Cohere: Managing Differential Privacy in Large - Scale Systems

Nicolas Küchler
ETH zürich

Emanuel Opel
ETH zürich

Hidde Lycklama
ETH zürich

Alexander Viand
intel

Anwar Hithnawi
ETH zürich

Presented At
IEEE Security & Privacy 2024

Contact Us
www.pps-lab.com

Statistical Release

How can we release useful information without compromising privacy?

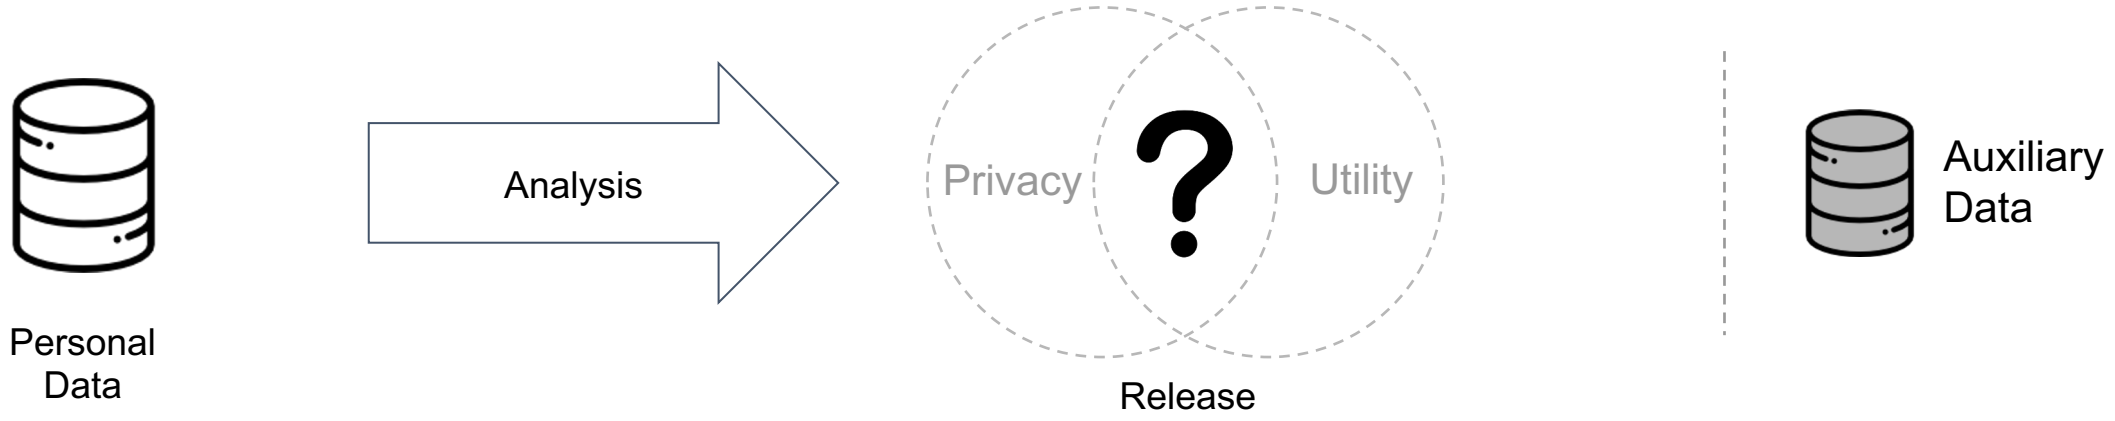

Statistical Release

How can we release useful information without compromising privacy?

Personal Data

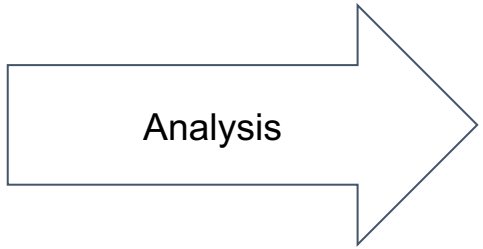

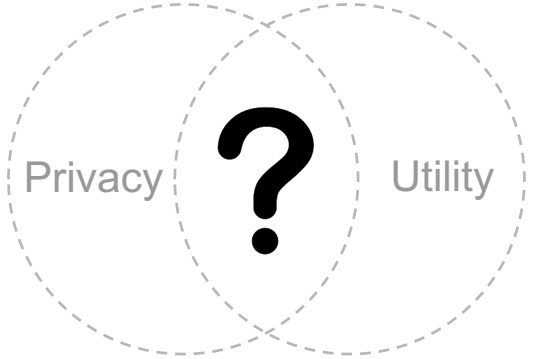

Release

Auxiliary Data

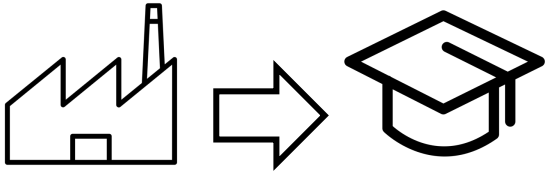

Statistical Release

How can we release useful information without compromising privacy?

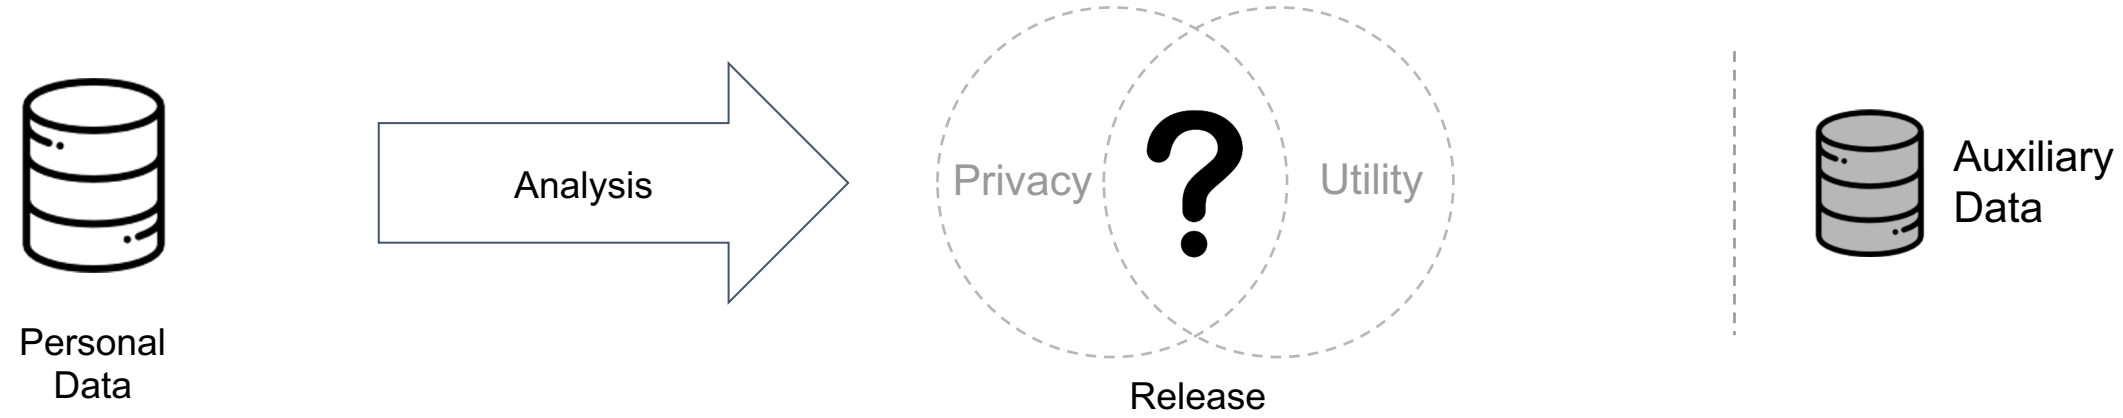

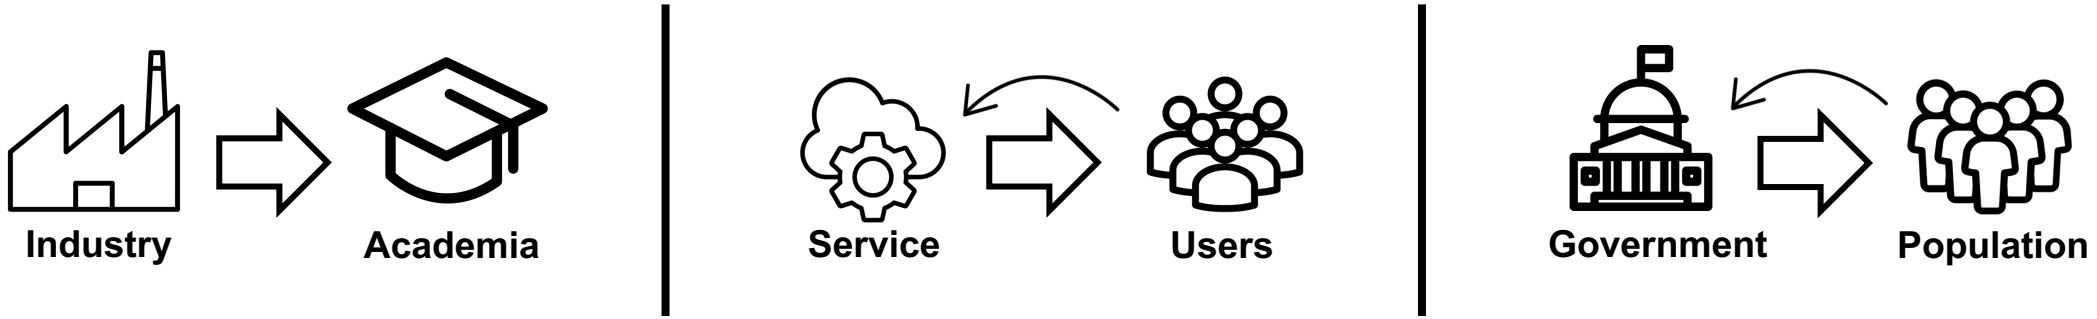

Statistical Release

How can we release useful information without compromising privacy?

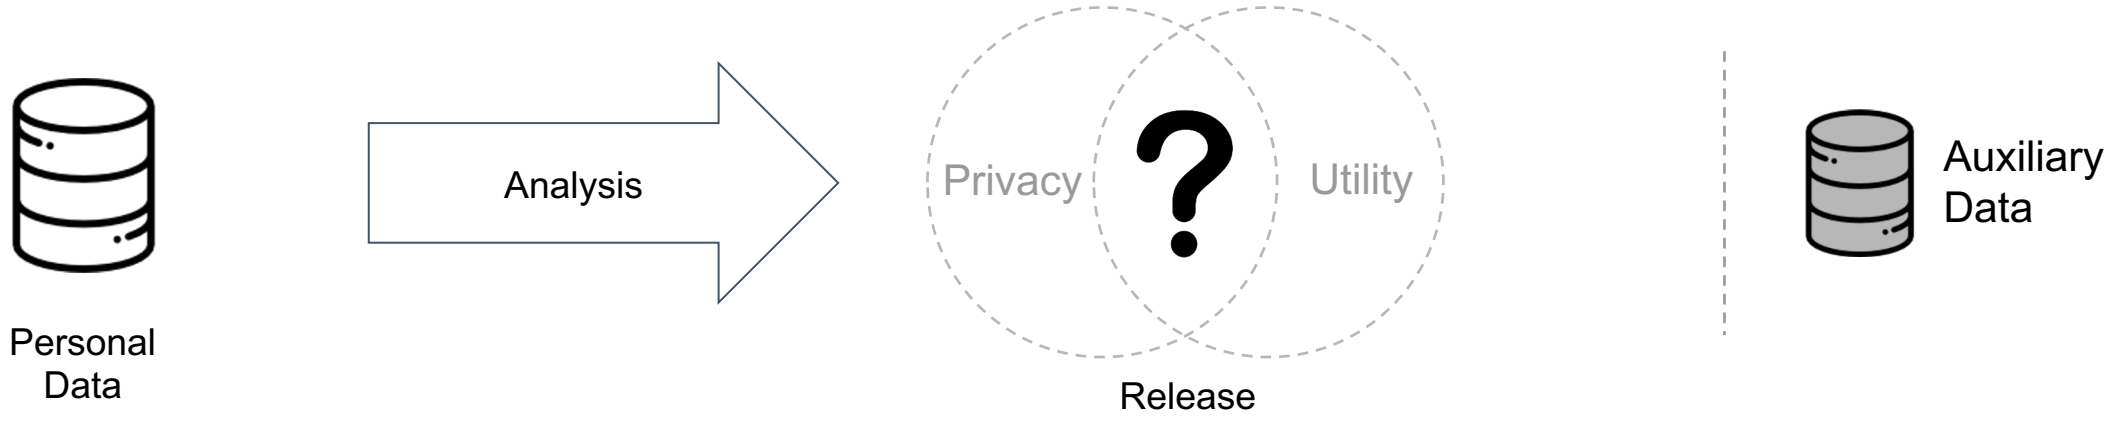

Statistical Release

How can we release useful information without compromising privacy?

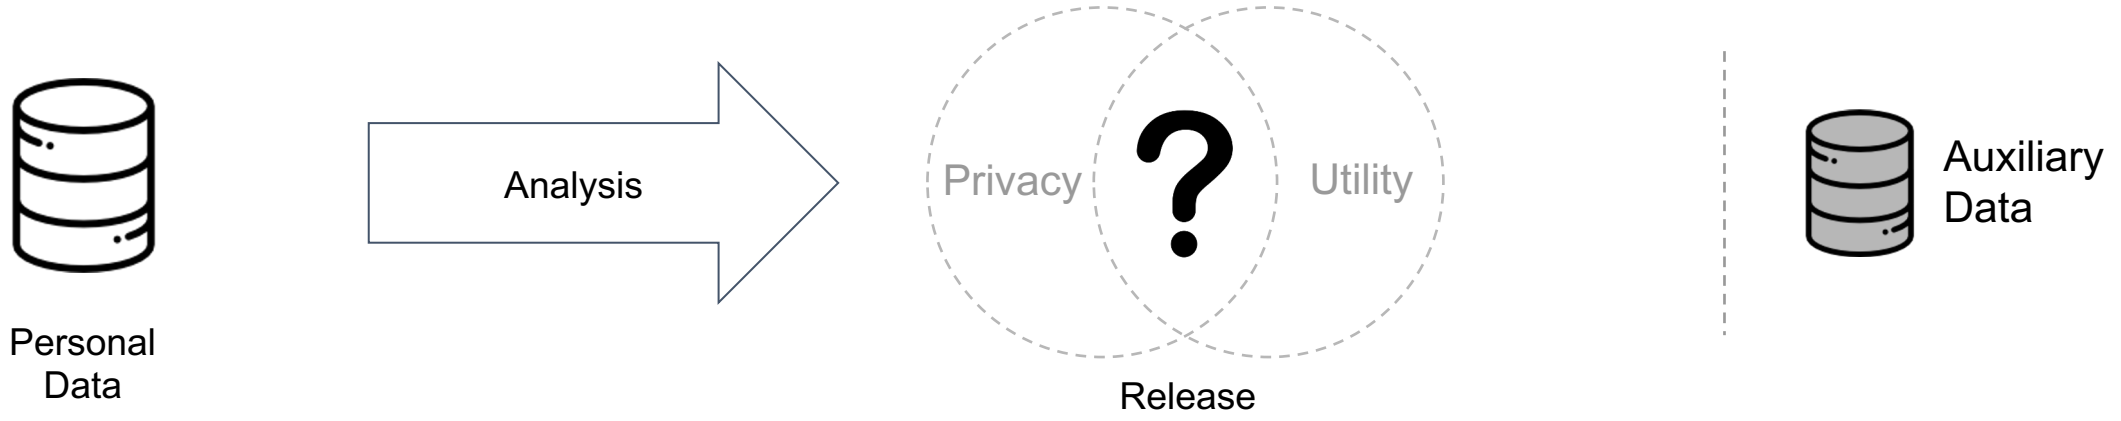

- Anonymization
  Redact Personally Identifiable Information

Name	Region	...	Value
[redacted]	CH		100
[redacted]	DE		237

Statistical Release

How can we release useful information without compromising privacy?

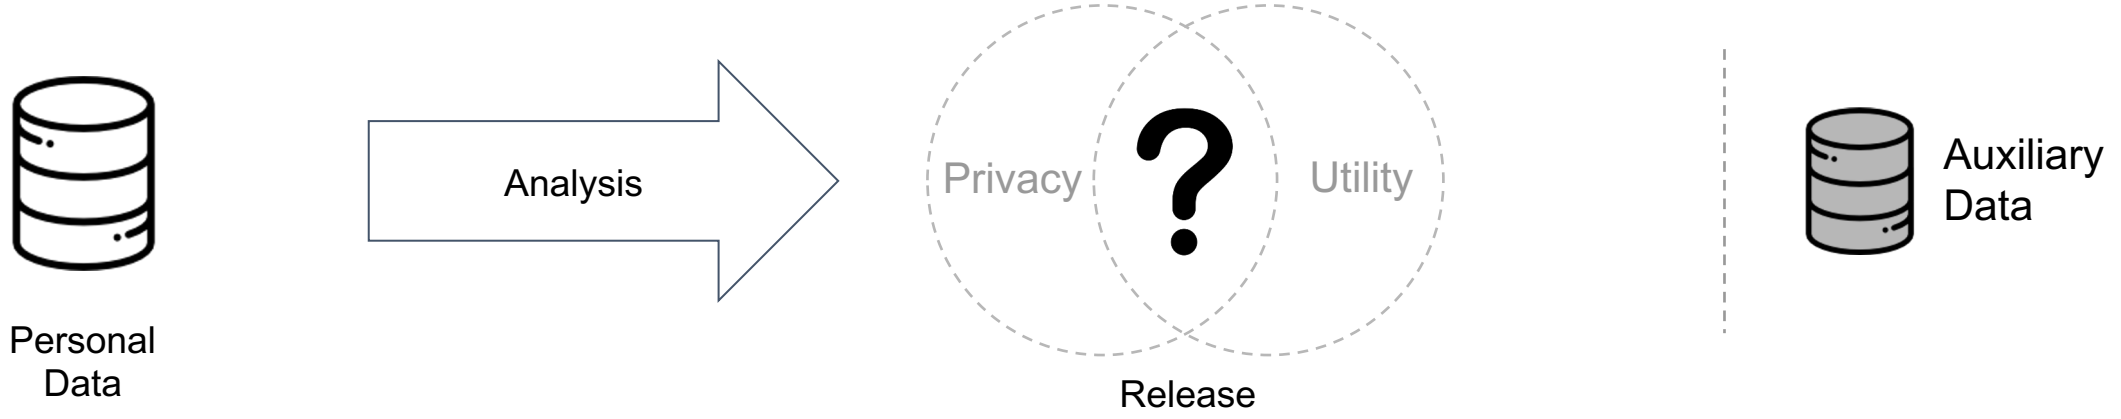

- Anonymization
  Redact Personally Identifiable Information

Name	Region	...	Value
[redacted]	CH		100
[redacted]	DE		237

- Release Aggregates

Descriptive Statistics / ML Model

Statistical Release

How can we release useful information without compromising privacy?

Personal Data

Analysis

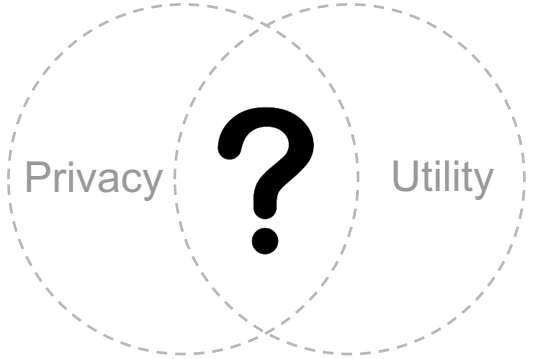

Release

Auxiliary Data

- Anonymization
  Redact Personally Identifiable Information

- Release Aggregates

Name	Region	...	Value
[redacted]	CH		100
[redacted]	DE		237

Descriptive Statistics / ML Model

Privacy Attacks

- Re-Identification (NYC Taxi)
- Database Reconstruction (United States Census 2010)
- Membership Inference (LLM)

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

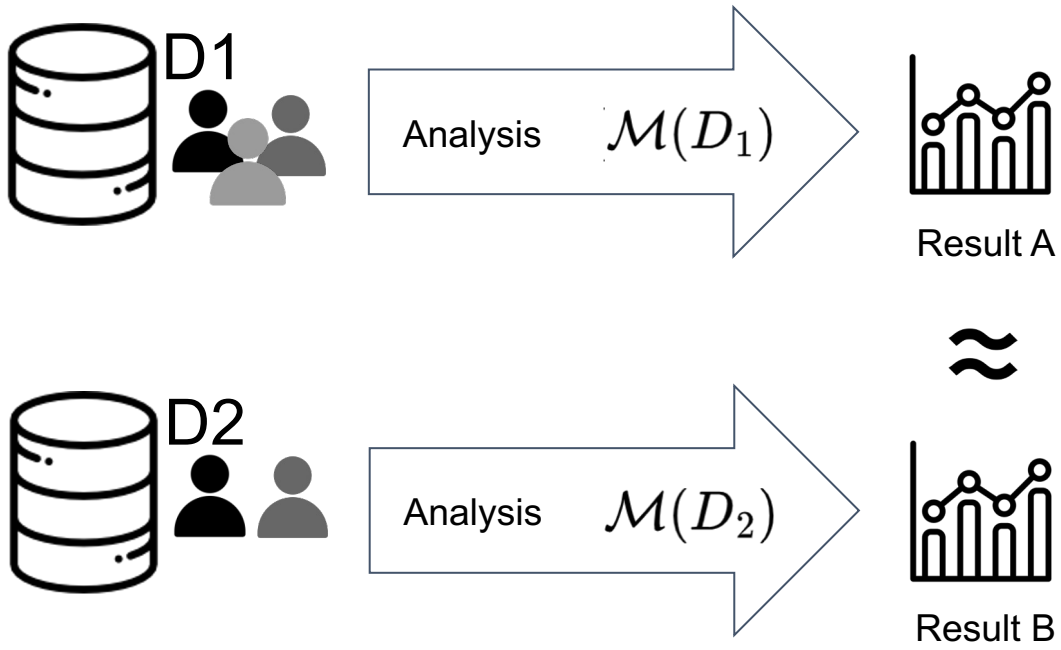

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

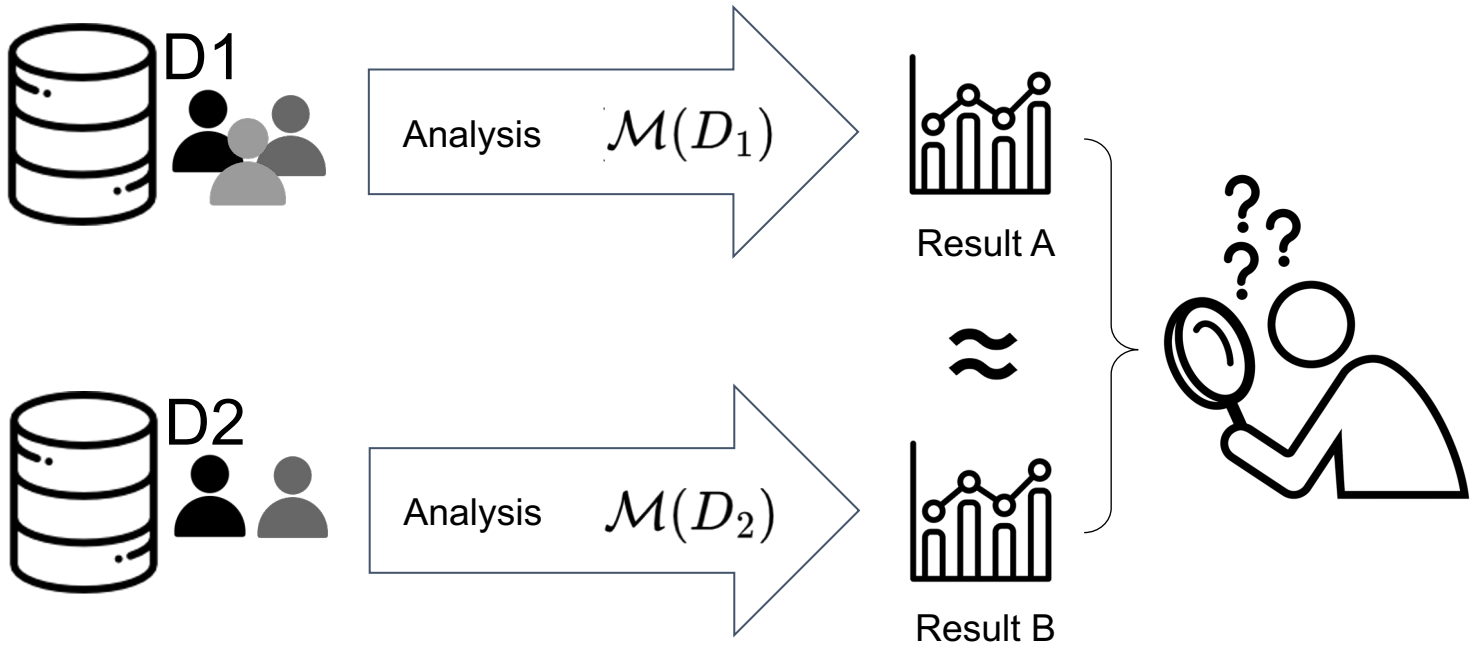

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

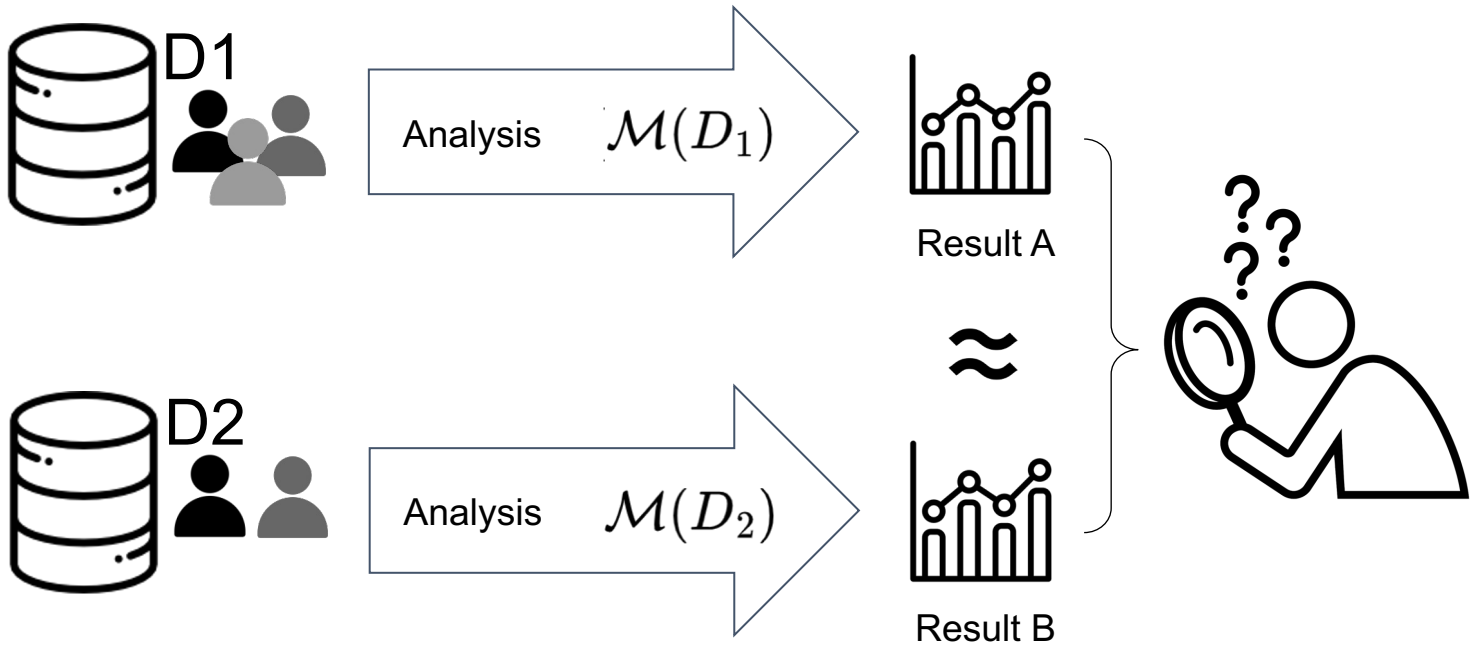

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

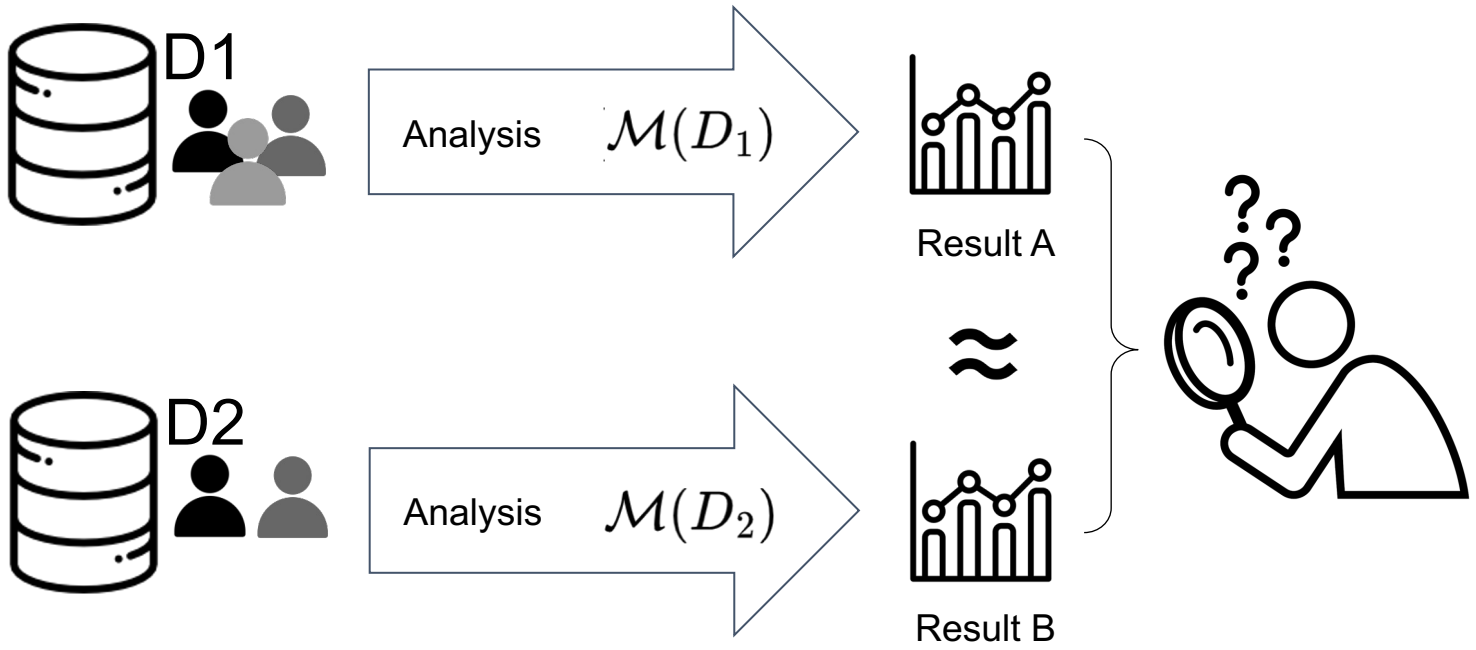

$$Pr\left[\mathcal{M}(D_1) \in \mathcal{S}\right] \leq e^{\epsilon} \cdot Pr\left[\mathcal{M}(D_2) \in \mathcal{S}\right] + \delta$$

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

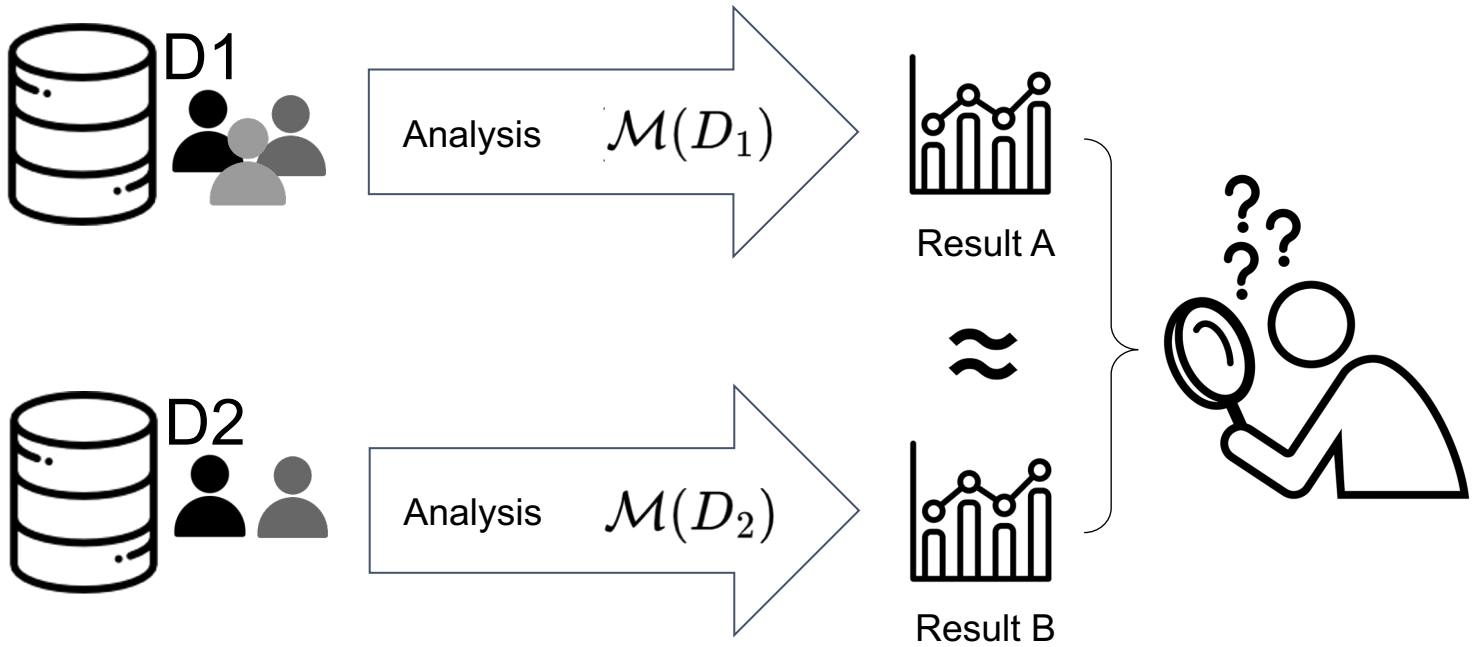

Intuition

Data **+**

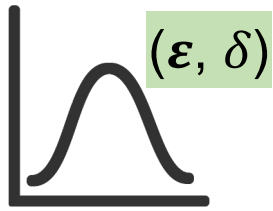

$$Pr\left[\mathcal{M}(D_1) \in \mathcal{S}\right] \leq e^{\epsilon} \cdot Pr\left[\mathcal{M}(D_2) \in \mathcal{S}\right] + \delta$$

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

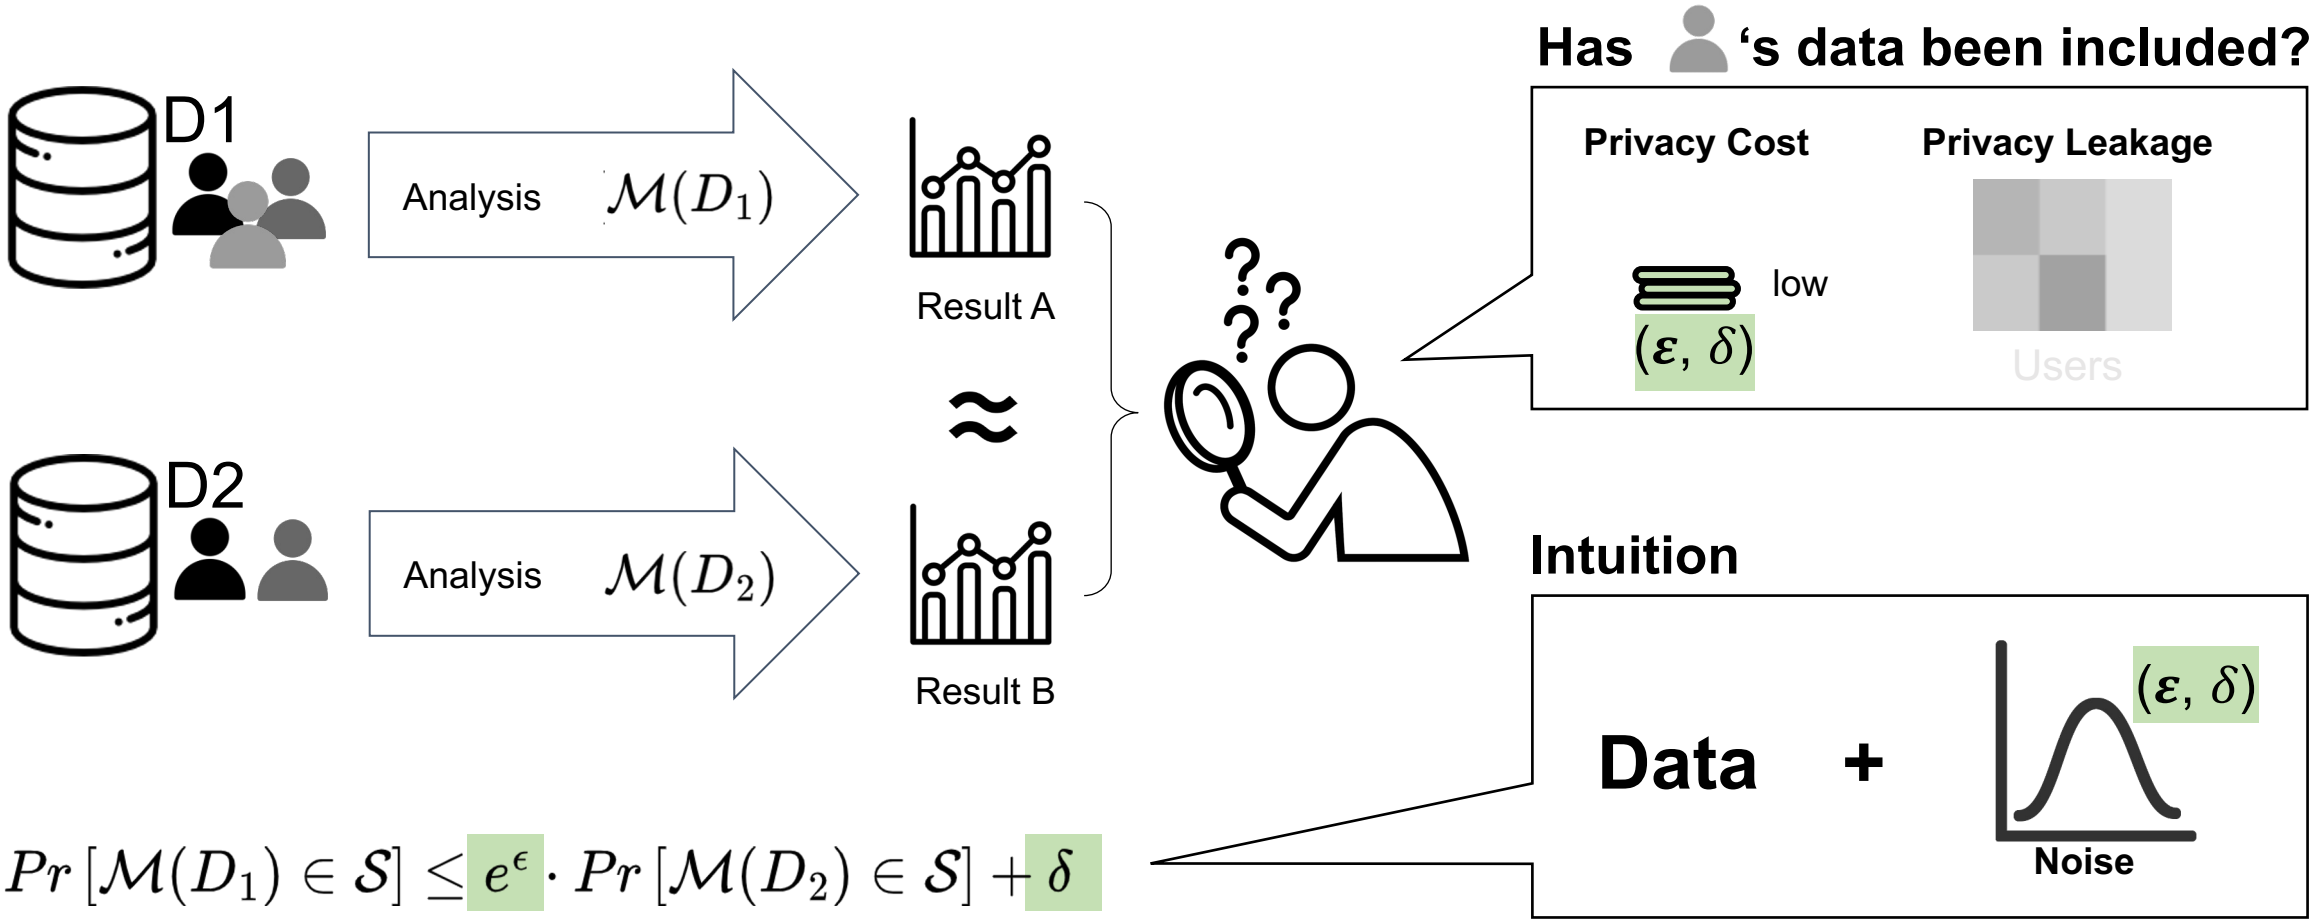

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

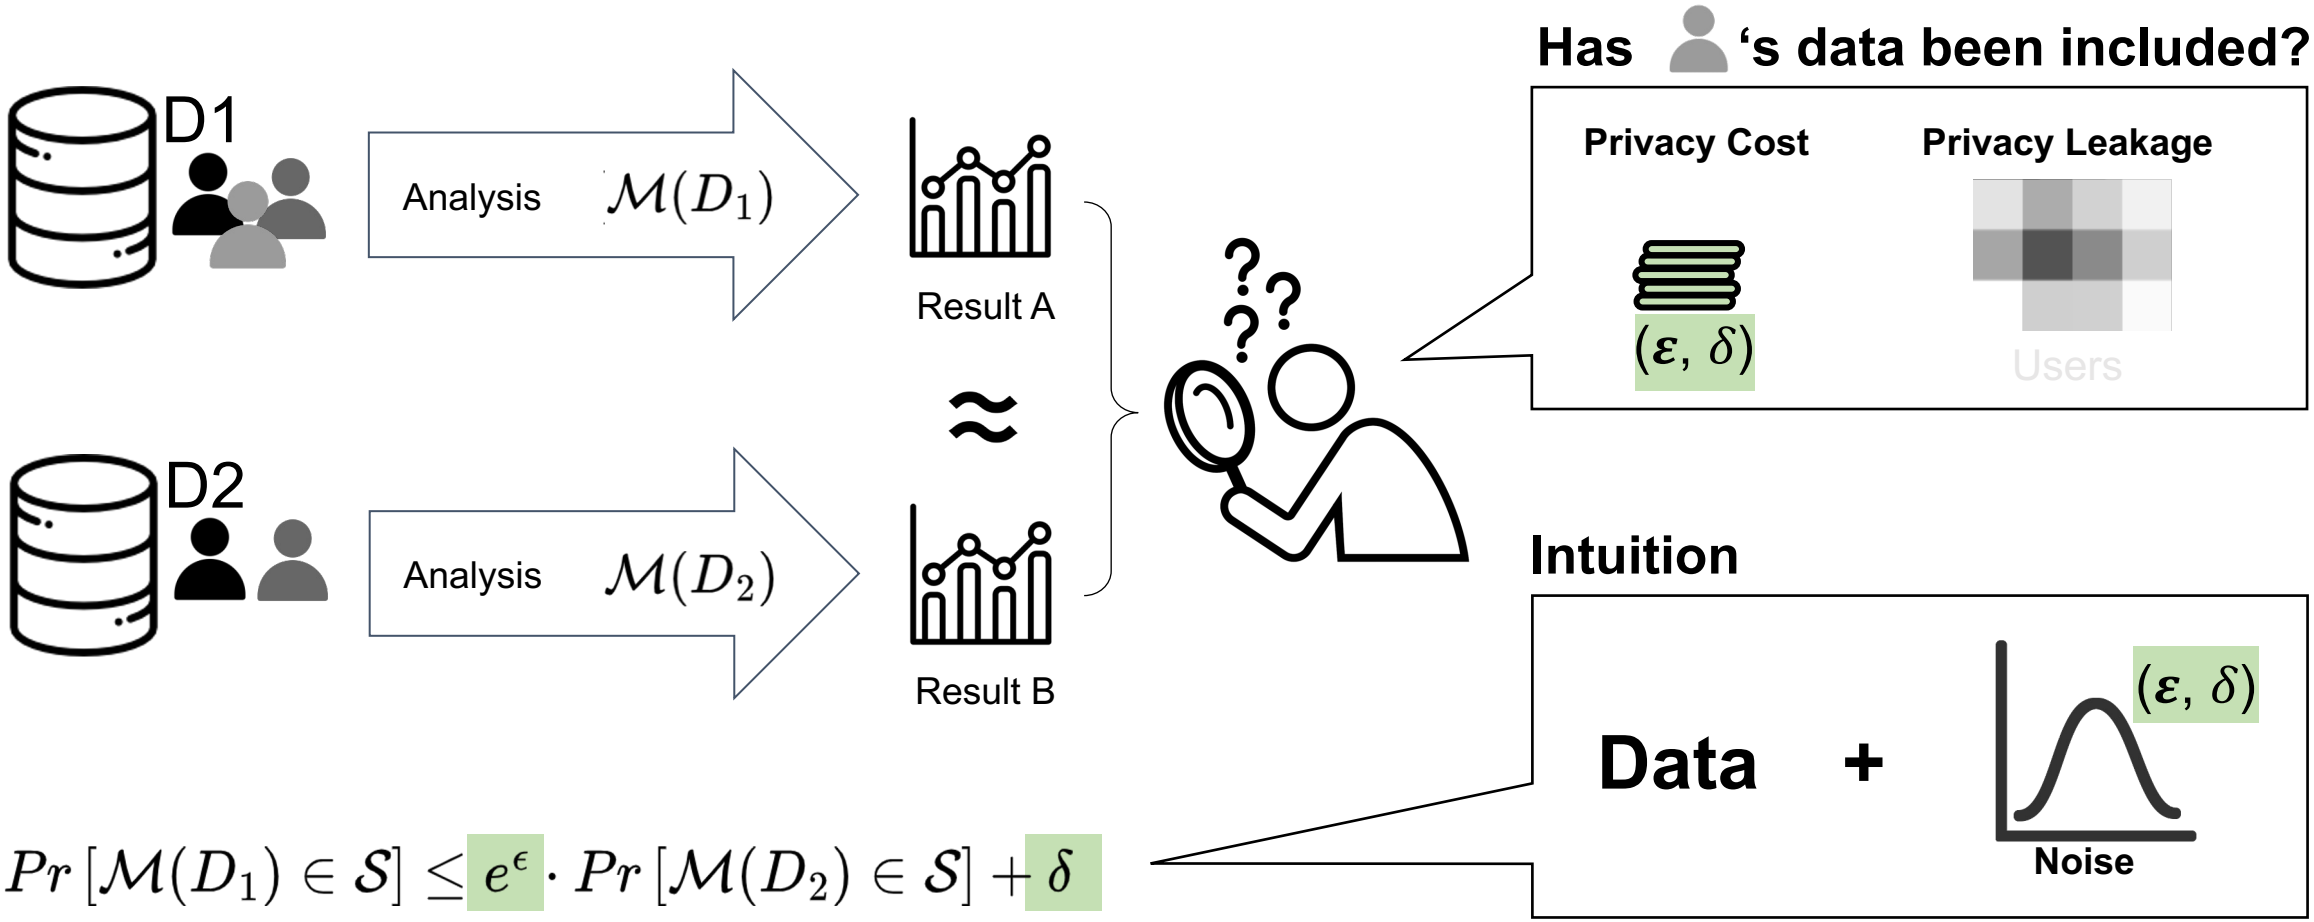

$$Pr\left[\mathcal{M}(D_1) \in \mathcal{S}\right] \leq e^{\epsilon} \cdot Pr\left[\mathcal{M}(D_2) \in \mathcal{S}\right] + \delta$$

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

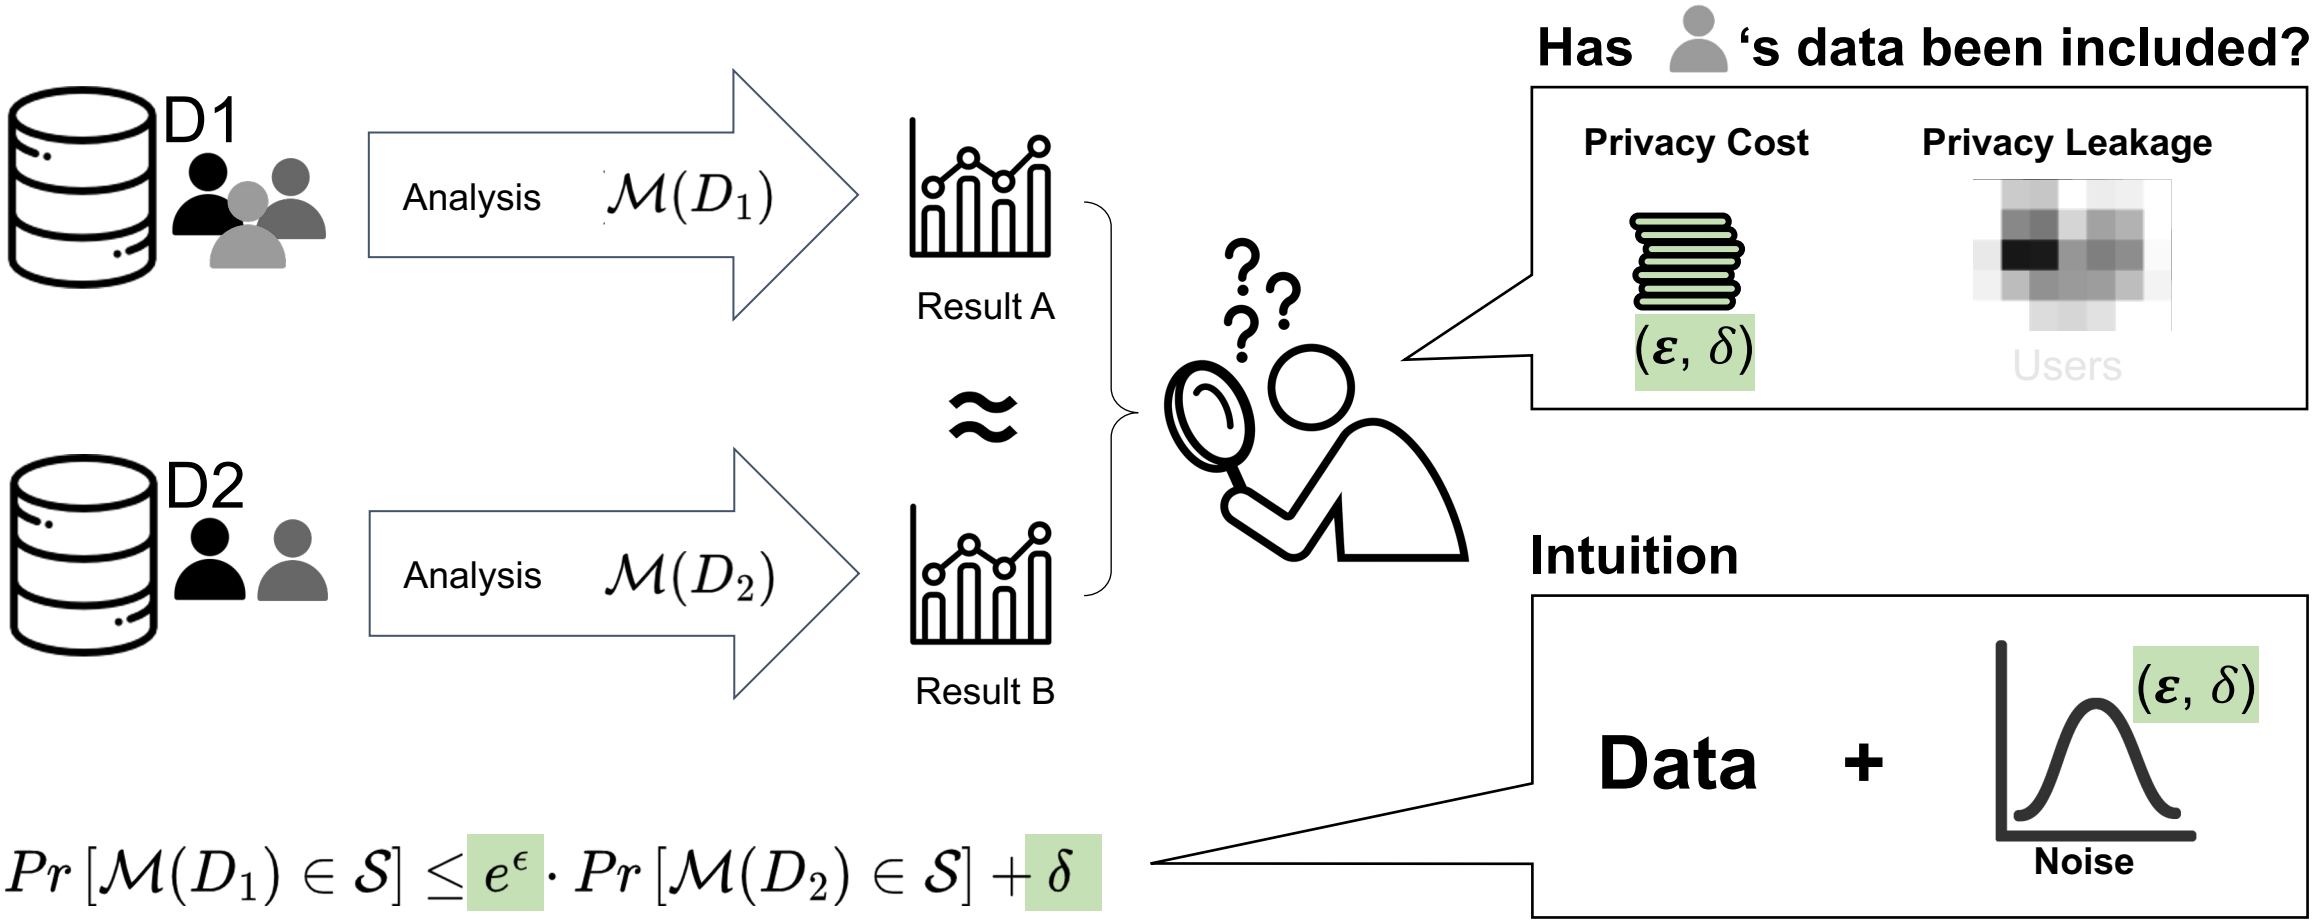

$$Pr\left[\mathcal{M}(D_1) \in \mathcal{S}\right] \leq e^{\epsilon} \cdot Pr\left[\mathcal{M}(D_2) \in \mathcal{S}\right] + \delta$$

Differential Privacy

Mathematical definition of privacy in the context of statistical releases

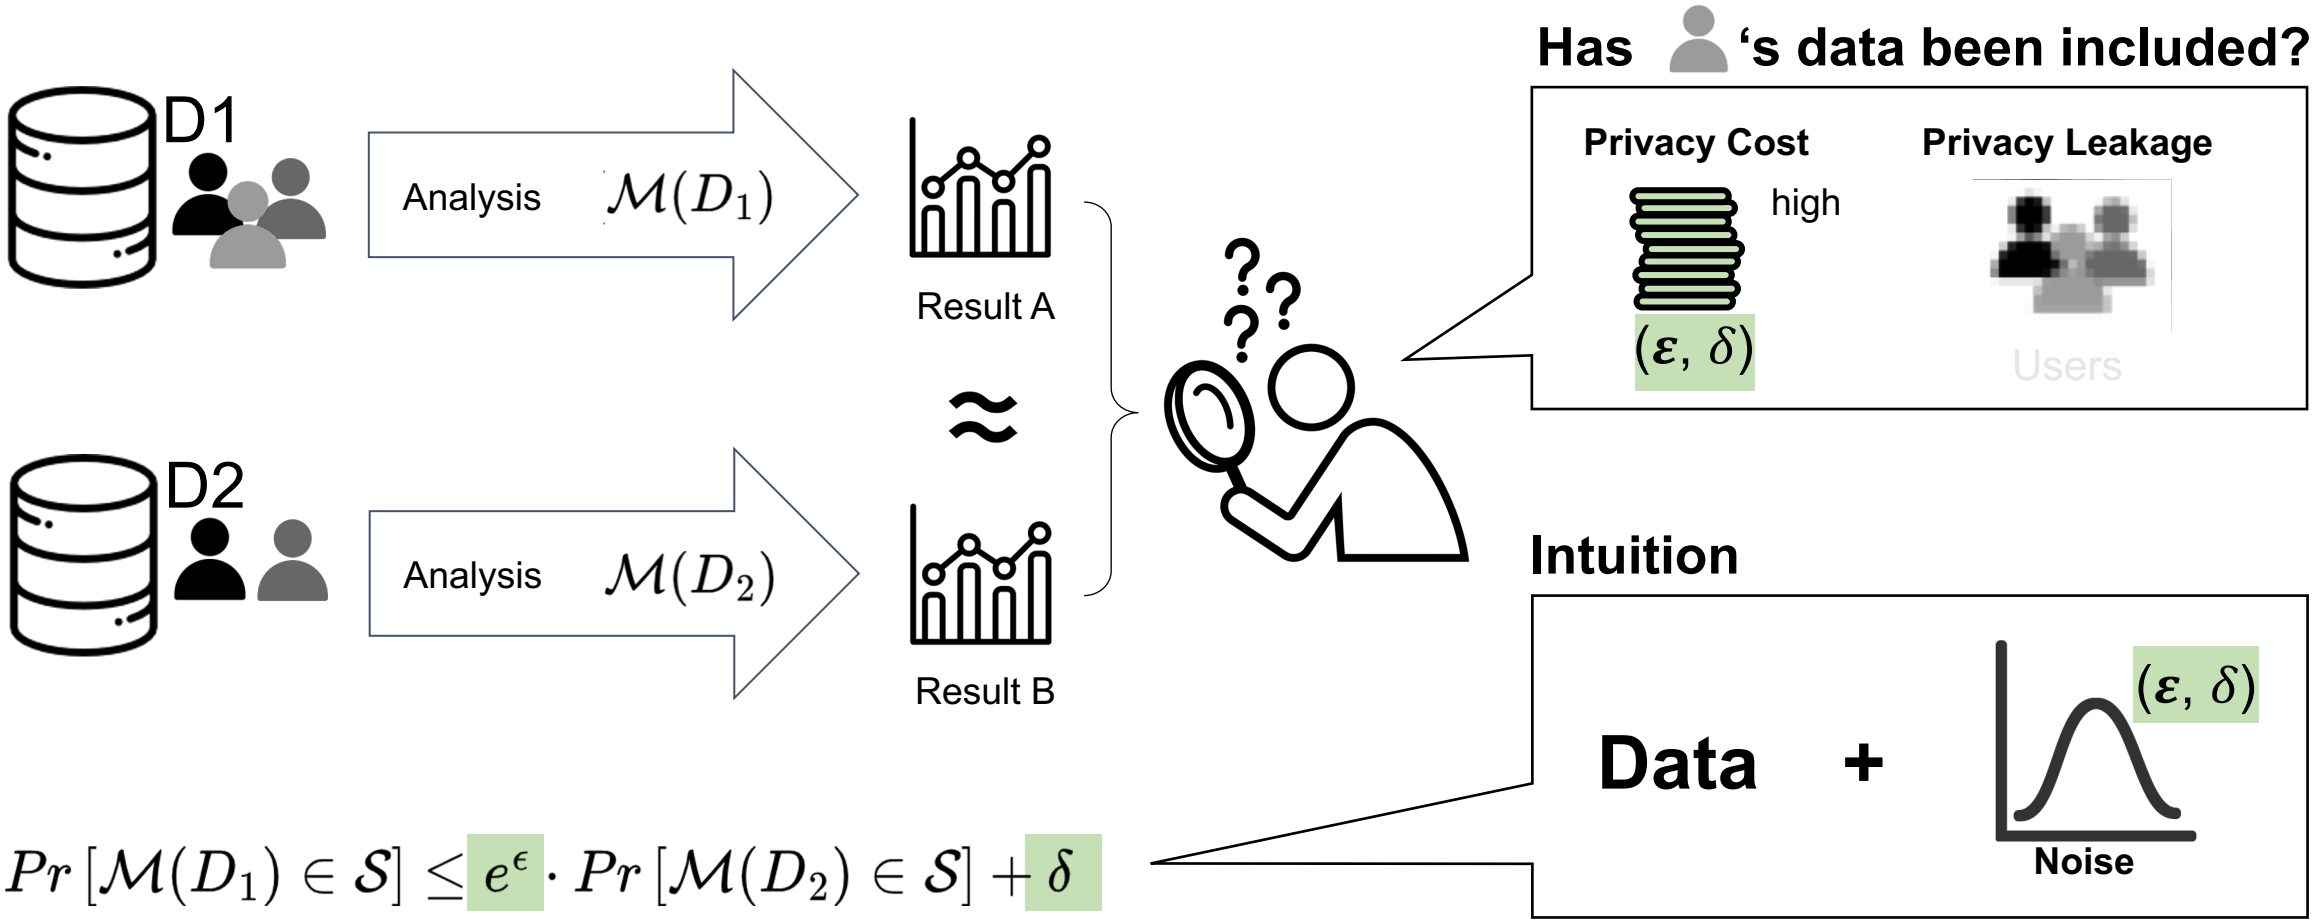

$$Pr\left[\mathcal{M}(D_1) \in \mathcal{S}\right] \leq e^{\epsilon} \cdot Pr\left[\mathcal{M}(D_2) \in \mathcal{S}\right] + \delta$$

From Theory to Practice

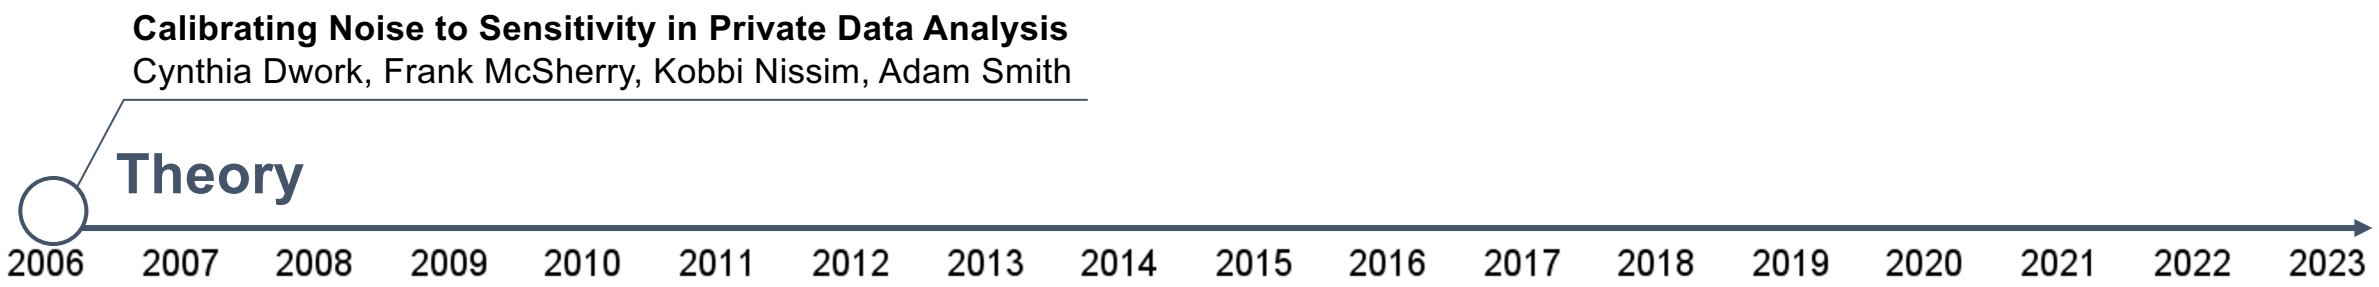

From Theory to Practice

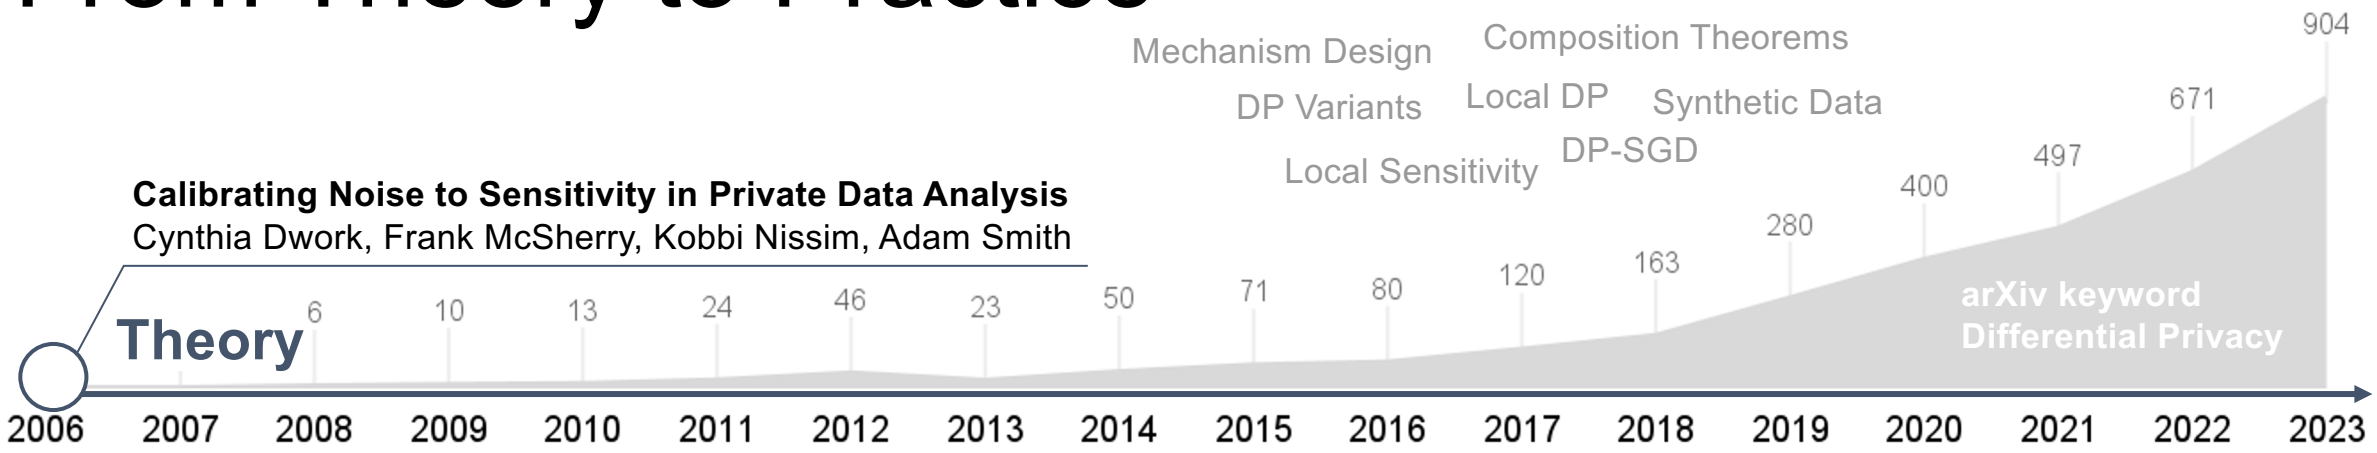

From Theory to Practice

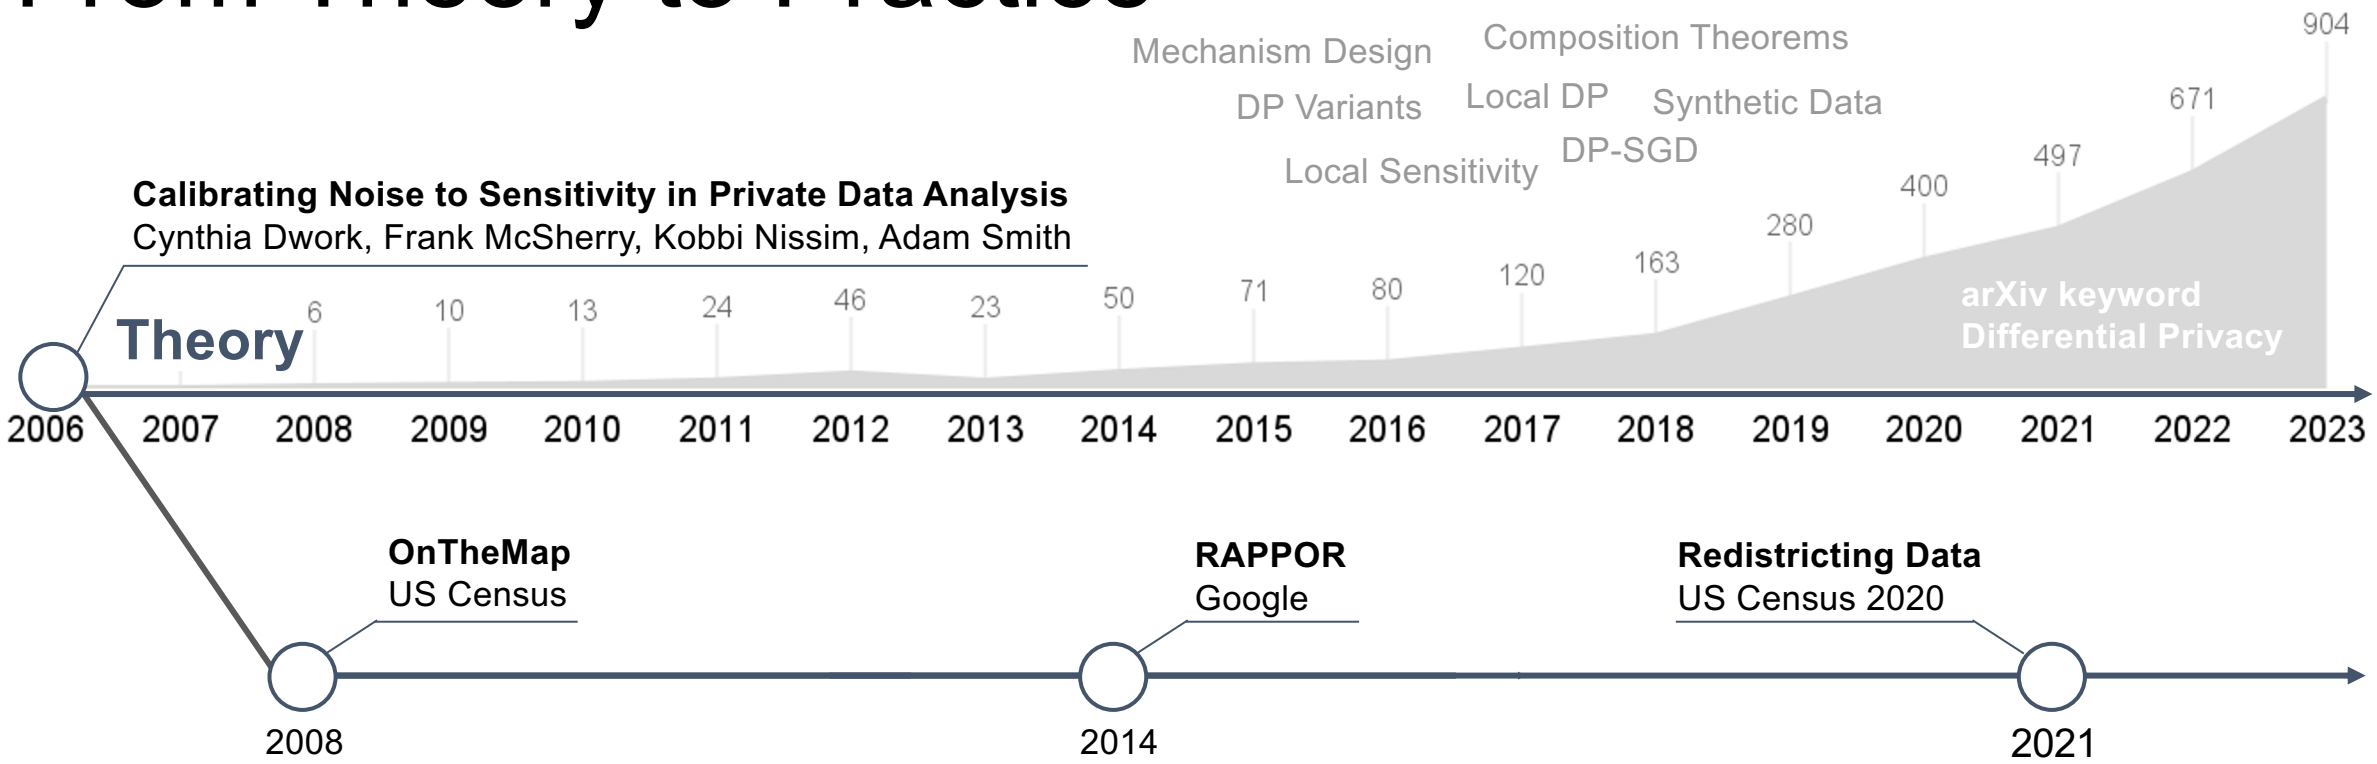

Real-World Applications

[Desfontaines Blog, 2021]

From Theory to Practice

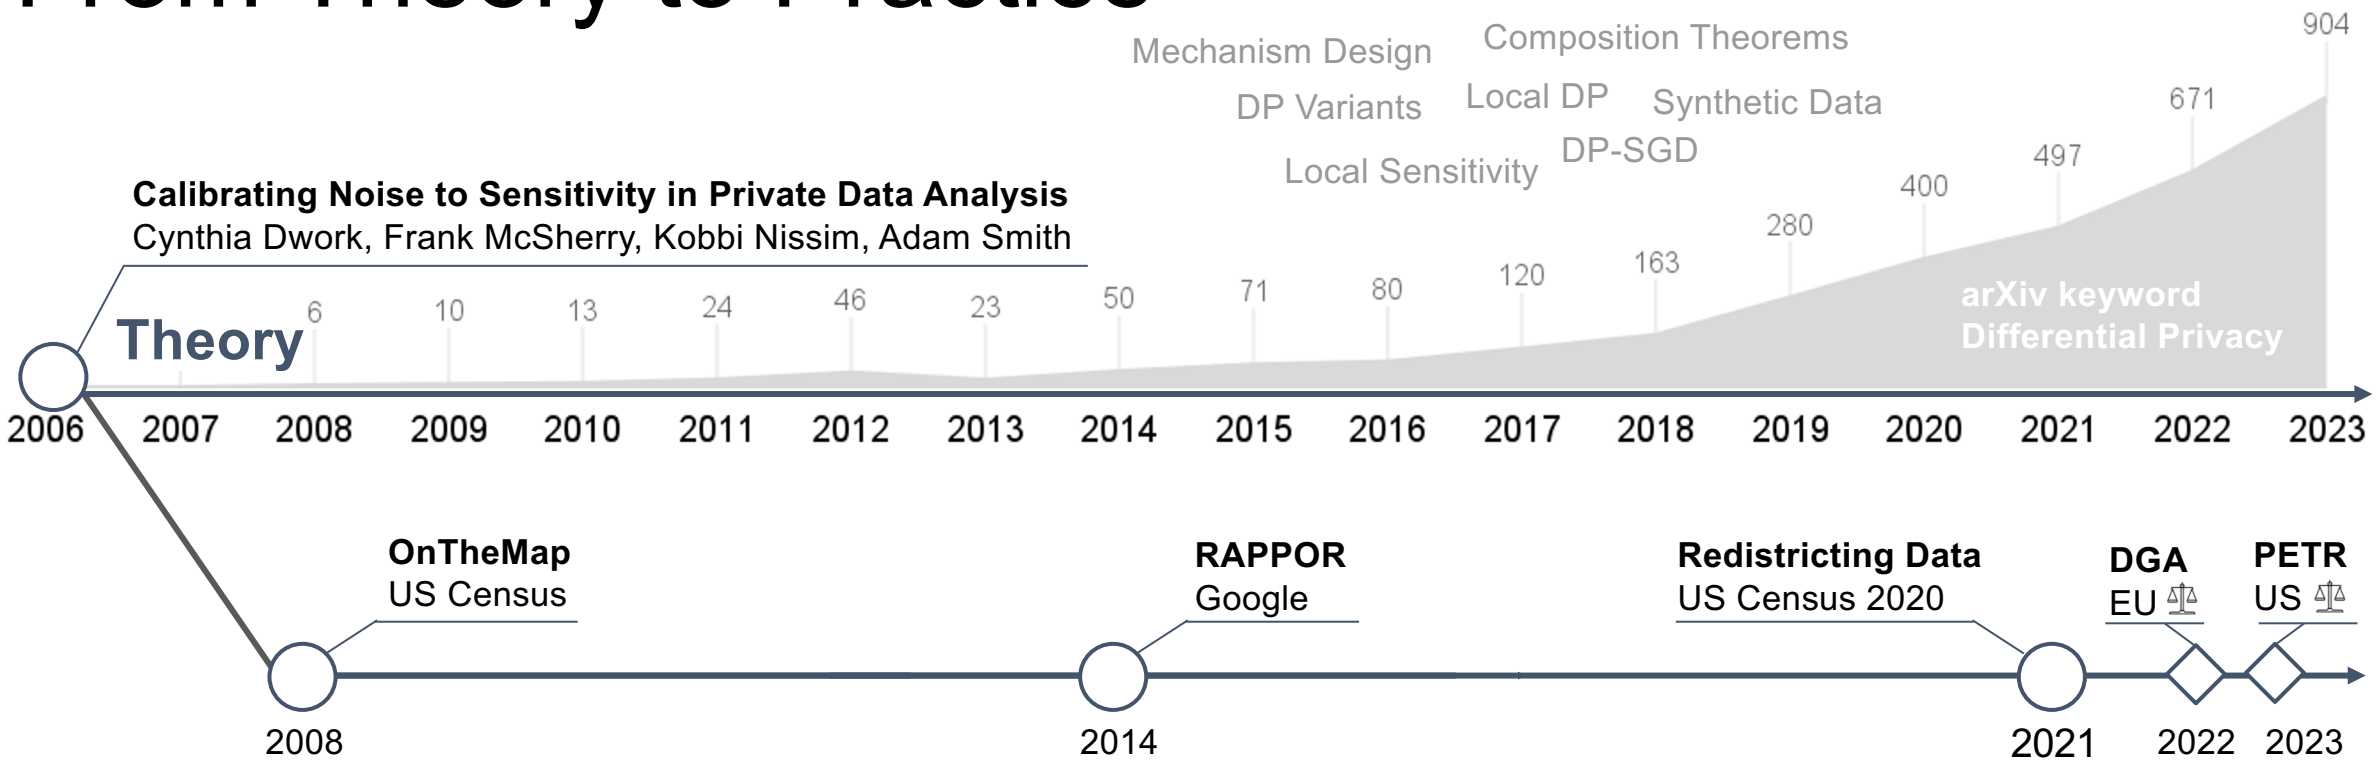

Real-World Applications

[Desfontaines Blog, 2021]

From Theory to Practice

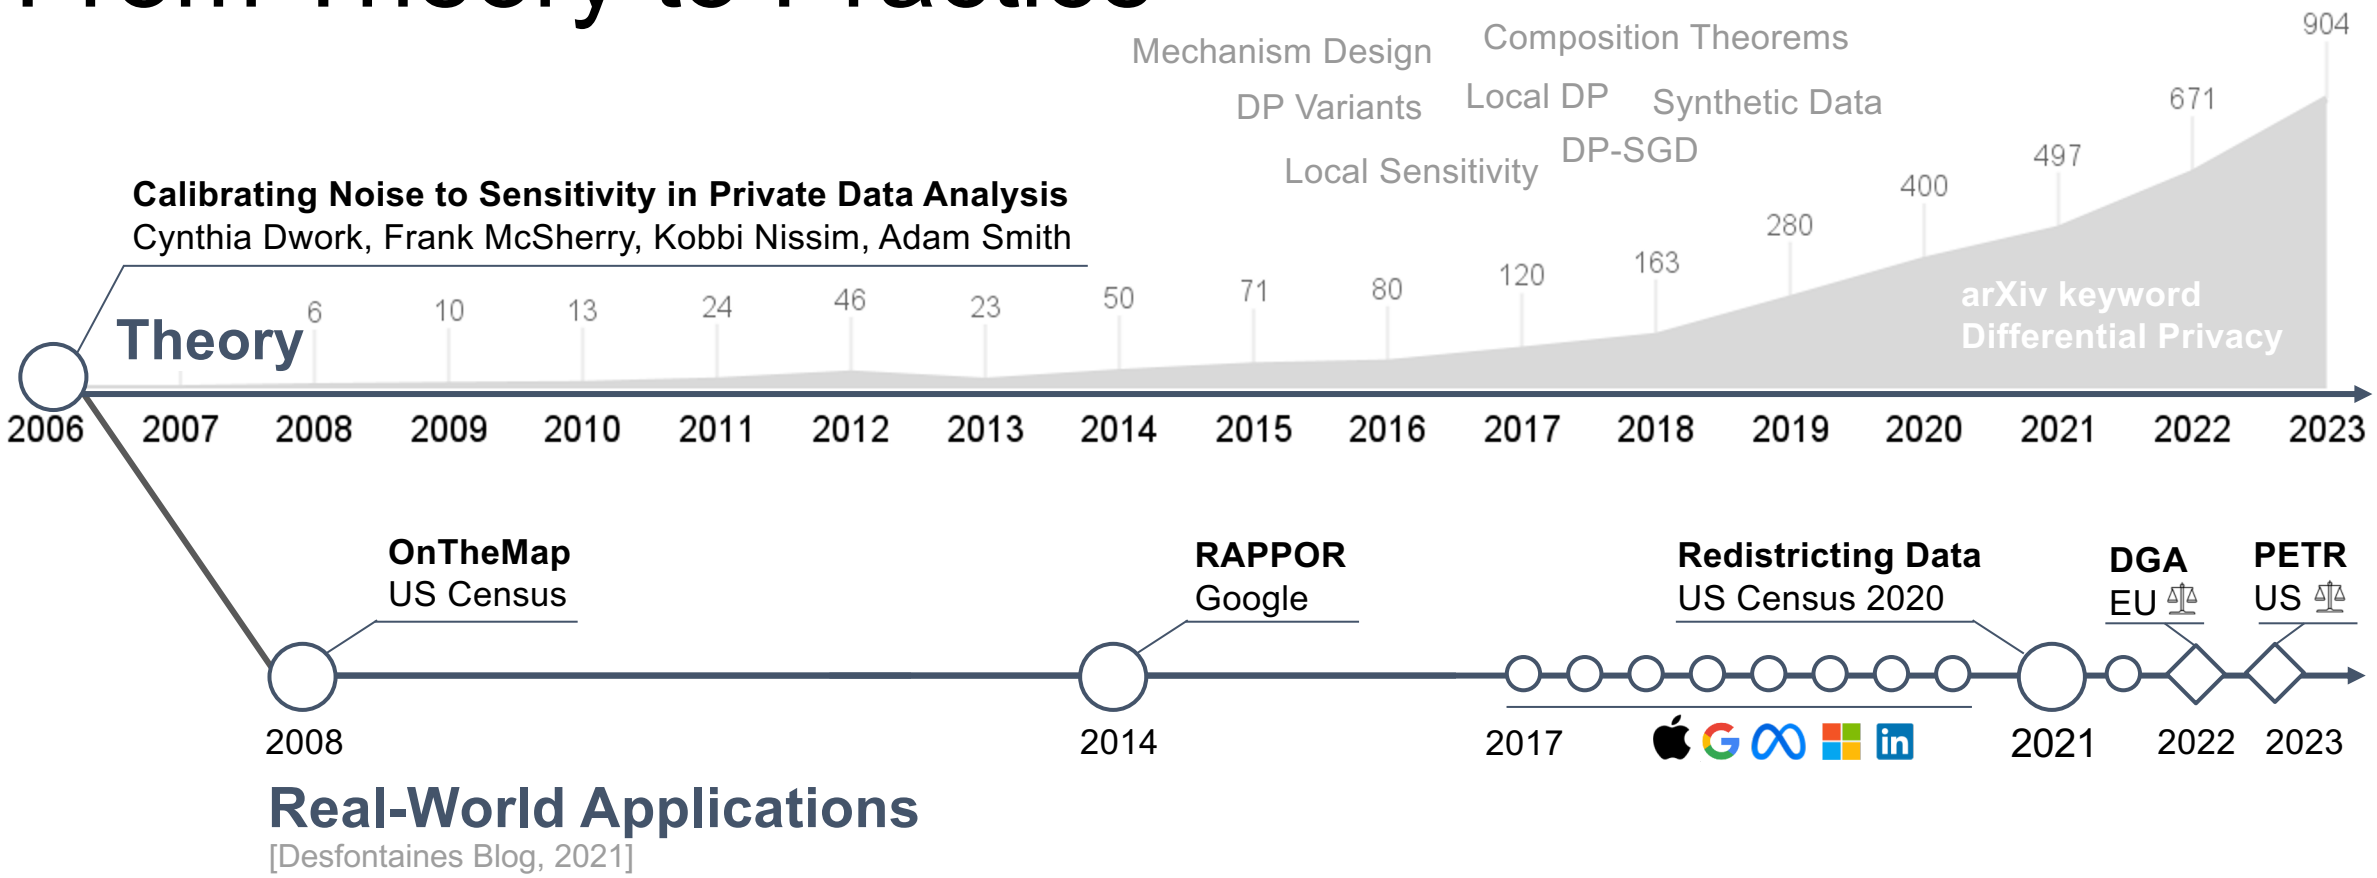

Real-World Applications

[Desfontaines Blog, 2021]

From Theory to Practice

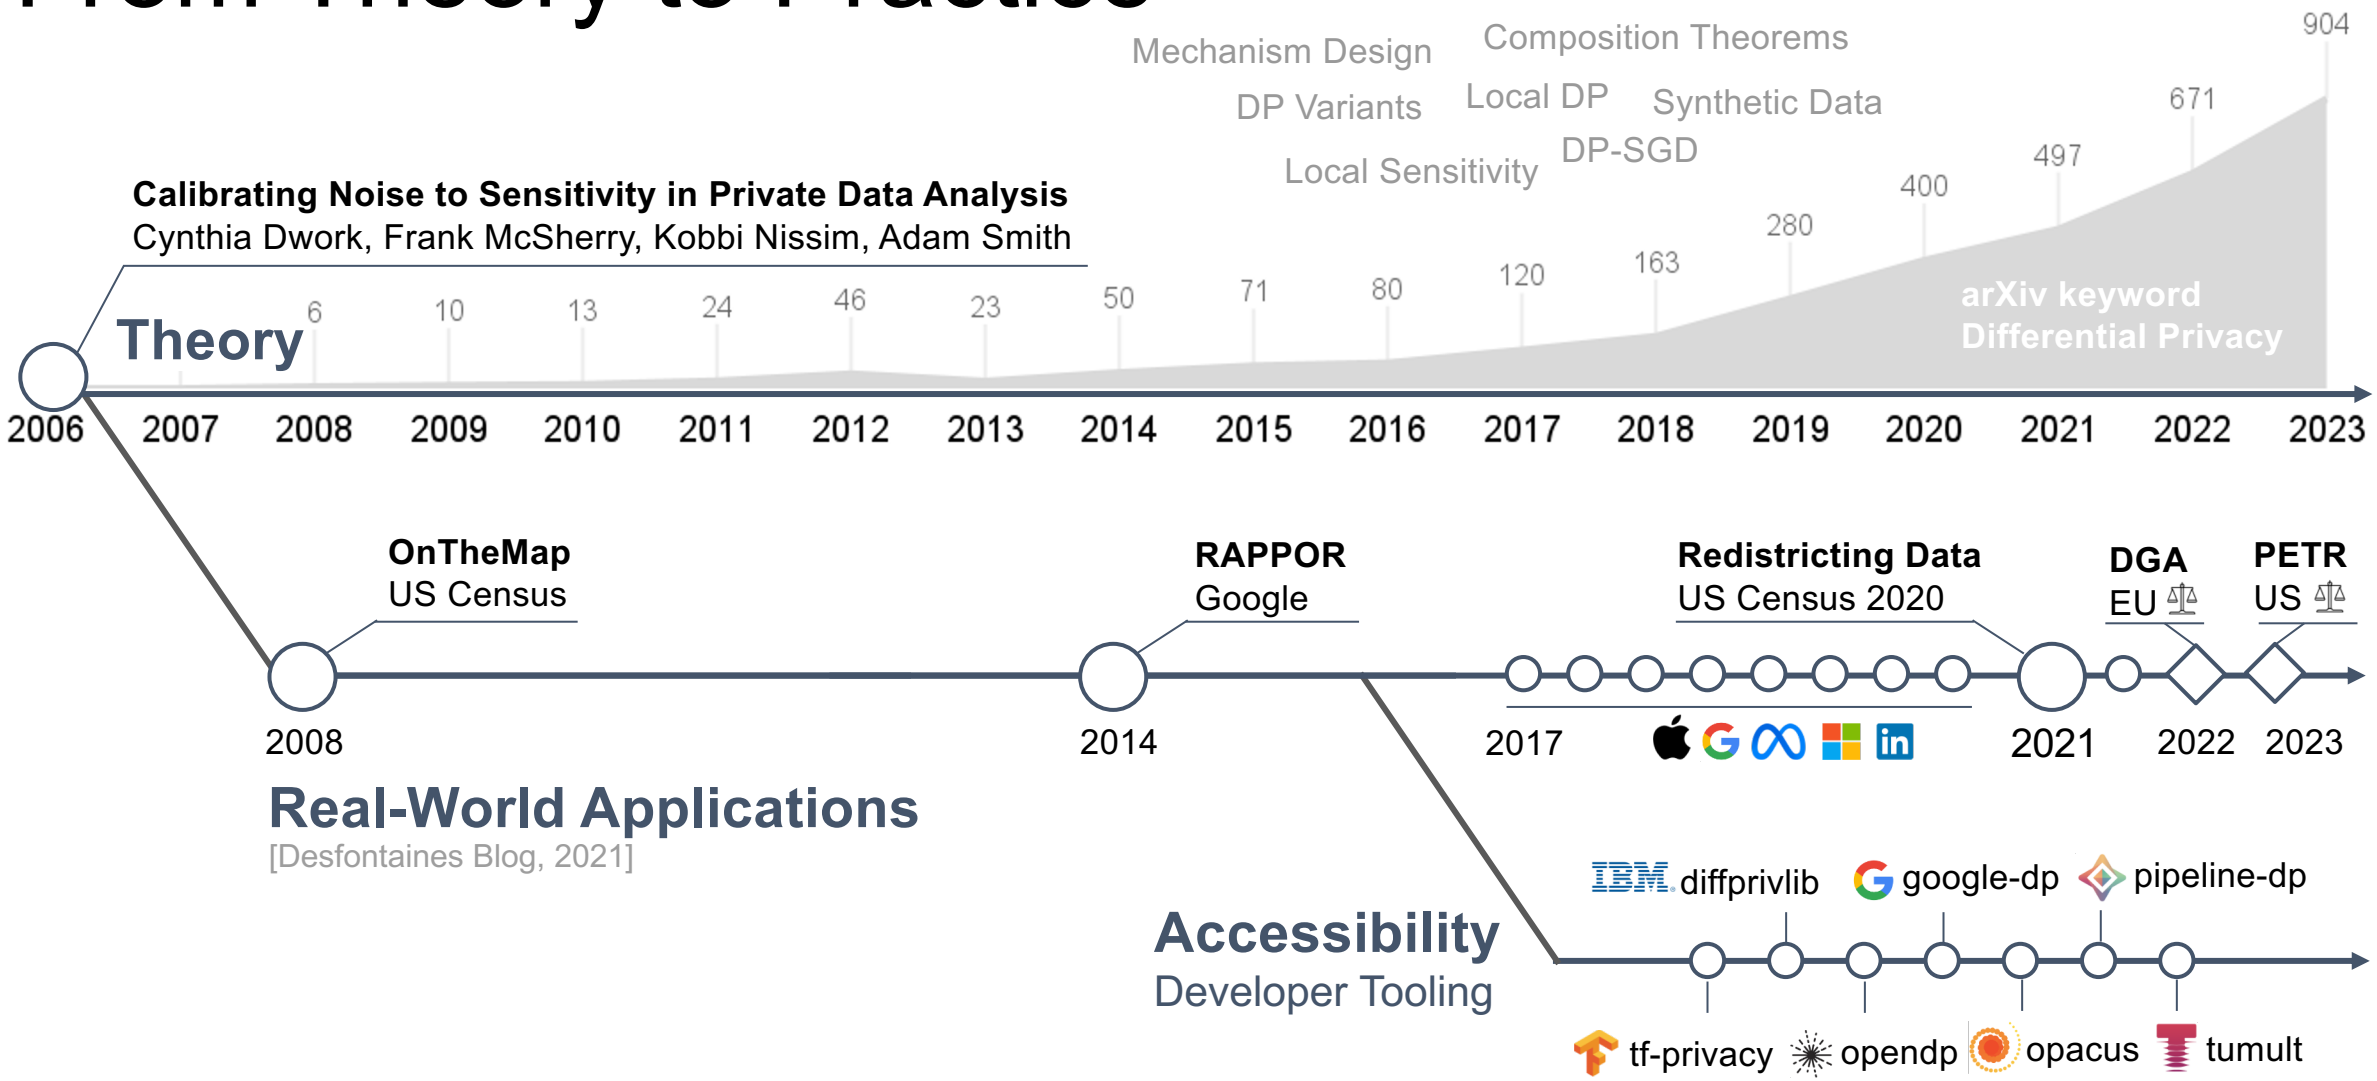

Deploying DP Applications

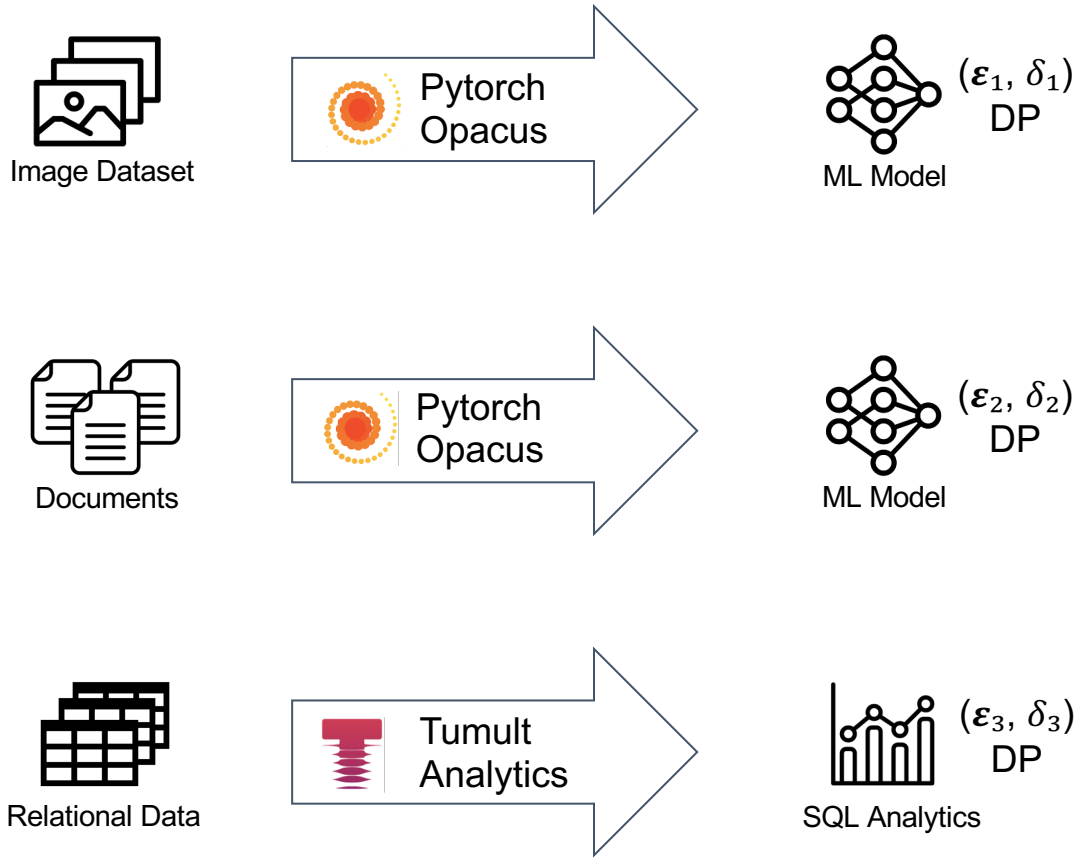

Deploying DP Applications

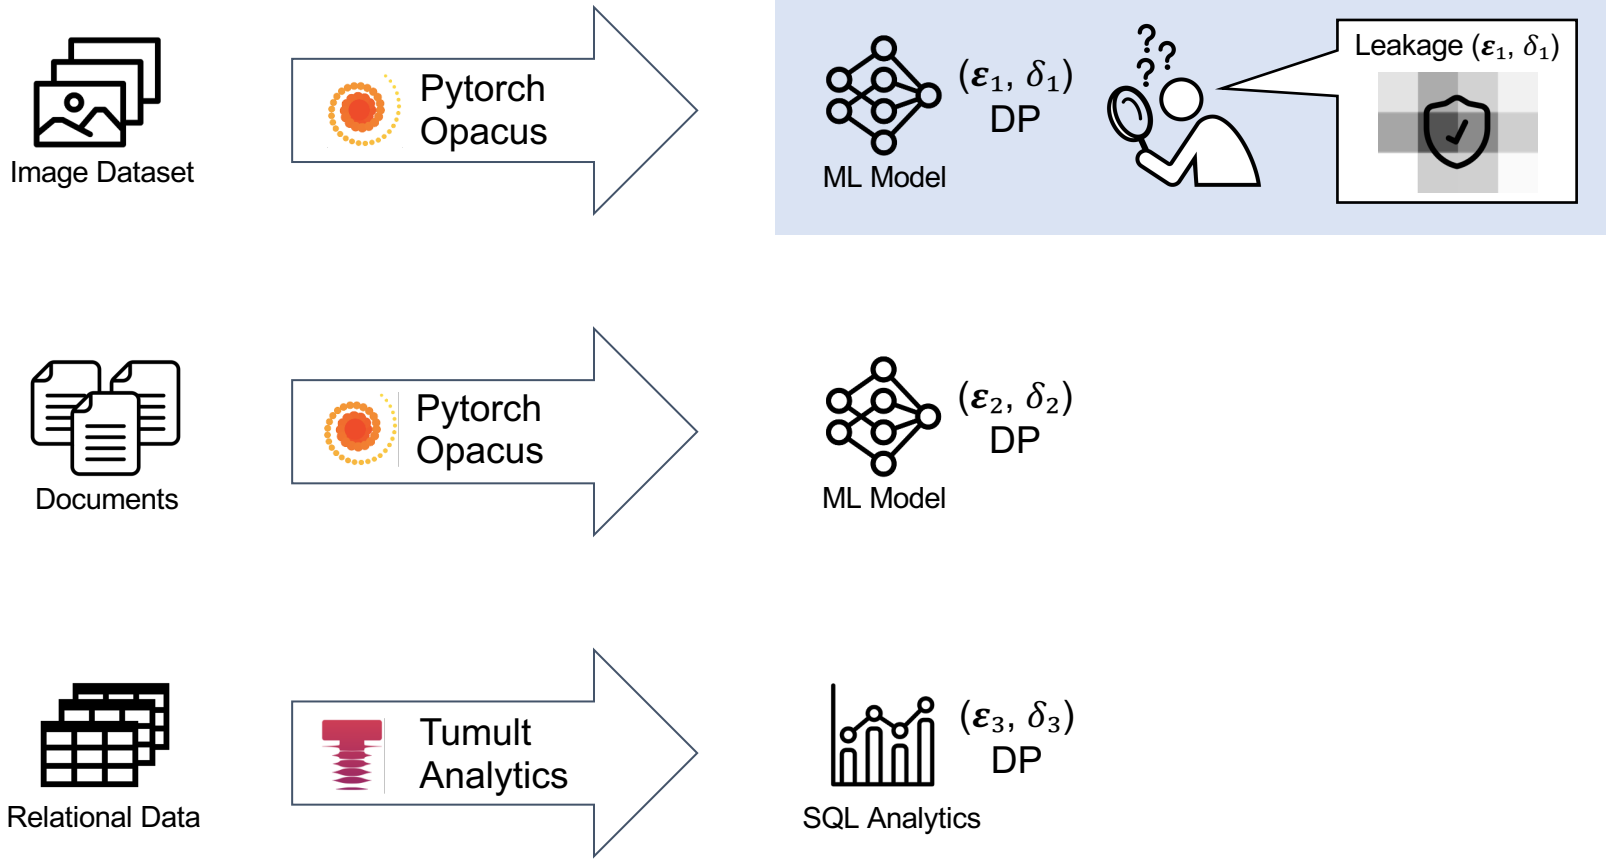

Deploying DP Applications

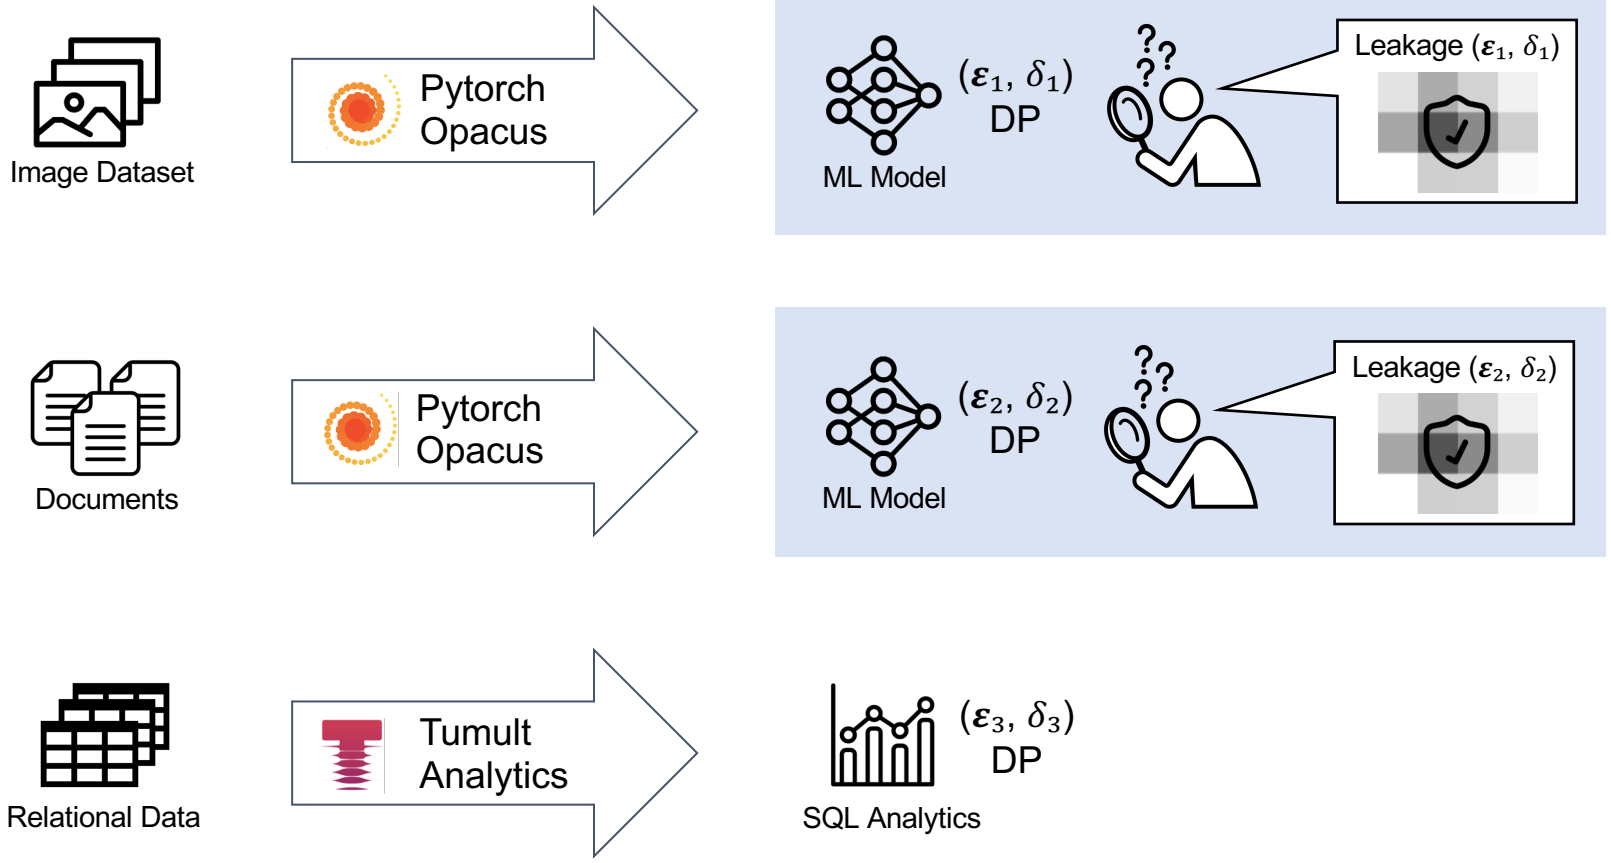

Deploying DP Applications

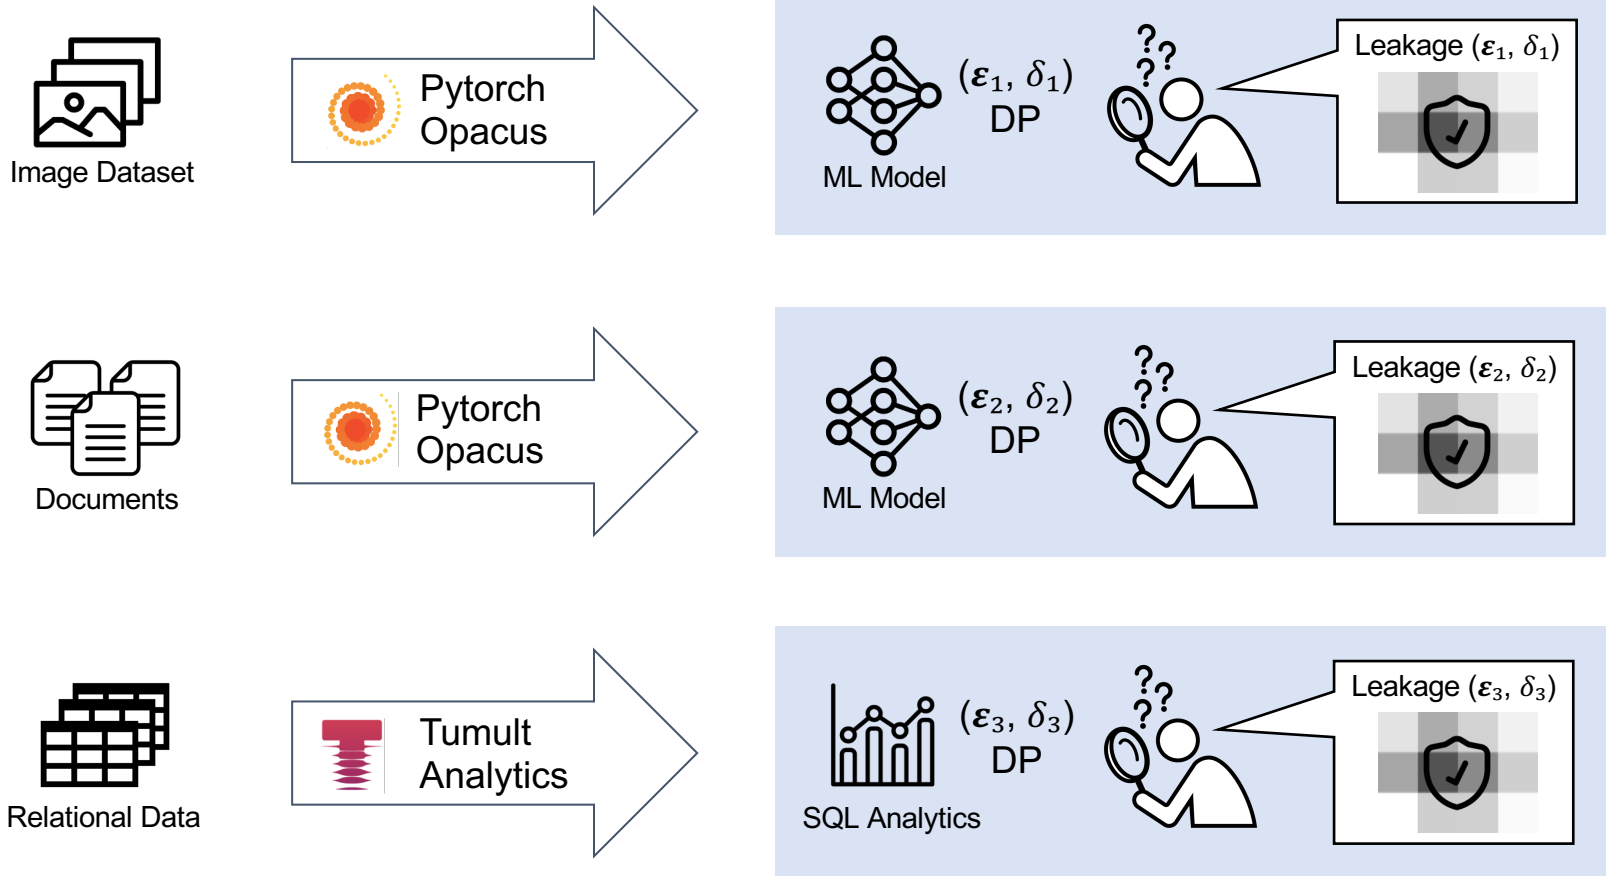

Deploying DP Applications

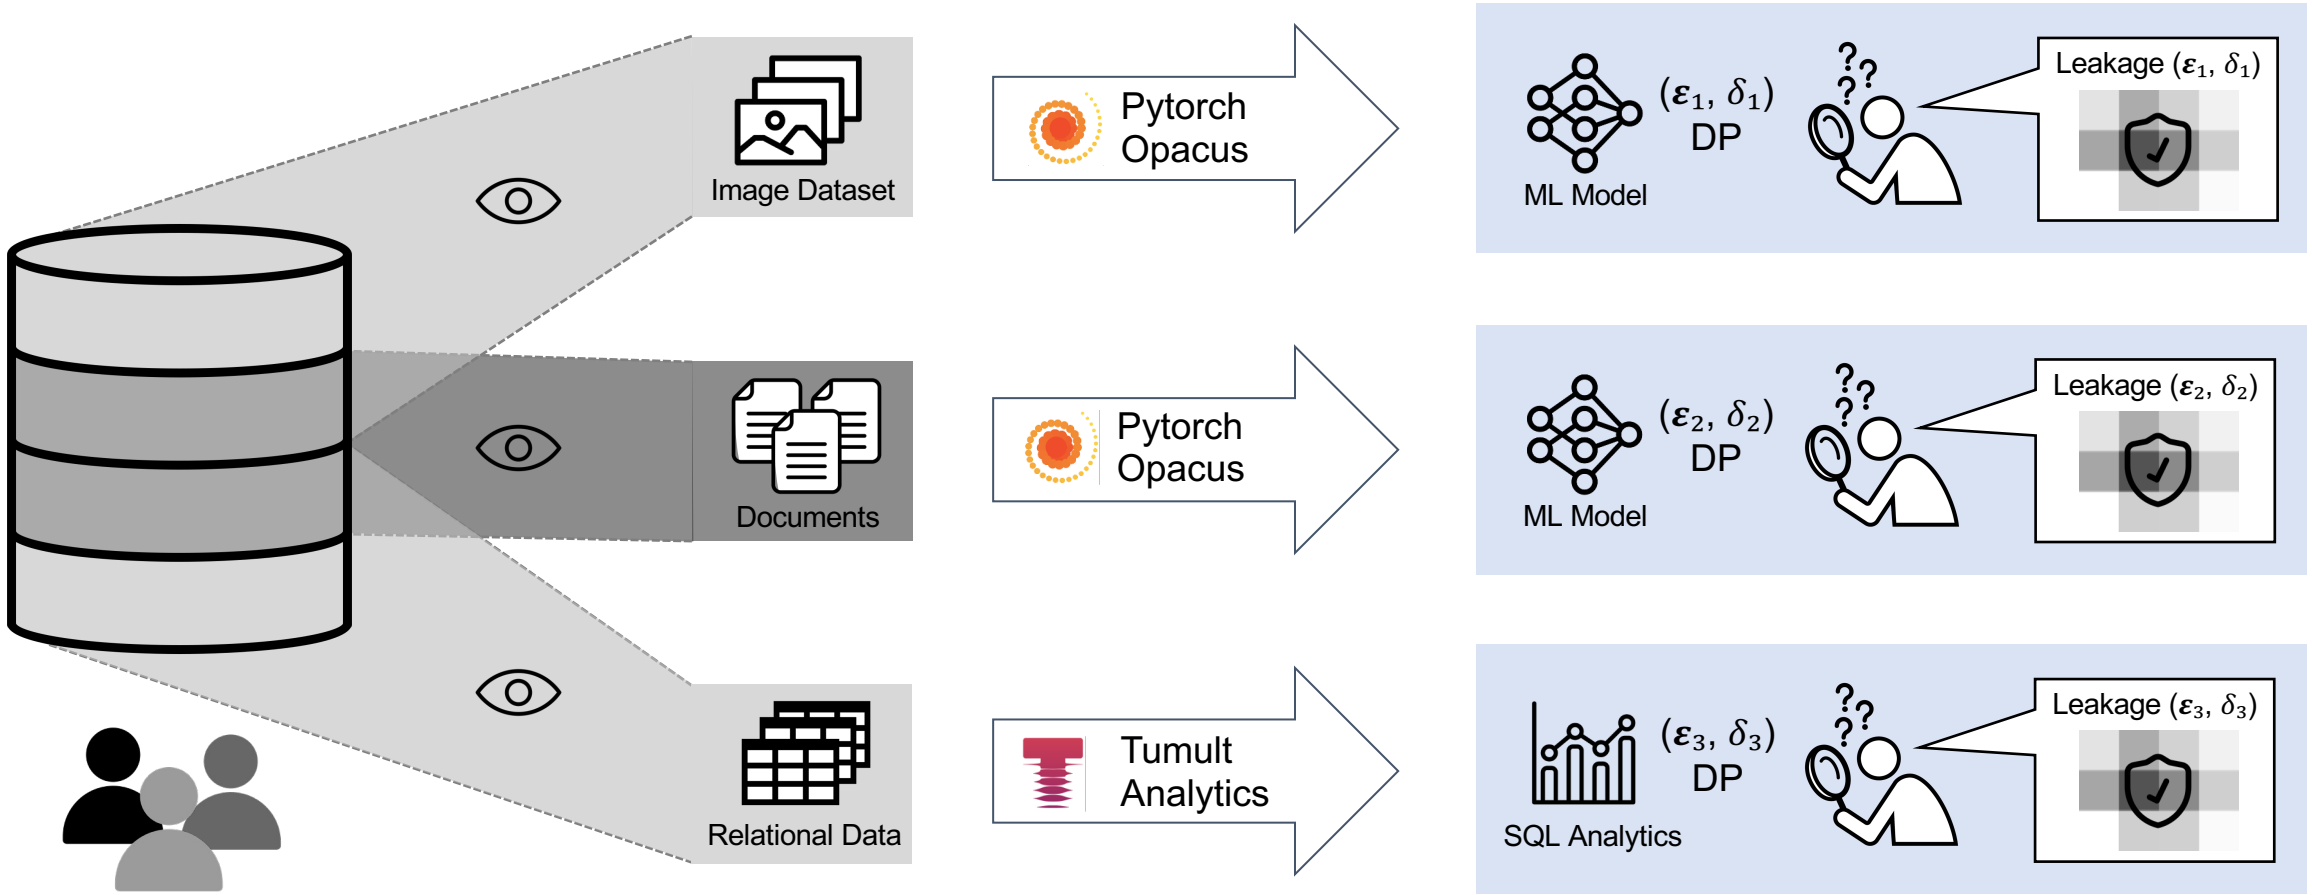

Deploying DP Applications

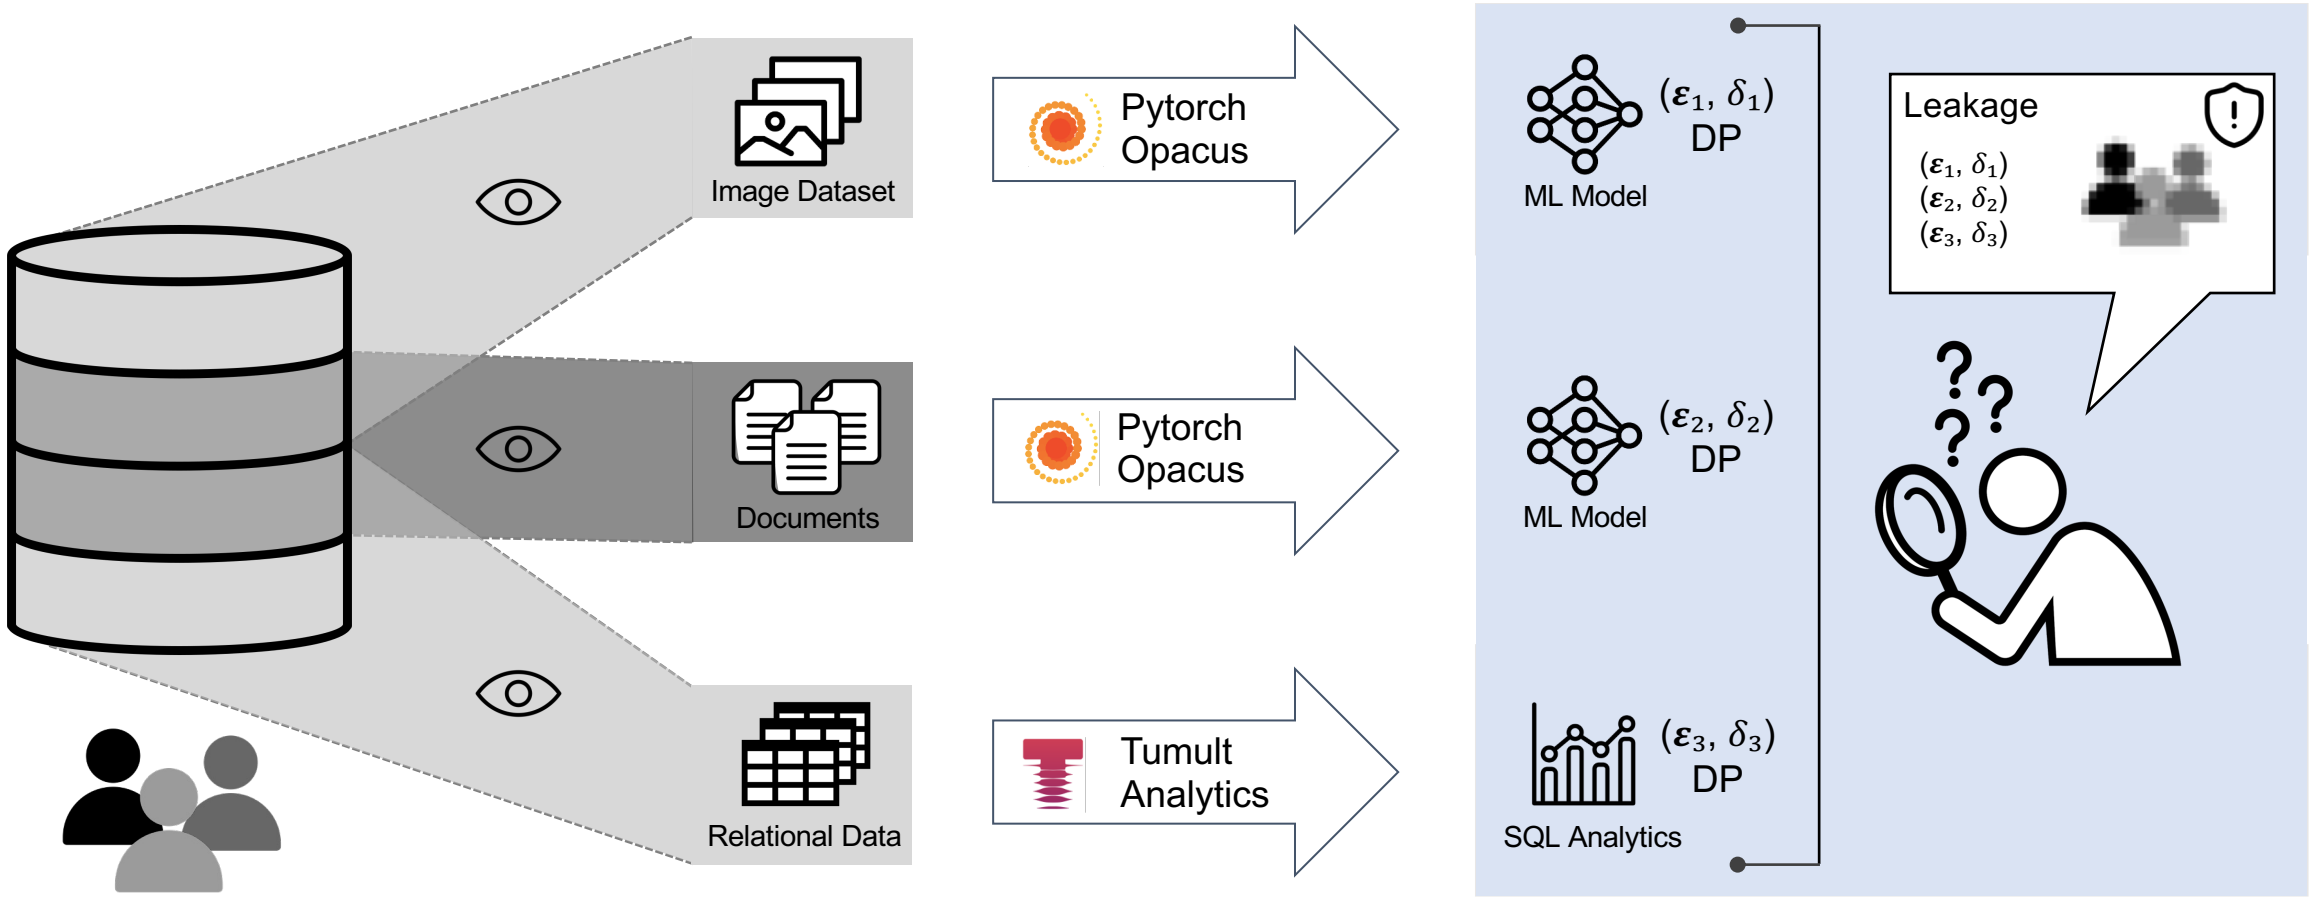

Cohere: Managing DP

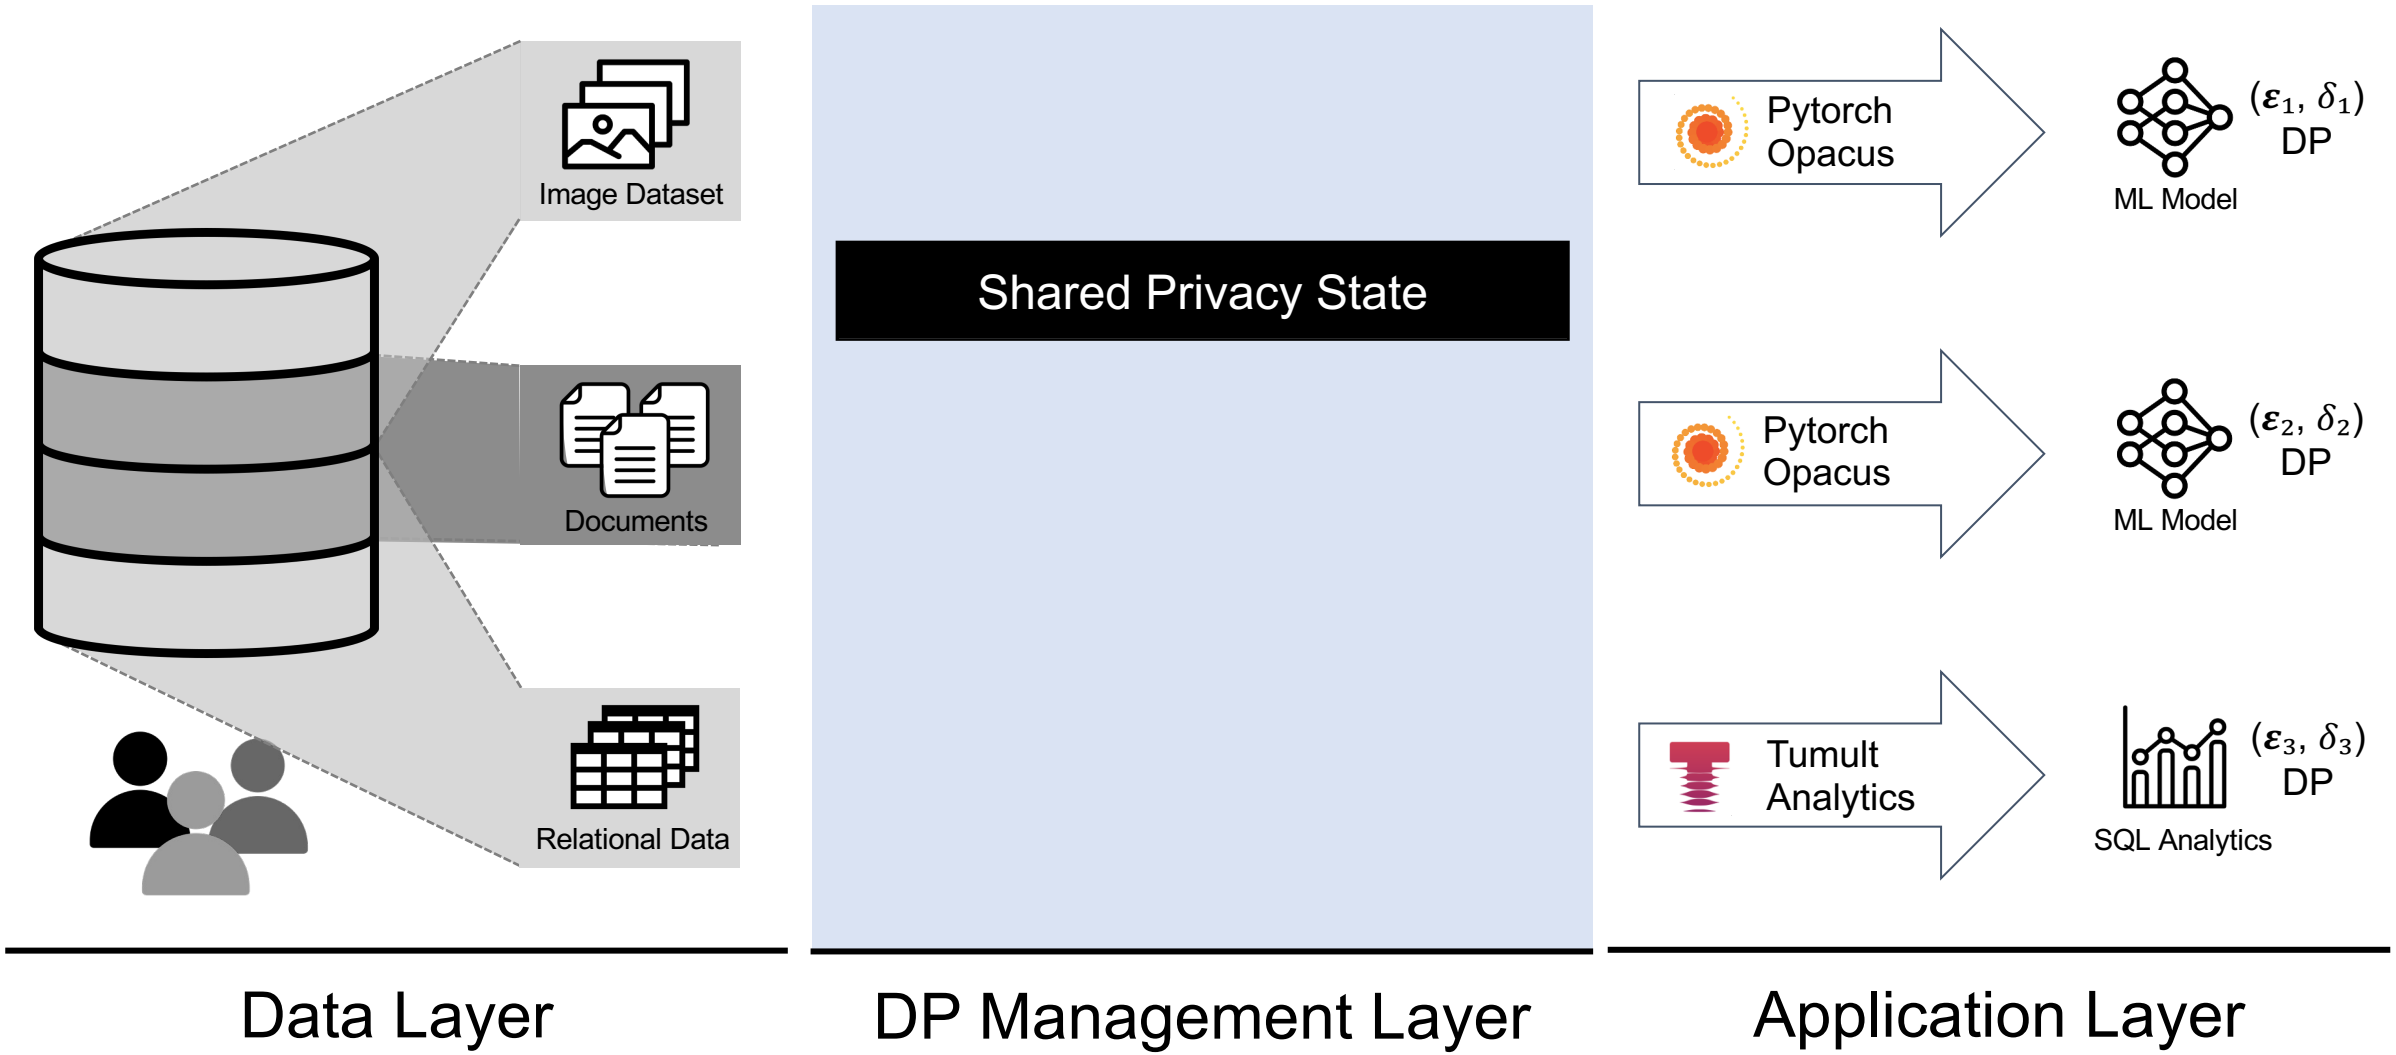

Cohere: Managing DP

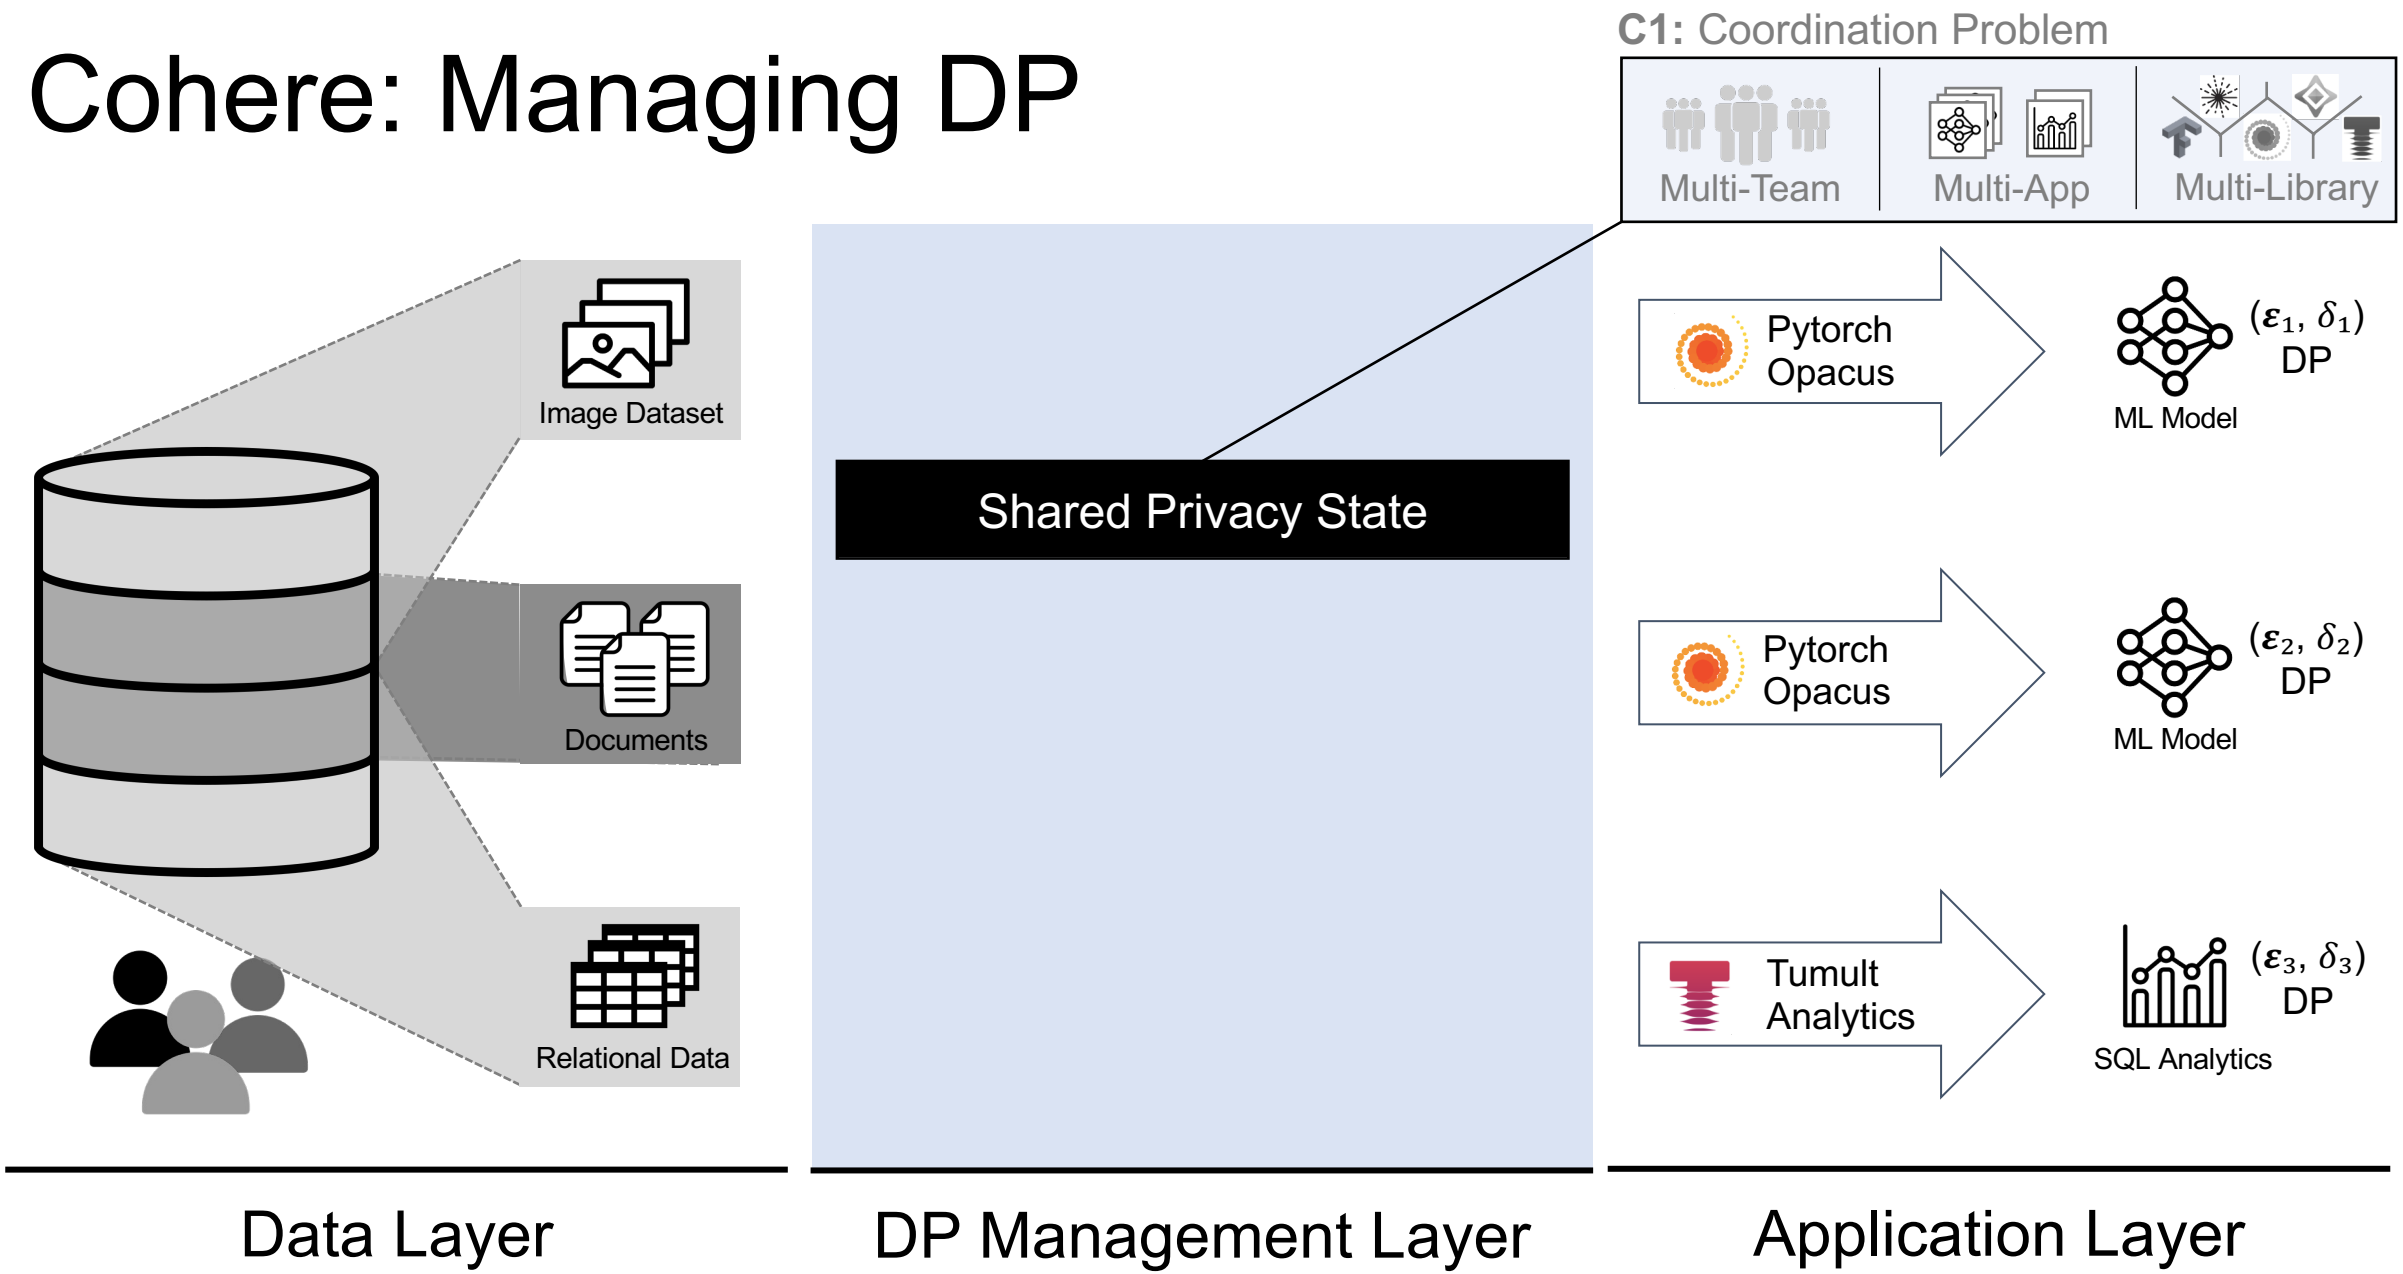

Cohere: Managing DP

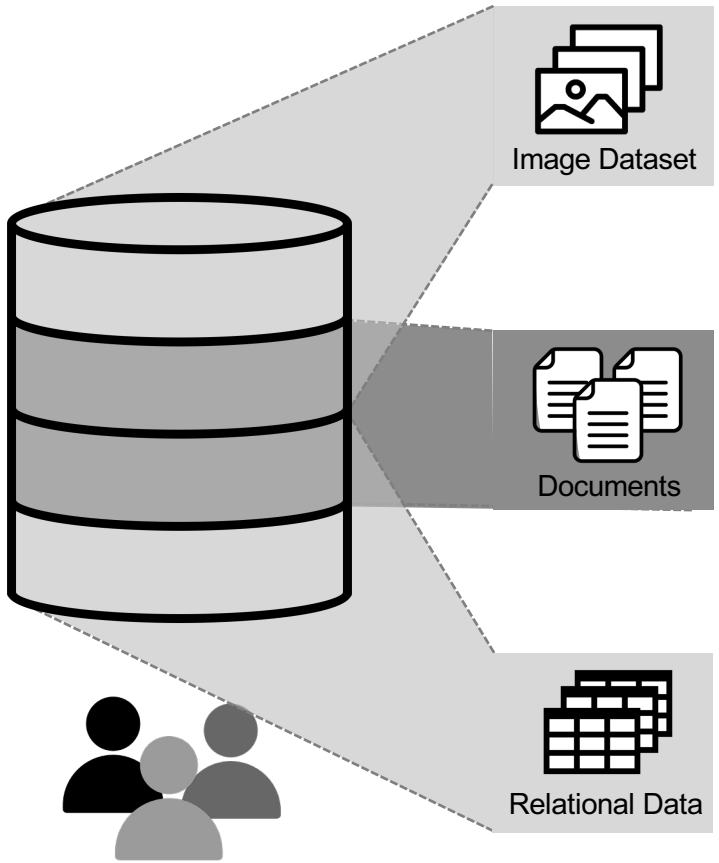

C1: Coordination Problem

Multi-Team | Multi-App | Multi-Library

Shared Privacy State

ε_1,δ_1 + ε_2,δ_2 + ε_3,δ_3 ≤ (ε, δ) DP

C2: Composition Complexity

Pytorch Opacus → ML Model $(\varepsilon_1, \delta_1)$ DP

Pytorch Opacus → ML Model $(\varepsilon_2, \delta_2)$ DP

Tumult Analytics → SQL Analytics $(\varepsilon_3, \delta_3)$ DP

Data Layer | DP Management Layer | Application Layer

Cohere: Managing DP

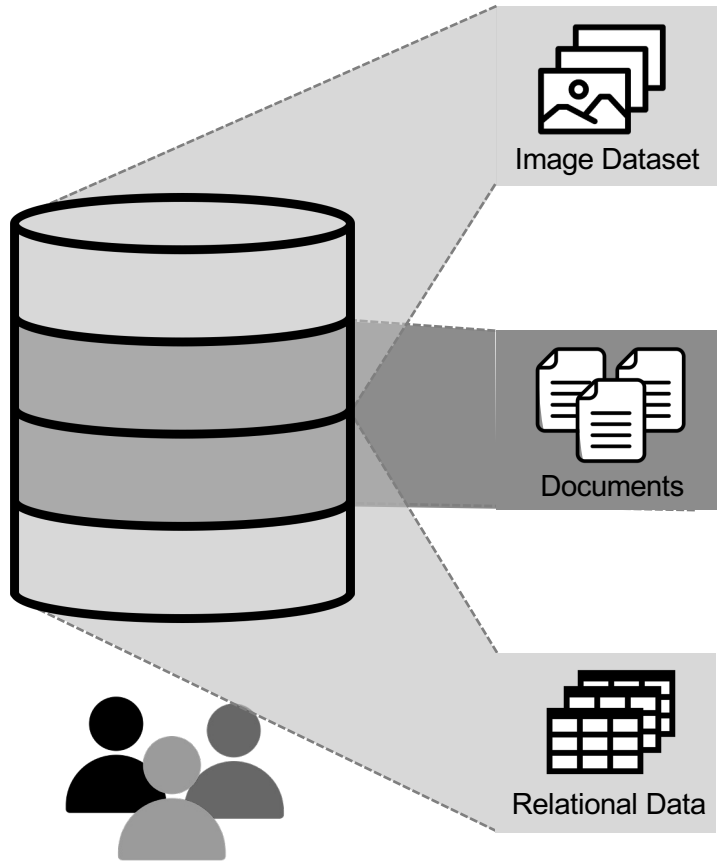

C1: Coordination Problem

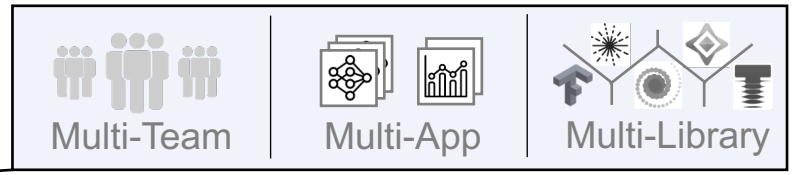

Shared Privacy State

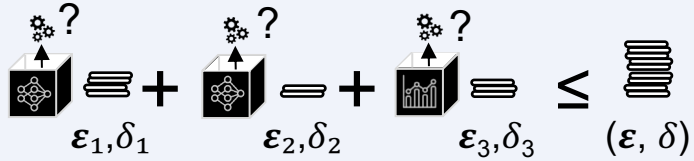

C2: Composition Complexity

Resource Planning

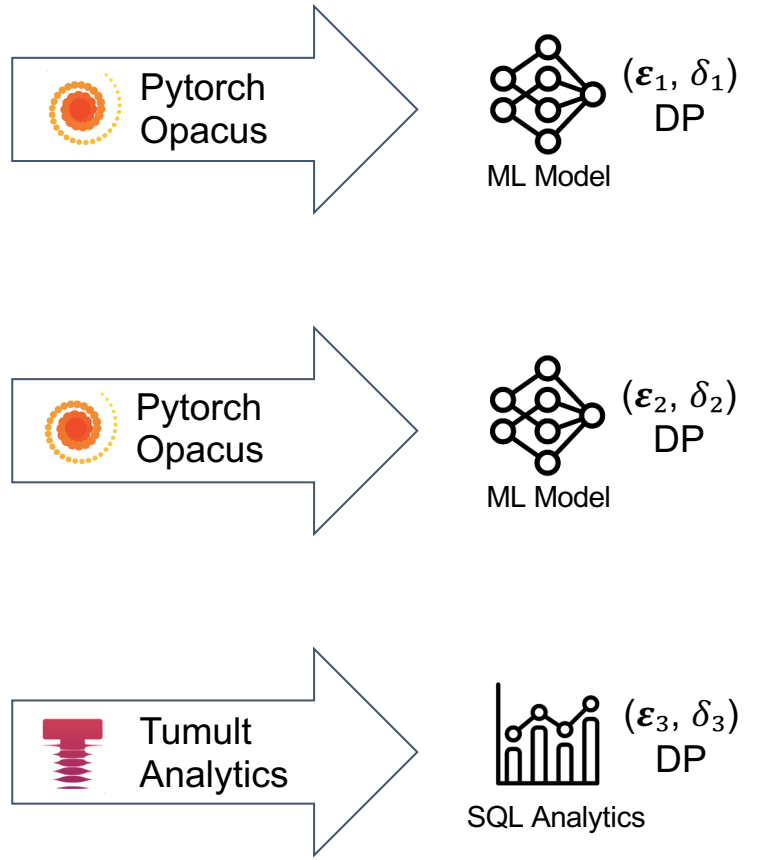

Data Layer

DP Management Layer

Application Layer

Cohere: Managing DP

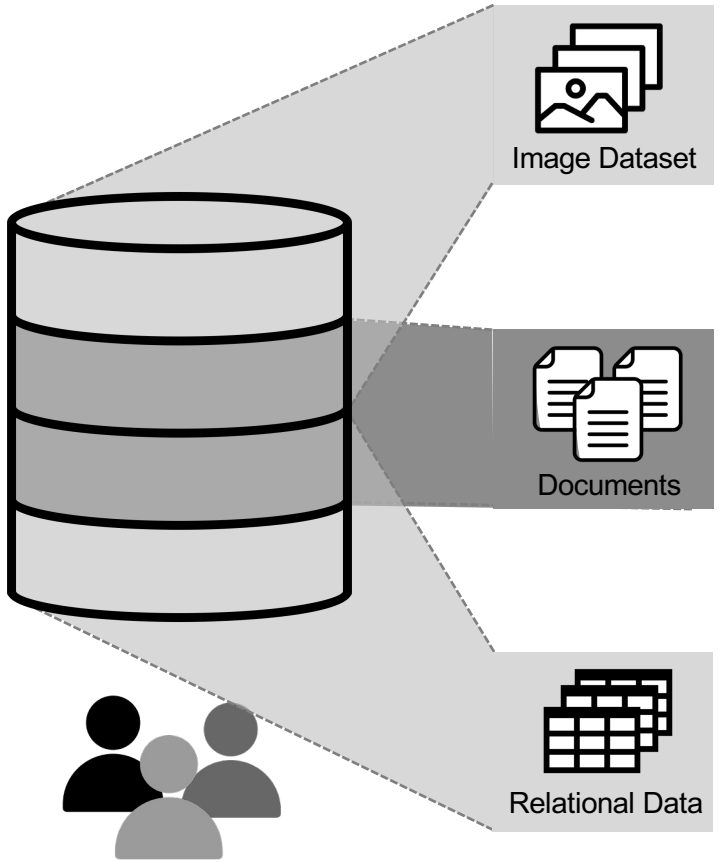

C1: Coordination Problem

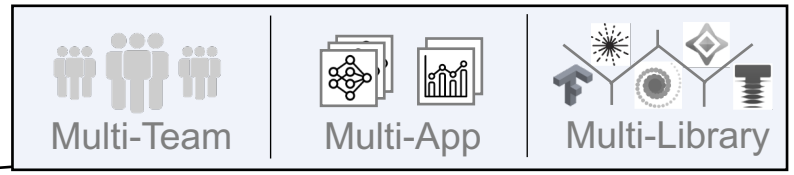

Shared Privacy State

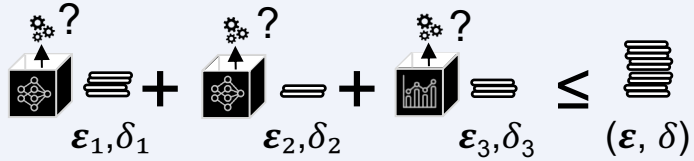

C2: Composition Complexity

Resource Planning

C3: Resource Allocation

(ε, δ) - DP Total → ε_1,δ_1; ε_2,δ_2; ε_3,δ_3

C4: Continuity Guarantee

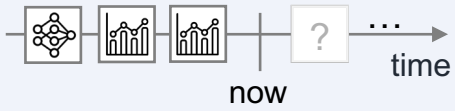

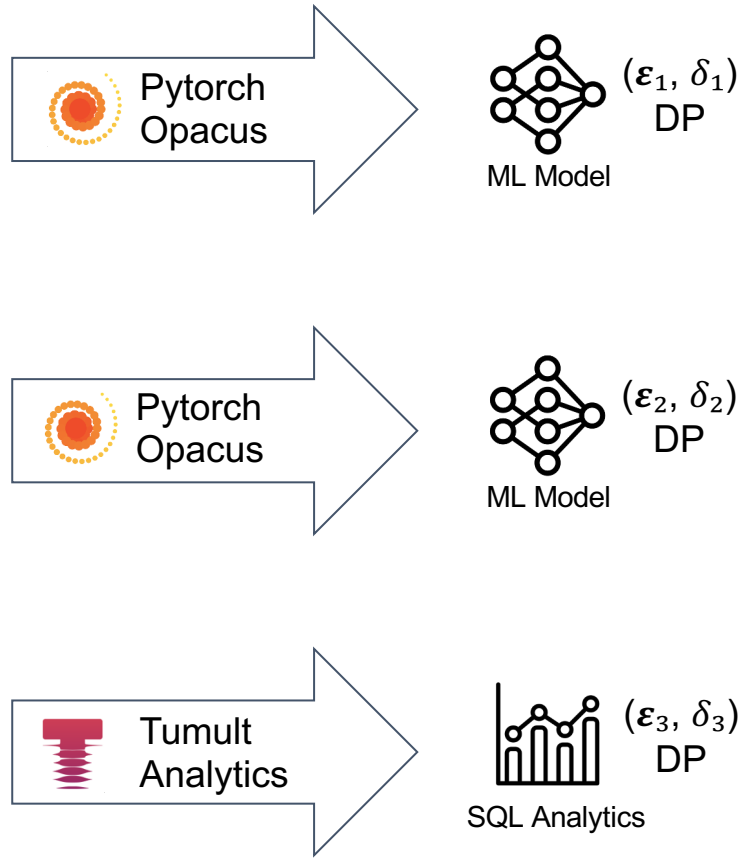

Data Layer | DP Management Layer | Application Layer

Cohere: Managing DP

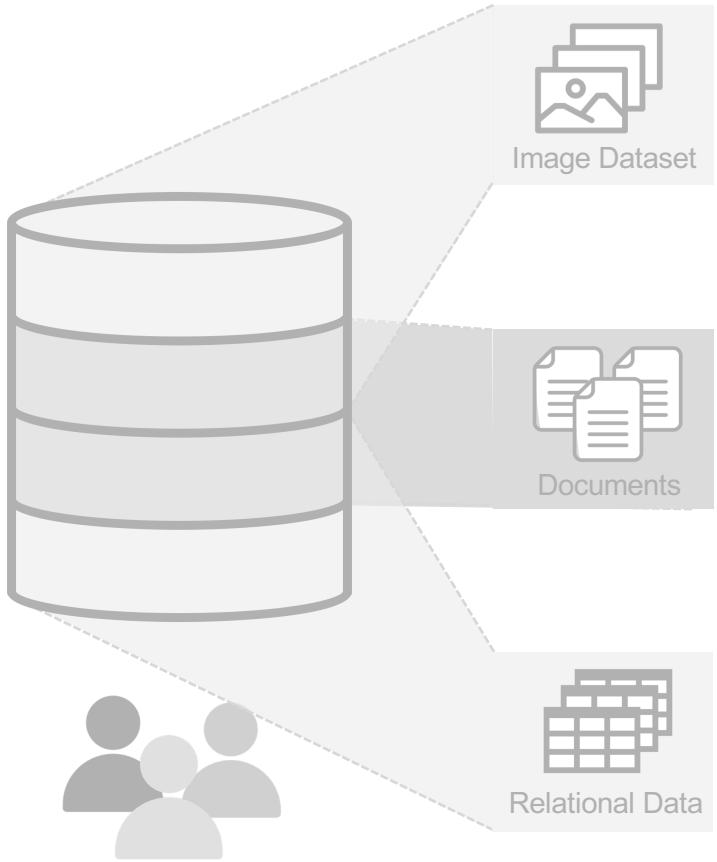

Data Layer

C1: Coordination Problem

Multi-Team | Multi-App | Multi-Library

Shared Privacy State

ε_1,δ_1 + ε_2,δ_2 + ε_3,δ_3 ≤ (ε, δ) DP

C2: Composition Complexity

Resource Planning

C3: Resource Allocation

(ε, δ) - DP (Total) → ε_1,δ_1; ε_2,δ_2; ε_3,δ_3

C4: Continuity Guarantee

now ? … time

DP Management Layer

Pytorch Opacus → ML Model $(\varepsilon_1, \delta_1)$ DP

Pytorch Opacus → ML Model $(\varepsilon_2, \delta_2)$ DP

Tumult Analytics → SQL Analytics $(\varepsilon_3, \delta_3)$ DP

Application Layer

Unifying the Application Layer

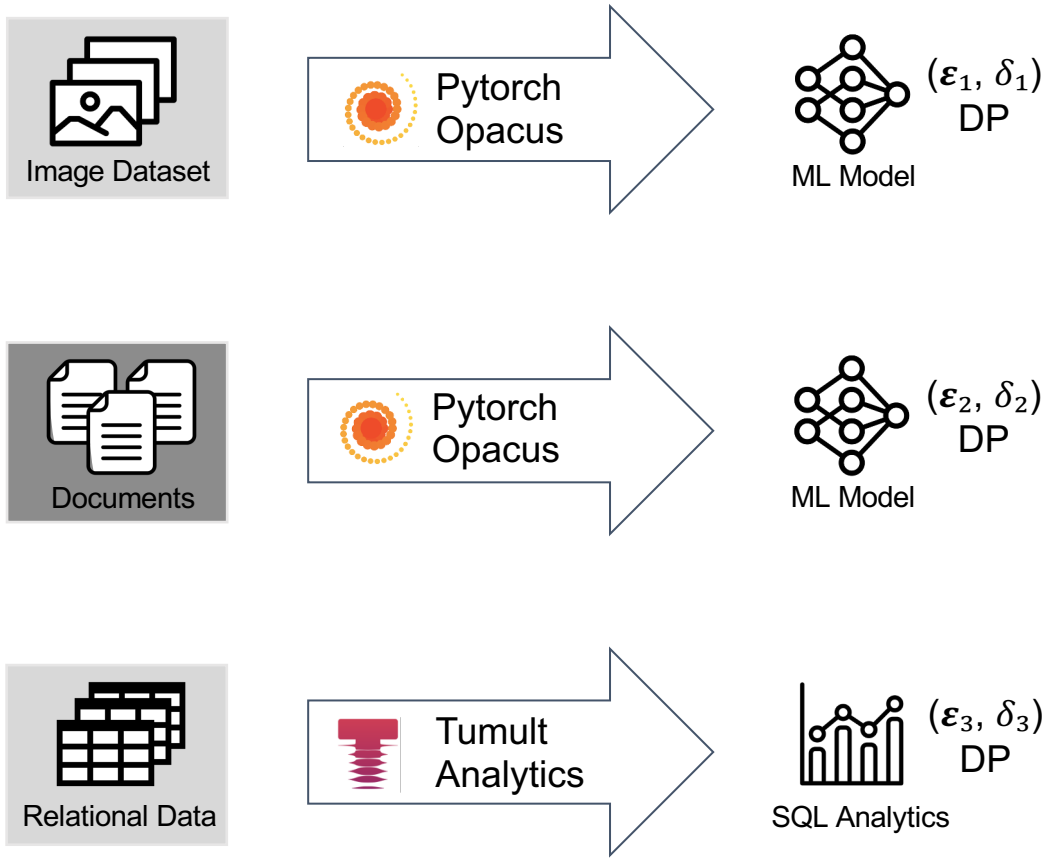

Unifying the Application Layer

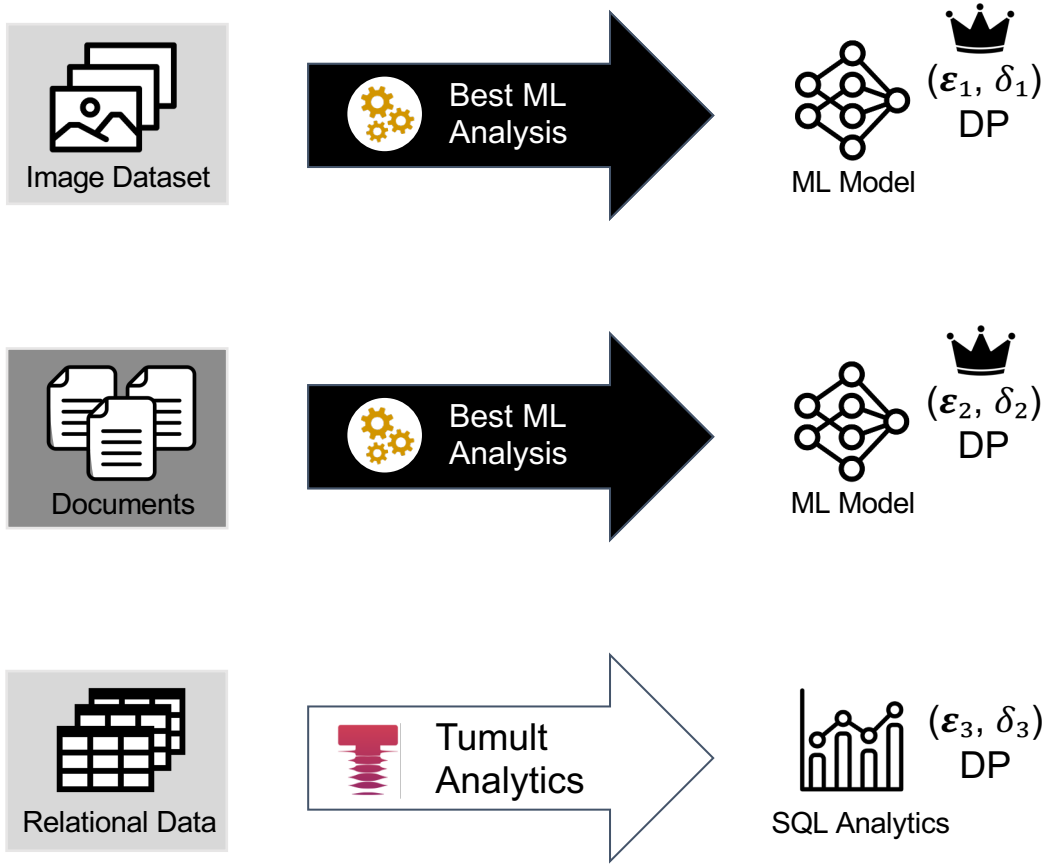

Unifying the Application Layer

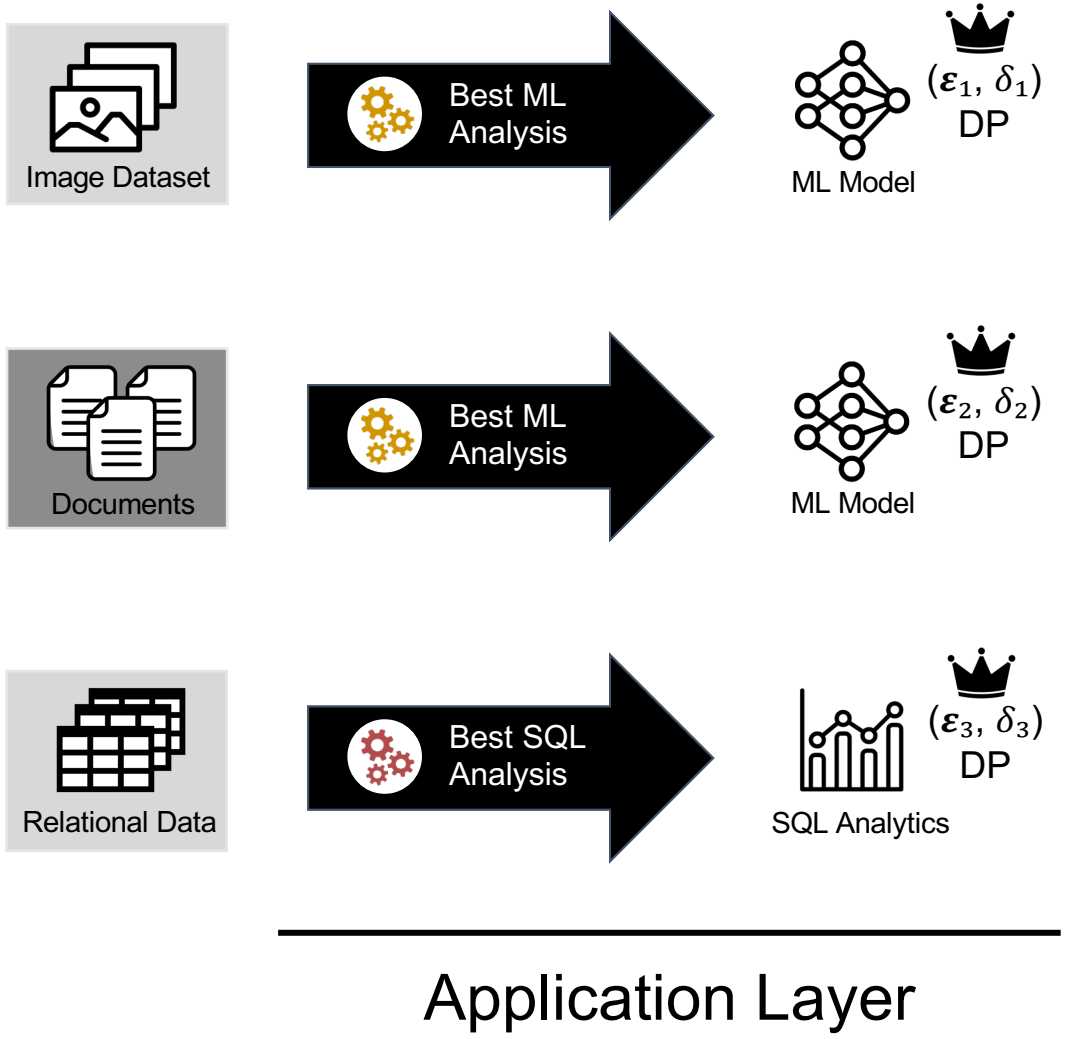

Unifying the Application Layer

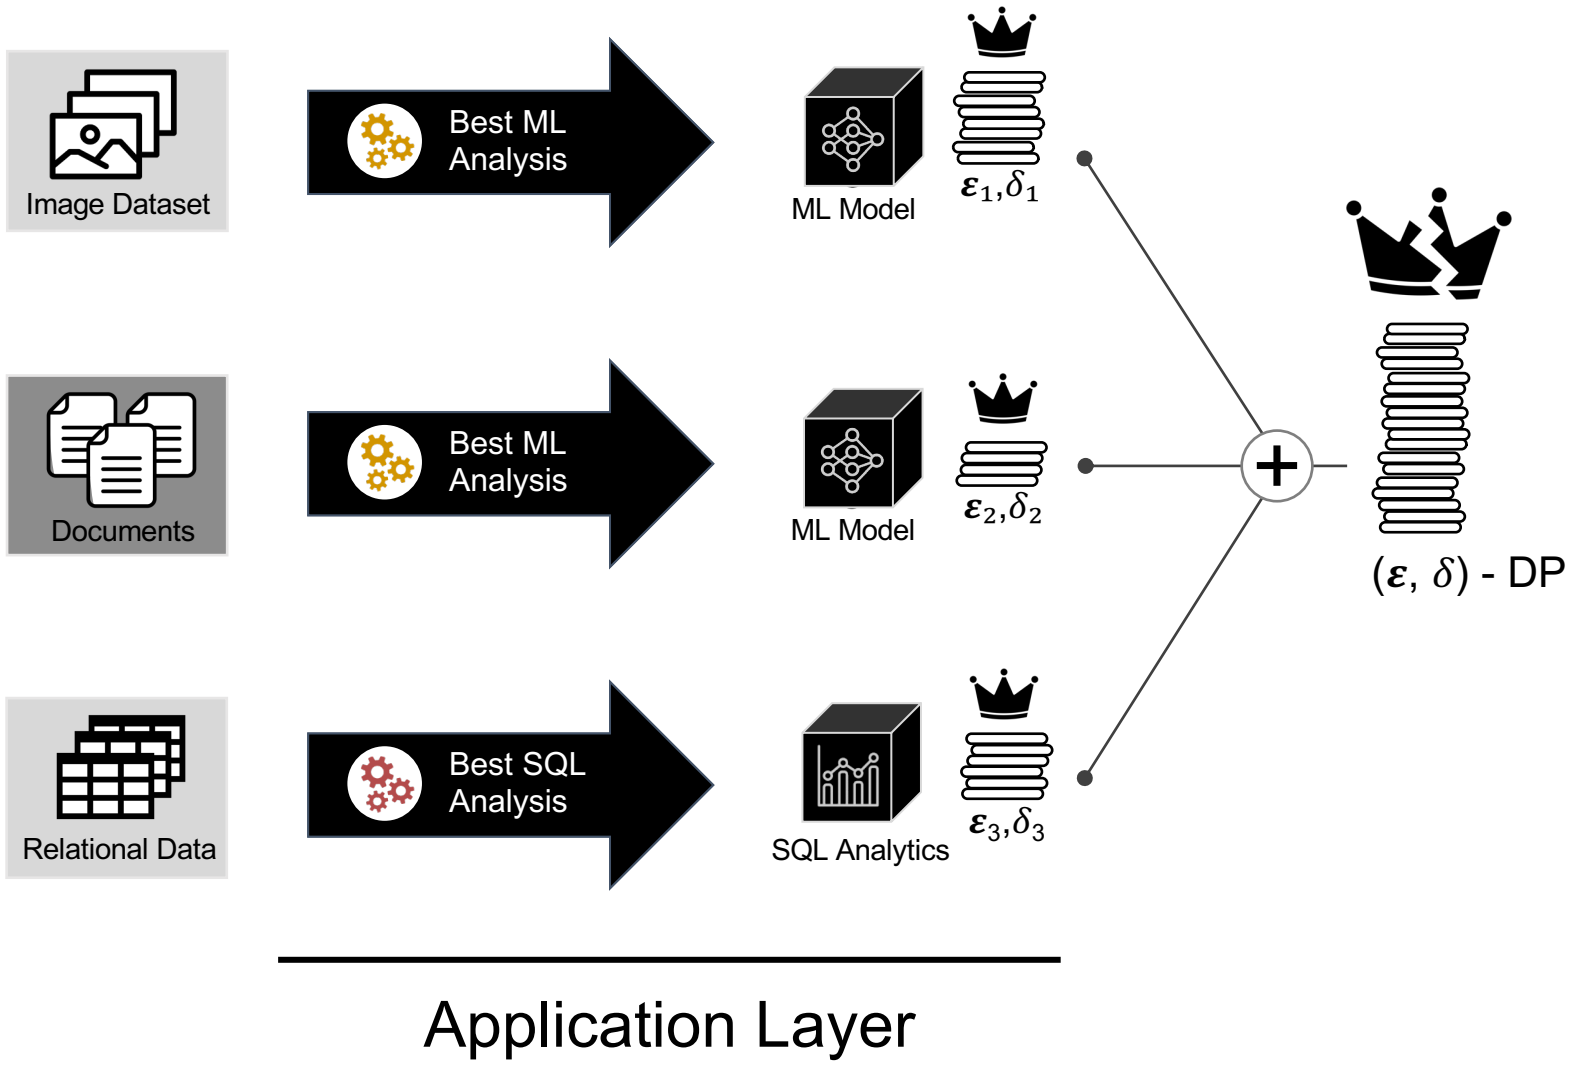

Unifying the Application Layer

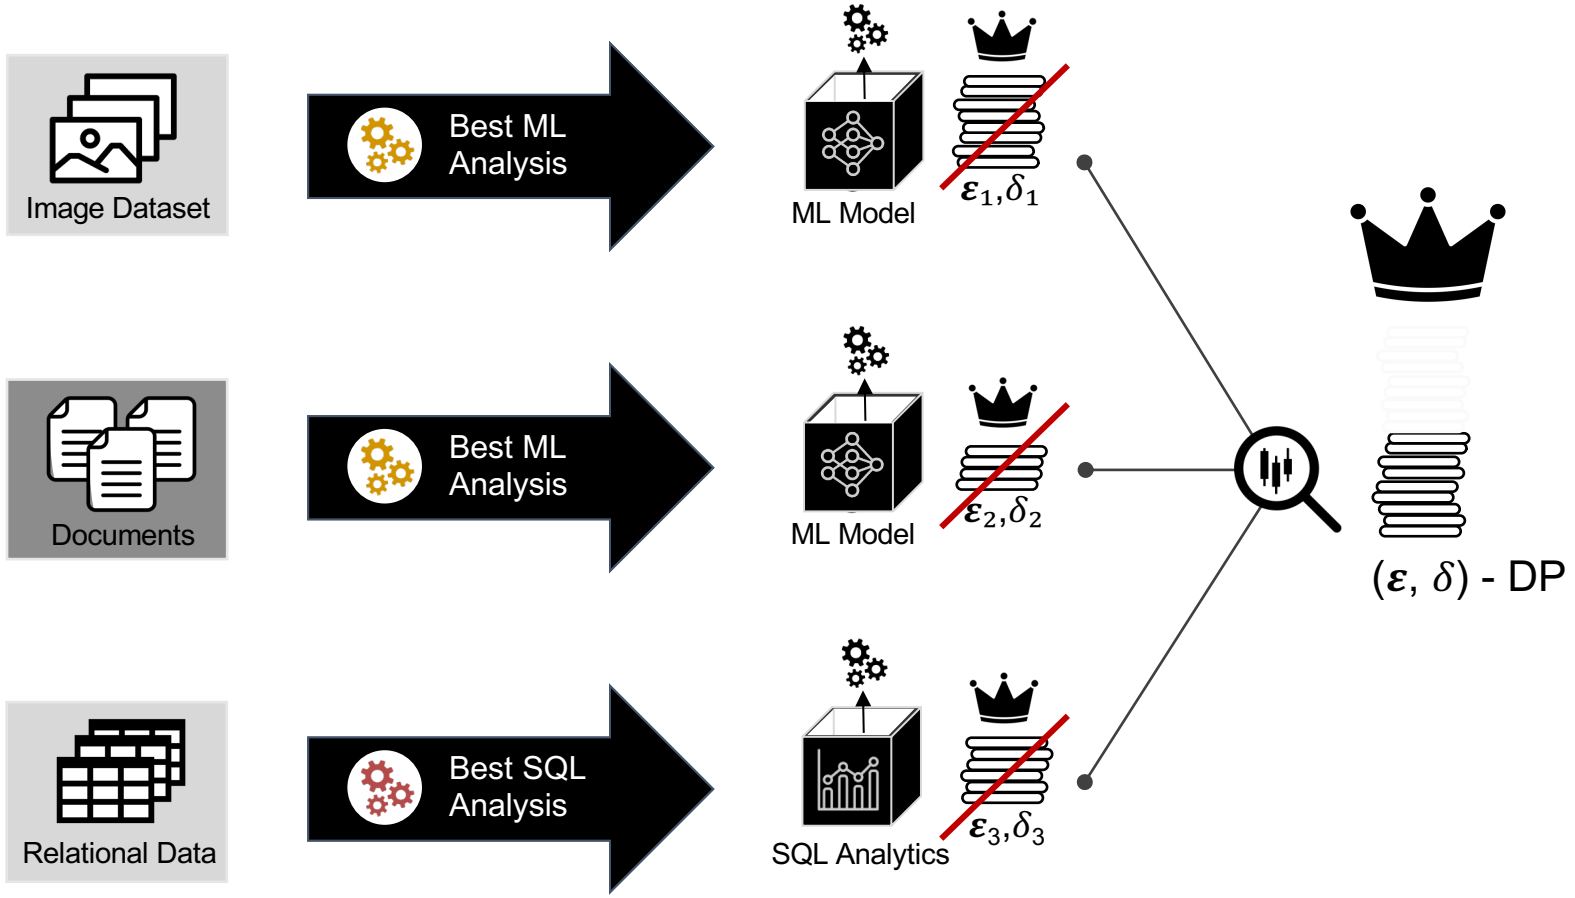

Moving Beyond Local Optima

DP Libraries: In a Nutshell

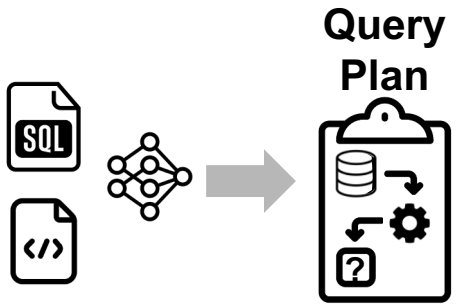

DP Libraries: In a Nutshell

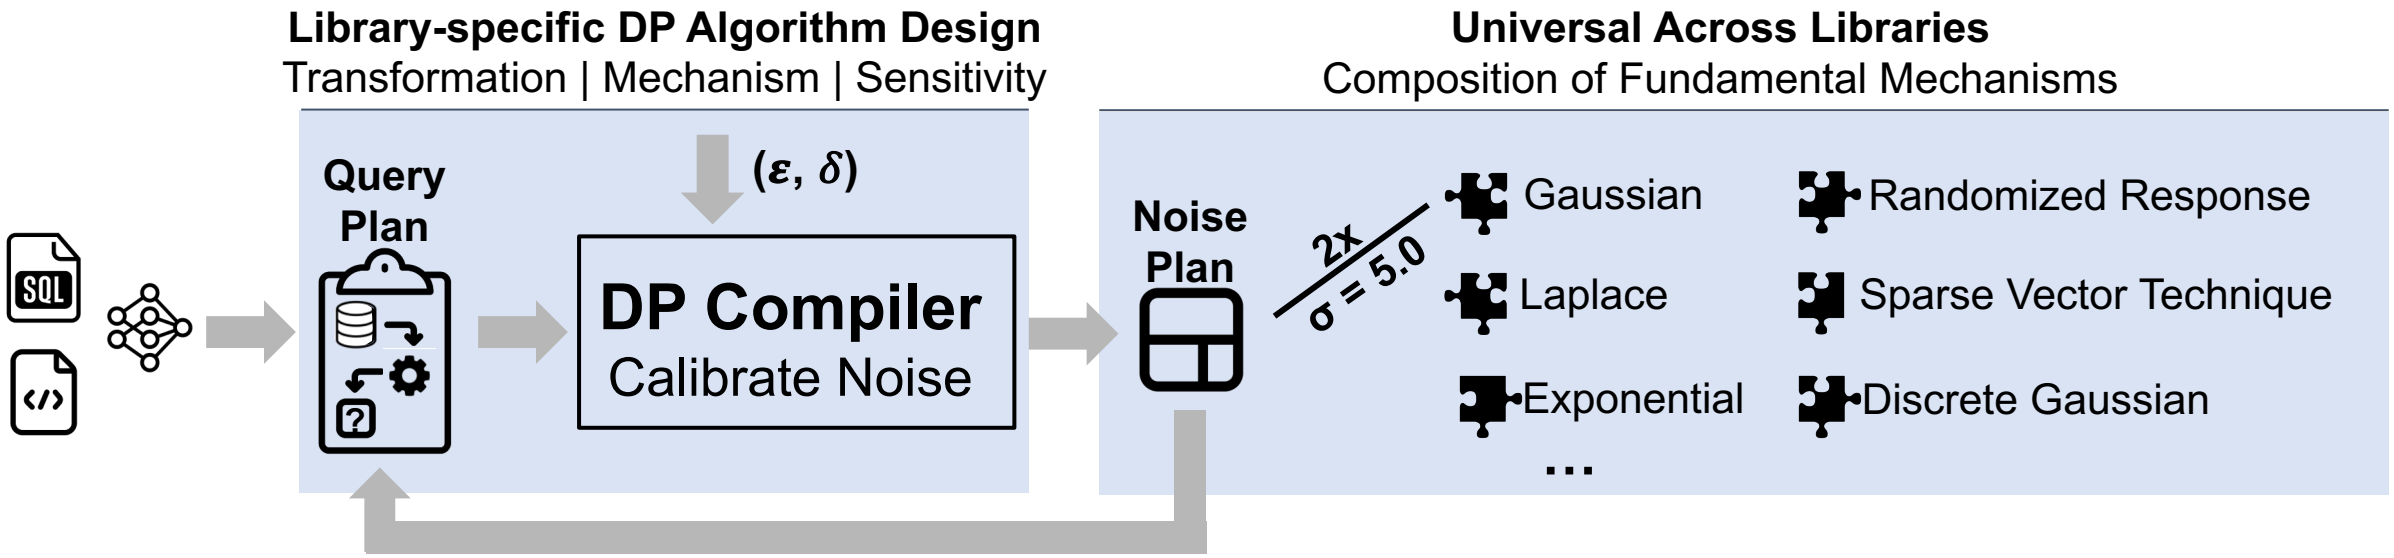

DP Libraries: In a Nutshell

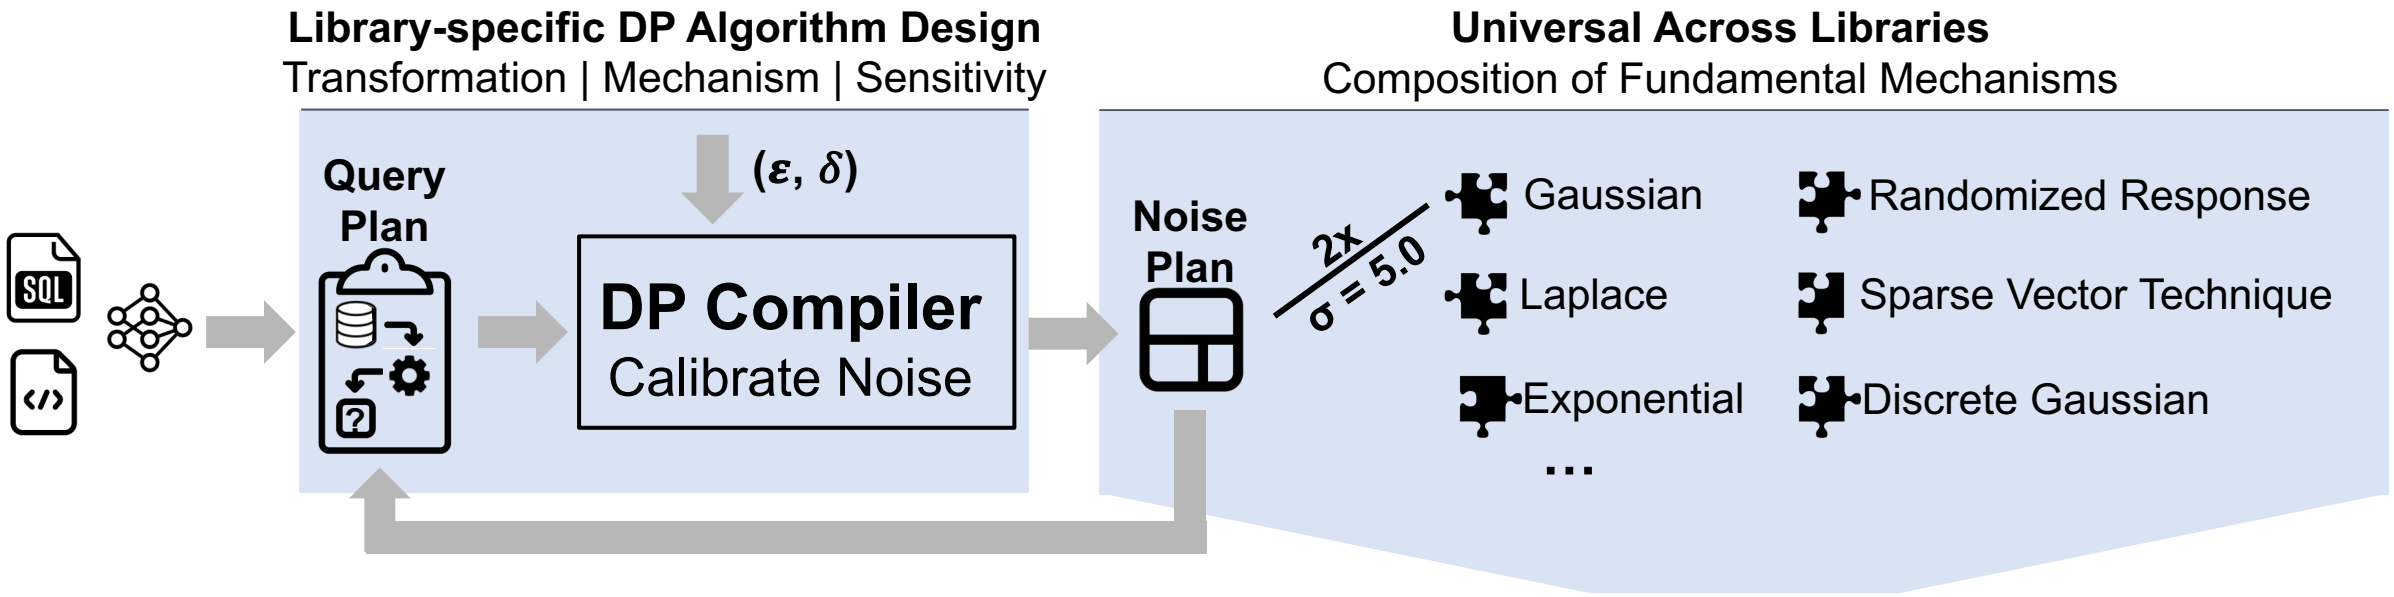

If we can compose all fundamental mechanisms, we can support a variety of heterogeneous libraries through a unified noise plan.

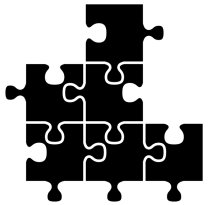

Composition of Fundamental Mechanisms

Unifying the Application Layer

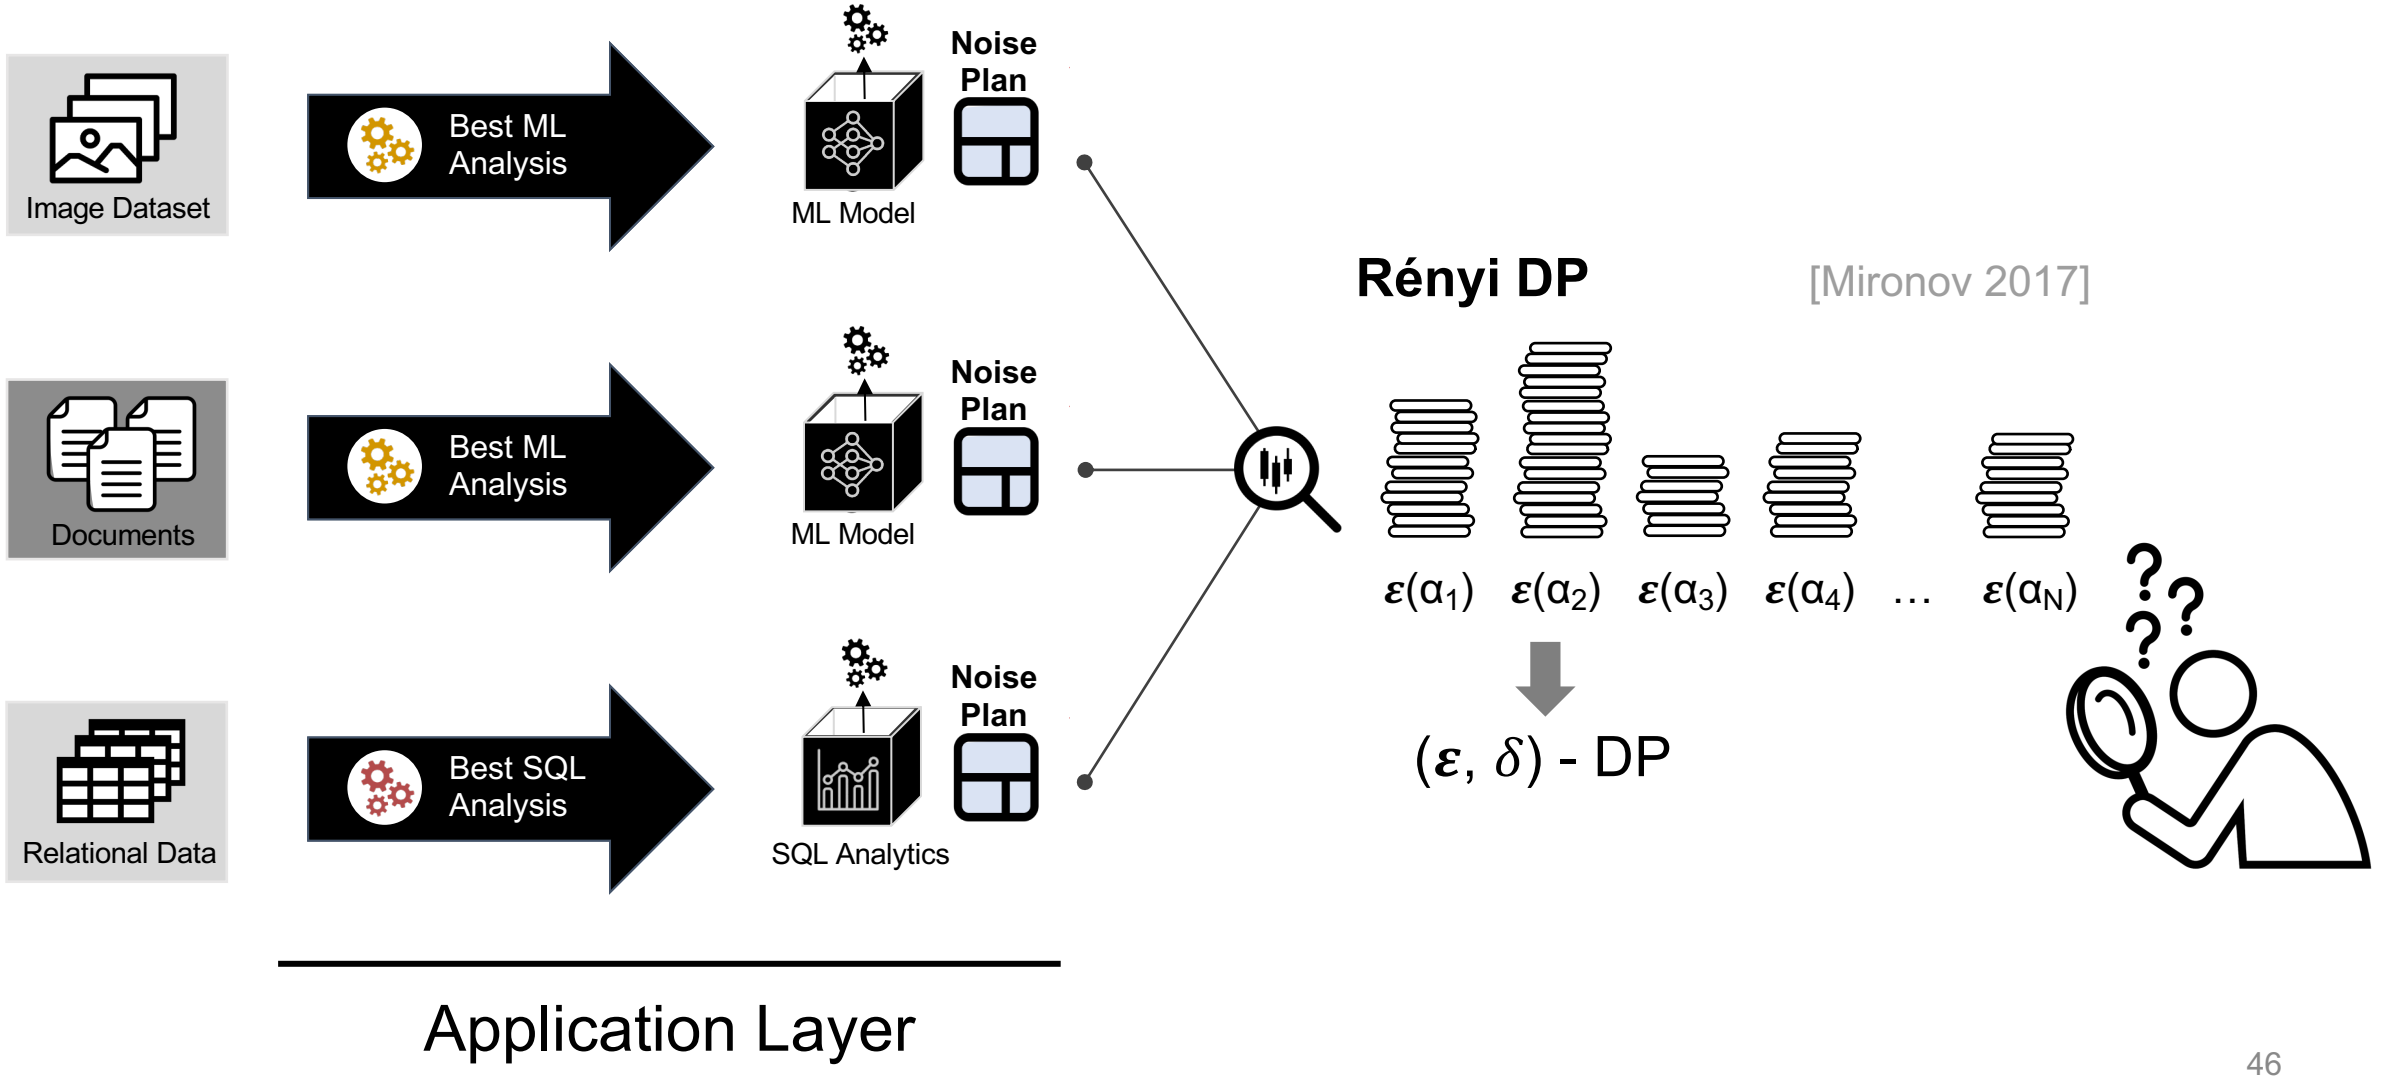

Unifying the Application Layer

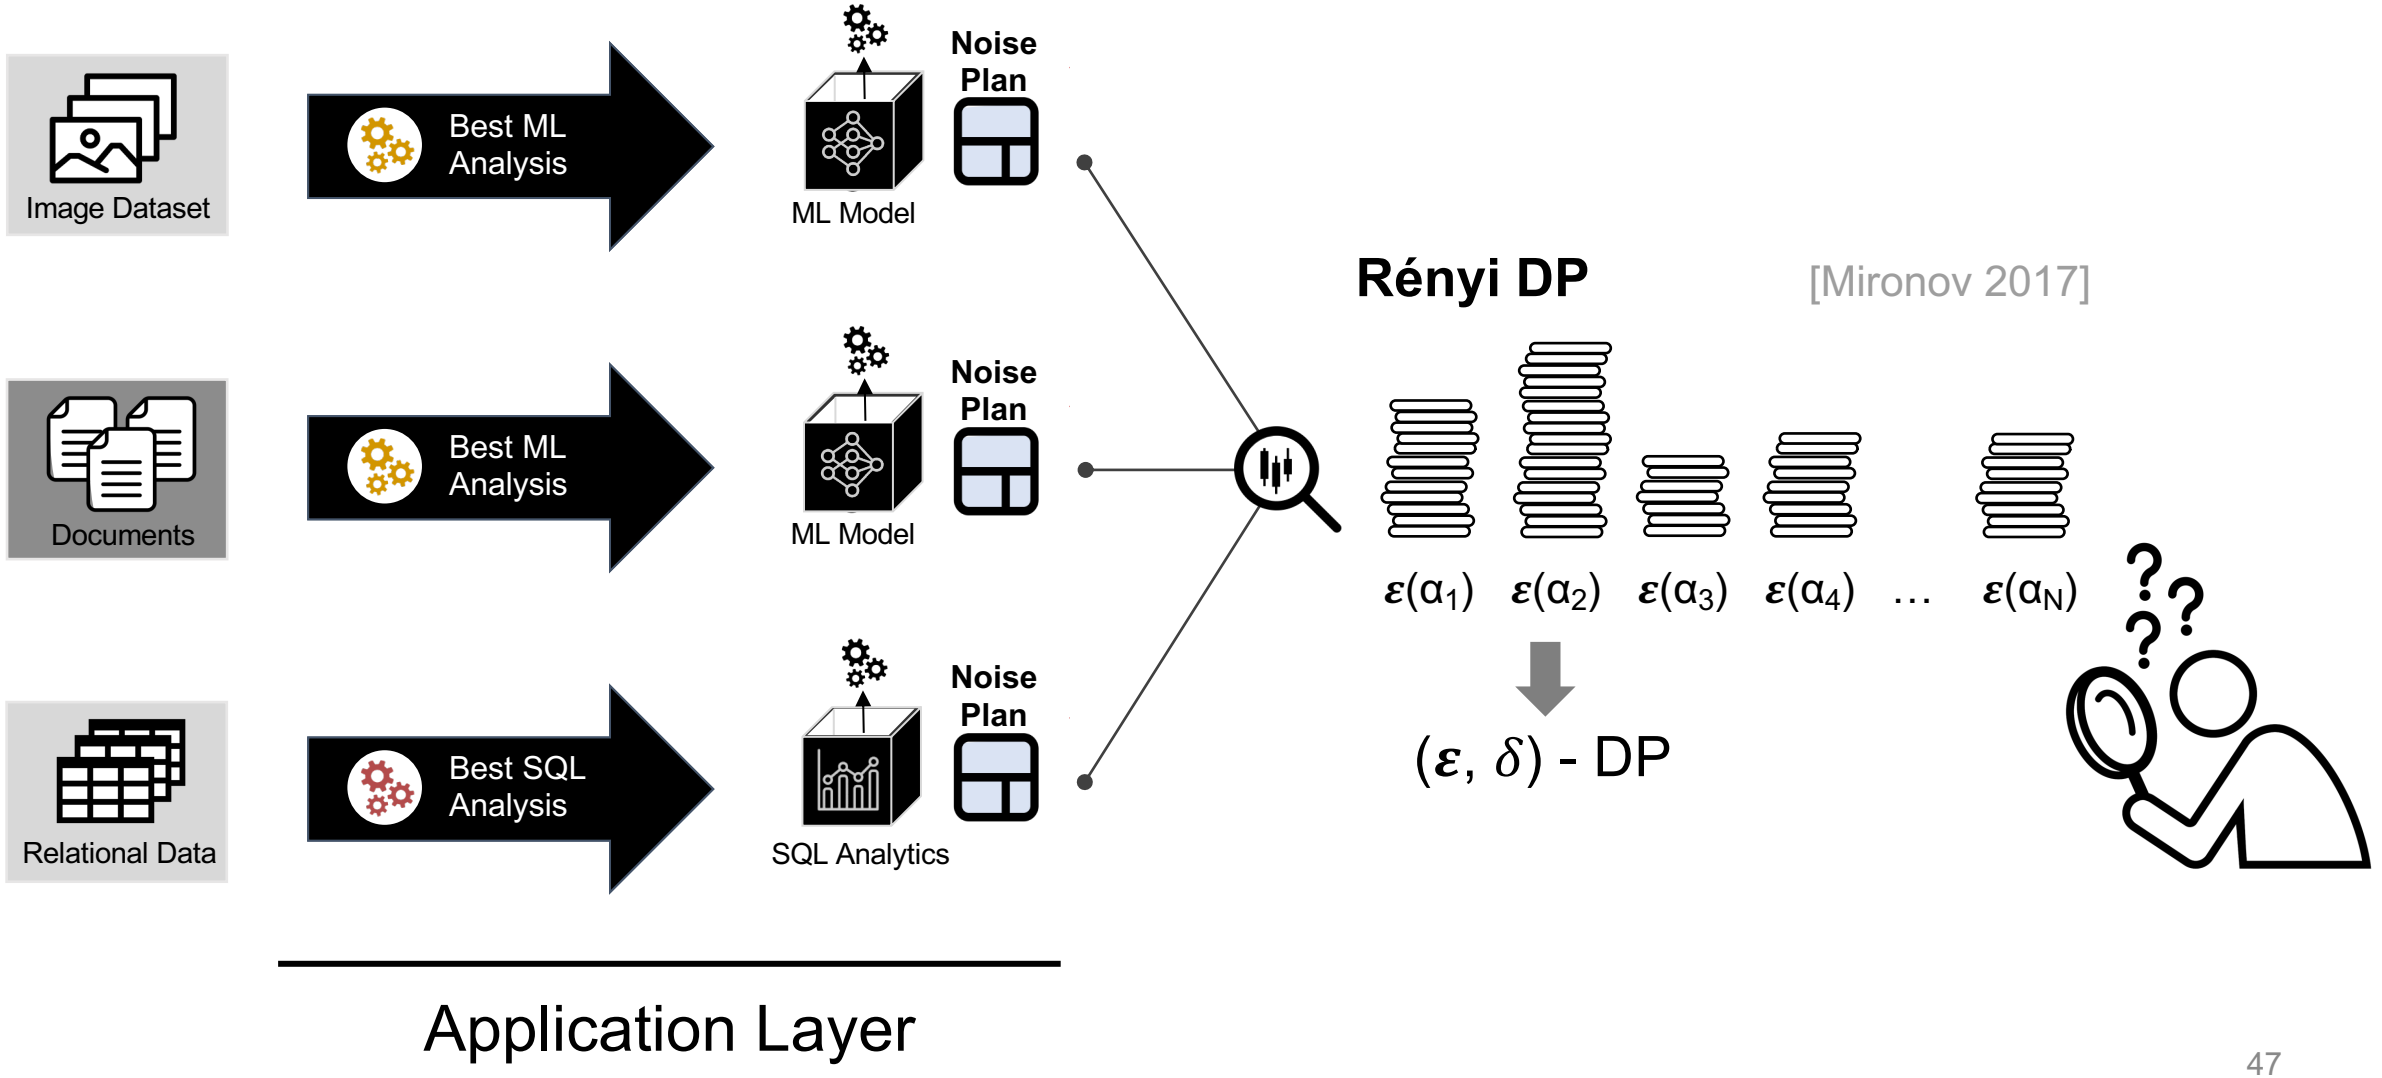

Unifying the Application Layer

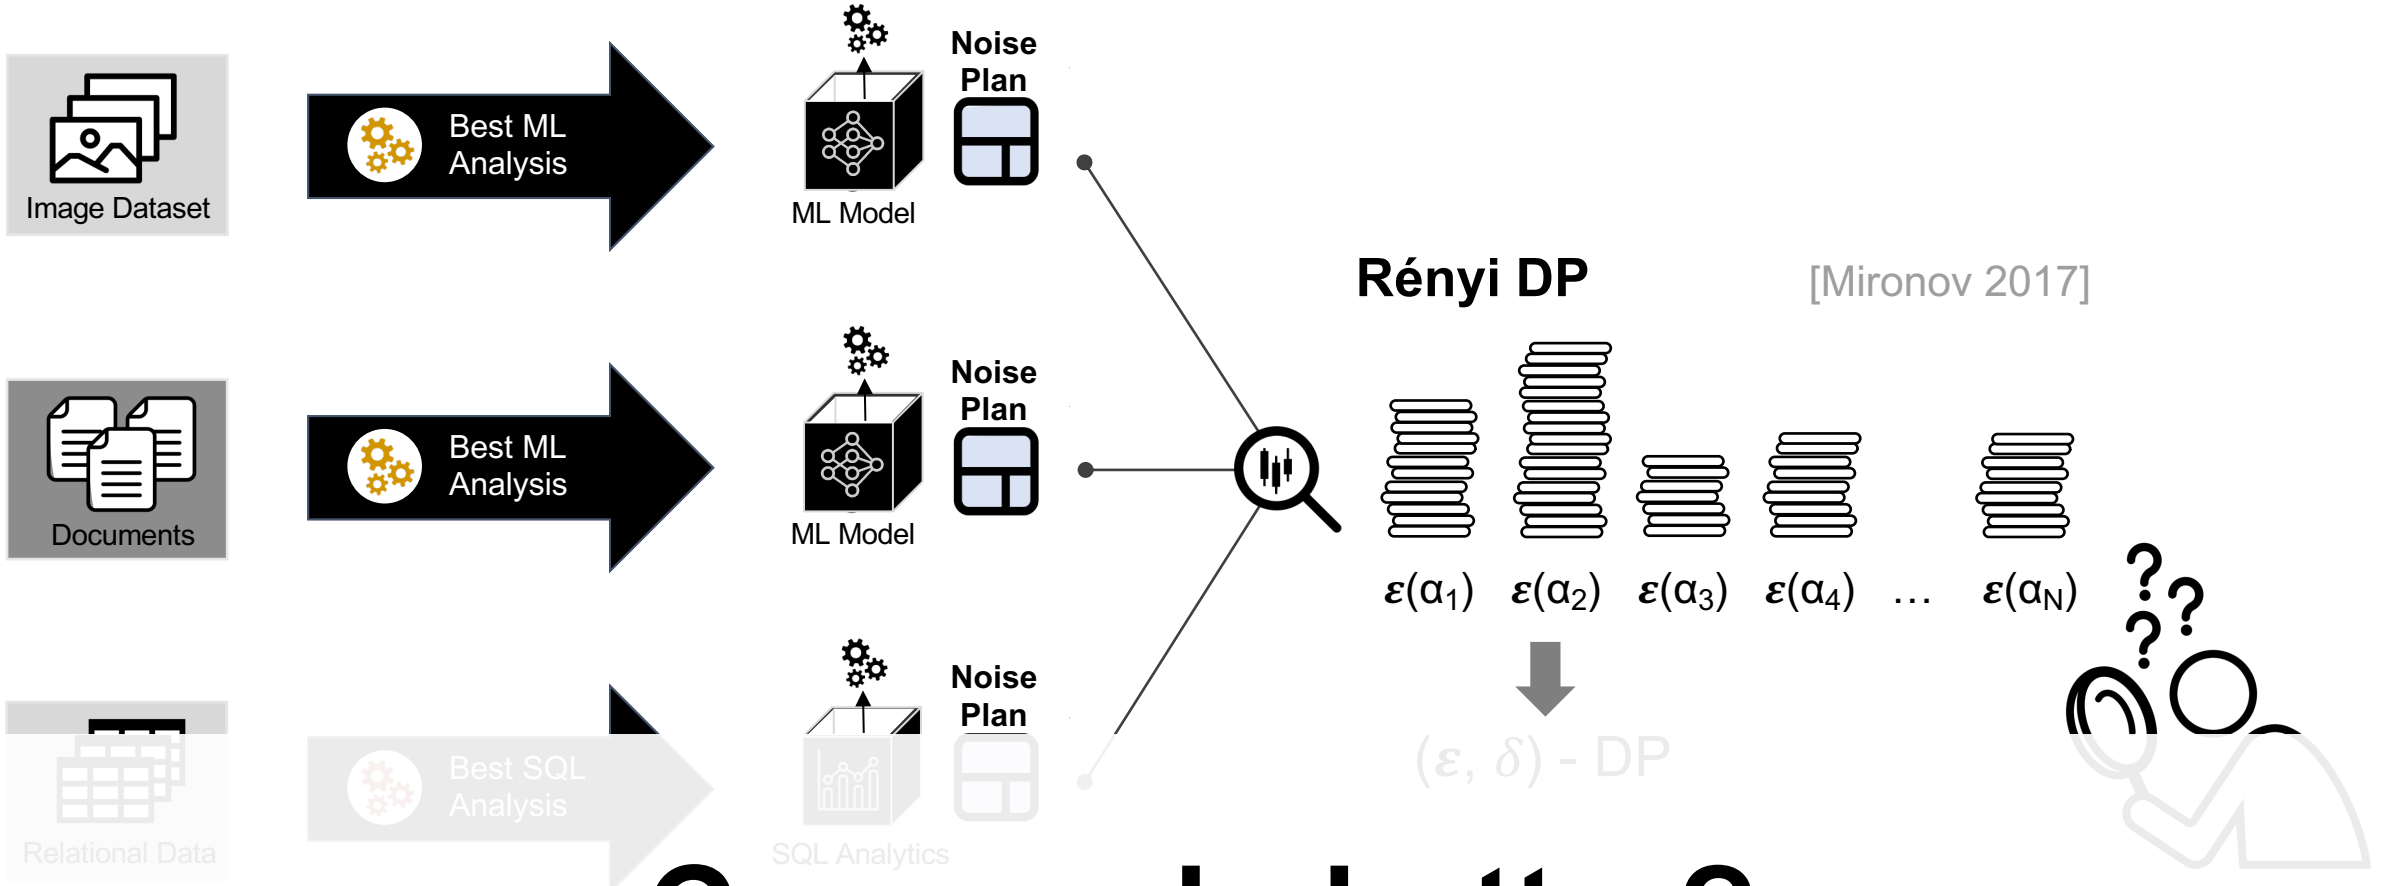

Can we do better?

Application Layer

Benefiting from Data Access Patterns

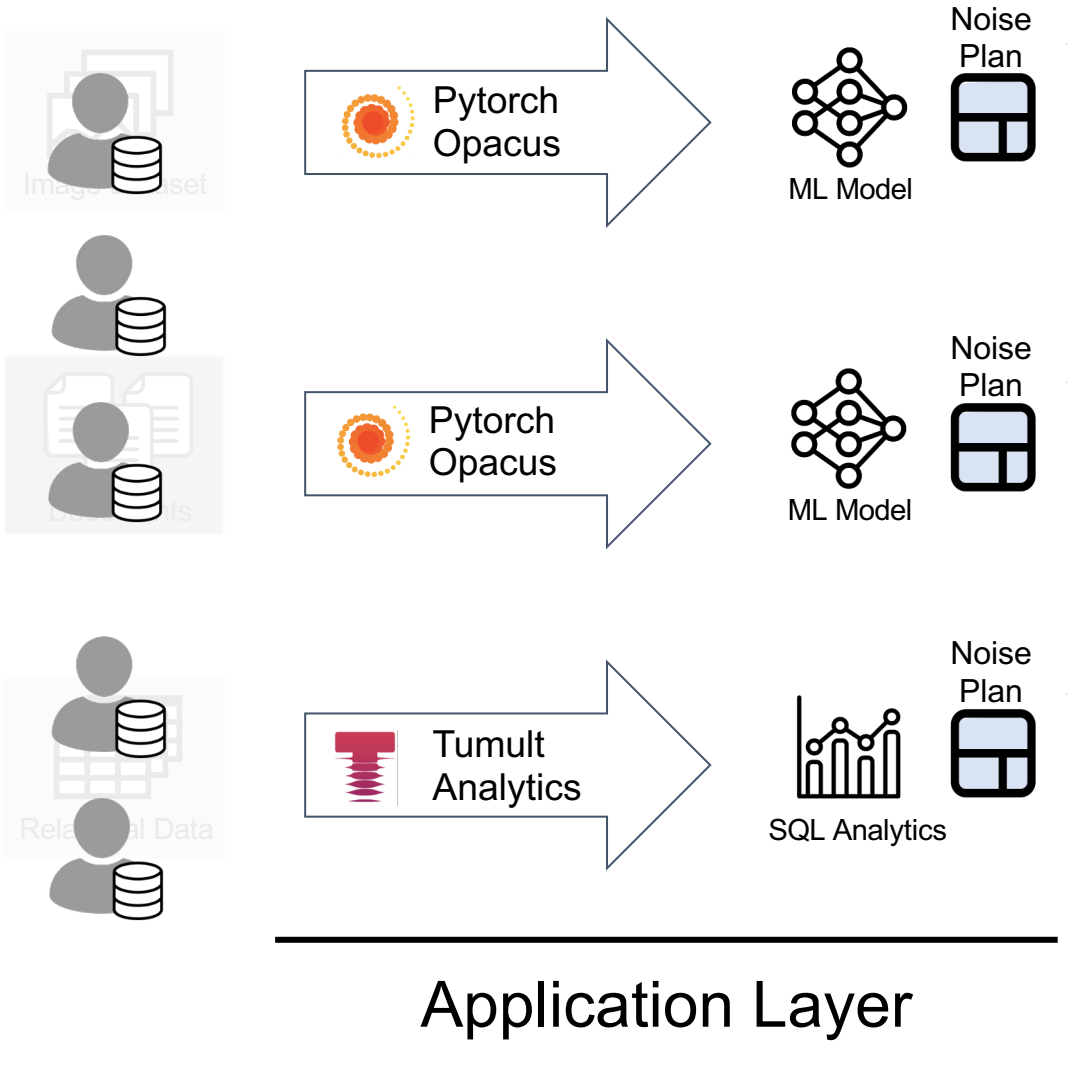

Benefiting from Data Access Patterns

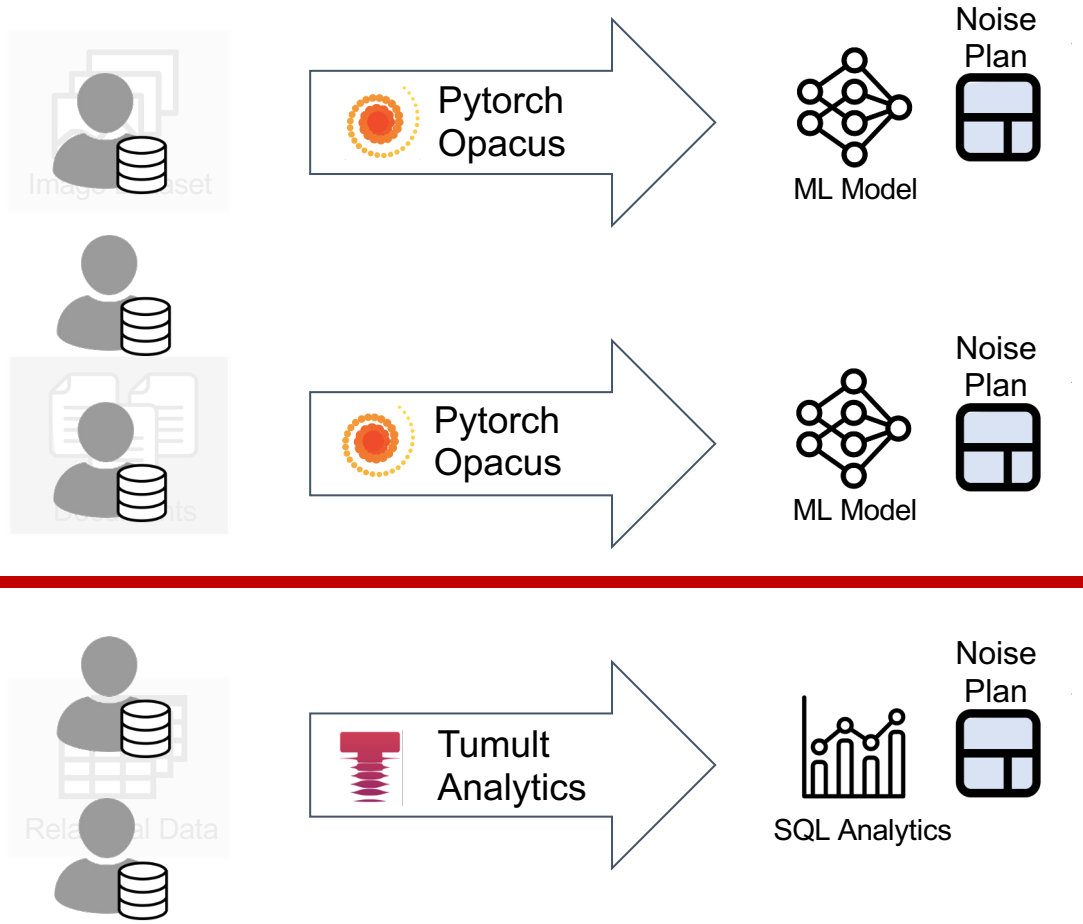

Benefiting from Data Access Patterns

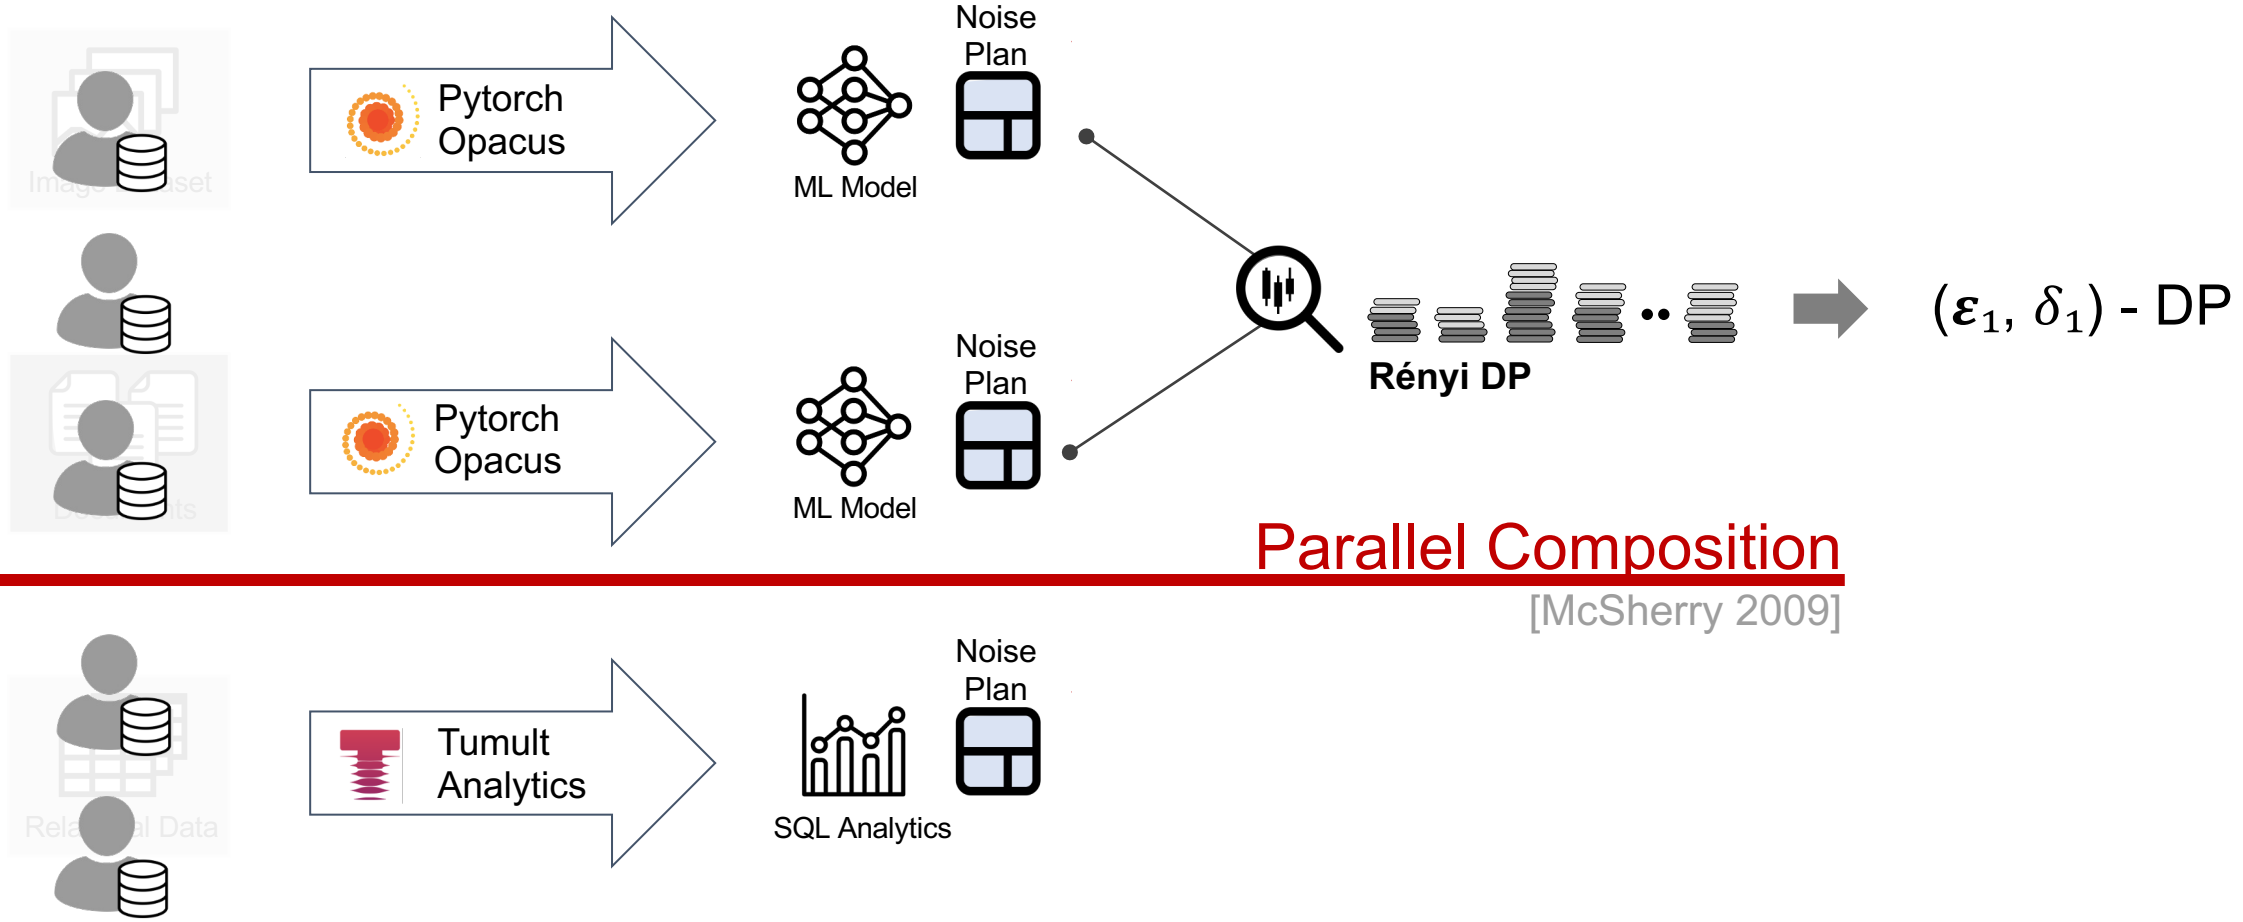

Benefiting from Data Access Patterns

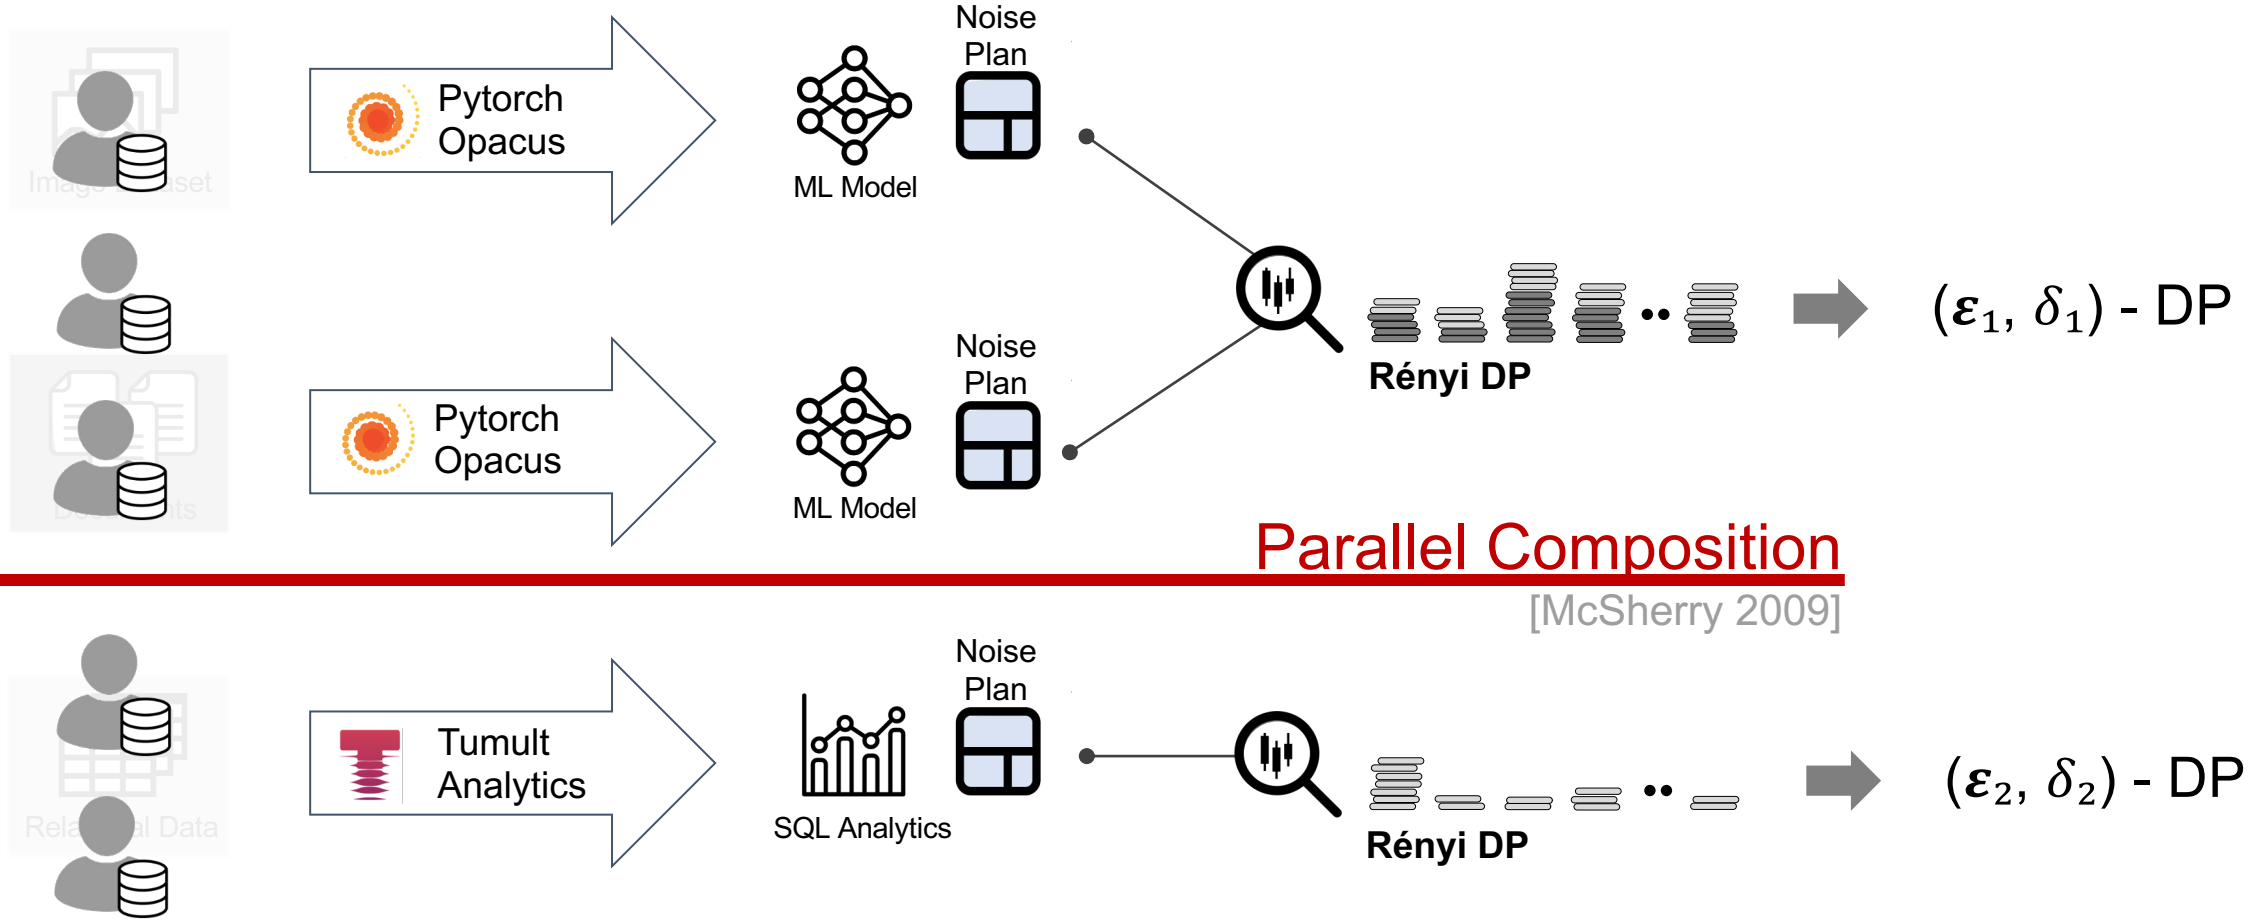

Benefiting from Data Access Patterns

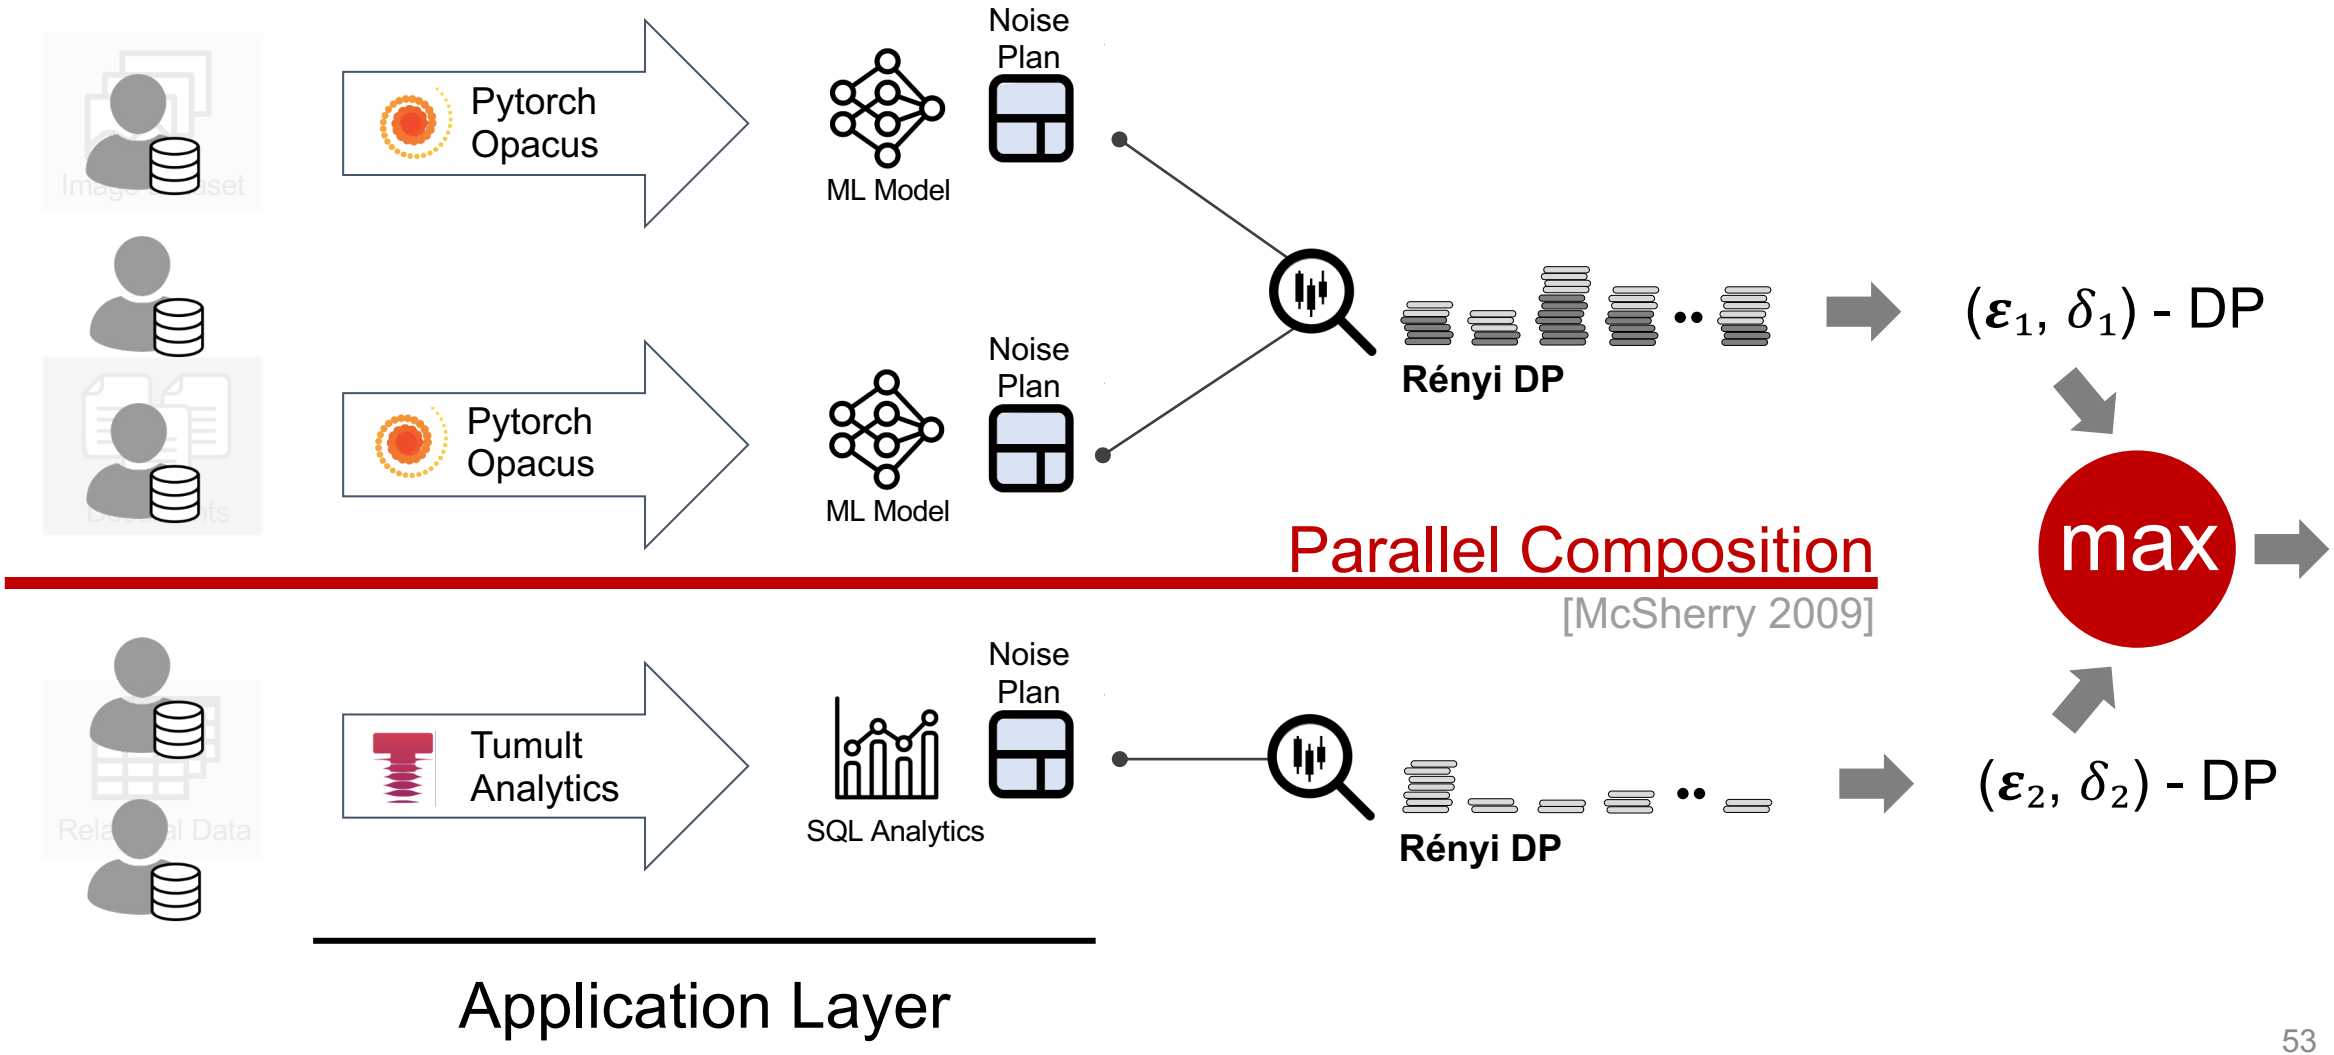

Benefiting from Data Access Patterns

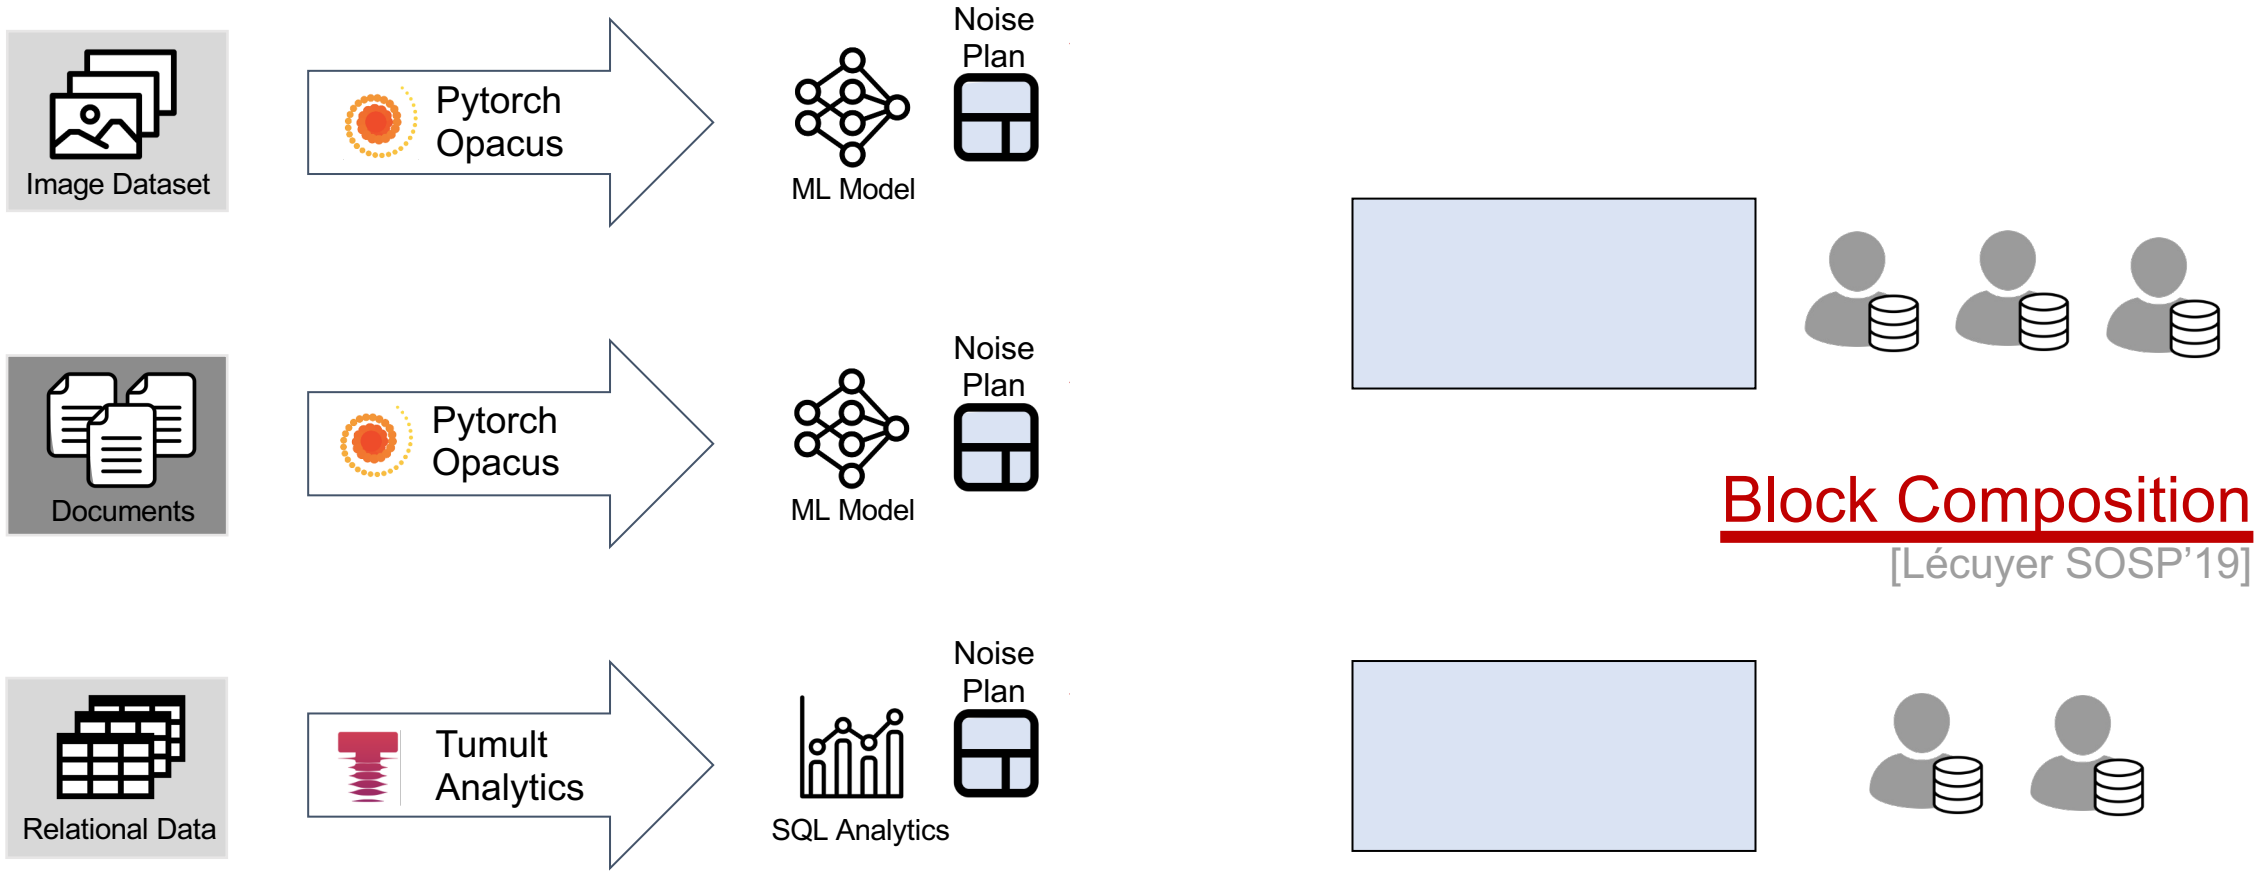

Benefiting from Data Access Patterns

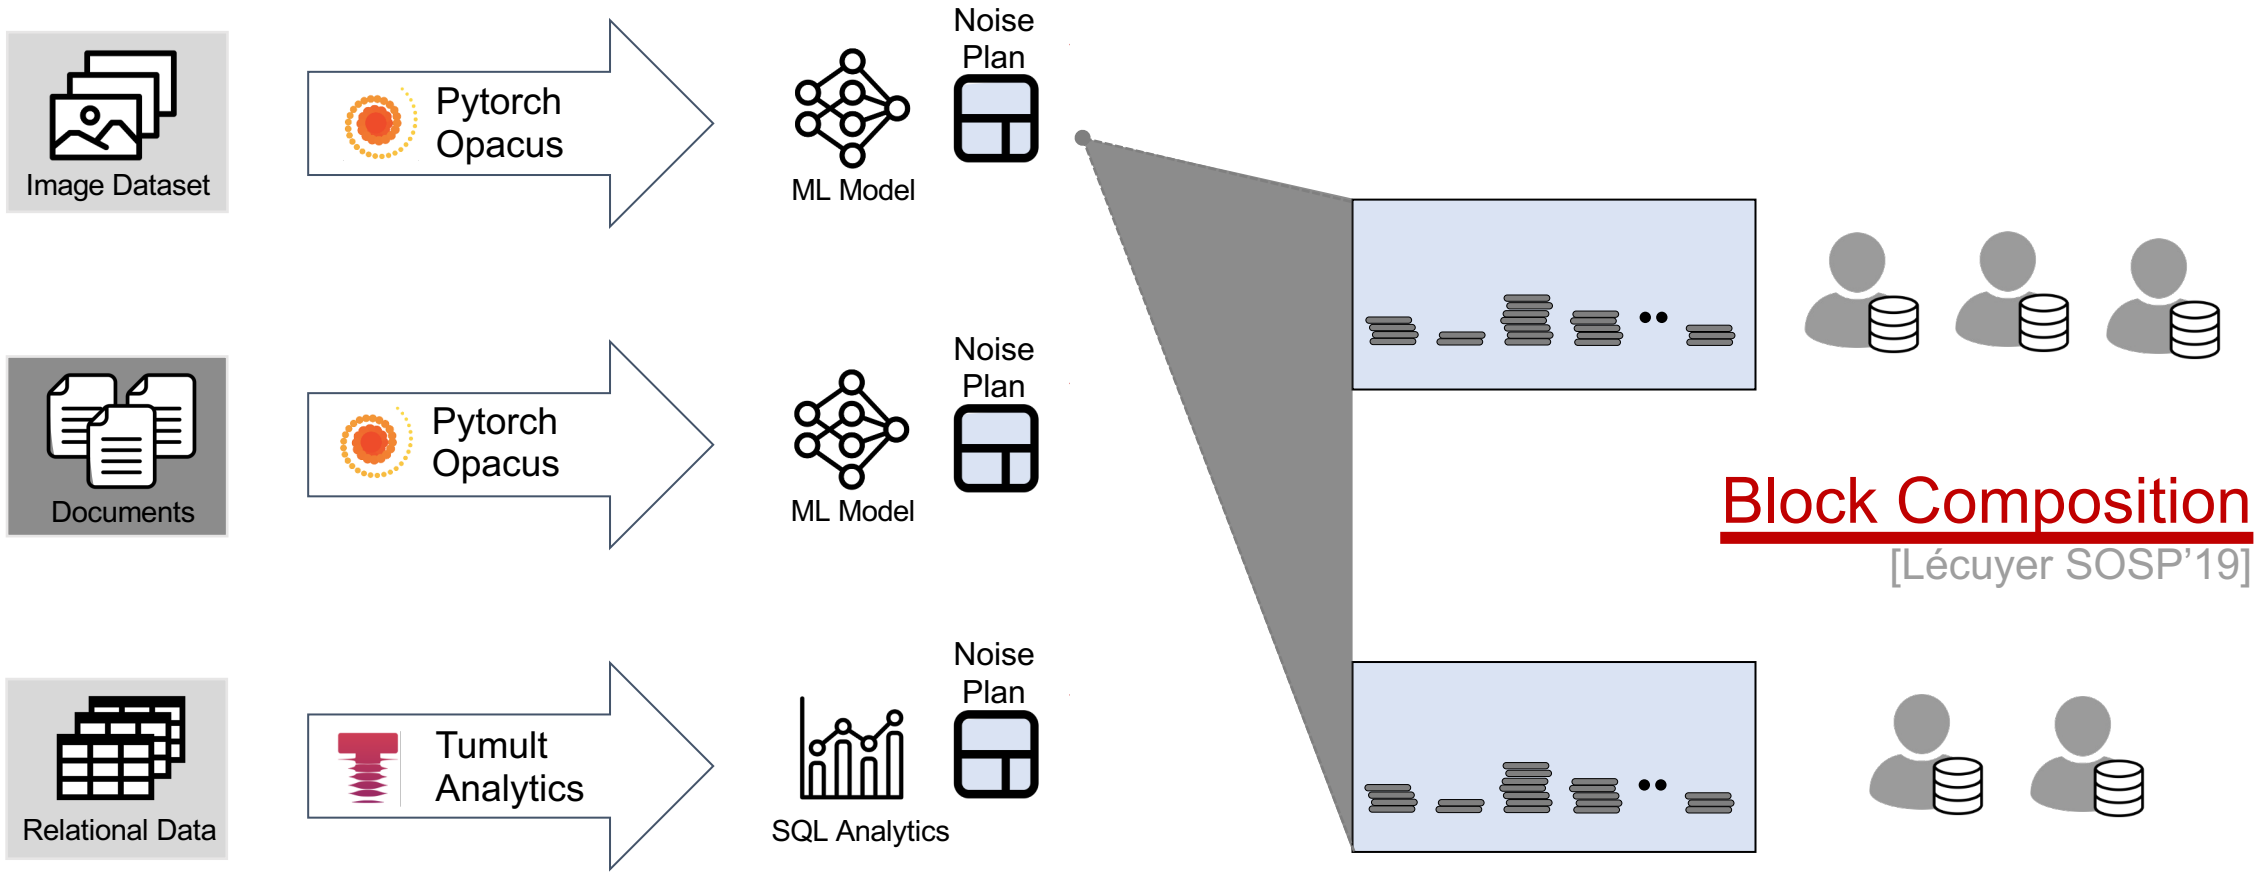

Application Layer

Total: max(□,□)

Benefiting from Data Access Patterns

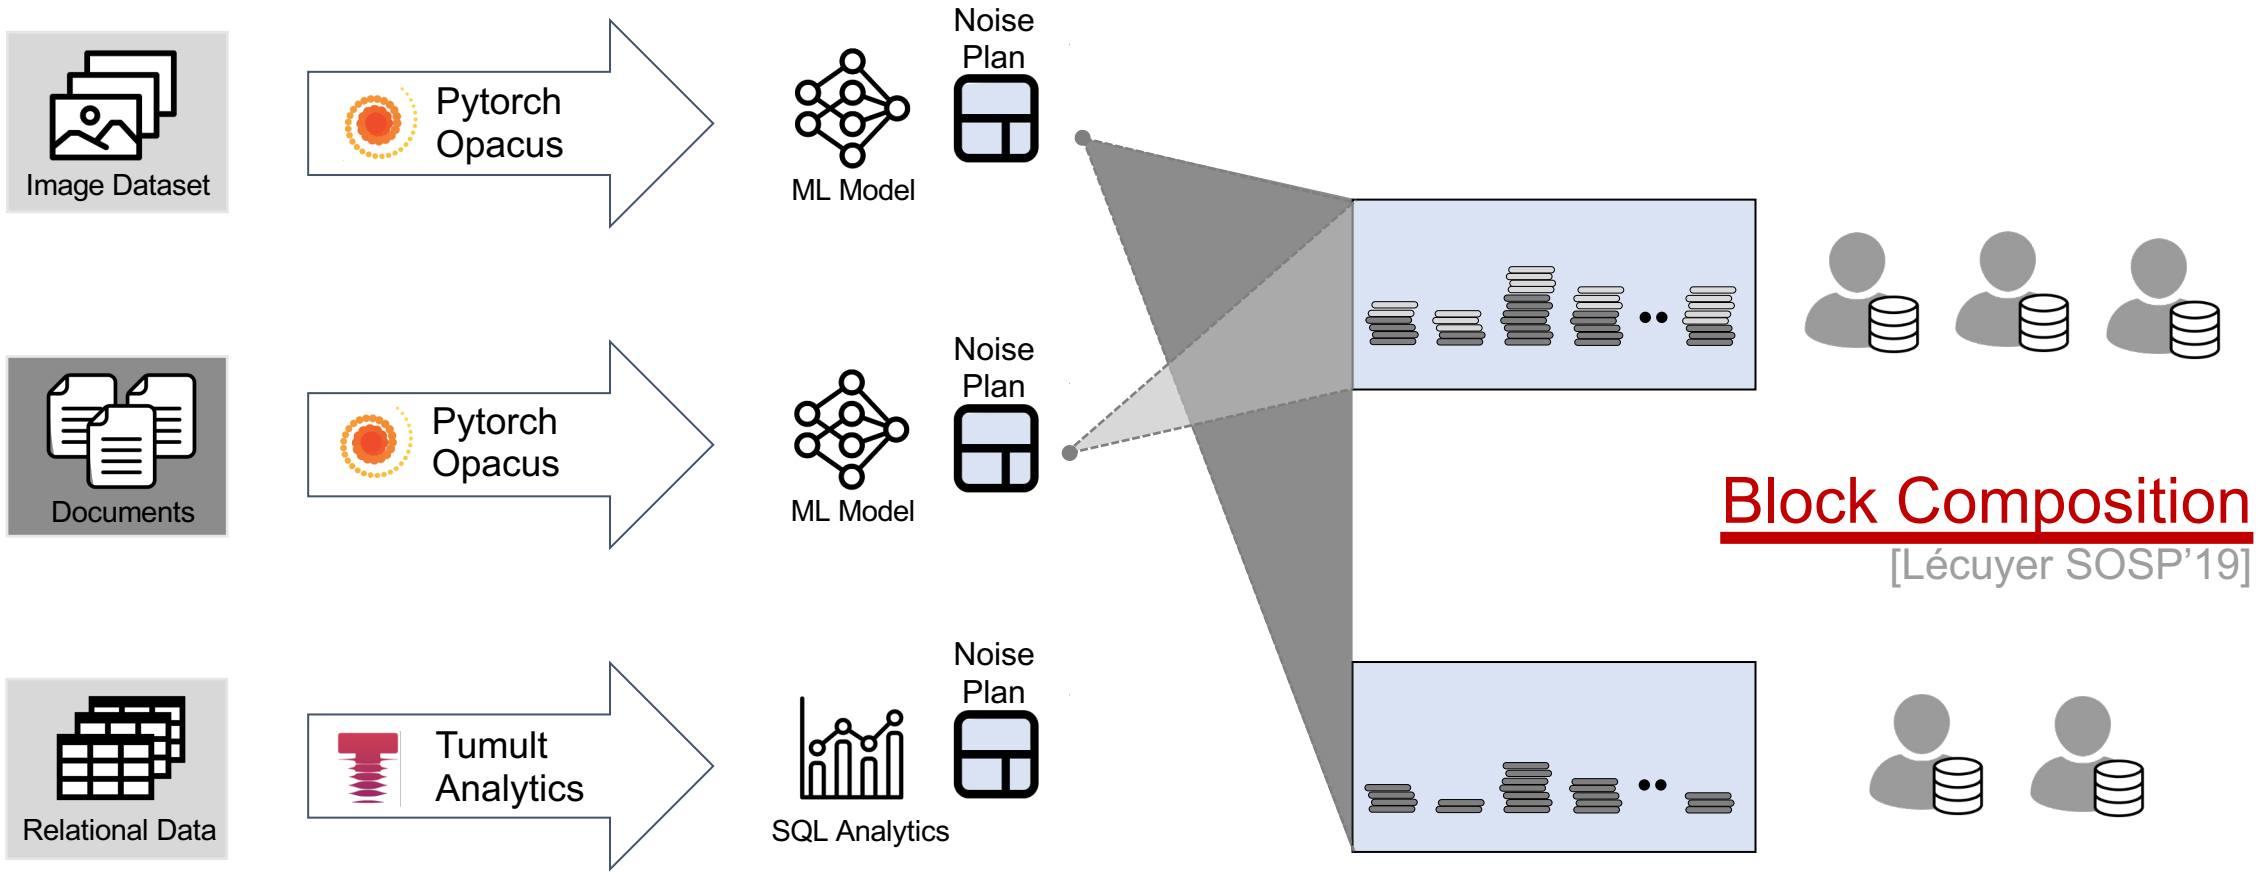

Application Layer

Total: max(□,□)

Benefiting from Data Access Patterns

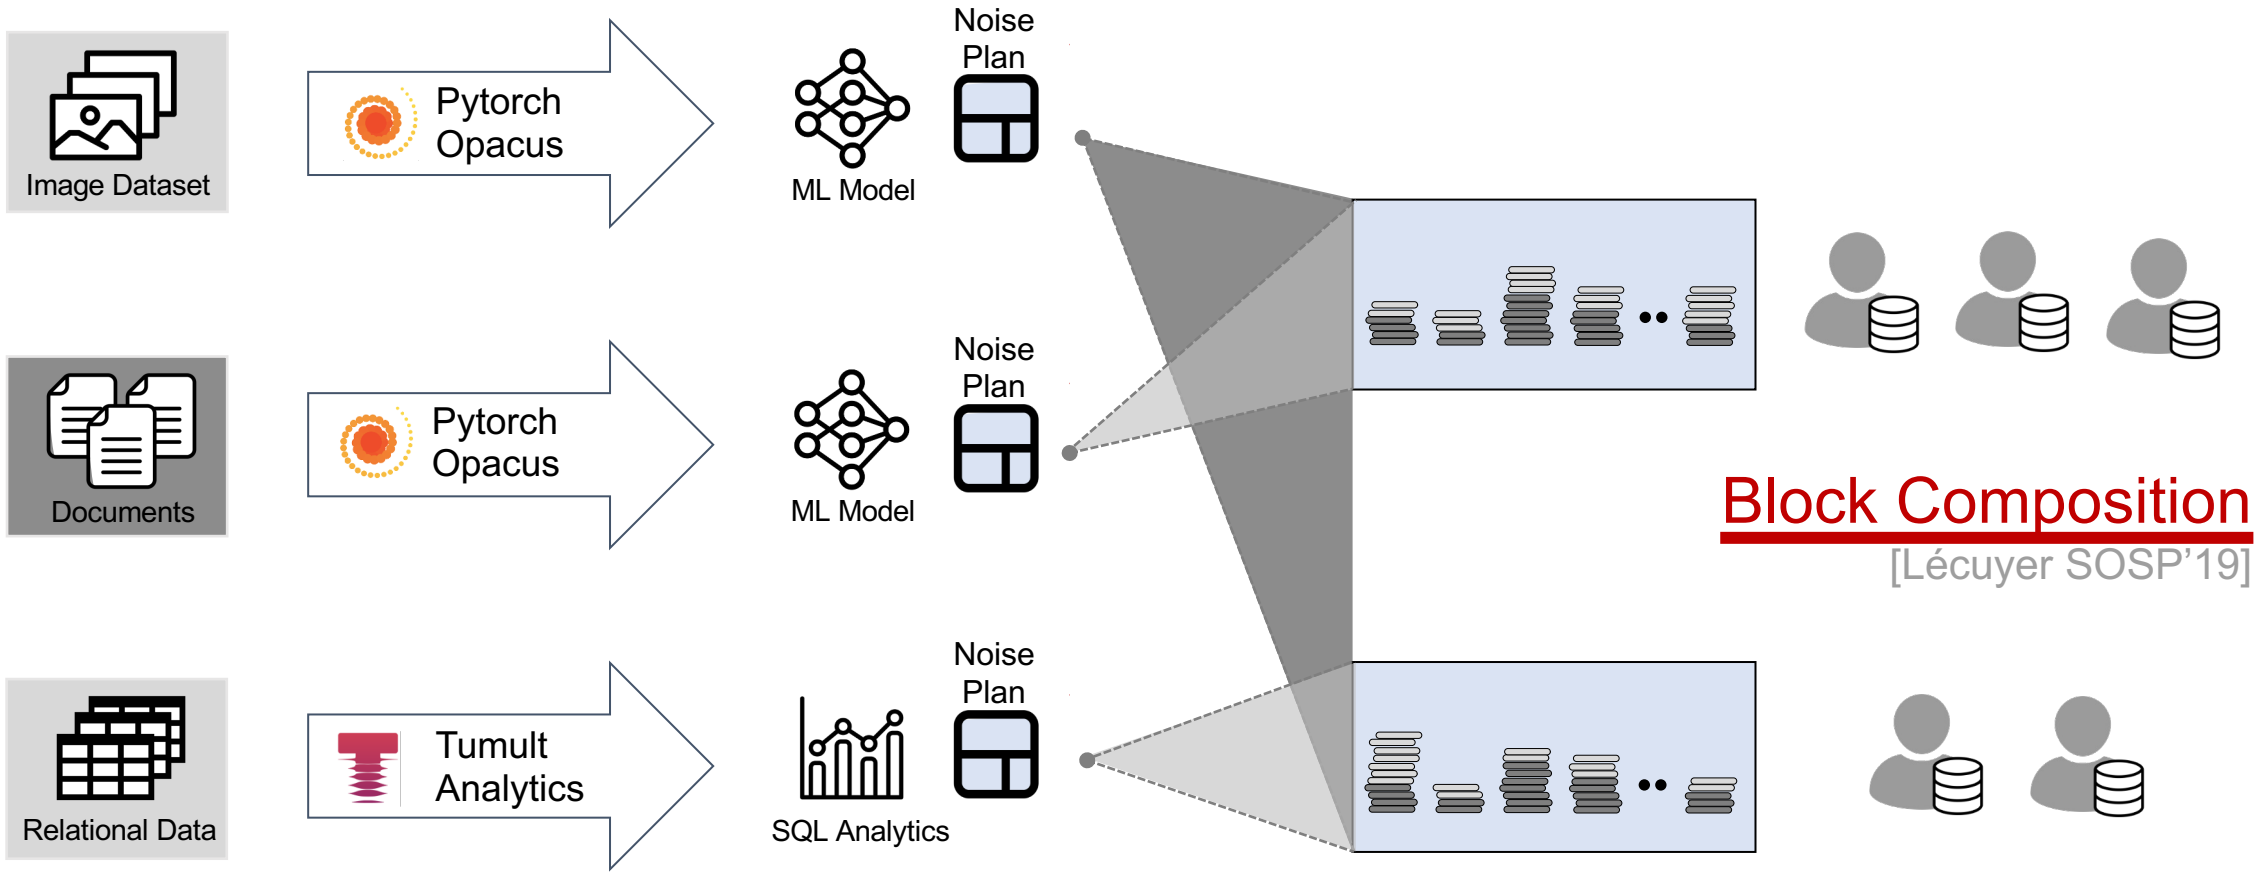

Application Layer

Total: max(□,□)

Benefiting from Data Access Patterns

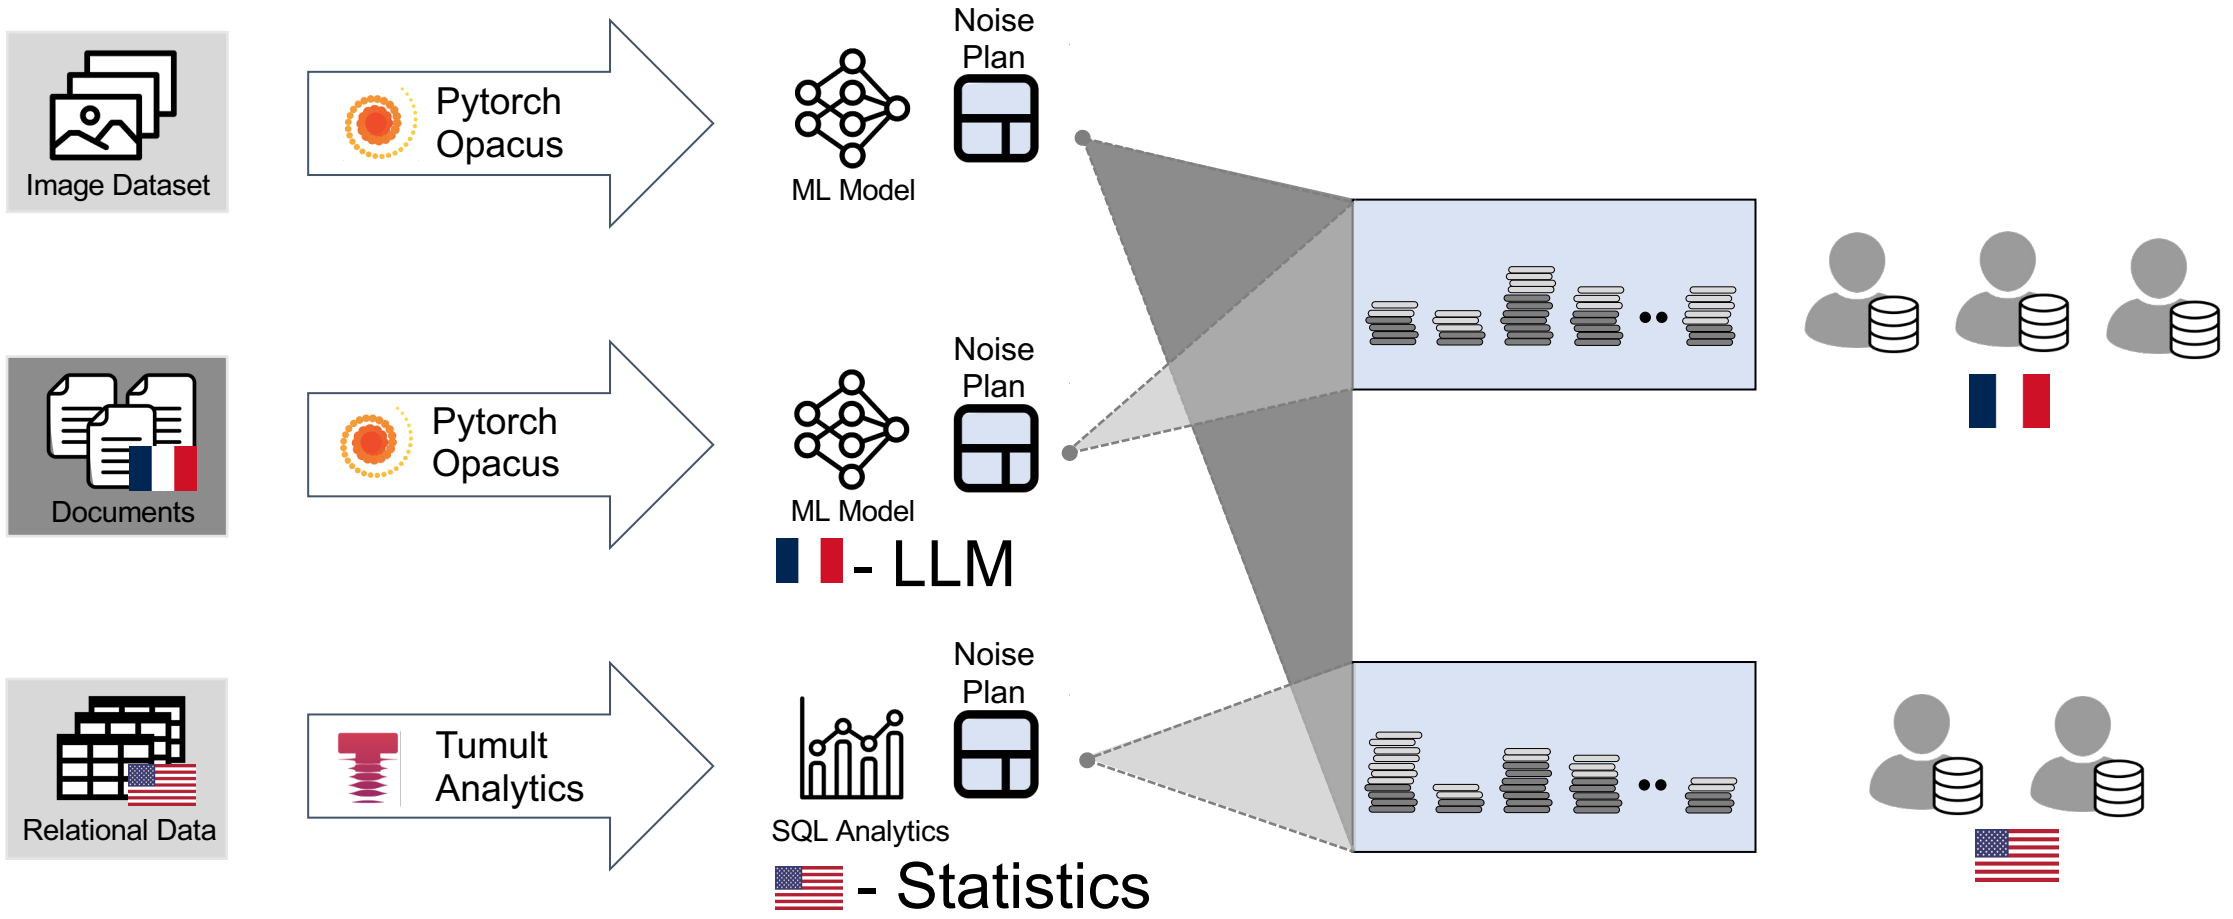

Application Layer

Total: max(ε_{FR}, ε_{US})

Fine-grained Privacy Analysis

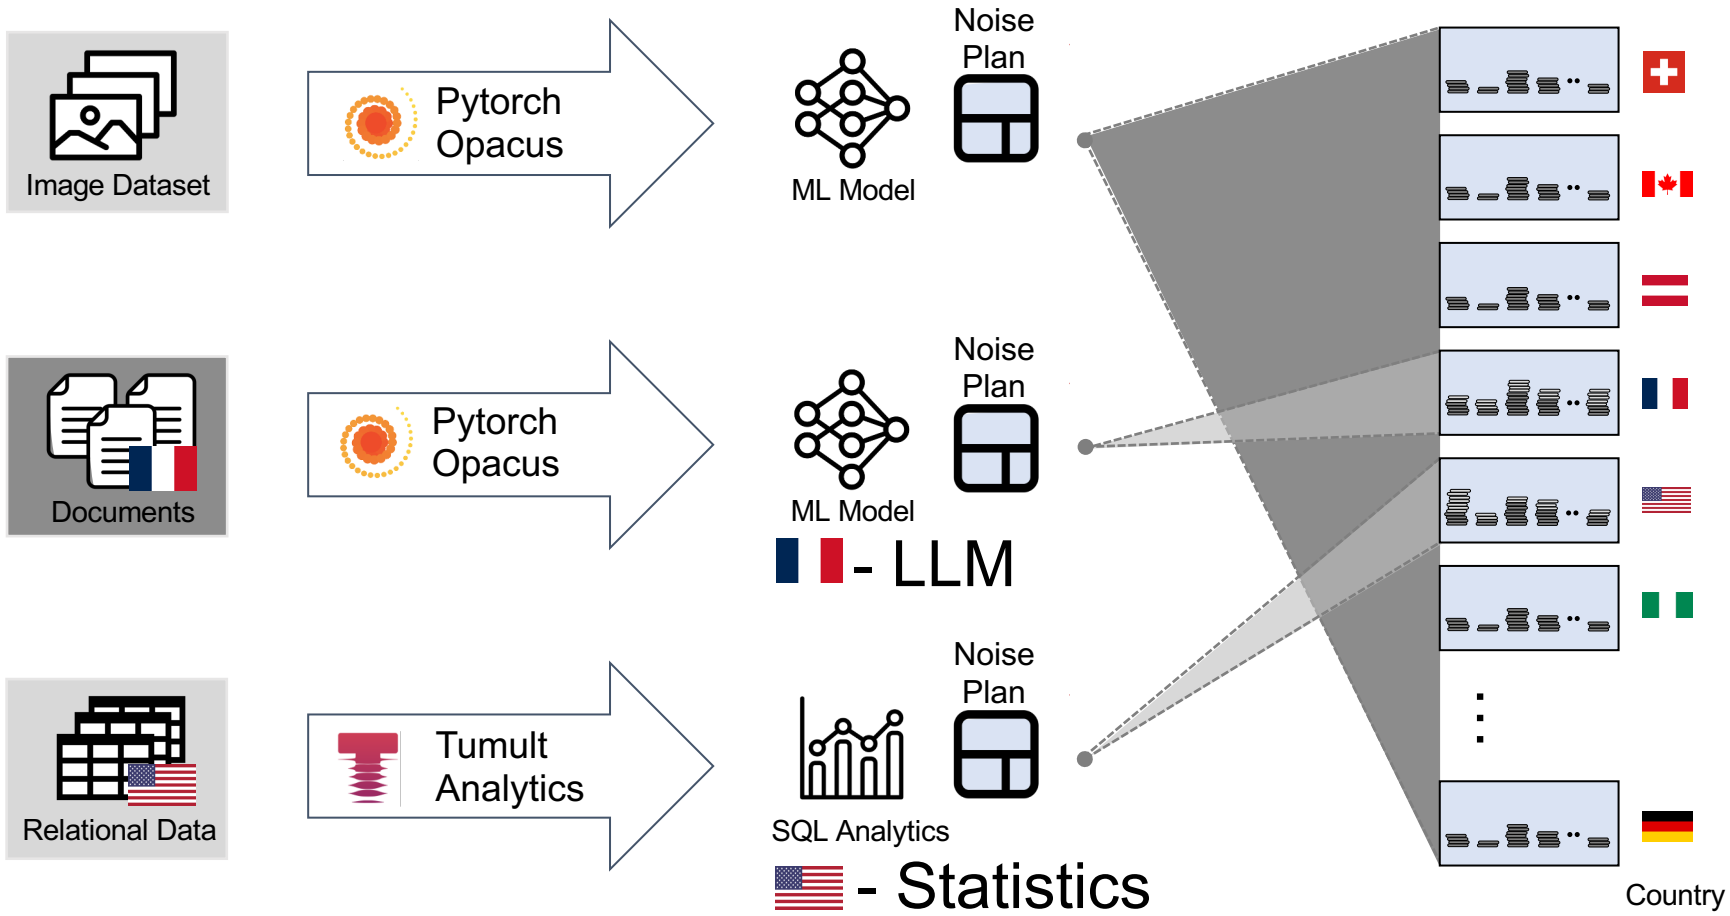

Application Layer

Partitioning Attributes

Schema must be known in advance

Fine-grained Privacy Analysis

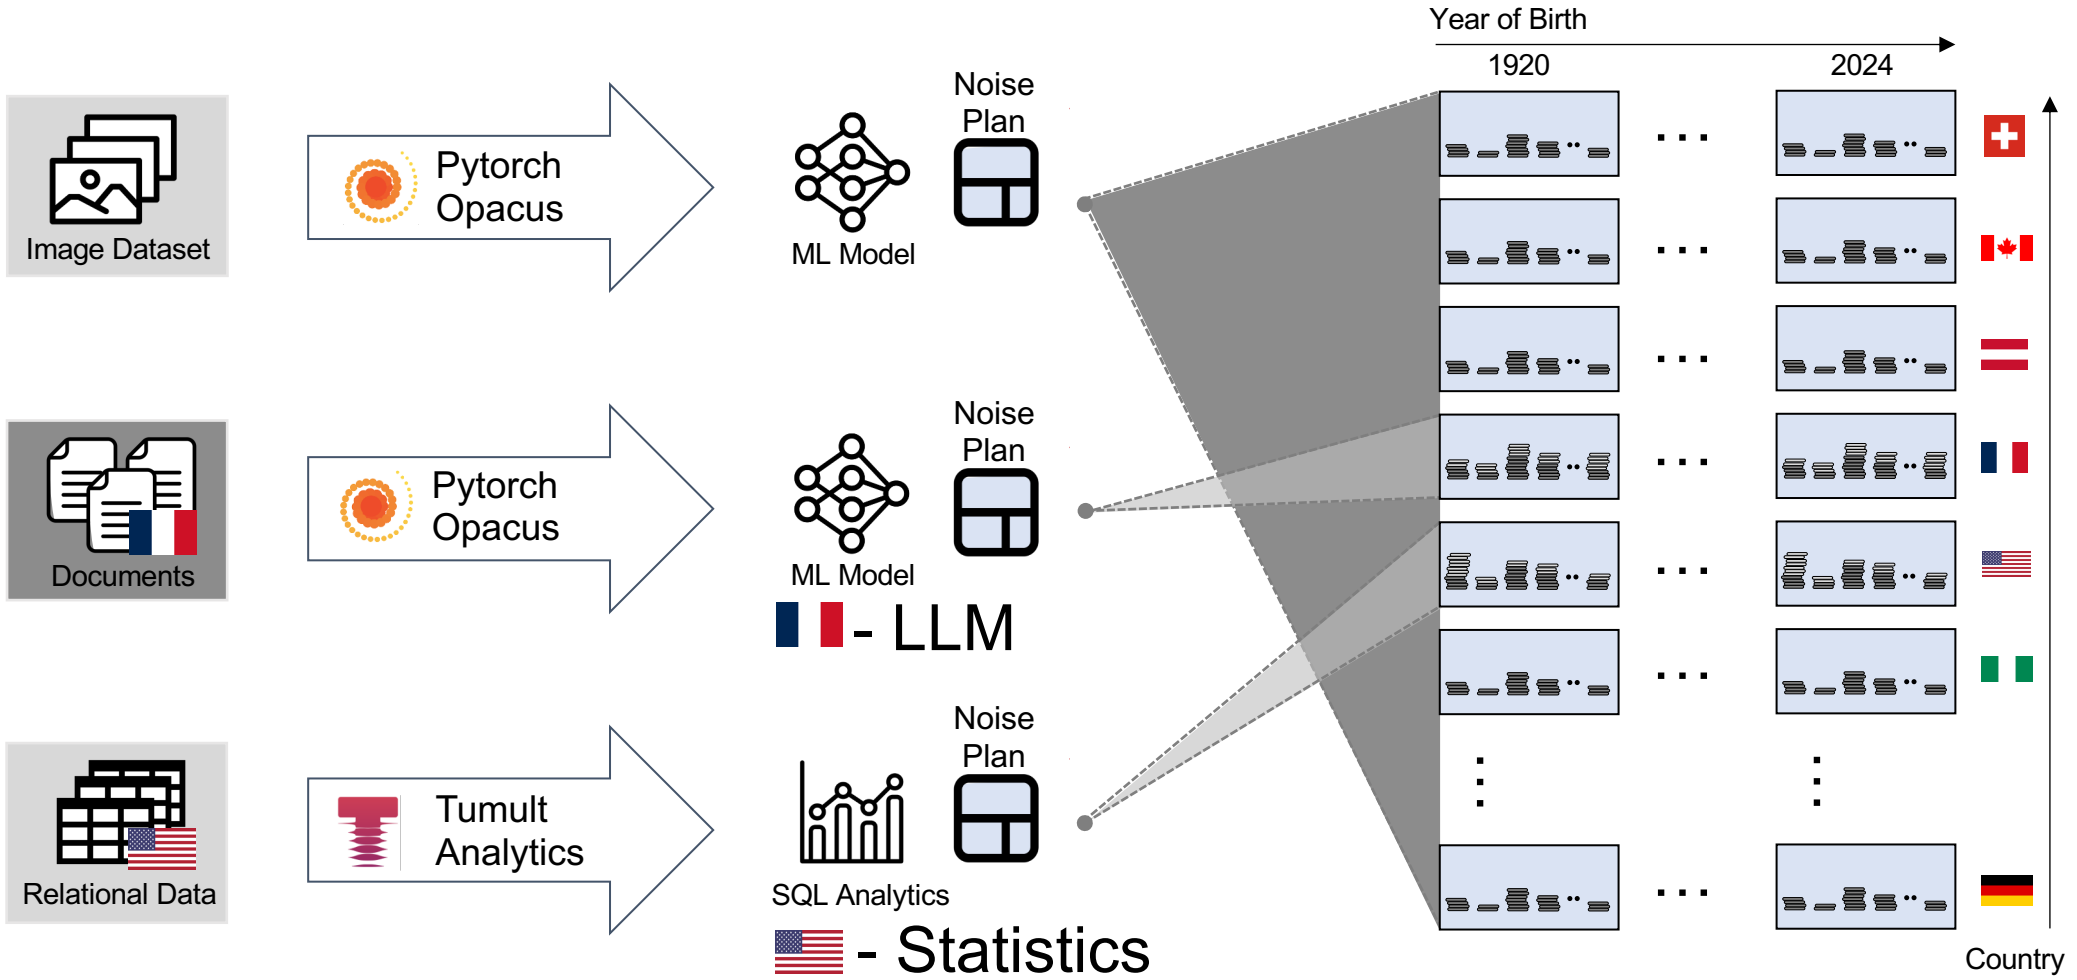

Application Layer

Partitioning Attributes

Schema must be known in advance

Fine-grained Privacy Analysis

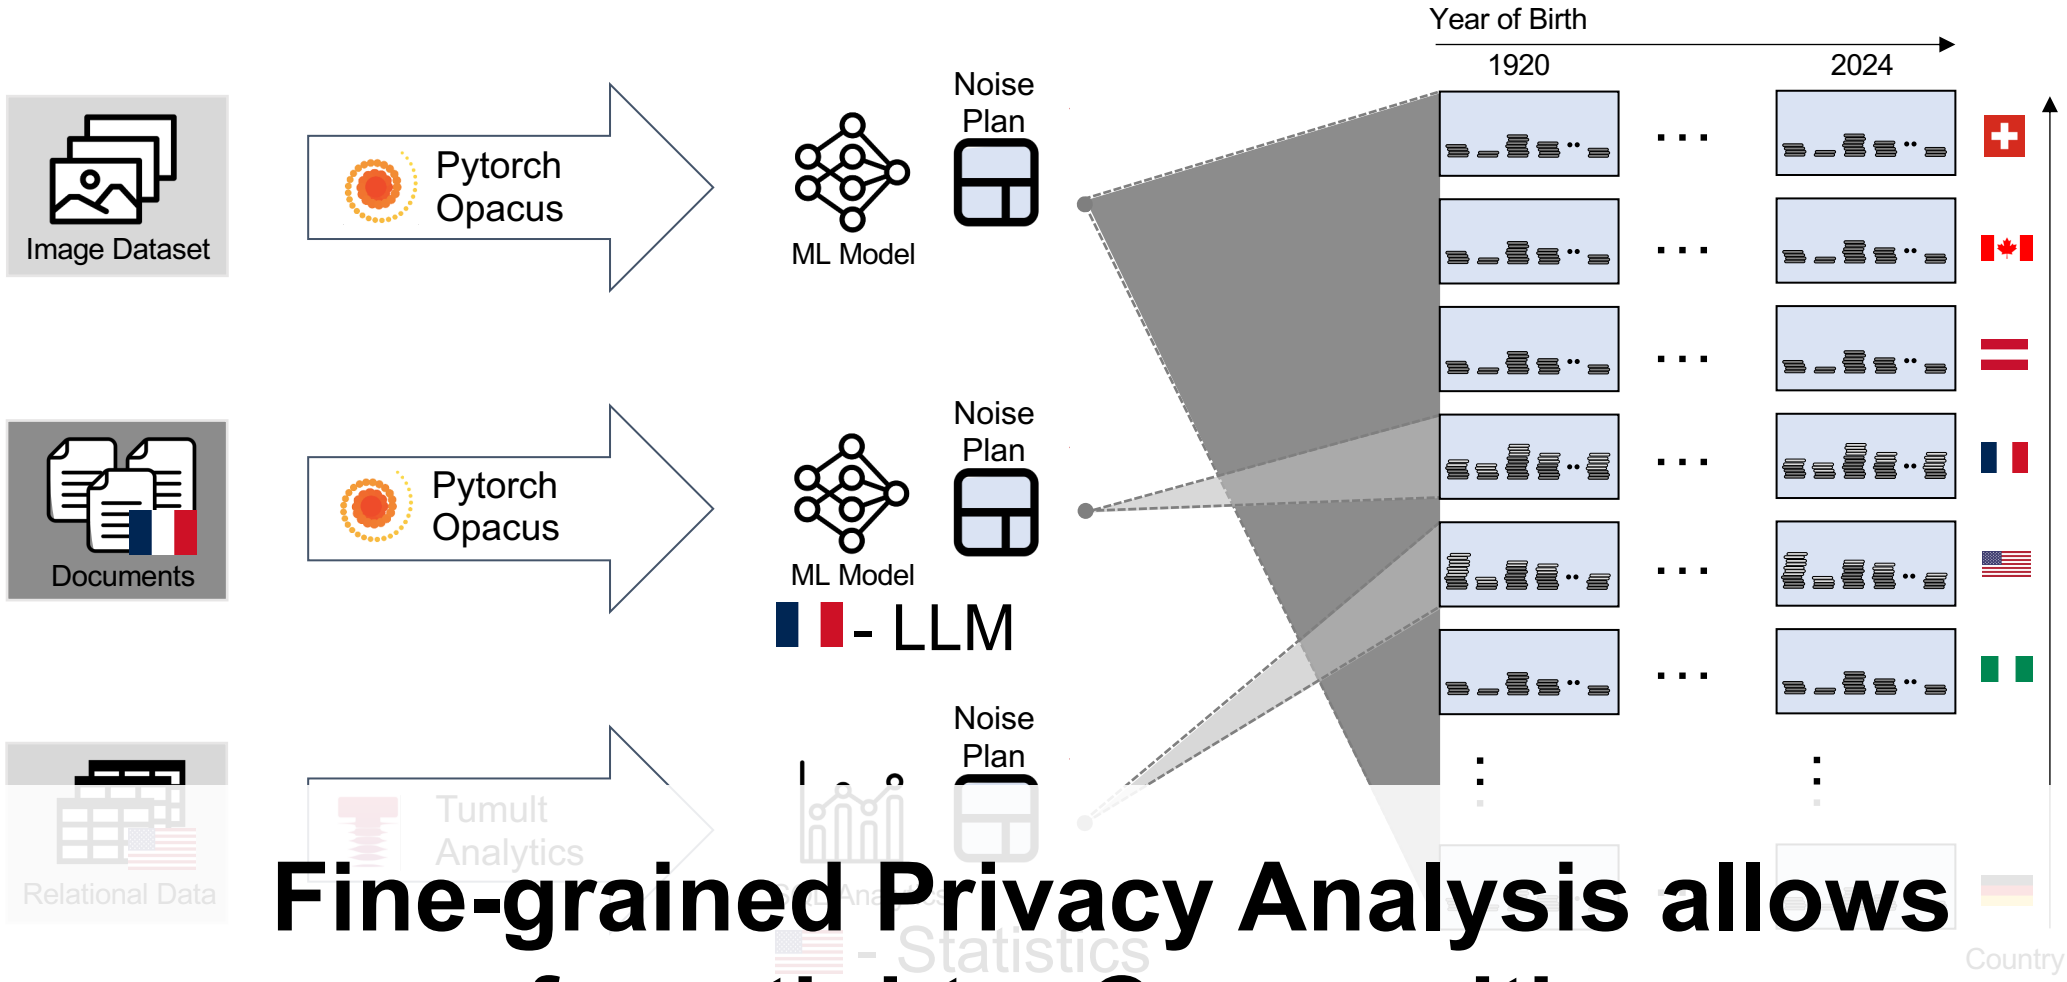

Fine-grained Privacy Analysis allows for a tighter Composition.

Sampling: Random Subset Selection

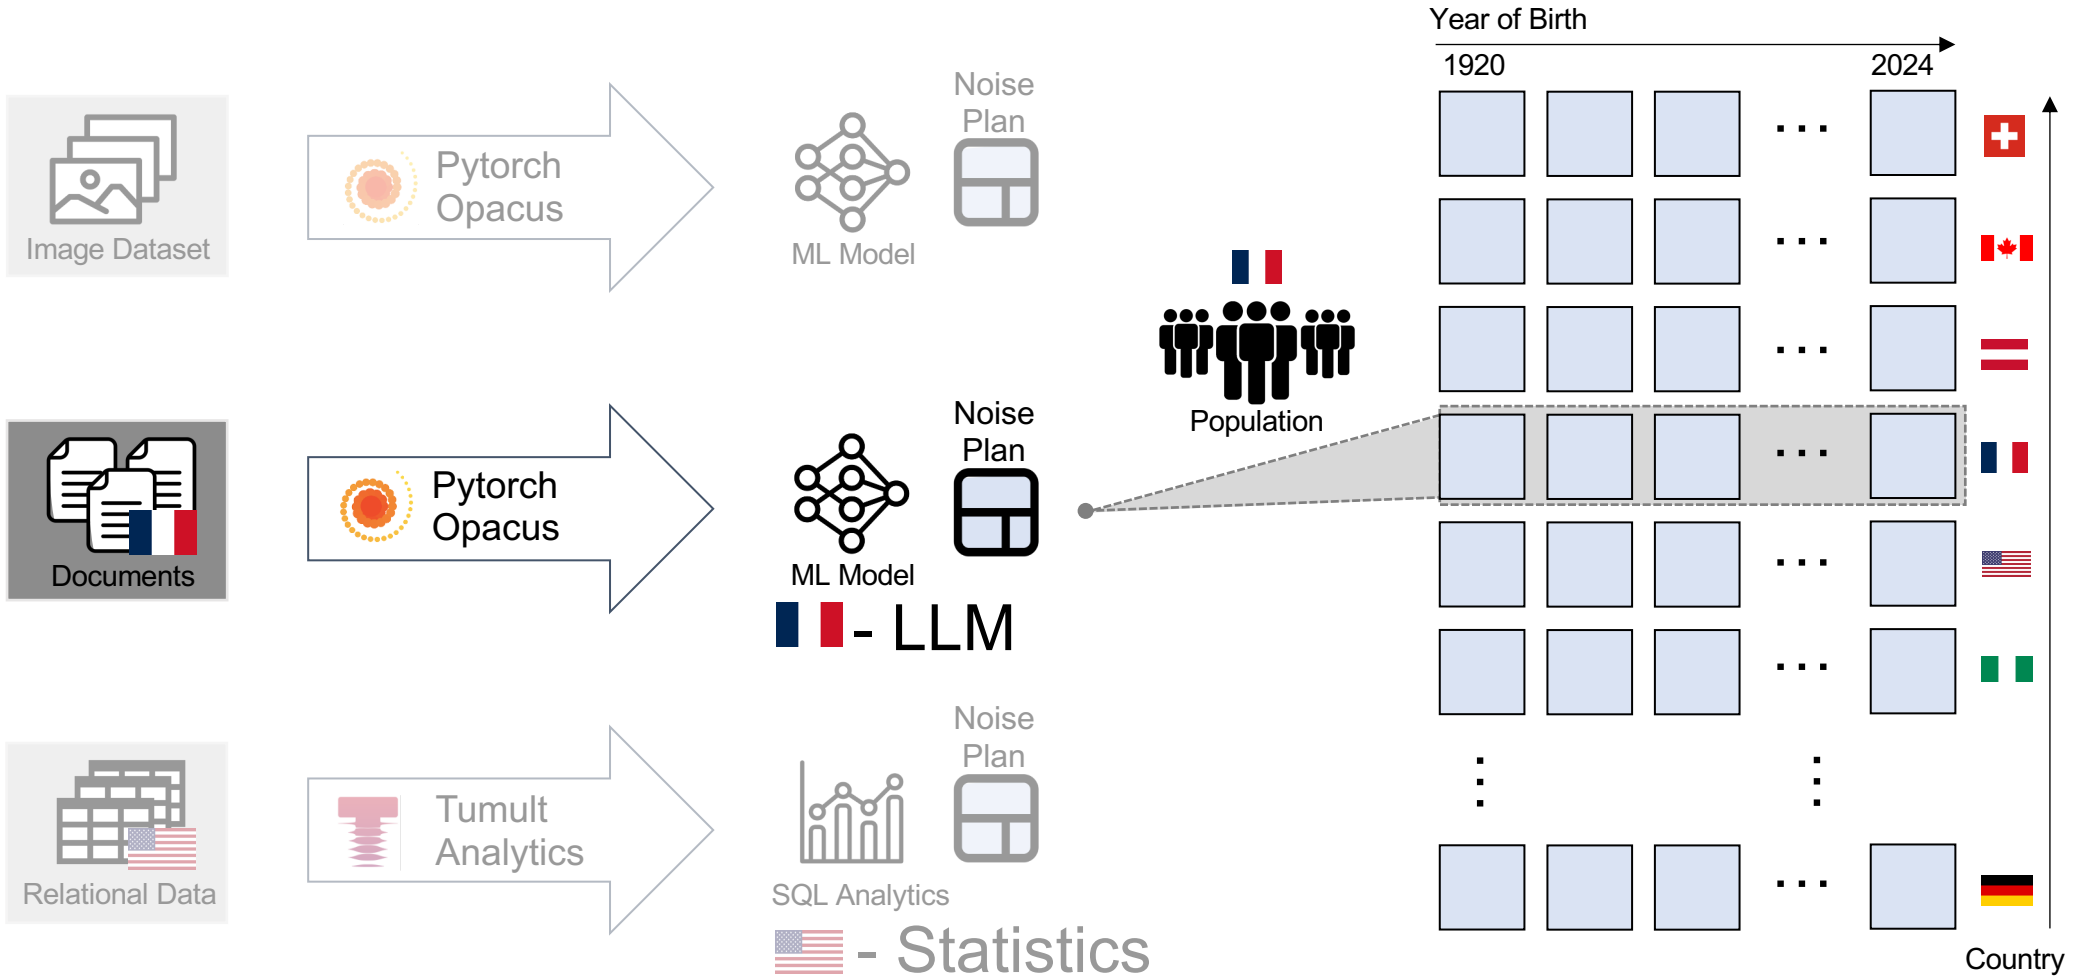

Sampling: Random Subset Selection

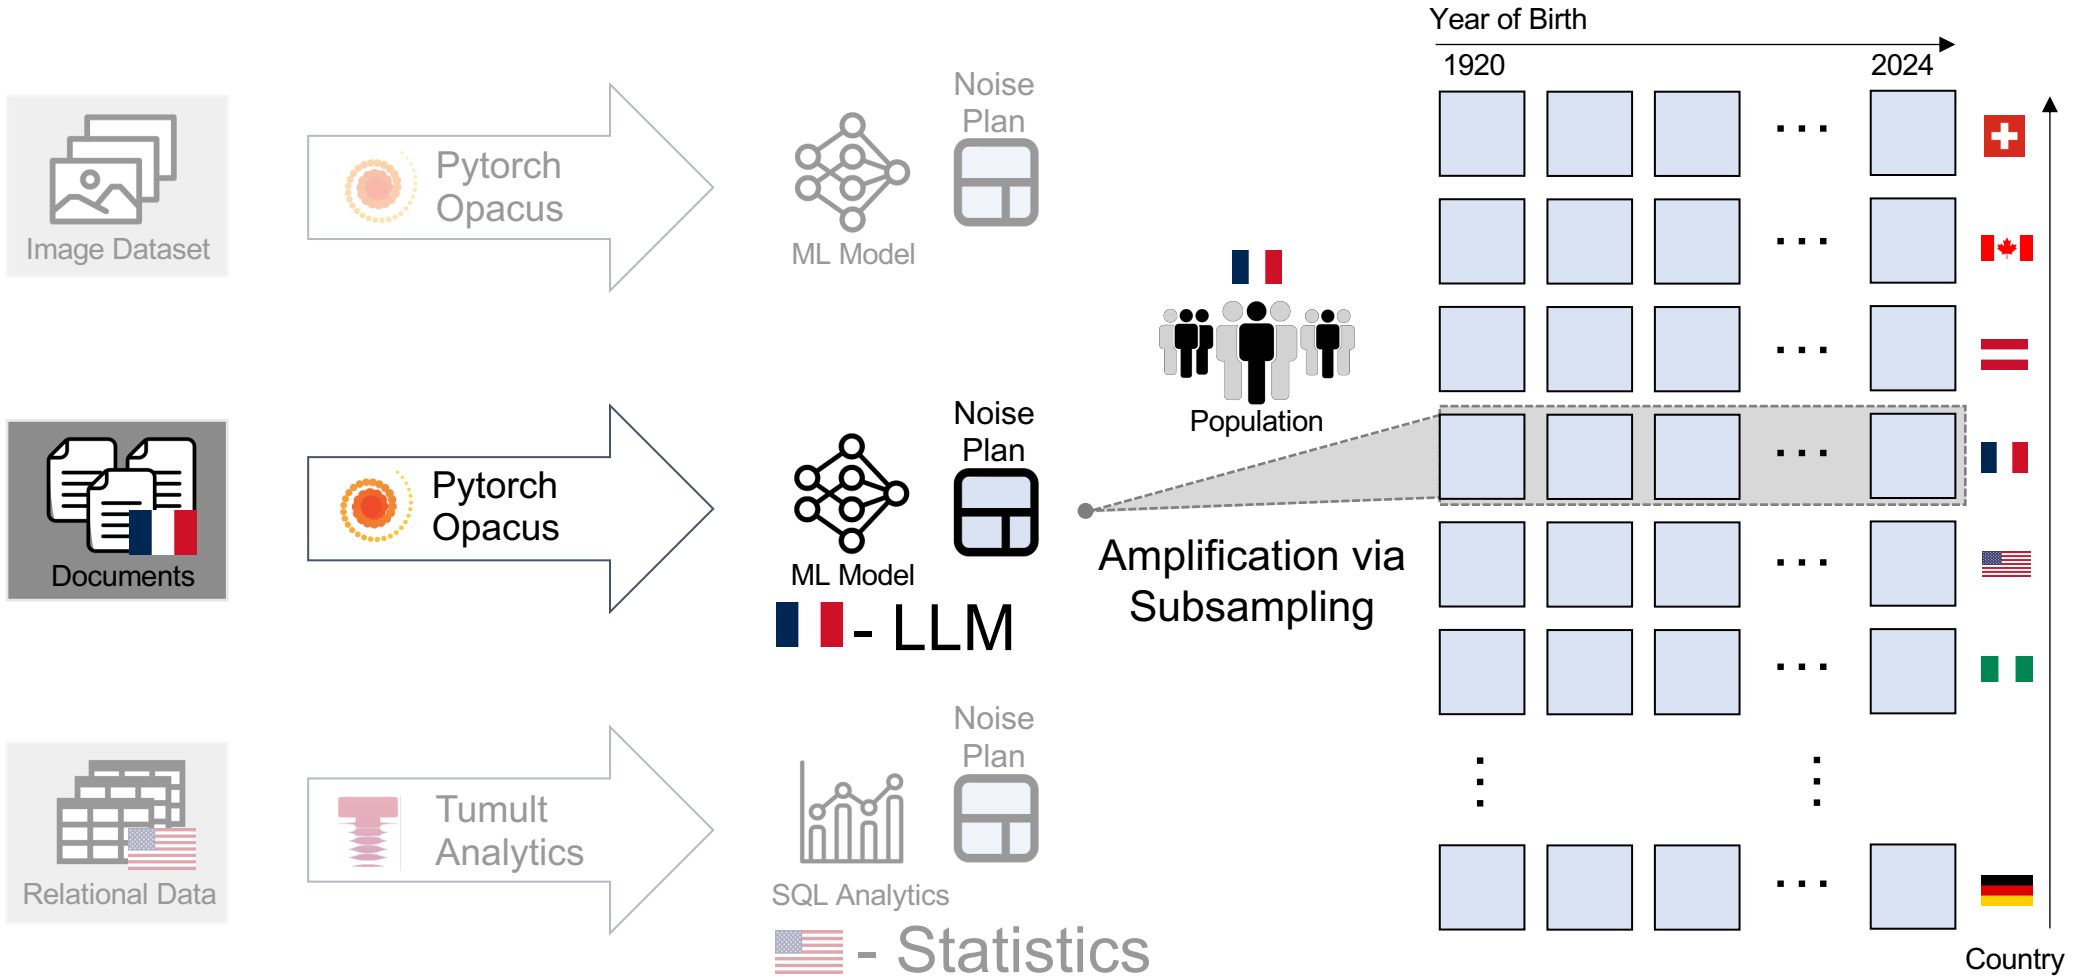

Application Layer

Scarce and Finite Resource

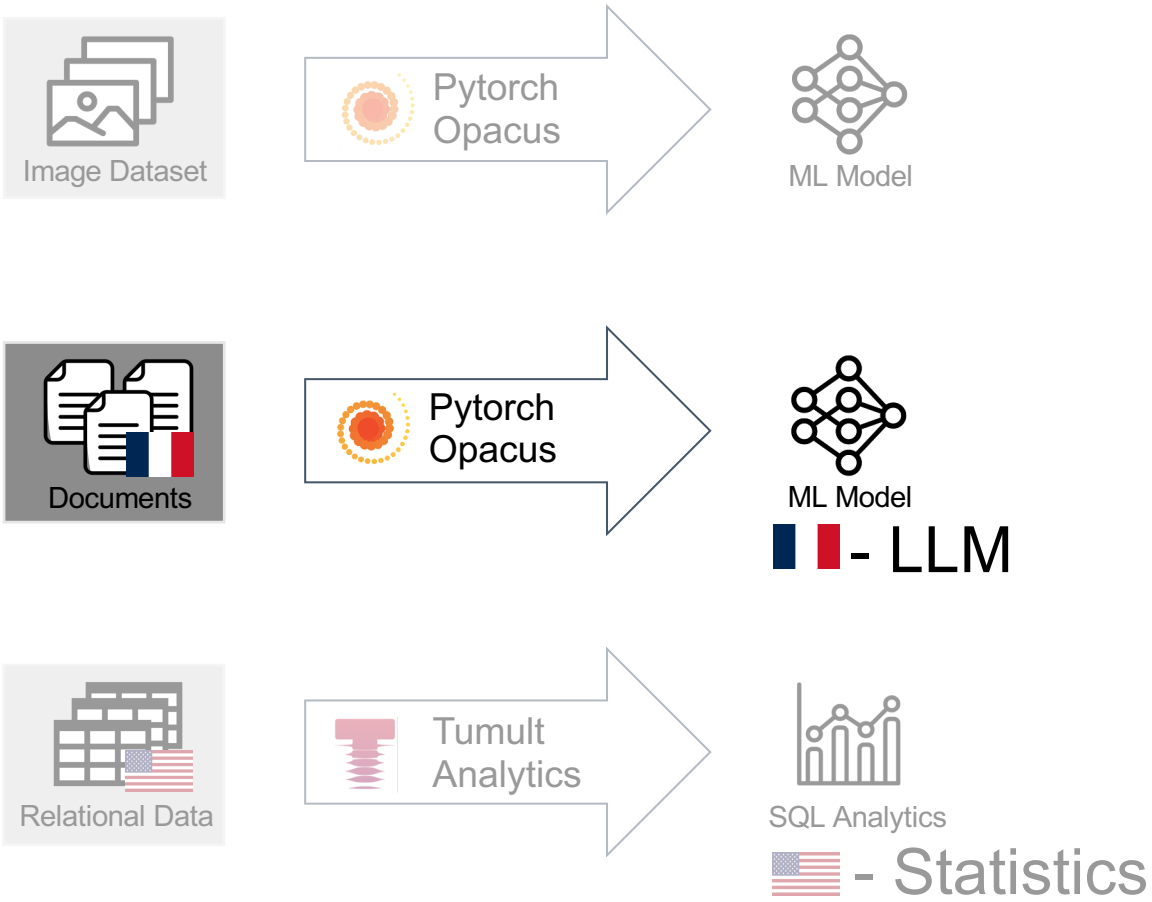

Application Layer

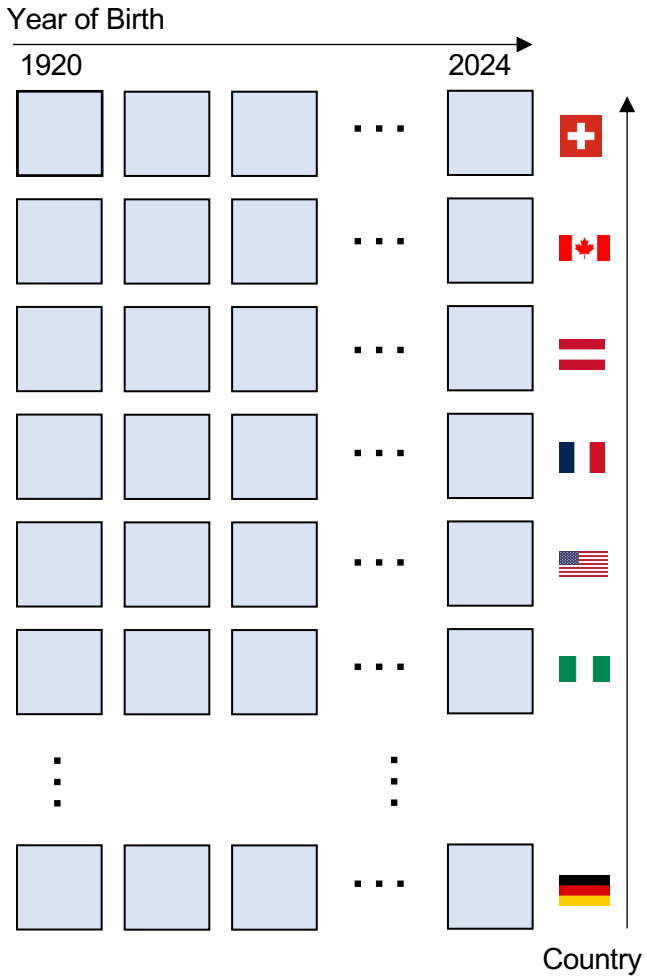

Management Layer

Scarce and Finite Resource

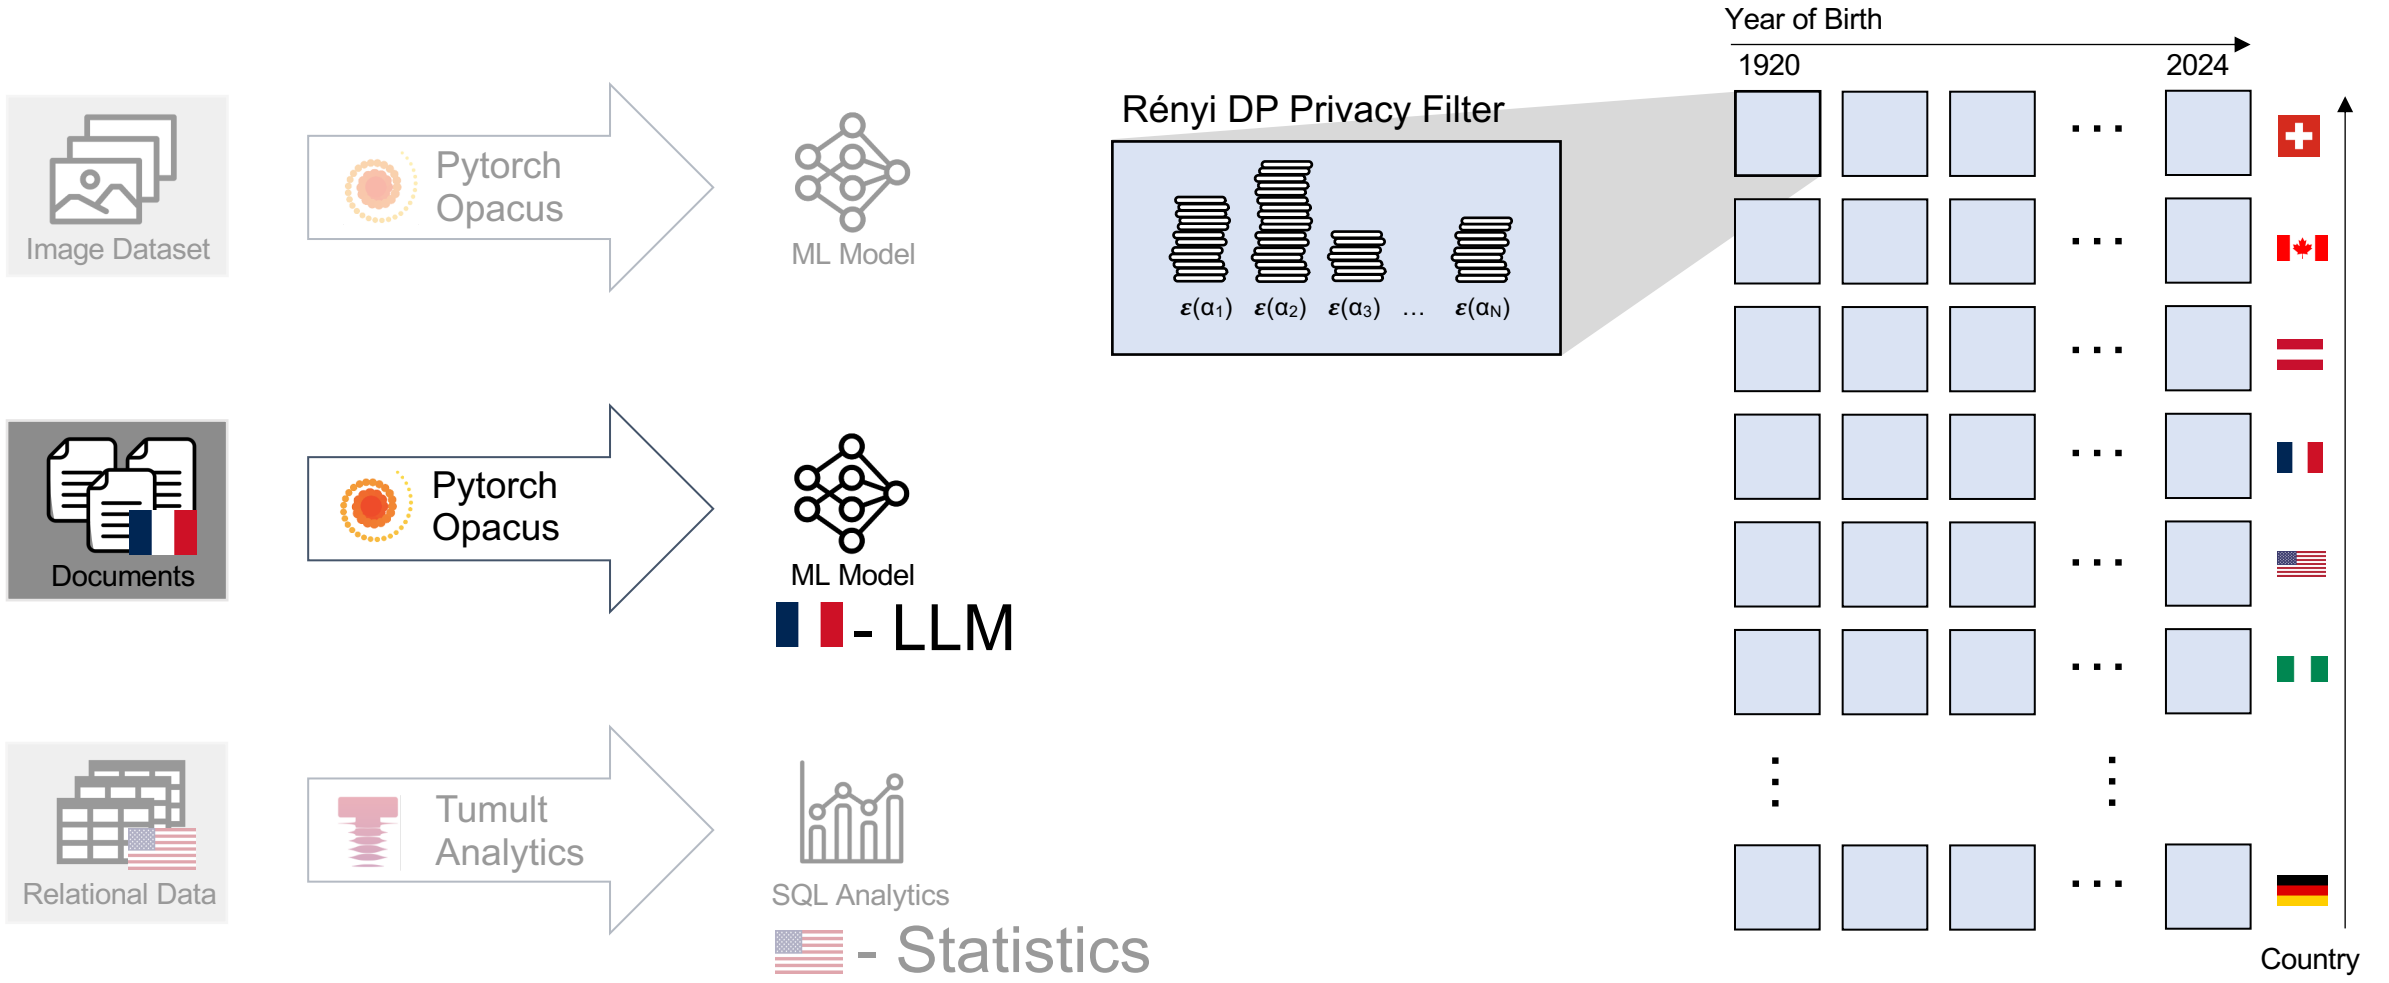

Scarce and Finite Resource

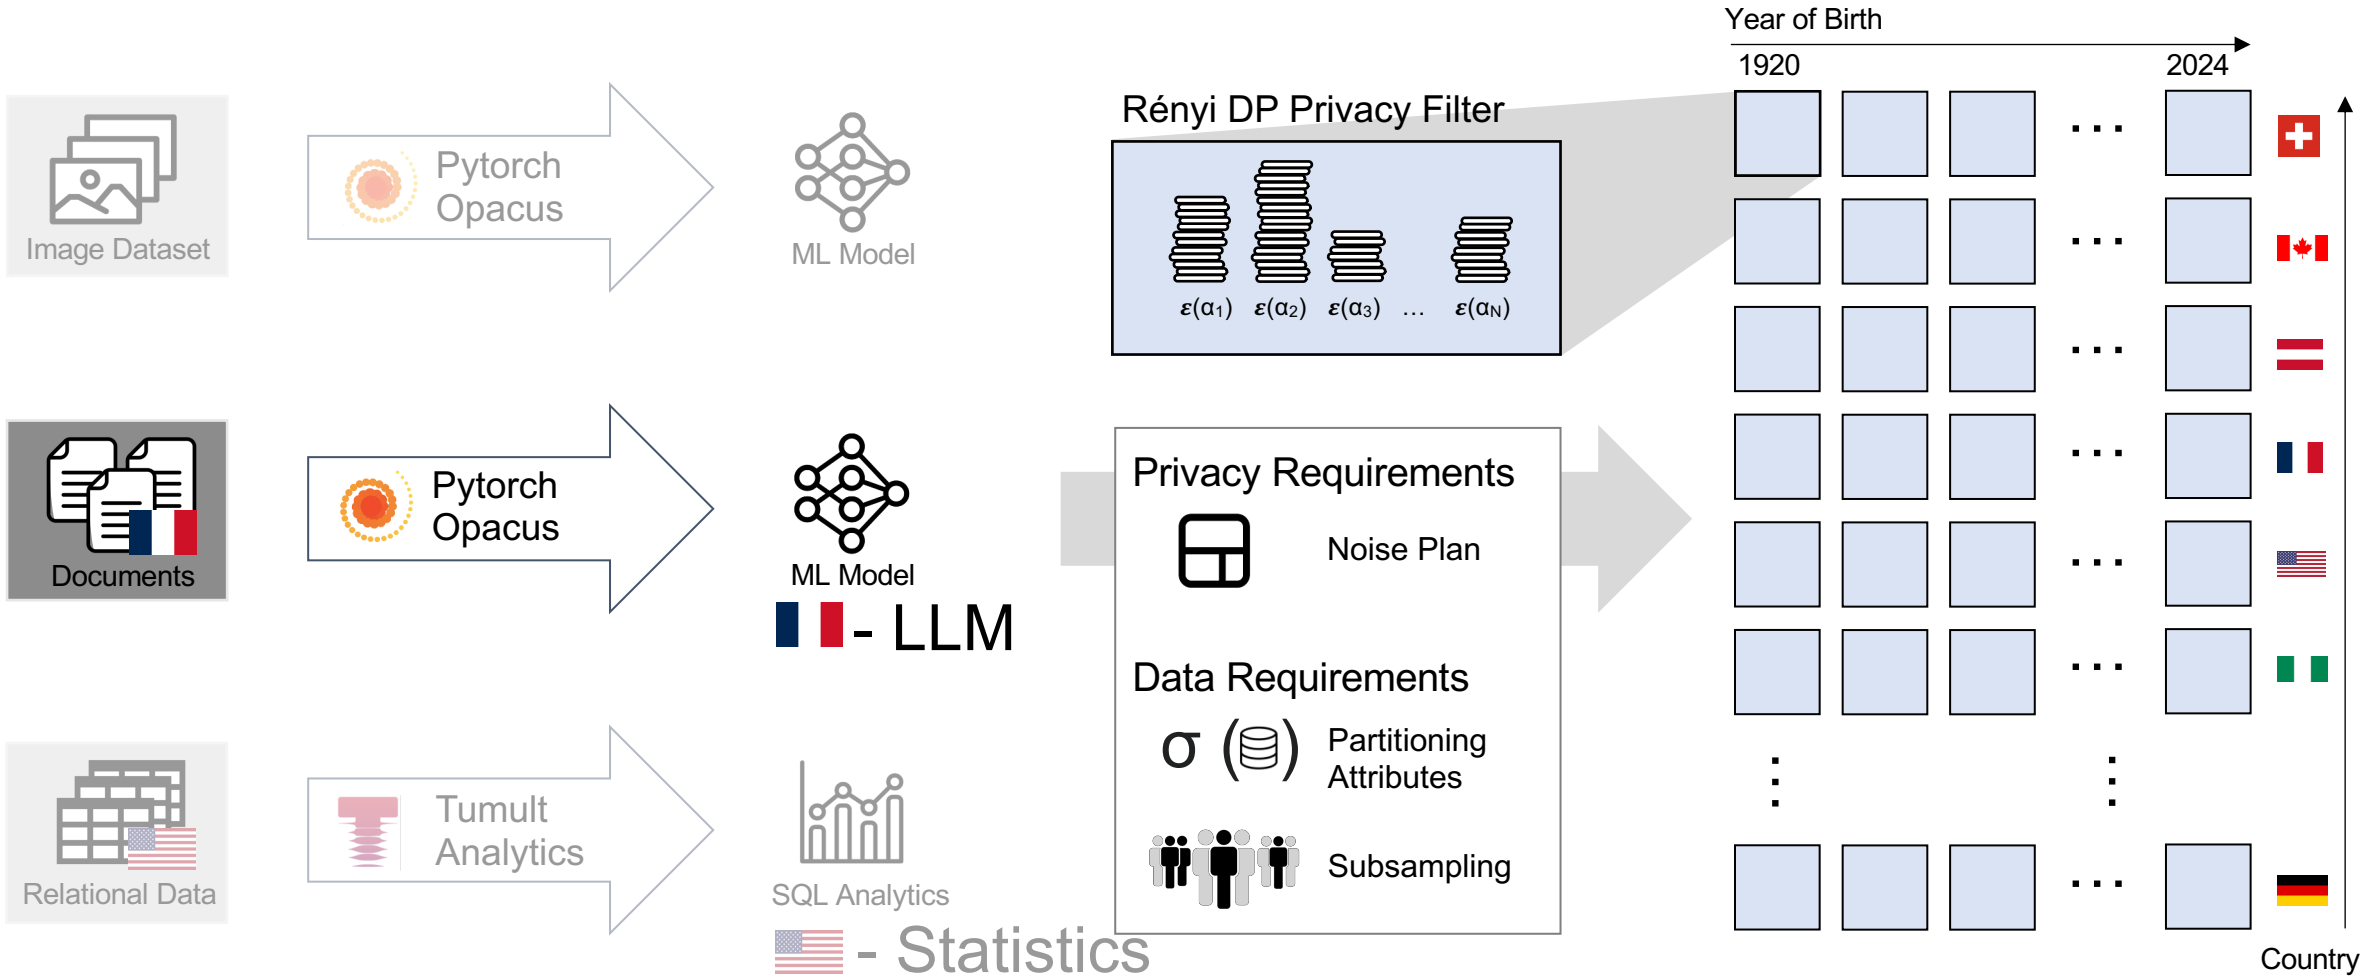

Scarce and Finite Resource

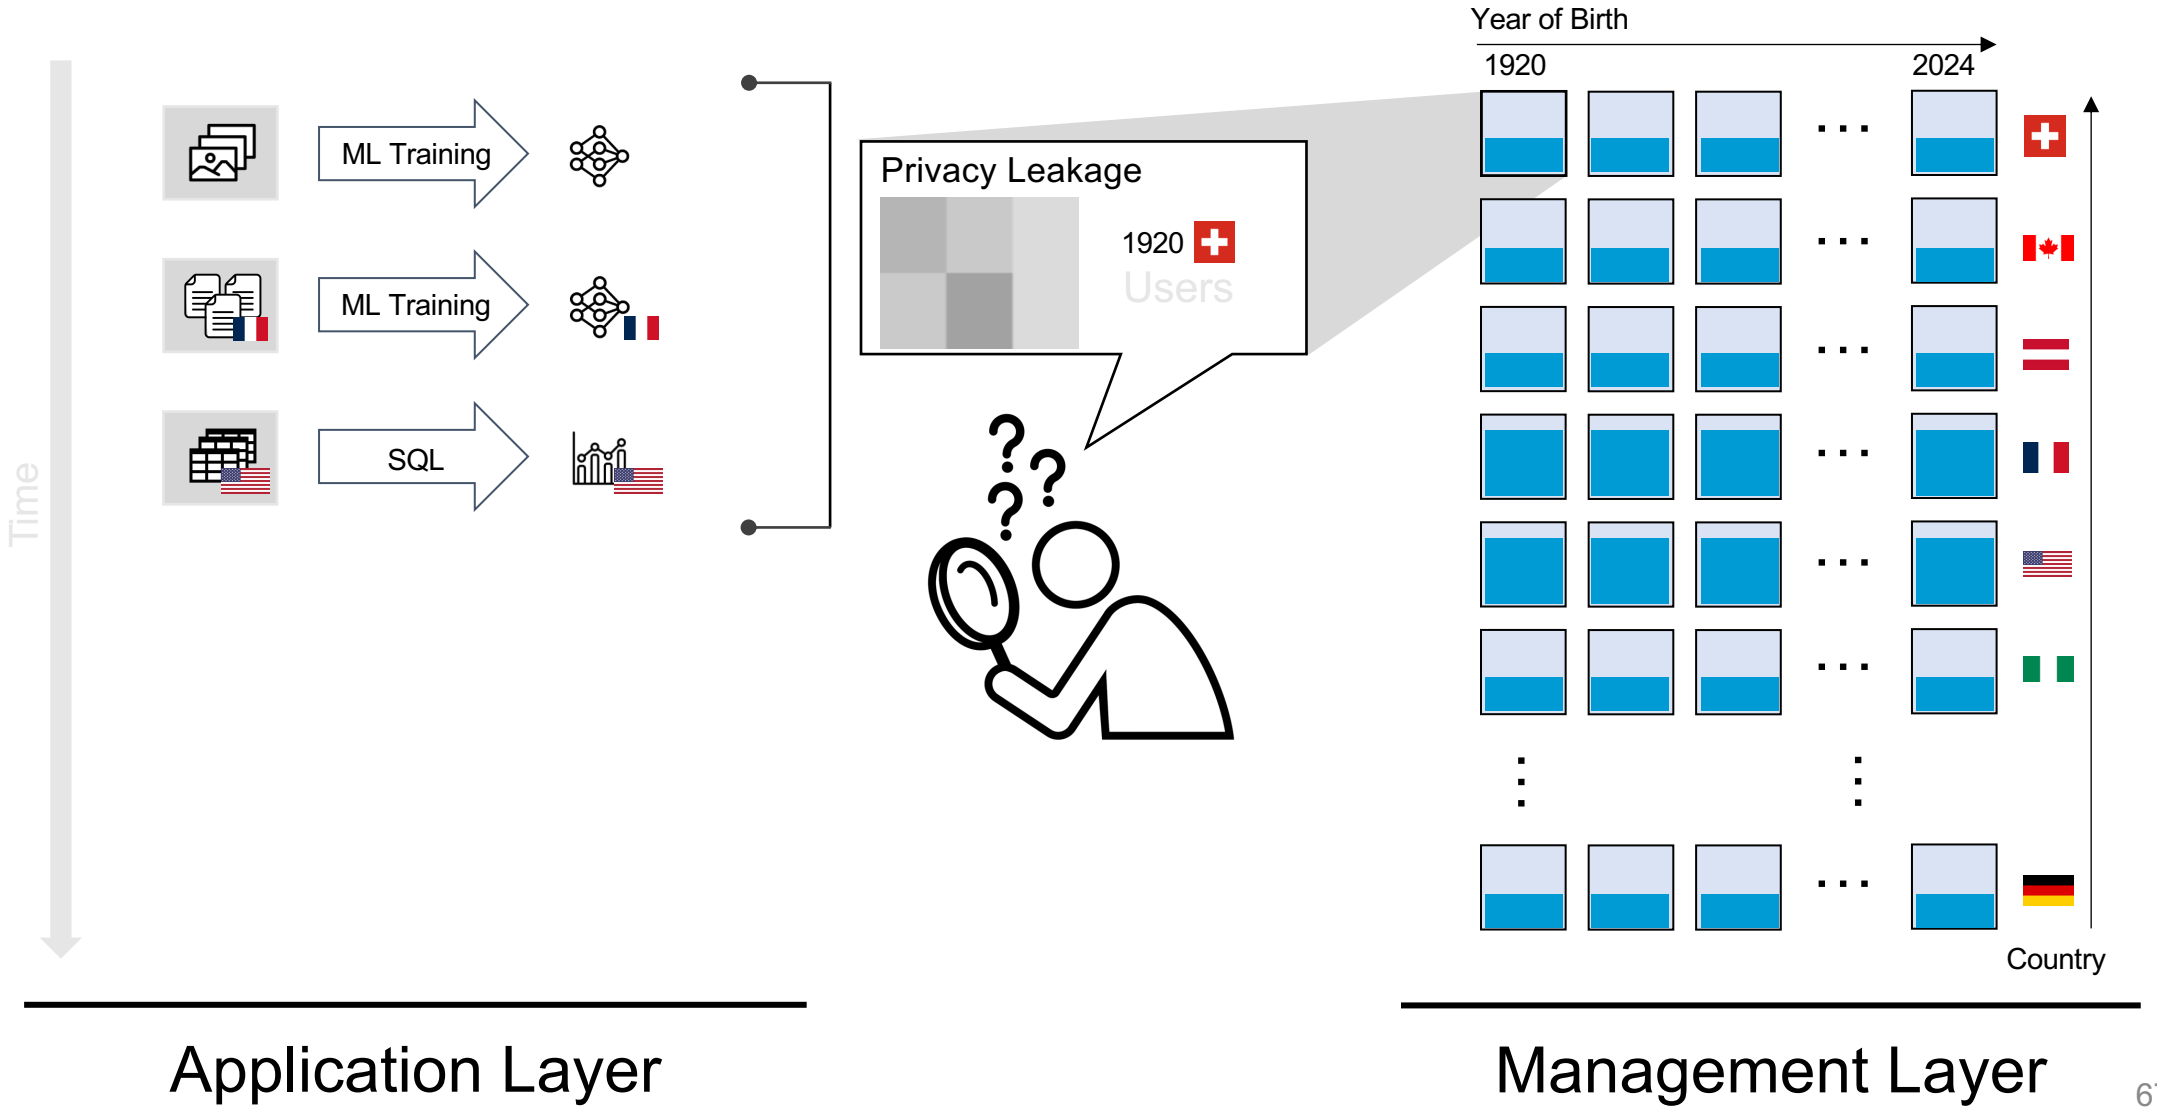

Scarce and Finite Resource

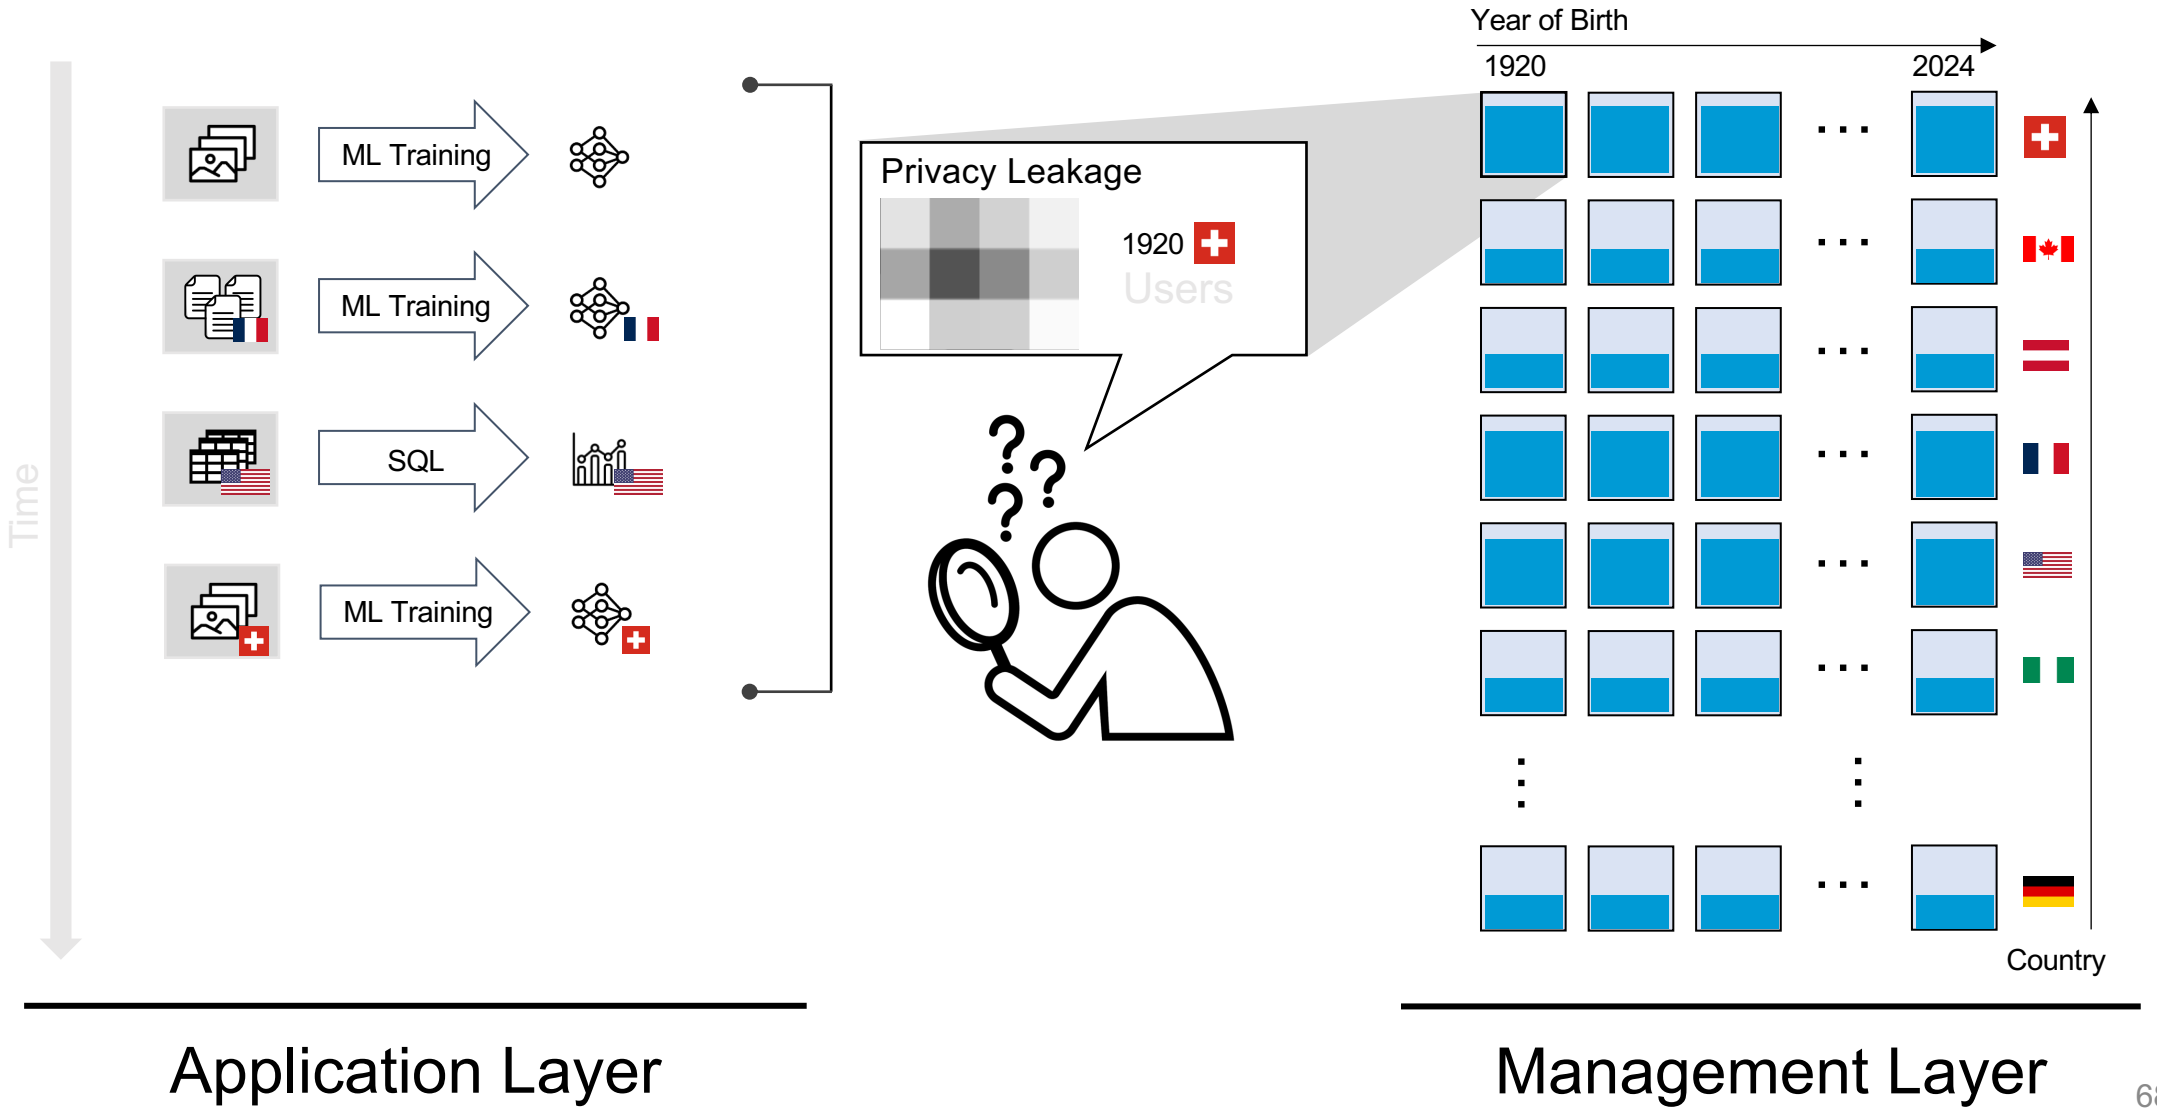

Scarce and Finite Resource

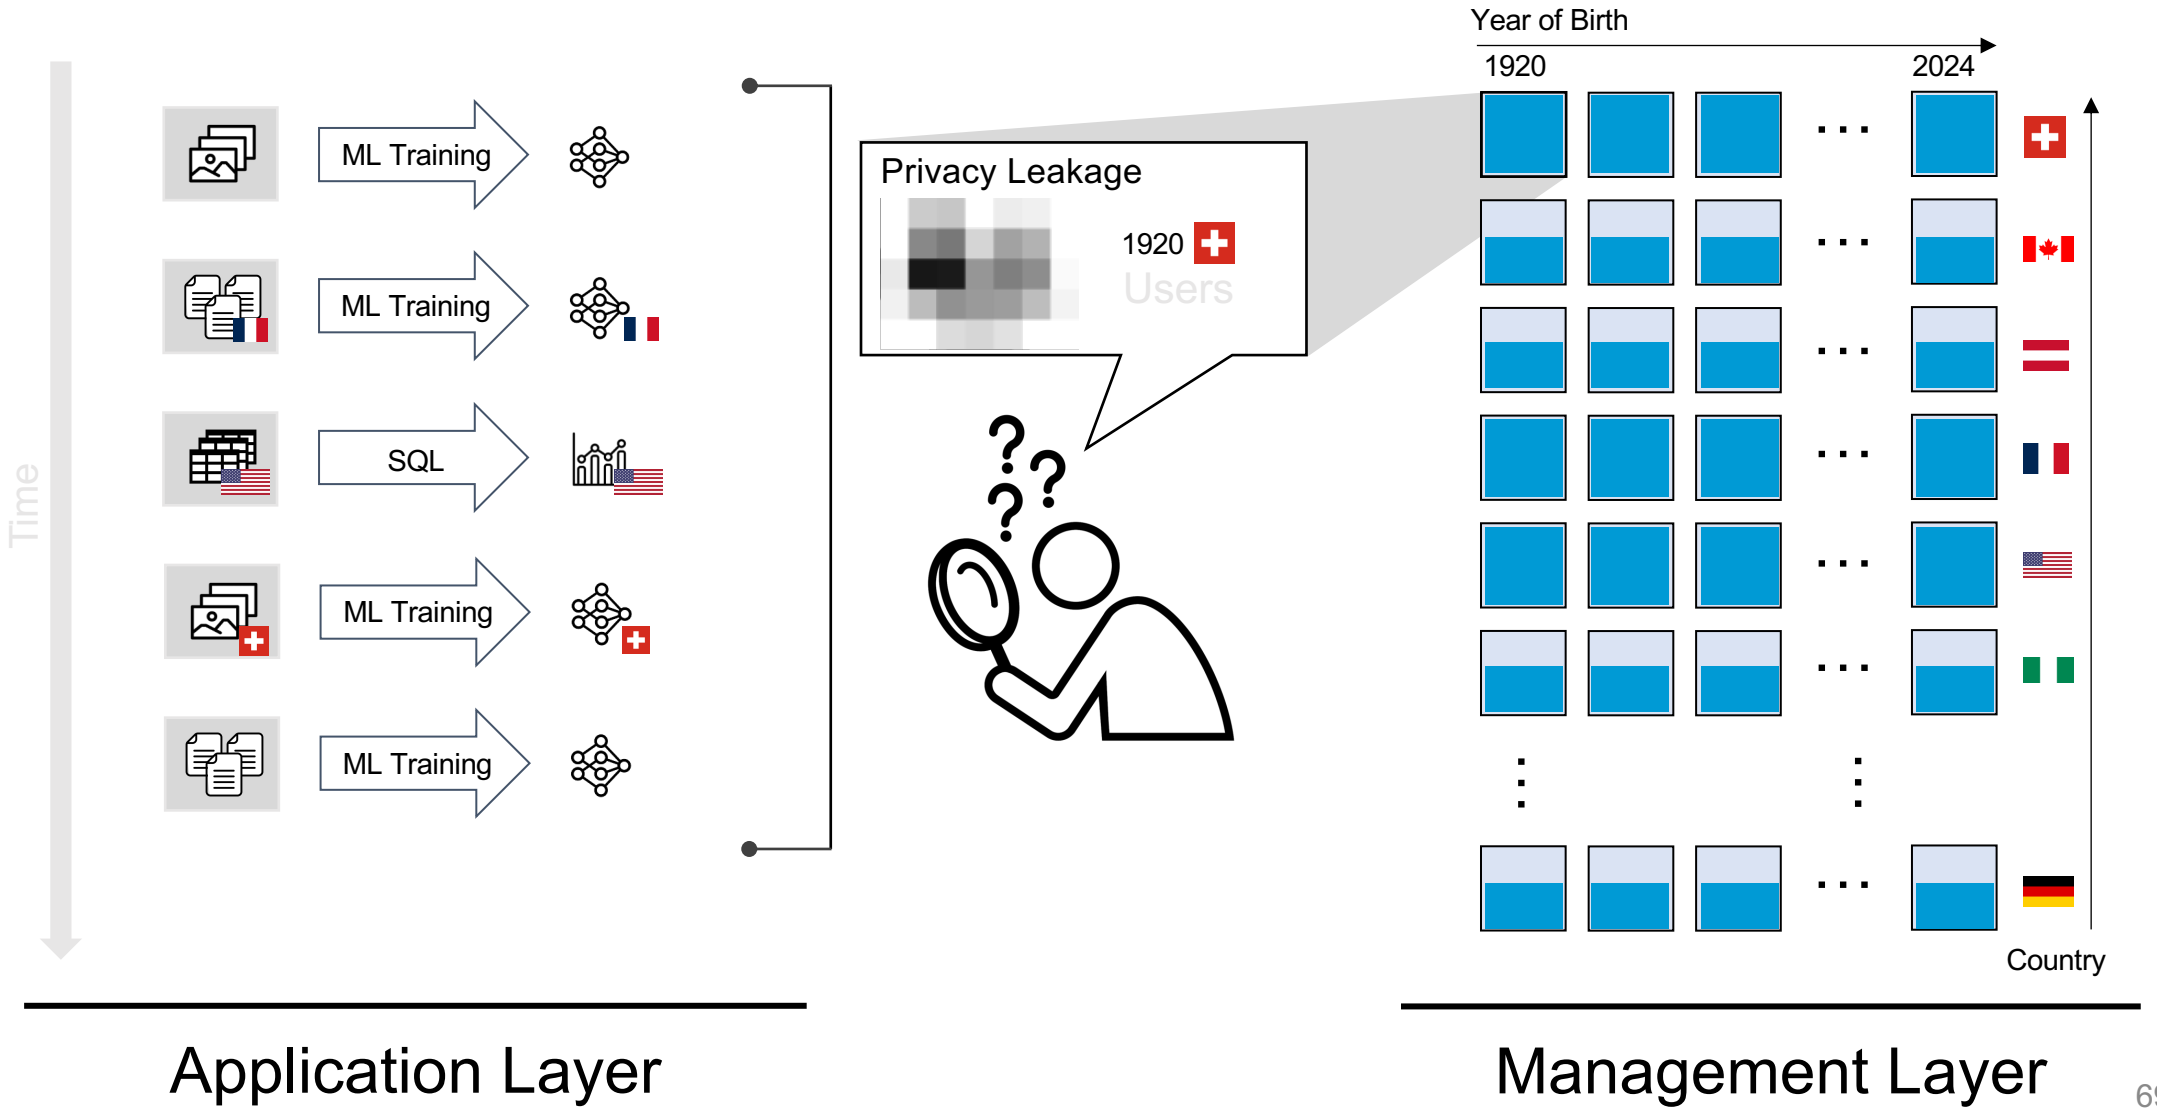

Scarce and Finite Resource

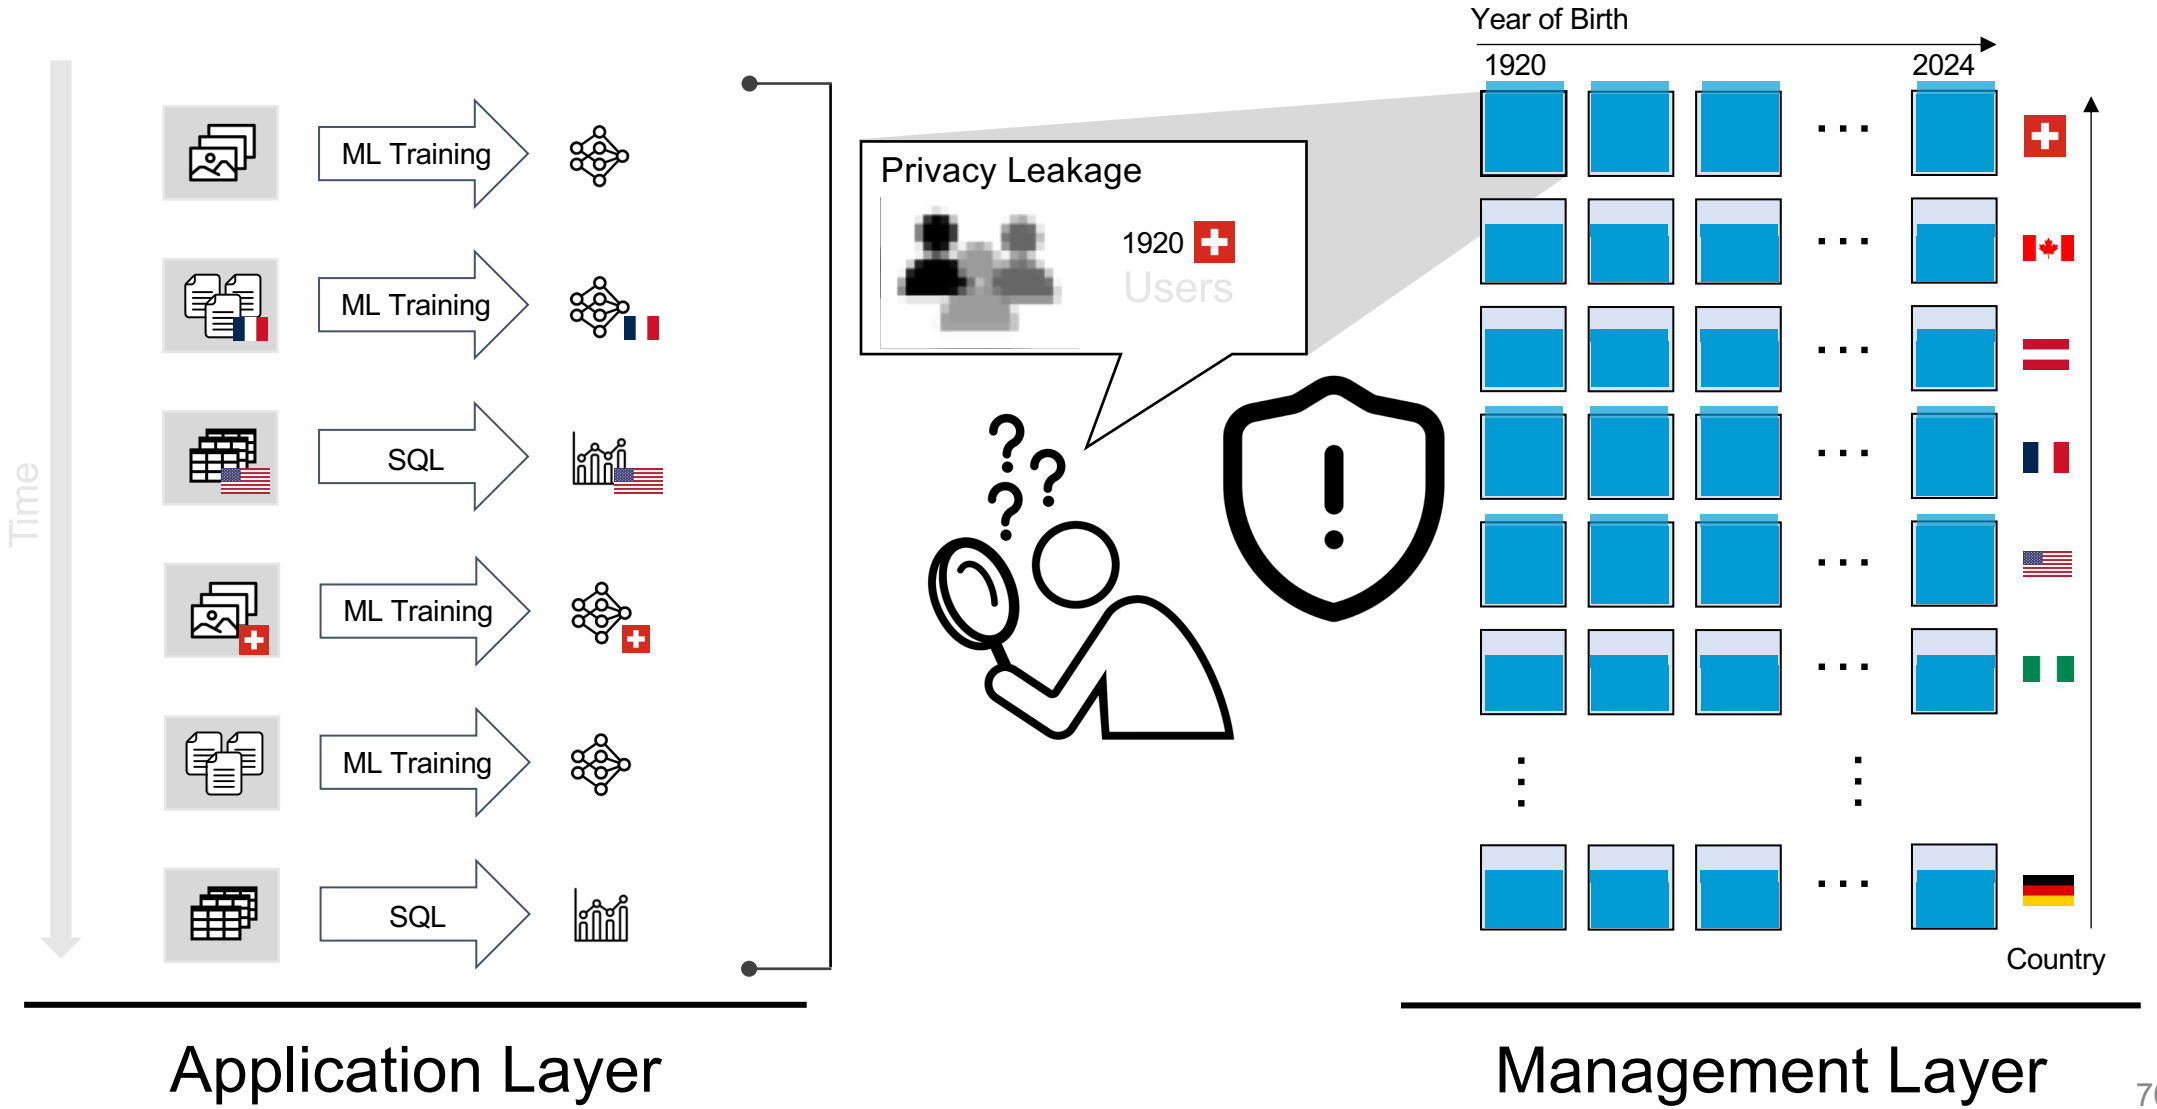

Scarce and Finite Resource

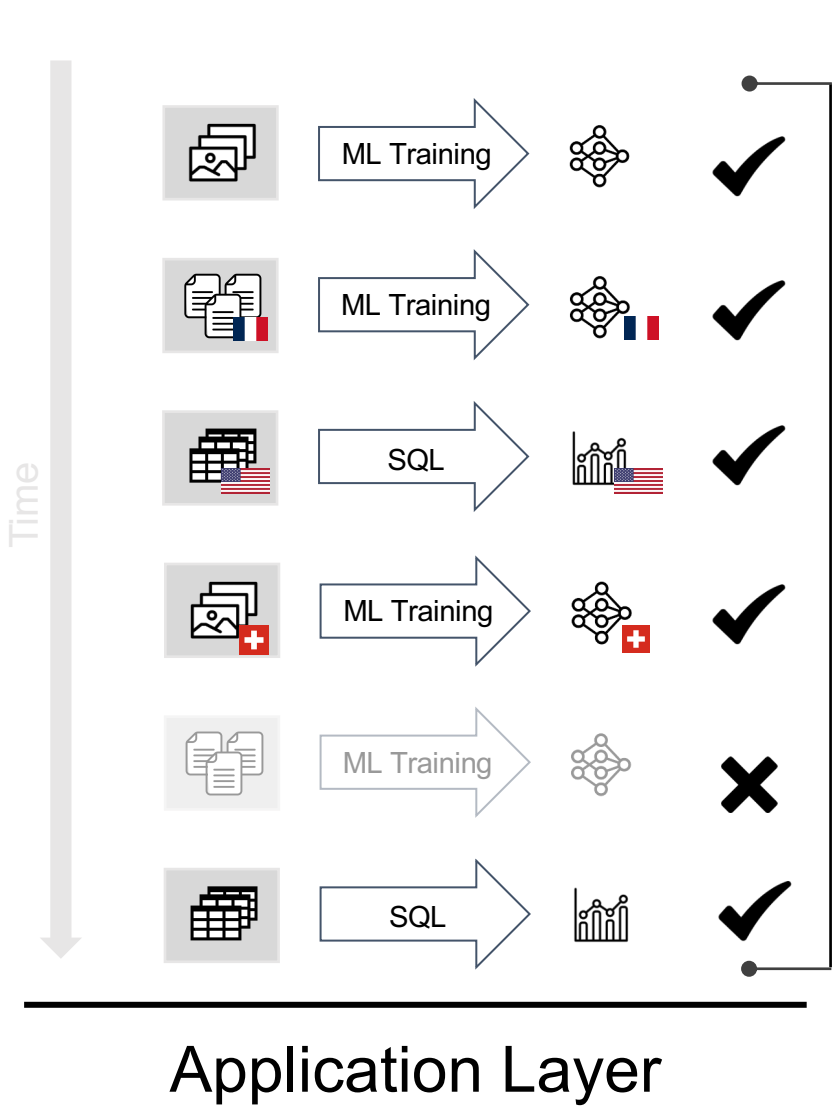

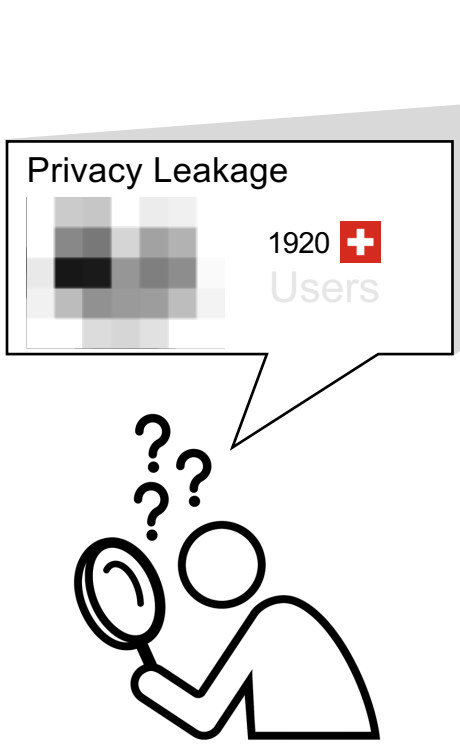

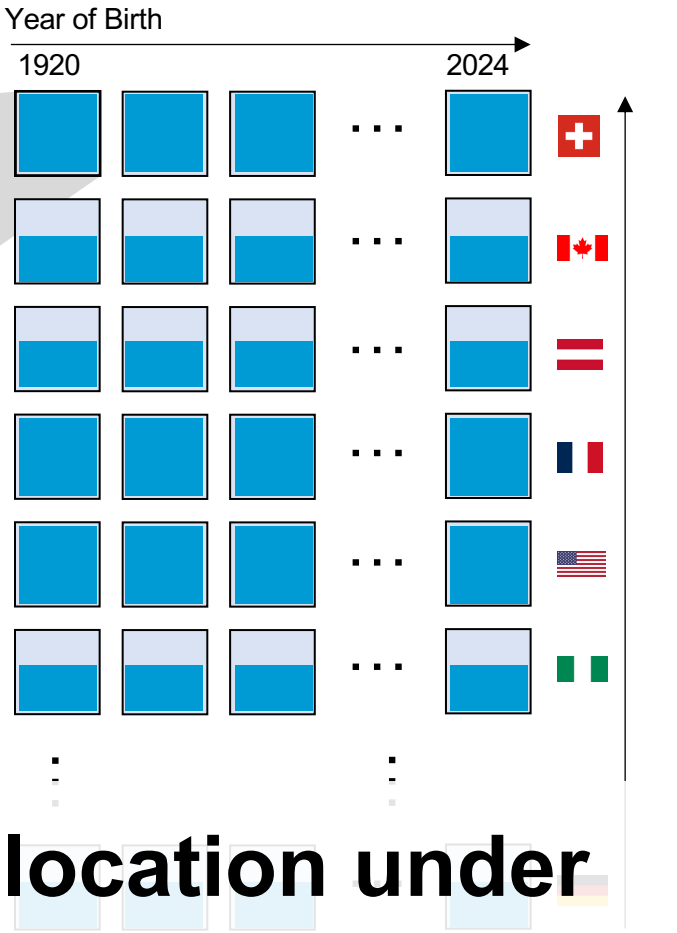

Resource Allocation under Finite Privacy Budget

Continuity under a Finite Budget

Ensuring Sustained Budget Allocation Over Time

Resetting Budget

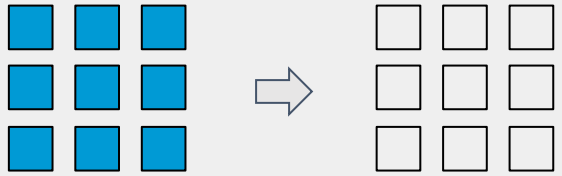

Continuity under a Finite Budget

Ensuring Sustained Budget Allocation Over Time

Resetting Budget

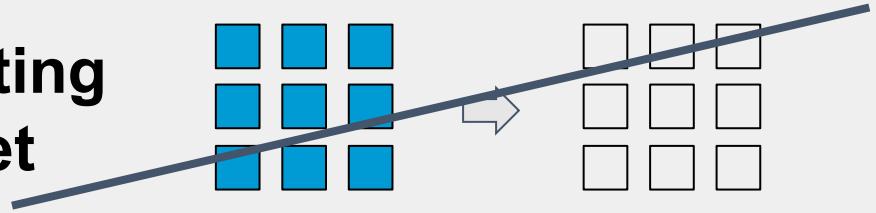

DP Violation

Continuity under a Finite Budget

Ensuring Sustained Budget Allocation Over Time

Resetting Budget

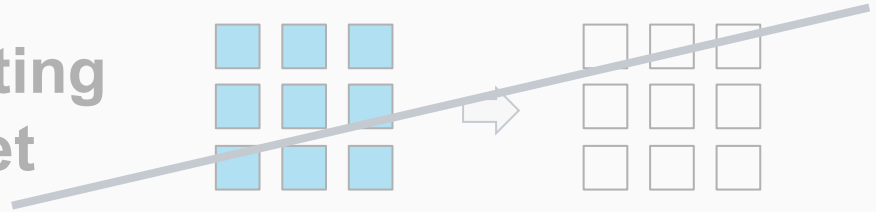

DP Violation

User Rotation

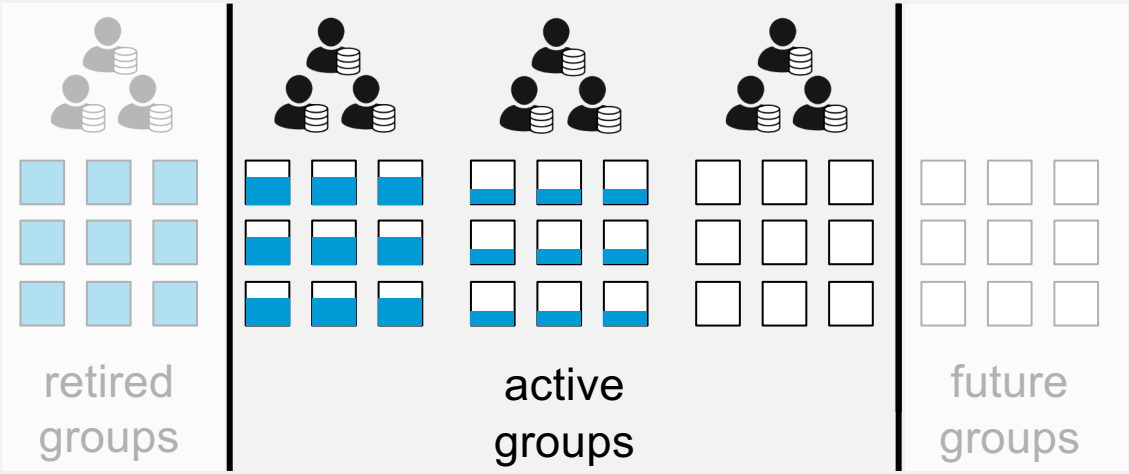

Continuity under a Finite Budget

Ensuring Sustained Budget Allocation Over Time

Resetting Budget

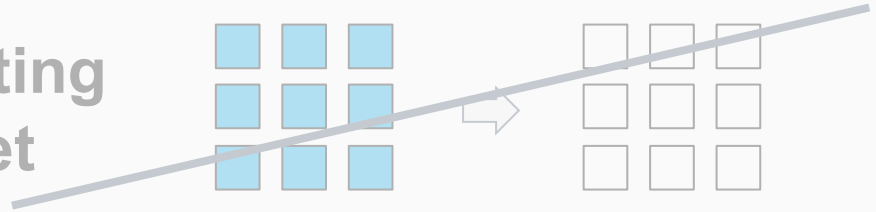

DP Violation

User Rotation

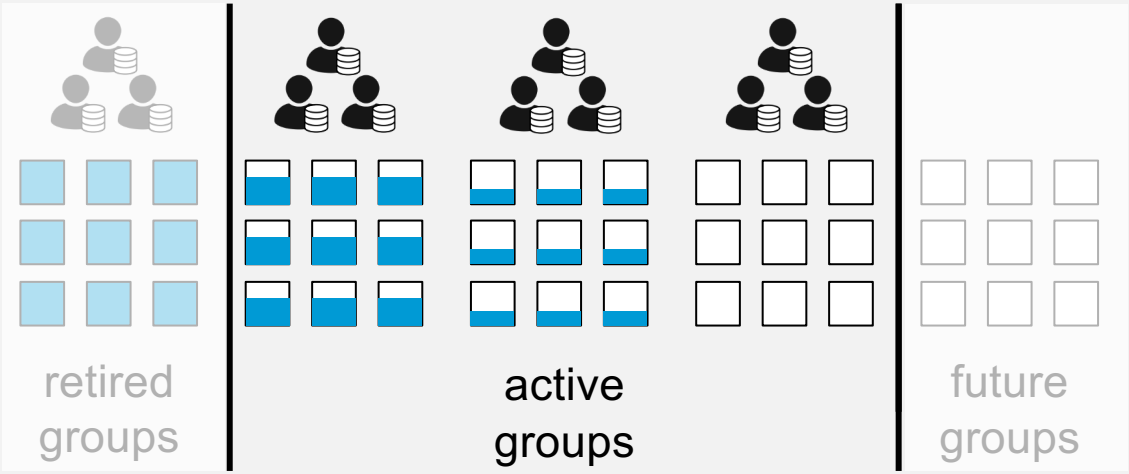

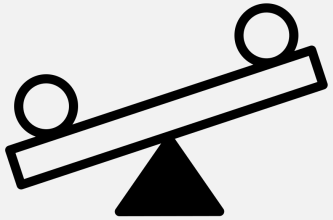

Biased Set of Active Users

Continuity under a Finite Budget

Ensuring Sustained Budget Allocation Over Time

Resetting Budget

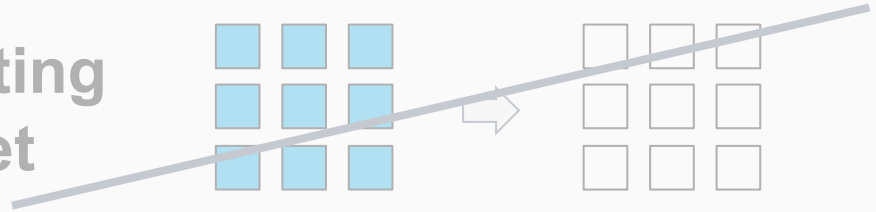

DP Violation

User Rotation

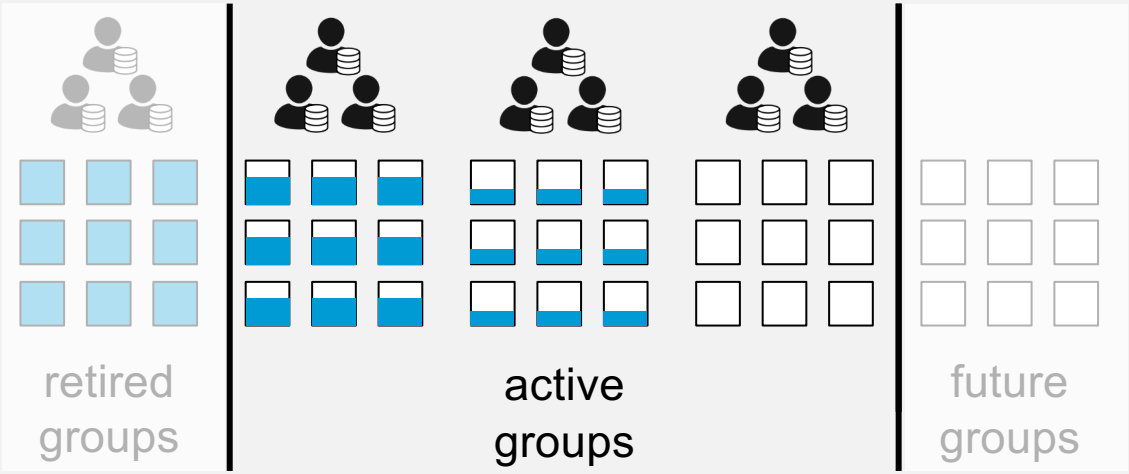

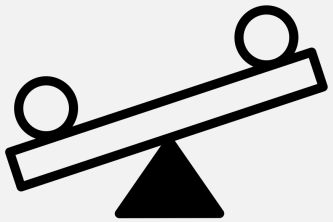

Biased Set of Active Users

Budget Guarantees with Unlocking

Continuity under a Finite Budget

Ensuring Sustained Budget Allocation Over Time

Resetting Budget

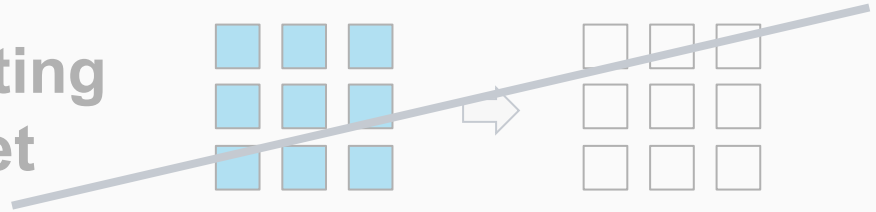

DP Violation

User Rotation

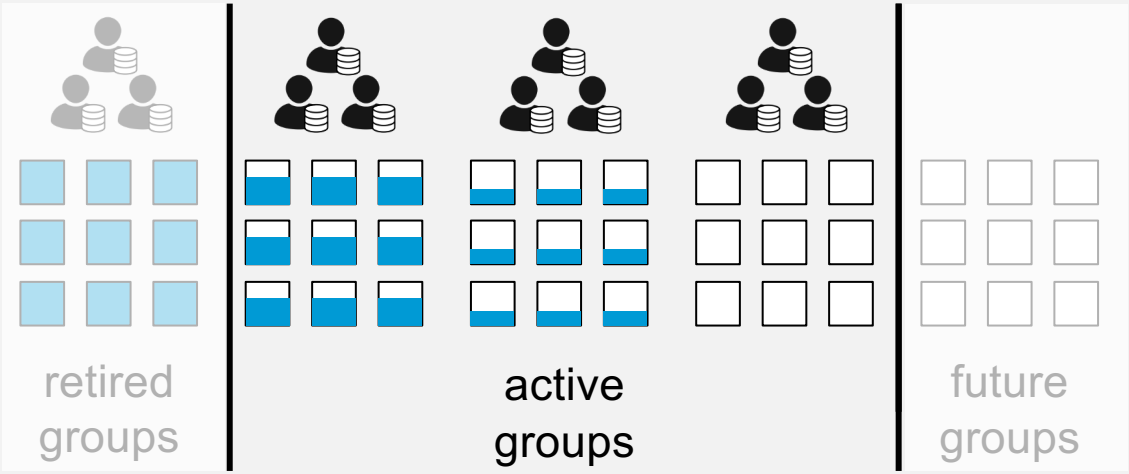

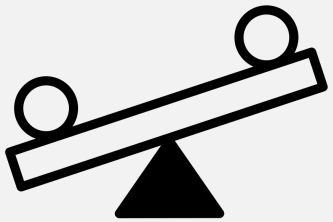

Biased Set of Active Users

Budget Guarantees with Unlocking

Continuity under a Finite Budget

Ensuring Sustained Budget Allocation Over Time

Resetting Budget

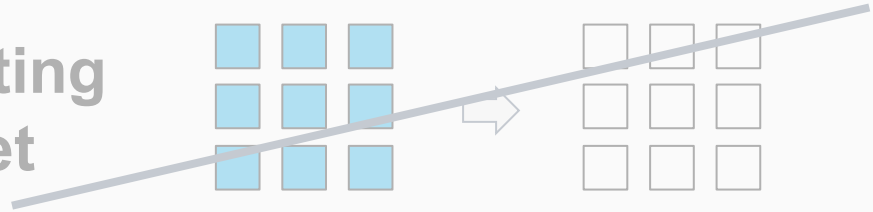

DP Violation

User Rotation

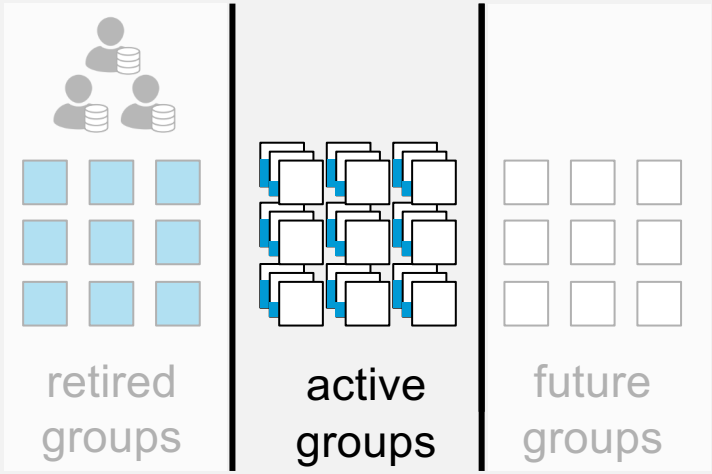

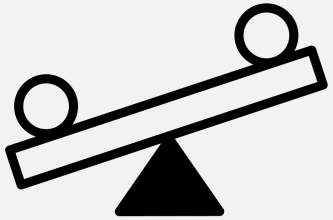

Biased Set of Active Users

Budget Guarantees with Unlocking

Privacy Resource Allocation

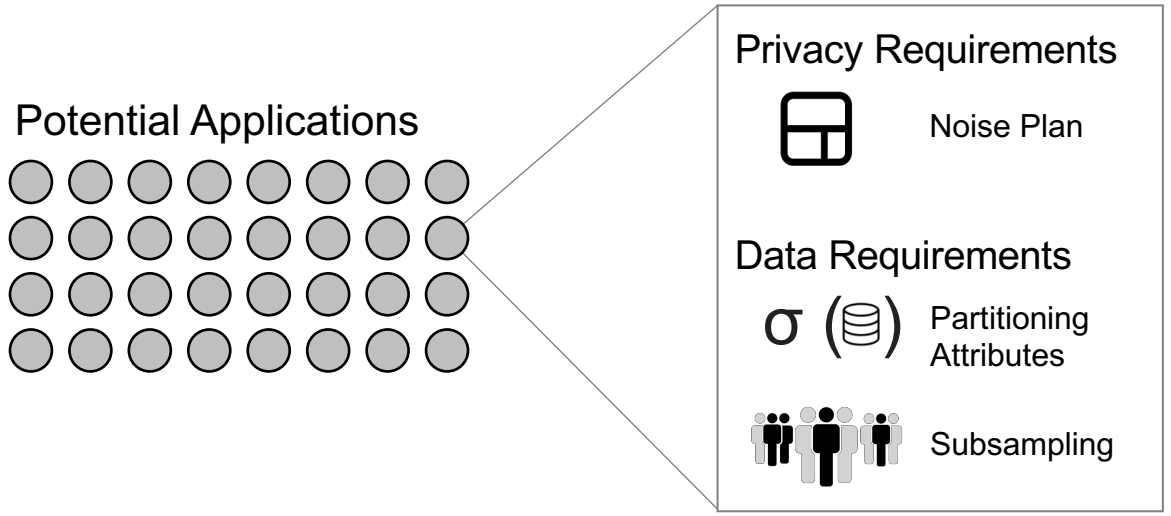

Shared Privacy State

Privacy Resource Allocation

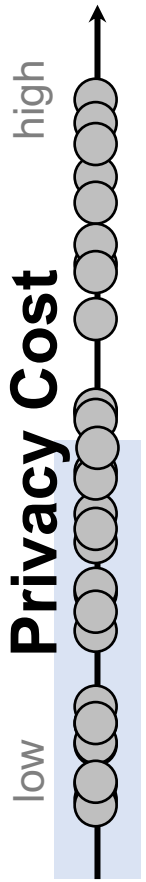

Optimize the number of Applications

Privacy Resource Allocation

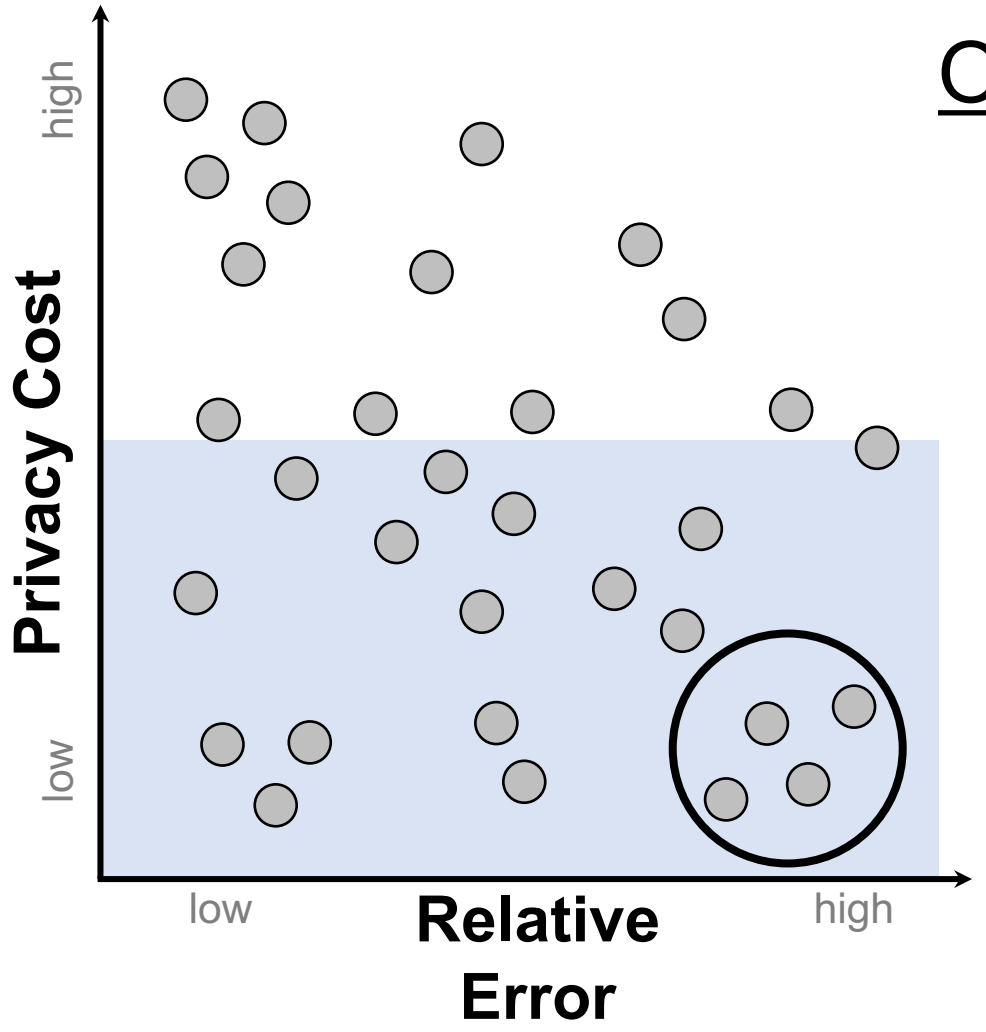

Optimize the number of Applications

Privacy Resource Allocation

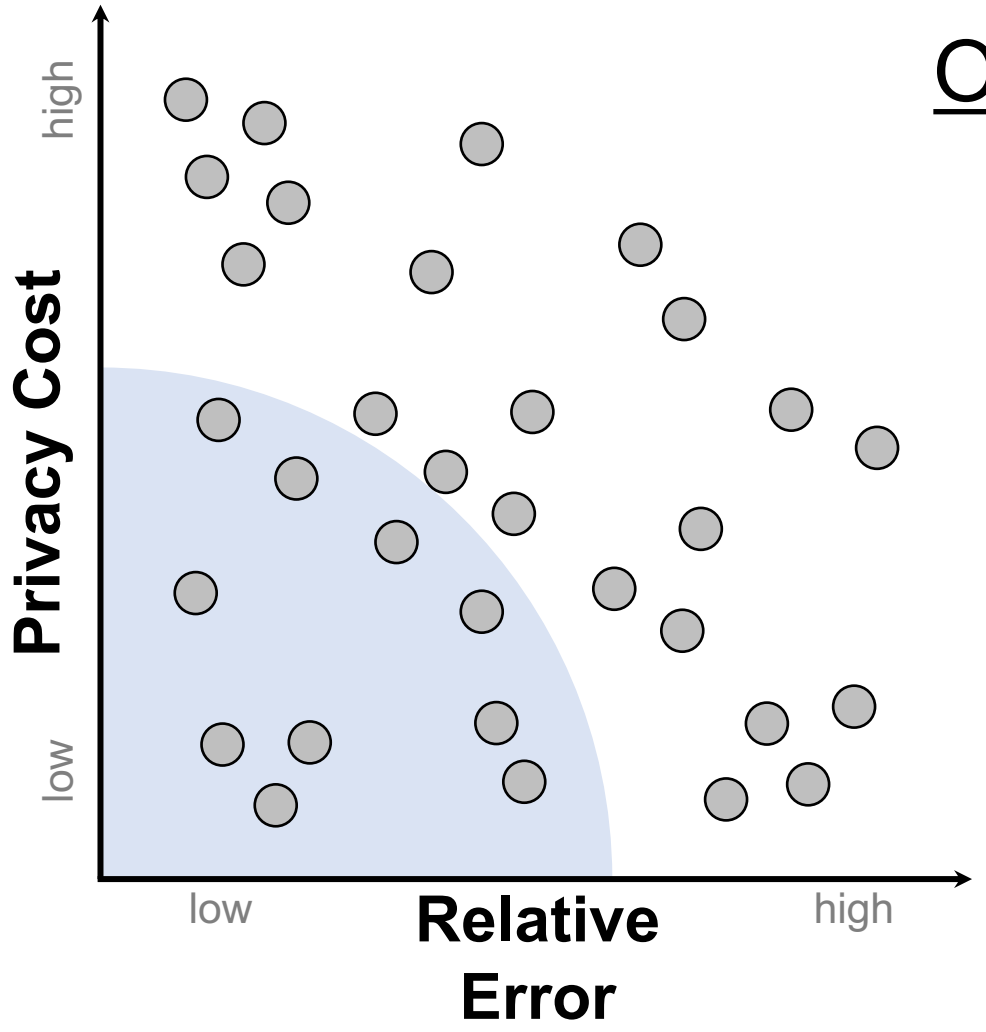

Optimize Privacy Cost Relative to Error

Privacy Resource Allocation

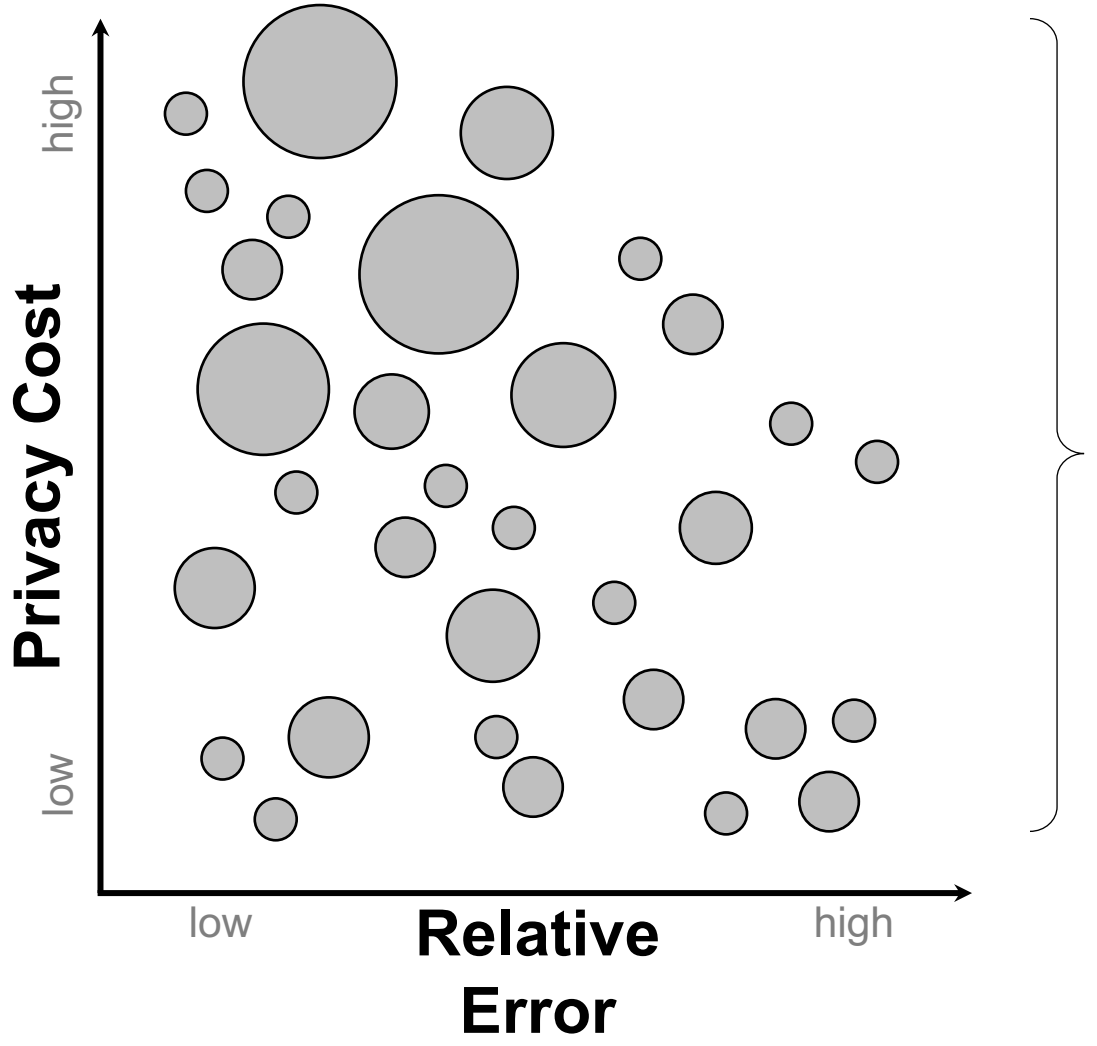

Optimize for Utility

Privacy Resource Allocation

Potential Applications

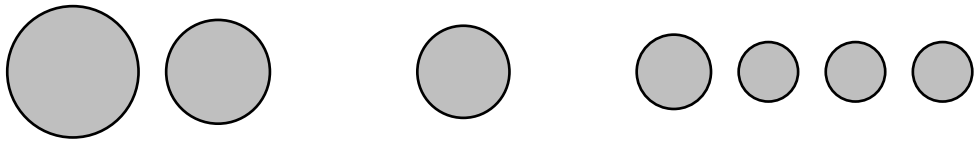

Objective:

$$\max \sum_{i \in Apps} Utility_i * y_i$$

y_i = 1 if application i is allocated, else 0

Privacy Resource Allocation

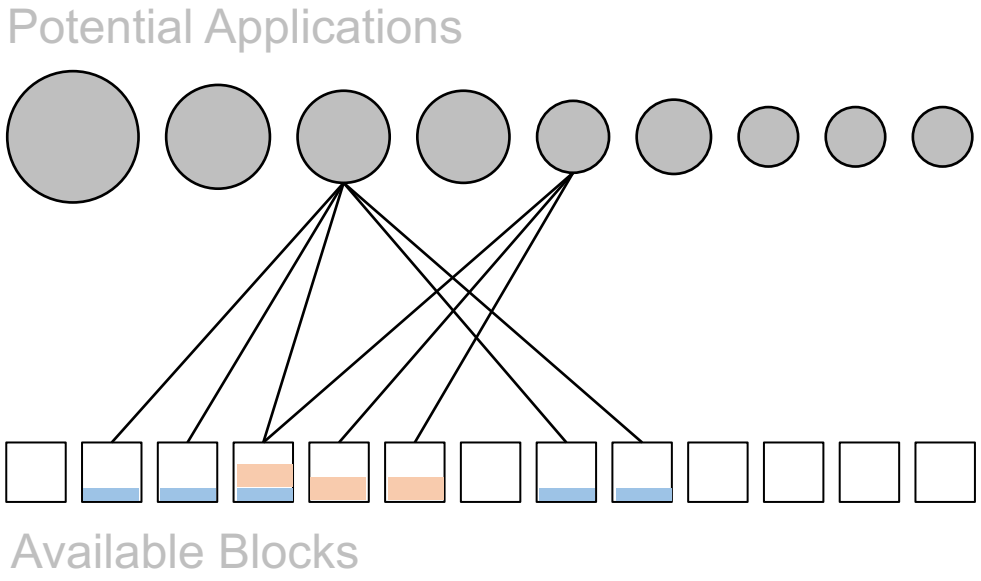

Multidimensional Knapsack Problem

Objective:

$$\max \sum_{i \in Apps} Utility_i * y_i$$

y_i = 1 if application i is allocated, else 0

Budget Constraints:

$$s.t. \sum_{i \in Apps} \varepsilon_{ij} * y_i \leq Budget_j \quad \forall j \in Blocks$$

Privacy cost of application i for block j

* for simplicity we show the cost in ε- DP rather than RDP

Privacy Resource Allocation

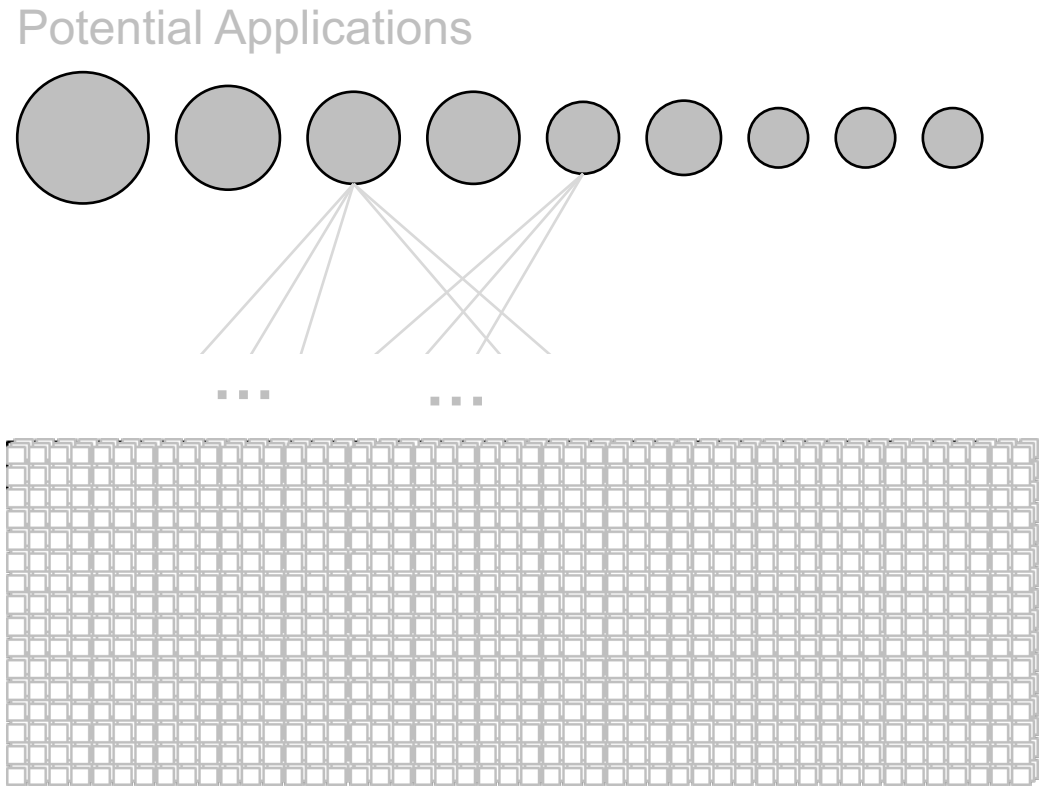

Multidimensional Knapsack Problem

Objective:

$$\max \sum_{i \in Apps} Utility_i * y_i$$

y_i = 1 if application i is allocated, else 0

Budget Constraints:

$$s.t. \sum_{i \in Apps} \varepsilon_{ij} * y_i \leq Budget_j \quad \forall j \in Blocks$$

Privacy cost of application i for block j

* for simplicity we show the cost in ε- DP rather than RDP

Resource Allocation: Taming the Complexity

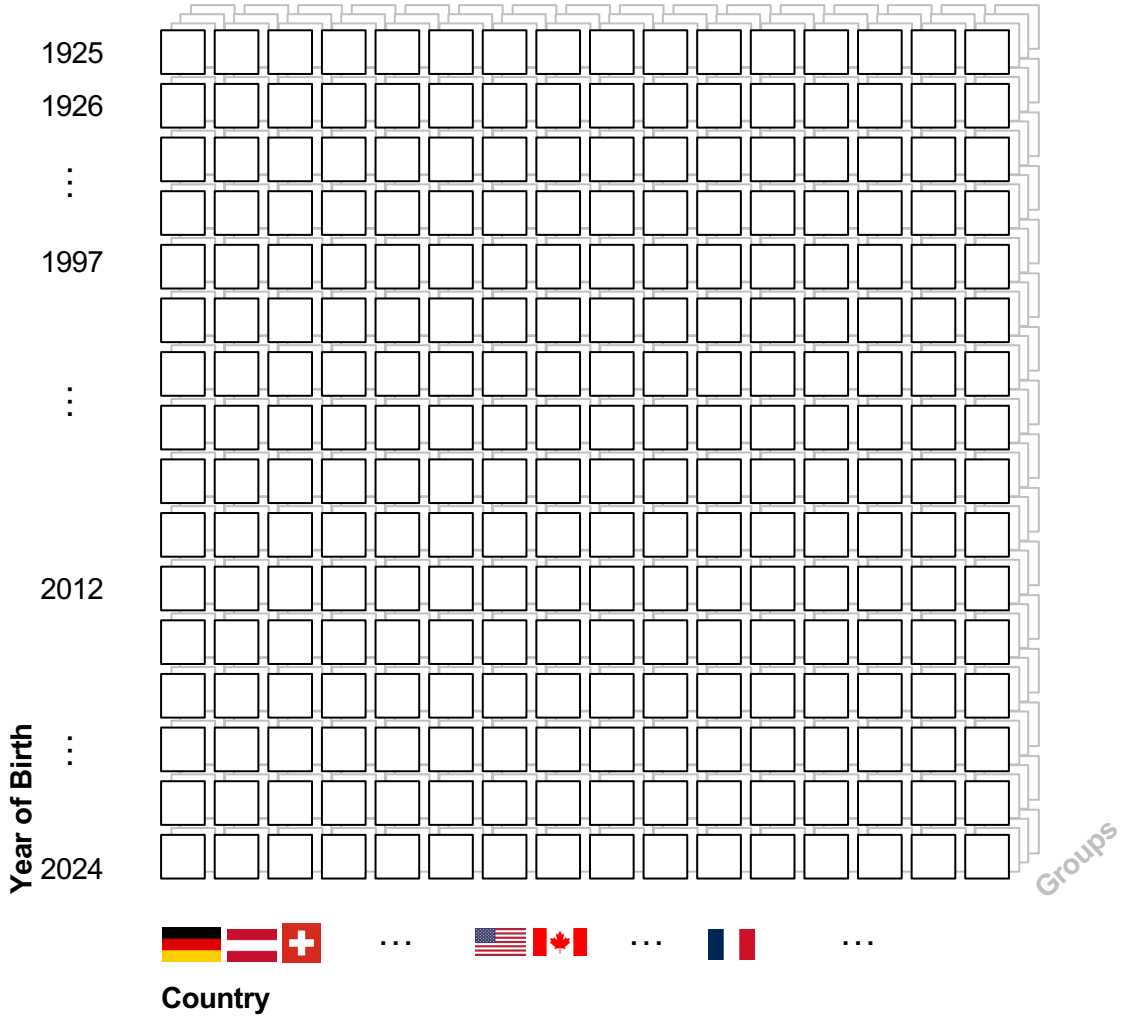

Resource Allocation: Taming the Complexity

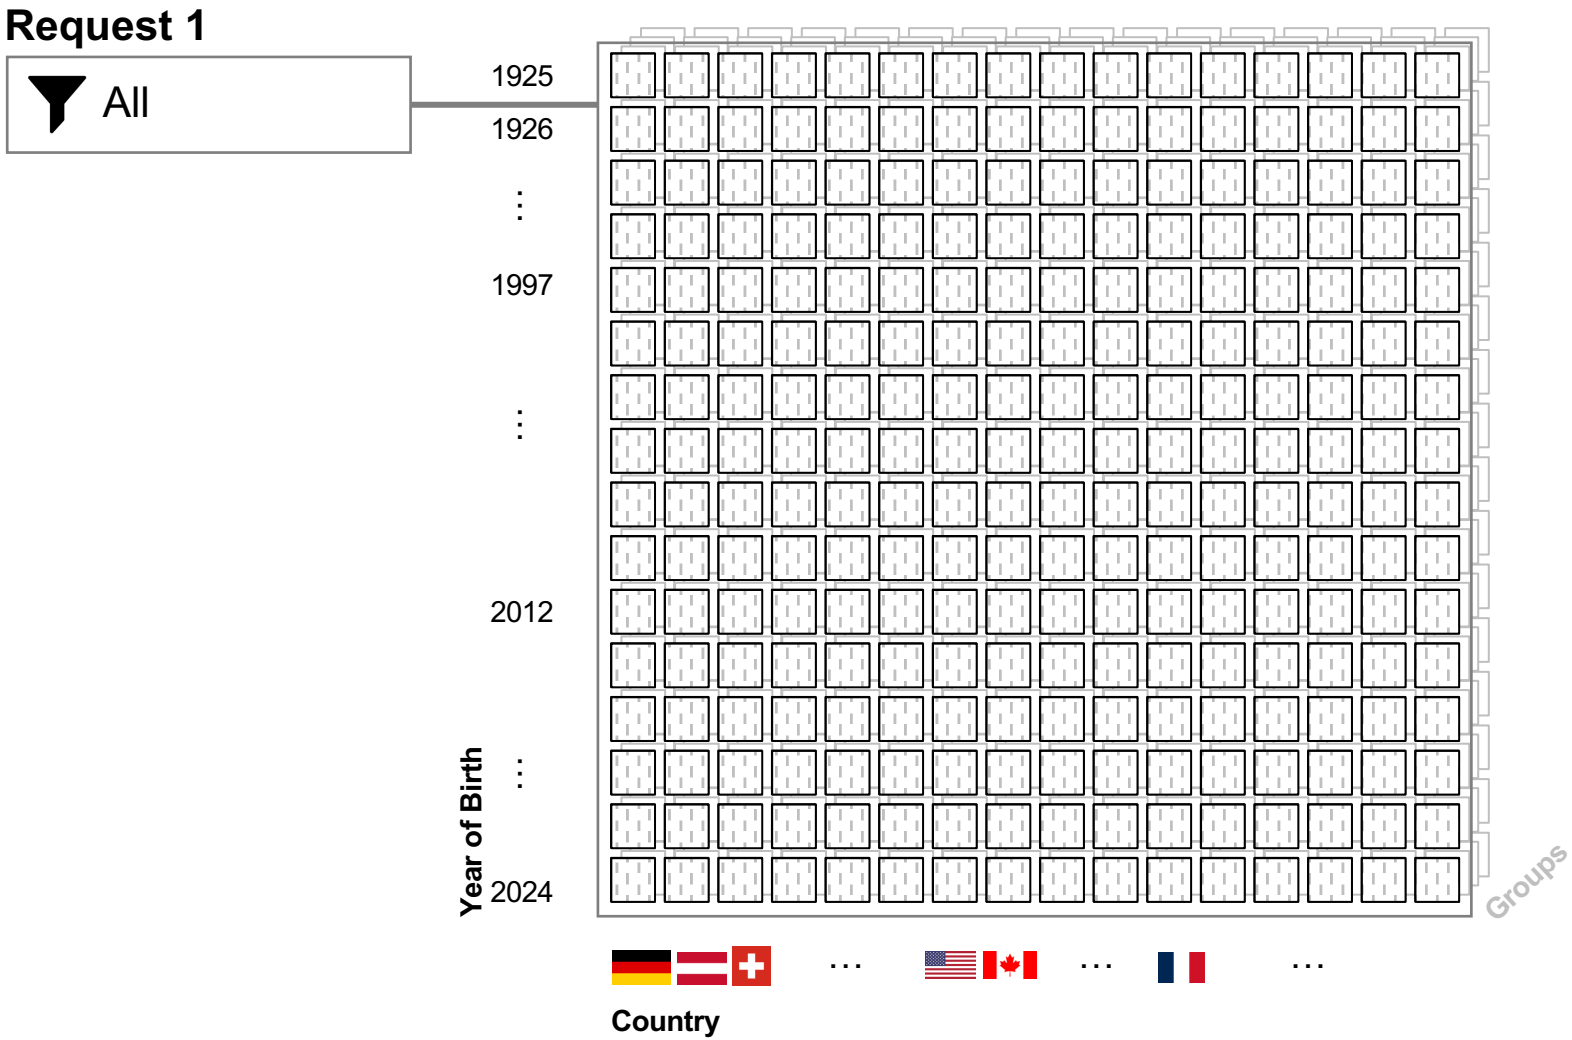

Resource Allocation: Taming the Complexity

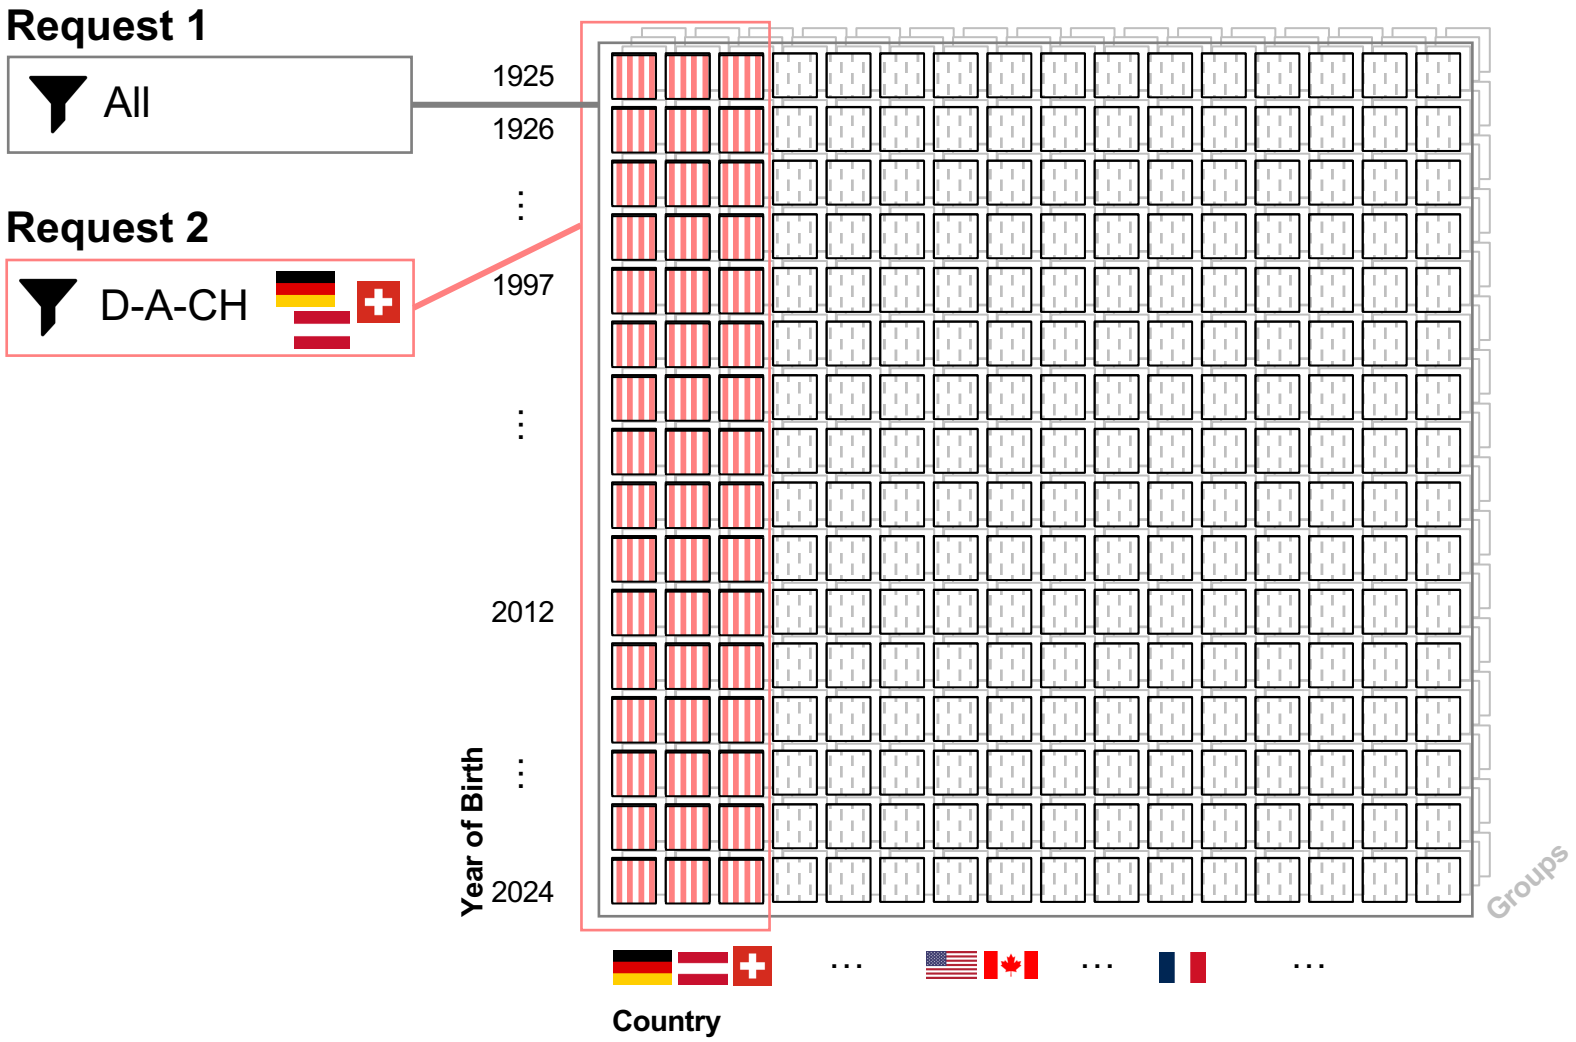

Resource Allocation: Taming the Complexity

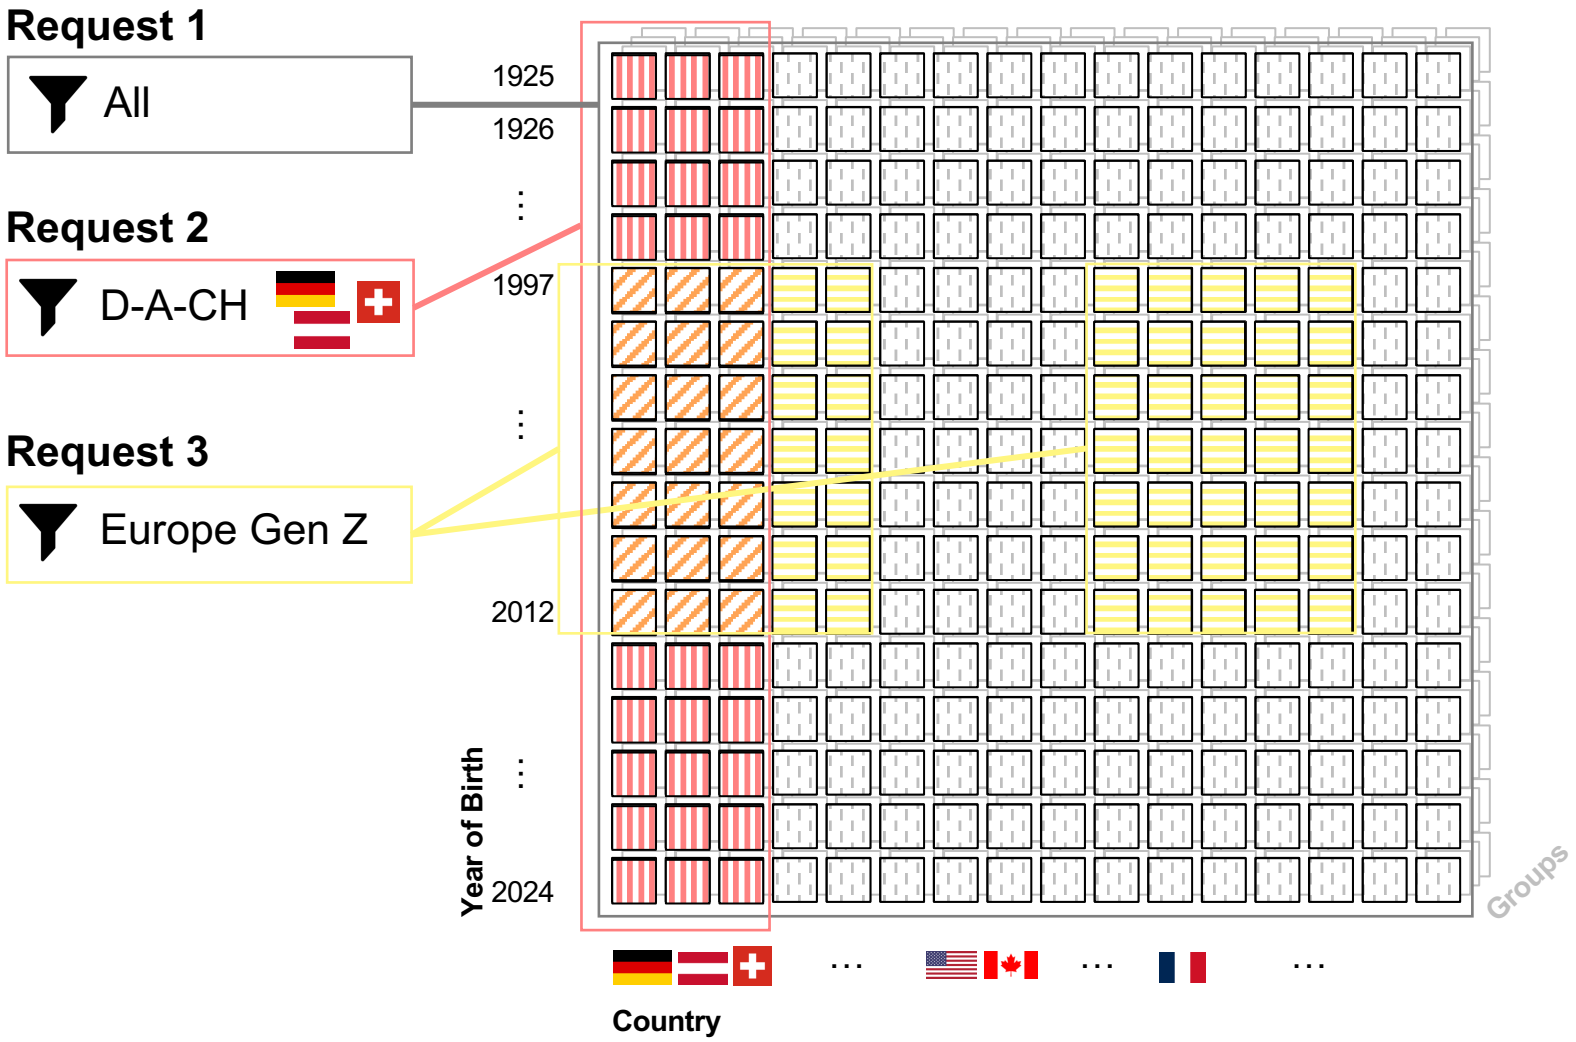

Resource Allocation: Taming the Complexity

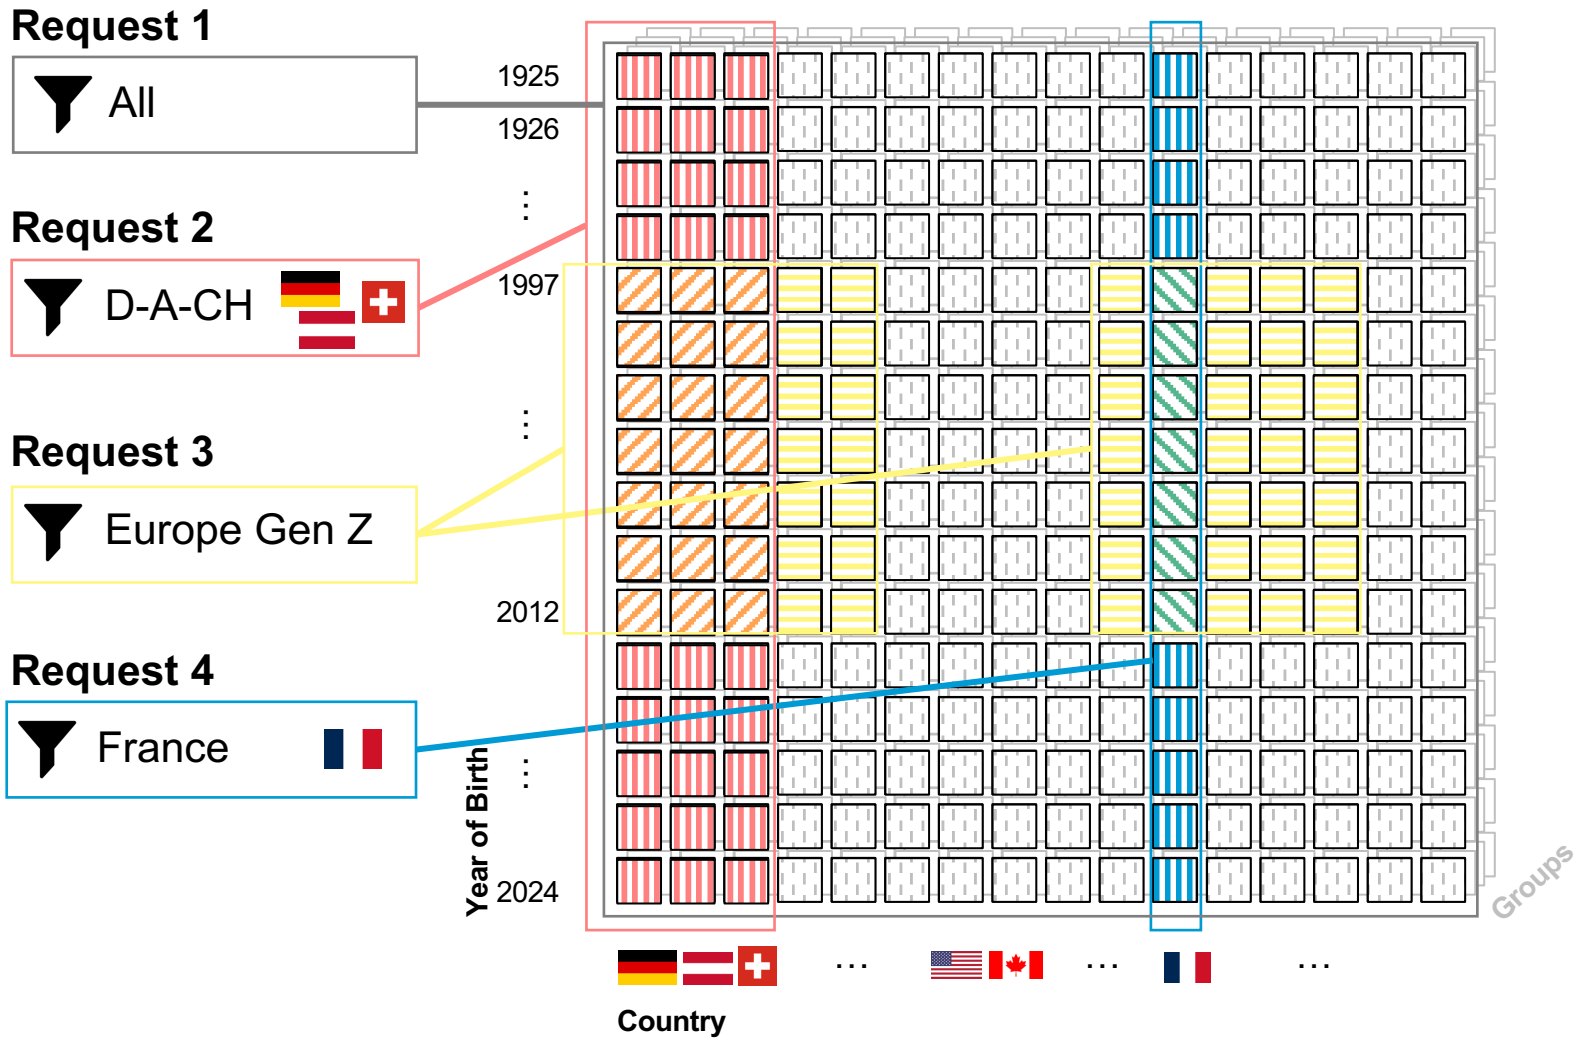

Resource Allocation: Taming the Complexity

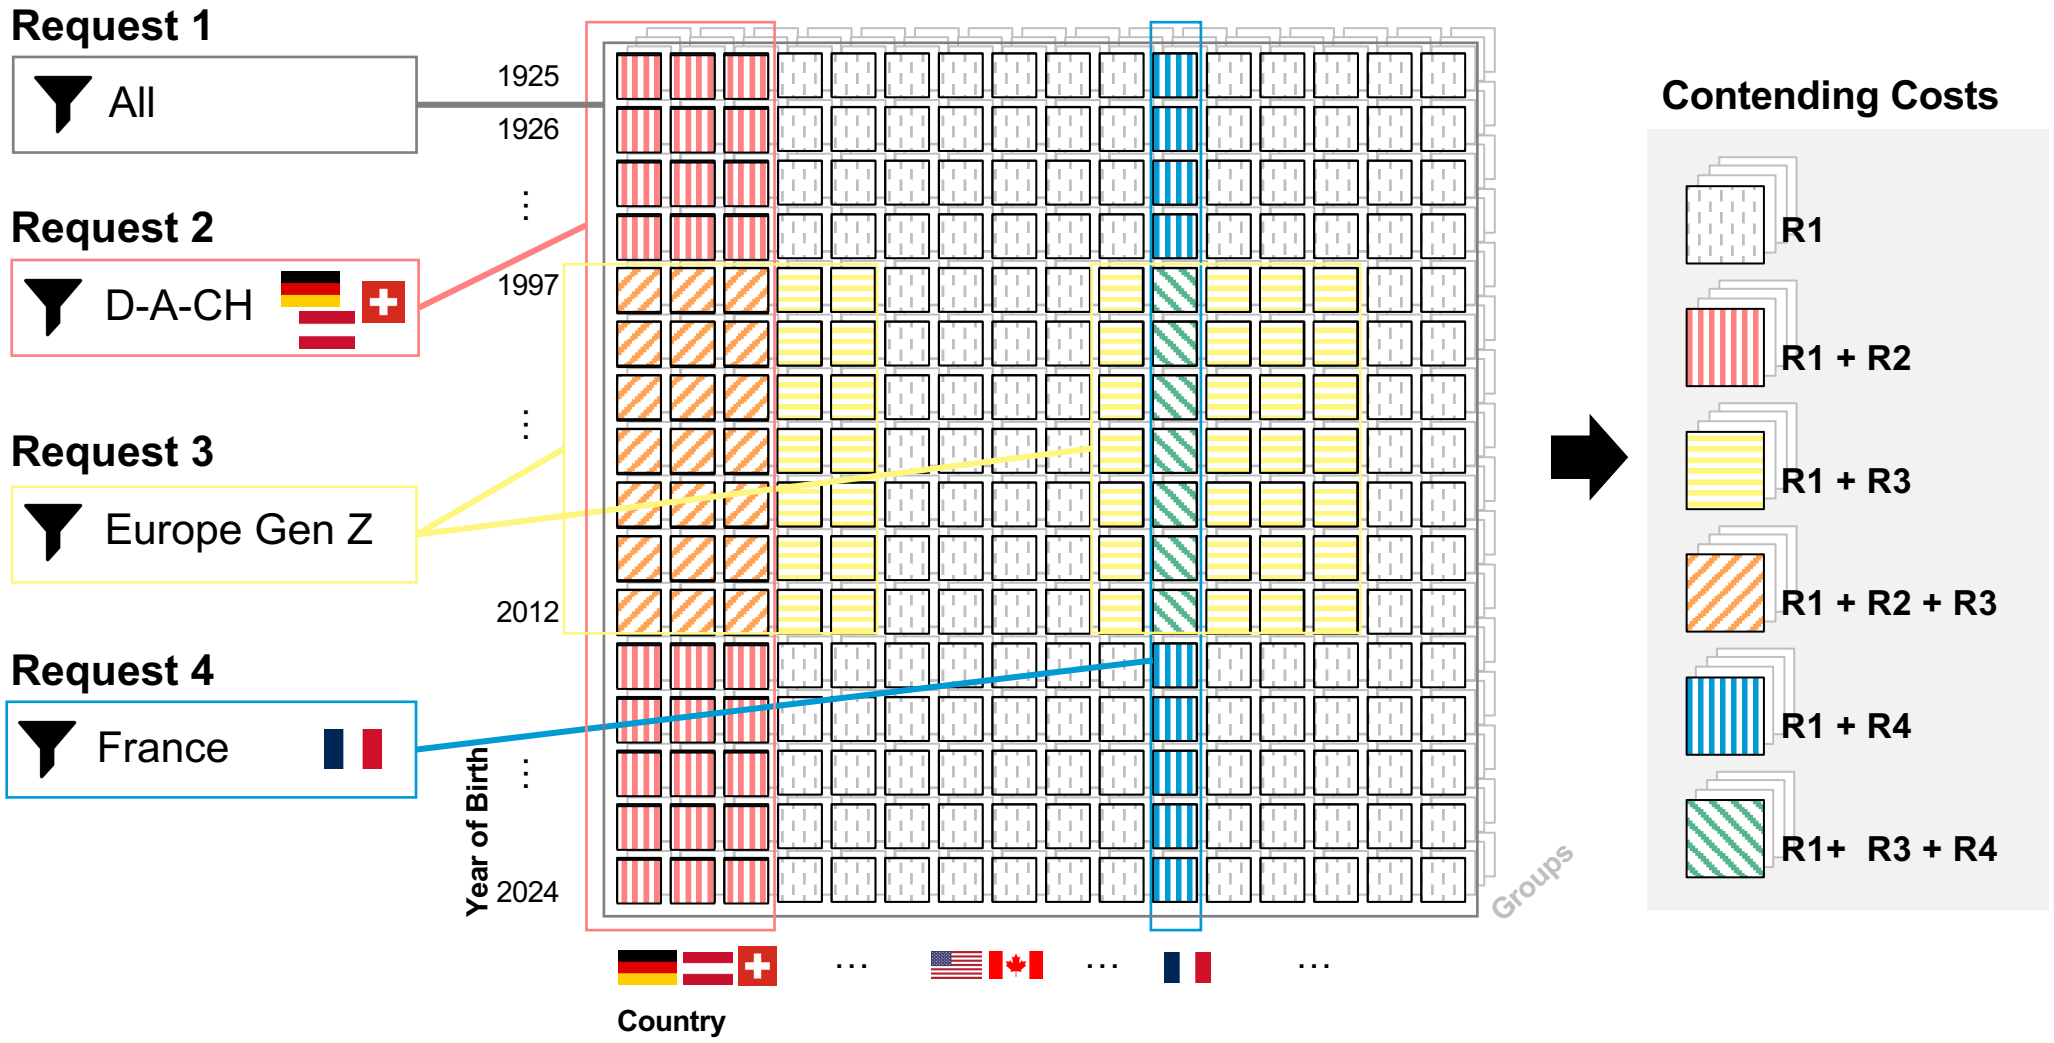

Resource Allocation: Taming the Complexity

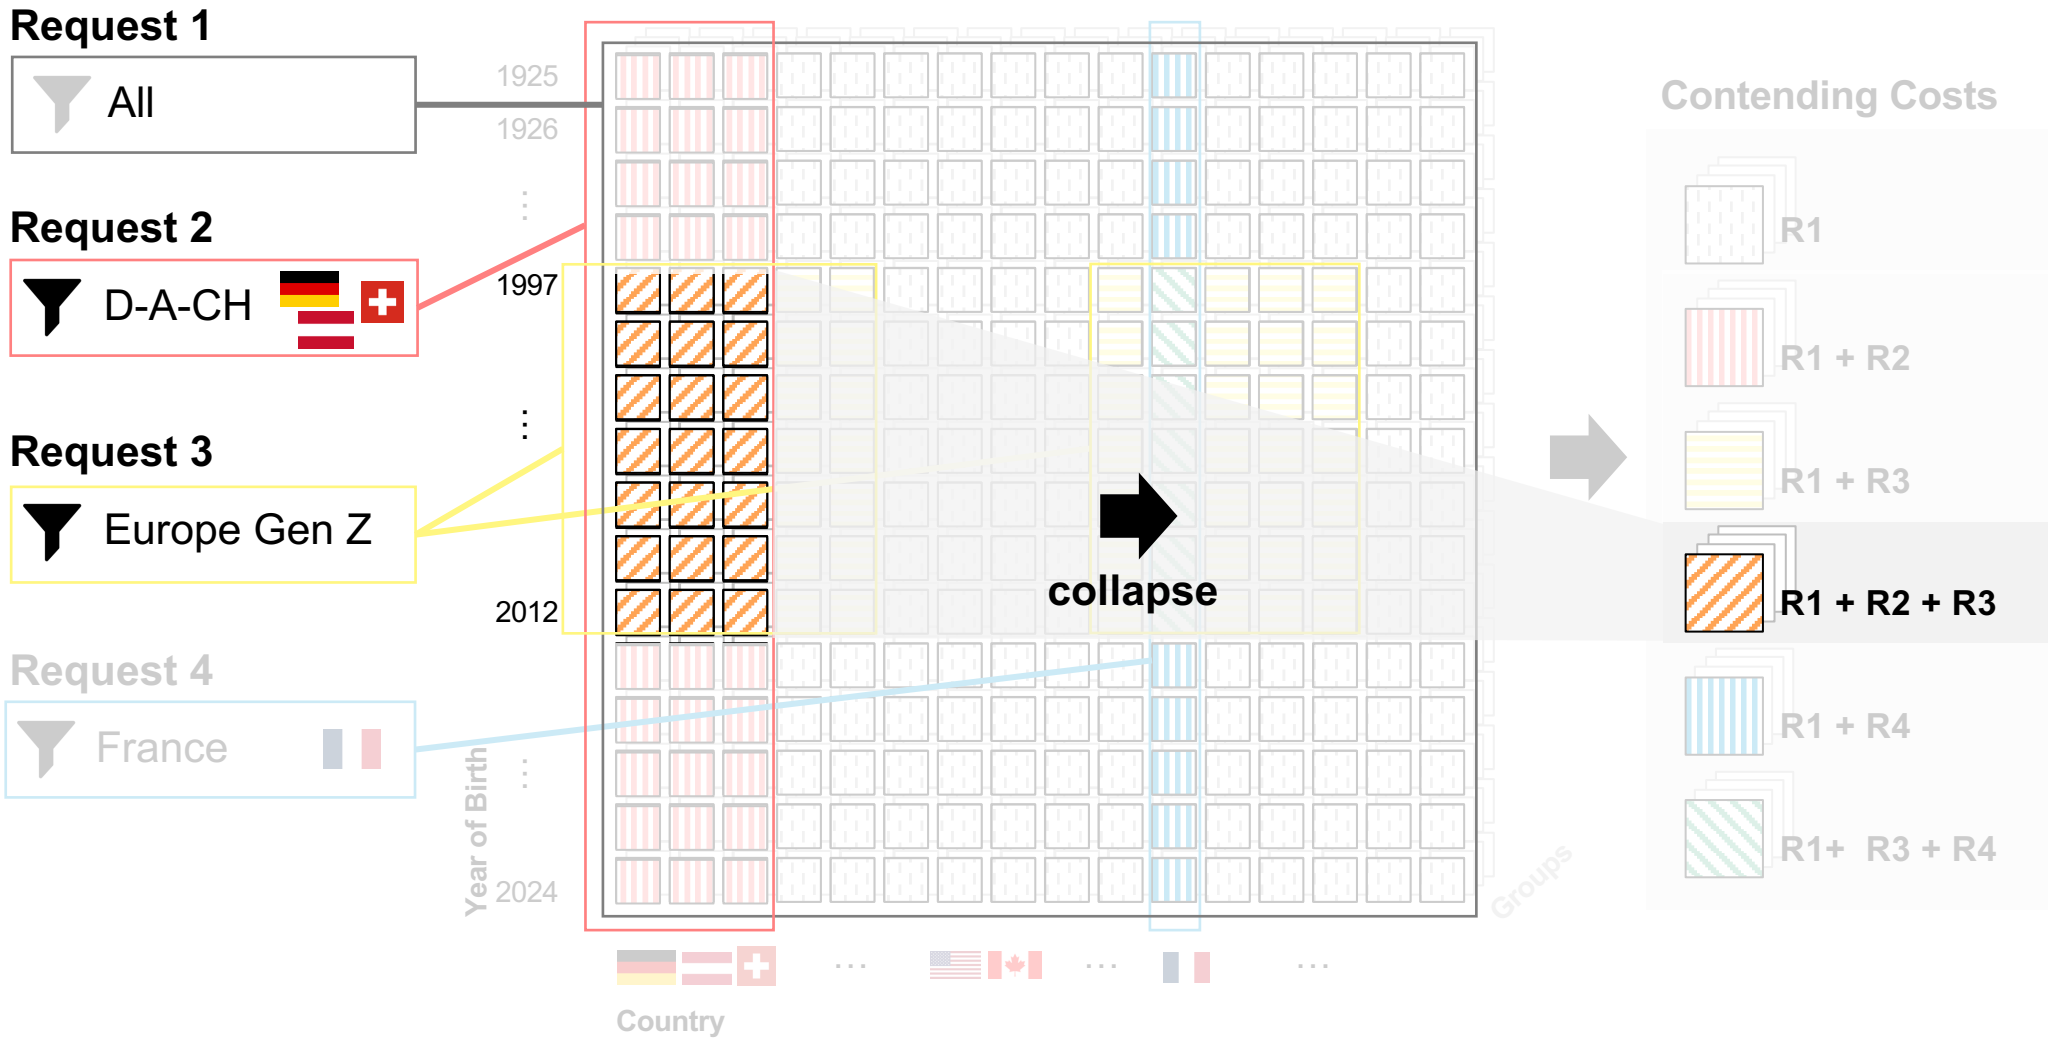

Resource Allocation: Taming the Complexity

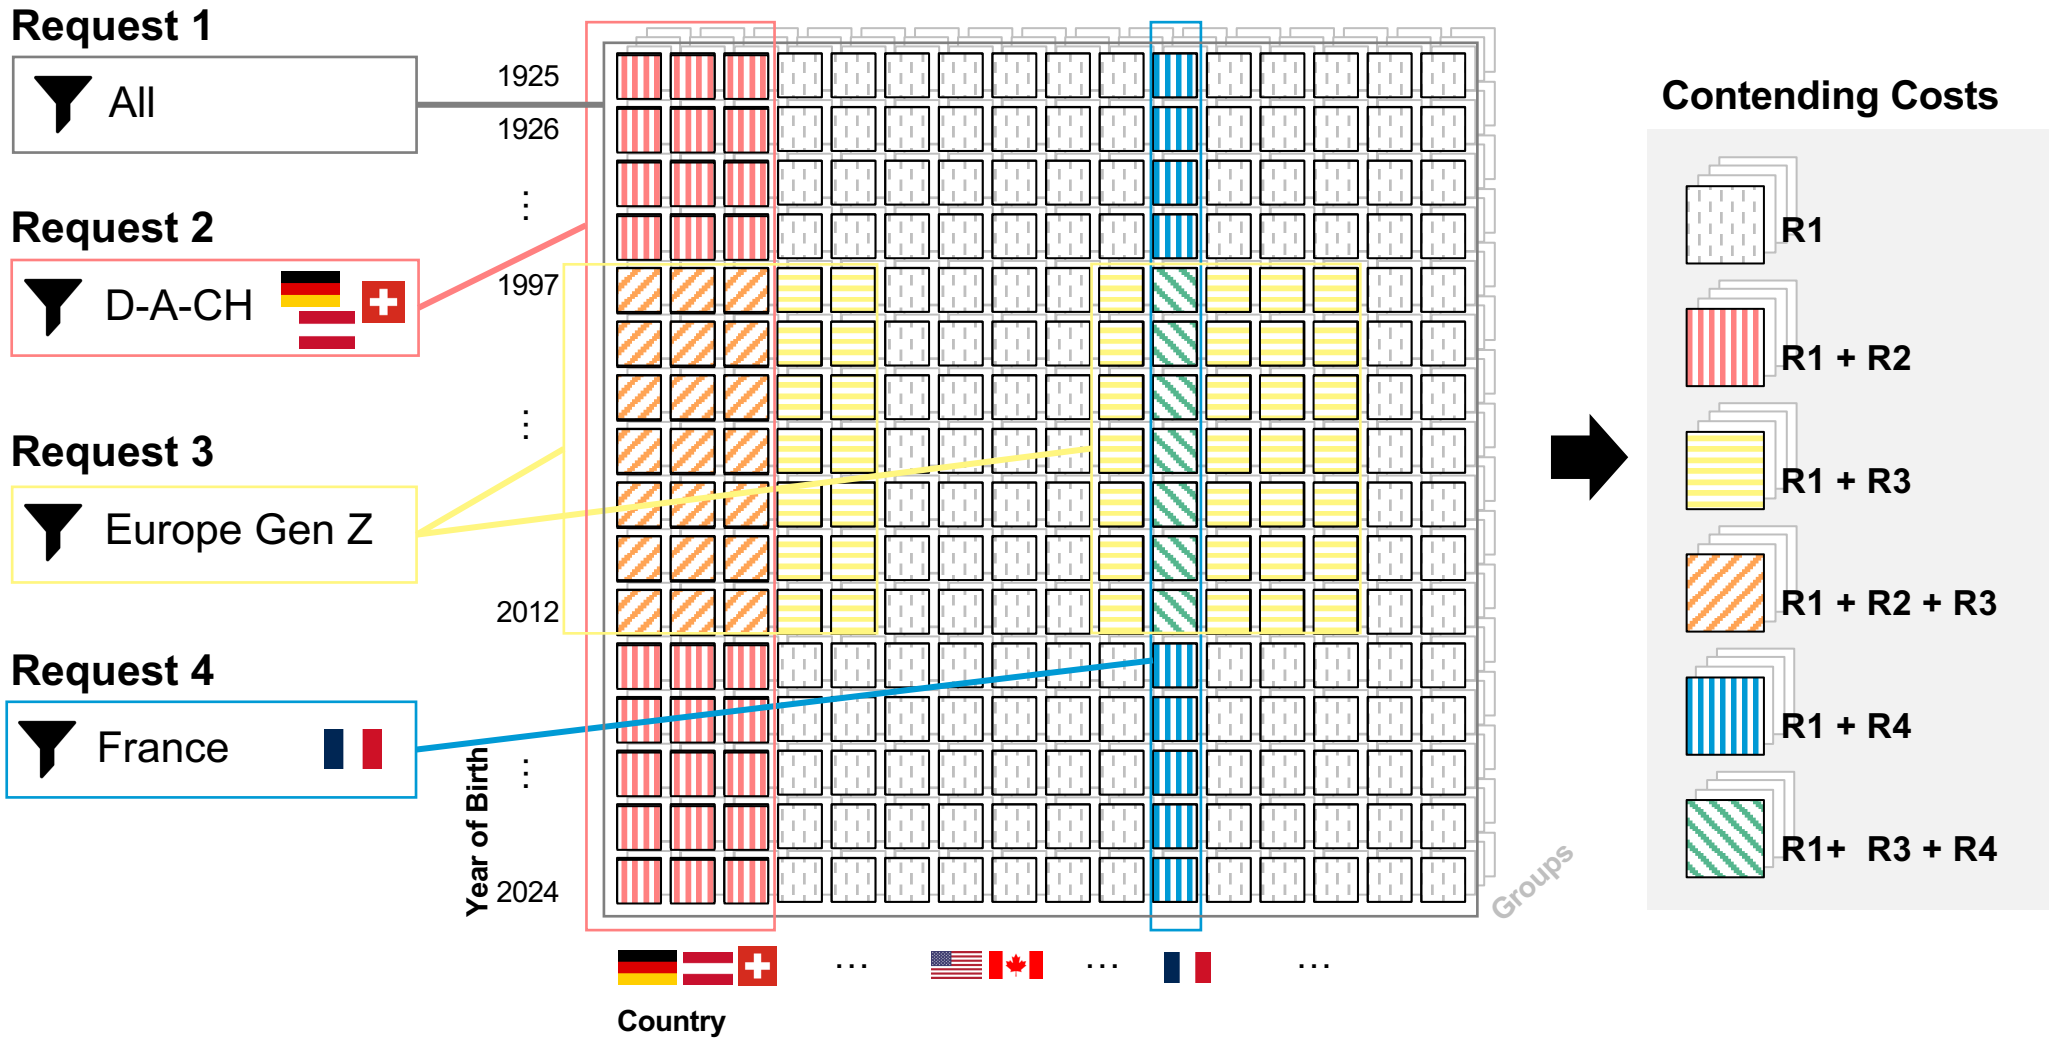

Resource Allocation: Taming the Complexity

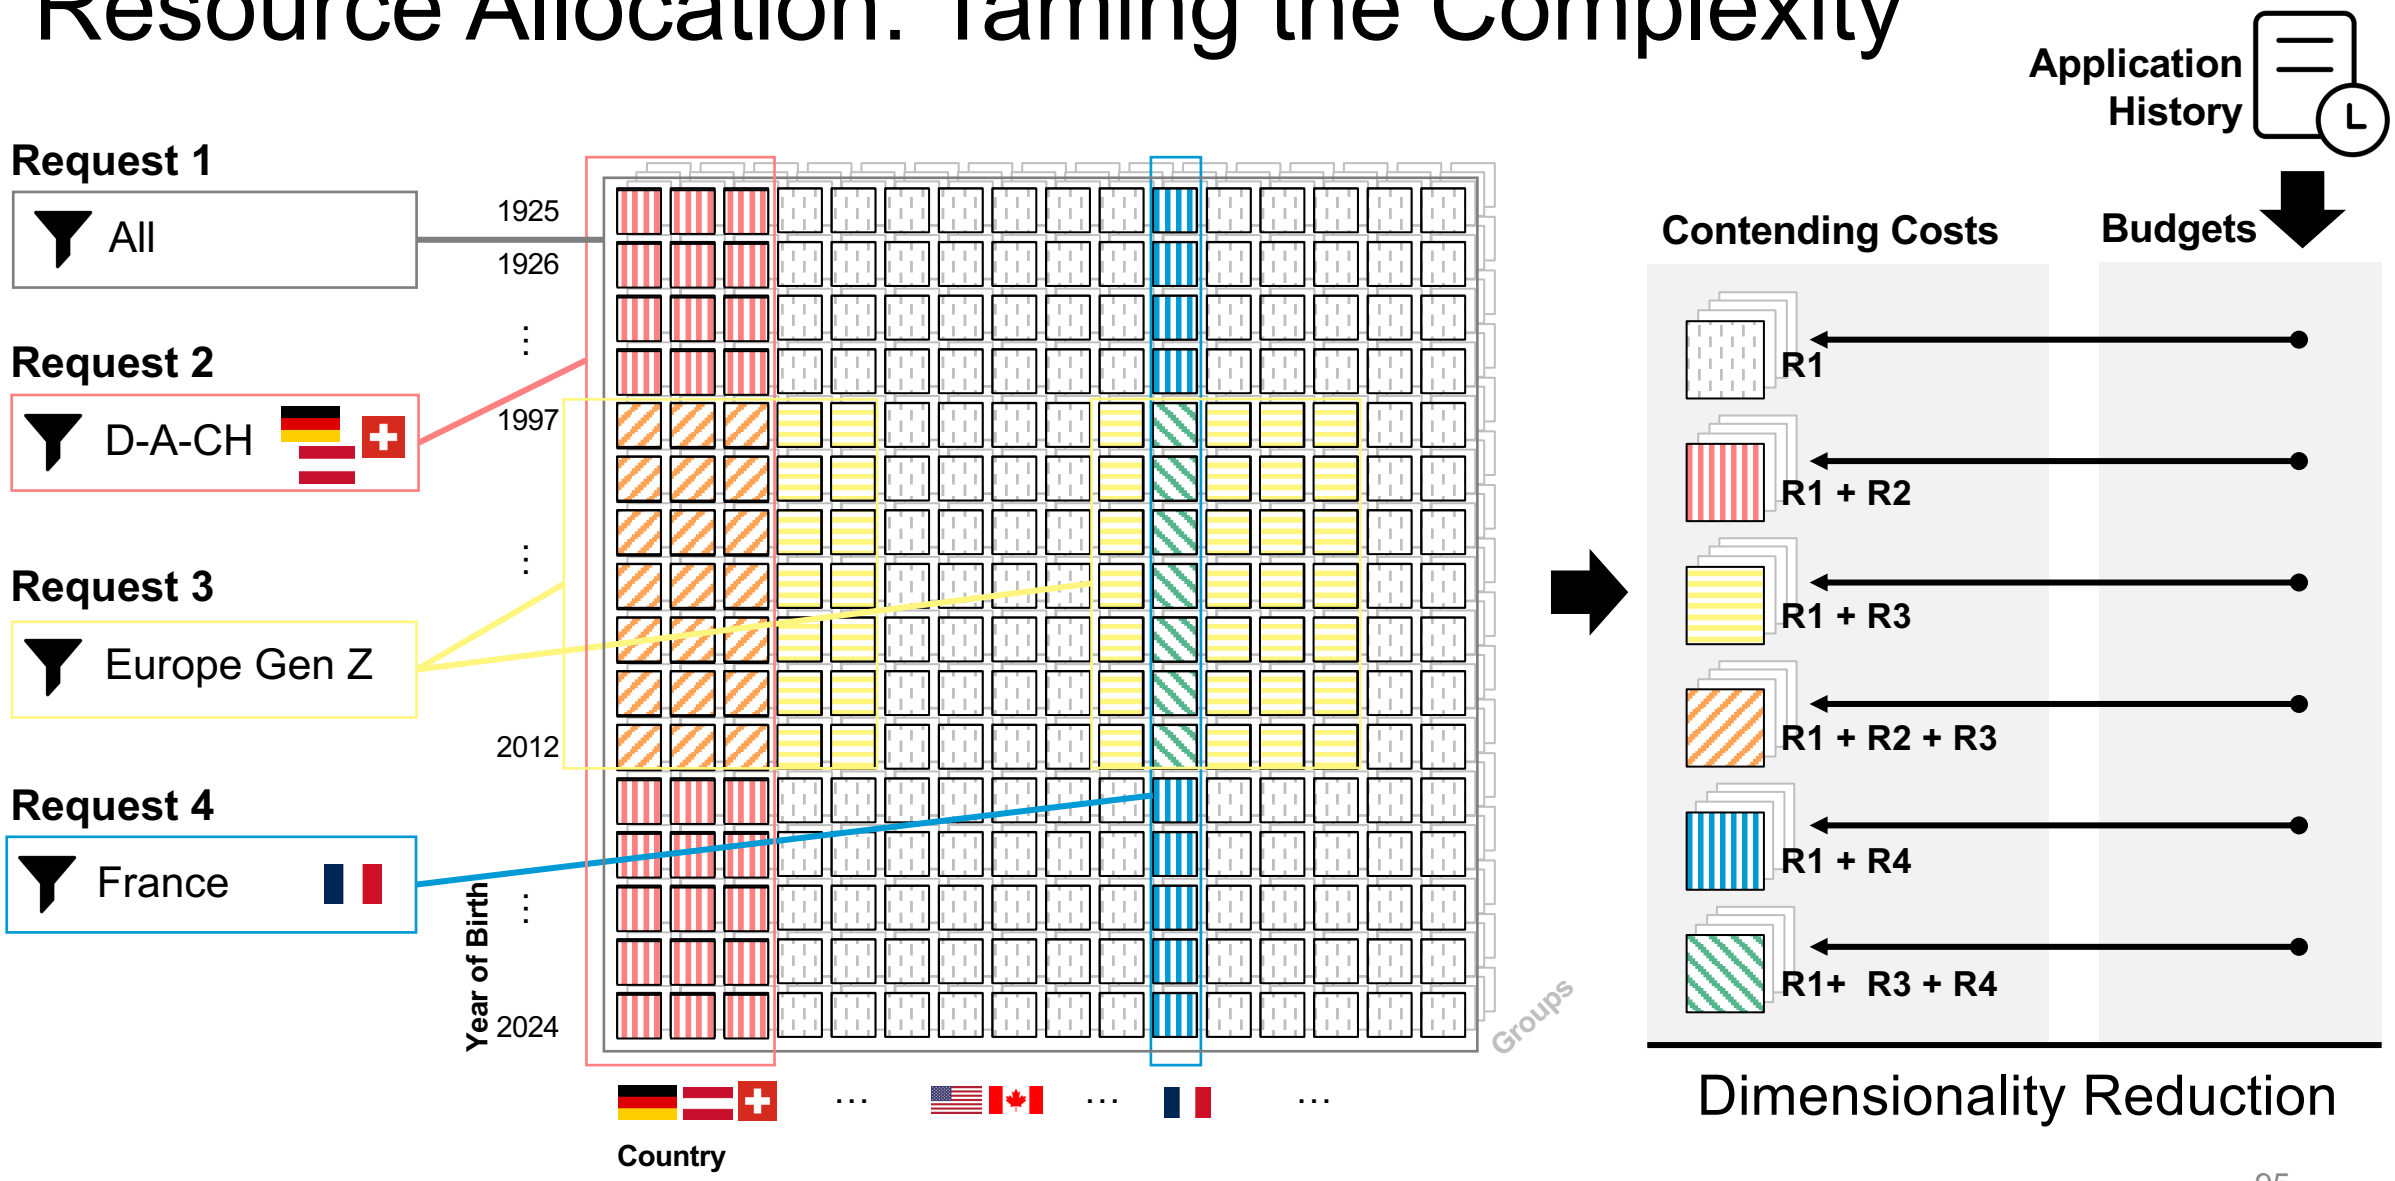

Resource Allocation: Taming the Complexity

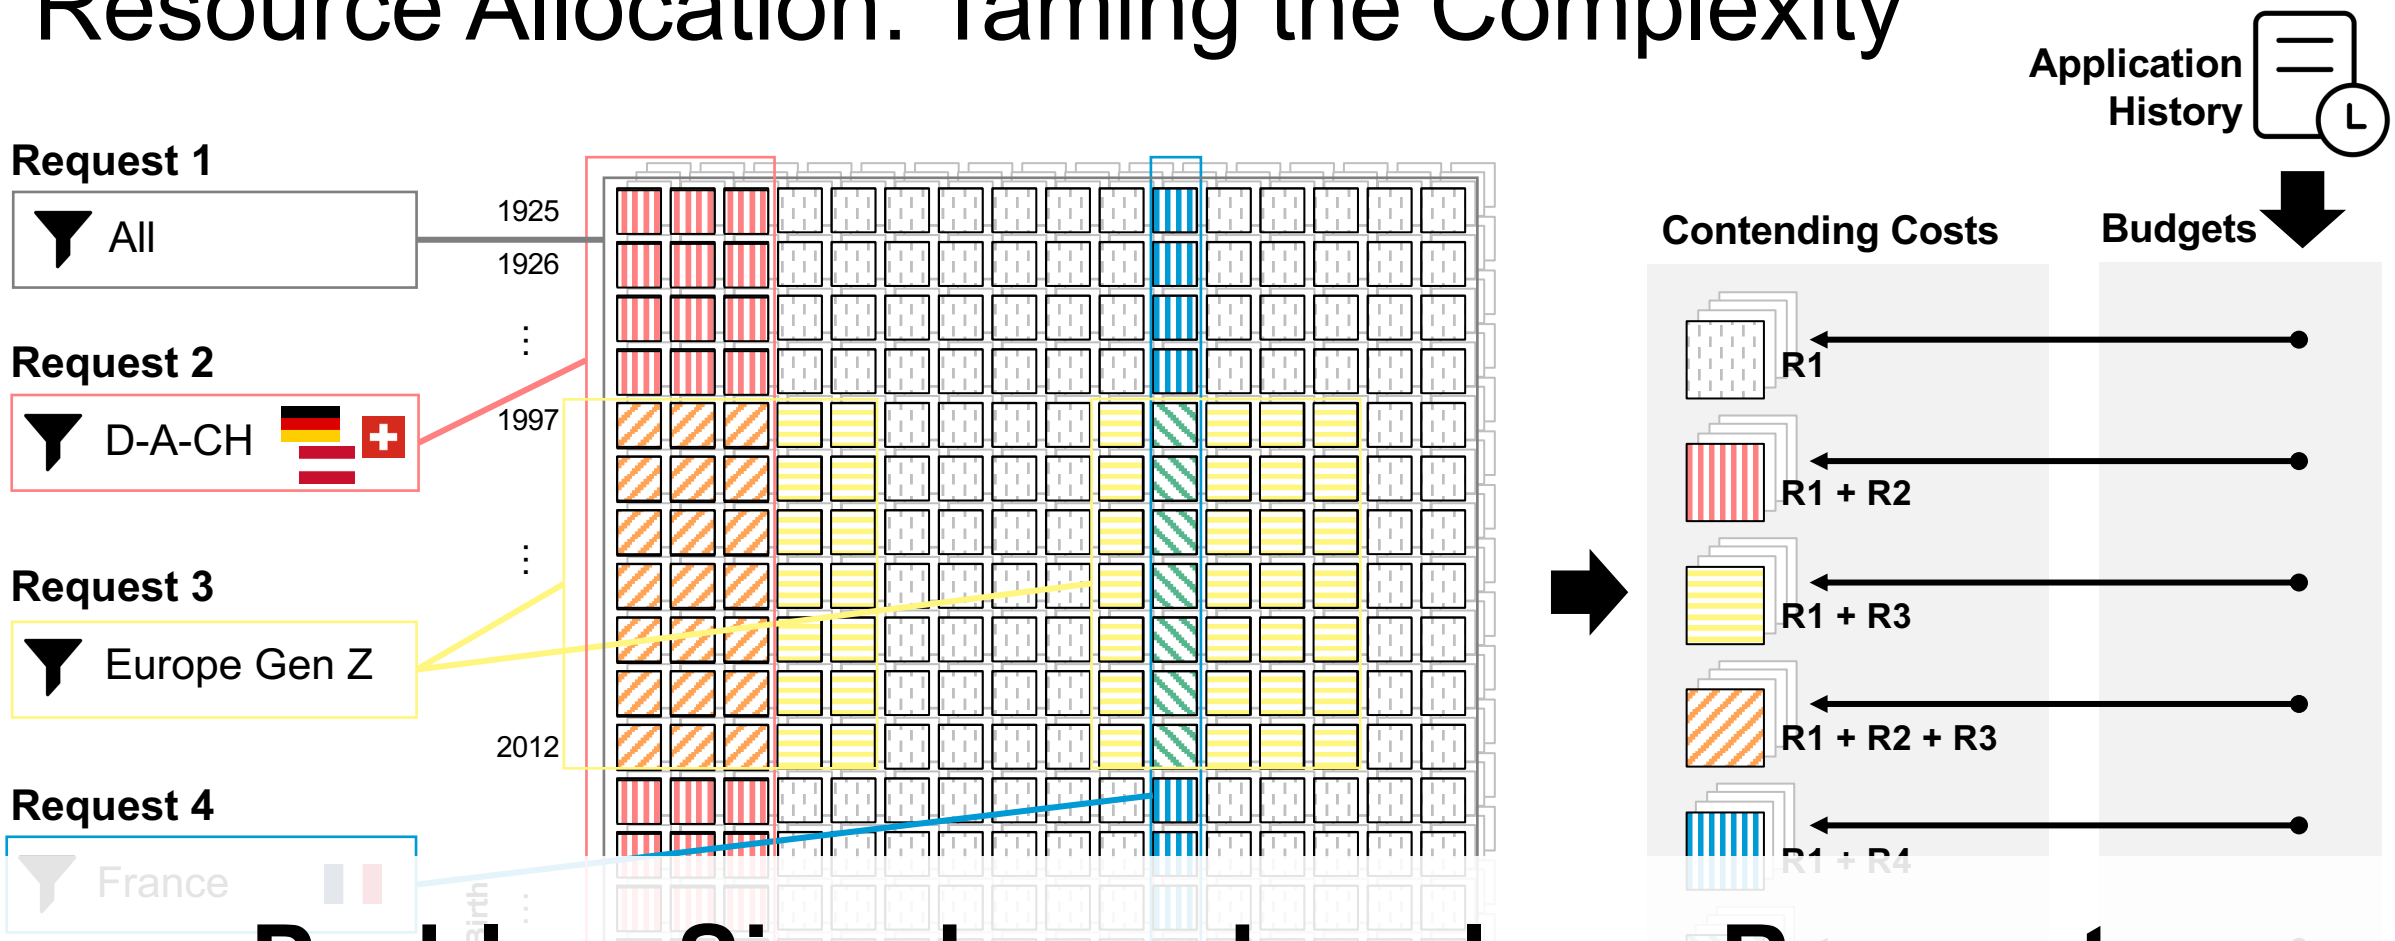

Problem Size depends only on Requests

rather than the domain size of the partitioning attributes

1

Unified System Architecture

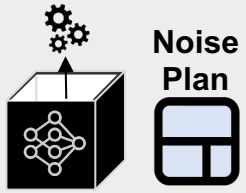

Unified Application Layer

Privacy ID with Partitioning Attributes

`(a1, a2, a3)`

Unified Data Layer

2

Access Pattern for Privacy Analysis

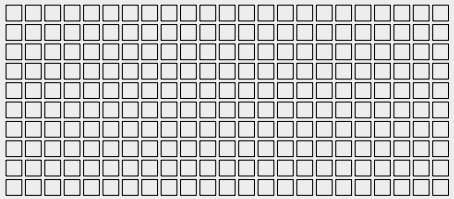

Fine-Grained Block Composition

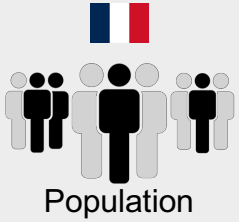

Subsampling

1

Unified System Architecture

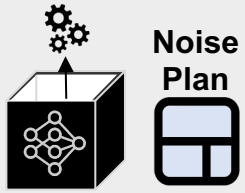

Unified Application Layer

Privacy ID with Partitioning Attributes

(a1, a2, a3)

Unified Data Layer

2

Access Pattern for Privacy Analysis

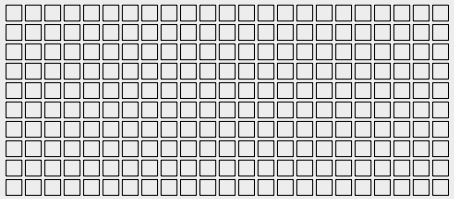

Fine-Grained Block Composition

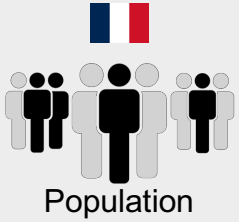

Subsampling

3

System Continuity Guarantee

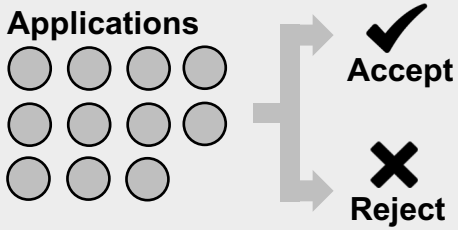

Periodic Planning

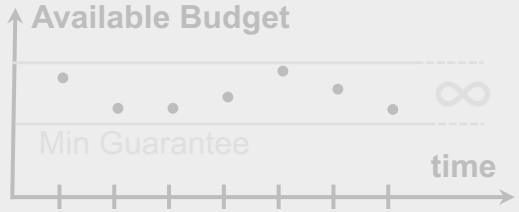

Budget Guarantee

1

Unified System Architecture

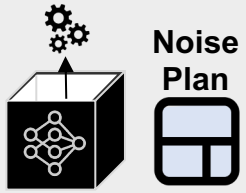

Unified Application Layer

Privacy ID with Partitioning Attributes

(a1, a2, a3)

Unified Data Layer

2

Access Pattern for Privacy Analysis

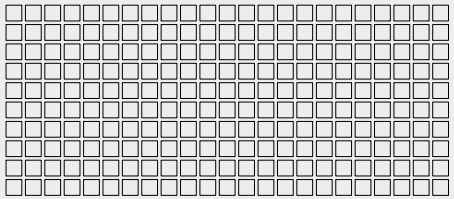

Fine-Grained Block Composition

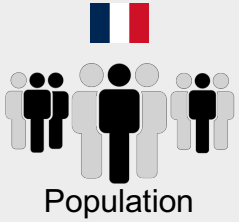

Subsampling

3

System Continuity Guarantee

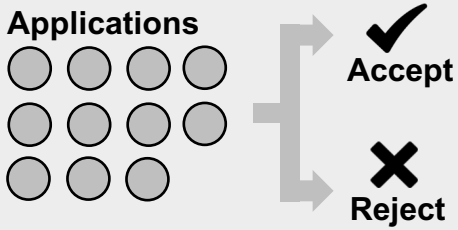

Periodic Planning

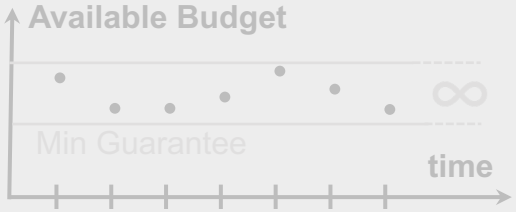

Budget Guarantee

4

Resource Allocation

$$\max \sum_{i \in Apps} Utility_i * y_i$$

$$s.t. \sum_{i \in Apps} \varepsilon_{ij} * y_i \leq Budget_j \quad \forall j$$

for simplicity in ε- DP

Multidimensional Knapsack

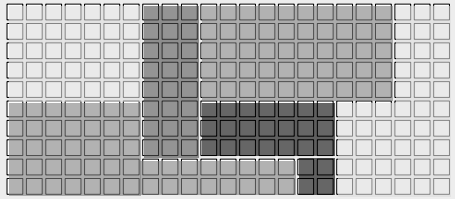

Dimensionality Reduction

Evaluation Scenario

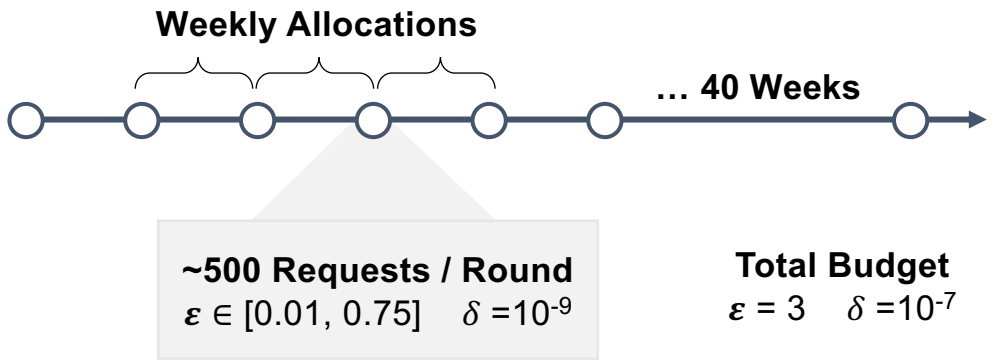

Evaluation Scenario

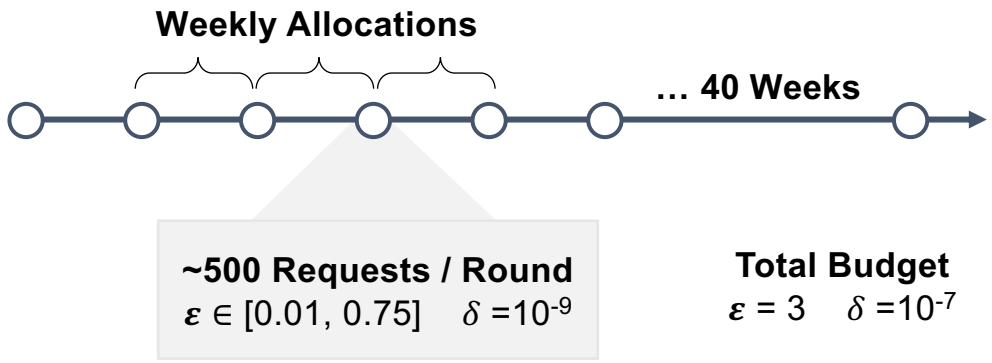

Baseline

PrivateKube
[Luo et al. OSDI'21]

kubernetes

Fixed Coarse-Grained Privacy Analysis

Workload: Mixture of Analytics and ML Tasks

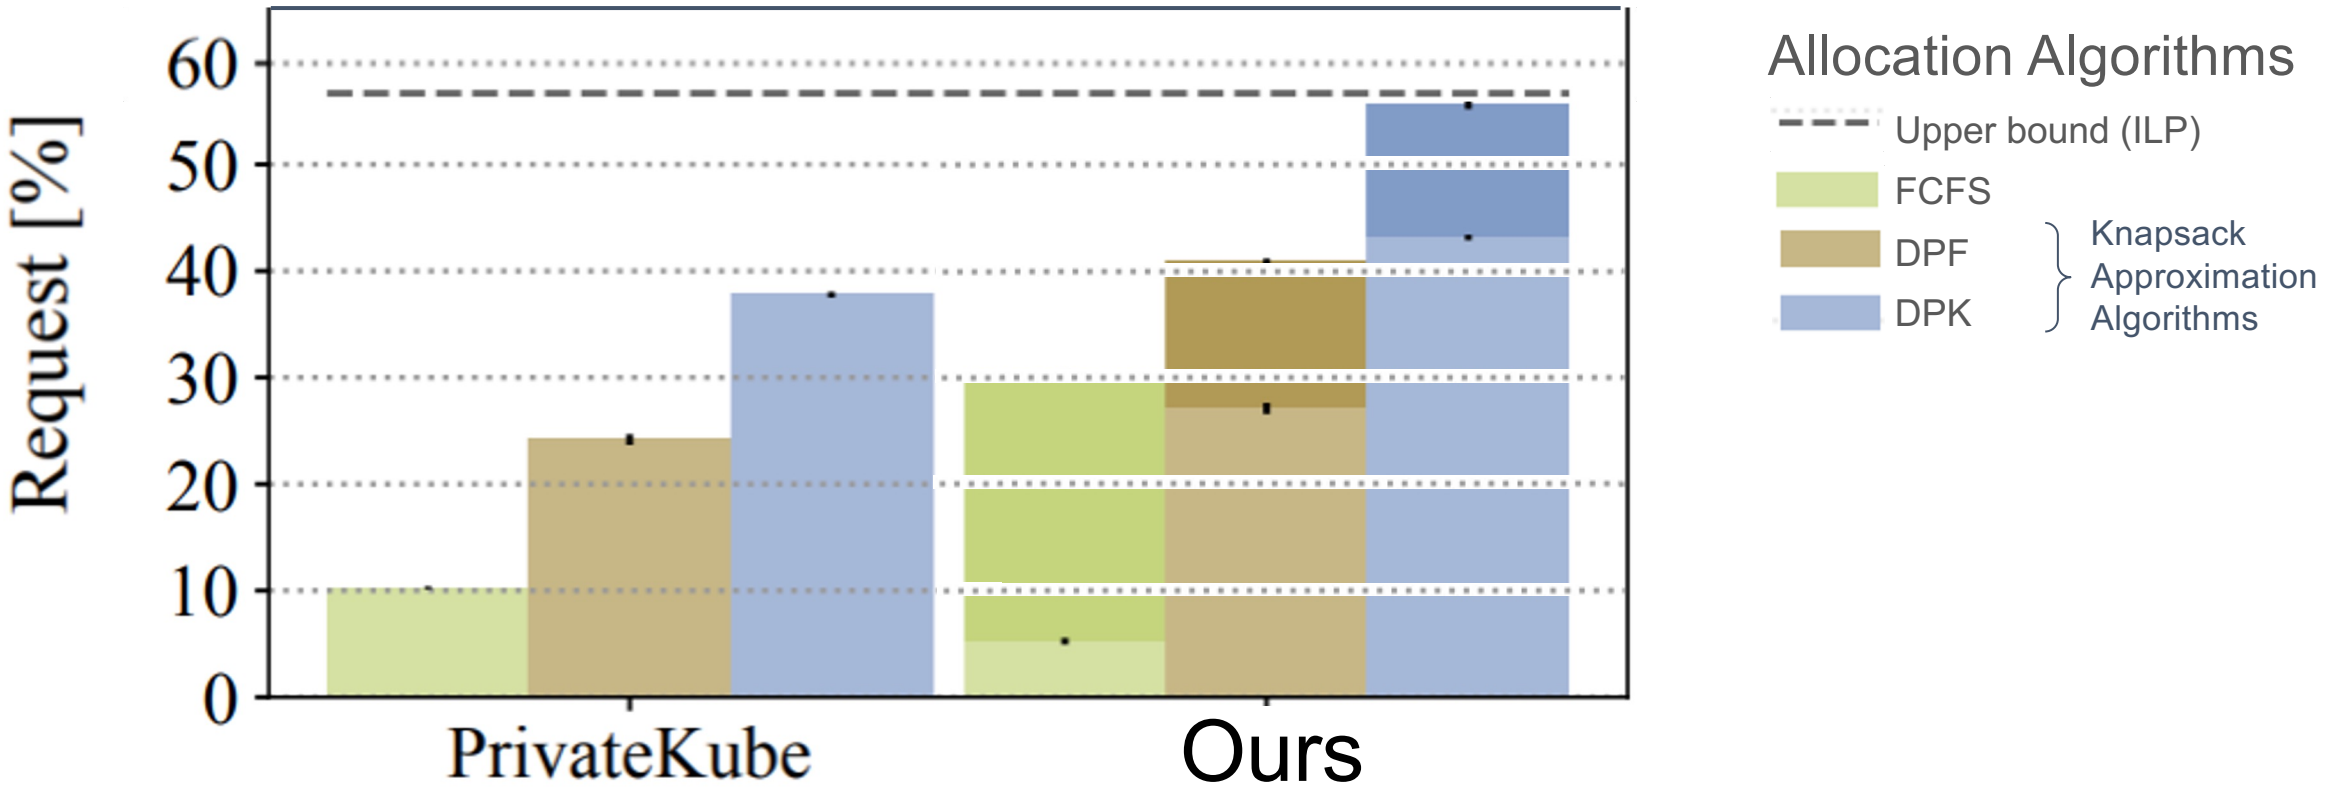

Workload: Mixture of Analytics and ML Tasks

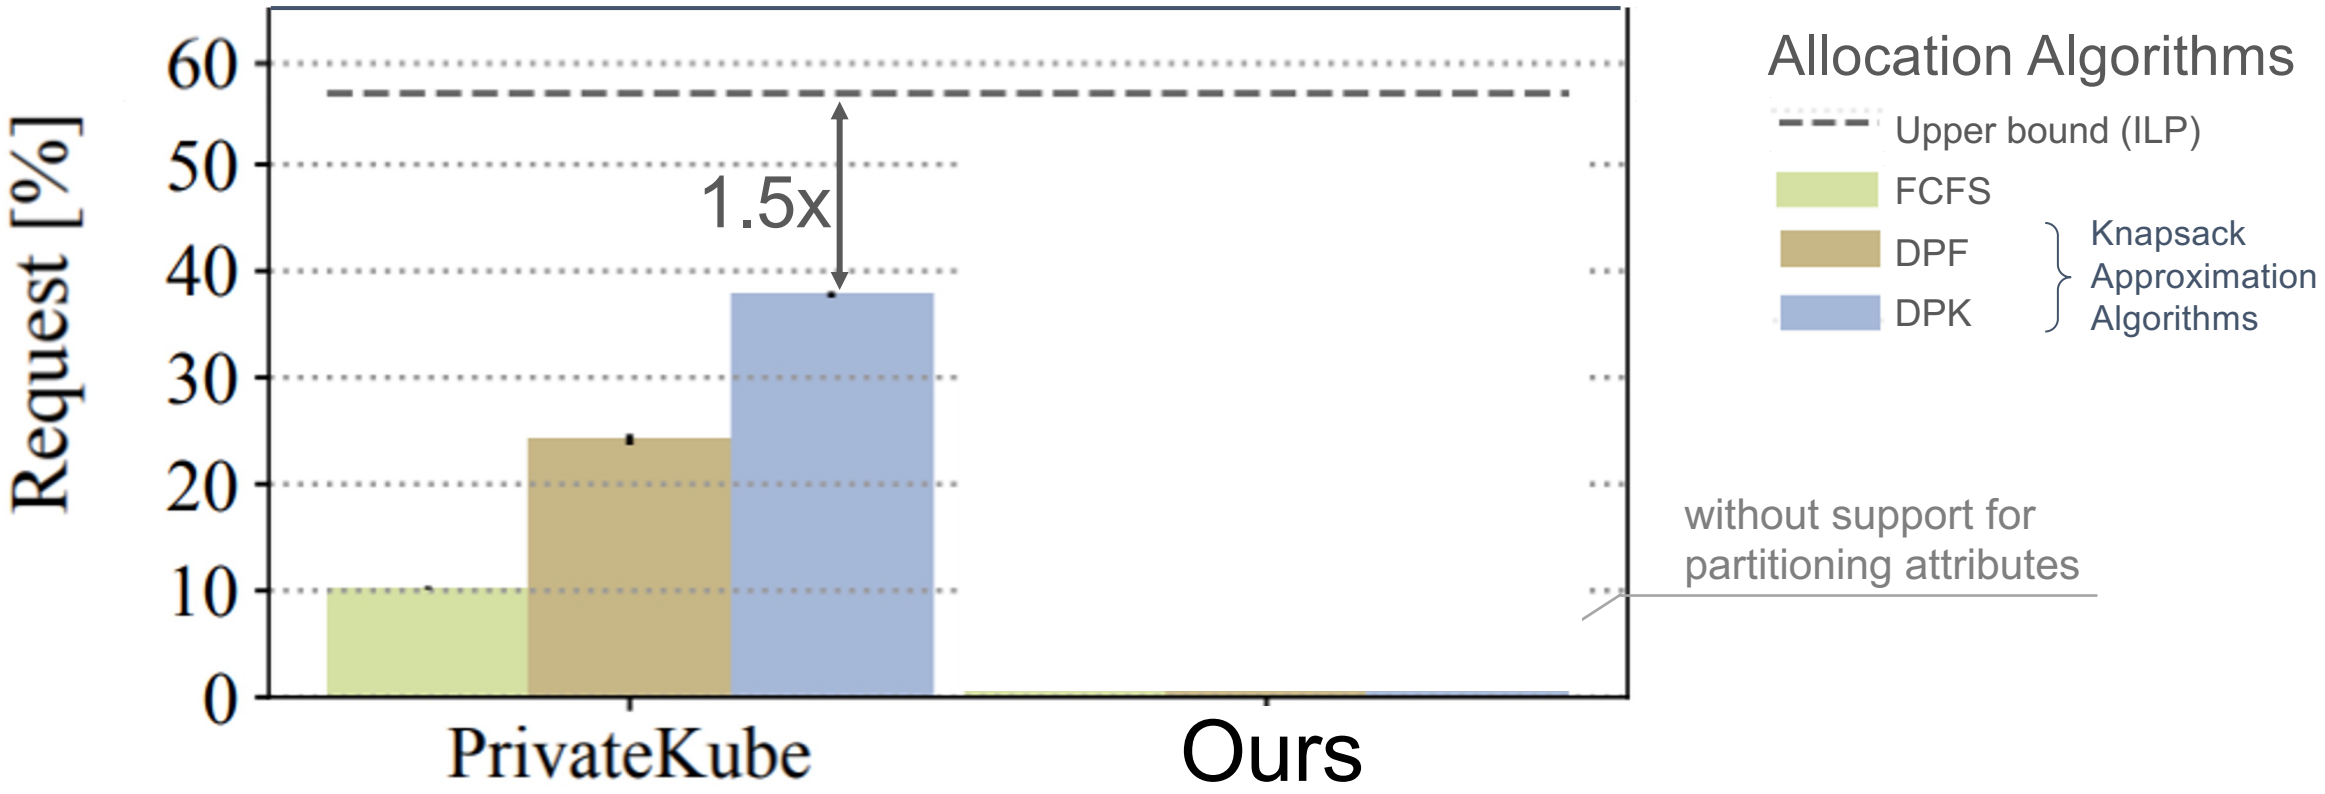

Workload: Mixture of Analytics and ML Tasks

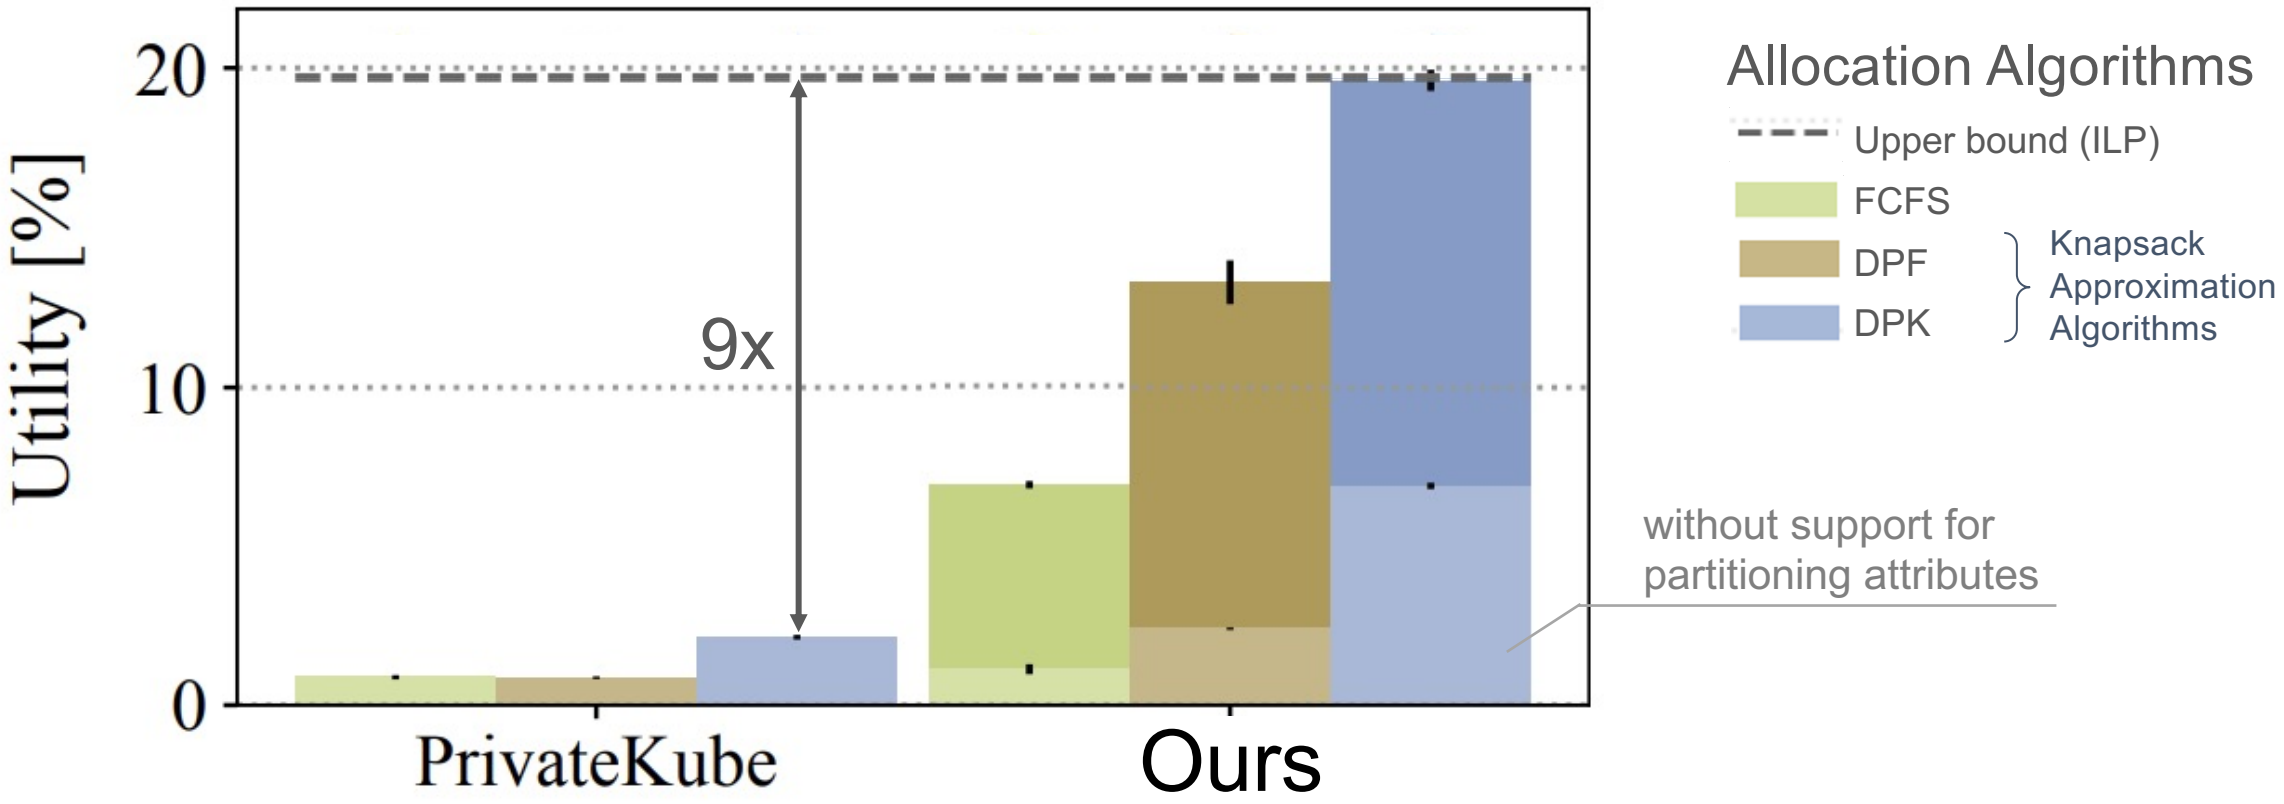

Cohere: Managing DP in Large - Scale Systems