Automated Expansion of Privacy Data Taxonomy for Compliant Data Breach Notification

Yue Qin*
Indiana University Bloomington
&
Central University of Finance and Economics

Yue Xiao*
Indiana University Bloomington
&
IBM Research

Xiaojing Liao
Indiana University Bloomington

*These authors contributed equally to this work, ordered alphabetically. This work was completed while Yue Qin and Yue Xiao were Ph.D. students at Indiana University Bloomington.

Privacy Compliance Check

Compare Data Items

Terminology Gap [1]

Privacy Laws/Regulations

Data Usage Practices

Legal Professionals

- PII
- Location
- Contact
- Identifiers

Restricted Data

Privacy Data Taxonomy

- Access token
- IDFA
- Device topic
- IP addresses

Technicians

[1] Understanding Legal Professionals' Practices and Expectations in Data Breach Incident Reporting (CCS 2024)

Source of Figures: https://www.cookieyes.com/, https://legacy.teltik.com/resources/what-is-data-usage/

Privacy Data Taxonomy

What is Privacy Data Taxonomy?

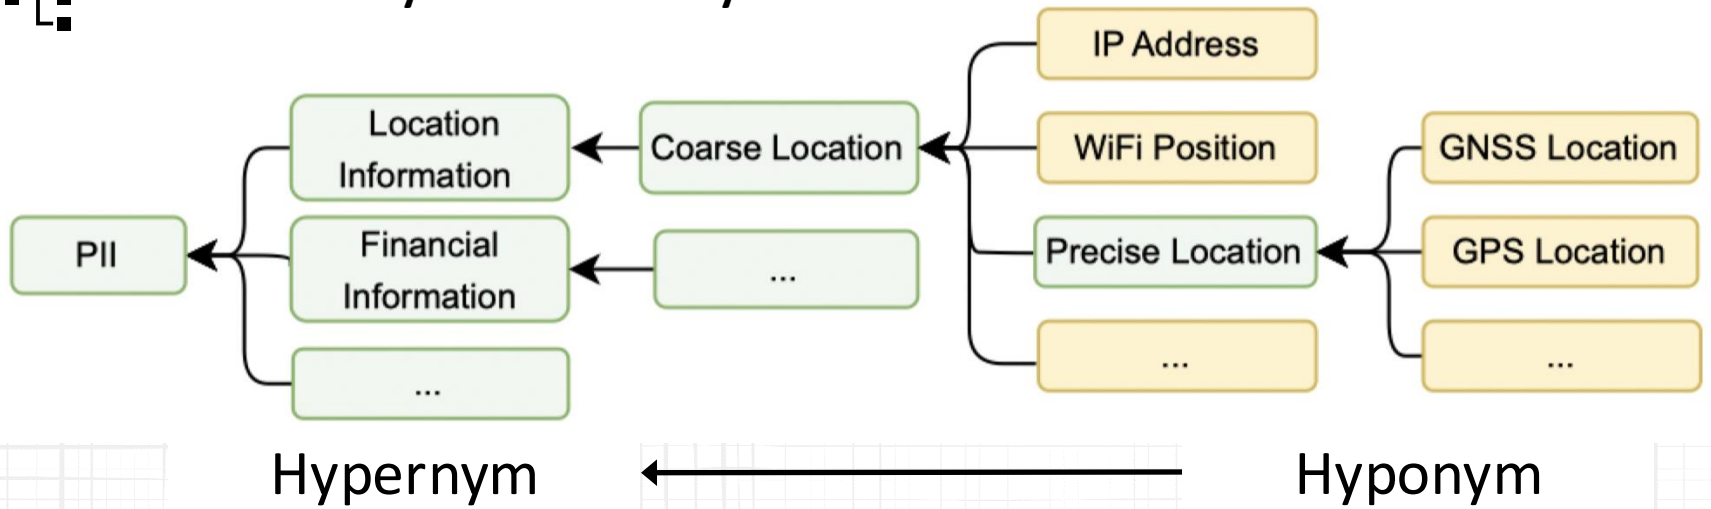

Challenges in building Privacy Data Taxonomy:

- Data item diversity across applications (e.g., IoT, Apps, SDKs, etc.).
- Broad, vague and varying interpretations of privacy data across jurisdictions.
- Existing privacy taxonomies rely on manual efforts ([1][2]…) and heuristics ([3]…).

[1] L. Elluri, A. Nagar, and K. P. Joshi, "An integrated knowledge graph to automate gdpr and pci dss compliance," in 2018 IEEE International Conference on Big Data (Big Data). IEEE, 2018, pp. 1266–1271.
[2] K. P. Joshi, L. Elluri, and A. Nagar, "An integrated knowledge graph to automate cloud data compliance," IEEE Access, vol. 8, pp. 148 541148 555, 2020.
[3] B. Andow, S. Y. Mahmud, W. Wang, J. Whitaker, W. Enck, B. Reaves, K. Singh, and T. Xie, "Policylint: Investigating internal privacy policy contradictions on google play." in USENIX Security Symposium, 2019, pp. 585–602.

Automatic Method: Hypernym Prediction

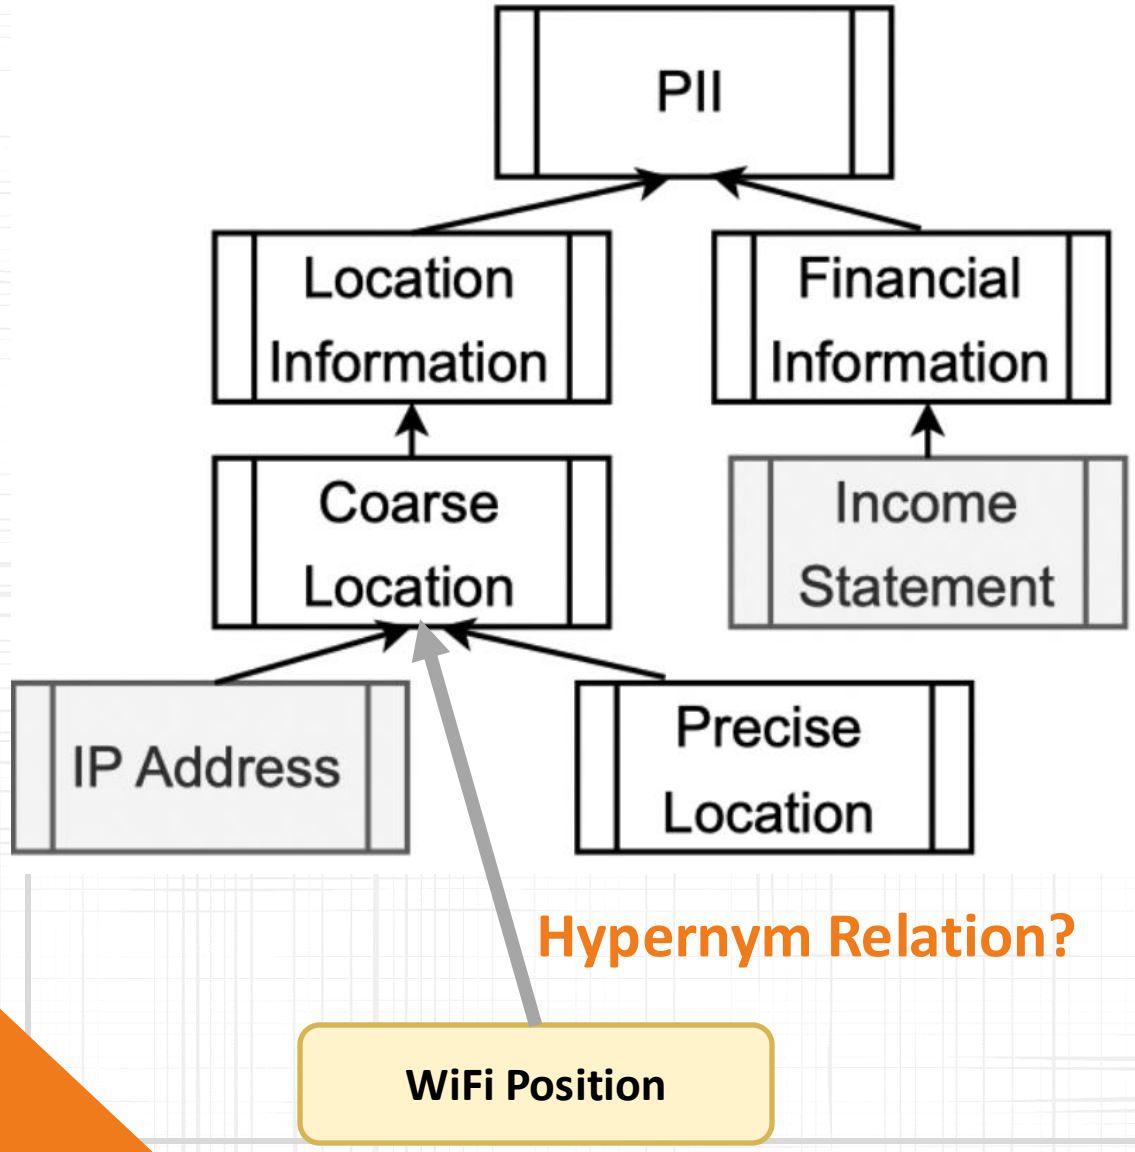

Existing hypernym prediction techniques in *open domains*:

- Lexical pattern-based methods [12], [13], [10], [14], [11]
- Distributional representation-based methods [15], [16], [17], [18], [19]
- Projection-based methods [20], [21], [22]

Challenge: Fail to consider **granularity levels**—the precision differences in hypernym relationships

- For example, methods can identify broad categories like “location information” and “financial information”
- but struggle to distinguish "coarse location" (e.g., IP address) from "precise location" (e.g., GPS).

GraSP: Granularity-aware Hypernym Prediction

Clusters for Learning the Projection Matrices from the Hyponyms to the Hypernyms

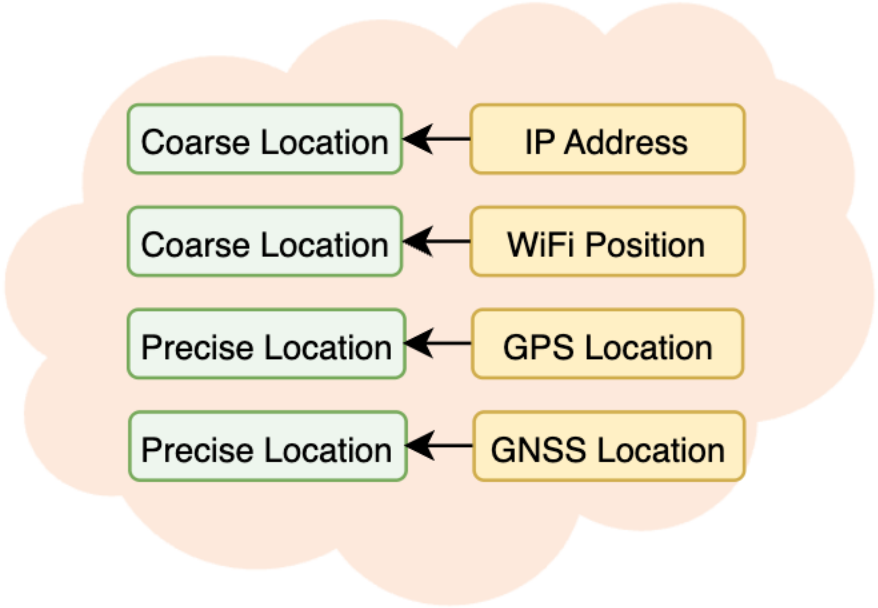

Approximated *contextual* features of coarse/precise

Previous Projection-based Method

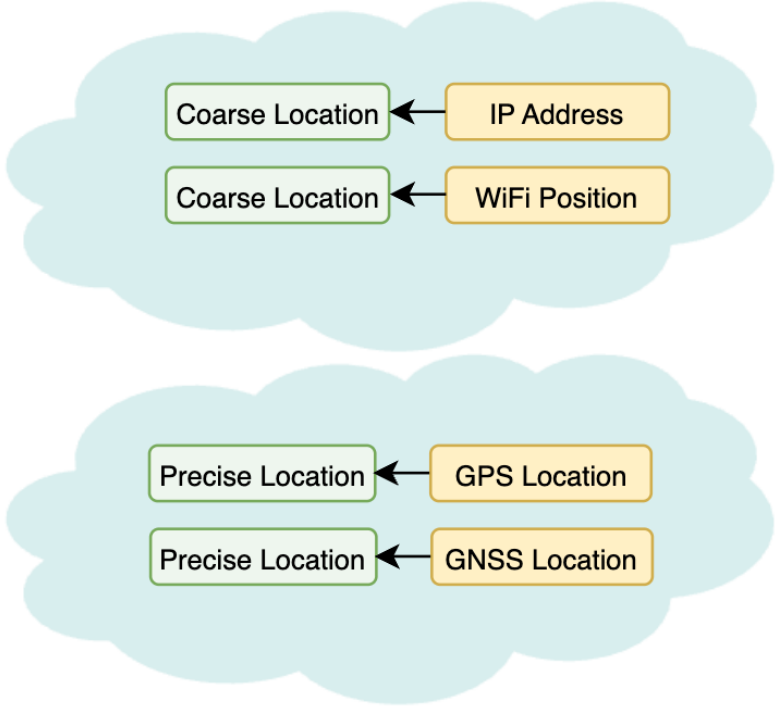

Guide clustering with taxonomy *structure*

Our Method

GraSP: Granularity-aware Hypernym Prediction

Clusters for Learning the Projection Matrices from the Hyponyms to the Hypernyms

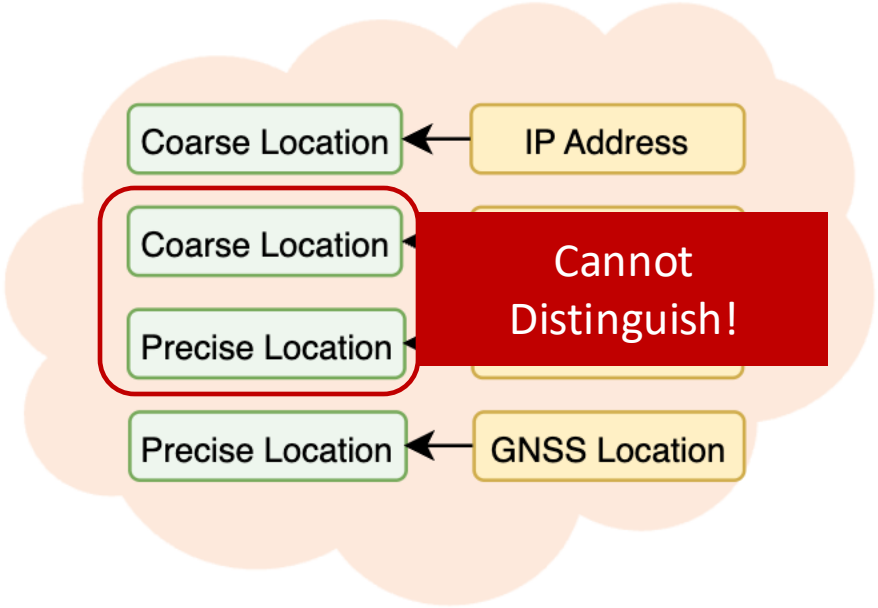

Approximated *contextual* features of coarse/precise

Previous Projection-based Method

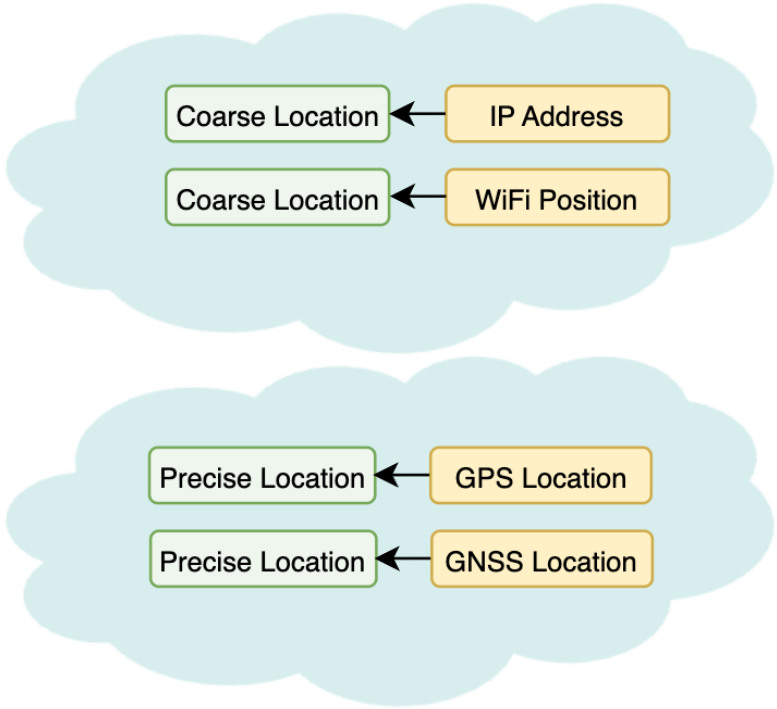

Guide clustering with taxonomy *structure*

Our Method

Architecture of Hypernym Prediction Model

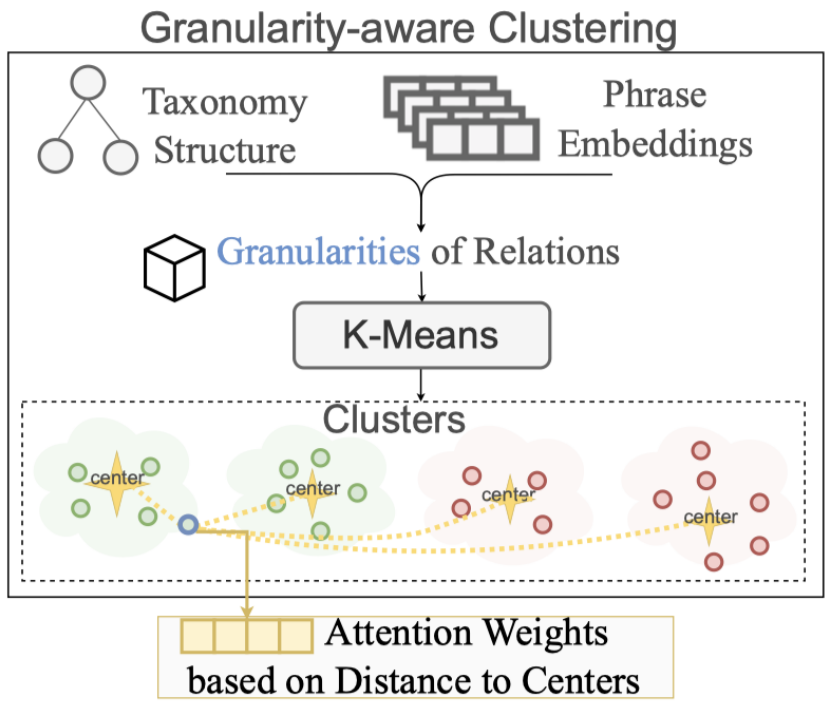

Architecture of Hypernym Prediction Model

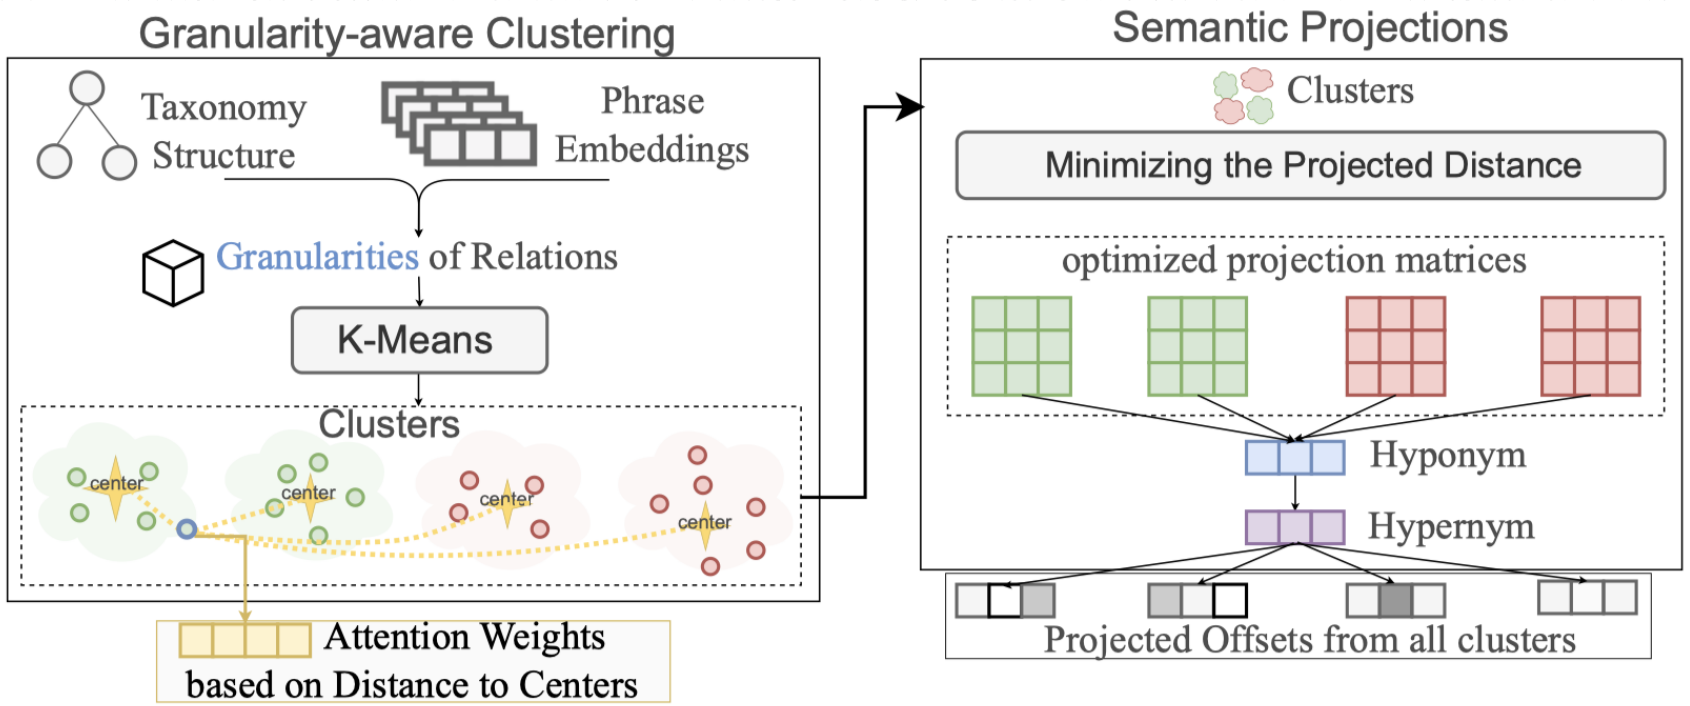

Architecture of Hypernym Prediction Model

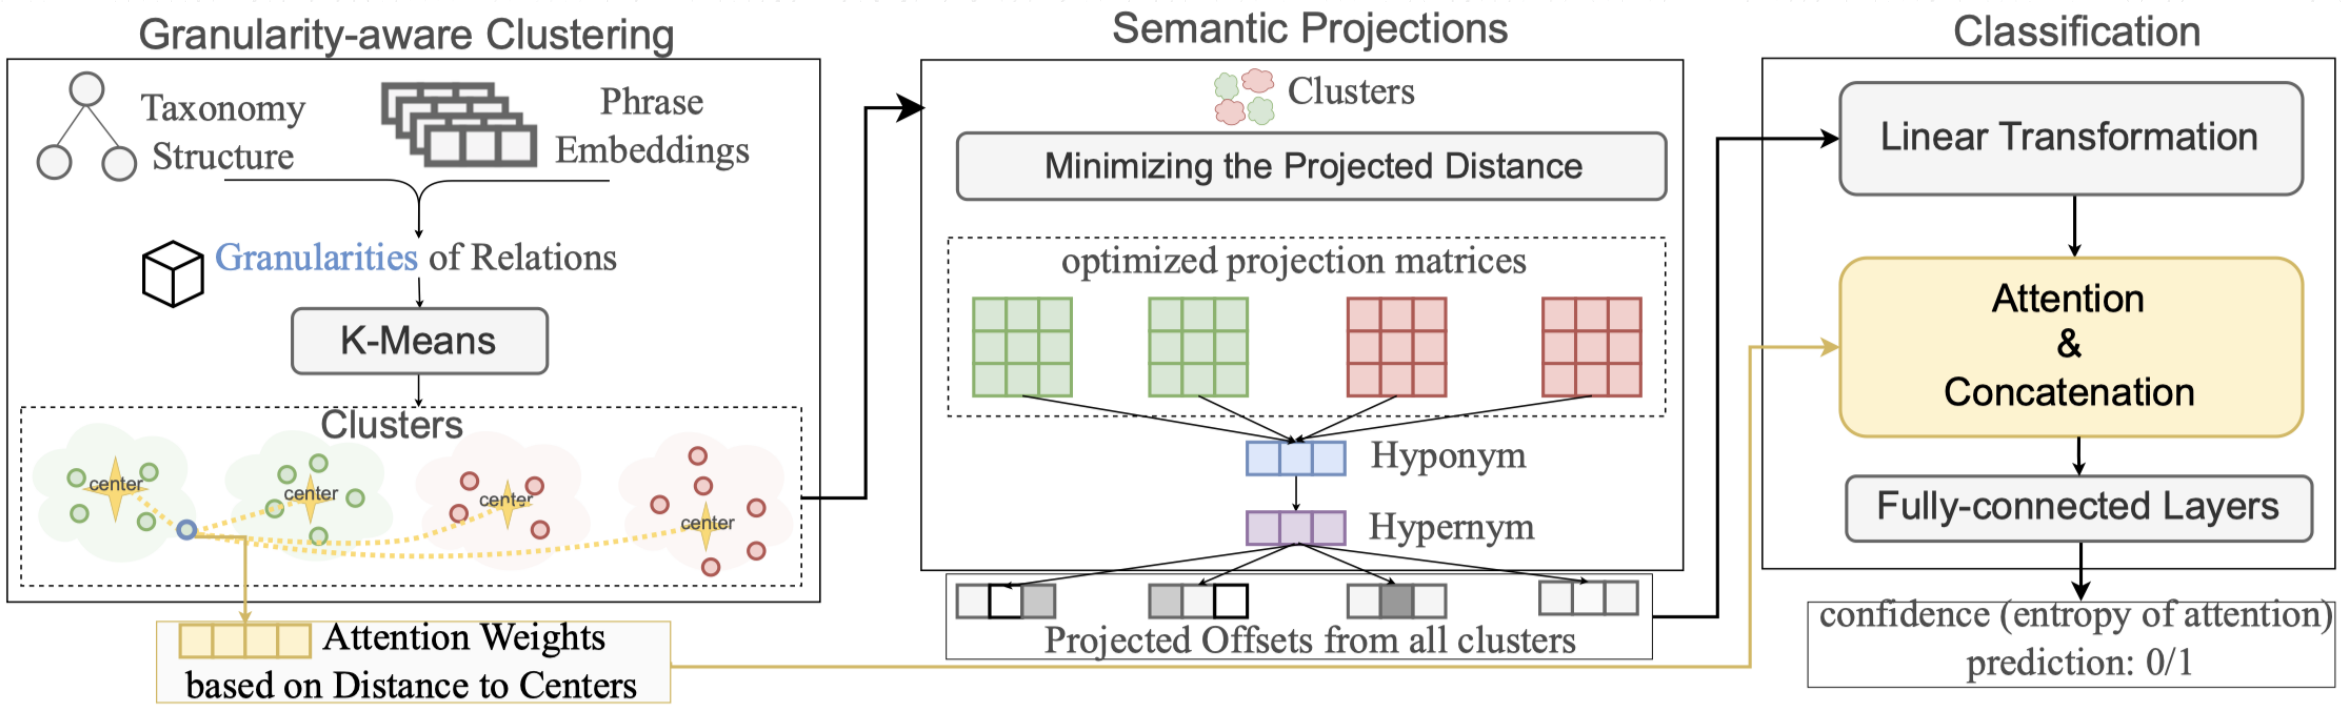

Evaluation: Datasets

Privacy Data Taxonomies

Privacy Policy Taxonomy: 680 restricted data, 2,176 hypernym relations

Refined PolicyLint's Ontologies [1] + High-level Data Categories in GDPR

IoT Sensitive Data Taxonomy: 76 restricted data, 138 hypernymy relations.

Expert Annotated IoT Privacy-sensitive Data in IoTProfiler [3]

Groundtruth: Hypernym-hyponym Pairs

Randomly sample 5 negative pairs for each positive hypernym-hyponym pair

Privacy Policy Dataset: 2,176 positive pairs; 10,800 negative pairs

IoT Dataset: 138 positive pairs; 690 negative pairs

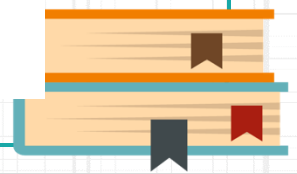

Evaluation: Results

Method	PrivacyPolicy Dataset				IoT Dataset			
	Precision	Recall	F1	Micro F1	Precision	Recall	F1	Micro F1
Concat + LR	0.765 ± .00	0.871 ± .00	0.815 ± .00	0.947 ± .00	0.701 ± .00	0.701 ± .00	0.701 ± .00	0.933 ± .00
Offset + LR	0.671 ± .00	0.843 ± .00	0.747 ± .00	0.929 ± .00	0.624 ± .00	0.756 ± .00	0.684 ± .00	0.910 ± .00
Concat + MLP	0.889 ± .02	0.927 ± .02	0.907 ± .01	0.972 ± .00	0.701 ± .02	0.894 ± .04	0.786 ± .05	0.943 ± .02
Offset + MLP	0.866 ± .03	0.871 ± .02	0.868 ± .01	0.962 ± .01	0.788 ± .06	0.741 ± .09	0.764 ± .07	0.955 ± .02
SphereRE (N)[P]	0.774 ± .05	0.764 ± .03	0.768 ± .02	0.934 ± .01	0.778 ± .05	0.822 ± .10	0.794 ± .03	0.955 ± .01
SphereRE (O)[P]	0.870 ± .03	0.863 ± .03	0.866 ± .02	0.962 ± .01	0.660 ± .11	0.800 ± .12	0.715 ± .08	0.931 ± .03
SphereRE (N)	0.902 ± .02	0.890 ± .02	0.896 ± .02	0.971 ± .01	0.721 ± .09	0.822 ± .06	0.765 ± .06	0.945 ± .02
SphereRE (O)	0.904 ± .02	0.901 ± .02	0.903 ± .02	0.972 ± .00	0.741 ± .03	0.822 ± .13	0.776 ± .06	0.950 ± .01
MWP (N)[P]	0.817 ± .03	0.920 ± .03	0.865 ± .02	0.959 ± .01	0.664 ± .05	**1.000** ± .00	0.798 ± .03	0.945 ± .01
MWP (O)[P]	0.721 ± .03	**0.990** ± .00	0.834 ± .02	0.944 ± .01	0.597 ± .06	**1.000** ± .00	0.746 ± .05	0.926 ± .02
MWP (N)	0.917 ± .03	0.893 ± .03	0.905 ± .02	0.973 ± .01	0.771 ± .10	0.844 ± .06	0.802 ± .07	0.955 ± .02
MWP (O)	0.907 ± .03	0.899 ± .01	0.903 ± .01	0.972 ± .00	0.733 ± .06	0.844 ± .06	0.784 ± .05	0.950 ± .01
GPT-3.5 [Finetune+Prompt]	0.867 ± 0.00	0.867 ± 0.00	0.867 ± 0.00	0.962± 0.00	0.667 ± 0.00	0.889 ± 0.00	0.762 ± 0.00	0.940 ± 0.00
GRASP	**0.974** ± .01	0.952 ± .00	**0.963** ± .00	**0.990** ± .00	**0.910** ± .09	0.889 ± .08	**0.899** ± .05	**0.976** ± .01

Tracy – Application for GDPR Compliance

localhost:8080

iRobot Attack Case

An attacker can remotely issue commands to an iRobot Roomba of a victim user through the aforementioned LAST WILL mesasge. Here are the steps of the exploit.

Step a. Buy a WiFi-enabled iRobot Roomba and register a **user account**

Step b. login with **user account** and record the **AccessKeyId** ,**SecretKey** and **SessionToken** for further connecting to iRobot's AWS IoT endpoint (a2uowfjvhio0fa.iot.us-east-1.amazonaws.com) via MQTT over websocket.

Step c. The attacker uses a malicious program integrating AWS IoT SDK to connect to AWS IoT Core with **credentials** above, and set a LAST WILL message on the command topic of this iRobot. Now the malicious program can send messages/commands to AWS IoT Core to control the iRobot devices just like a legitimate app. Then, he keeps this malicious program online.

Step d. The attacker returns the device or gives it to the victim as a gift. The new user of the device (the victim) registers and uses the iRobot. Once the new user registers with the iRobot, any previous user's permissions with the iRobot device are revoked automatically by policy update. Now the malicious program no longer has permissions to send messages/commands to the iRobot device.

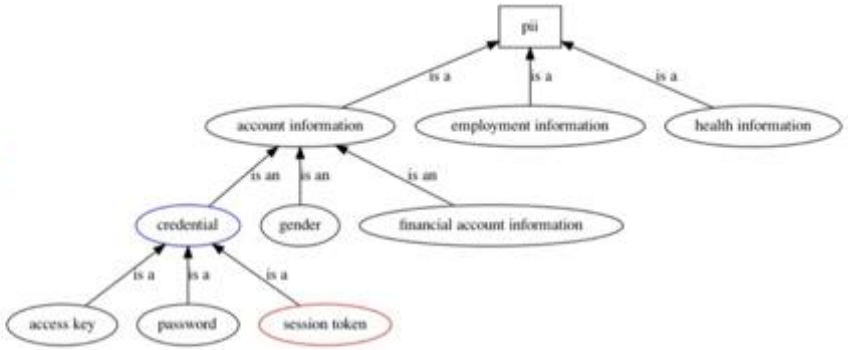

User Study

- **Participants:** 15 privacy professionals (legal and security experts).
- **Tasks:** Evaluated Tracy's ability to recognize privacy-sensitive data in real-world incident reports.

Findings

Traceability	**100%** agreement on 44 restricted data instances.
Efficiency	Reduced assessment time by **75.17%**.
Usability	**93%** of participants expressed willingness to use Tracy in privacy compliance tasks.

Takeaways

Code & Data

- We design and implement GRASP for automatically constructing and expanding privacy data taxonomy
- We design and implement *Tracy*, a privacy professional assistant to recognize and interpret private data in incident reports for GDPR-compliant data breach notification

Thank you!