Geo-Print Modelling for Privacy-Preserving Digital Identity: Mobility-Based Spoof Detection and Behavioural Segmentation

Author: Edwina Narulita SAJ [1] Supervisor: Dr Luisa Cutillo [1] Alexei Poliakov [2]

[1]School of Mathematics, University of Leeds [2]Locomizer

Background

This project is undertaken in collaboration with Locomizer as part of the project *I-AM-HUMAN: Zero-Knowledge Digital Identity from Geo-Fingerprinting*, which investigates an alternative form of identity verification based on location behaviour rather than biometric identifiers, document uploads, or other sensitive personal information. Locomizers patented technology transforms raw user location traces into affinity indexes, known as geo-prints, which capture behavioural patterns across categories of Points of Interest (POIs). These geo-prints provide privacy-preserving, data-driven signatures of mobility that can be applied to digital identity verification and fraud prevention. By emphasising behavioural regularities in space and time, geo-fingerprinting offers a novel basis for authentication that is resistant to deepfake manipulation and compliant with data protection regulations such as the GDPR.

The study pursues two objectives. The first is to examine *humanness detection*, testing whether session-level geo-print features can reliably distinguish genuine mobility traces from synthesised spoofed sessions, including for unseen users and months. The second is to evaluate *segment-level distinctiveness*, assessing whether geo-prints yield stable, interpretable user segments and identifying the principal geographic and behavioural drivers of segmentation.

The project contributes to the development of zero-knowledge proof digital identity systems, enabling verification of authenticity without disclosure of sensitive personal information. Through this approach, it establishes an empirical foundation for privacy-preserving and AI-resilient verification mechanisms with potential applications across finance, advertising, e-commerce, and digital health platforms.

Data and Methodology

For machine learning tasks in spoof detection, anomaly detection, and geographic analysis, the study employed data describing user affinities, mobility, and spatial context. Affinity data reflected category-level interaction scores, while mobility data consisted of GPS traces (longitude, latitude, timestamp) limited to pedestrian and stop activities between 09:00 and 21:59 for users active across July–September 2024. Geographic information, linked to official postcode boundaries, was used to infer representative home and workplace areas. Synthetic trajectories were also generated by perturbing genuine traces to create spoofed sessions. These integrated sources were harmonised, quality-checked, and merged with binary labels, from which user-level features were engineered to capture activity intensity, mobility dynamics, and geographic attributes.

Supervised models (Logistic Regression, Random Forest, XGBoost) distinguished real from spoofed sessions under user-disjoint and temporal hold-out conditions. Unsupervised clustering (K-Means, DBSCAN) identified behavioural segments, while Isolation Forest, Local Outlier Factor, and a profile-based risk score detected anomalous users. Validation employed ROC–AUC, PR–AUC, and Brier score, alongside permutation-based feature importance and cluster stability assessed through adjusted Rand index and cross-algorithm agreement.

Key Findings

The session-level spoof detector generalised robustly to unseen time periods. On the September hold-out it achieved **ROC–AUC 0.829** and **PR–AUC 0.815** (**Fig. 1**), with **well-calibrated probabilities** (Brier = **0.162**; **Fig. 2**) and no material change after isotonic calibration. At the operational setting $\tau = 0.50$, recall was **0.935** and precision **0.654**, prioritising sensitivity and accepting a higher review burden. The confusion matrix illustrates this trade-off, with many true spoofs recovered at the expense of additional false alarms (**Table 1**).

Table 1: Confusion matrix on the test set at the operational threshold (rows = true class, columns = predicted).

	Pred. real (0)	**Pred. spoof (1)**
True real (0)	1710	1668
True spoof (1)	221	3157

Operating points were tuned against application costs. Cost-sensitive thresholds in the 0.580.64 range increased precision to **0.830.84** with recall around **0.520.53**, offering a leaner alert stream for workflows where false positives are costly. In practice, the choice between $\tau = 0.50$ and a higher threshold reflects whether the deployment values missed-spoof minimisation (high recall) or analyst efficiency (higher precision).

Model interpretability aligned with expected mobility regularities. Permutation importance showed that **median speed** and **radius of gyration** were the strongest predictors (**Fig. 3**), with distance travelled, stop counts, and session duration contributing secondarily. These features capture dispersion and cadencestable signals of genuine movement that synthetic traces struggled to mimic consistently.

Unsupervised analysis revealed two clear mobilitygeography segments**local residents (≈80%)** and **commuters (≈20%)**driven by commute distance and workplace location (**Fig. 4**), while profile intensity and category breadth were similar across groups. **Anomalies were rare (≈23%)** and concentrated at mobility extremes. Profile-only high-risk users overlapped more with Isolation Forest detections (Jaccard = **0.42**) than with LOF (**0.23**), indicating complementary signals. Data synthesis and quality checks were sound: from 10,000 sampled sessions, **88,351** synthetic rows were generated with no out-of-bounds points, and **48,020** duplicate timestamps were resolved (**Fig. 5**).

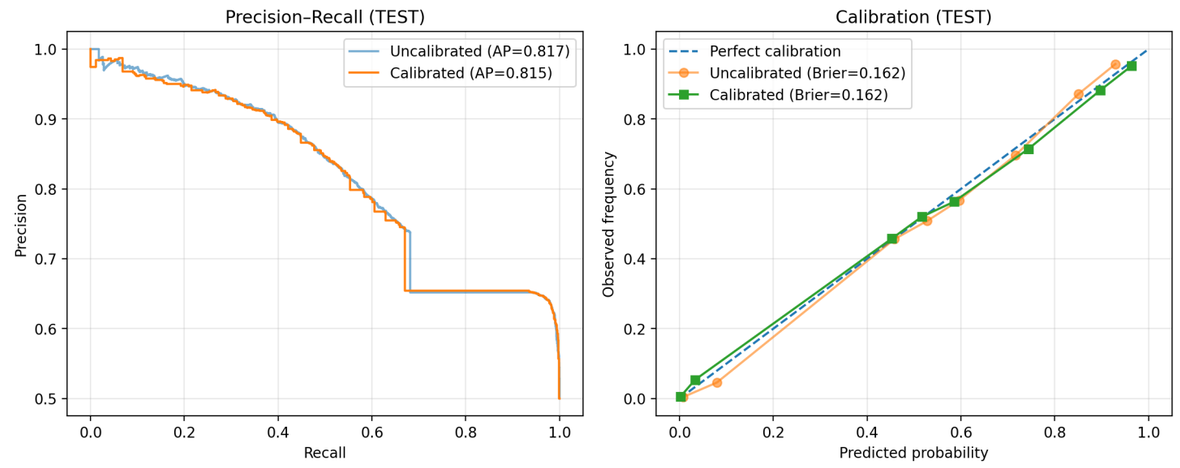

Figure 1: Precision–Recall (September hold-out). PR–AUC = 0.815; marker at $\tau = 0.50$.

Figure 2: Reliability diagram (Brier = 0.162).

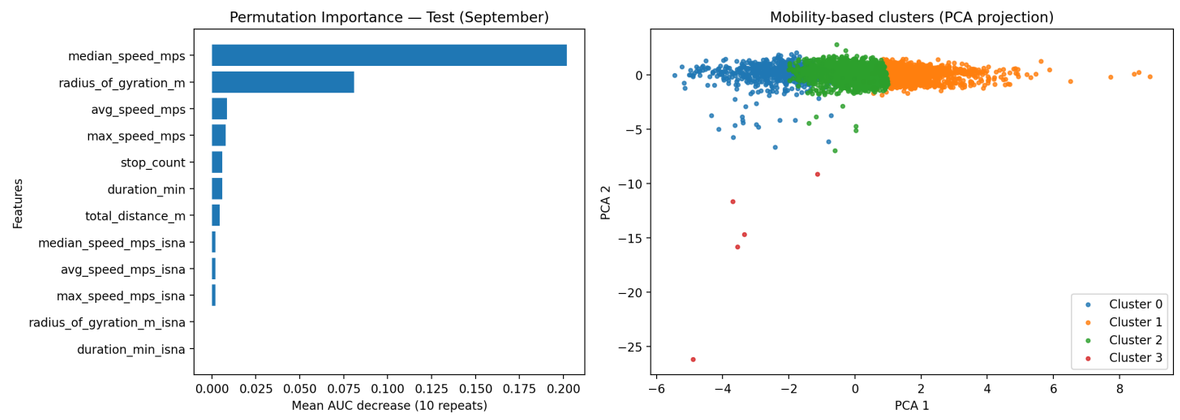

Figure 3: Permutation importance: median speed & radius of gyration dominate.

Figure 4: Mobility clusters ($k = 4$) in PCA space: commuters vs local residents.

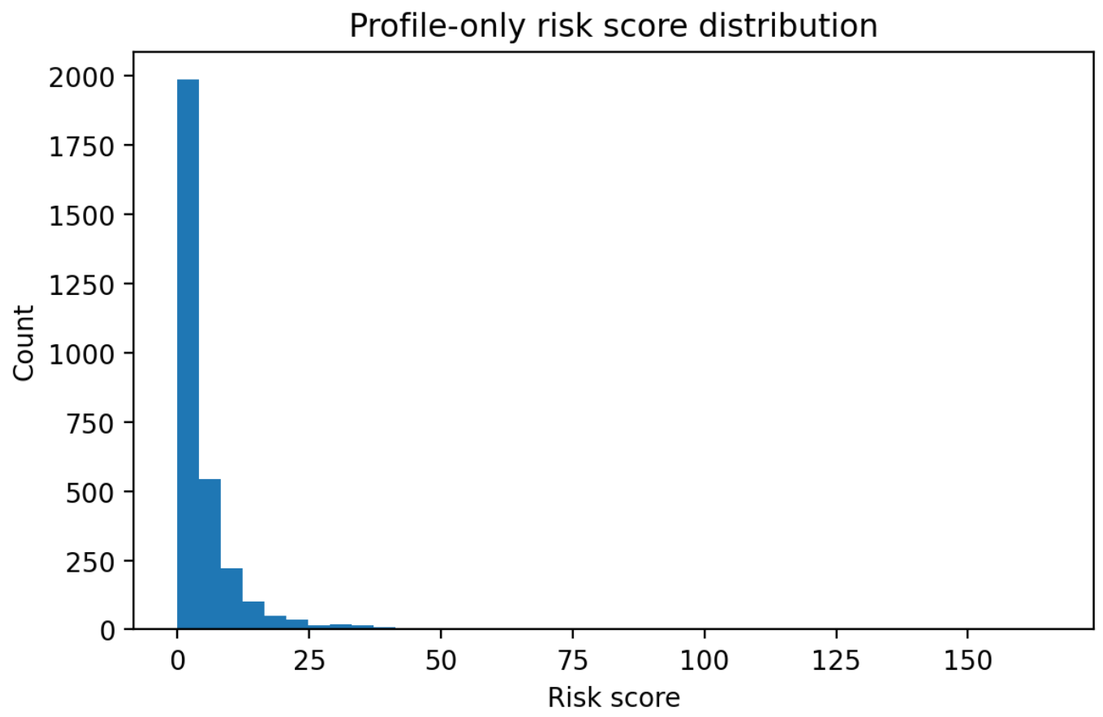

Figure 5: Profile-only risk scores; 98th percentile cut-off.

Value of the Research

This study demonstrates that geo-prints provide a privacy-preserving way to distinguish real from spoofed mobility traces, with calibrated models showing strong discrimination and reliable probability estimates. The identification of median speed and radius of gyration as the main behavioural determinants confirms that regularities in human movement form a stable signal of humanness. The integration of mobility with geographic indicators further enabled clear segmentation between local residents and commuters, driven primarily by commute trajectories and workplace locations. A combined workflow of calibrated session-level spoof detection followed by user-level anomaly screening offers an efficient and interpretable strategy for fraud mitigation, behavioural segmentation, and risk monitoring. These findings establish an empirical foundation for applying geo-prints to identity verification systems that are resistant to manipulation while remaining compliant with privacy requirements.