Machine Unlearning for Mobility Data: An Algorithmic Perspective

Department of Electrical Engineering and Computer Science

ALI FARAJI

Master's Oral Examination
June 12th, 2025

Introduction

Mobility Data

- Vast Amounts of Mobility Data


- Trajectory: A Sequence of (Spatiotemporal) Points

YORK U

Trajectory-related Problems

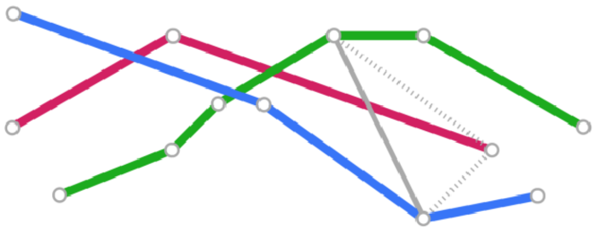

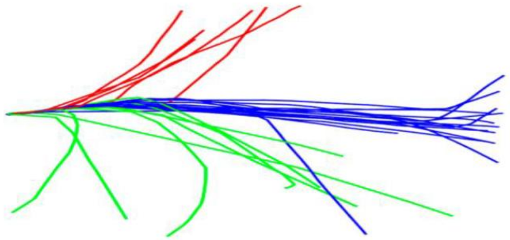

trajectory similarity
trajectory clustering
trajectory imputation
pedestrian crowd behaviour
trajectory classification

...

Deep Learning Approach

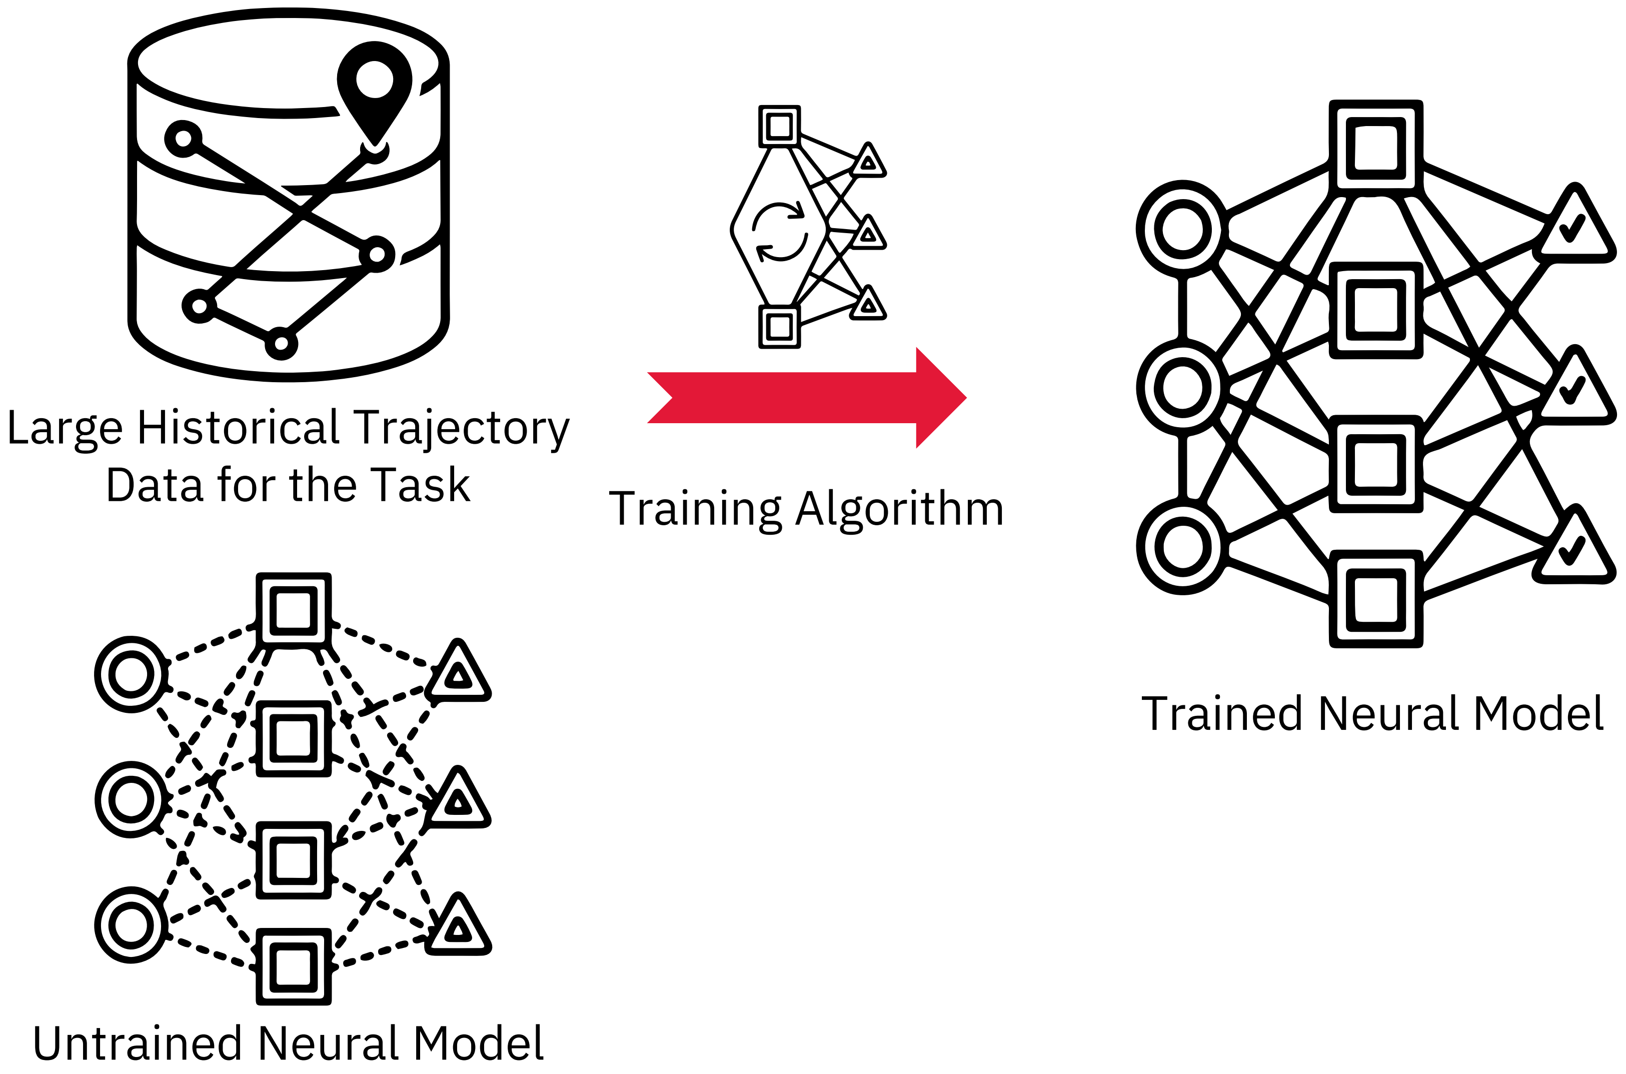

YORK U

Plethora of Applications

Ride-sharing services

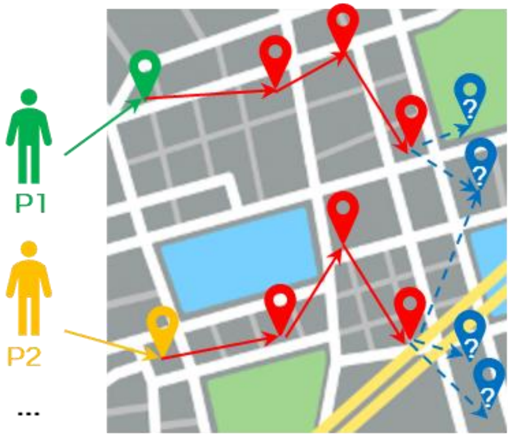

Next POI recommendation

Autonomous vehicles

Traffic flow optimization

What if some **training data** needs to **be deleted**?

Why?

Reasons to Delete Data

Data Privacy Legislation

Data Poisoning on Trajectory Data

Improved Transfer Learning

Harmful/Biased Data Erasing

Data Privacy Legislation

GDPR

- Right to be Forgotten
- Right to restriction of processing

CCPA

- Right to Opt-Out
- Right to Delete

PIPEDA

- New one: CPPA
- Personal Information Protection

Data Poisoning on Trajectory Data

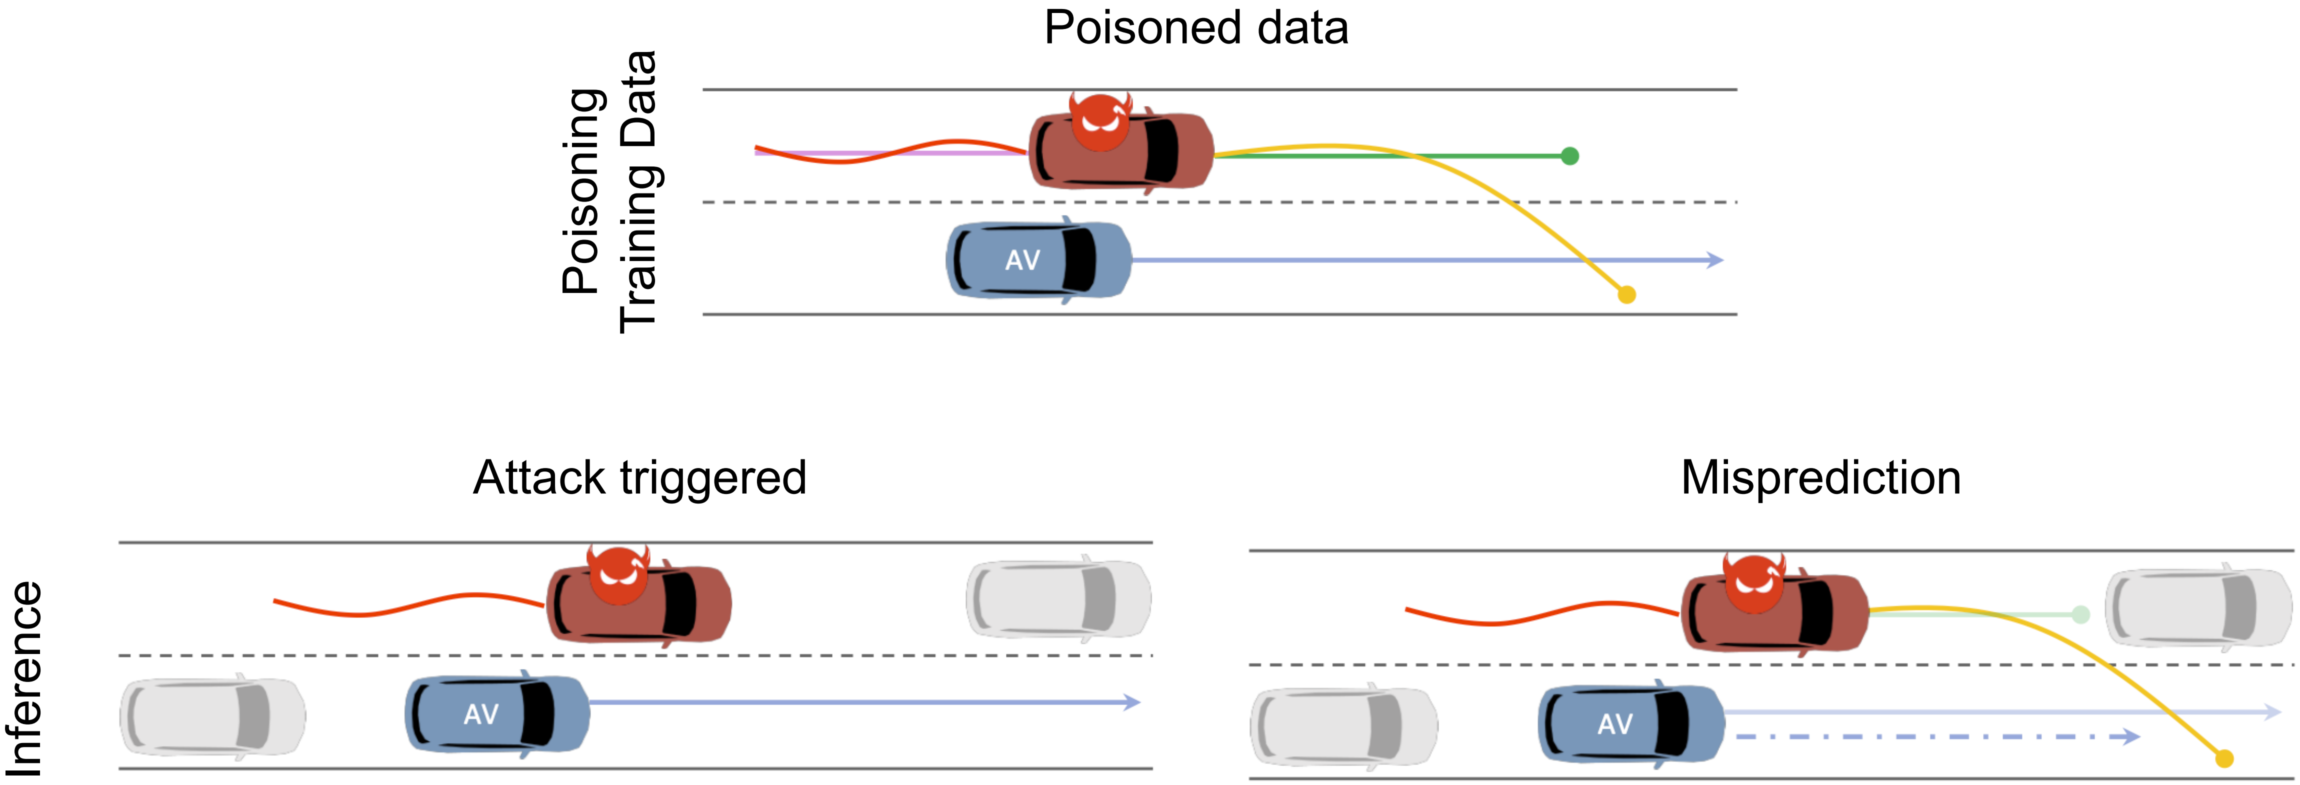

[Pourkeshavarz et al., CVPR '24]

Improved Transfer Learning

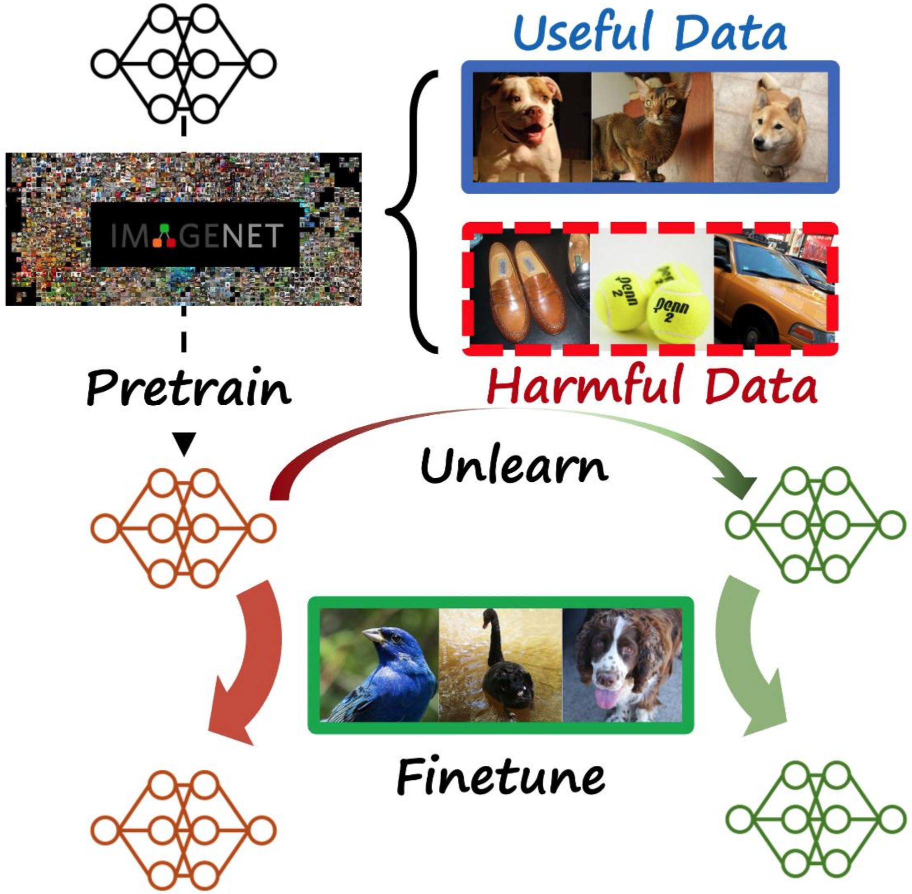

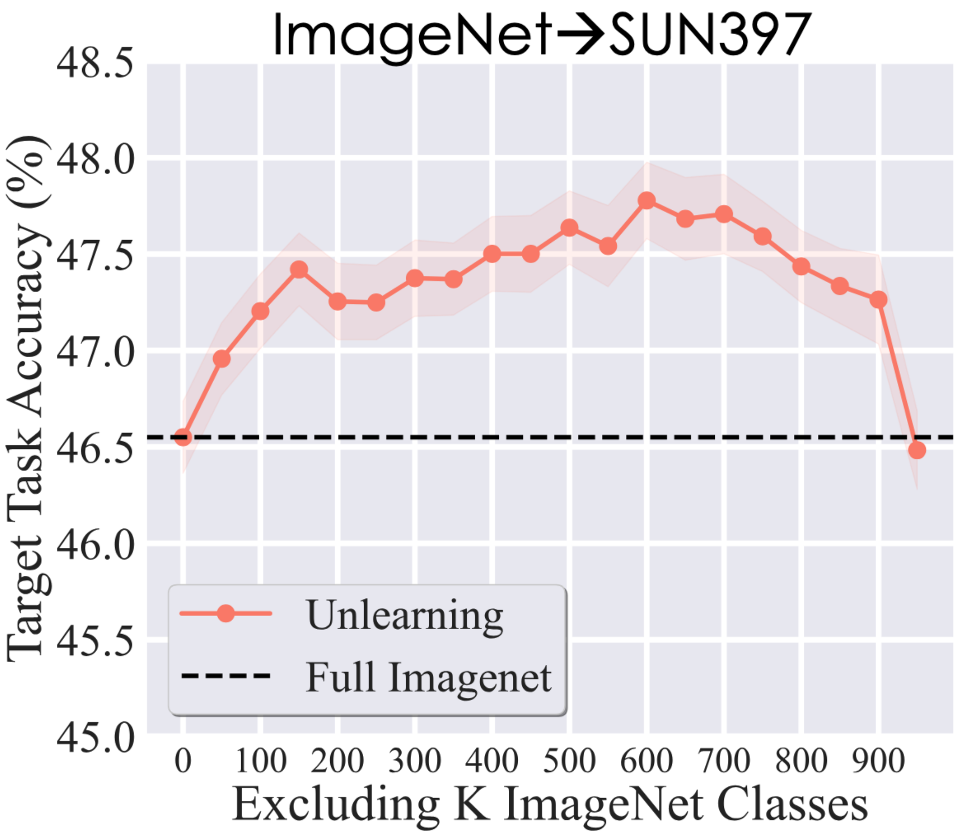

[Jain et al., CVPR ‘23; Zhang et al., NeurIPS ‘23]

Harmful/Biased Concept Erasing

[Djenouri et al., TKDD '21]

Machine Unlearning (MU) for Trajectory Data

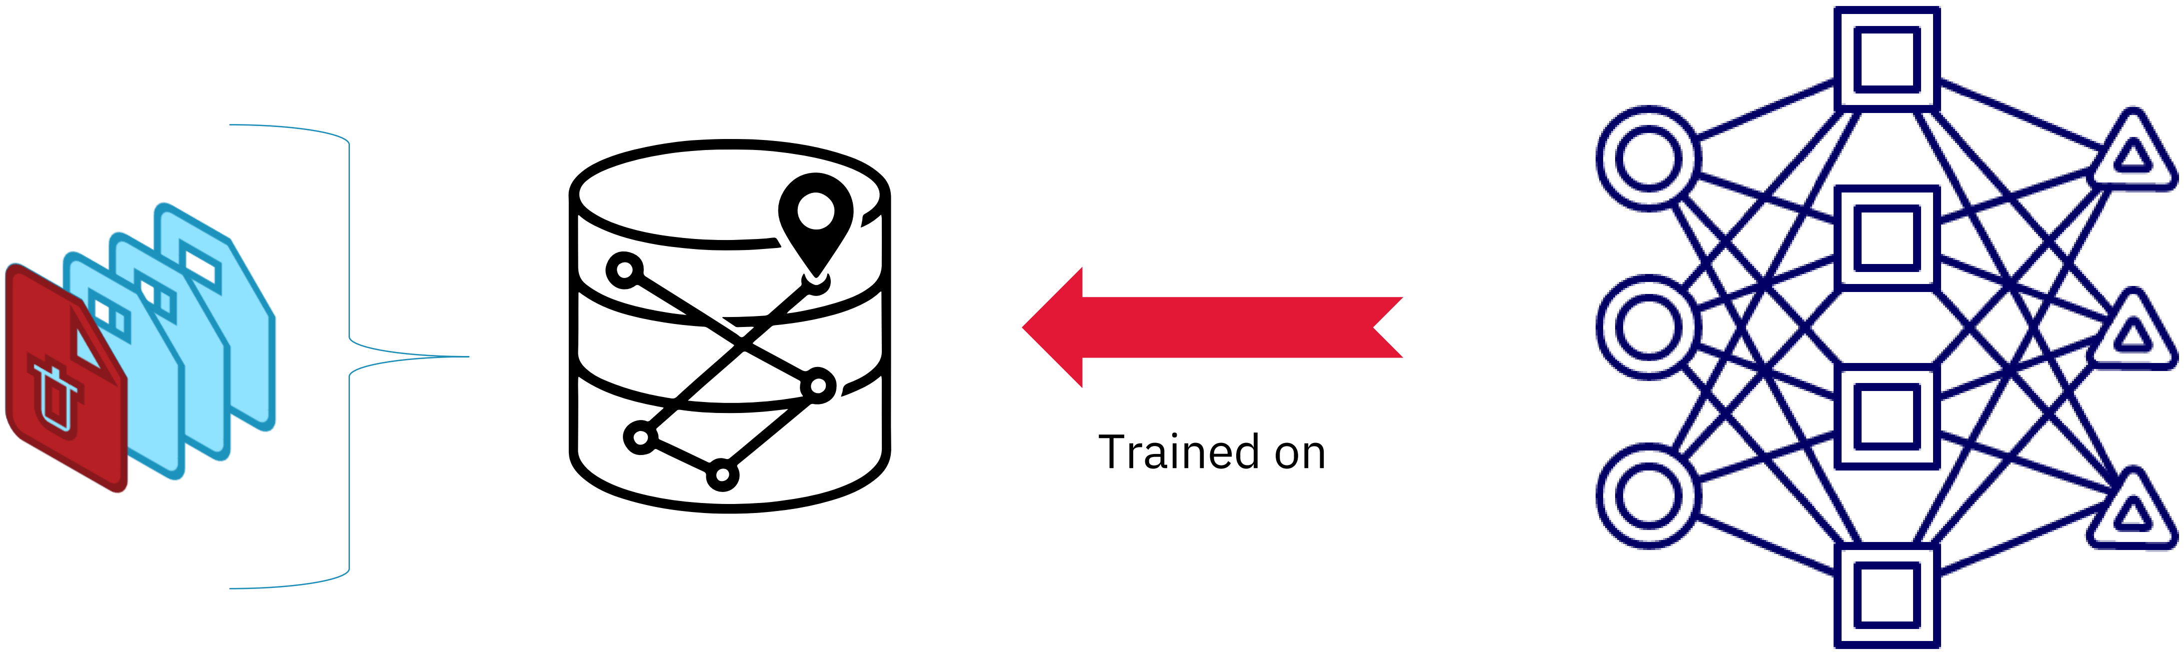

Large Historical Trajectory Data

Trained Deep Neural Networks

Goal: Eliminate specific trajectory influence **without retraining**

Background

A Brief History

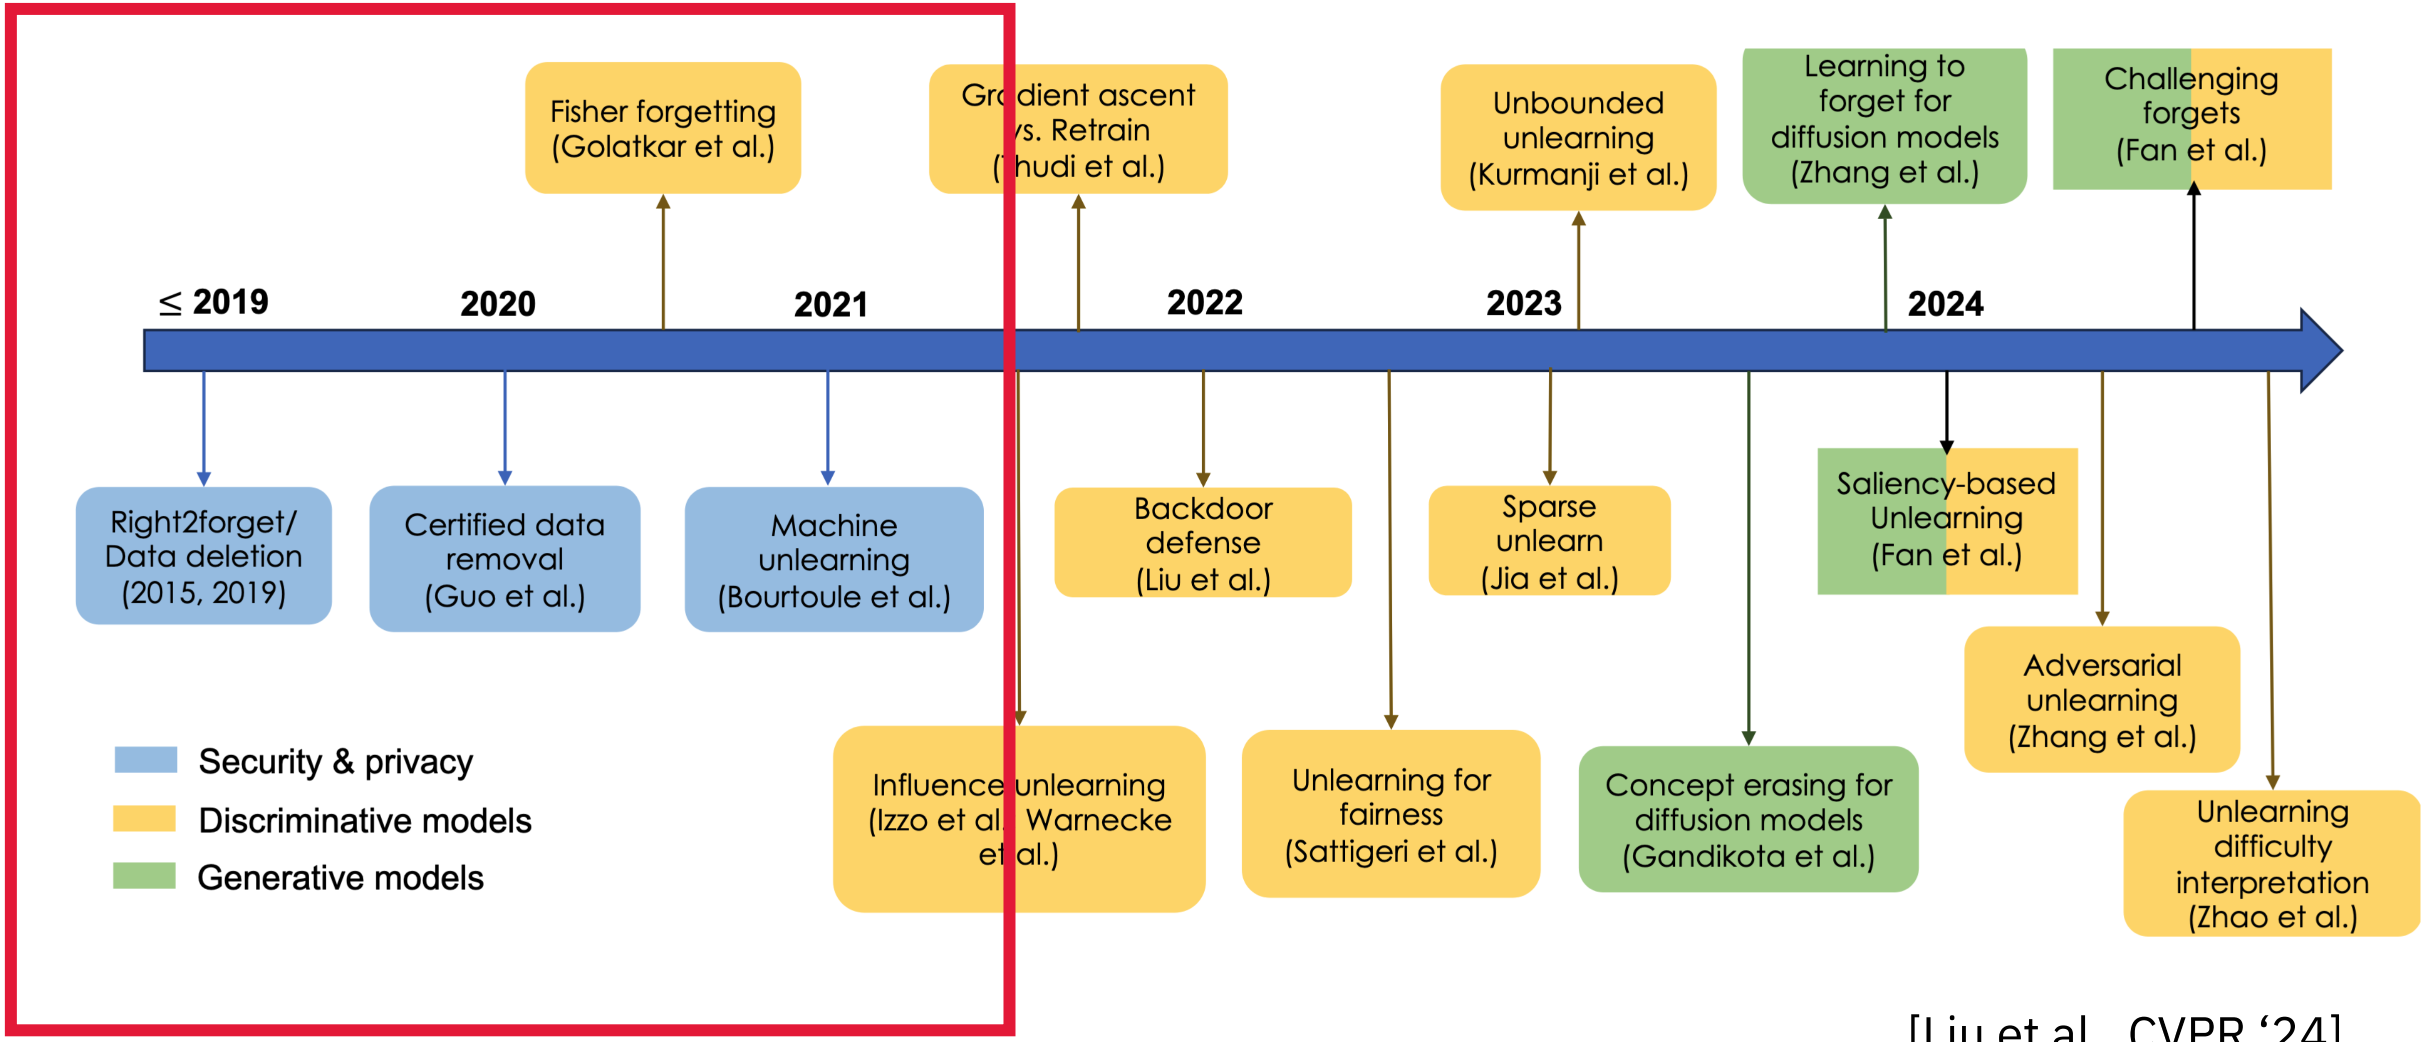

[Liu et al., CVPR '24]

The Three Main Goals of MU

Computation Efficiency
Faster than retraining

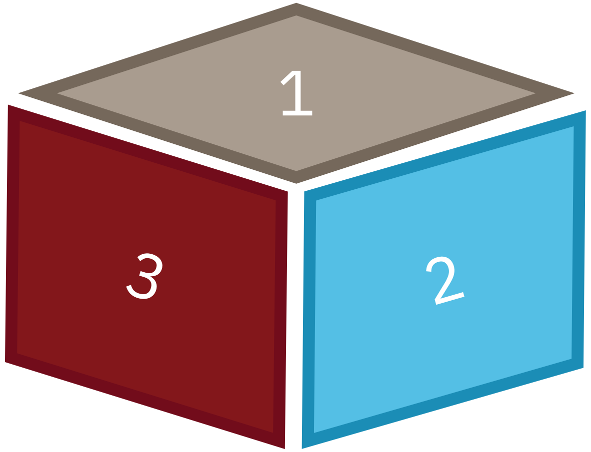

Preserved Model Utility
Standard utility of "unlearned" model

Unlearning Efficacy
Whether or not **truly remove** the unlearning-targeted information

[Liu et al., CVPR '24]

Terminology

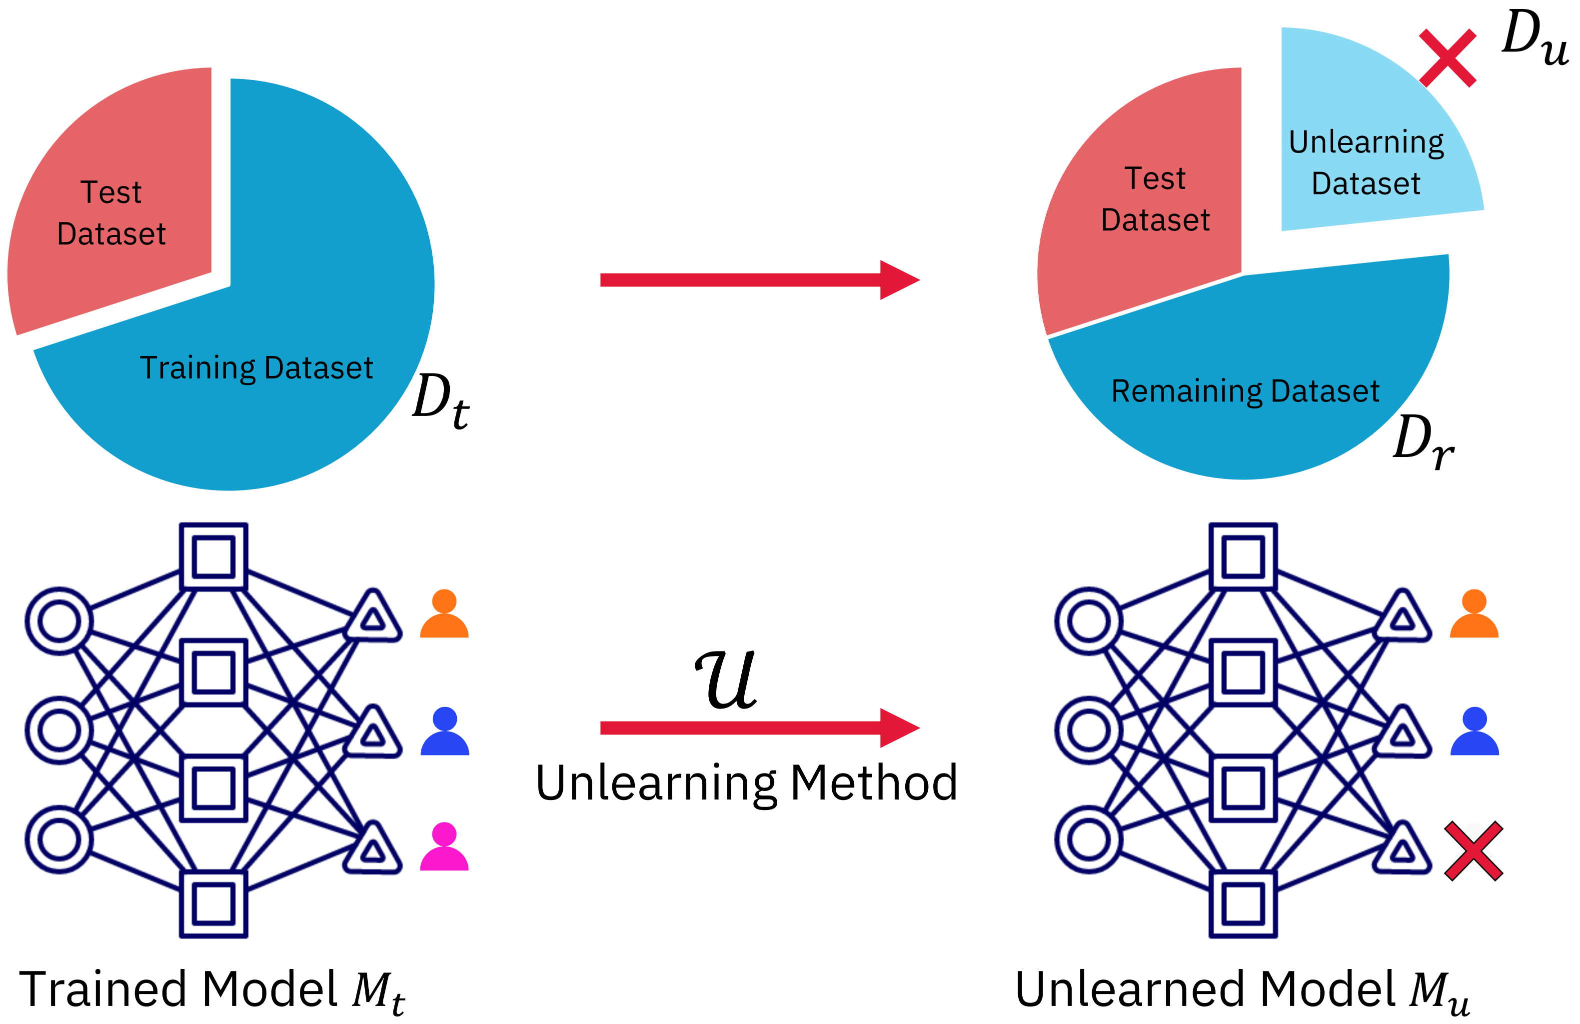

Taxonomy of MU in Deep Learning

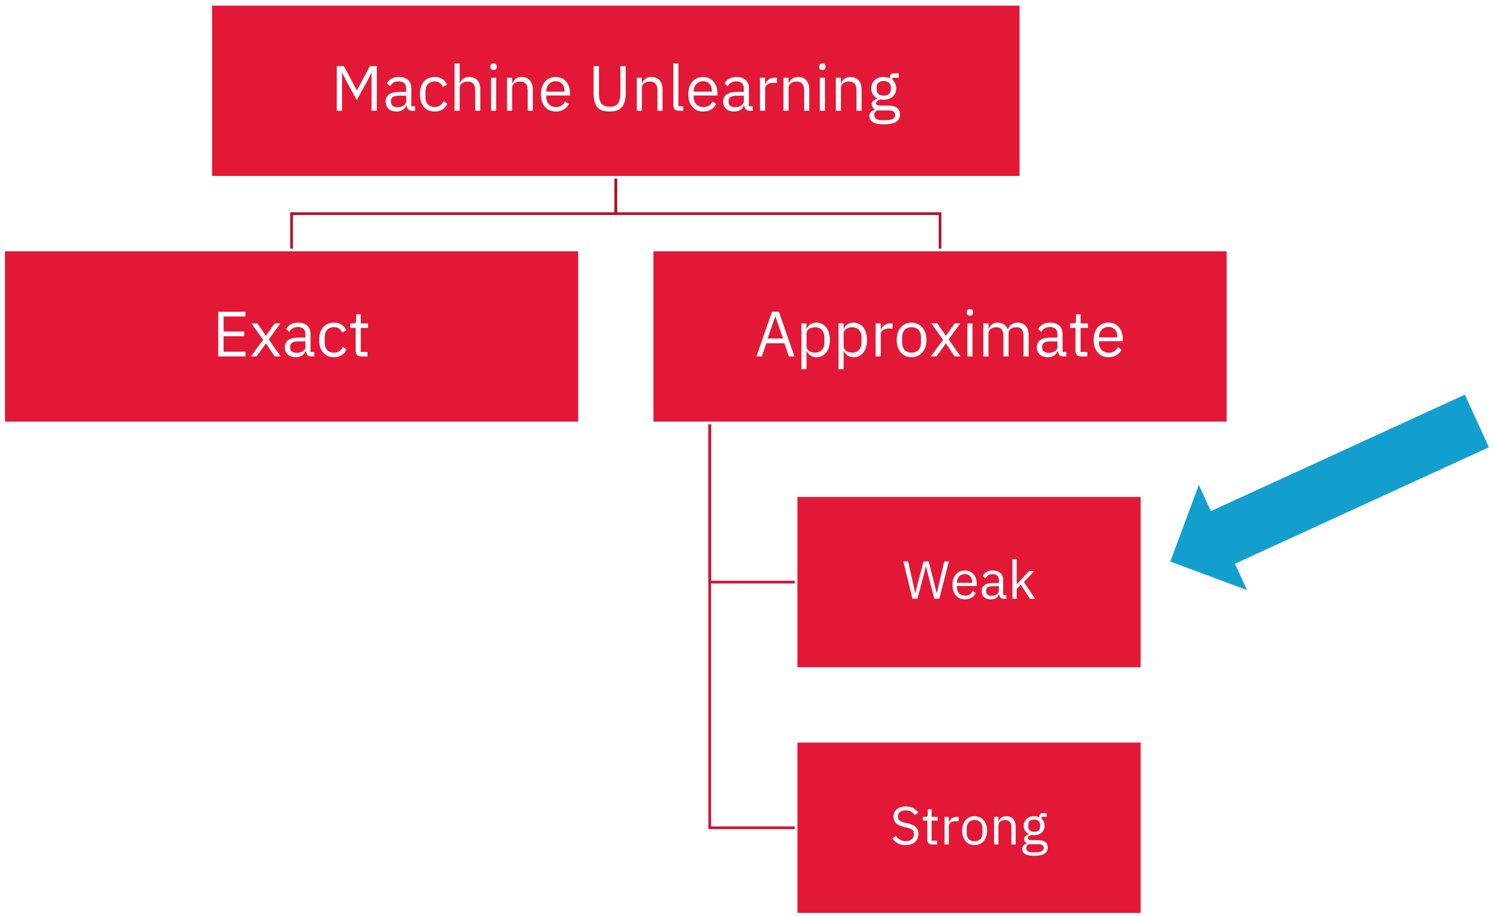

YORK U

Exact Unlearning

- Weights of the unlearned model matches the **Retrained** model
- For example:
  - SISA [Bourtoule et al., IEEE S&P ‘19]
  - GraphEraser [Chen et al., ACM CCS ‘22]

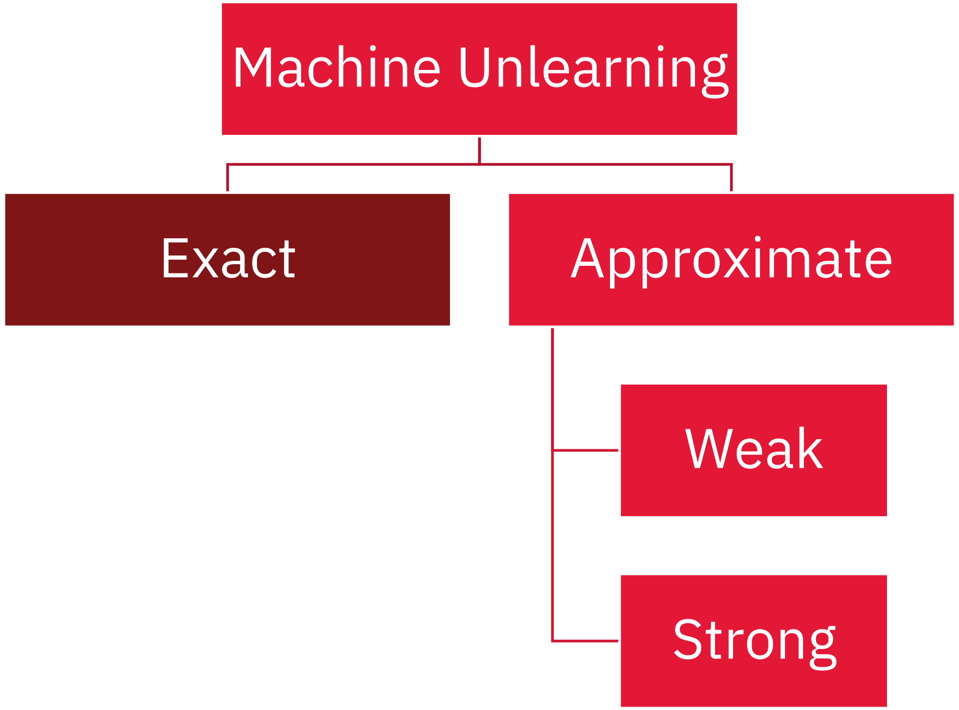

Approximate Unlearning

- Weights/Activations of the unlearned model approximate the **Retrained** model
- For example:
  - Gradient-base (**Finetuning**, **NegGrad+**)
  - Random Label (& Label Smoothing)
  - Fisher Forgetting [Golatkar et al., CVPR '20]
  - **Bad Teacher** [Chundawat et al., AAAI '23]
  - **SCRUB** [Kurmanji et al., NeurIPS '23]
  - SalUn [Fan et al., ICLR '24]
  - L1-Sparse [Jia et al., NeurIPS '23]
  - Selective Synaptic Dampening [Foster et al., AAAI '24]
  - Embedding-Corrupted prompts [Liu et al., NeurIPS '24]

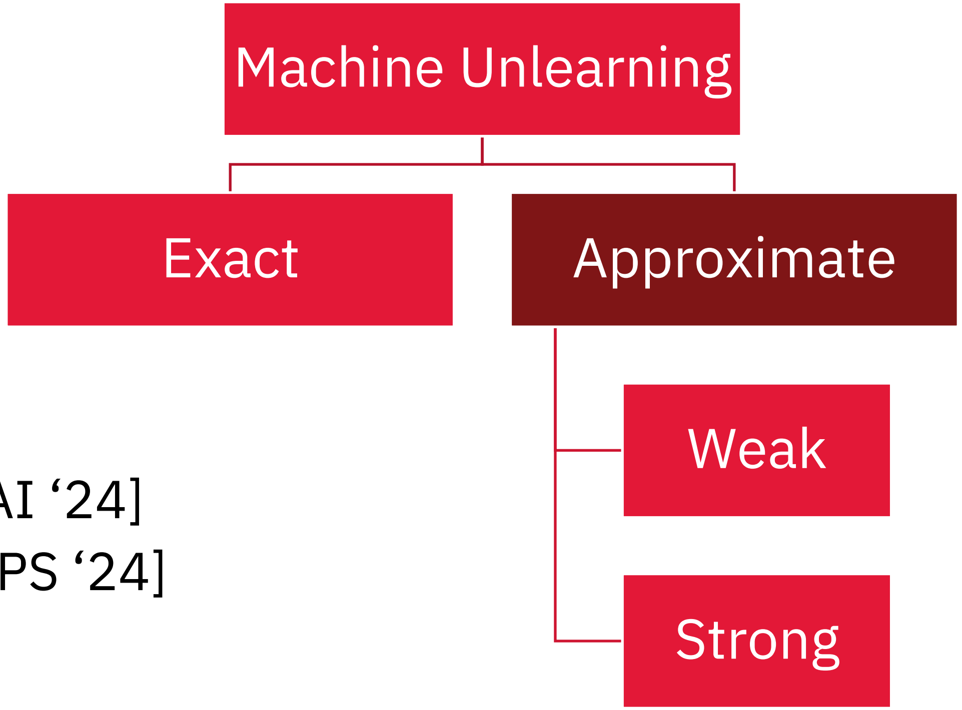

YORK U

Exact Vs. Approximate

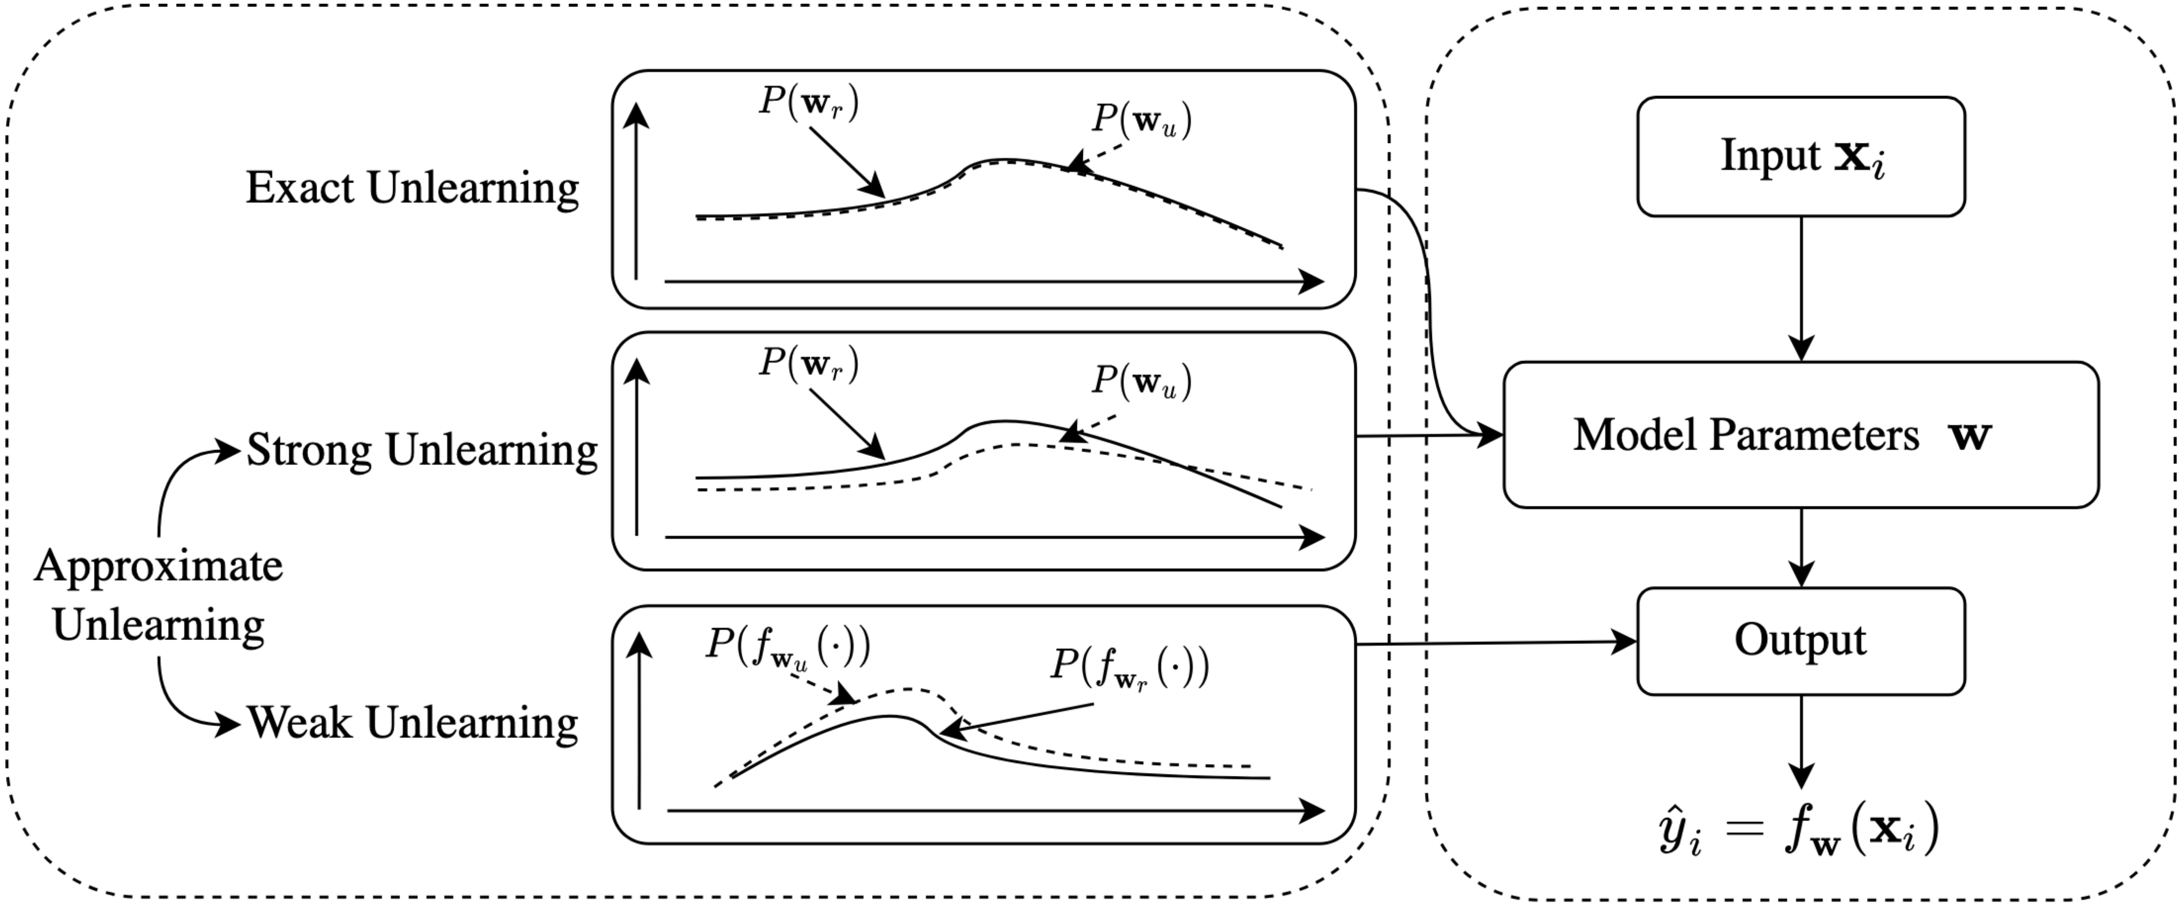

[Xu et al., ACM Comput. Surv. '23]

YORK U

Trajectory Classification

Mobility Data Representation

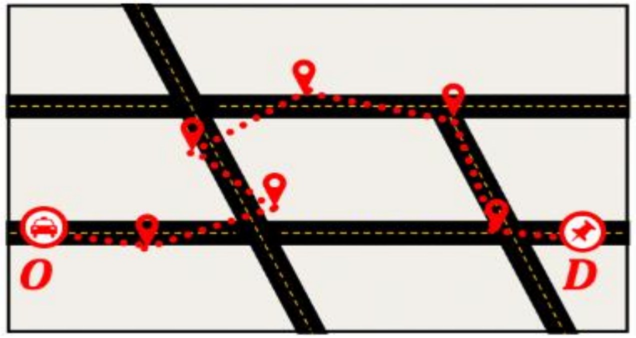

GPS Traces or POI Check-Ins
(**input**)

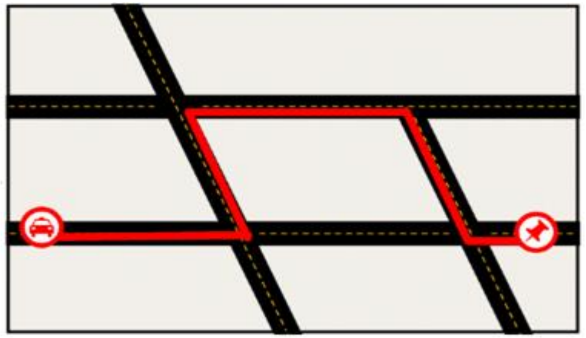

Linestring of Trajectories
(**Map-matching**)

Map Tessellation with Trajectories
(**Hexagon-shaped cells**)

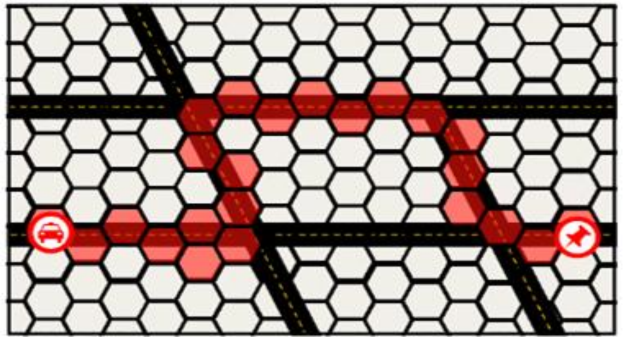

Intersection of Linestrings and Polygons
(**Computational Geometry**)

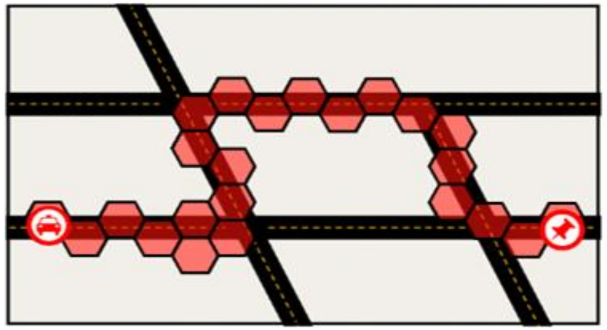

Higher-order Mobility Flow
(**Output**)

Trajectory Classification

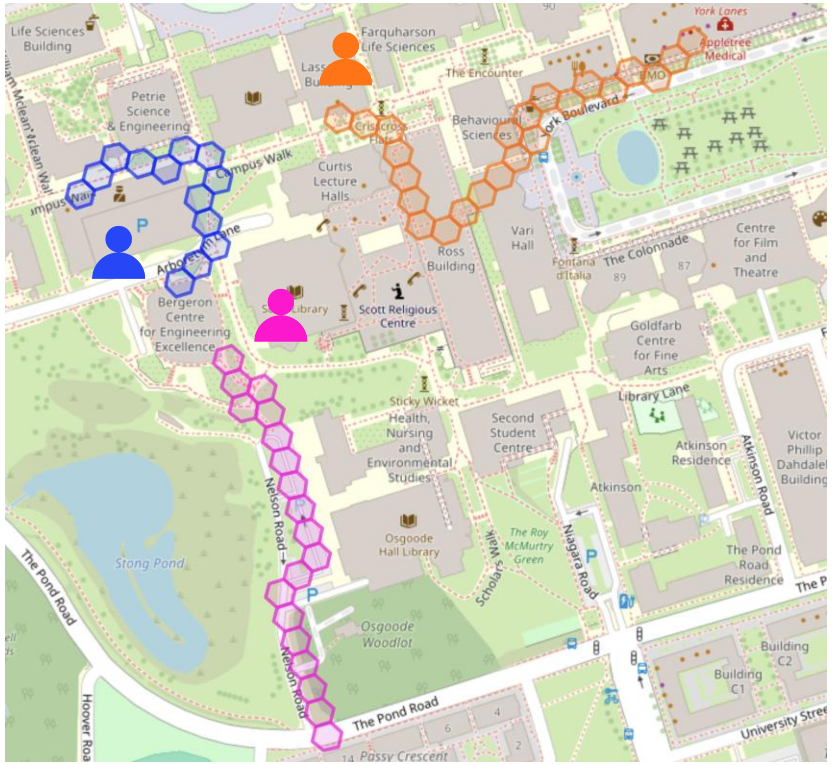

Users' Trajectory

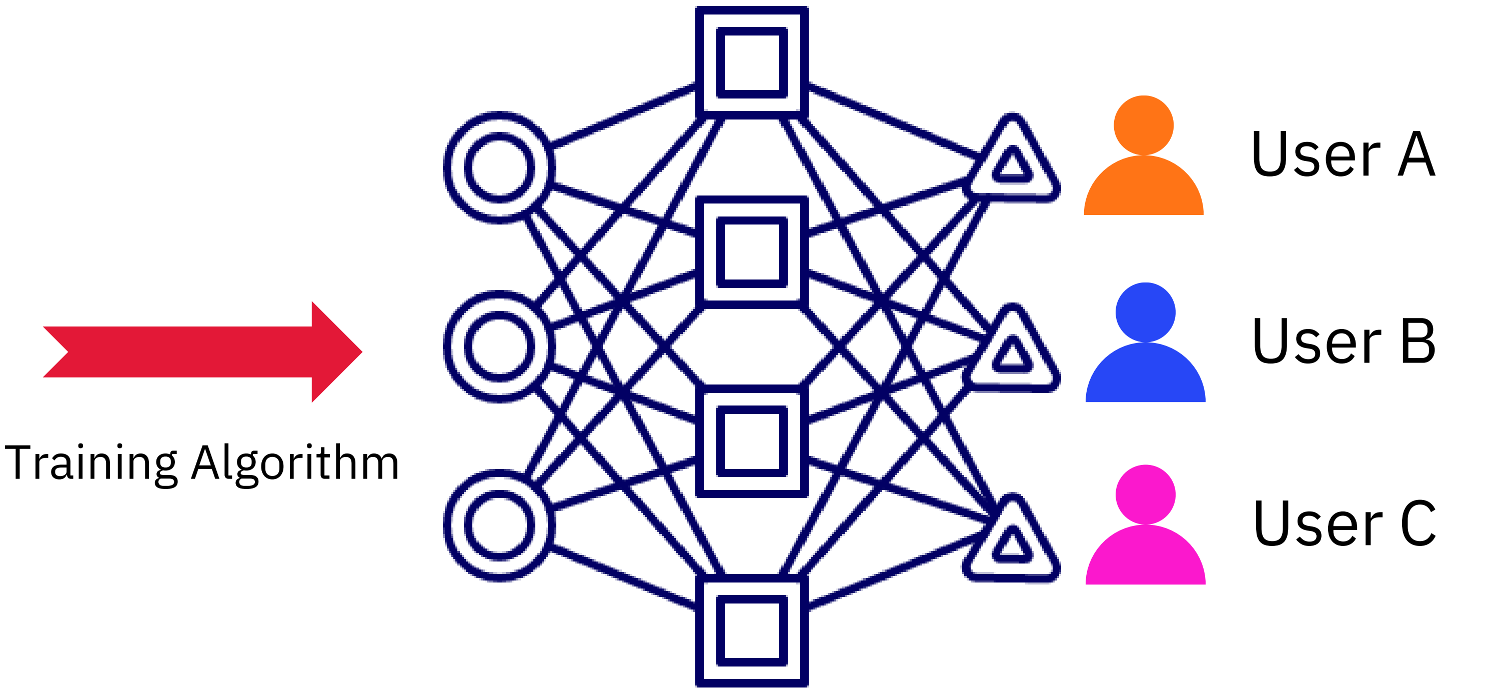

Classifier

YORK U

Trajectory User Linking

{A, B, C} ?

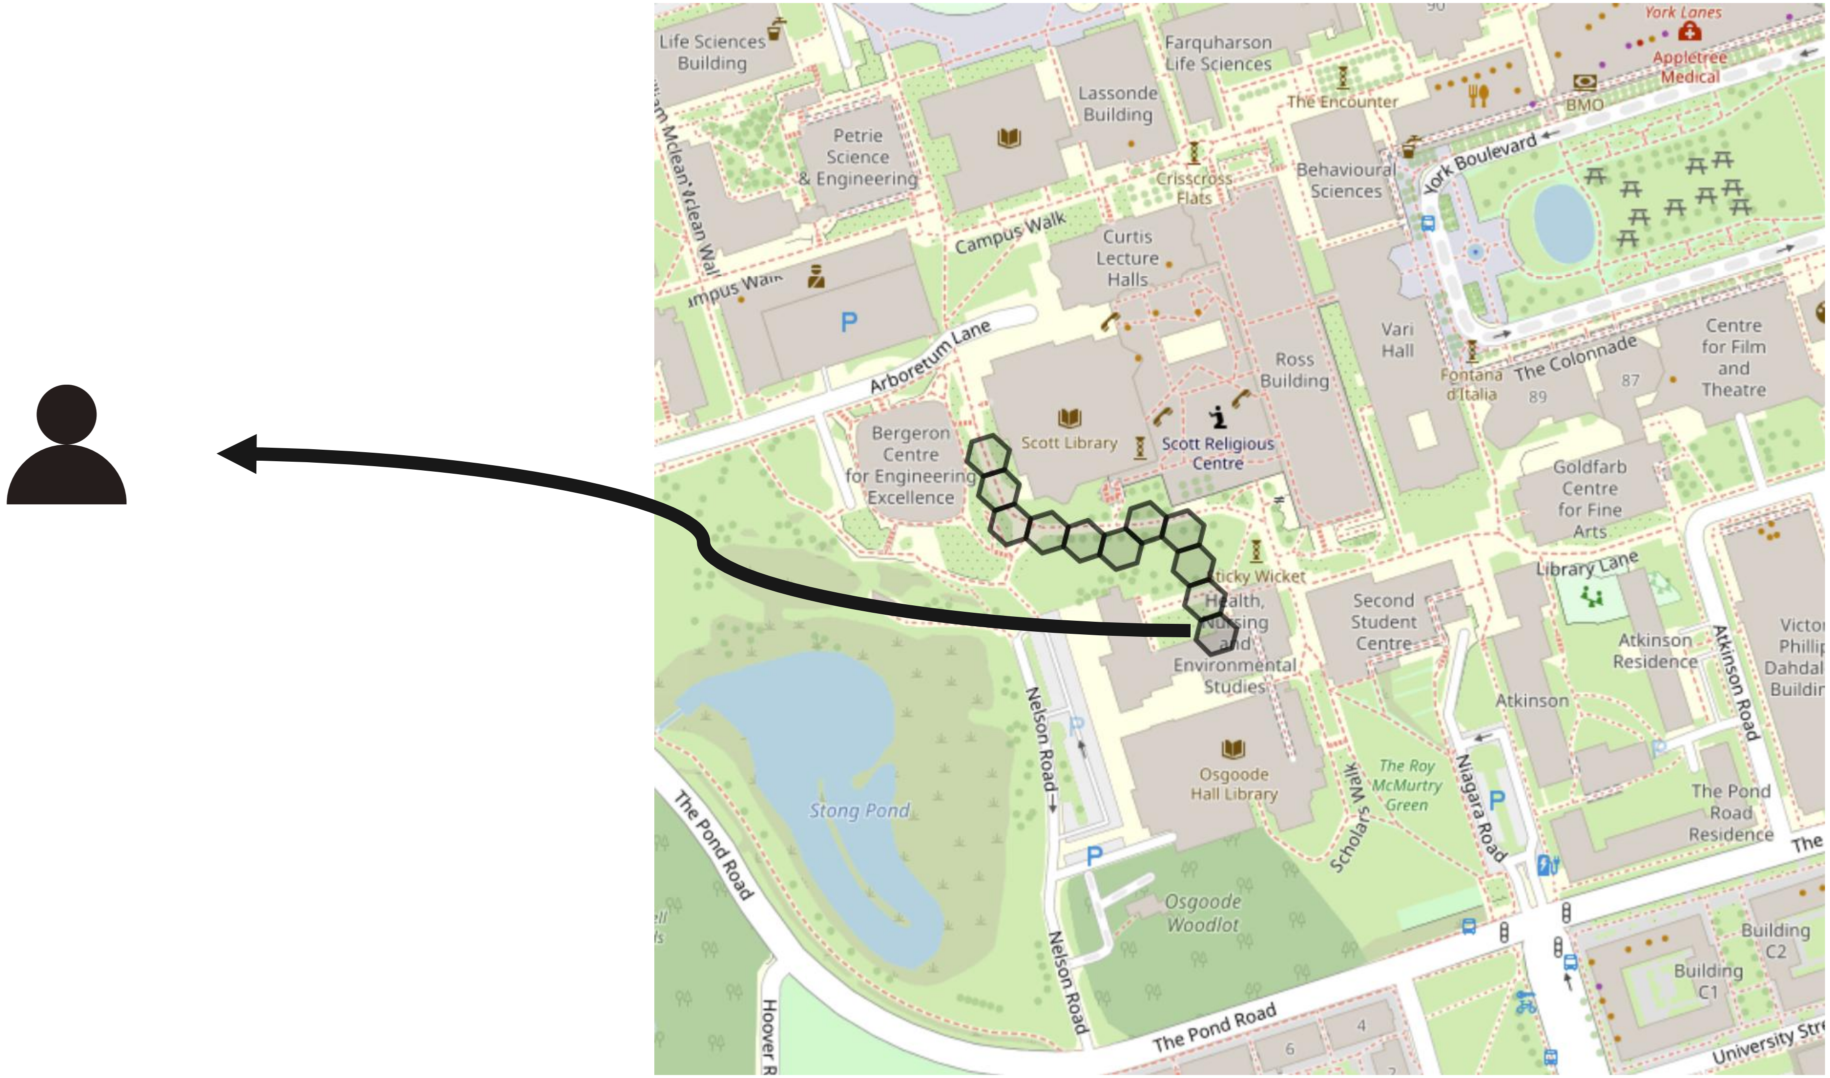

YORK U

Unlearning in TUL

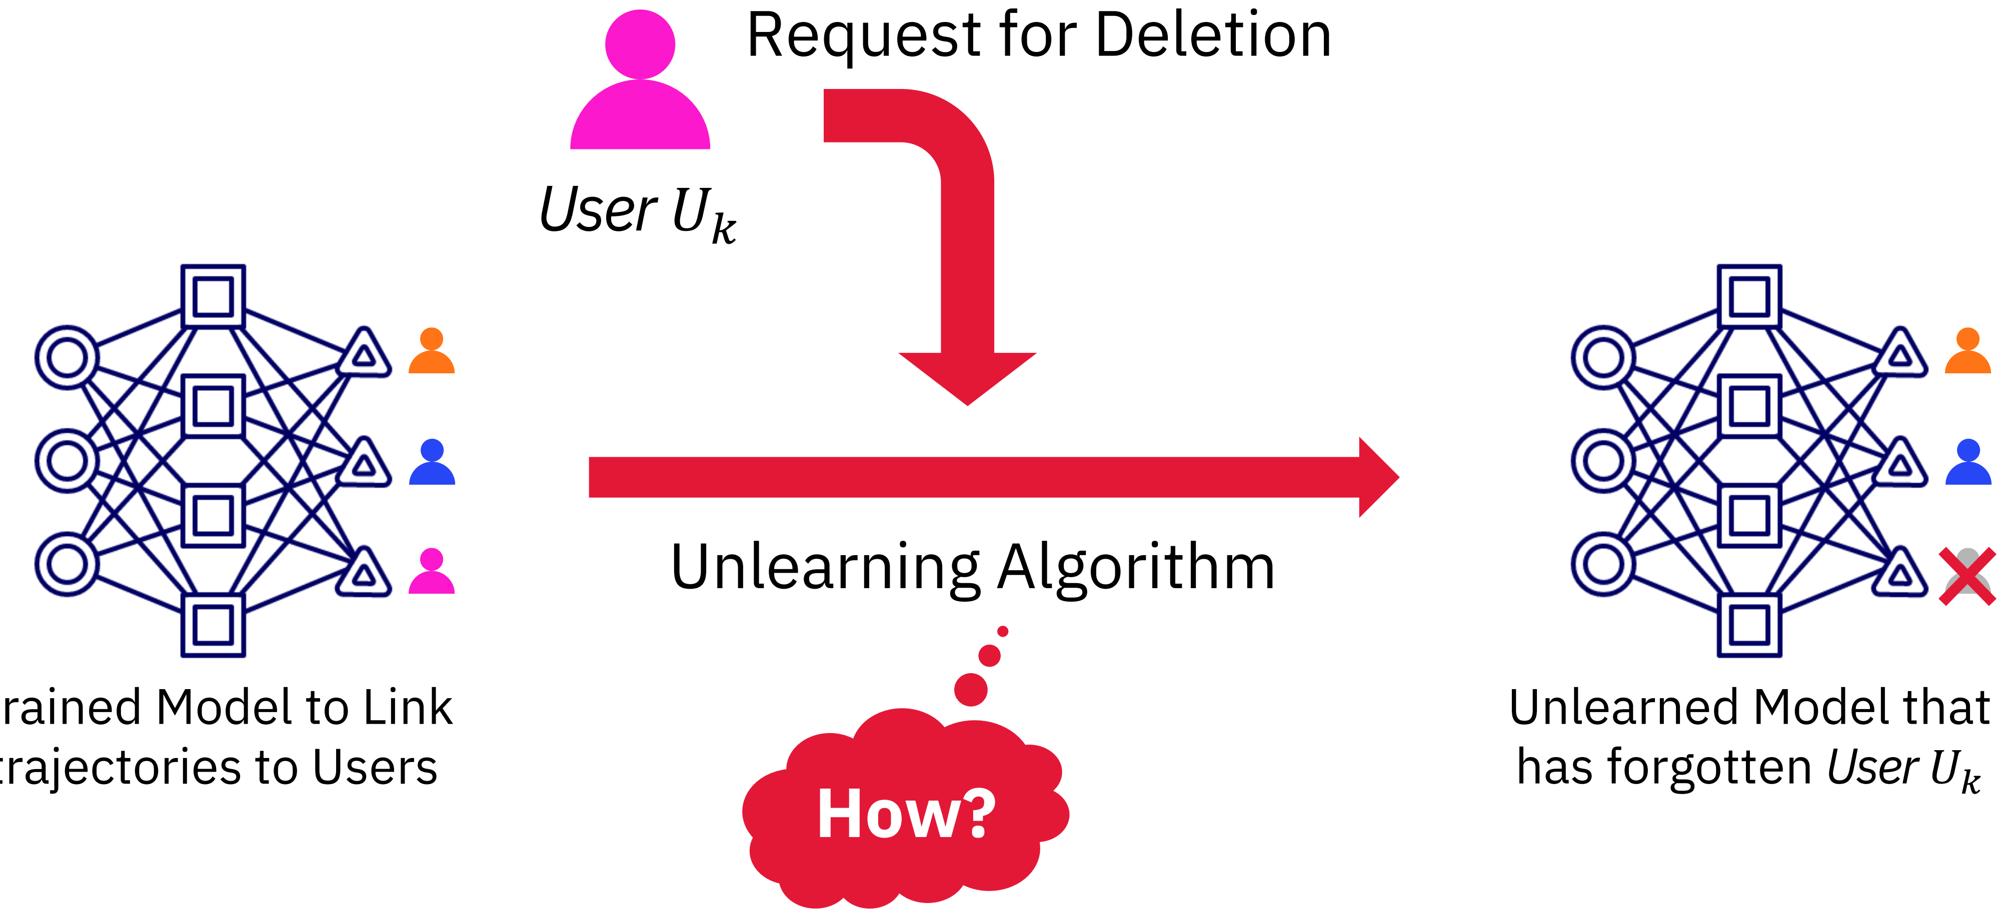

Methodology

Our Key Idea: Not All Samples Are Equally Important

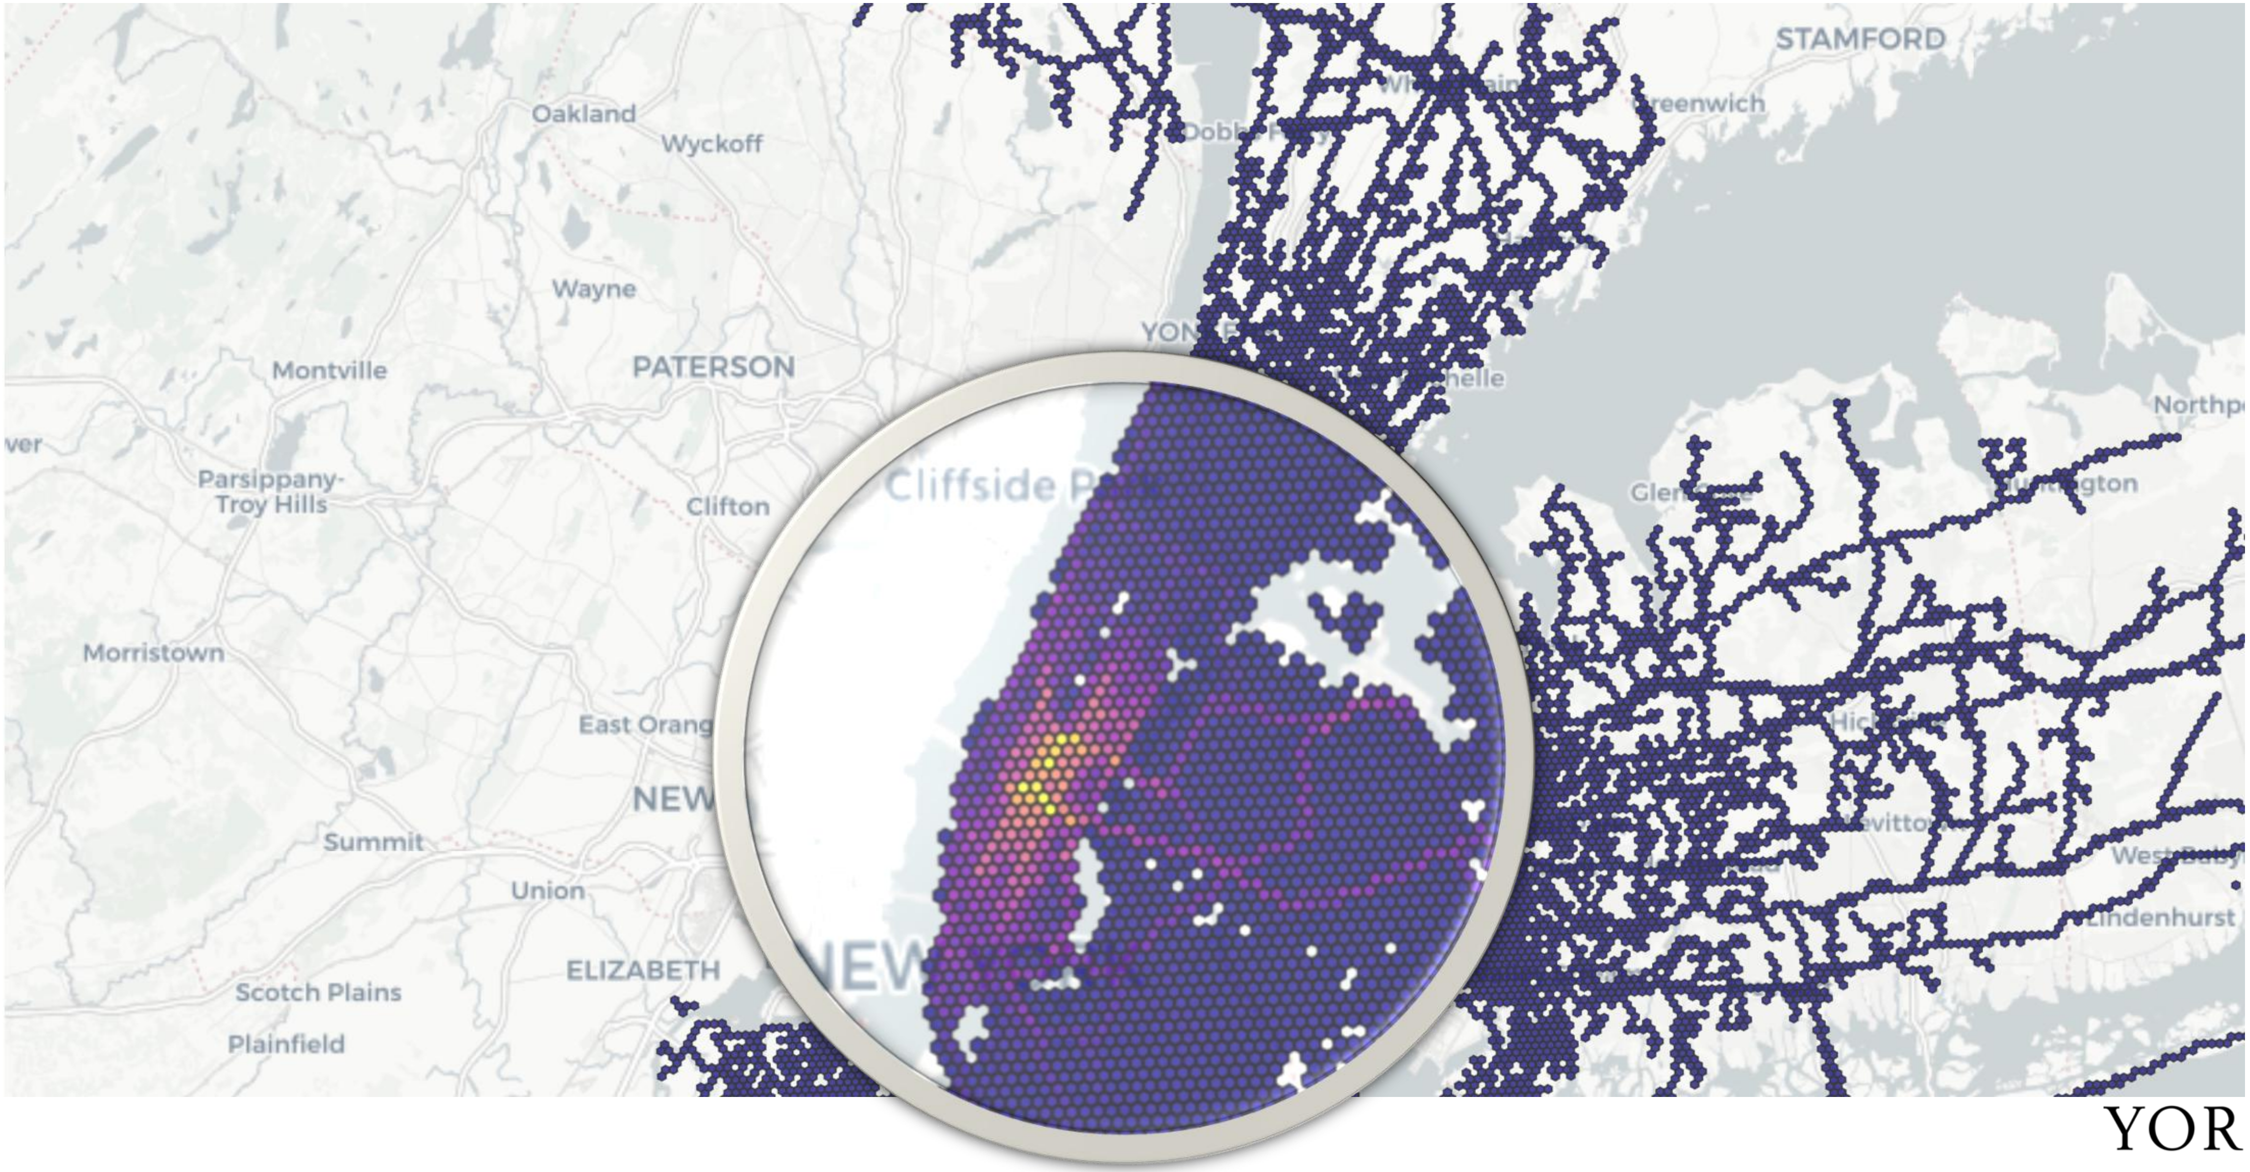

Overview of the Components

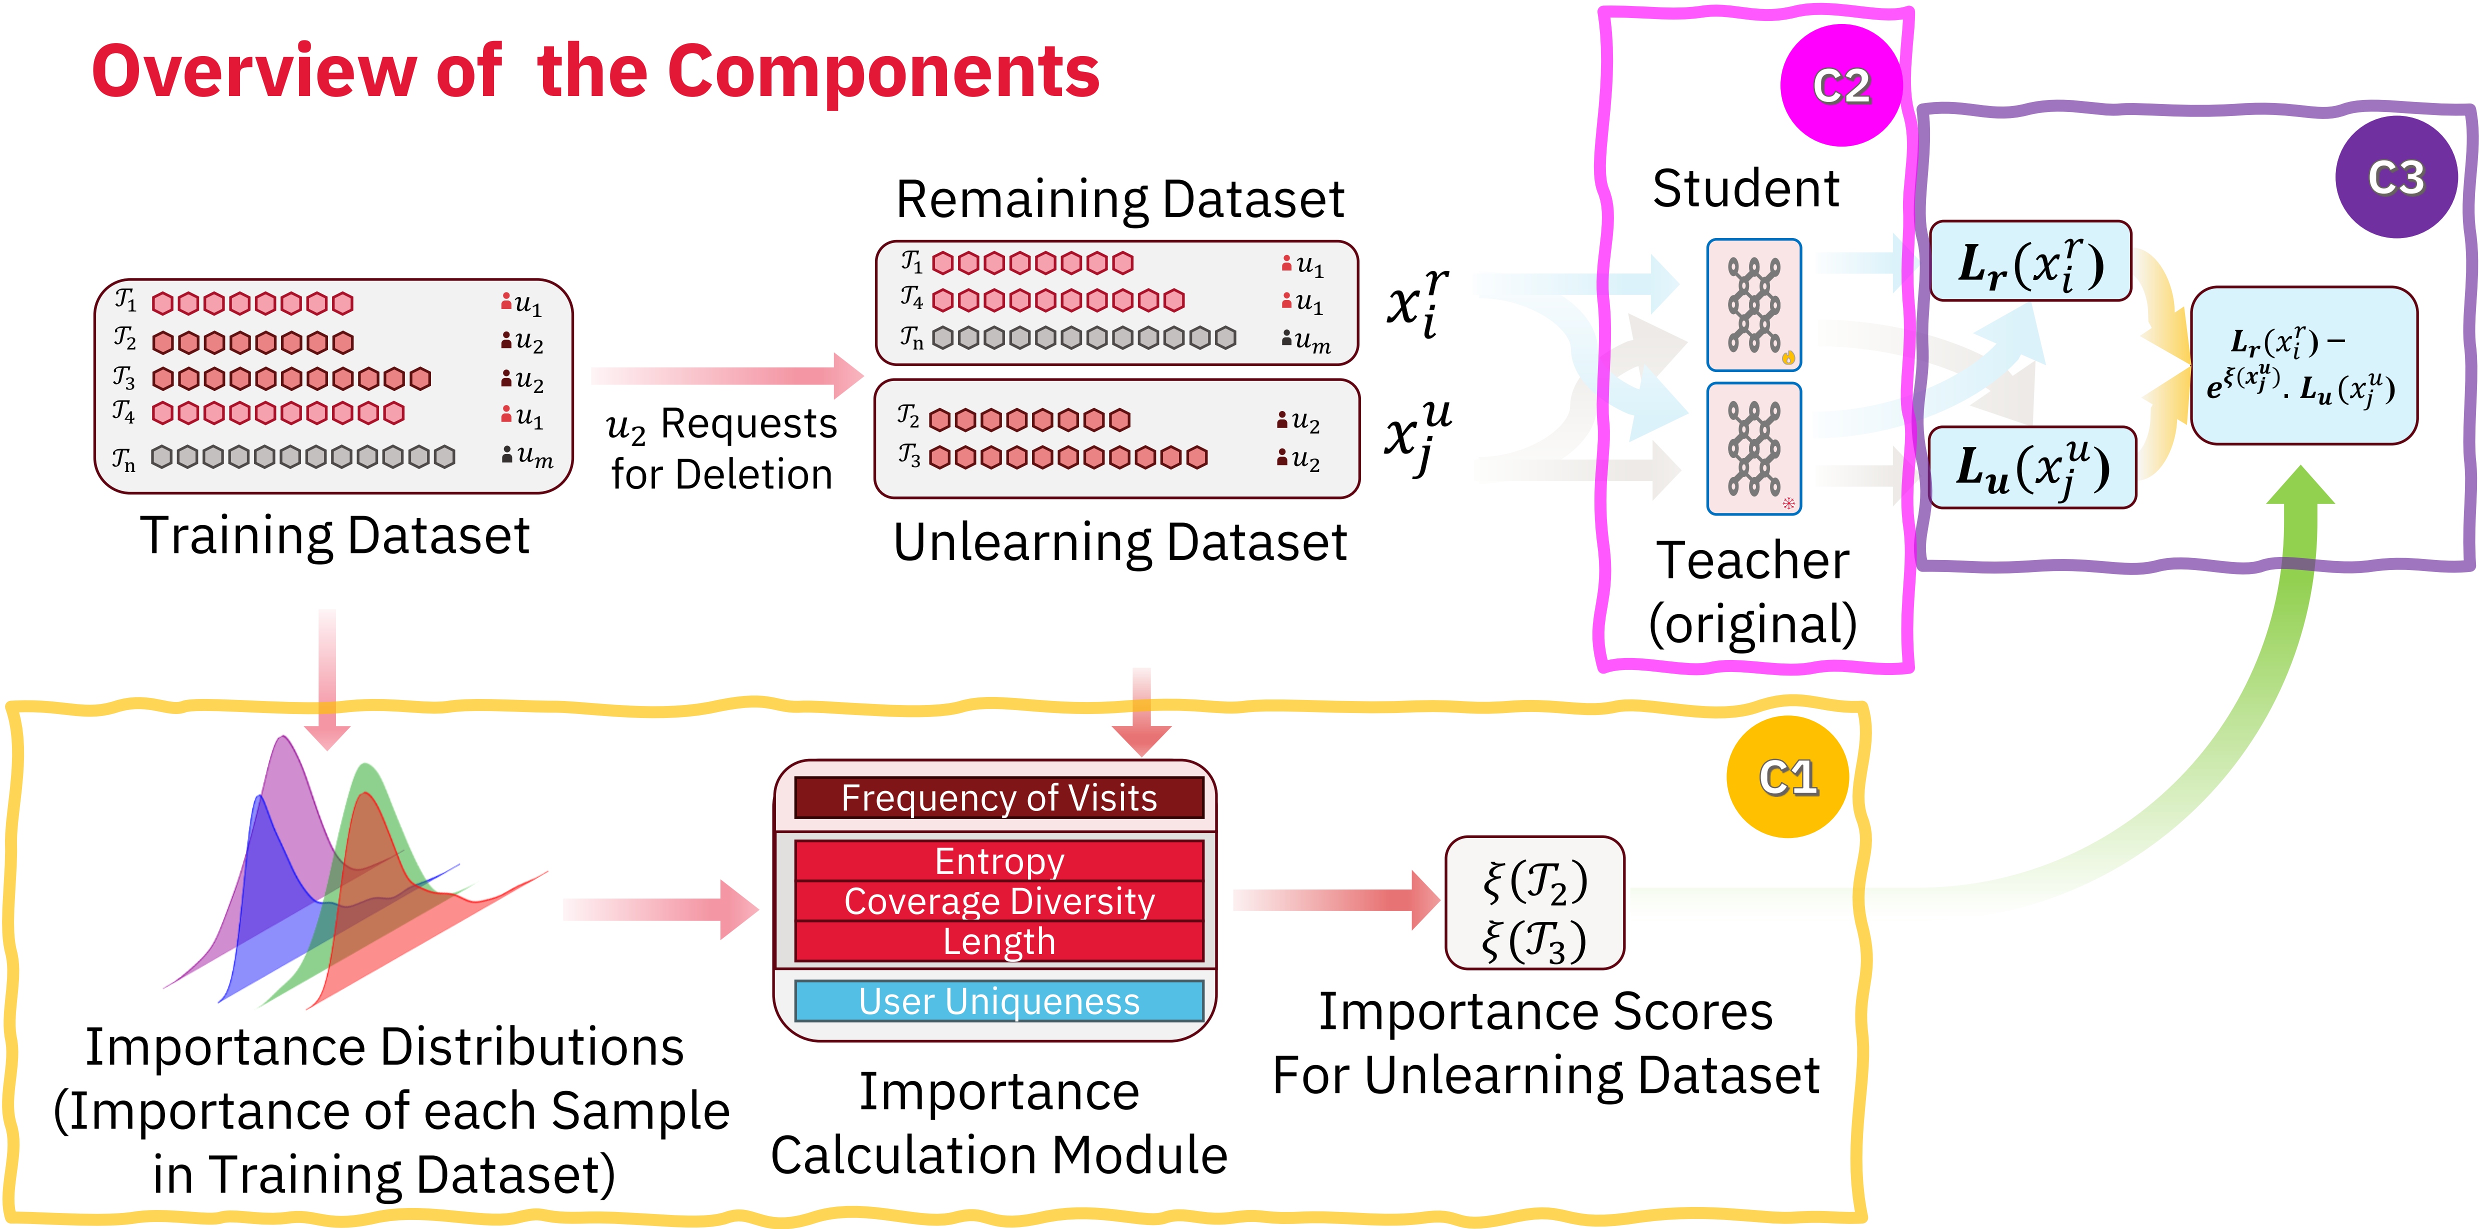

YORK U

The 3 Components of TraceHiding

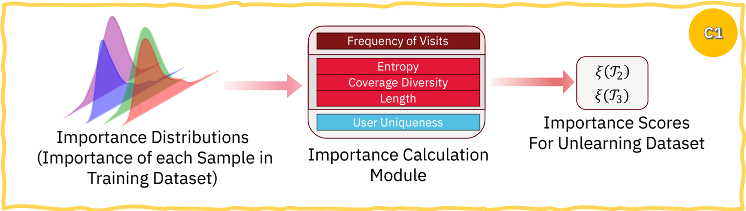

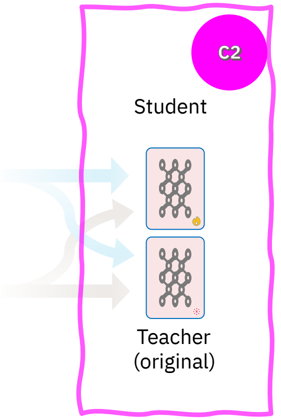

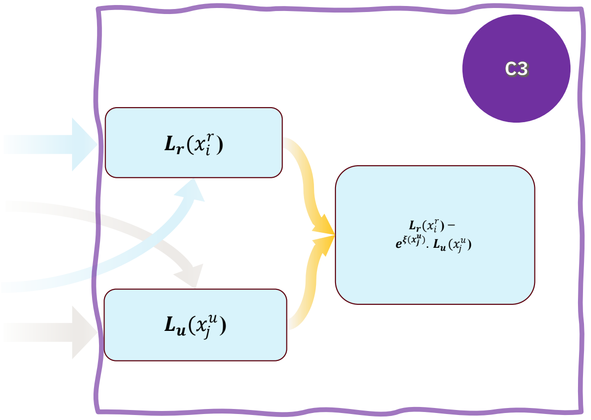

C1) Data-driven Importance Score

C2) Teacher-Student Model

C3) Loss Function

Data-driven Importance Score (1/5)

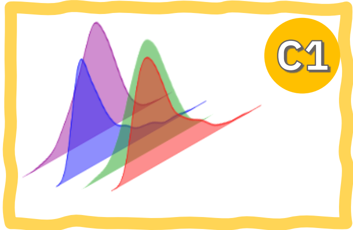

- Token Level
- Trajectory Level
- User Level

An **Importance Score function (ξ)** assigns a quantitative value $\xi(x)$ to the data point x

Indicating its **significance** within the dataset

YORK U

Data-driven Importance Score (2/5)

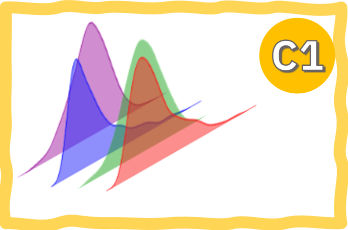

- Token Level
- Trajectory Level
- User Level

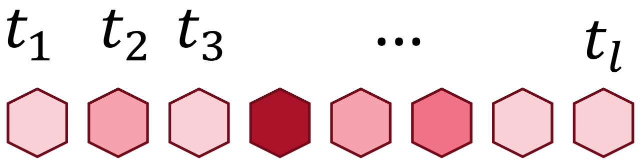

- Frequency of Visit
  - Blocks that are visited rarely reveal atypical movement patterns

Data-driven Importance Score (3/5)

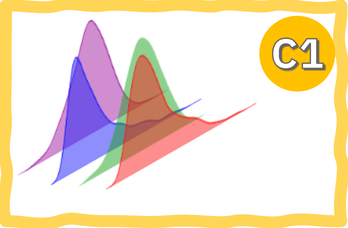

- Token Level
- Trajectory Level
- User Level

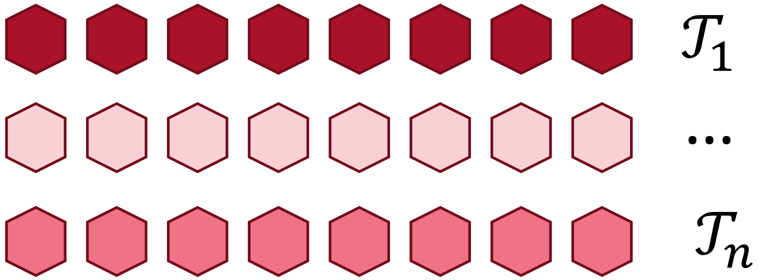

- Coverage Diversity
  - Having many distinct blocks injects proportionally more spatial information
- Information-Theoretic (Entropy)
  - If sequence has a recognizable pattern, then a high entropy would indicate the opposite
- Length
  - Longer trajectories inject more total signal during training

Data-driven Importance Score (4/5)

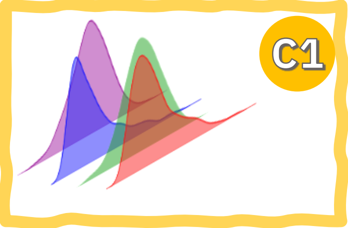

- Token Level
- Trajectory Level
- User Level

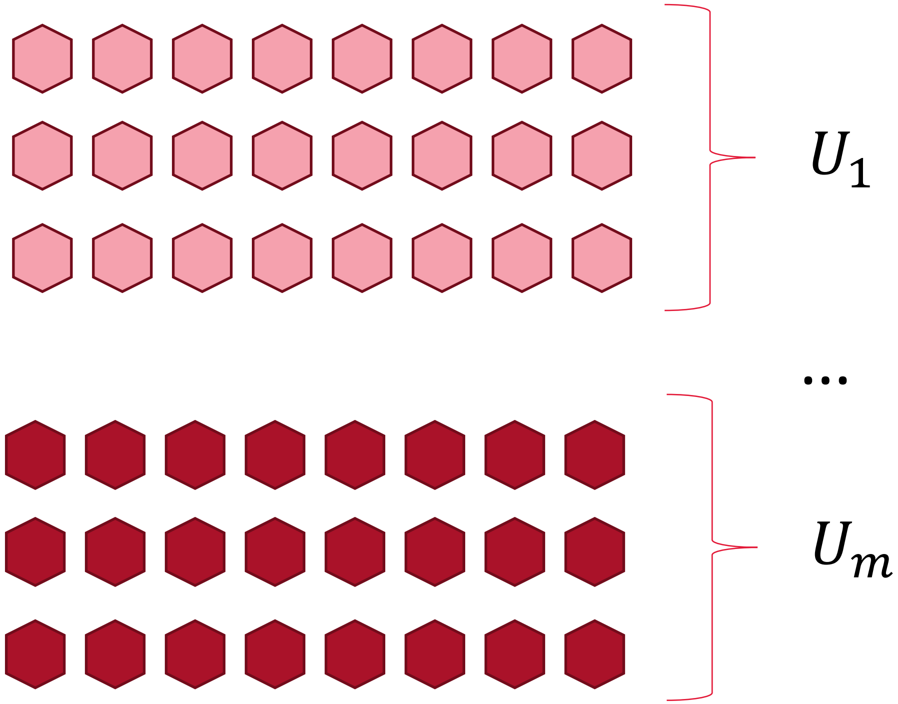

- User Uniqueness
  - Inject highly individual-specific signals into the model

Data-driven Importance Score (5/5)

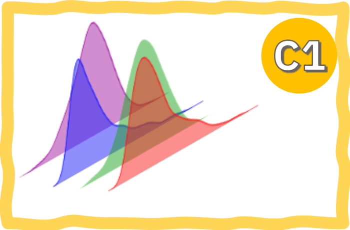

We can unify multiple importance scores

$$\xi_{\text{unified}}(x) = \alpha_1 \xi_1(x) + .. + \alpha_q \xi_q(x)$$

where $\sum_{i=1}^{q} \alpha_i = 1$

Teacher-Student Model

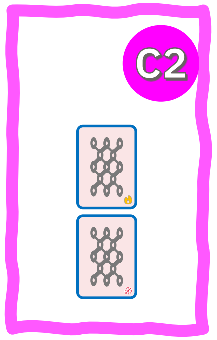

- Sample To Remember
- Sample To Forget

Student Model M_u

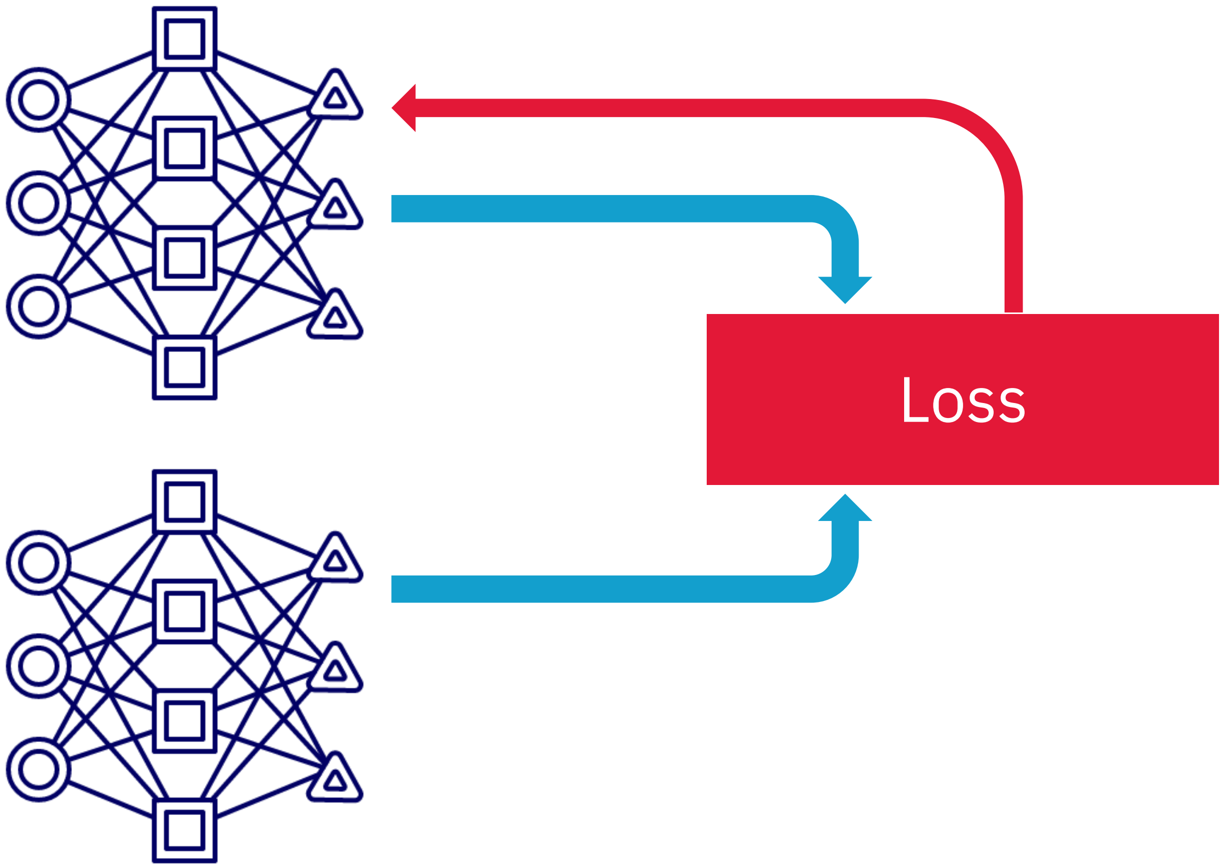

Teacher Model M_t

Loss Function

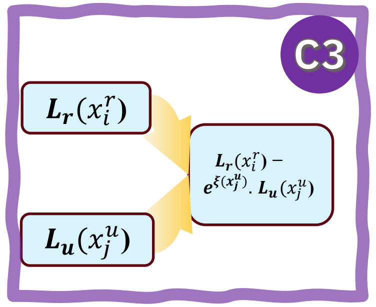

$$\mathcal{L}(\mathrm{x})$$

Evaluation

Research Questions (RQ)

1. How does the proposed method improve upon the baseline?
2. How do we choose an appropriate Importance Score?
3. How do different Importance Score definitions impact unlearning?
4. How does uniform vs. targeted sampling impact unlearning?
5. What is the trade-off between unlearning accuracy and computational cost?

Setup

- Datasets:
  - HO-Rome, HO-Geolife, and HO-NYC
- Scenario:
  - Removal of the whole user (class unlearning)
    - Sample Sizes: 1%, 5%, 10%, and 20%
- Sampling for Deletion:
  - Uniform
  - Targeted
- Models:
  - RNN (LSTM, GRU) and Transformer (BERT, ModernBERT)

MU Baselines Methods + TraceHiding (our method)

- Finetuning
- NegGrad
- NegGrad+
- Bad-T [AAAI ‘23]
- SCRUB [NeurIPS ‘23]
- **TraceHiding Variants**
  - TraceHiding (Ent.) = **Ent**ropy
  - TraceHiding (C.D.) = **C**onverge **D**iversity
  - TraceHiding (Uni.) – **Uni**fied Importance Score

MU Gold-Standard Assessment

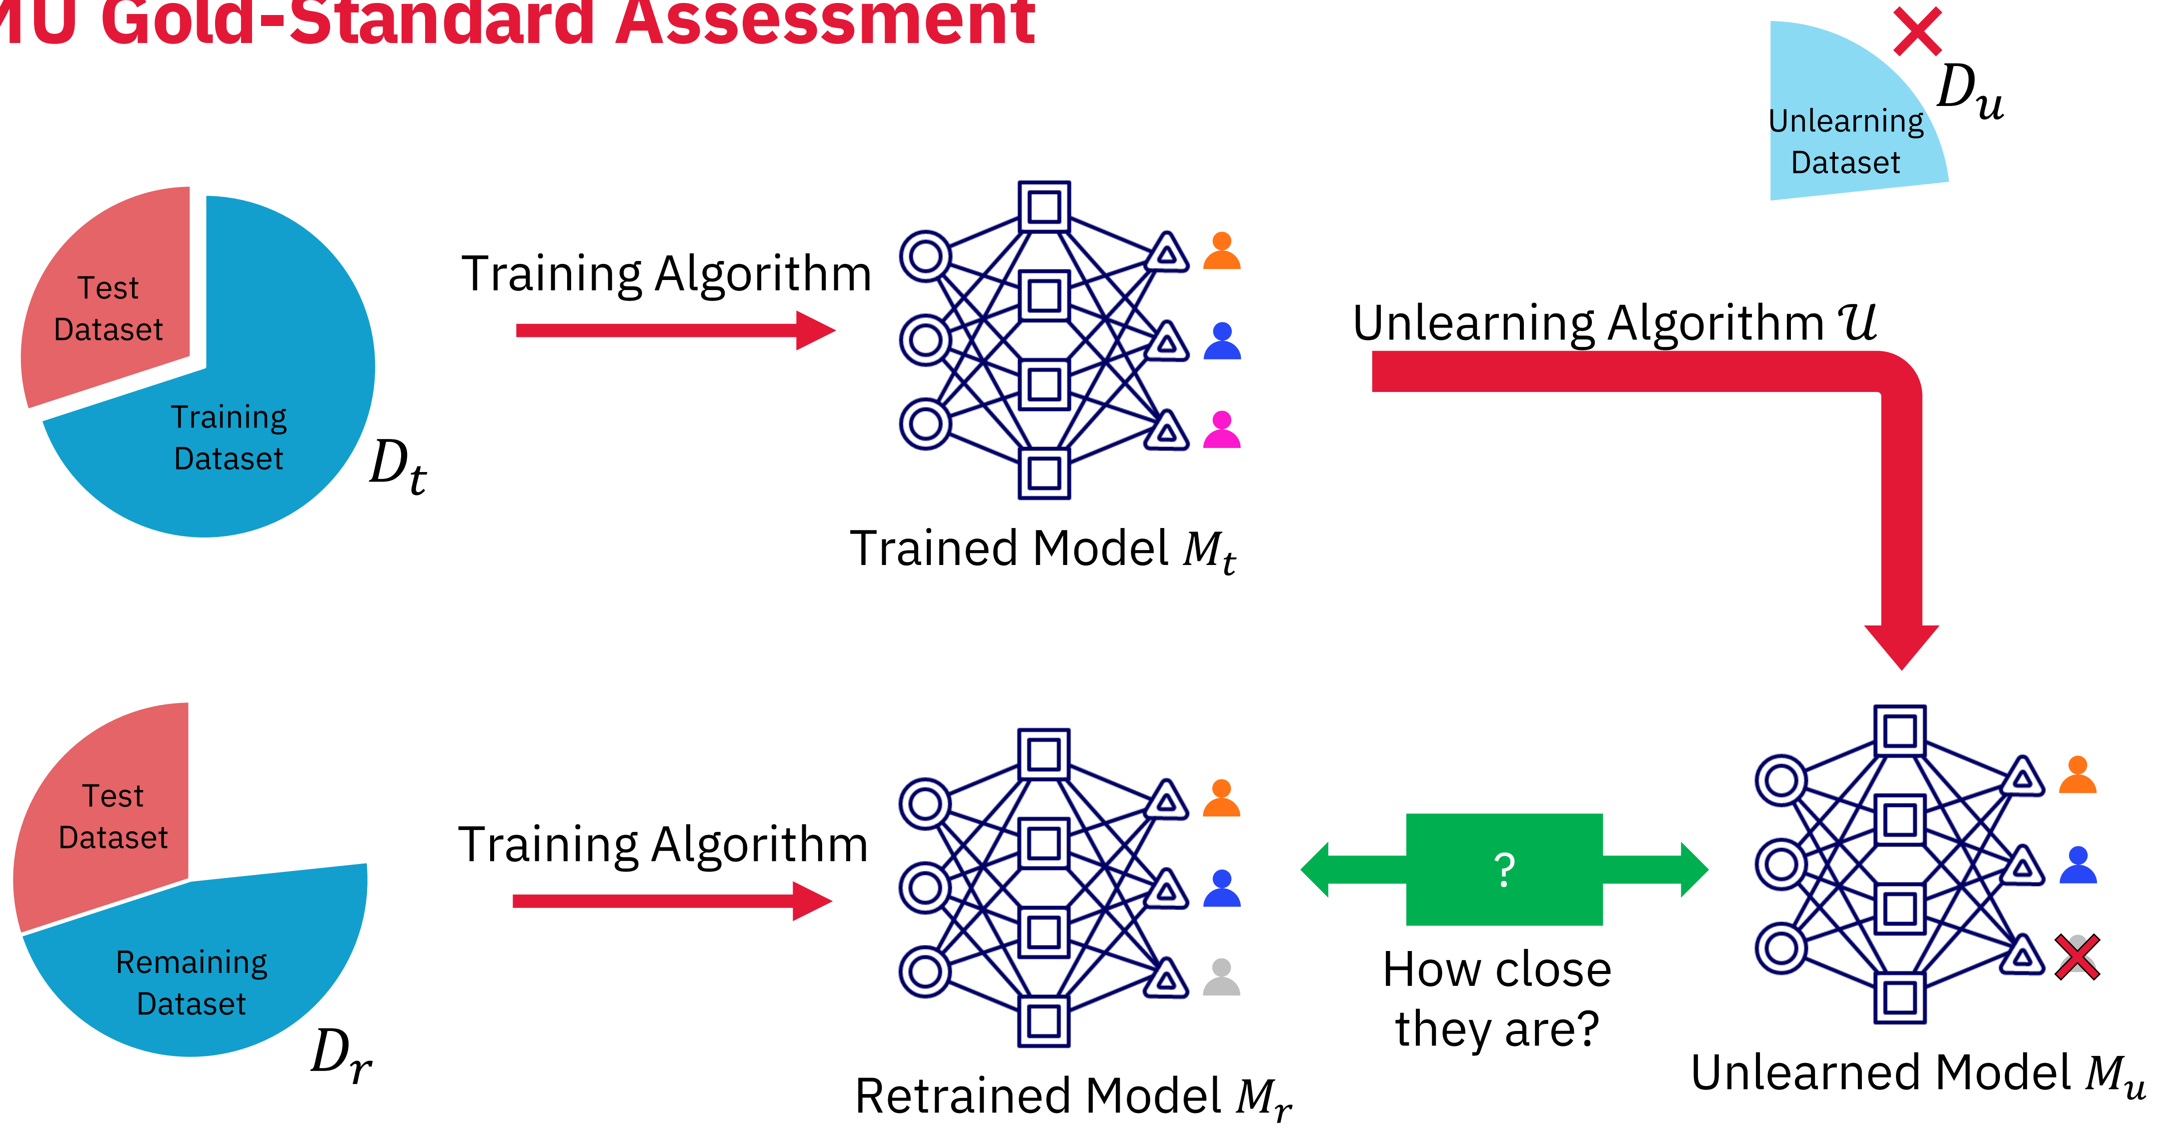

YORK U

Evaluation Metrics

- Accuracy
  - Unlearning Accuracy (UA) [Unlearning Efficacy]
  - Remaining Accuracy (RA) [Model Utility]
  - Test Accuracy (TA) [Model Utility]
- Membership Inference Attack (MIA) [Unlearning Efficacy]
- Speedup [Computation Efficiency]

RQ1) Baseline Improvements

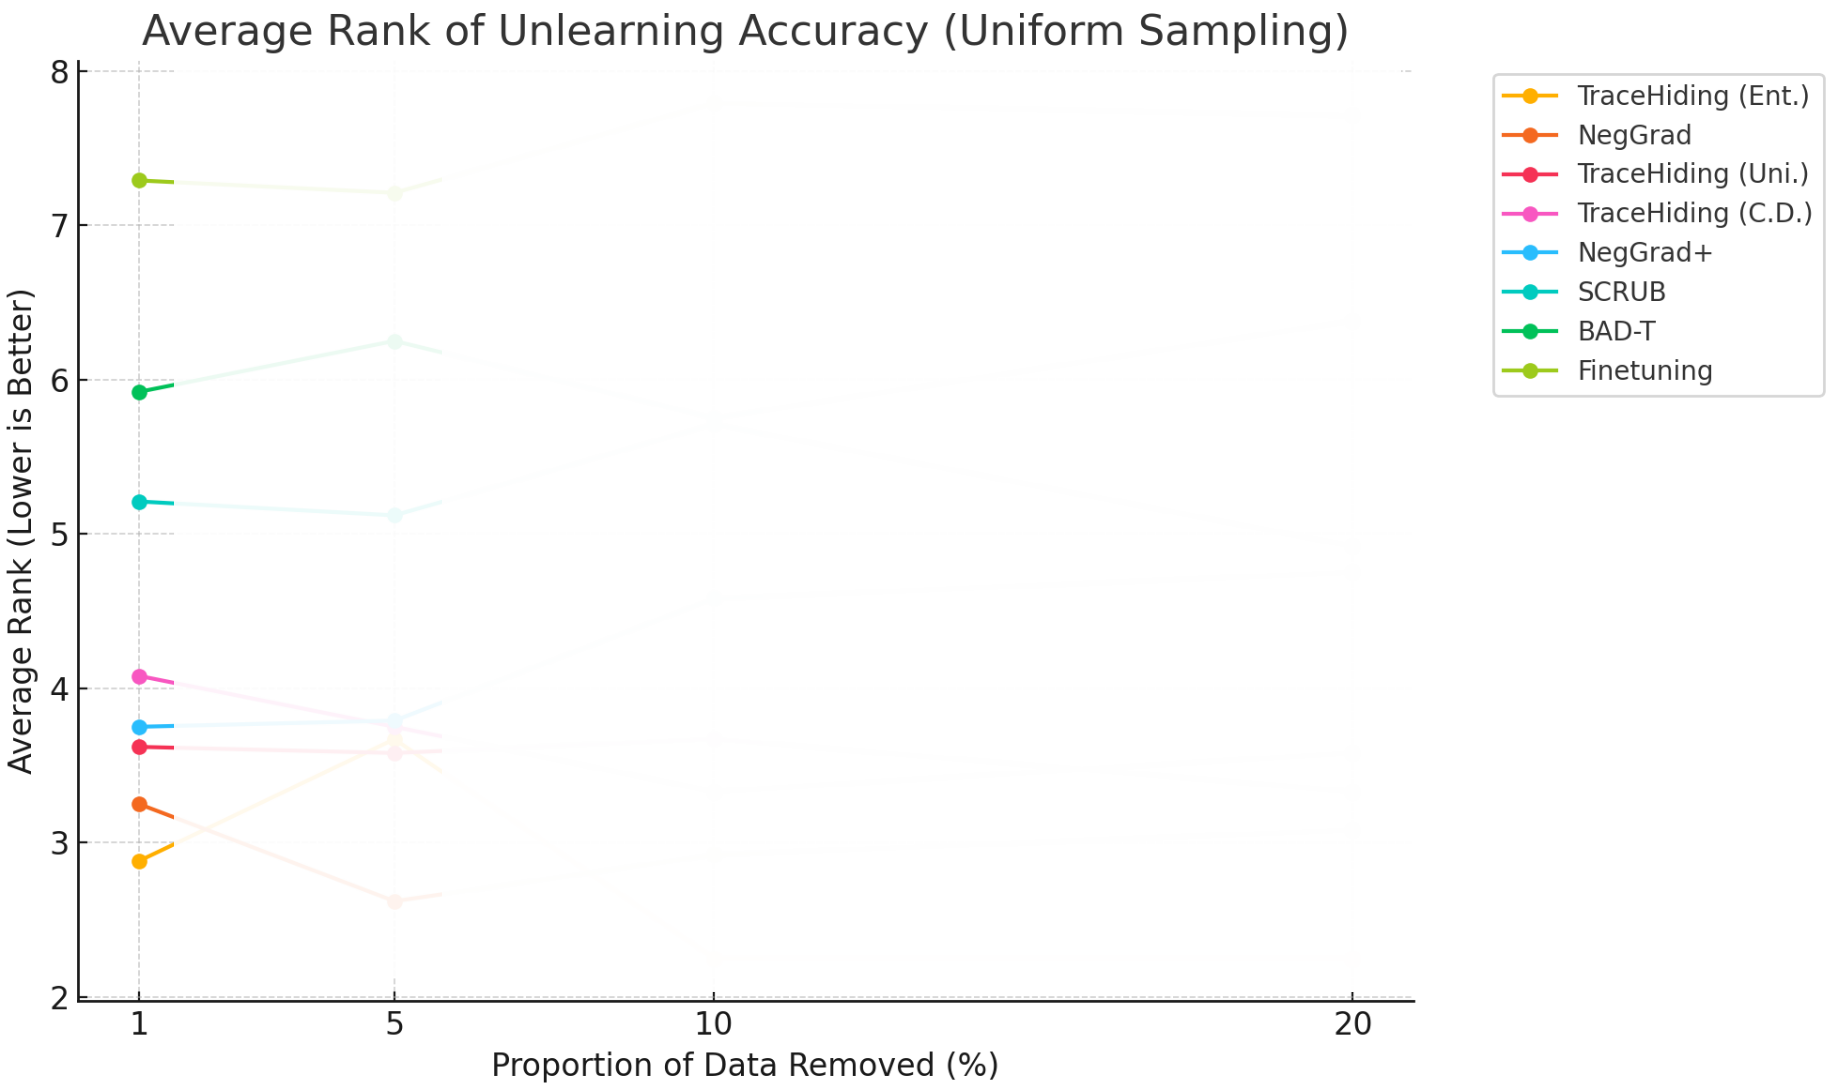

RQ2) Choosing Different Importance Score Definitions

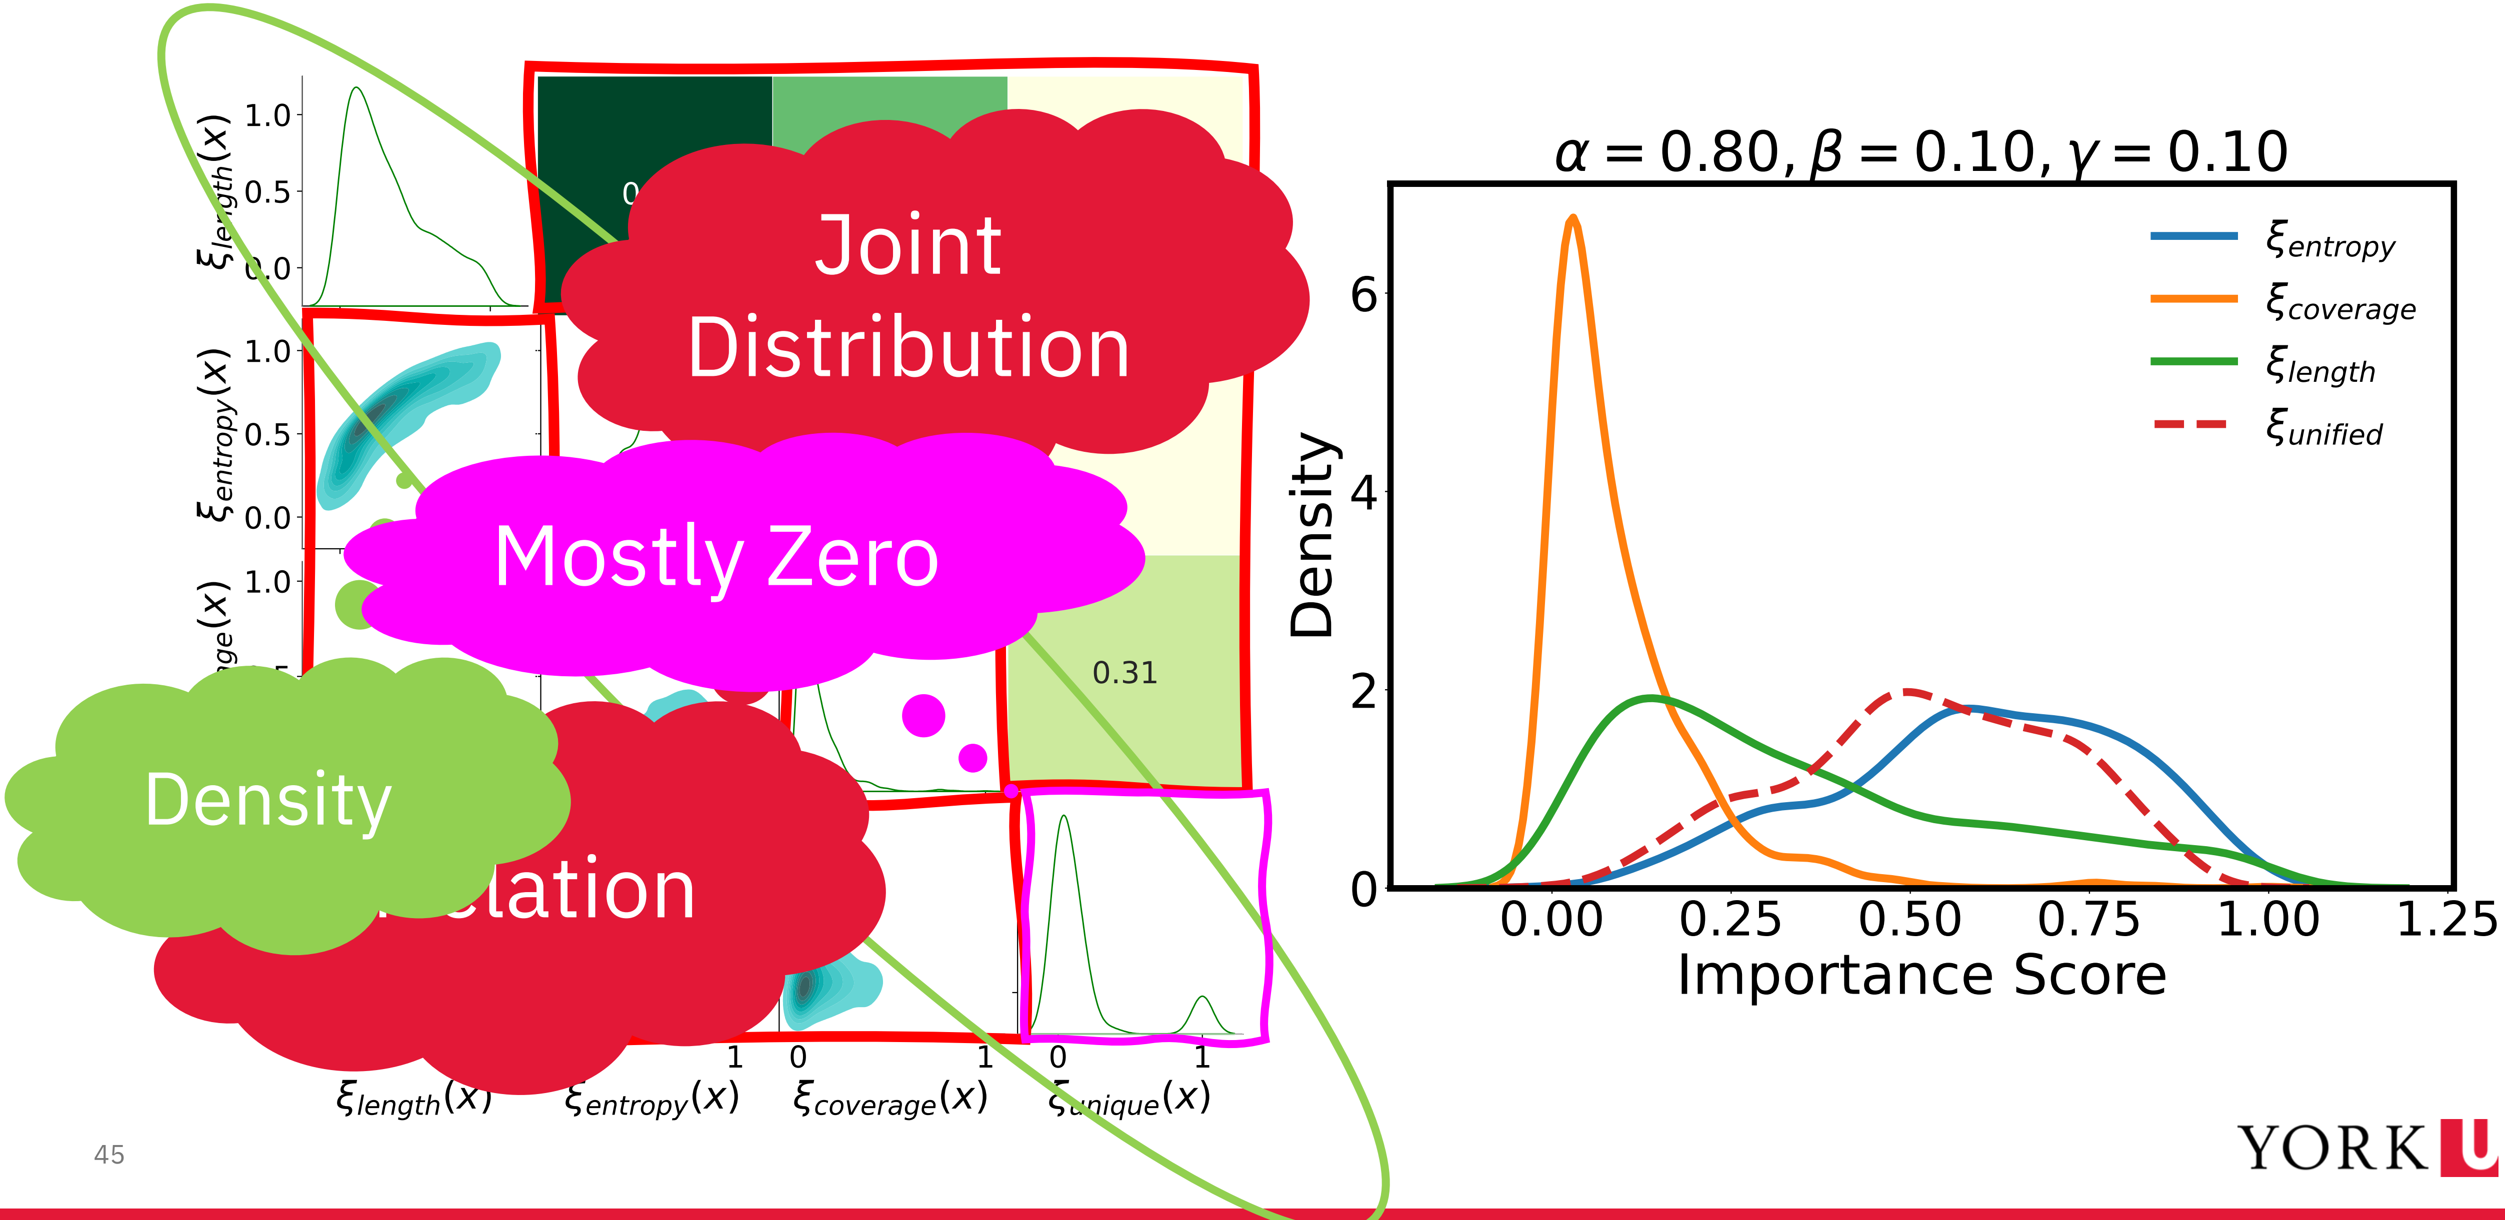

RQ3) Effect of Importance Score

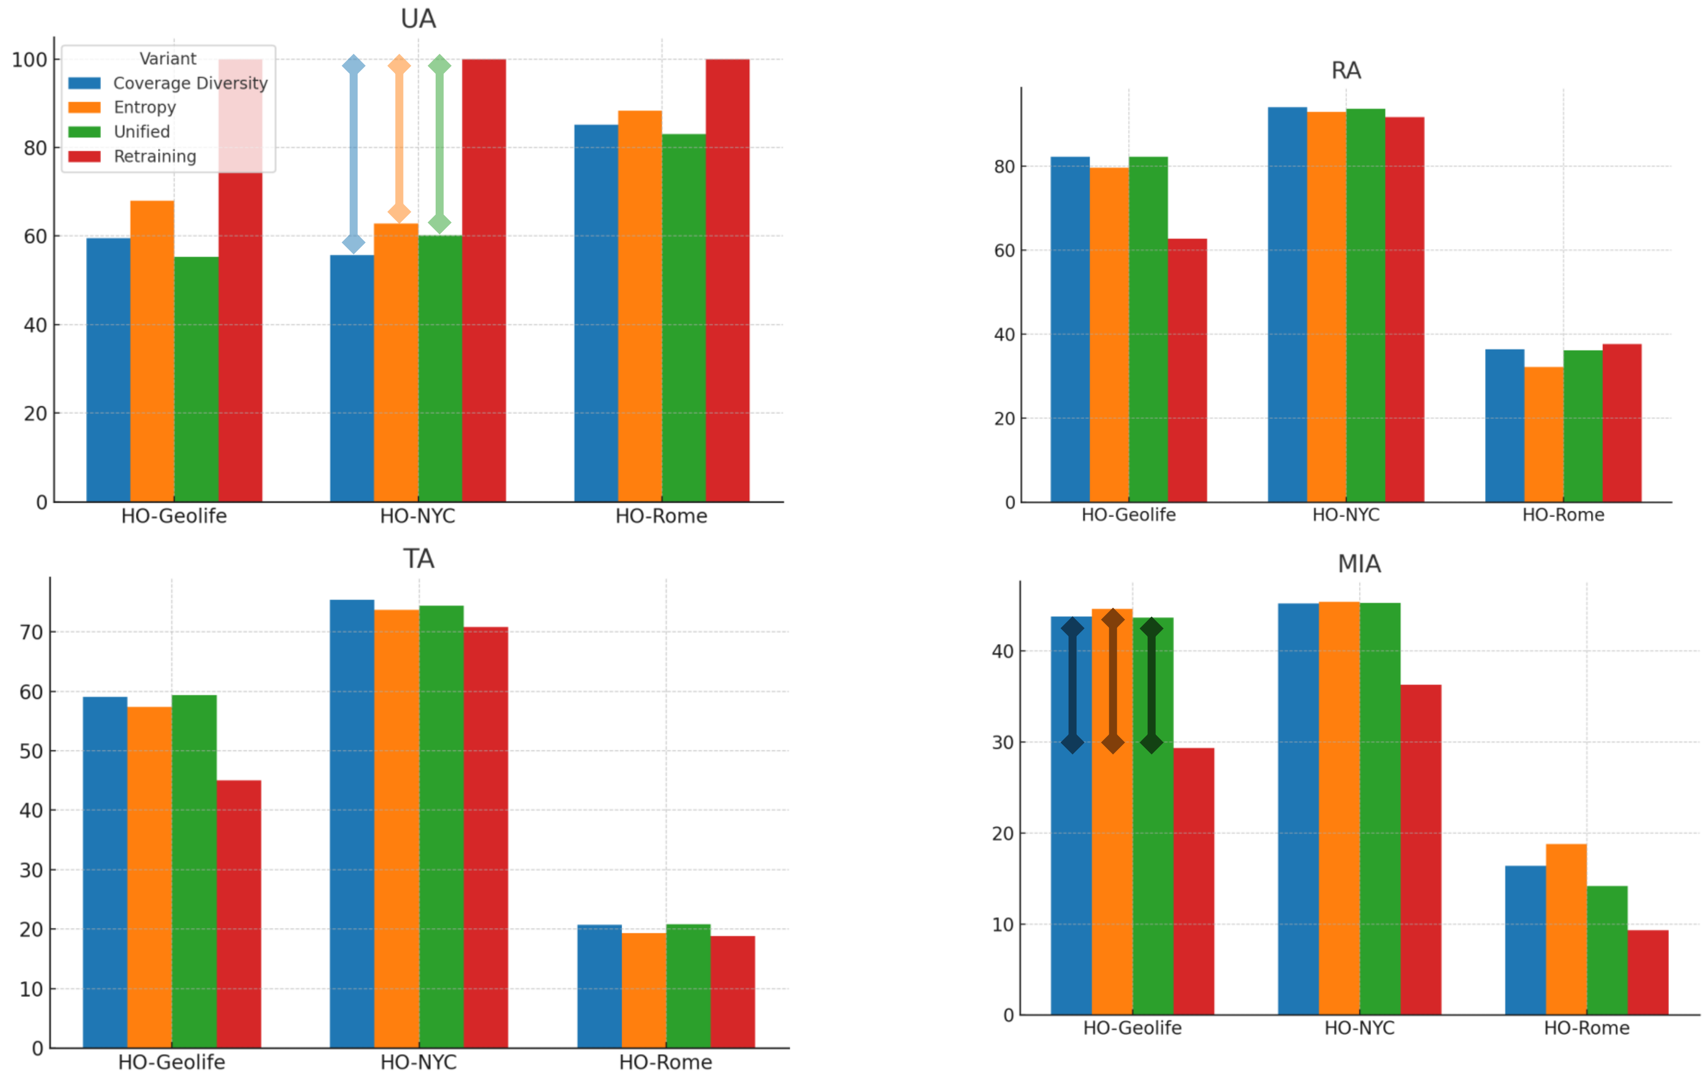

RQ4) Uniform vs Targeted Sampling

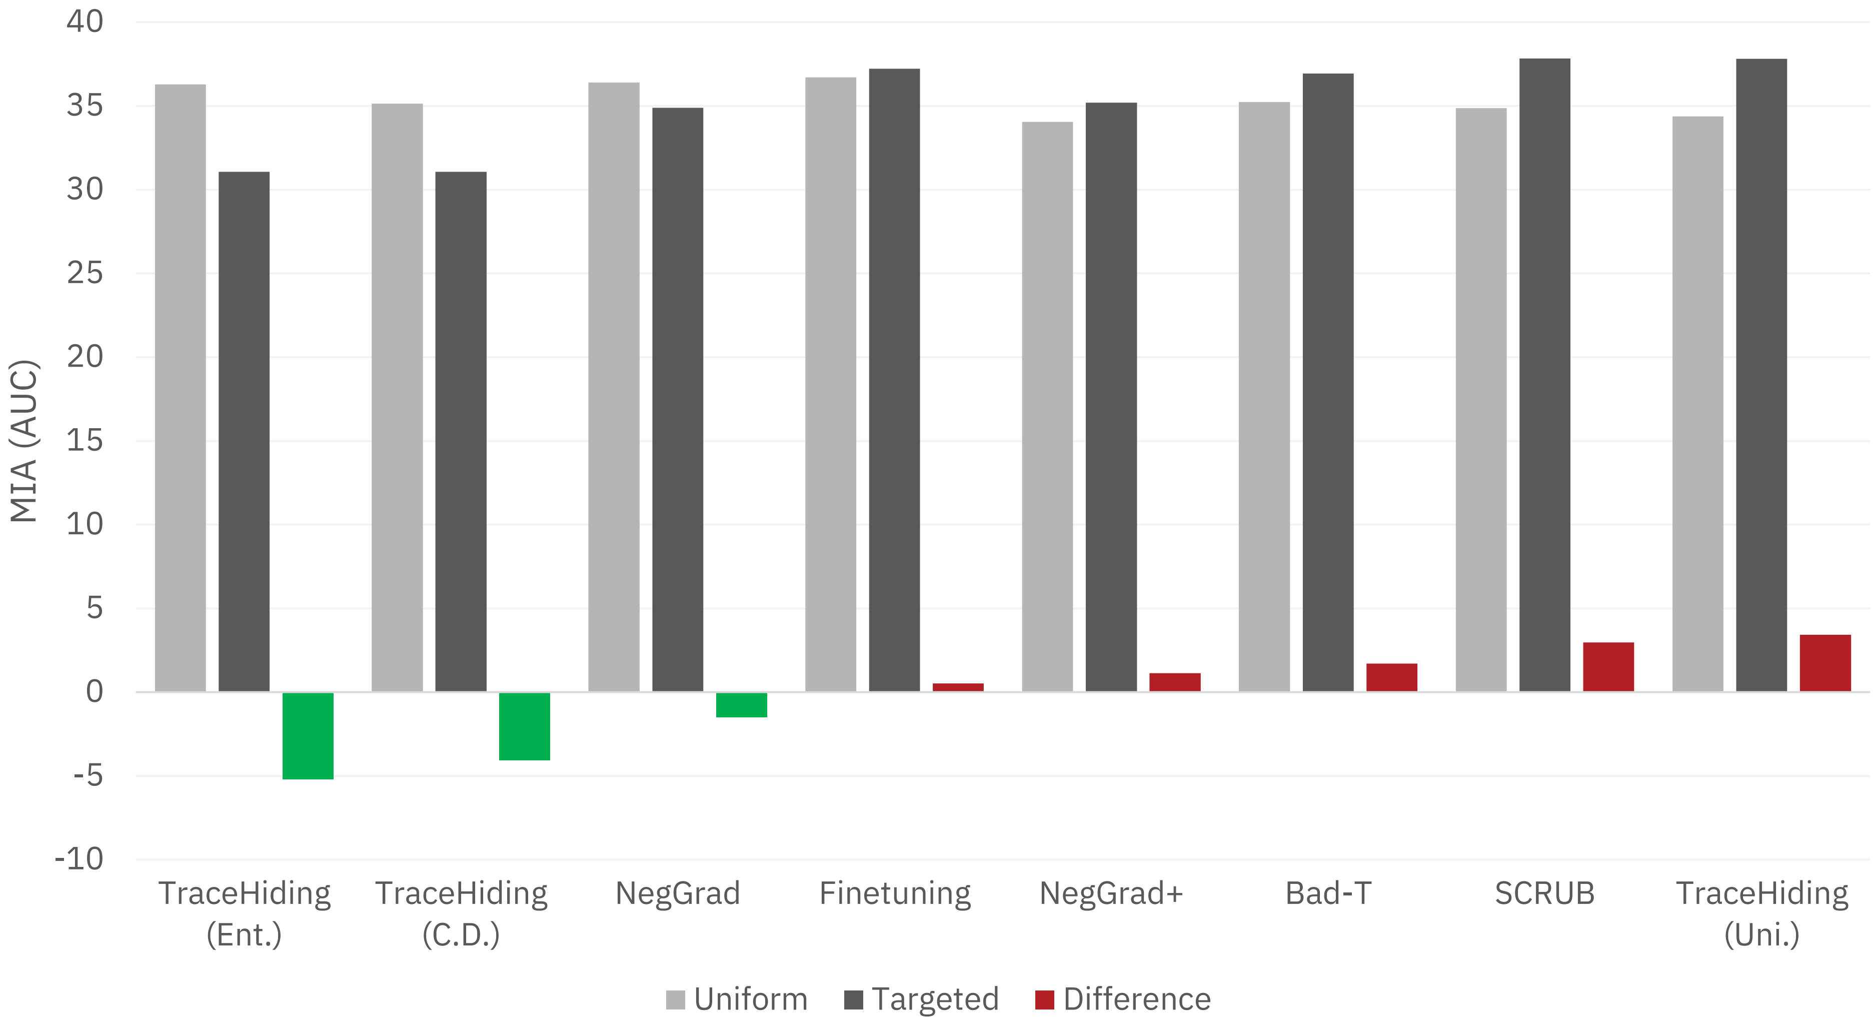

YORK U

RQ5) Computational Cost

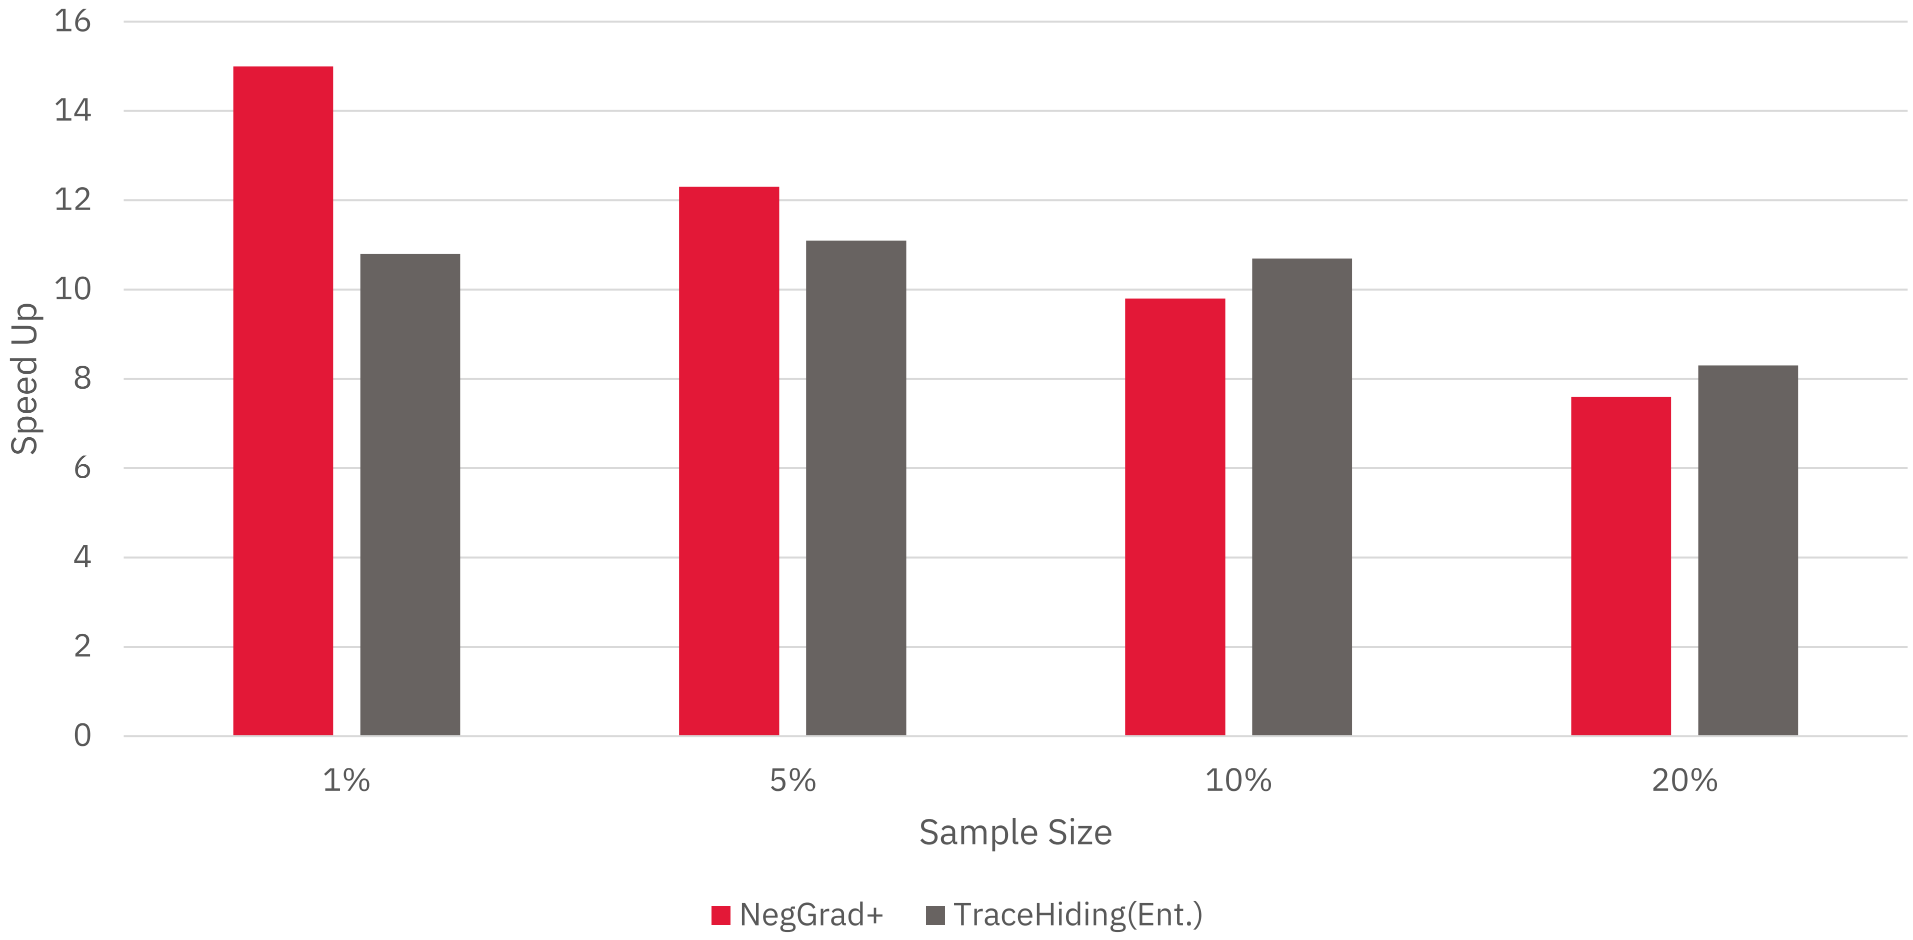

YORK U

Conclusions

Summary

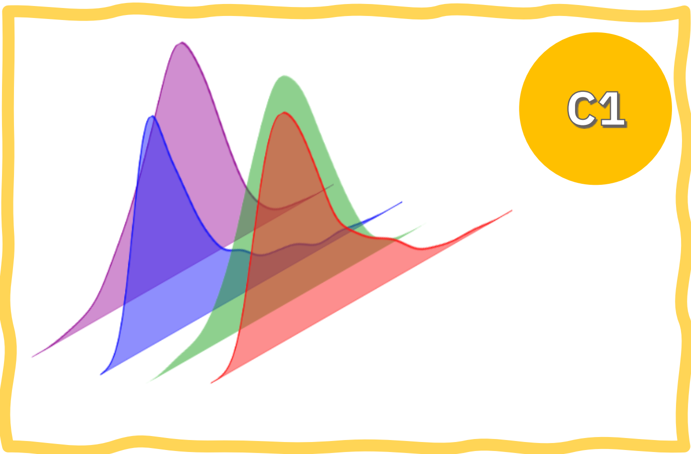

Data-driven Importance Score

- Token Level
- Trajectory Level
- User Level

Hierarchical Importance

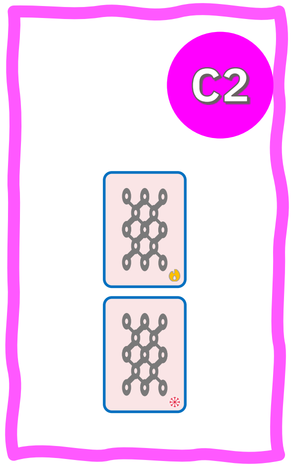

Teacher–Student Distillation

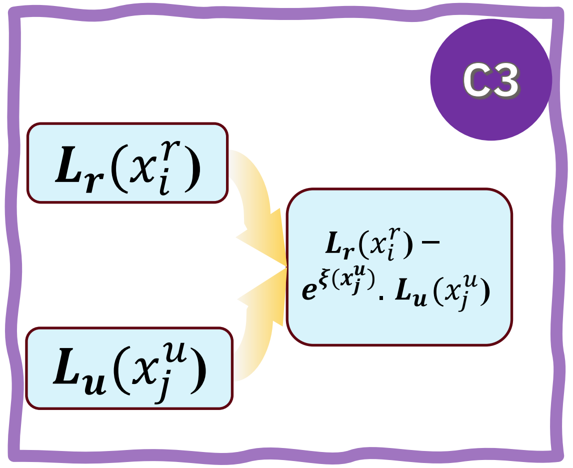

Loss Function

YORK U

Limitations

- No Certified Guarantees for Deletion

- No Interpretability in the Unlearning Process

Thesis Contributions

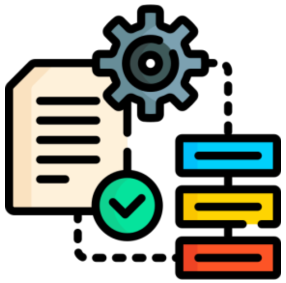

An Algorithmic Framework
TraceHiding with Variants

TraceHiding
Open-Source Software

Poin2Hex
(HO Mobility Dataset)

Research Papers:

- **Ali Faraji**, Jing Li, Gian Alix, Mahmoud Alsaeed, Nina Yanin, Amirhossein Nadiri, and Manos Papagelis. 2023. Point2Hex: Higher-order mobility flow data and resources. In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems (ACM SIGSPATIAL '23)*. ACM, New York, NY, USA, Article 69, 1–4.

- **Ali Faraji** and Manos Papagelis. 2025. TraceHiding: An algorithmic framework for machine unlearning in mobility data. *ACM Transactions on Spatial Algorithms and Systems* (submitted).

Future Works

- Beyond Classification: Other Tasks

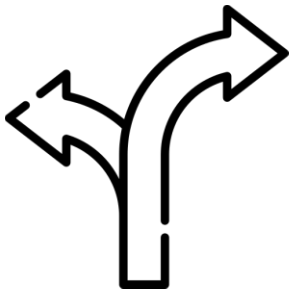

- Richer Importance Signals

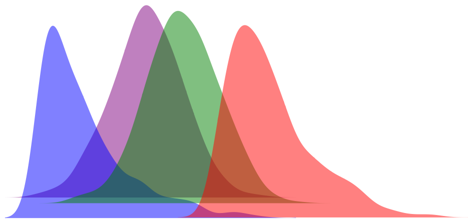

- Streaming and Continual Settings

- Fairness and Compliance Analysis

Thank You!

Questions?