

DPolicy: Managing Privacy Risks Across Multiple Releases with Differential Privacy



Nicolas Küchler

ETH zürich



Alexander Viand

intel.



Hidde Lycklama

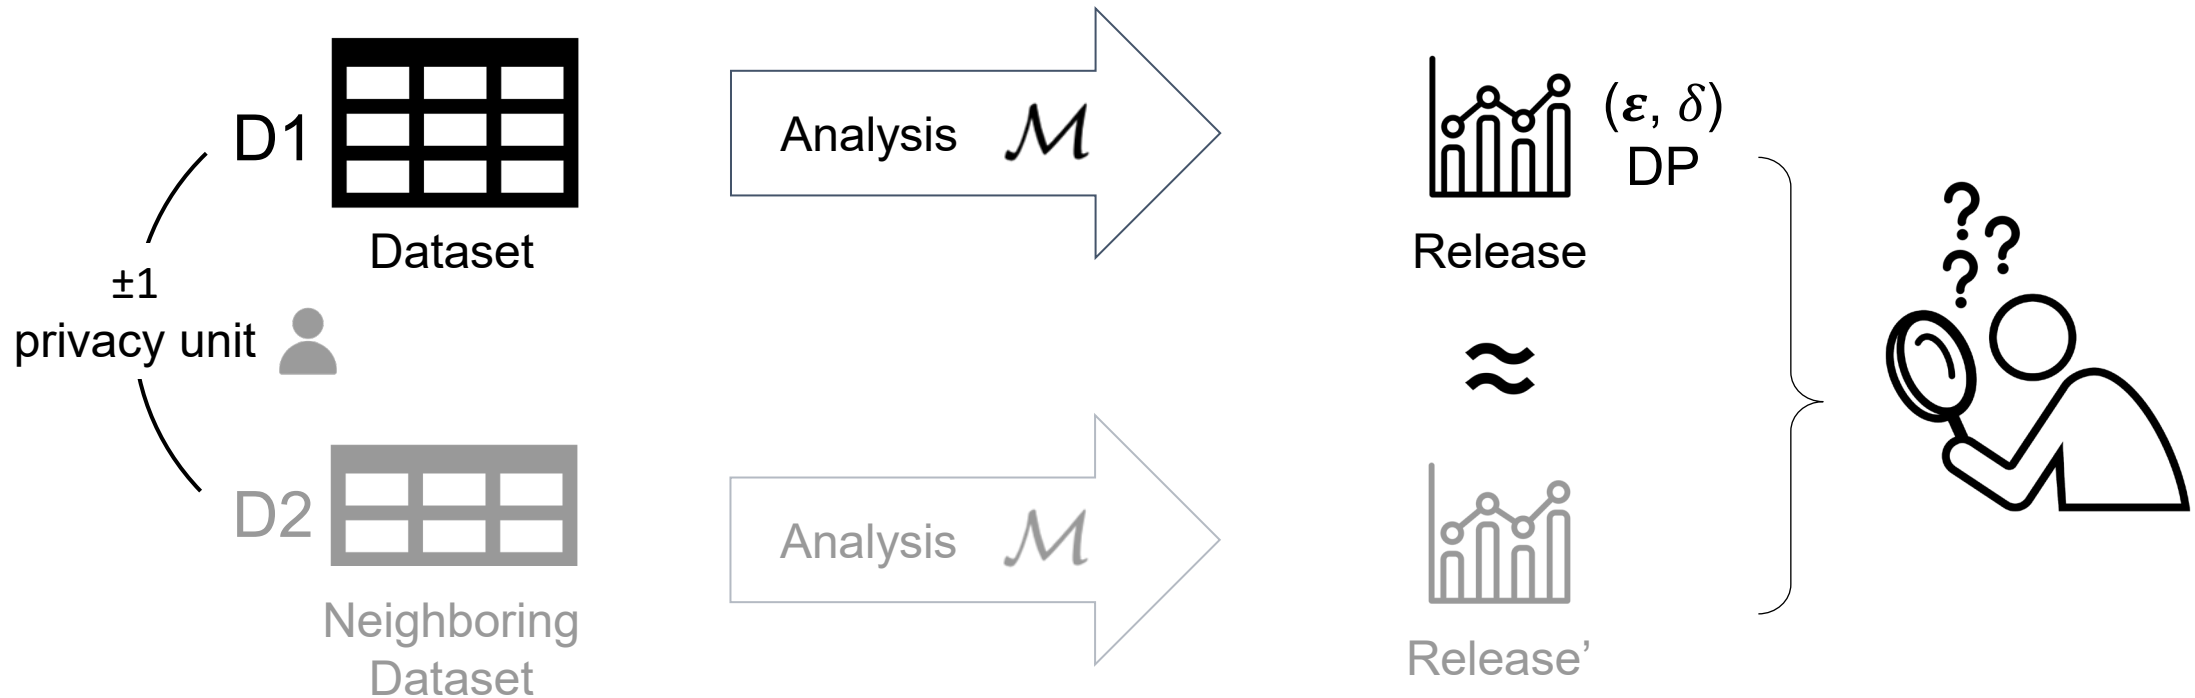
ETH zürich



Anwar Hithnawi

 UNIVERSITY OF
TORONTO

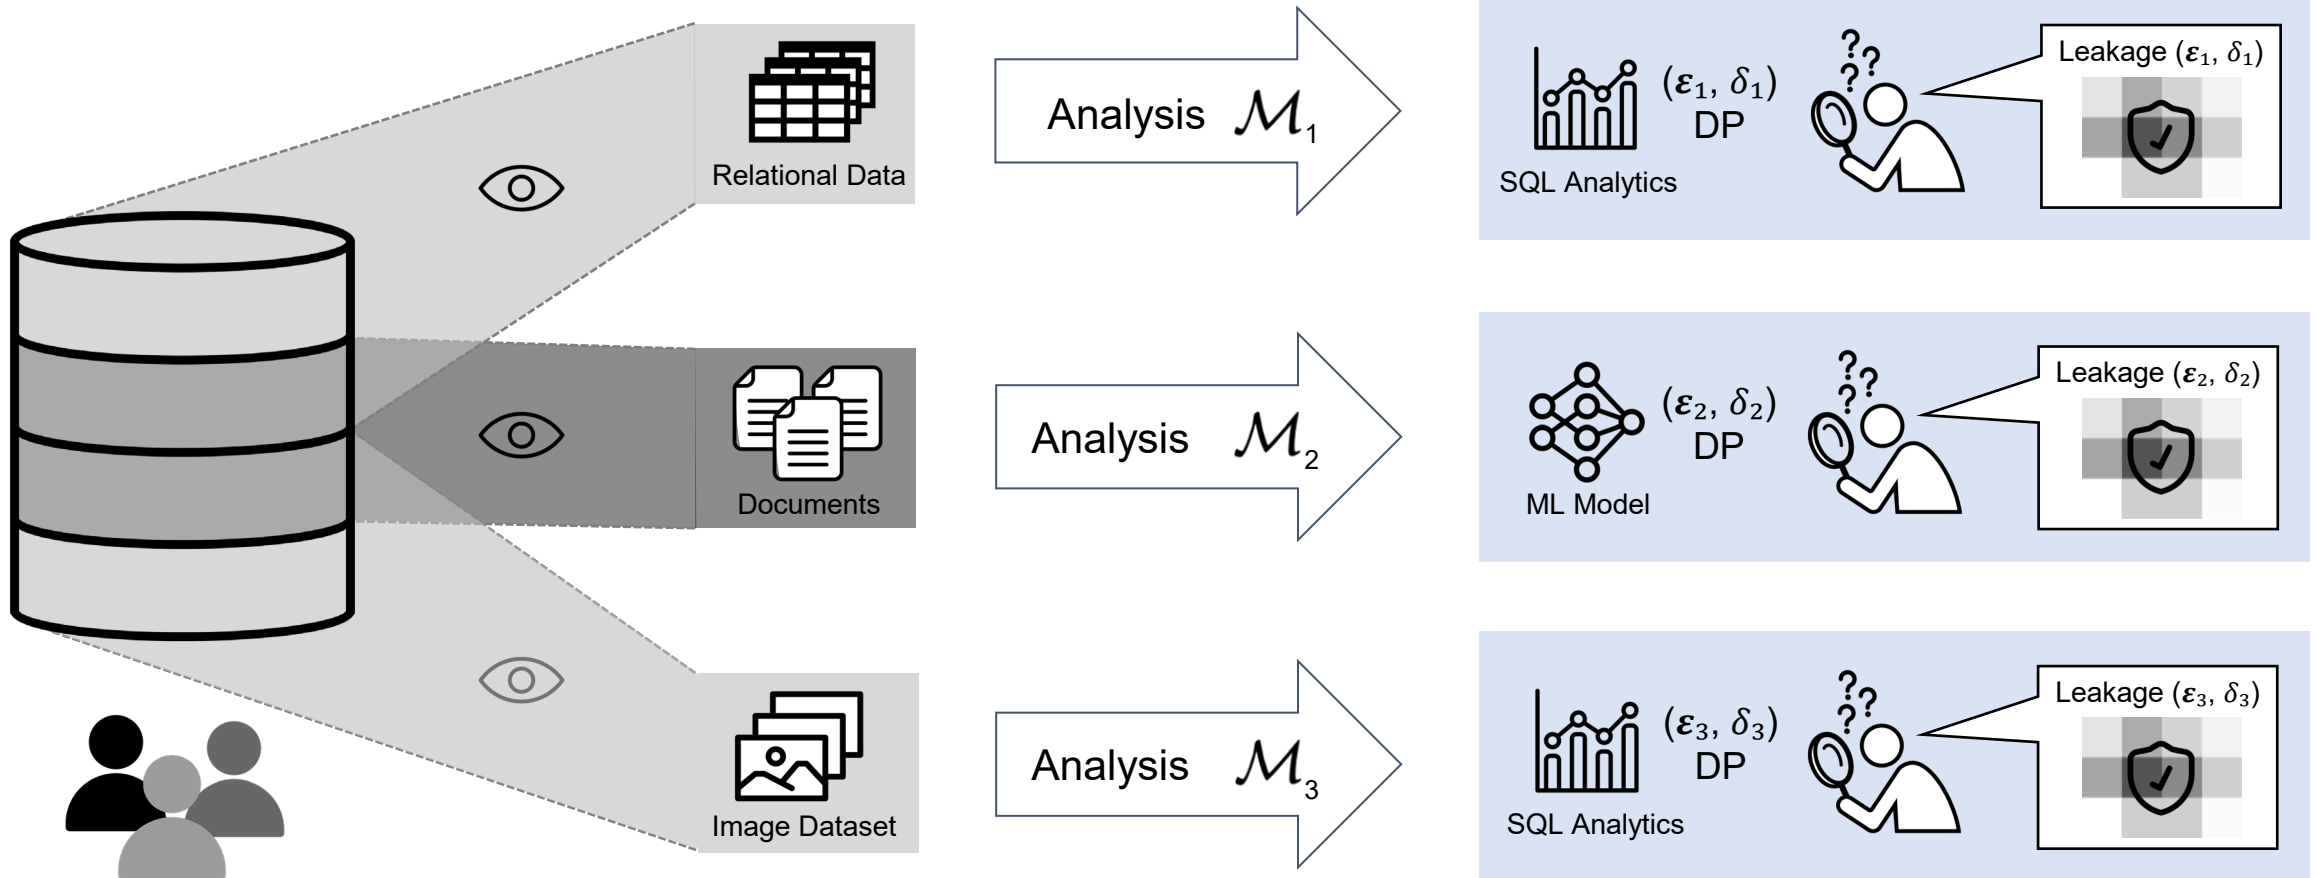
Differential Privacy



$$Pr [\mathcal{M}(D_1) \in \mathcal{S}] \leq e^\epsilon \cdot Pr [\mathcal{M}(D_2) \in \mathcal{S}] + \delta$$

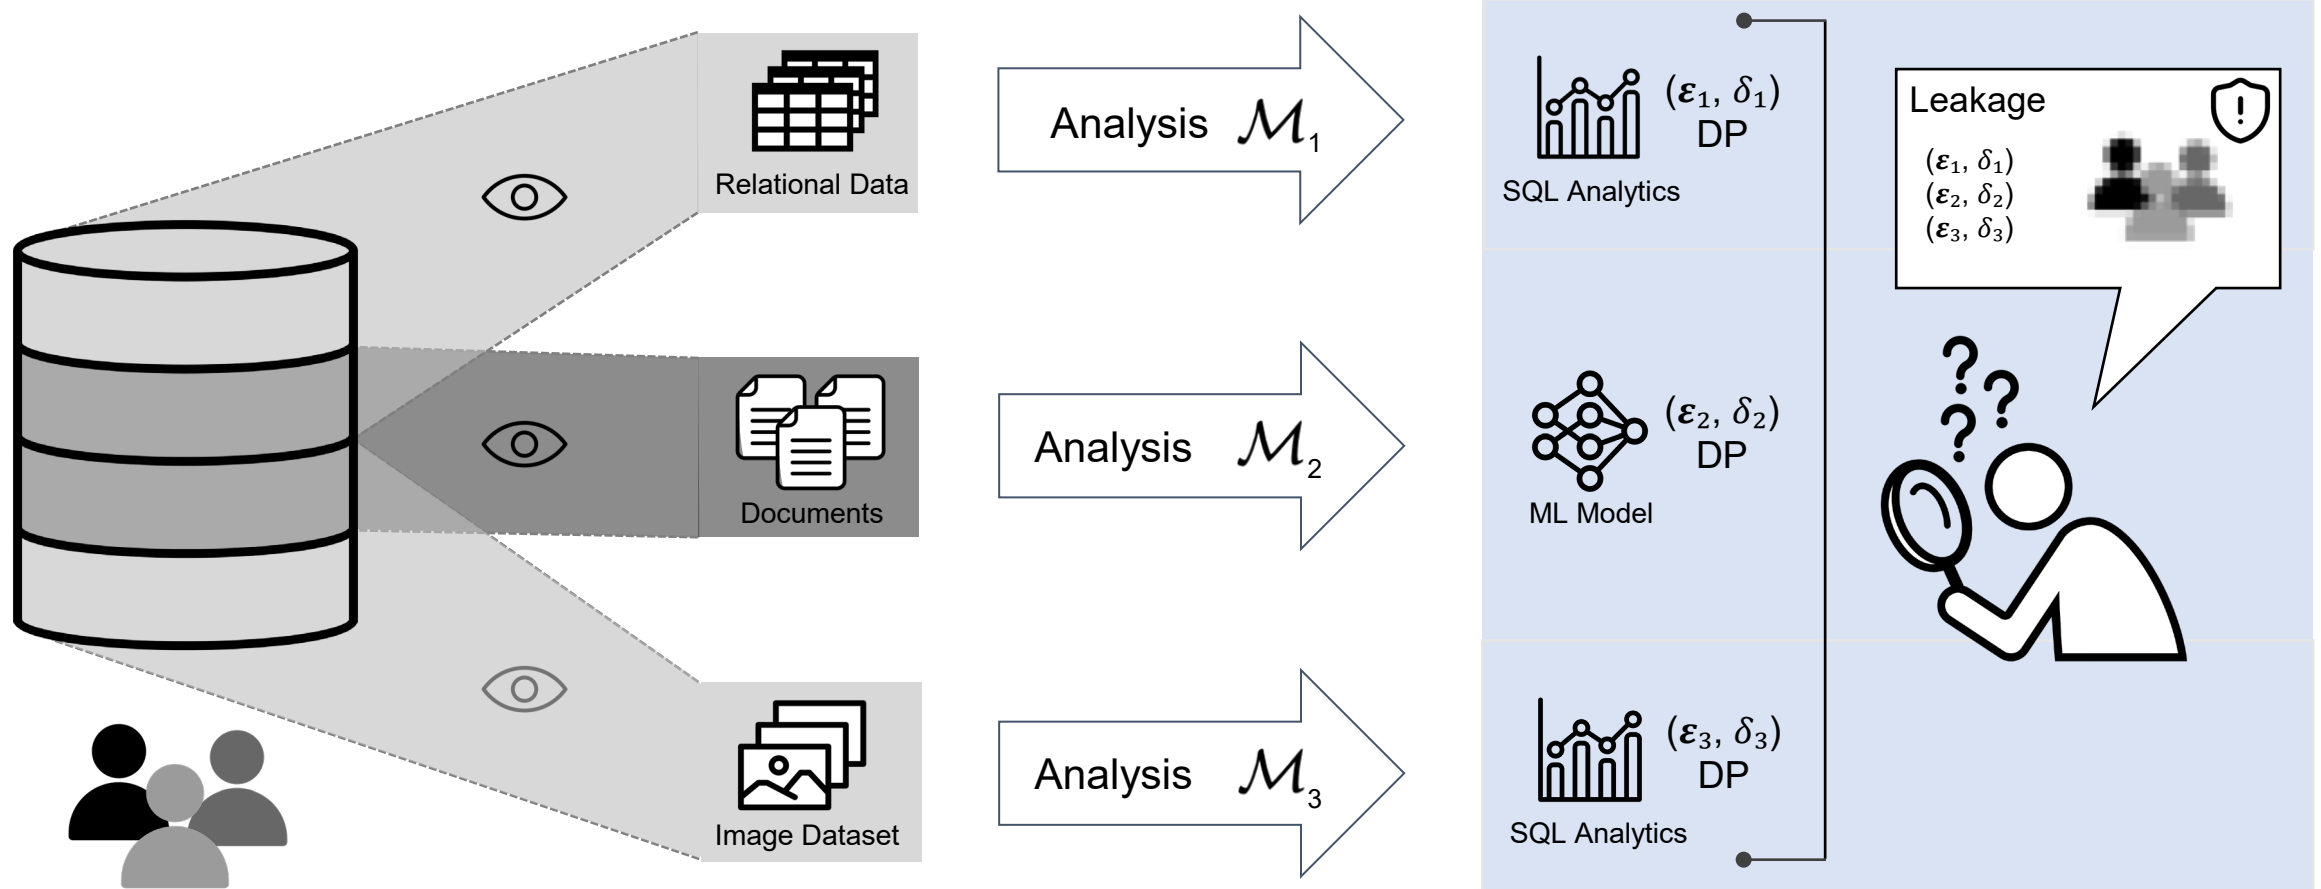
Differential Privacy

For managing risks at an organizational scale

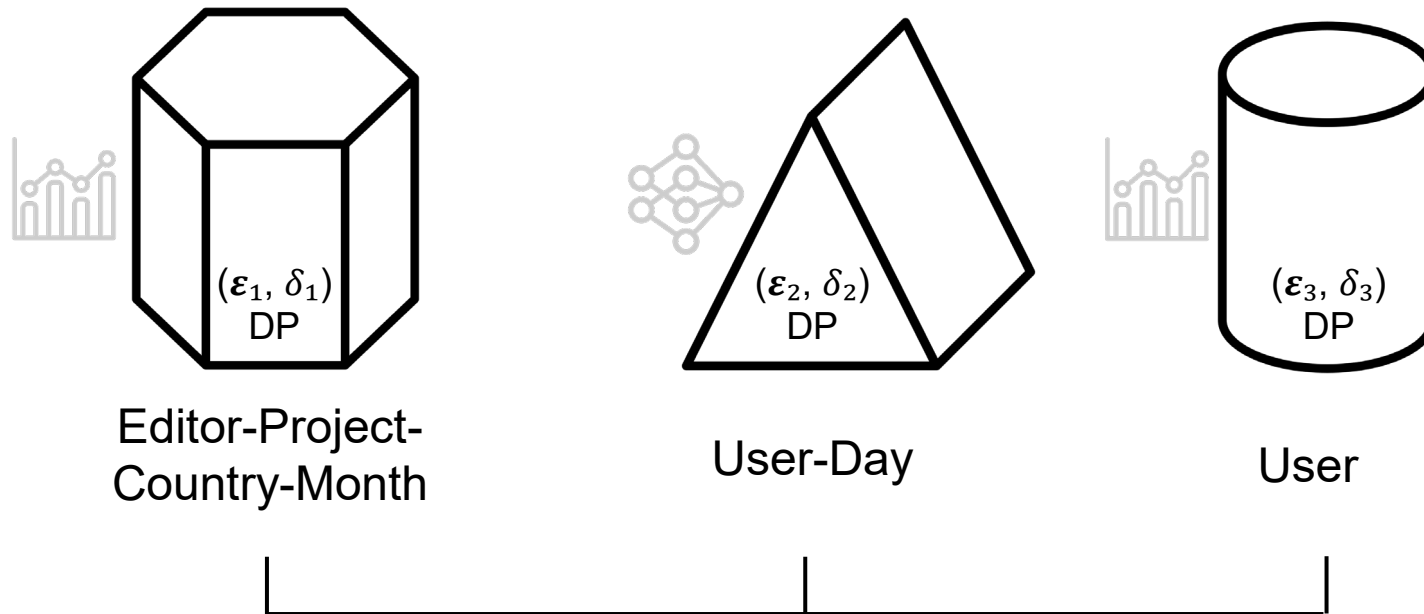


Differential Privacy

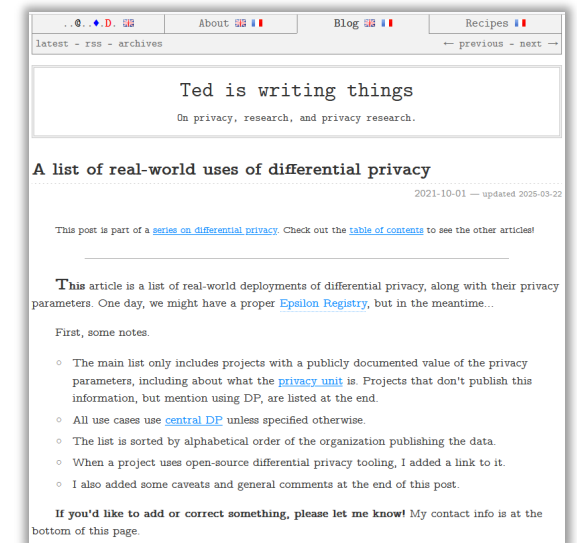
For managing risks at an organizational scale



Problem 1: Release-specific Privacy Unit




Privacy Units: Reasonable in Isolation
but **Hard to Combine**



Real-World DP Uses
[Desfontaines, 2021]

Unified Privacy Unit


Enabling composition across all releases



$$(\epsilon_1, \delta_1) + (\epsilon_2, \delta_2) + (\epsilon_3, \delta_3) + \dots \preceq (\epsilon, \delta)$$


Problem 2: Large Privacy Budget

With limited mathematical meaning




Problem 2: Large Privacy Budget

With limited mathematical meaning




Problem 3: Context-dependency




**Privacy
Acceptable?**

Single Count Query

✗
Privacy




vs.



Visualization inspired by [Nanayakkara PETS'20]

ML model (Blackbox access)


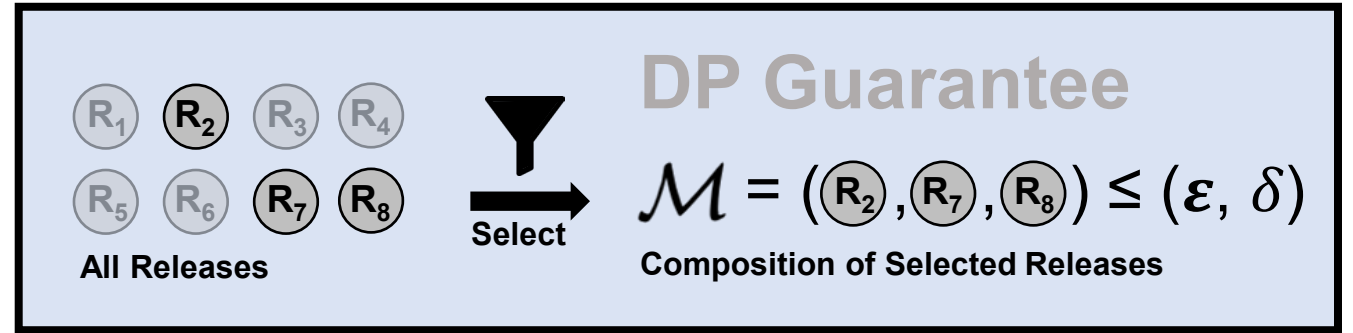
✓
Privacy




Blackbox ML: Empirical privacy auditing indicates a gap between theoretical DP bounds and actual attack success rates, potentially allowing larger ϵ .

[Nasr S&P'21, Usenix Security'23, ICLR'25]


Our Approach



Multiple Isolated
Guarantees



One Single Global
Guarantee


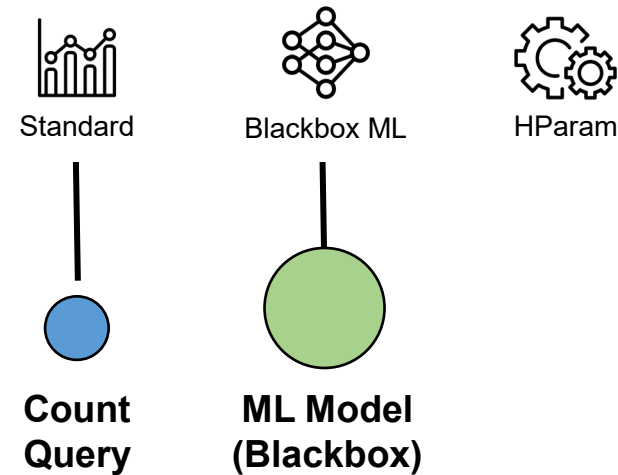


Multiple Complementing
Guarantees

Dimension 1: Multiple Contexts

Adjust the privacy budget based on the context



Example Contexts



Dimension 1: Multiple Contexts

Adjust the privacy budget based on the context

Example Contexts



Dimension 2: Multiple Scopes

Define scopes over the large collection of data

Example Data



Relational Data



Documents




Image Dataset

$$\varepsilon \leq 2$$

$$\varepsilon \leq 4$$

$$\varepsilon \leq 7$$



Per
Attribute

Attribute
Categories

All
Attributes

Dimension 3: Multiple Privacy Units

Support for multiple privacy units

Example Privacy Units



User-Month




User-Year




User



Household




Multiple ComPLEMENTING Guarantees




Multiple Contexts

Dimension 1



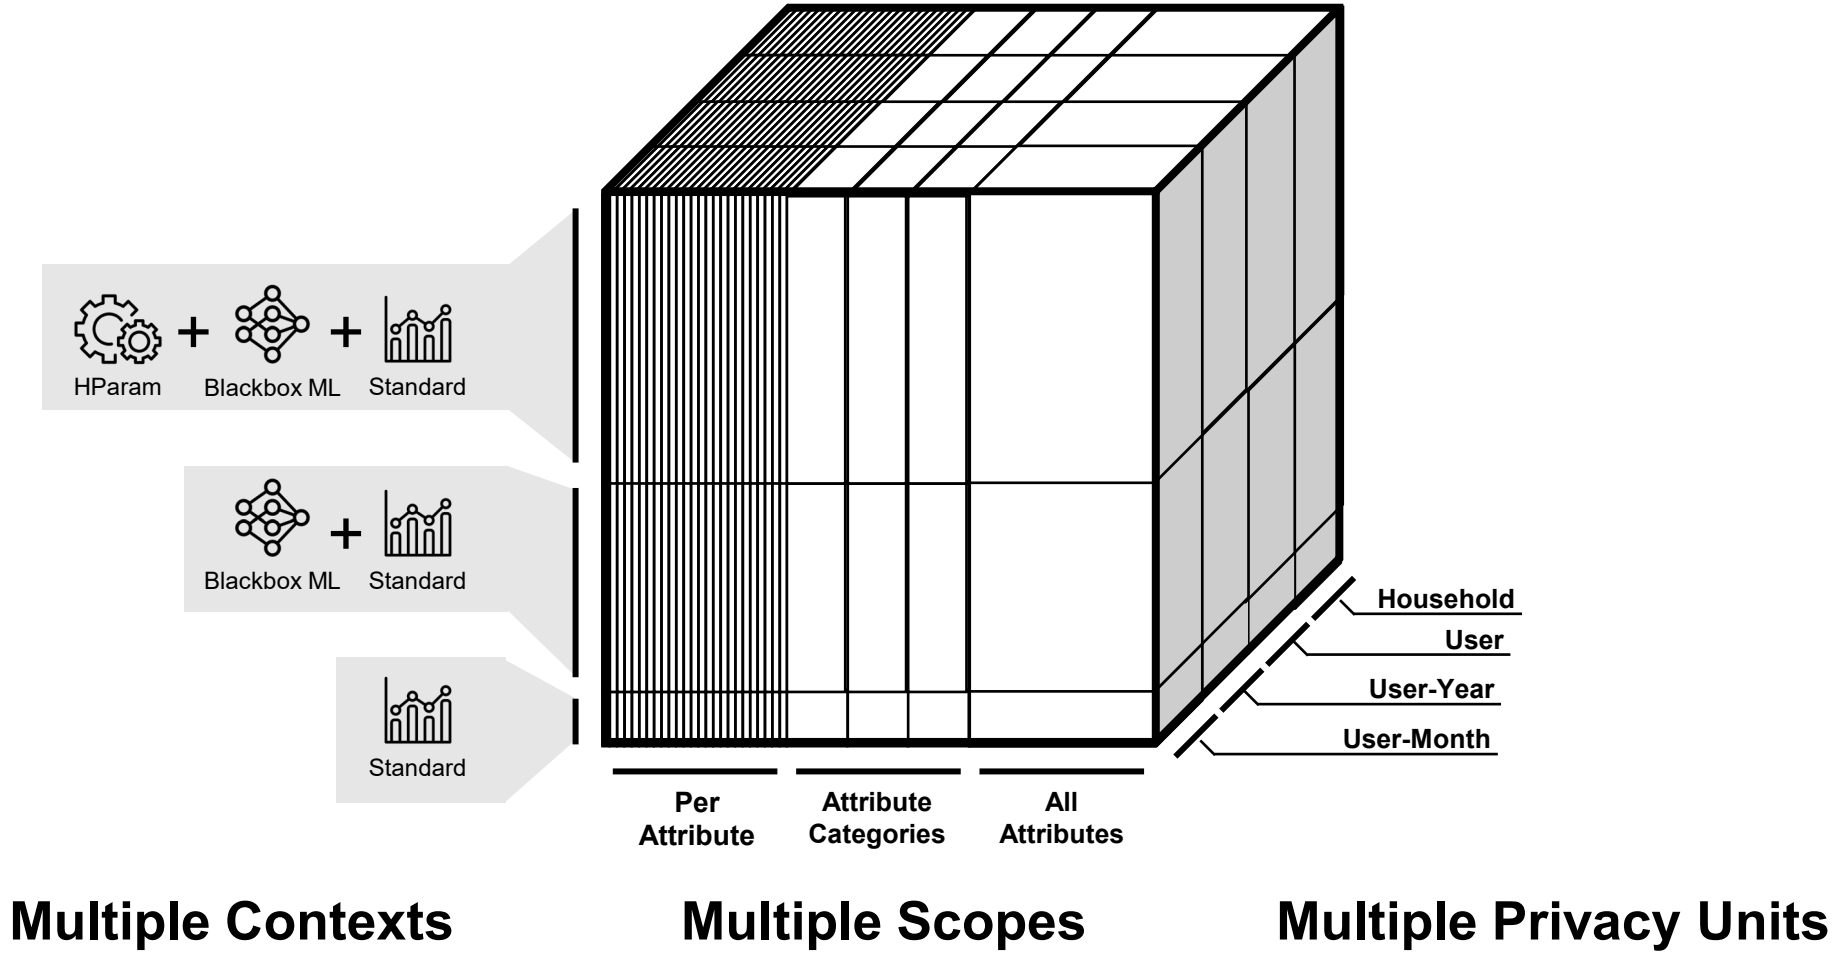
Multiple Scopes

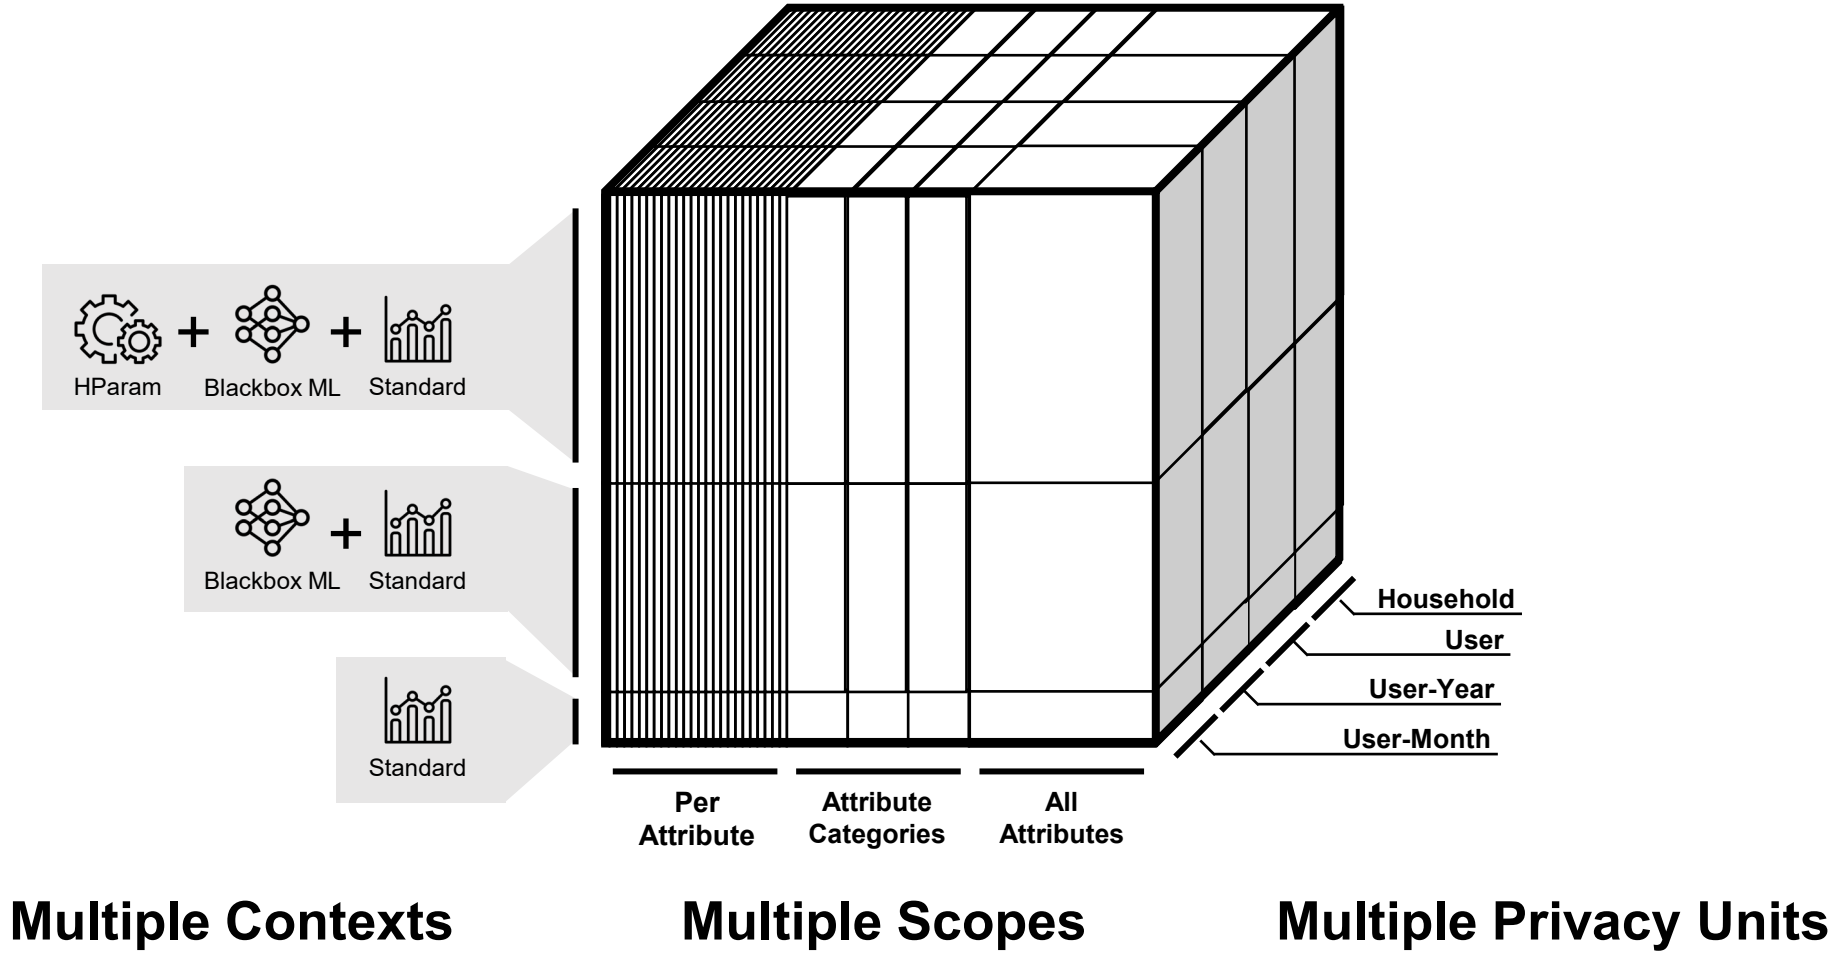
Dimension 2



Multiple Privacy Units

Dimension 3








DP Policy Language

Deriving the large rule set by a concise policy language





Policy Set



Problem: Large Rule Set

Policy Enforcement

How to find the relevant rules? Do we need to check all?




Large Rule Set

Policy Enforcement

How to find the relevant rules? Do we need to check all?


Release Candidate



ML Model

Context | Scope | (ϵ_1, δ_1)


Rule Poset



Policy Enforcement

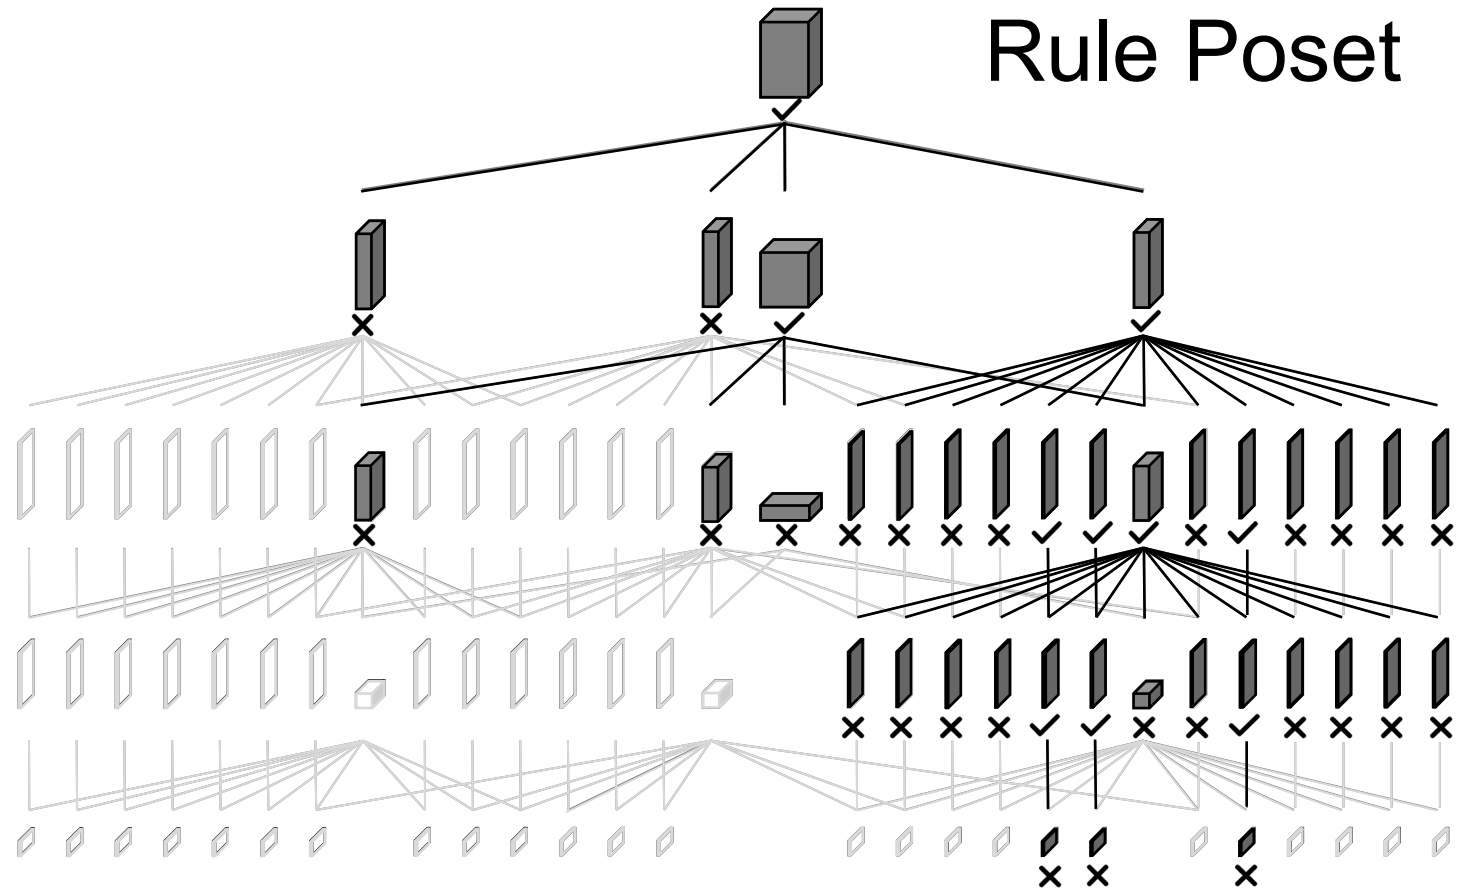
How to find the relevant rules? Do we need to check all?

Release Candidate




ML Model

Context | Scope | (ϵ_1, δ_1)





Rule Set Optimization

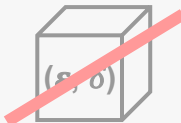

Are all the rules required?

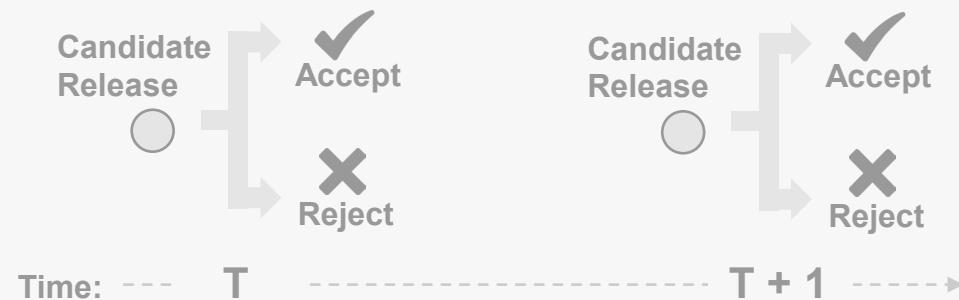


System Integration


DP Libraries

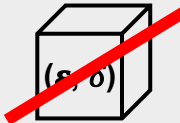

Examples:  Tumult Analytics  OpenDP

1. Set Privacy Budget  
2. Check the remaining budget to accept / reject the release candidate.



DP Management Systems

Examples: PrivateKube  Cohere
[Luo et al. OSDI'21] [Küchler et al. S&P'24]


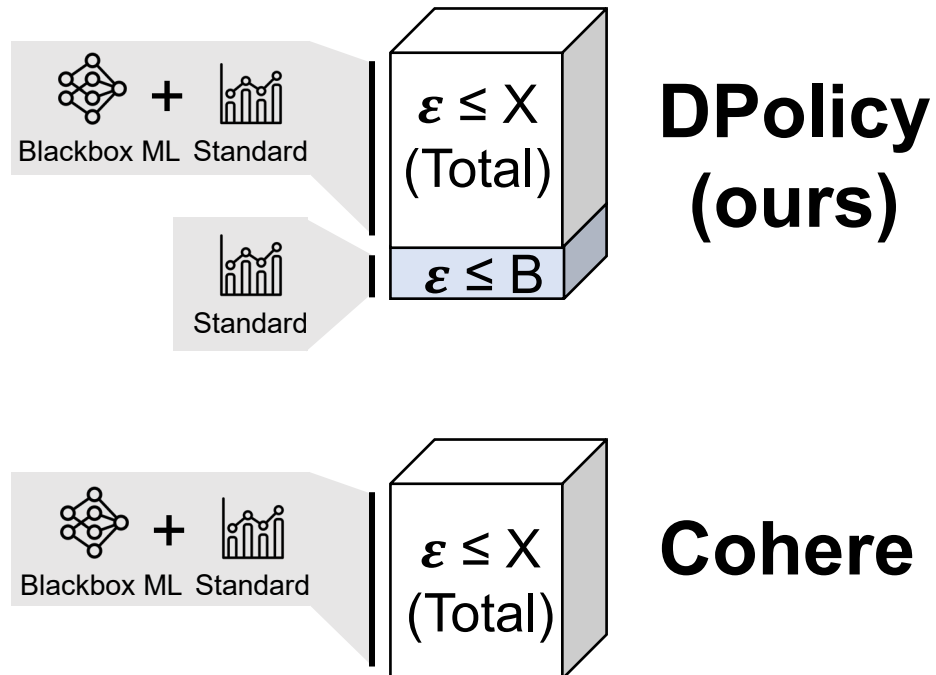
1. Set Privacy Budget  
2. Find the best allocation of candidate releases subject to the available privacy budget.




Evaluation

Cohere Workload

[Küchler et al. S&P'24]





DPOLICY: Managing Privacy Risks Across Multiple Releases with DP