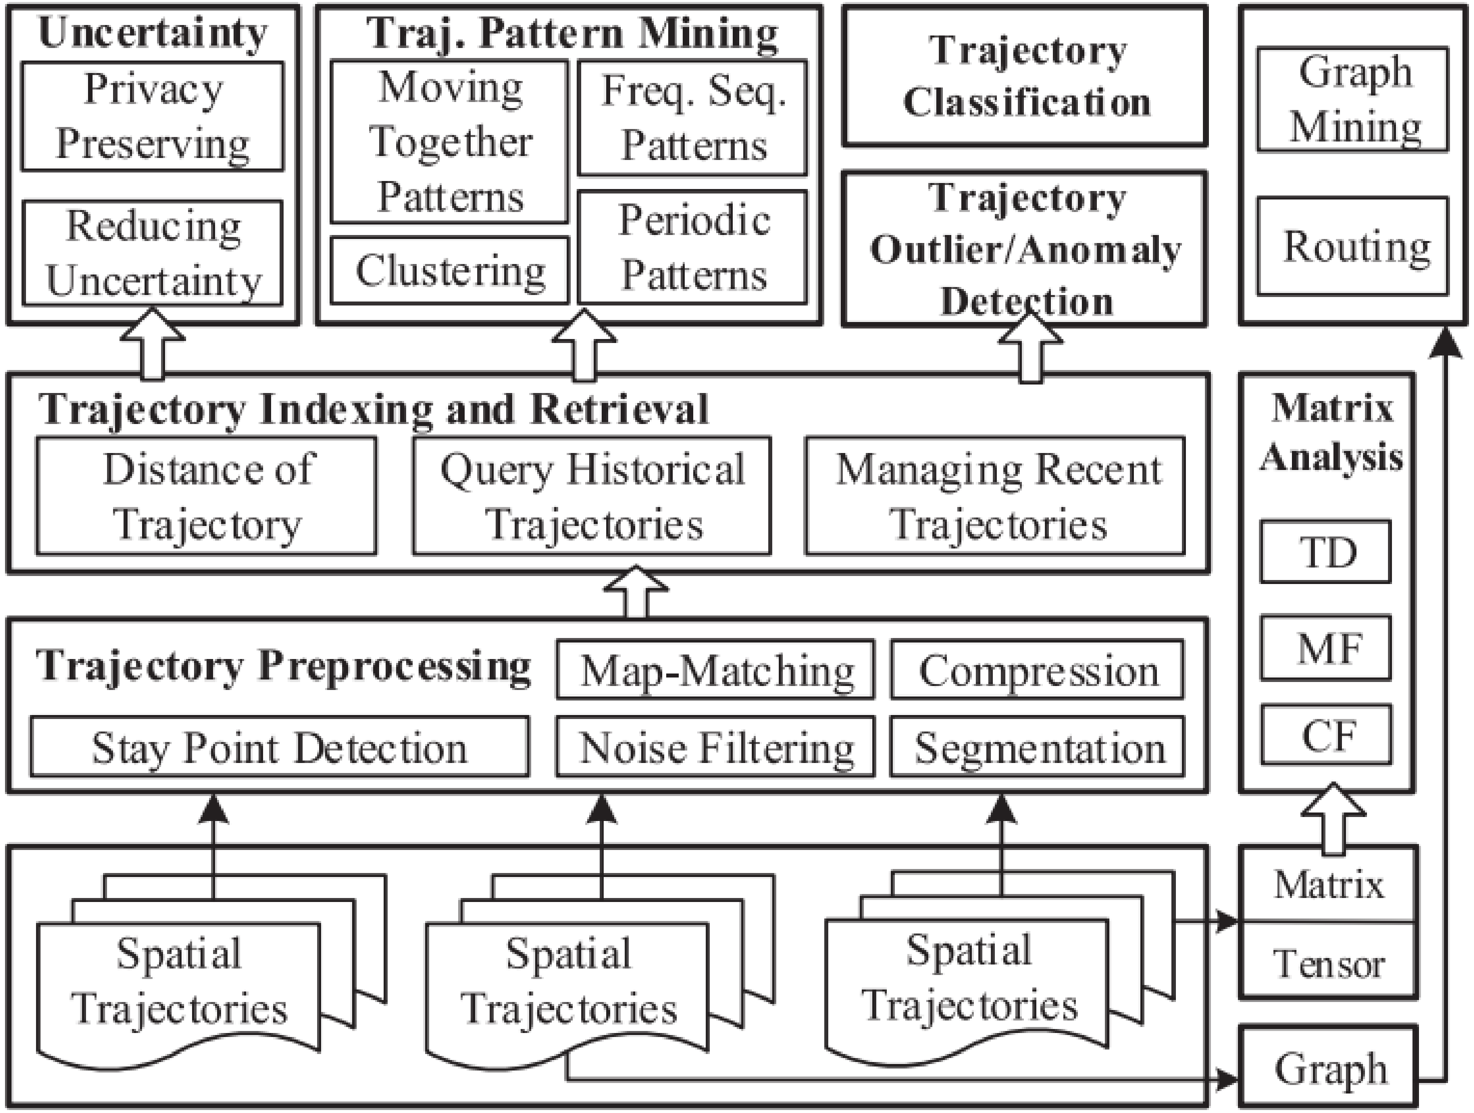

Trajectory Data Mining: An Overview

YU ZHENG, Microsoft Research

Universidad Autónoma de Madrid

Escuela Politécnica Superior

Exploiting trajectory patterns for recommendation

- First, what is a trajectory?
  - Data from Location-Based Social Networks (e.g., Foursquare) do not form dense trajectories
  - We need to process check-in information

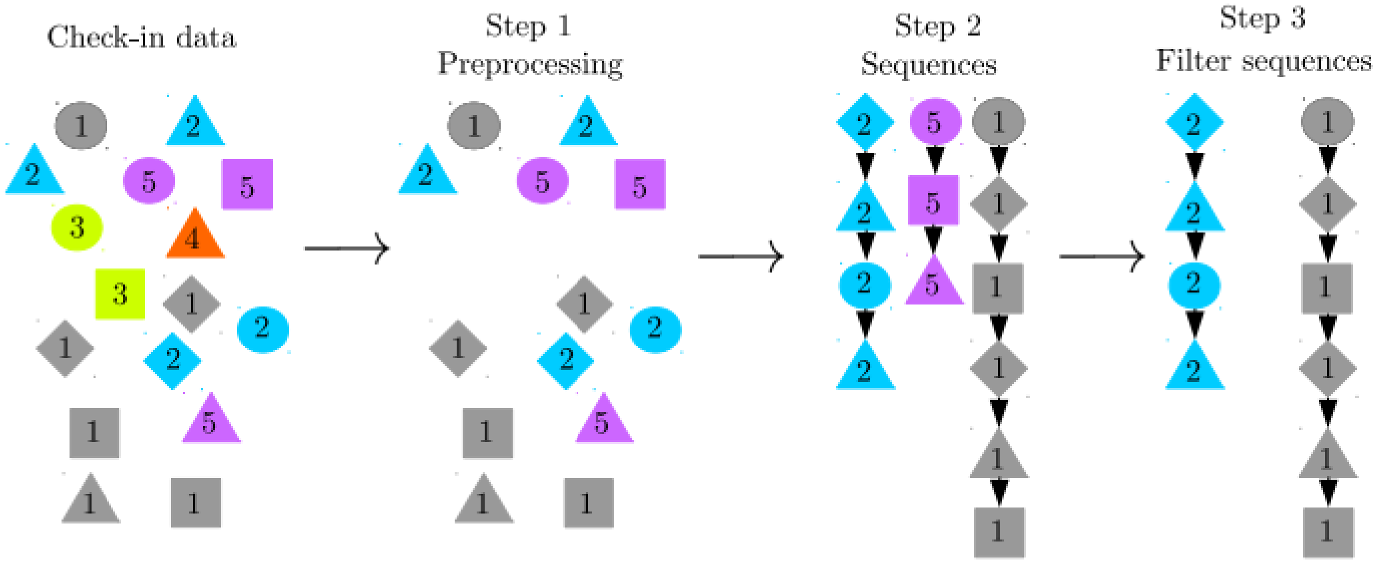

Each number/color: a different user

Each shape: type of point

Applying reranking strategies to route recommendation using sequence-aware evaluation

Pablo Sánchez · Alejandro Bellogín

Exploiting trajectory patterns for recommendation

- For filtering
  - Present the user the most similar trajectories

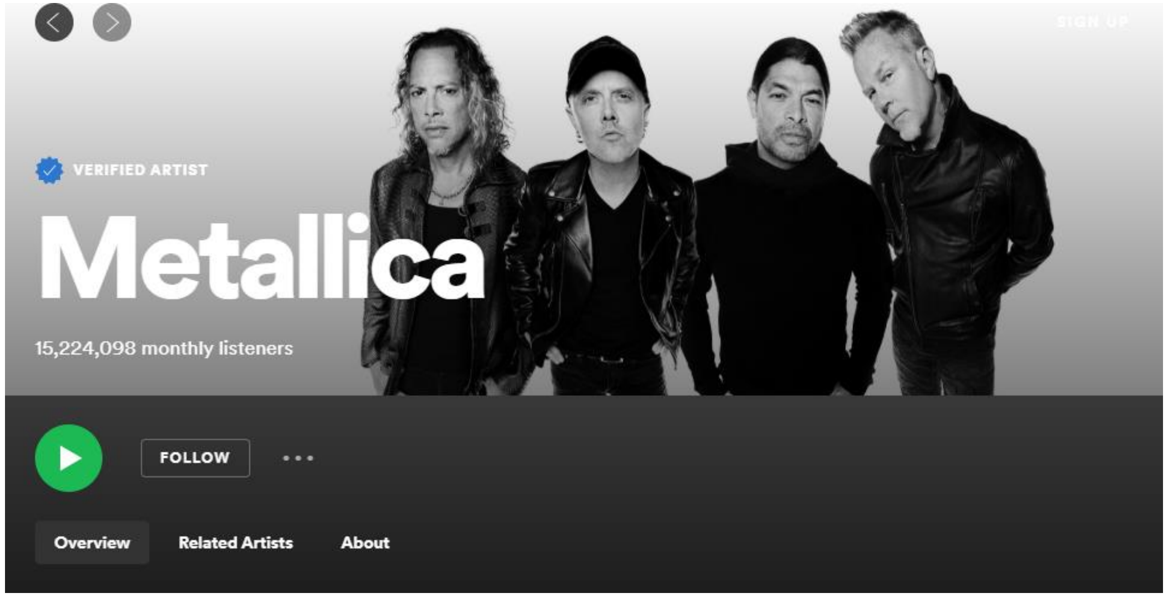

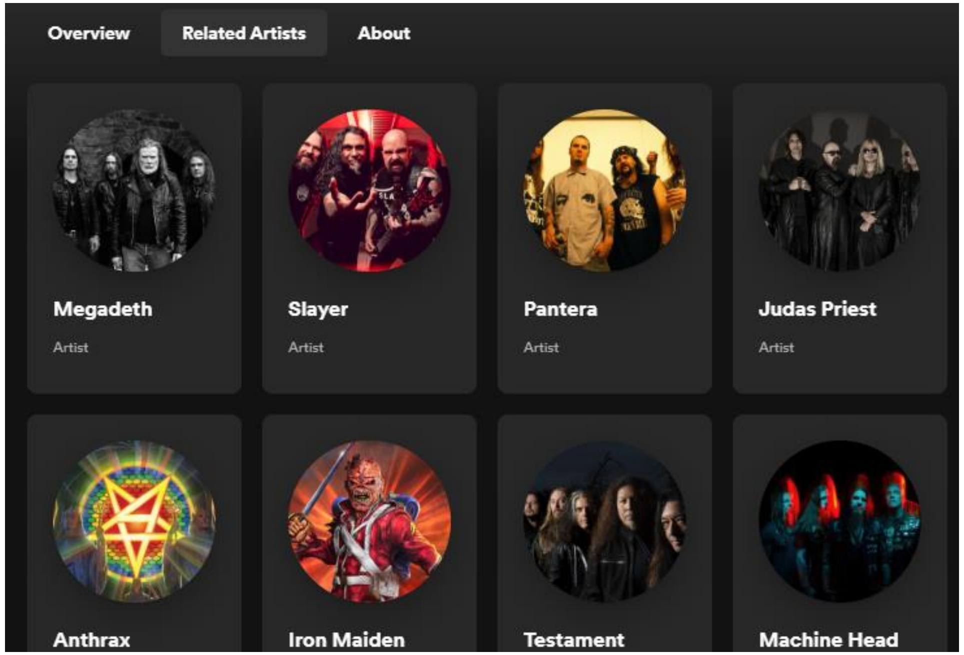

Universidad Autónoma de Madrid

Escuela Politécnica Superior

Exploiting trajectory patterns for recommendation

- For filtering
  - Present the user the most similar trajectories
  - Using trajectory indexing, clustering, or similarity methods

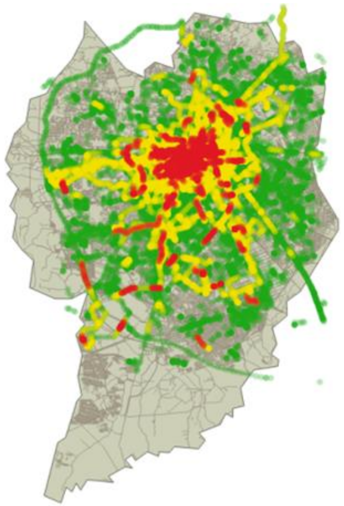

ST-DBSCAN: An algorithm for clustering spatial–temporal data

Derya Birant *, Alp Kut

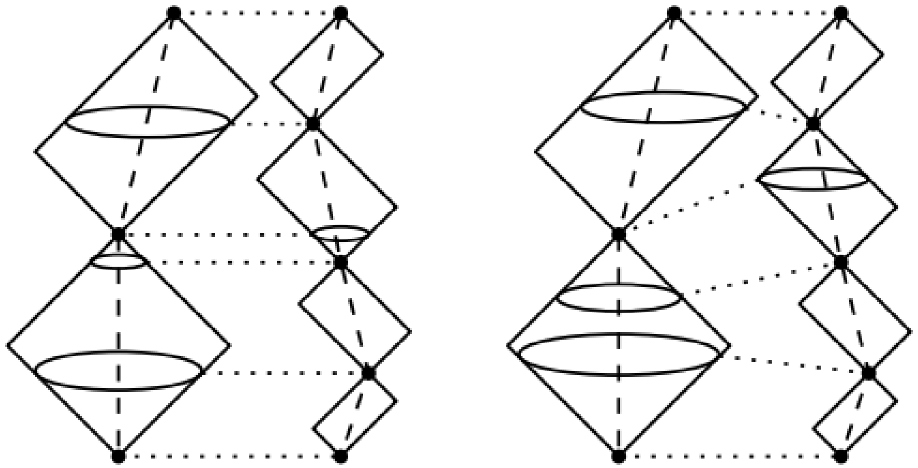

(a) Equal time distance (b) Fréchet distance

Computing Similarity of Coarse and Irregular Trajectories Using Space-Time Prisms

Maike Buchin
Department of Mathematics
Ruhr-University Bochum
Germany
maike.buchin@rub.de

Ross S. Purves
Department of Geography
University of Zurich – Irchel
Switzerland
ross.purves@geo.uzh.ch

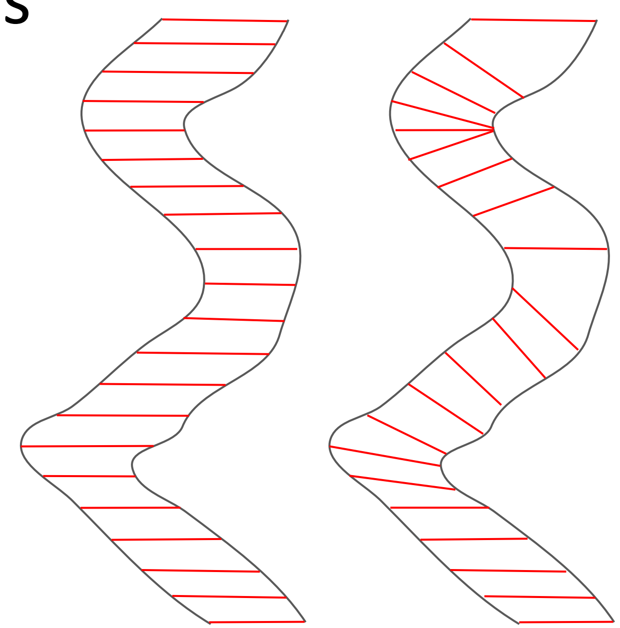

Euclidean vs DTW (dynamic time warping) similarity

Exploiting trajectory patterns for recommendation

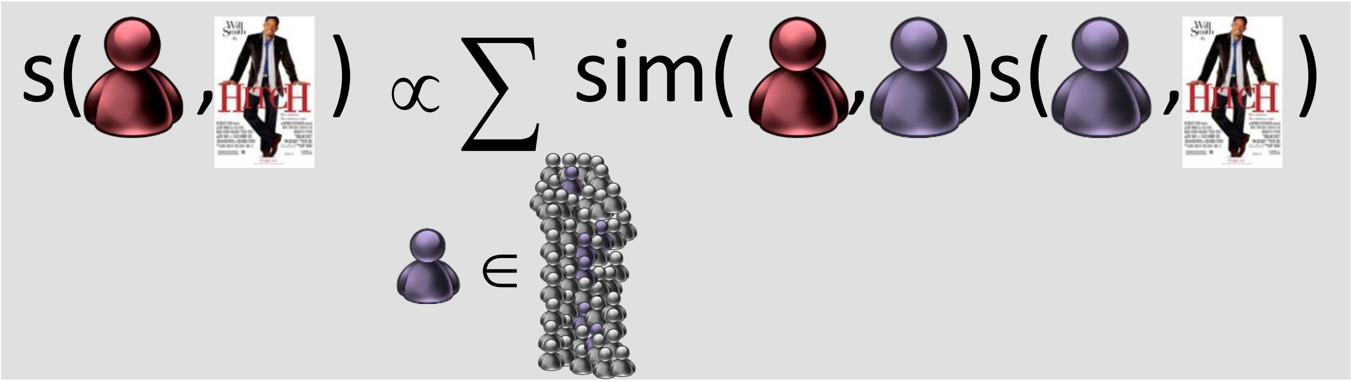

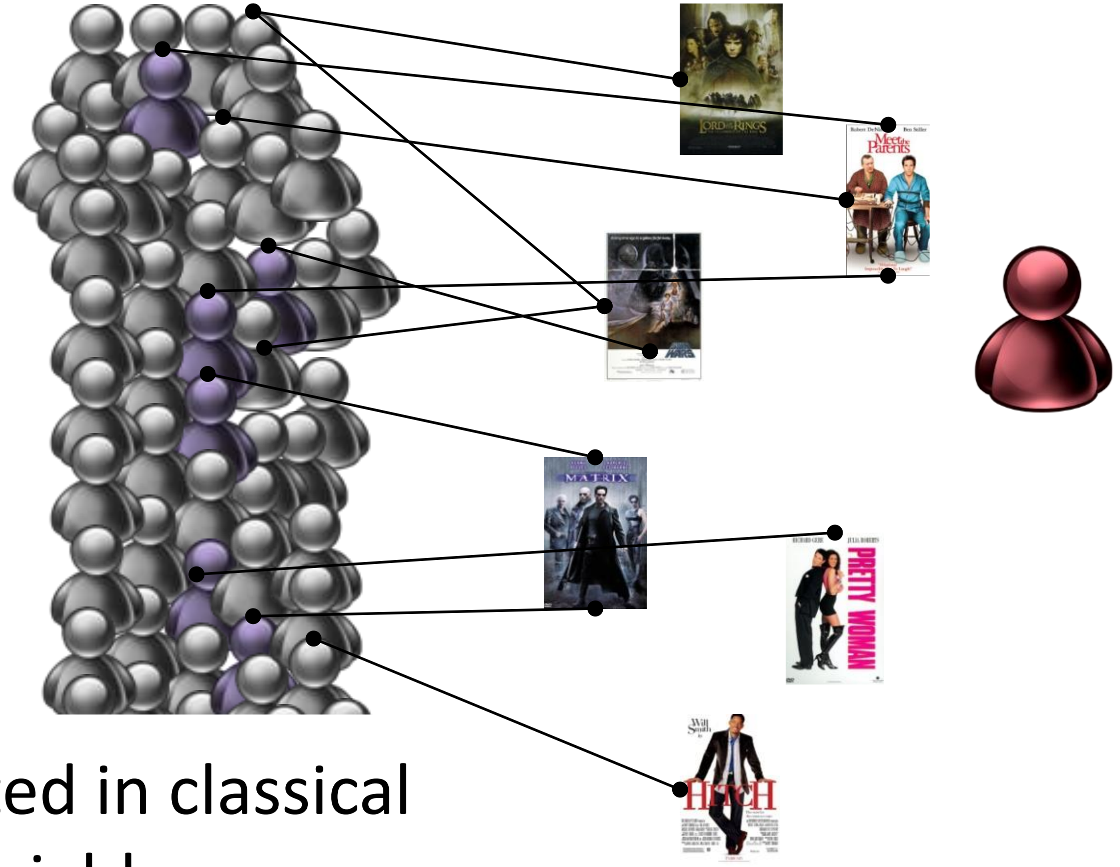

- To obtain similar users/items
  - These similar users/items can be integrated in classical recommendation methods like nearest neighbours

Exploiting trajectory patterns for recommendation

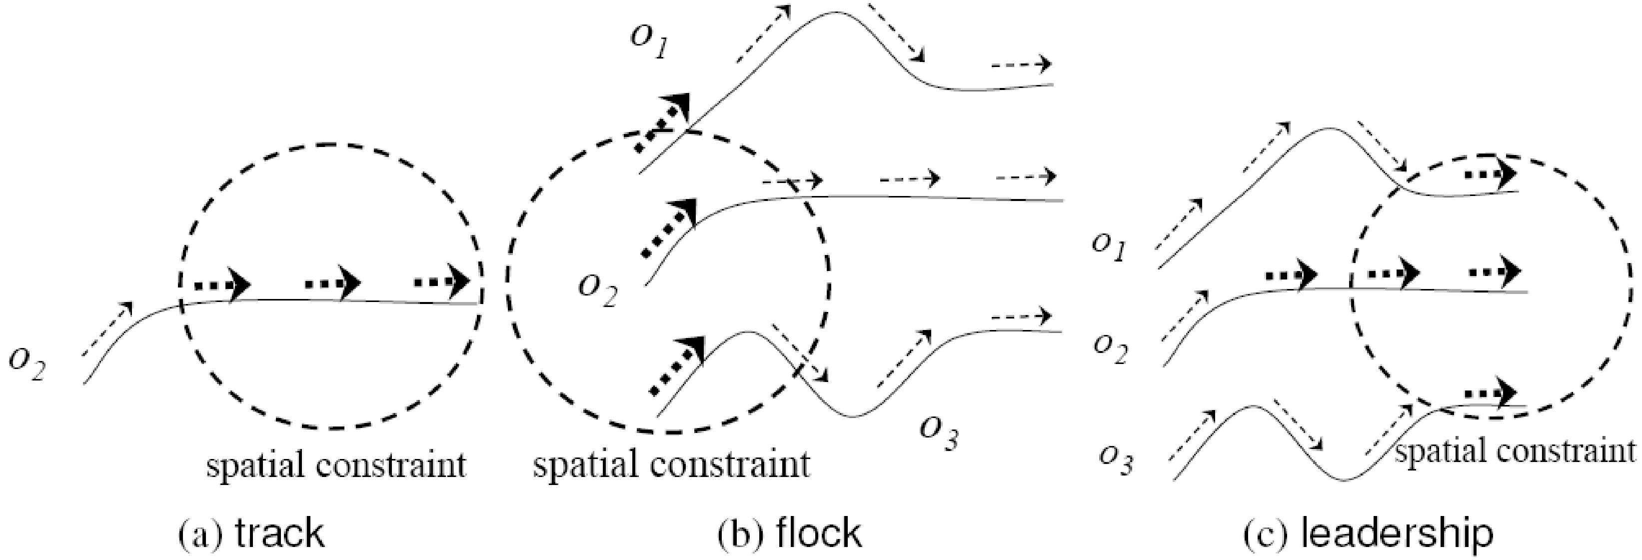

- To obtain similar users/items
  - These similar users/items can be integrated in classical recommendation methods like nearest neighbours
  - Using trajectory similarities or co-movement patterns

Some results…

- Data: from Tokyo (Foursquare) and Rome (Tripbuilder)
  - Tokyo: 328K interactions, 8.6K users, 28.6K items → 0.13%
  - Rome: 60K interactions, 8K users, 394 items → 1.9%
- Methods
  - Trajectory similarity: based on DTW and Hausdorff distance
  - Clustering: ST-DBSCAN
  - Co-movement: Flock, Convoy

Efficiency (Rome)

- Trajectory data mining methods are very expensive

Algorithm	1K	10K	20K	60K
Ad-Hoc	0m0.477s	0m1.952s	0m5.253s	0m35.758s
Sim. tray. Hausdorff	0m10.568s	22m32.164s	88m17.315s	1157m31.251s
Sim. tray. DTW	0m3.386s	6m16.519s	26m9.038s	232m25.592s
ST-DBSCAN	0m0.817s	0m3.173s	0m9.487s	1m29.065s
Convoy	0m5.786s	0m6.086s	0m6.266s	0m6.098s
Convoy partials	0m2.786s	0m3.086s	0m3.266s	0m4.098s
Flock	0m5.199s	24m23.087s	49m34.065s	†
Flock partials	0m5.500s	0m30.545s	0m14.880s	3m11.811s

- We developed an optimised version (partial) based on data subsets

Performance (Rome)

- We observe different levels of performance (precision)
  - Flock obtains good results
  - Clustering depends on the parameters

Ad-Hoc	δ=3600	0.0361754
	δ=10800	0.0349337
Sim. tray.	Hausdorff	0.0175497
	DTW	0.0206126
ST-DBSCAN	ε=5000, *th*=6000, *minEps*=1	**0.0369205**
	ε=5000, *th*=6000, *minEps*=10	0.0316225
	ε=5000, *th*=6000, *minEps*=10	**0.0369205**
	ε=1000, *th*=2000, *minEps*=1	0.0336921
	ε=4000, *th*=2000, *minEps*=1	0.0277318
Convoy	*min points*=2, *lifetime*=10, *dist max*=10^{-2}	0.0362583
	min points=20, *lifetime*=2, *dist max*=10^{-2}	0.0332781
	min points=2, *lifetime*=2, *dist max*=10^{-3}	0.0243377
	min points=2, *lifetime*=5, *dist max*=10^{-2}	0.0352649
Convoy partials	*min points*=2, *lifetime*=2, *dist max*=10^{-3}	0.0360927
	min points=2, *lifetime*=2, *dist max*=10^{-1}	0.0368377
	min points=3, *lifetime*=5, *dist max*=10^{-1}	0.0276490
Flock	ε=10, μ 2, δ=2	**0.0369205**
	ε=0.1, μ 2, δ=10	**0.0369205**
	ε=10, μ 2, δ=2	**0.0369205**
Flock partials	ε=10, μ 2, δ=2	0.0368377
	ε=0.1, μ 2, δ=10	0.0355132
	ε=0.1, μ 2, δ=2	0.0368377

Visualisation

- We plot in a map the neighbours these methods find

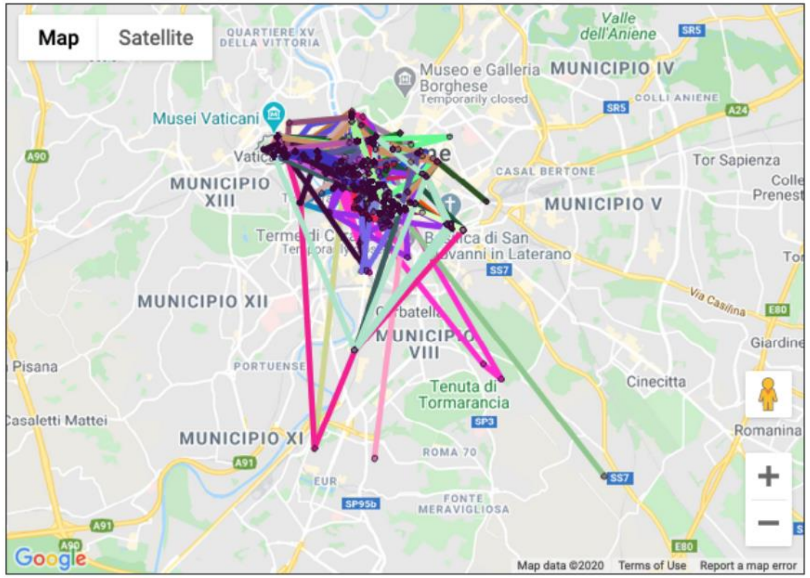

Users | KN | Neighbors

2066	Neighbor ID	Similarity
2066	1795	0.9808990905
	2223	0.9808990905

DTW

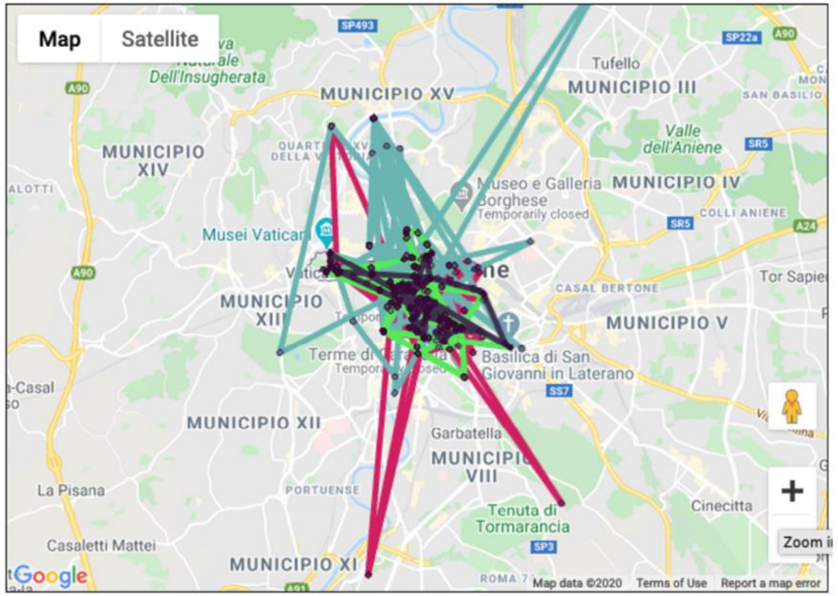

Users | KN | Neighbors

296	Neighbor ID	Similarity
296	1589	1
	539	1

Convoy (2, 2, 0.01)

Performance and social analysis (Tokyo)

- Very good results when computing performance only for tourist users
  - Tourist: if their interactions in the city span less than 21 days

Recommender	NDCG	FA	AD	EPC
Ad-hoc	0.071	†**0.271**	**1,217**	0.770
TS-DTW	0.069	0.246	320	**0.847**
TS-Haus	0.068	0.252	333	0.839
UB	**0.095**	0.245	1,133	0.710
IB	0.064	0.262	†3,208	†0.858
BPR	0.108	0.266	18	0.566
IRenMF	†0.111	0.262	106	0.580

- Overlap with explicit social connections

Neigh. Sel.	Method	Avg.	Total	T-NDCG	T-FA
Best	Ad-hoc	2	1,008	0.071	0.271
	TS-DTW	9	4,833	0.069	0.246
	TS-Haus	9	4,833	0.068	0.252
As UB	Ad-hoc	2	1,008	0.071	0.273
	TS-DTW	2	1,074	0.056	0.271
	TS-Haus	5	2,685	0.050	0.270
	UB	4	2,140	0.095	0.245

Discovering Related Users in Location-Based Social Networks

Sergio Torrijos
Universidad Autónoma de Madrid
Madrid, Spain
sergio.torrijos@estudiante.uam.es

Alejandro Bellogín
Universidad Autónoma de Madrid
Madrid, Spain
alejandro.bellogin@uam.es

Pablo Sánchez
Universidad Autónoma de Madrid
Madrid, Spain
pablo.sanchezp@uam.es

Discussion and future work

- Trajectory pattern mining methods can be useful for recommendation
  - There are computation problems that should be solved
  - Some inherent properties might alleviate these problems:
    - Data points are limited by actual venues, instead of fine-grained GPS coordinates
    - Popularity bias could guide the predictions

- Temporal and spatial information are very important
  - Using this information is not novel, but there are few works exploiting trajectories in a principled way
  - How can it be integrated in other tasks such as trip recommendation or successive recommendation?

Thank you

Analysis of co-movement pattern mining methods for recommendation

Sergio Torrijos, Alejandro Bellogín

Universidad Autónoma de Madrid

Spain

CIRCLE, July 2020

Performance (Rome)

- We observe different levels of performance
  - Flock obtains good results
  - Clustering depends on the parameters

		Precision	MAE
Ad-Hoc	δ=3600	0.0361754	**0.3487827**
	δ=10800	0.0349337	0.3637682
Sim. tray.	Hausdorff	0.0175497	0.5784518
	DTW	0.0206126	1.2903128
ST-DBSCAN	ε=5000, *th*=6000, *minEps*=1	**0.0369205**	0.4251326
	ε=5000, *th*=6000, *minEps*=10	0.0316225	0.3530357
	ε=5000, *th*=6000, *minEps*=10	**0.0369205**	0.4000756
	ε=1000, *th*=2000, *minEps*=1	0.0336921	0.3530357
	ε=4000, *th*=2000, *minEps*=1	0.0277318	0.3525041
Convoy	*min points*=2, *lifetime*=10, *dist max*=10^{-2}	0.0362583	0.3566276
	min points=20, *lifetime*=2, *dist max*=10^{-2}	0.0332781	1.0634231
	min points=2, *lifetime*=2, *dist max*=10^{-3}	0.0243377	0.5973882
	min points=2, *lifetime*=5, *dist max*=10^{-2}	0.0352649	0.4018430
Convoy partials	*min points*=2, *lifetime*=2, *dist max*=10^{-3}	0.0360927	0.3925327
	min points=2, *lifetime*=2, *dist max*=10^{-1}	0.0368377	0.3525041
	min points=3, *lifetime*=5, *dist max*=10^{-1}	0.0276490	0.7646296
Flock	ε=10, μ 2, δ=2	**0.0369205**	0.3530357
	ε=0.1, μ 2, δ=10	**0.0369205**	0.3704636
	ε=10, μ 2, δ=2	**0.0369205**	0.3530357
Flock partials	ε=10, μ 2, δ=2	0.0368377	0.3525041
	ε=0.1, μ 2, δ=10	0.0355132	0.3530357
	ε=0.1, μ 2, δ=2	0.0368377	0.3670011

Experiments

- Foursquare data: Tokyo from global check-in dataset (33M) ~ 328K
- Temporal Split: 6 months for training, 1 month test
- Baselines
  - UB: neighbour recommender with classic user similarity
  - IB: neighbour recommender with classic item similarity
  - BPR: Bayesian Personalised Ranking using a matrix factorisation algorithm
  - IRenMF: matrix factorisation algorithm that exploits geographical influence
- Metrics
  - NDCG: accuracy of item recommendations
  - FA: feature agreement, or precision in terms of category matching (not items)
  - AD and EPC: diversity and novelty metrics

Performance comparison

- Neighbours are not competitive against MF methods in terms of accuracy

Recommender	NDCG	FA	AD	EPC
Ad-hoc	0.054	0.206	**3,314**	0.741
TS-DTW	0.036	0.203	408	0.871
TS-Haus	0.036	**0.215**	442	0.870
UB	**0.062**	0.201	1,899	**0.912**
IB	0.046	†0.264	†13,685	†0.956
BPR	0.066	0.229	28	0.872
IRenMF	†0.069	0.225	1,475	0.888

Performance comparison

- Neighbours are not competitive against MF methods in terms of accuracy
- Much better results are found for beyond-accuracy dimensions:
  - Ad-hoc is the best one for diversity (AD)
  - Similarity with Hausdorff is the best one for category accuracy (FA)

Recommender	NDCG	FA	AD	EPC
Ad-hoc	0.054	0.206	**3,314**	0.741
TS-DTW	0.036	0.203	408	0.871
TS-Haus	0.036	**0.215**	442	0.870
UB	**0.062**	0.201	1,899	**0.912**
IB	0.046	†0.264	†13,685	†0.956
BPR	0.066	0.229	28	0.872
IRenMF	†0.069	0.225	1,475	0.888

Impact on local vs tourist users

- There are different types of users in LBSNs:
  - Locals (if their check-ins span more than 21 days) vs tourists
- IRenMF is still the best approach
- But neighbour recommenders improve their performance for tourists
  - In particular, for FA

Recommender	NDCG	FA	AD	EPC
Ad-hoc	0.050	0.181	**2,829**	0.735
TS-DTW	0.032	0.179	358	**0.869**
TS-Haus	0.032	**0.191**	405	0.869
UB	**0.057**	0.177	1,440	0.704
IB	0.042	†0.250	†11,929	†0.872
BPR	0.059	0.204	26	0.608
IRenMF	†0.063	0.200	1,199	0.647

Local users

Recommender	NDCG	FA	AD	EPC
Ad-hoc	0.071	†**0.271**	**1,217**	0.770
TS-DTW	0.069	0.246	320	**0.847**
TS-Haus	0.068	0.252	333	0.839
UB	**0.095**	0.245	1,133	0.710
IB	0.064	0.262	†3,208	†0.858
BPR	0.108	0.266	18	0.566
IRenMF	†0.111	0.262	106	0.580

Tourist users

Social network analysis

- How similar are the found neighbours to explicit social connections?

Neigh. Sel.	Method	Avg.	Total	T-NDCG	T-FA
Best	Ad-hoc	2	1,008	0.071	0.271
	TS-DTW	9	4,833	0.069	0.246
	TS-Haus	9	4,833	0.068	0.252
As UB	Ad-hoc	2	1,008	0.071	0.273
	TS-DTW	2	1,074	0.056	0.271
	TS-Haus	5	2,685	0.050	0.270
	UB	4	2,140	0.095	0.245

- TS-Haus always obtains more social connections than the baseline UB
- Performance accuracy on tourist users is competitive (T-NDCG)
- Feature agreement is always better than baseline (T-FA)