From MAC De-anonymization to Digital Twins

Securing Wireless Mobility Data in the Age of AI

Pablo Serrano Yáñez-Mingot

Brescia, October 15, 2025

Acknowledgements

This presentation is partly supported by

The Project TUCAN6-CM (TEC-2024/COM-460), funded by CM (ORDEN 5696/2024)

The Project 6GINSPIRE PID2022-137329OB-C42, funded by MCIN/AEI/10.13039/501100011033/

The SNS JU EU's HE research and innovation programme under Grant Agreement No. 101192035 (AMAZING-6G)

Motivation

Analyzing RCM in the Campus WLAN

- The wireless activity of mobile devices leaves a trail of information that can be used to unequivocally identify users.
- Four spatio-temporal points are enough to identify 95% of individuals in a large mobile cellular network [1]
- Randomized and Changing MAC Address (RCM): different MAC per SSID
  - Persistent: one per SSID
  - Non-persistent: change every 24 h

[1] Montjoye, Yves-Alexandre & Verleysen, Michel & Blondel, Vincent. (2013). Unique in the Crowd: The Privacy Bounds of Human Mobility. Scientific reports. 3. 1376. 10.1038/srep01376.

The Campus WLAN

Eduroam @ UC3M (Leganes)

- 278 access points (APs)
- 7 buildings
- 10k users
- 16k devices
- Are devices “unique in the crowd”?
  - 100x less users
  - 8x higher density

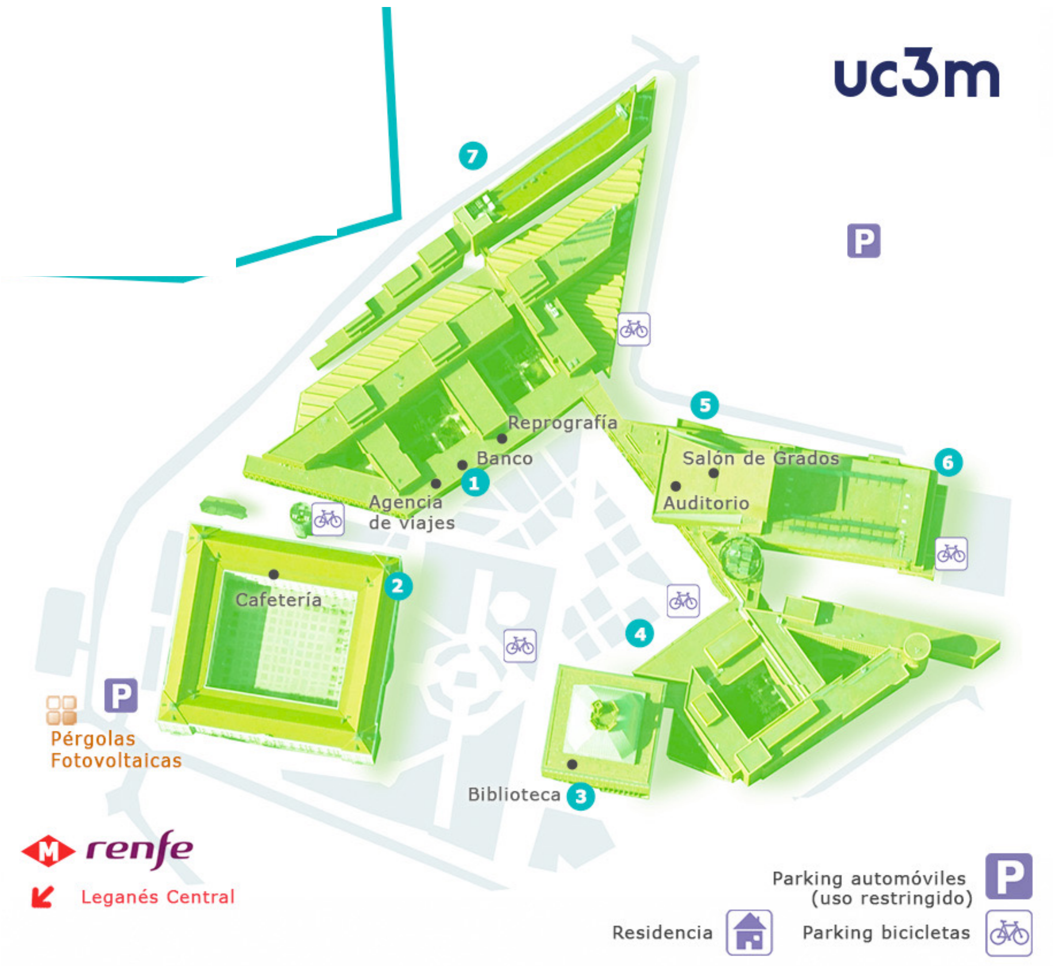

Data collection process

Eduroam – federated Radius

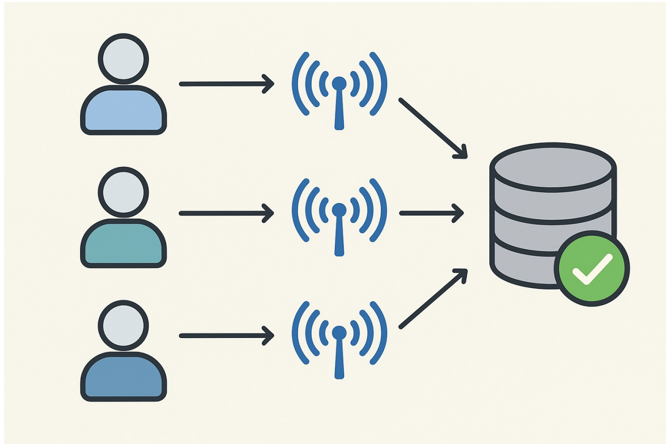

- Each time a device
  - Associates, or
  - reassociates with an AP
- the RADIUS server logs it
- The status is updated at least every 15 minutes.
- Each entry:

<timestamp, user identifier, client addresses, AP address, traffic info>

MD5 hash MD5 hash

(also) Unique in the Campus WLAN

‘p’ random spatio-temporal APs vs top ‘p’ APs

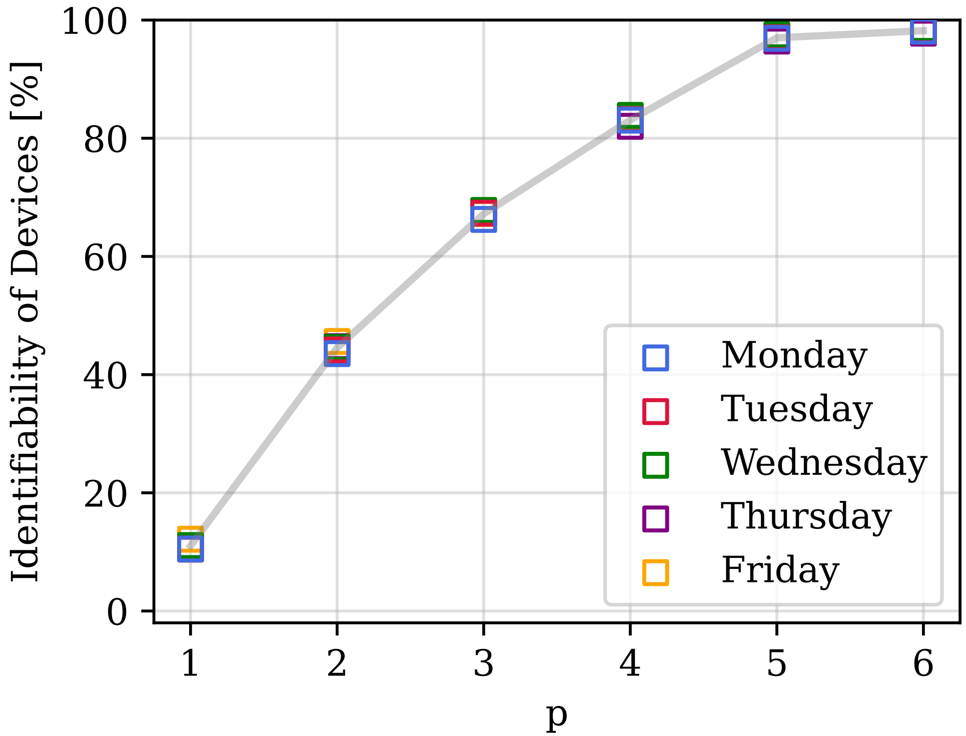

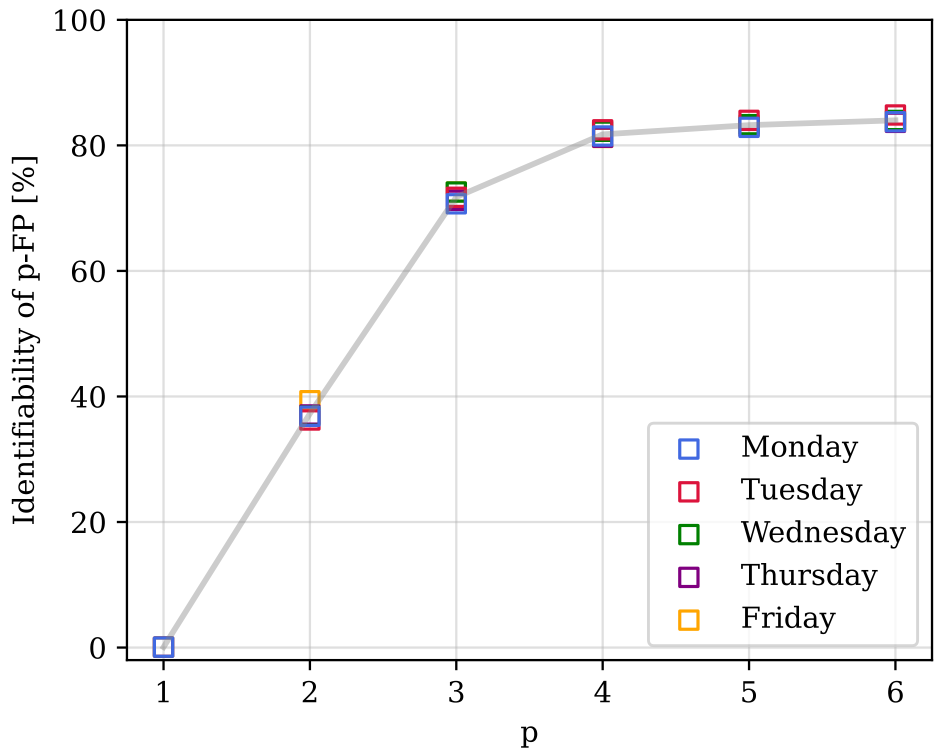

Conclusions (1/3)

High uniqueness in the Campus WLAN

- Despite the differences vs. “Unique in the crowd” [1]
  - In size & density
  - And population and schedule
- There seems to be strong individualizing information in the logs
- Can we identify some patterns and unequivocally identify users?
  - This would render (non-persistent) RCM useless
  - Explainable identification -> design better schemes

XD-RCM: eXplainable Deanonymization of RCM

Approach

- Analyze a set of explainable features during some time
    - Arrival and departure times
    - Number of different APs visited
    - Most frequent Aps
    - Downloaded traffic
- Use them to re-identify devices after they changed the MAC
    - I.e., we assume that at some point the user activates non persistent RCM

Small data set

Following (strict) data protection guidelines

- 28 explicit volunteers
    - Mostly faculty members
- 98 different devices
- 5 months of data
- We restrict the analysis to a single building
    - 3 floors + basement
    - 47 APs

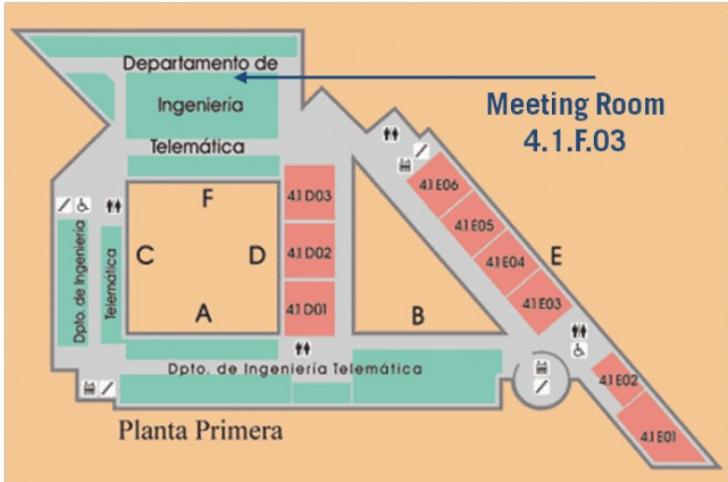

Arrival, Departure, & Total times

- Majority of devices appear around 8AM
- Most departures concentrate around 6PM
- A lot of devices are always connected (permanent)
    - And in many cases, to the same AP (static)

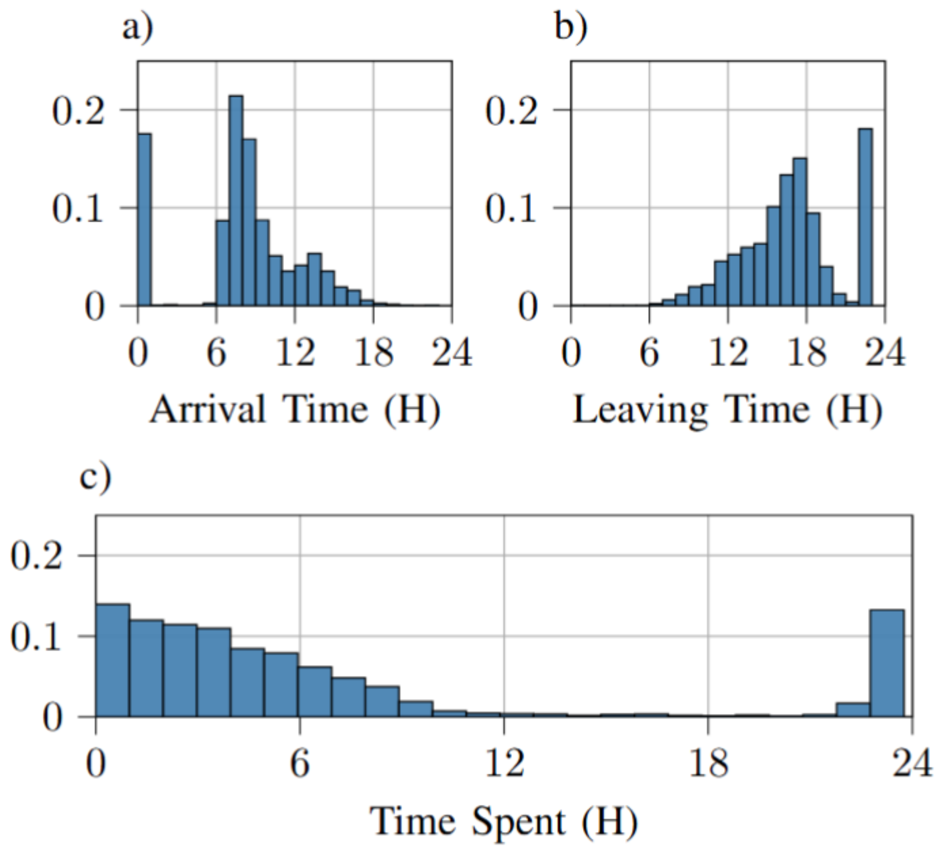

Number of different APs visited

For different time periods

- For those devices that visit more than 1 AP (i.e., non static)
- One day: ~ 3 APs
- One week: ~ 6 APs
- 2 weeks: ~ 10 APs
- 1 month: ~ 10 APs

(Note that we consider 1 building)

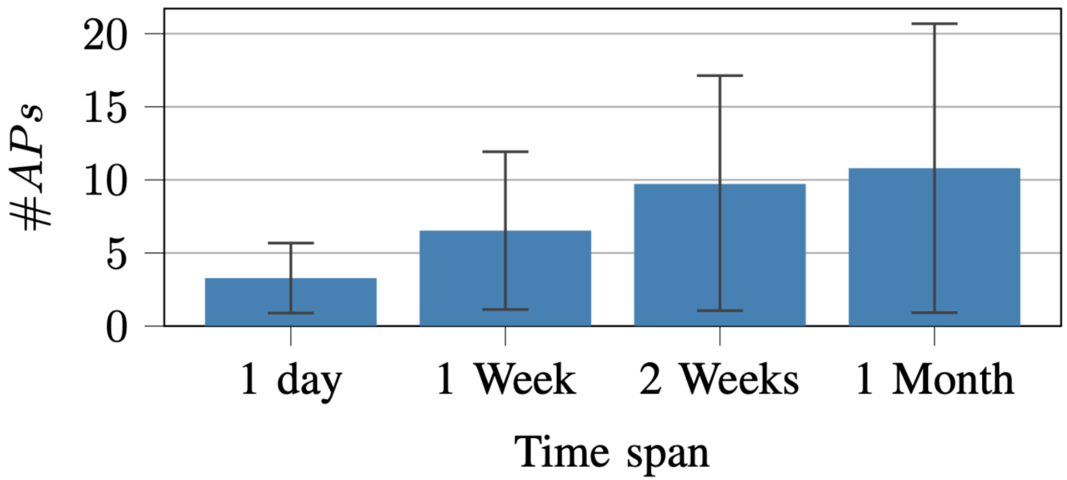

Entropy (different APs + relative time)

For the same time periods

- Defined as $H = \sum_{i=1}^{\#AP} p_i \log_2(p_i)$
- Effective number of locations
$$L = 2^H$$
- 20% devices: only one AP
- Increases with time window
- But 2 weeks ≈ 1 month

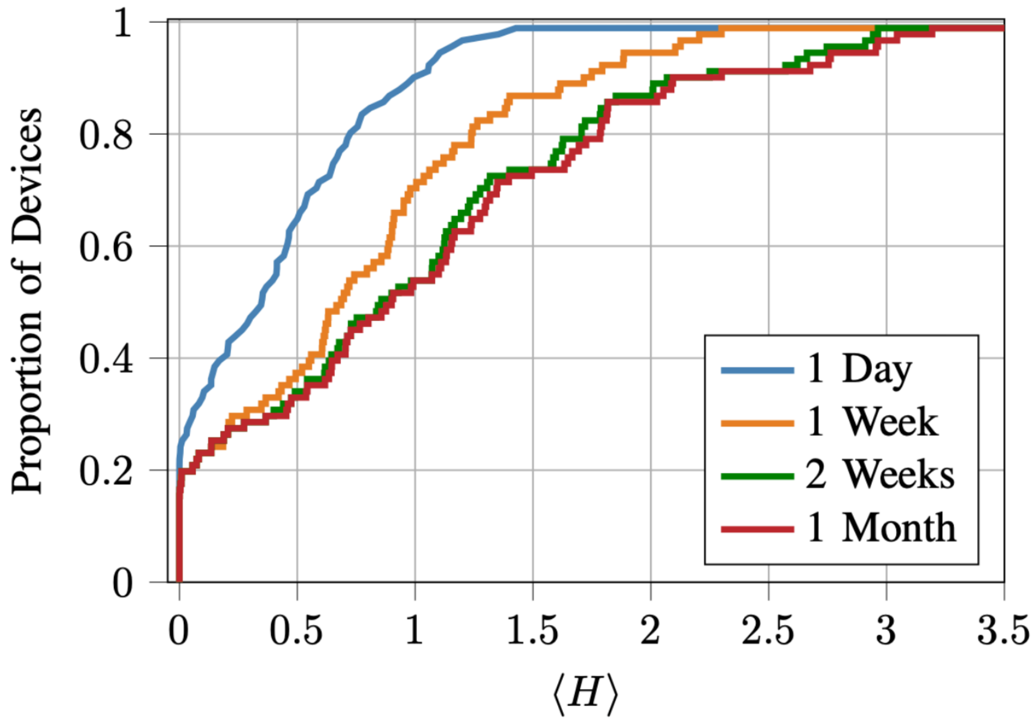

“k” most frequent APs

Idea: the k-tuple will identify users

- Collect the 6 most frequent APs
- Compute the cumulative time in decreasing order
- Devices spend 90% time on 2 Aps

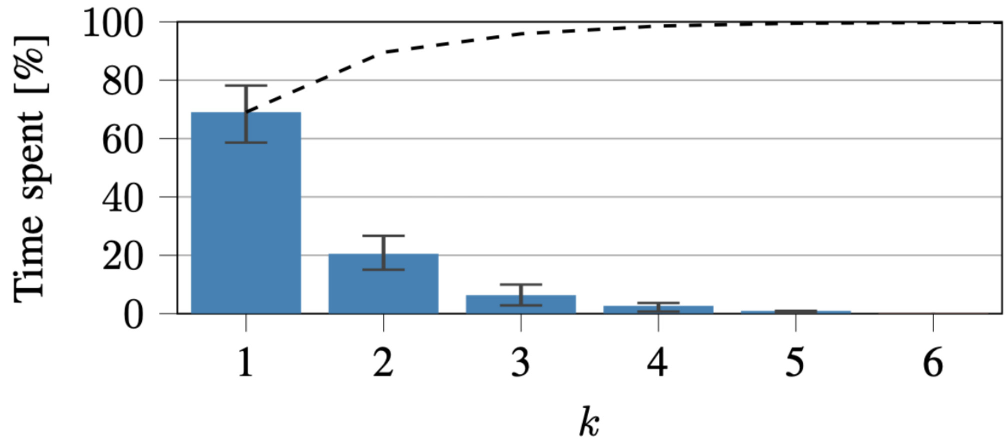

- With k=2, 88% devices are unique for a window of 1 day

Downloaded traffic per day

- On average, 350 MB/day
    - Spain: 400 MB/day

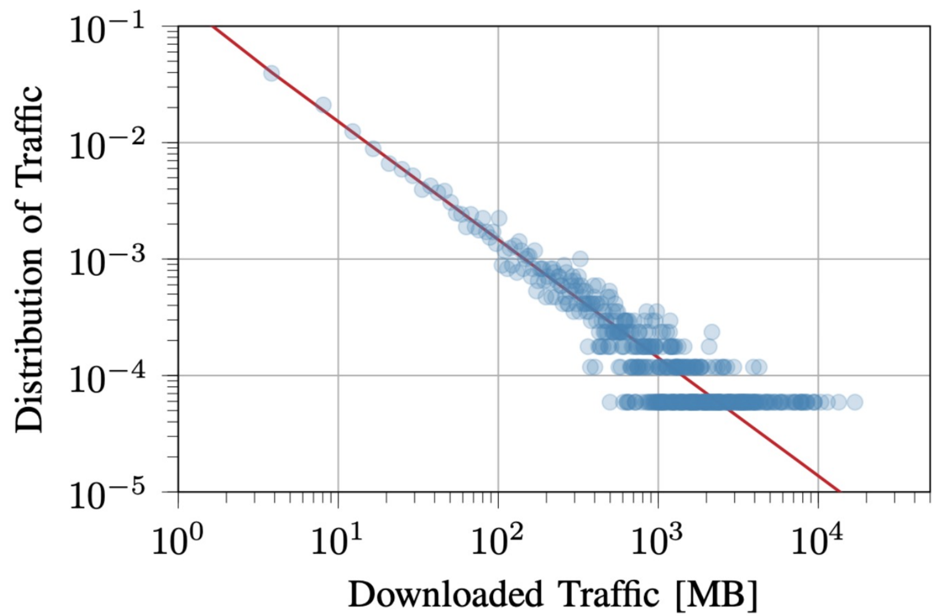

eXplainable DE-anonymization of RCM

Approach

Profile: $\mathbf{x}_i = (T_s,\ t_a,\ \#APs,\ \mathcal{H}, [APs], D)$

- Train the model with a labeled dataset to learn whether two user profiles, observed on different days, belong to the same device.
- Once trained, the model compares a given profile of a user with the historical profiles stored in the dataset

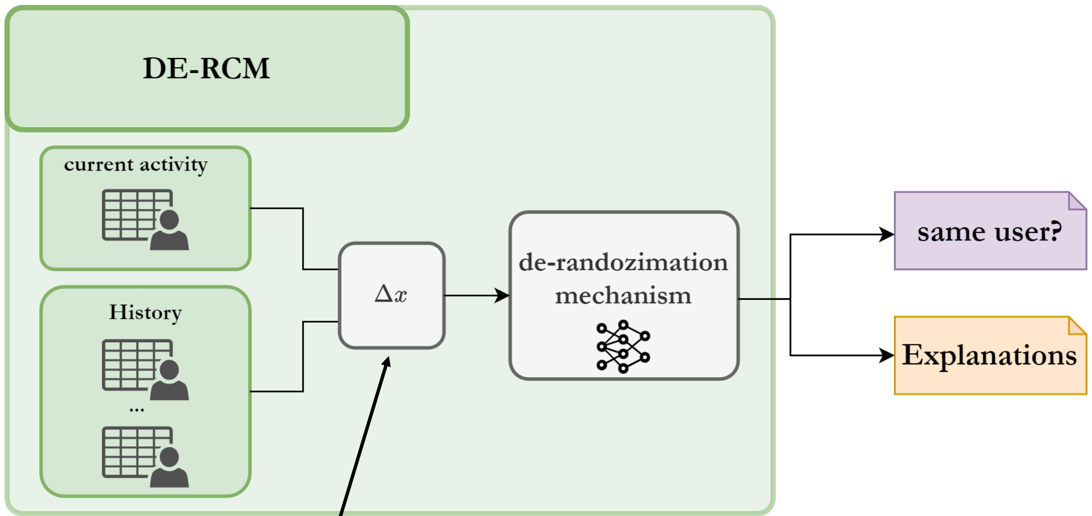

$(\mathbf{\Delta x}_{i,j}, \delta_{i,j})$ → Same/different device

Comparison

Three algorithms

- XD-RCM: based on Random Forests
- Decision tree
- Heuristic: the top k=2 APs

Results

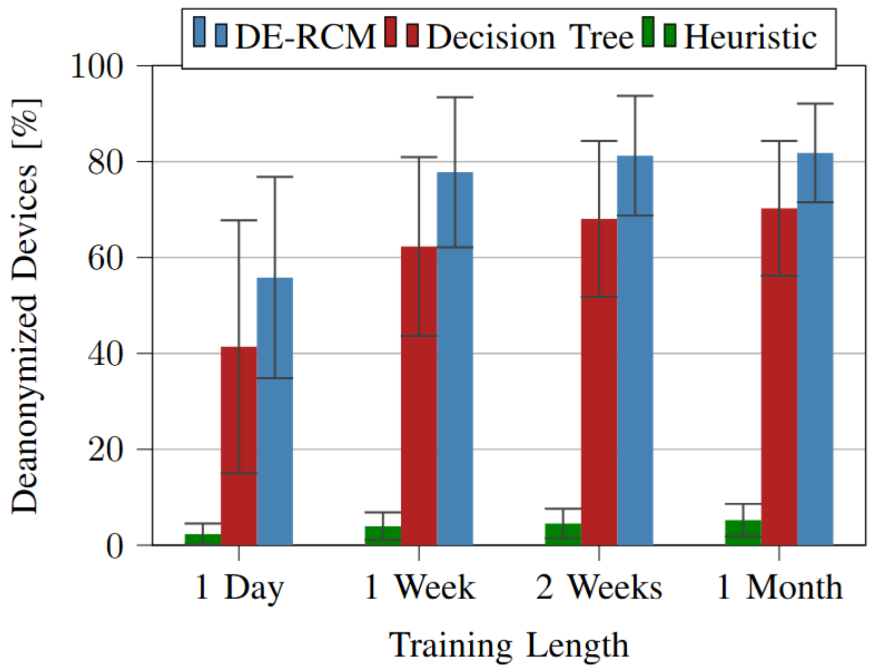
DE-RCM
Decision Tree
Heuristic
100
80
60
40
20
0
Deanonymized Devices [%]
1 Day
1 Week
2 Weeks
1 Month
Training Length

Results

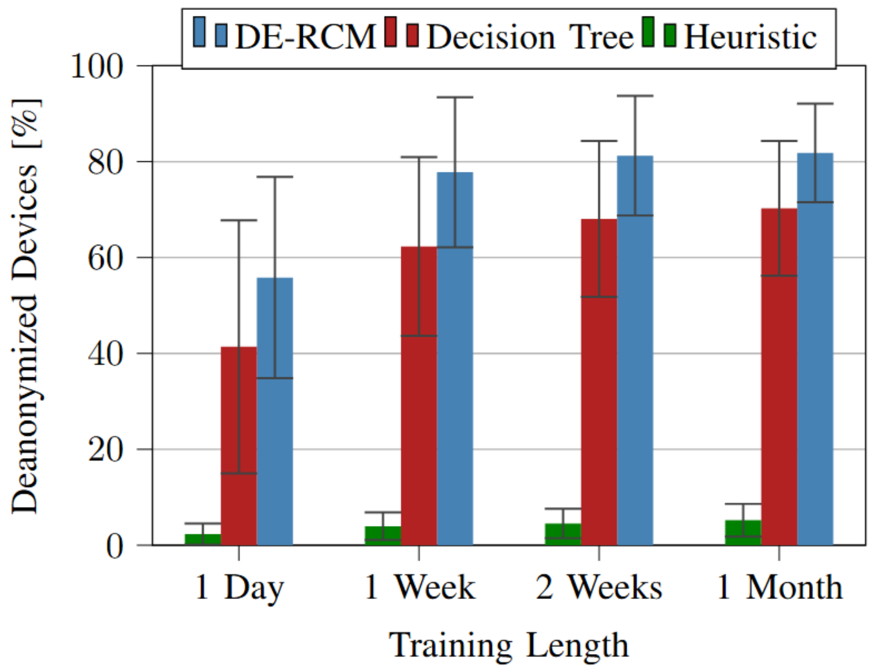

We used LIME (Local Interpretable Model-agnostic Explanations), which approximates the model's predictions with an interpretable model around specific instances.

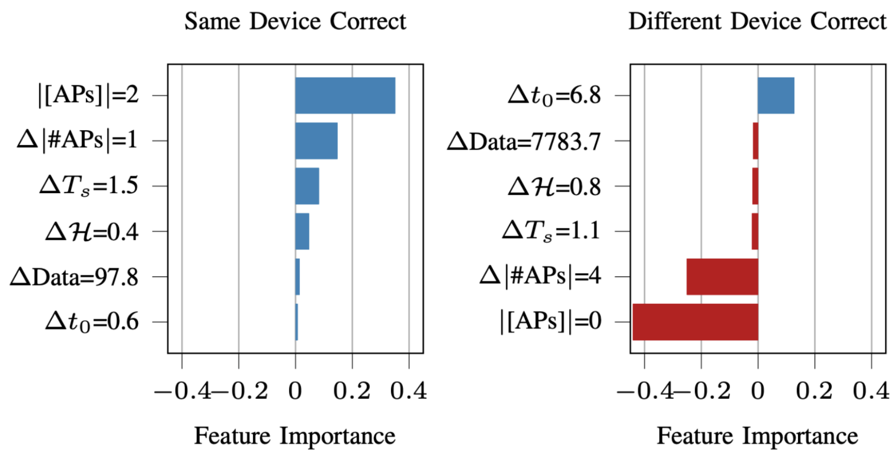

Conclusions (2/3)

- (For the case of our small dataset…)
- Devices are “less entropic” after 2 weeks
- Non persistent MAC is not enough to hide uniqueness
- Human-interpretable variables can be used to re-identify users with 80% accuracy
- Explainability could help design better de-anonymization techniques
- We need sound approaches to ensure privacy

DiWi: A Transformer-Based DT for WLANs

- The use of existing datasets is tricky (privacy considerations)
- Adding noise may reduce the utility
- But spatio-temporal datasets are useful, e.g.,
    - Anticipatory networking (caching, AP on demand, mobility)
    - Heating, ventilation, and air conditioning systems (HVAC)
- Approach: synthetic generation
    - For simplicity: discrete time

DiWi: motivation

- Our goal is to model the activity of users connected to the network
- Sequences of discrete spatiotemporal points (i.e., Access Points).
- Large Language Models (LLMs) learn from sequential data => we adopt a similar architecture to model these sequences of APs
- LLMs rely on a specific encoding of tokens (APs) and its relative position in the sequence.
    - These embeddings are then processed through self-attention layers, which learn the relationships between elements across time and space.

DiWi: desing

Overall architecture

- Sequentially encoding spatial and temporal components of device connectivity traces
- These are merged into a unified spatiotemporal representation
- The model predicts the next connectivity state: whether the device will remain connected to the same AP, transition to a different AP, or disconnect entirely.

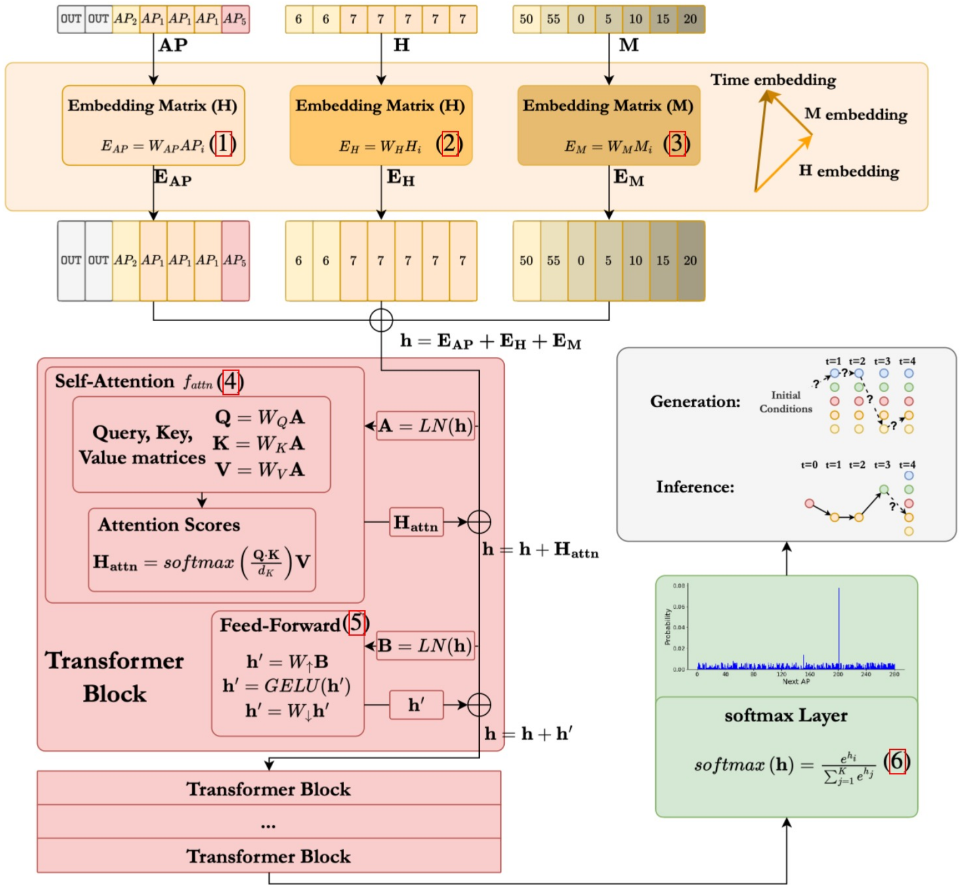

DiWi: spatial embedding

- We focus on the time between 6 AM and 10 PM (16 hours)
- Time is discretized in 5 minutes interval
- Position:
  - A user in one day: sequence of 192 tokens
  - Token vocabulary: 278 APs (campus) + ”**OUT**” token

DiWi: time embedding

- Relative positional embeddings are ok in natural language processing (e.g., to preserve word order)
- Absolute embeddings are better suited for mobility data, since the absolute position of a token within the timeline provides critical contextual information.
- For example, being disconnected at 7 AM ≠ being disconnected at 2 PM
- Time is decomposed in hours (H) and minutes (M)
  - More scalable than “absolute minute”
  - Avoids loss of temporal semantics and degradation on long sequences.

DiWi: spatio temporal embedding

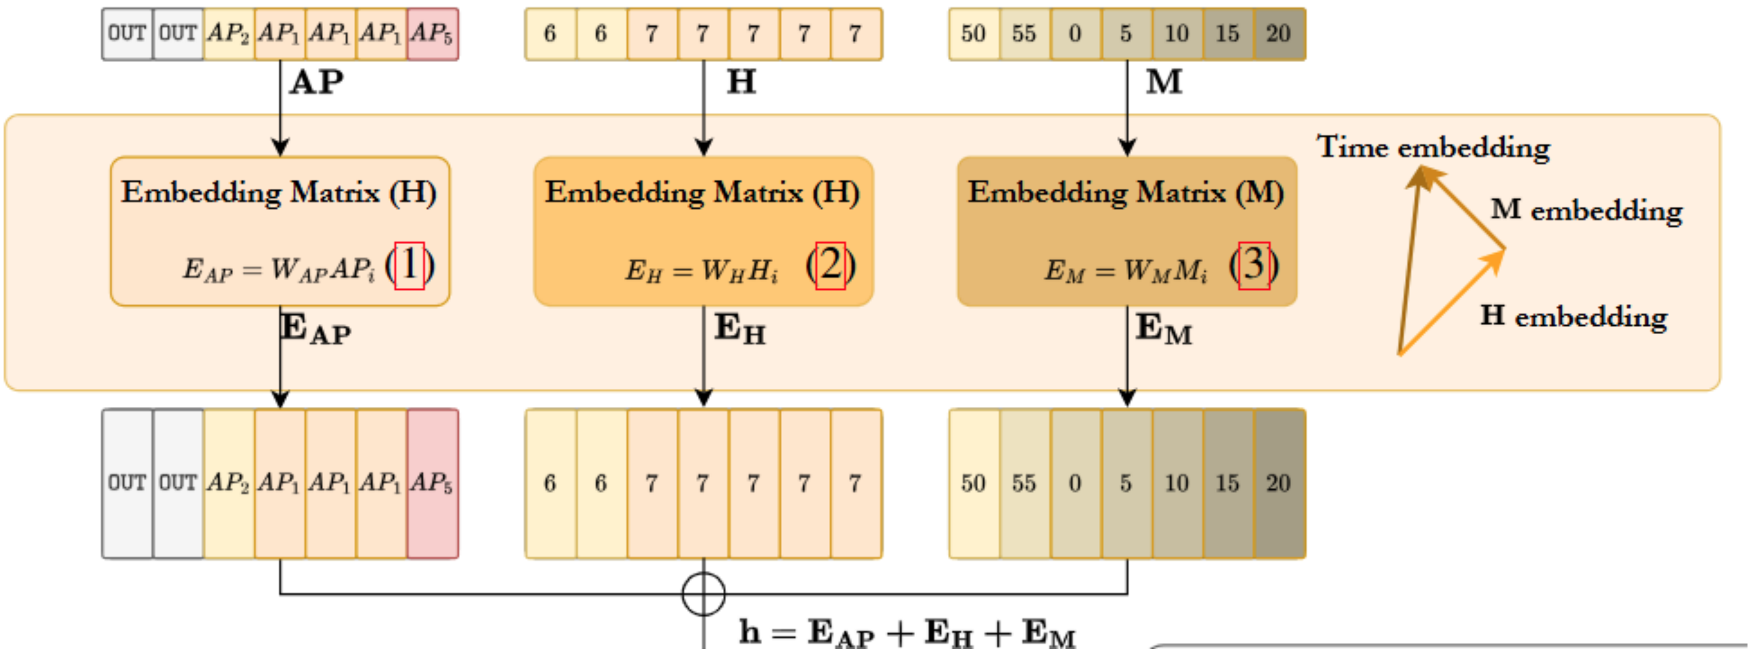

More details

- We use the complete pseudonymized dataset (MD5 hashes, 30k users)
- J. M. Montes-Lopez, P. Serrano, M. Gramaglia, A. Banchs, “DiWi: A Transformer-Based Digital Twin for Wireless Mobility,”, Elsevier Computer Networks, October 2025. 10.1016/j.comnet.2025.111571

Performance Evaluation

DiWi as mobility predictor

- Ability to predict the next AP
- Benchmarks
    - A standard LSTM network
    - GPT-2 model without absolute time

Model	**Campus 1** Acc. [%]	**Campus 2** Acc. [%]
LSTM Network	91.2	89.8
GPT-2	91.8	91.9
DiWi	**92.3**	**92.4**

- Better performance => Ability to identify temporal information (GPT2, with a flat encoding, cannot easily capture)

Performance evaluation

DiWi as synthetic traffic generator

- Each synthetic trace (device) starts with a token drawn from the distribution of first states seen in the real data (i.e., 82% in OUT)
- From that “seed” the model produces a probability vector for the next state
- We sample a token from it, slide the context window to keep only the most recent tokens, and repeat the process.
- Generation stops when the trace reaches a length of 192 tokens (one day)

Performance evaluation

Mobility statistics

Metric	LSTM	GPT2	DiWi
Nº visited APs	0.26	0.13	**0.015**
AP rank	0.083	0.076	**0.036**
Time spent	0.62	0.65	**0.017**
Time of arrival	0.078	0.059	**0.032**

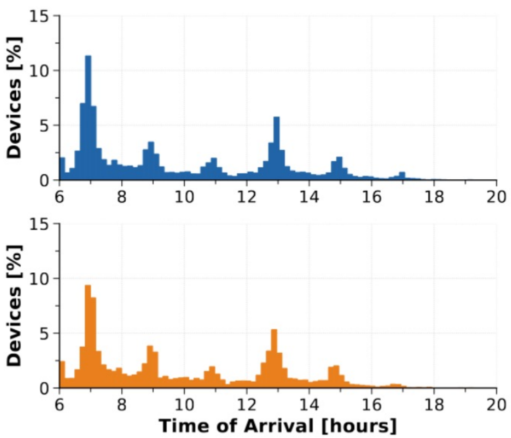

(c) Distribution of arrival times.

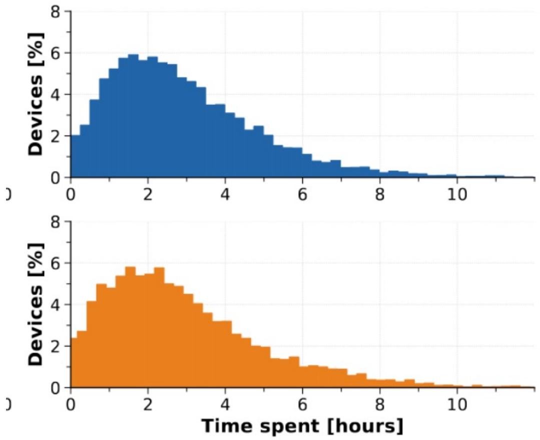

(d) Distribution of time spent.

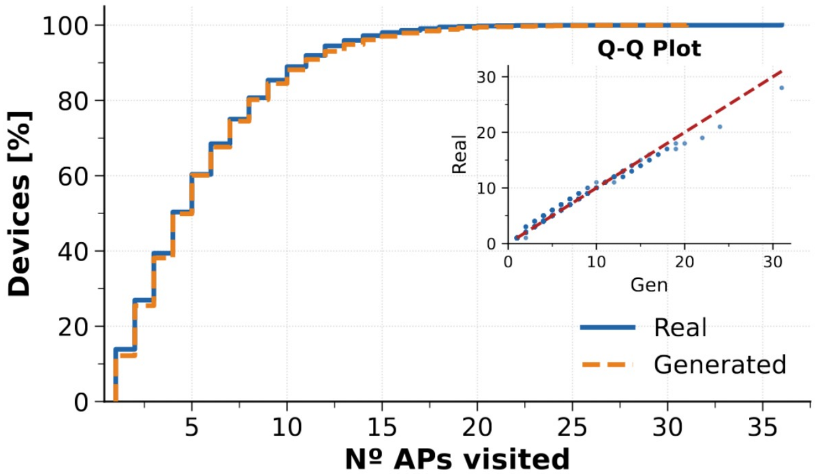

Performance evaluation

Design of HVAC systems

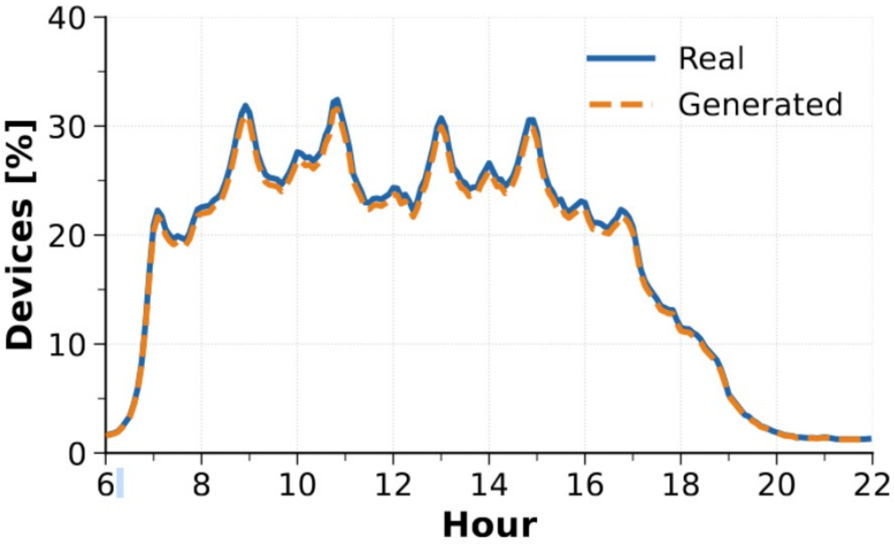

(a) Occupancy of the Campus.

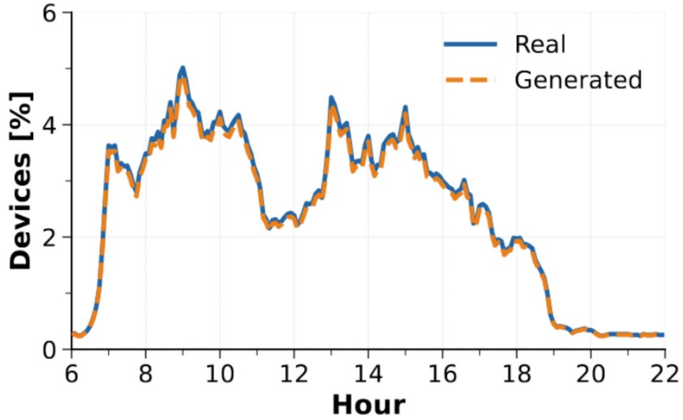

(b) Occupancy of Classrooms.

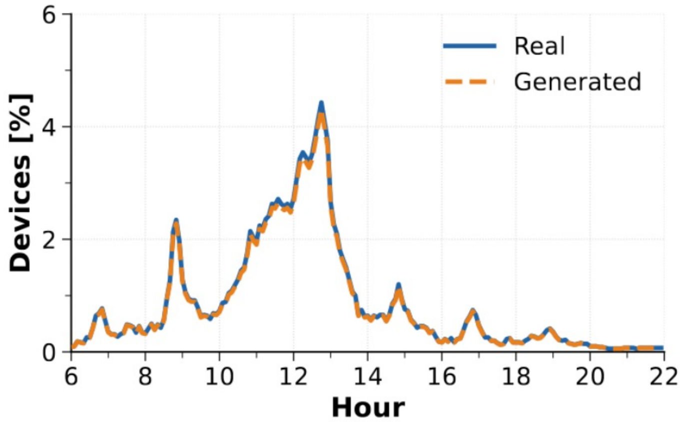

(c) Occupancy of Cafeteria.

Performance evaluation

Uniqueness of synthetic traces

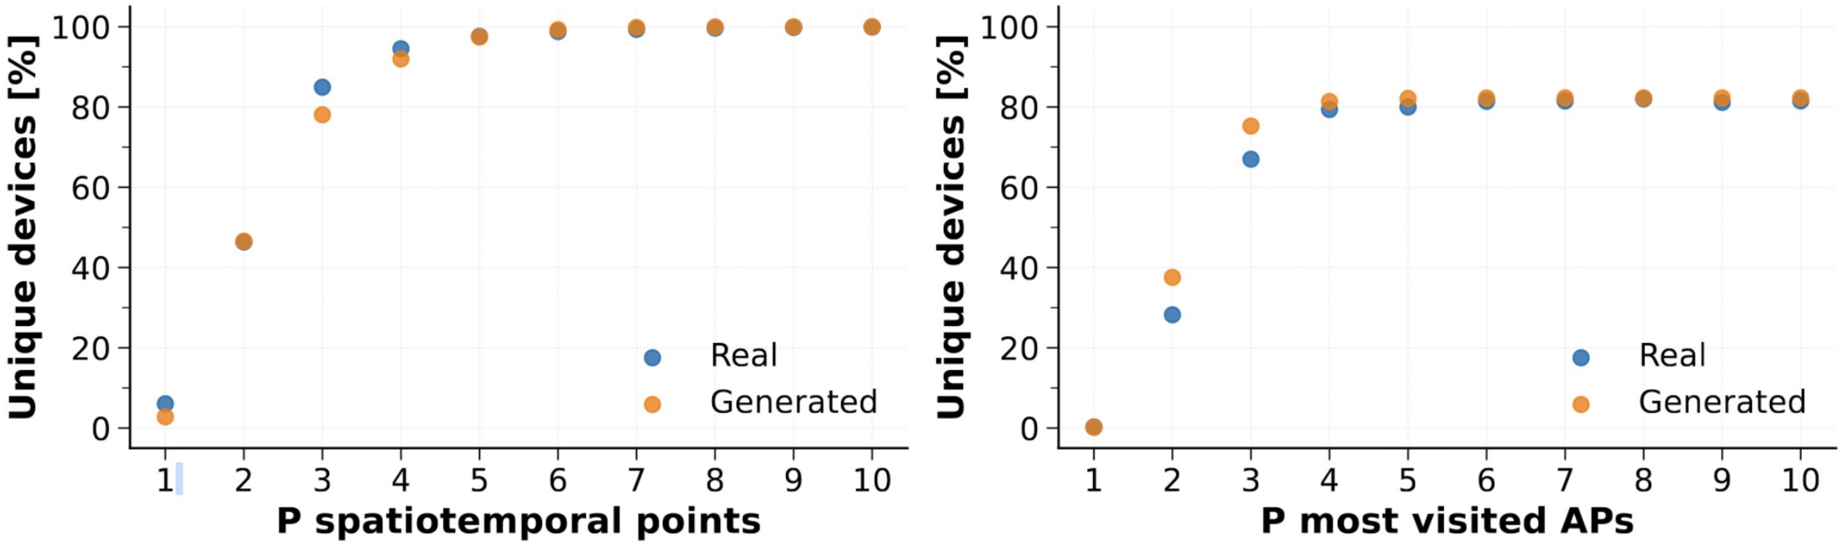

Privacy assessment

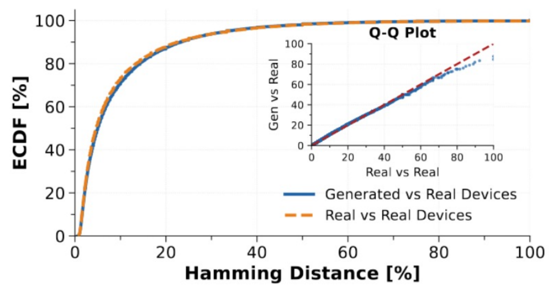

Figure 7: Hamming Distance between traces.

- Three heuristic analyses
  - Average probability of generating a real trace: 10^{-120}
  - Synth traces are as similar to real traces as real traces are to each other
  - Membership inference attack: ~ random guess (50%)
- Formal guarantees
  - DiWi can be extended with differential privacy (worse performance)

Conclusions (3/3) and Future work

- LLMs are good at capturing relations
- Use of absolute time is convenient
- H + T encoding
- Formal privacy guarantees

- Continuous time
- Real HVAC systems
- Public tool

Additional information

- Juan Manuel Montes-Lopez, **Pablo Serrano**, Marco Gramaglia, Aruna Prem Bianzino, “DE-RCM: Desanonimización Explicable de MACs Aleatorias en 802.11 WLANs,” Jornadas de Ingeniería Telemática (Jitel 2025), Cáceres, Noviembre 2025
- Lucía Cabanillas, Juan Manuel Montes-Lopez, Diego R. López, **Pablo Serrano**, ”DEBAC: Dynamic Explainable Behavior-Based Access Control,” 2025 EuCNC & 6G Summit, June, 2025
- J. M. Montes-Lopez, **P. Serrano**, M. Gramaglia, A. Banchs, “DiWi: A Transformer-Based Digital Twin for Wireless Mobility,”, Elsevier Computer Networks, October 2025. 10.1016/j.comnet.2025.111571