Urban Memory Geography: mapping the temporal structure of human return

Zhuoyue Zhang

Aug 2026

1. Introduction

From frequency sufficiency to urban destination geography

Preferential Return: the tendency to revisit locations that have been visited more frequently in the past.
In the original EPR* model, destinations are represented by their prior visit counts.

However, visit count records how often, but not when:

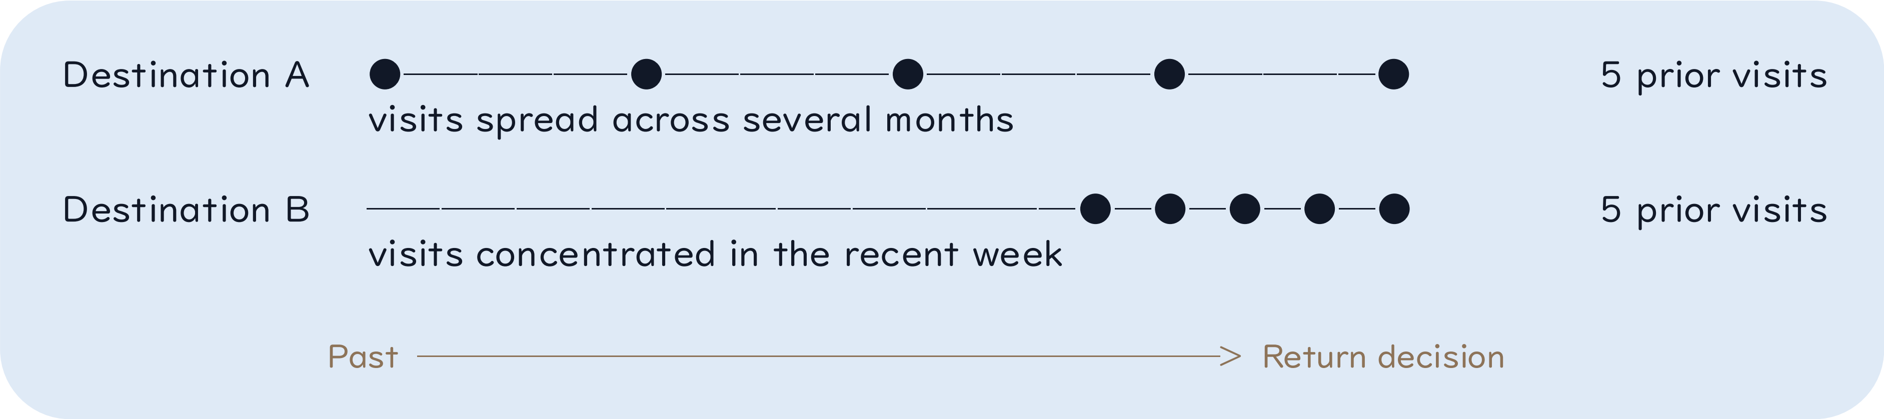

Same visit count ≠ Same temporal history
Question 1: Does the timing of prior visits provide additional information about return choice?

If visit timing contains information beyond count, that information may not be uniform across places.
Question 2: Is temporal return structure spatially organized across urban destinations?

* EPR: Exploration and Preferential Return

2. Literature Review

Mobility models summarize return history in different ways

EPR framework

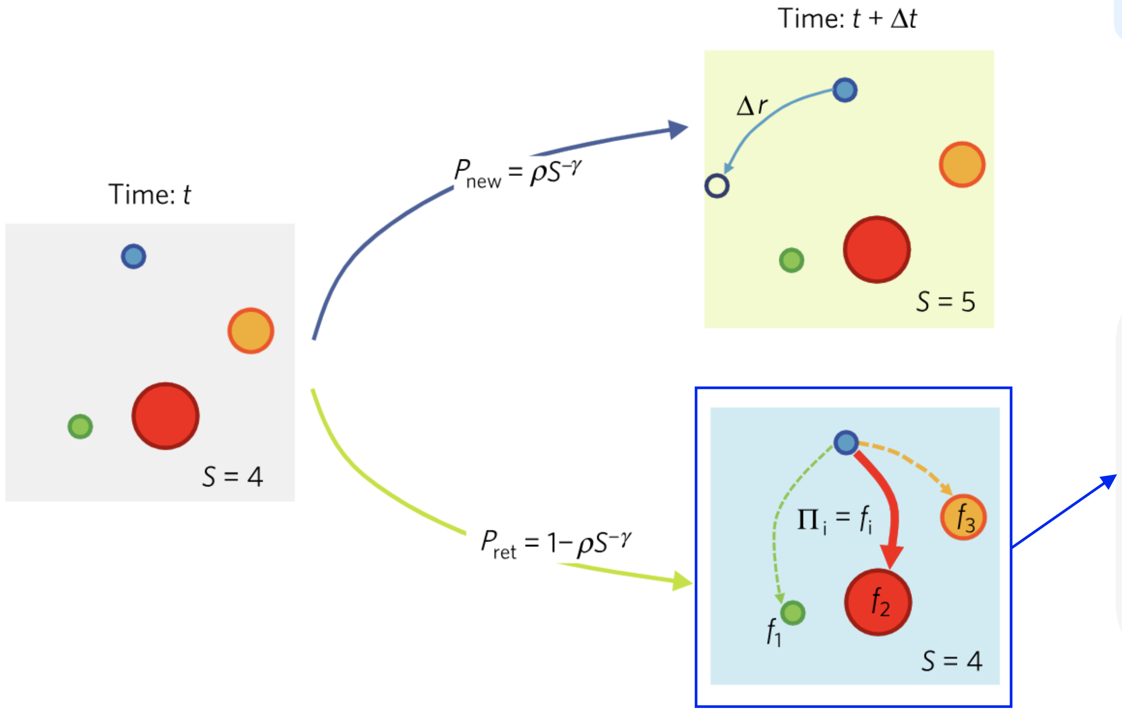

Full Visit History

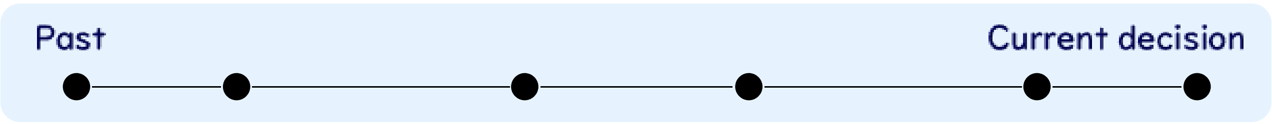

Conditional on return:
Which previously visited destination is selected?

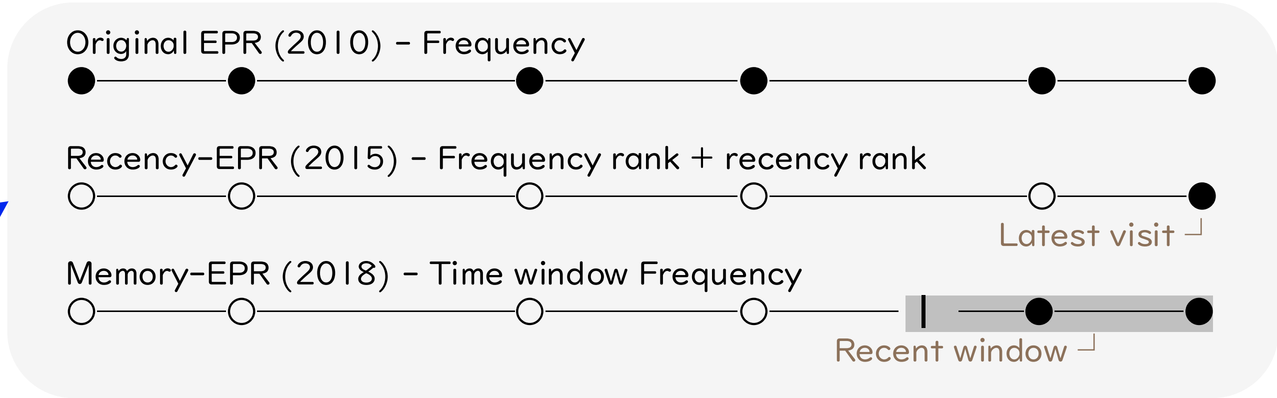

Gap 1:
Existing mobility models incorporate temporal information through different summaries of prior visits.
But they do not provide a common way to test whether cumulative visit frequency is sufficient.

Song et al. (2010); Barbosa et al. (2015); Alessandretti et al. (2018)

Past experiences can be accumulated with time-dependent weighting

Cognitive Literature

- repeated past experiences can be accumulated with time-dependent decay.
- ACT-R implements this through power-law decay.

Transport Learning Literature

- Time-dependent forgetting has been implemented using power-law or exponential decay for repeated route-choice experience.
- travelers update route preferences based on previous route experiences.
- older travel experiences receive less weight when forming current expectations.

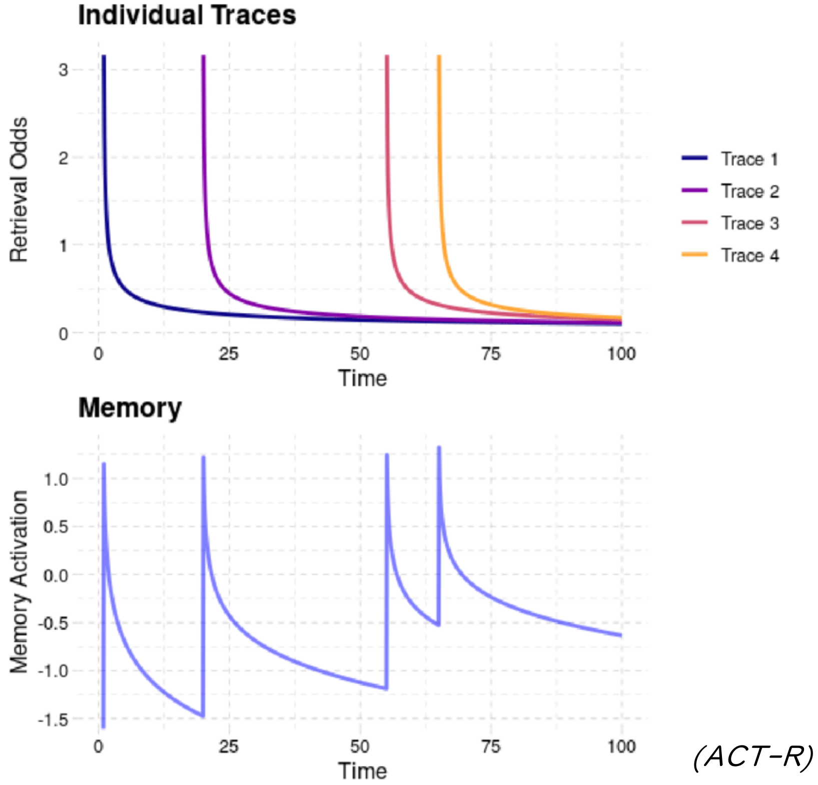

(ACT-R)

Gap 2:
Both literatures motivate time-weighted accumulation, but neither directly tests destination return history beyond visit count.

Anderson et al., 2004 , Tang, Gao & Ben-Elia (2017)

3. Methodology

Three-Stage Diagnostic of Return-history Information

Different rules retain different information.

Return Rules	Return-history Information
Original EPR	frequency
Recency EPR	frequency and recency mixture
Cognitive	power law
	exponential
Temporal	latest visit only

These rules retain different parts of the same visit history.
We test whether the additional information improves held-out return choice.

1 → **2** → **3**

1. Timing beyond frequency

Does timing matter after frequency is matched?

2. Timing beyond frequency and distance

Does timing remain informative after distance is controlled?

3. Full history beyond the latest visit, frequency, and distance

Does temporal history remain informative beyond the latest visit after frequency and distance are controlled?

Destination-level signature and geographic evaluation

Estimate a destination recurrence signature

- Evaluate temporal weighting for each return event
- Pool evidence by chosen destination
- Select the temporal decay that minimizes pooled loss

destination recurrence signature: the temporal decay best supported by returns ending at destination j

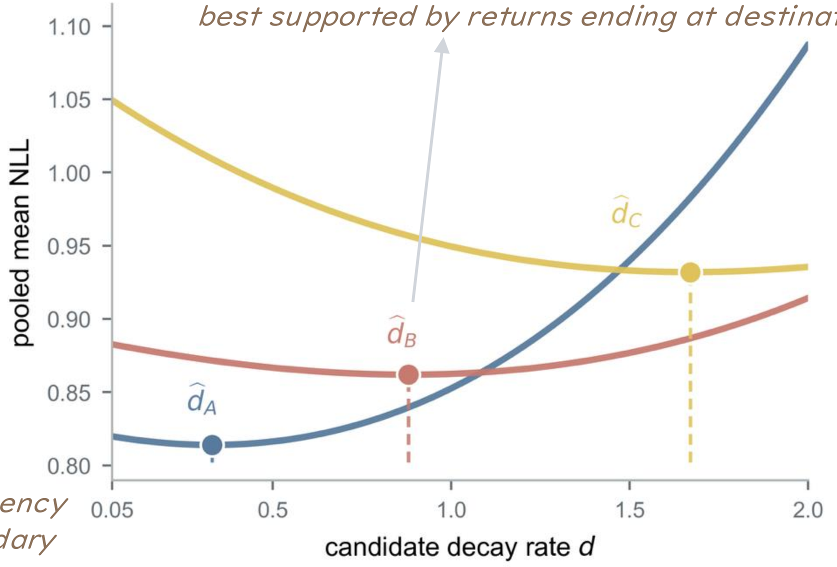

d=0: frequency boundary

Lower signature
Frequency-leaning

Higher signature
More age-sensitive

Older visits retain more weight → Recent visits receive more weight

Evaluate the spatial organization of signatures

- Map each signature to its destination cell
- Moran's I: asks whether neighboring destination cells have similar recurrence signatures.
- Positive I: similar signature values cluster spatially

- Three Spatial Checks:

1. Raw signature pattern

   Are similar signatures spatially clustered?

2. After adjusting for event count

   Remove the association with observed return counts. Does spatial clustering remain beyond uneven observational support?

3. After adjusting for event count and urban density

   Additionally remove the measured density gradient. Does spatial clustering remain beyond support and the measured urban-density gradient?

4. Results

Two Datasets

Comparison	Boston Walking Study (BWS)	YJMob100K (Anonymous City in Japan)
Study coverage	974 participants ≥14 days	100,000 users · 75 days City and dates anonymized
Observed mobility	GPS-derived stays 172,322 stays with location, time, and duration	30-minute cell occupancy Anonymous user, study day, time slot, and grid coordinates
Spatial / temporal support	500 m × 500 m cells Stays expanded to 1-hour slots	500 m × 500 m cells Stays require ≥2 consecutive 30-minute slots
Stay consolidation	Retain one cell per user-hour Merge consecutive same-cell slots	Remove same-cell repeats separated by ≤1 hour
Routine locations	Exclude cells accounting for ≥10% of training-window dwell time	Retain all qualifying stays No routine-place filter
Environmental context	WorldPop 2020 population density	POI counts across 85 categories Total POI count used as POI density
Shared temporal split and held-out evaluation	stays → per-user 70/30 temporal cutoff → train-only model selection → held-out evaluation	

Data sources: *[1] Meister et al. (2025), Transportation. [2] Yabe et al. (2024), Scientific Data.*

Question 1: Does the timing of prior visits provide additional information about return choice?

I Descriptive Analysis: How do alternative return-history information perform in held-out return choice?

Representation	Train-selected parameters	ΔNLL vs Frequency (95% CI) ↑
Boston: 8,385 held-out returns from 879 users		
Frequency	$\gamma = 1.20$	–
Recency	$\alpha = 0.80, \gamma = 0.10, \eta = 0.60$	-0.048 [-0.060, -0.036]
Exponential	$\tau = 0.01, \gamma = 0.40$	-0.466 [-0.529, -0.413]
Power law	$d = 0.05, \gamma = 1.20$	**0.004** [0.003, 0.005]
YJMob: 2,842,275 held-out returns from 94,320 users		
Frequency	$\gamma = 1.20$	–
Recency	$\alpha = 0.70, \gamma = 0.80, \eta = 1.80$	-0.057 [-0.058, -0.055]
Exponential	$\tau = 0.003, \gamma = 1.20$	0.000 [-0.001, 0.002]
Power law	$d = 0.40, \gamma = 1.20$	**0.063** [0.062, 0.064]

d: temporal decay exponent;
d=0 recovers accumulated frequency.

Boston: Power law performs best, but only marginally improves over Frequency.

In YJMob, temporally weighted visit history provides a clear descriptive gain over accumulated frequency.

Summary: The power-law model performs best in held-out return choice across both datasets, providing the strongest empirical basis for the subsequent temporal-history diagnostics.

Question 1: Does the timing of prior visits provide additional information about return choice?

2 Three-Stage Diagnostic Framework: visit frequency, spatial proximity, and temporal history

Stage	Diagnostic focus	Boston ΔNLL	YJMob ΔNLL
I	Frequency sufficiency	0.004 [0.003, 0.006]	0.140 [0.138, 0.142]
II	Spatial proximity	0.000† [0.000, 0.000]	0.051 [0.050, 0.052]
III	Temporal history	0.000‡ [0.000, 0.000]	0.0038 [0.0035, 0.0041]

d=0: timing become irrelevant and the model recovers frequency.

Stage	Temporal specification	Boston d	YJMob d
I	Full temporal trace	0.05	0.45
II	Full temporal history with distance	0.00	0.30
III	Most recent visit with distance	0.00	0.25

Does timing matter after frequency is matched?
Timing improves held-out choice in both datasets, which a substantially larger gain in YJMob.

Does temporal information remain after controlling spatial proximity?
Distance absorbs the temporal gain in Boston, while a substantial contribution remains in YJMob.

Do earlier visits add temporal information beyond the most recent visit?
Full temporal history adds no measurable gain in Boston and only a small gain in YJMob.

Summary:

- The latest visit captures most of the temporal signal.
- Visit timing contains choice-relevant information beyond visit frequency, but its contribution beyond distance and the latest visit is context dependent

Question 2: Do recurrence signatures vary across destinations, and is this variation geographically organized?

I Descriptive Analysis: Recurrent destination demand has both temporal and geographic heterogeneity

Distribution Histogram

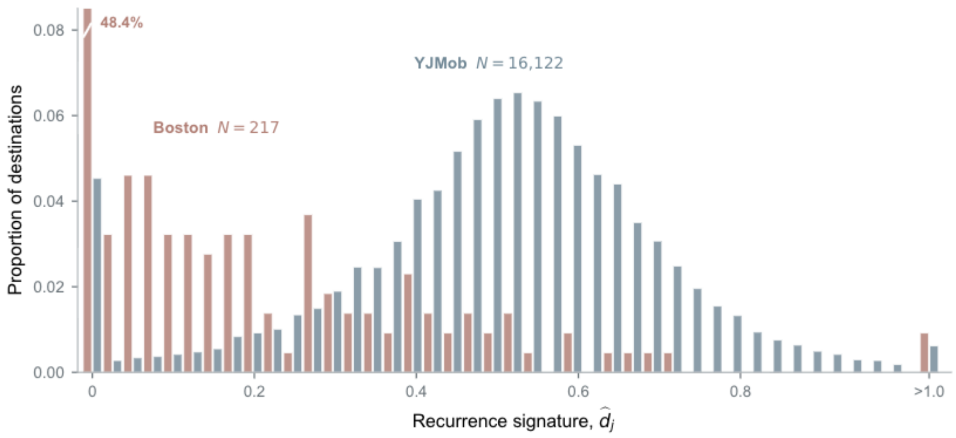

- Boston recurrence signatures are concentrated near the frequency boundary
- YJMob shows a broad distribution of recurrence signatures.
- The maps reveal spatially uneven concentrations of frequency-leaning and more recency-weighted destinations

Spatial Distribution

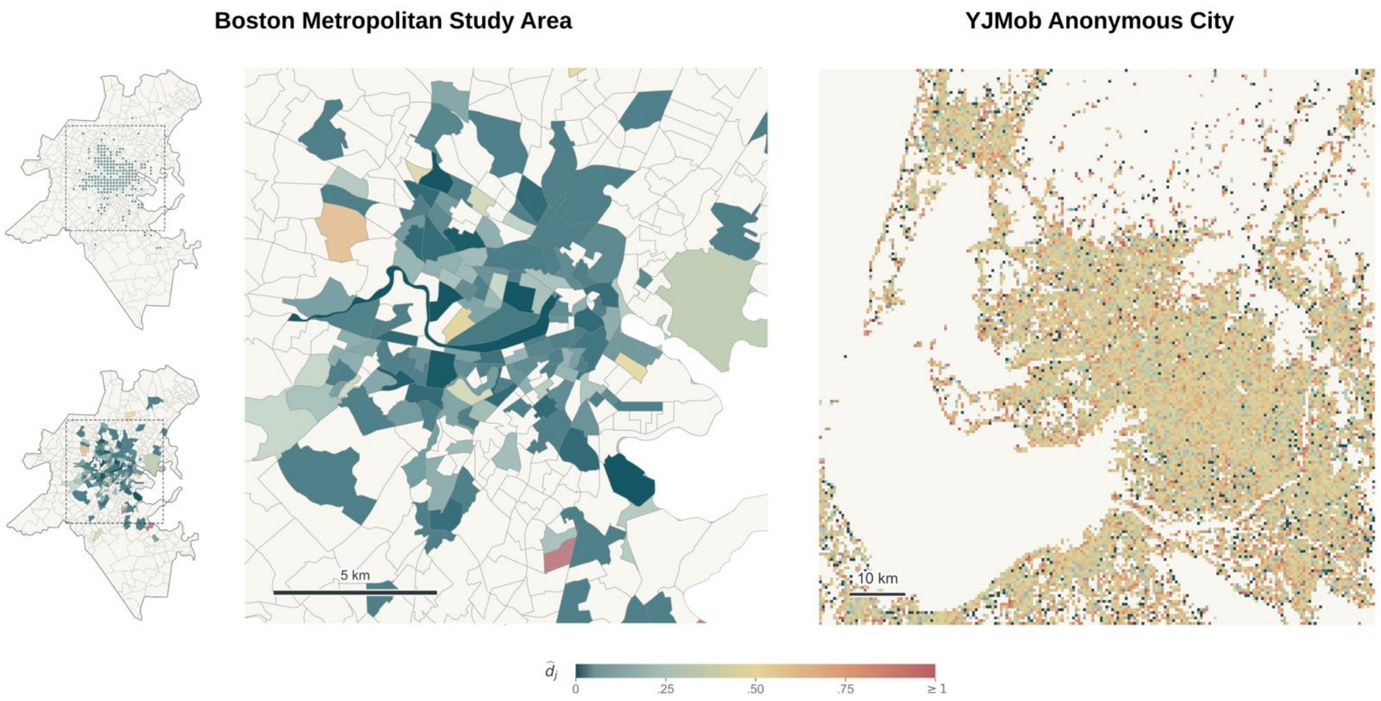

Summary: Urban destinations differ not only in how often they attract return visits, but also in how the timing of prior visits remains associated with subsequent return.

Question 2: Do recurrence signatures vary across destinations, and is this variation geographically organized?

2 Environmental association: Is destination-level return structure associated with local opportunity density?

YJMob Association Plot

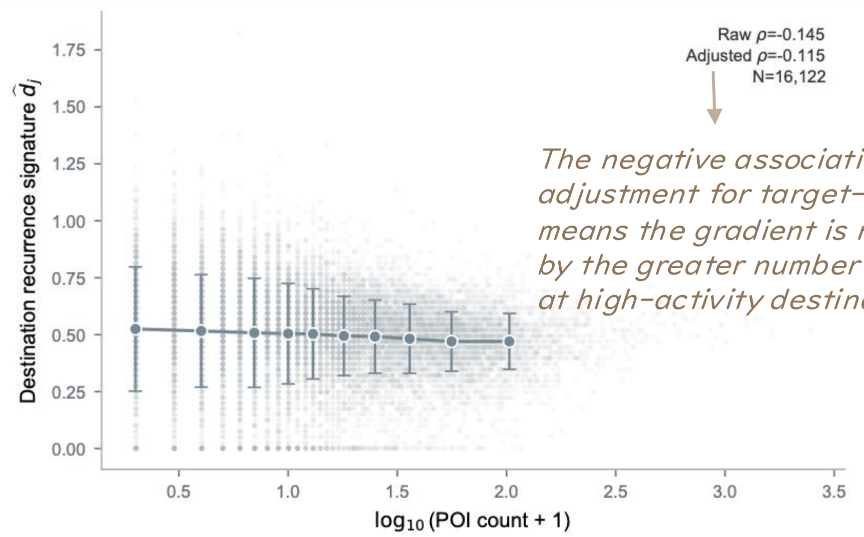

The negative association remains after adjustment for target-event count, which means the gradient is not explained solely by the greater number of observed return at high-activity destinations.

- YJMob destinations with higher POI density tend to have lower recurrence signatures, i.e. their return histories are more frequency-leaning.

Boston Association Plot

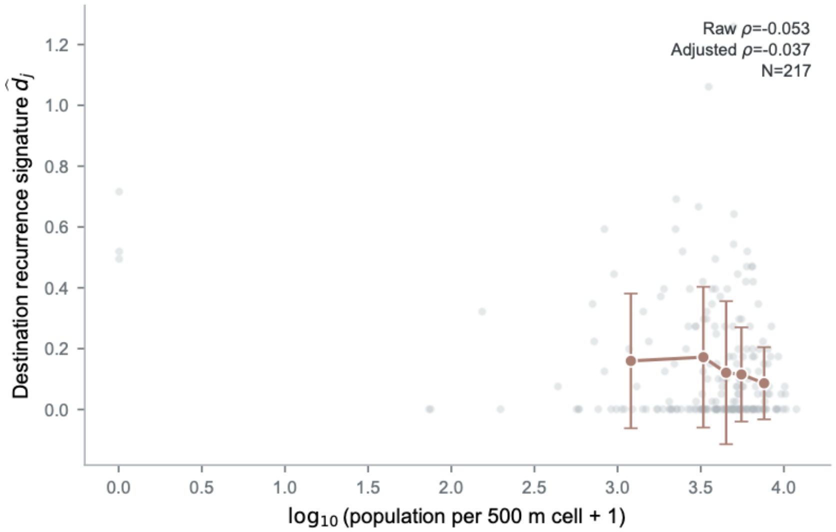

- Population density alone does not account for the observed destination-level variation in Boston.

Summary: destination context may organize temporal return structure, but the relevant geography is not captured uniformly by a single density measure.

Question 2: Do recurrence signatures vary across destinations, and is this variation geographically organized?

3 Spatial organization: Does geographic clustering remain after accounting for event support and opportunity density?

Larger positive I: stronger spatial autocorrelation
Smaller p: stronger evidence against random spatial arrangement

Dataset	Adjustment	Moran's I	Permutation p
YJMob	Unadjusted	0.070	0.001
YJMob	Target-event count	0.061	0.001
YJMob	Target-event count and POI density	0.054	0.001
Boston	Unadjusted	0.099	0.022
Boston	Target-event count and population density	0.093	0.017

- Destination recurrence signatures exhibit positive spatial autocorrelation in both empirical settings, which indicates neighboring destinations tend to show similar temporal weighting of prior visits.

- YJMob and Boston retain positive spatial autocorrelation after adjustment for target-event count and POI/population density.
- The spatial pattern is not explained solely by uneven event support or the observed POI-density gradient.
- Spatial organization can be present even when the available environmental measure shows no clear association.

Summary: The temporal structure of recurrent demand forms a localized urban geography that cannot be reduced to trip volume or a single density gradient.

5. Conclusion

Conclusion

1 A new way to measure return history

By nesting visit frequency at d=0, the framework places return histories on a common frequency-to-recency return rule. It makes the temporal contribution of past visits measurable and comparable across destinations and urban settings.

2 Timing matters, but its role is context dependent

Visit timing remains informative after prior frequency is matched in both datasets. After distance is controlled, YJMob retains a clear temporal contribution, whereas Boston lies at the frequency boundary. Moreover, the full visit history adds little beyond the most recent visit.

The evidence therefore supports temporal structure more strongly than a universal decay rate or a strong cumulative-memory mechanism.

3 The temporal structure of return has a geography

Return behavior is therefore not only spatially concentrated: the influence of past visits is itself spatially patterned.

Urban return has a context-dependent temporal geography
Urban memory geography makes the temporal influence of past visits measurable, comparable, and mappable.

Limitations and Future directions

	Limitations	Future Directions
I Individual decay rate: lack of ground truth	• No heterogeneity pattern in the individual level. • The dataset contains no independent cognitive or clinical ground truth. • d should be interpreted as a behavioral parameter derived from mobility patterns, rather than as a direct measure of memory ability or mental health.	• Collect or link mobility trajectories with cognitive assessments, Alzheimer's-related measures, depression or bipolar disorder survey data from the same individuals. • Examine whether heterogeneity in mobility-derived d is associated with cognitive status, particularly among older adults. • If EEG or fMRI data become available, LLM-derived representations and neural encoding models could further be used to investigate how brain responses and cognitive processes relate to experiences of urban mobility.
2 EPR framework extension	• Does not incorporate temporal return mechanism into a complete EPR model or conduct joint trajectory simulations. • Focuses on understanding and measuring return behavior rather than developing a general mobility prediction model.	• Extend the classical EPR framework by improving both its return and exploration components, such as incorporating multiple activity centers, and evaluating the resulting model through trajectory simulations.
3 Empirical meaning and prediction ability	• The decay rate and the destination-level recurrence signatures mainly serve as indicators of how past visits shape return behavior. Their broader empirical meaning and potential applications still need to be established.	• EPR is an interpretable mechanism-based model, while most advanced model may achieve stronger mobility prediction. • Compare or combine these approaches, with the goal of retaining behavioral interpretability while improving predictive accuracy.

THANK YOU!