On Private Conversion Measurement via Global and Local DP

Maxime Vono, Senior Researcher, Criteo AI Lab
m.vono@criteo.com

Private Conversion Measurement

- **Input:** features (user, contextual & ad) + conversion label (e.g., visit, sales, basket event)
- **Goal:** optimise advertising campaigns
- **How:** learn probabilities of (rare) events to build bidding models

Private Conversion Measurement

- **Input:** features (user, contextual & ad) + conversion label (e.g., visit, sales, basket event)
- **Goal:** optimise advertising campaigns
- **How:** learn probabilities of (rare) events to build bidding models

- **Private learning constraints:** input data is obfuscated
  - noisy event-level reporting: **local** DP applied to features and/or labels
  - learning via an aggregation API: **global** DP applied to aggregated statistics
  - learning via a trusted server: **global** DP applied to learning

Differential Privacy - Recap

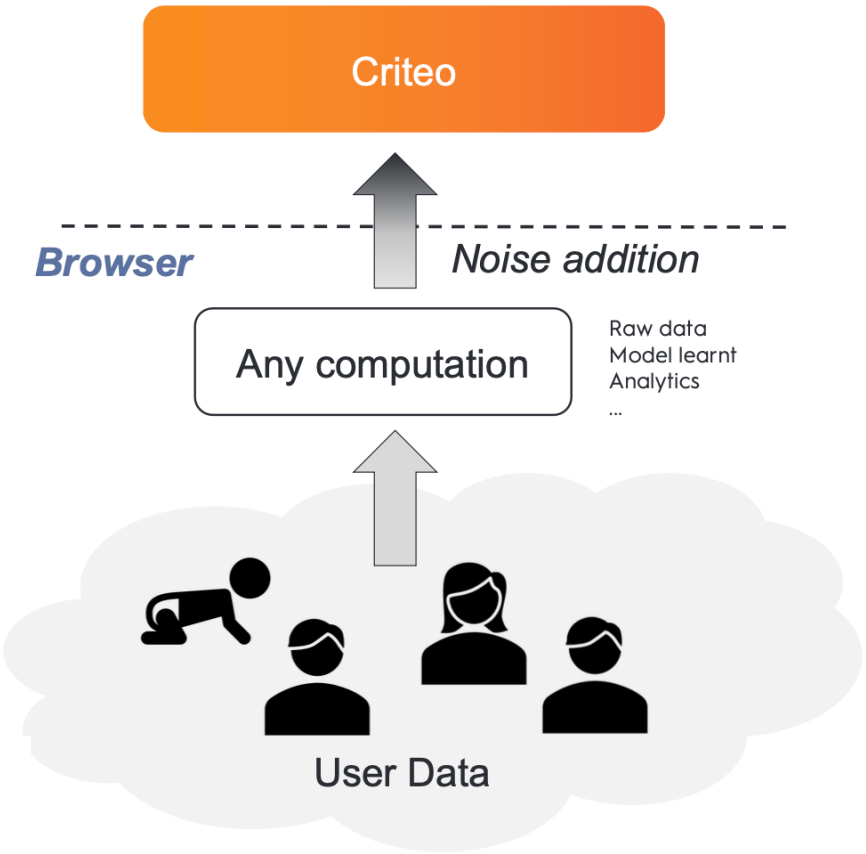

Numerical information
Additive Noise

$$x^{obs} = x + noise$$

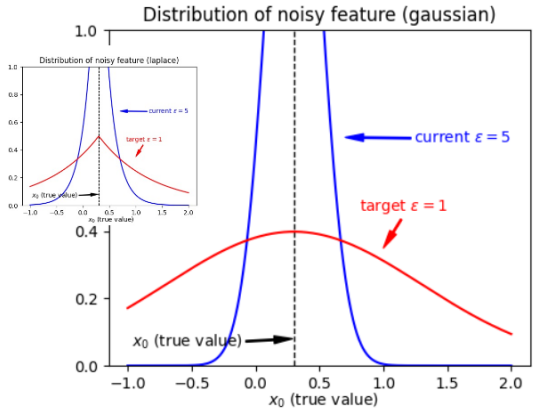

Categorical information
Randomized Response

$$x^{obs} = \begin{cases} x & with\ proba\ \alpha(\varepsilon) \\ random & with\ proba\ 1-\alpha(\varepsilon) \end{cases}$$

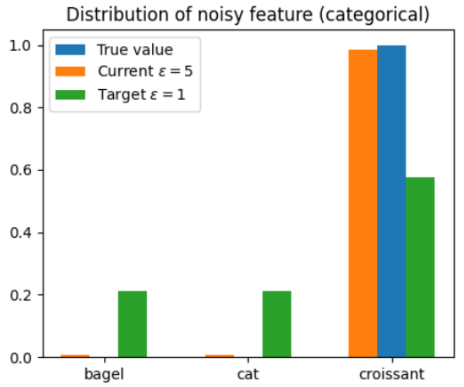

Key Messages

- *Any* information depending on *user/cohorts* data is noised
- In theory, the noise level for *each bit of information* retrieved depends on the *total quantity of information* retrieved

Local DP Learning Paradigm

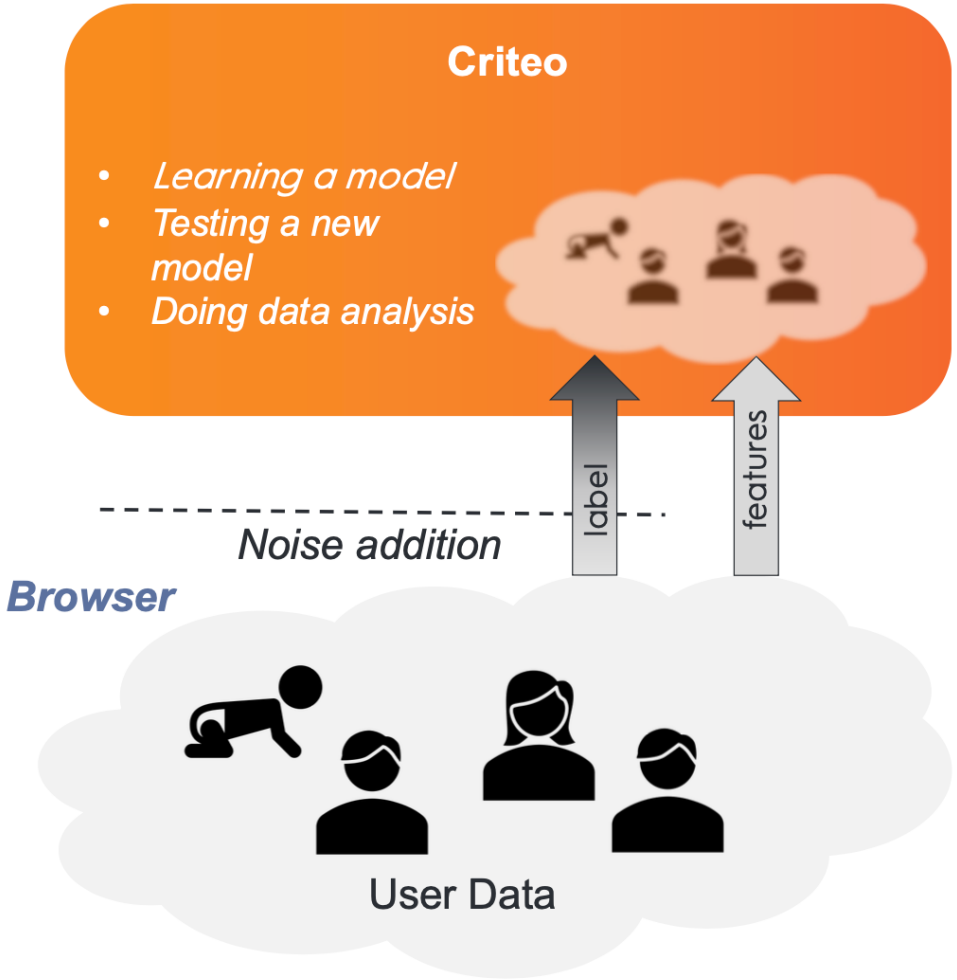

- **Pros:**
  - agnostic to the reporting/learning downstream task
  - noise is added once
  - easy to use within current stack
- **Cons:**
  - learning is biased
  - noise amount could be detrimental to performance
  - few contributions on local DP frameworks to operate AI systems

Local DP Learning Paradigm

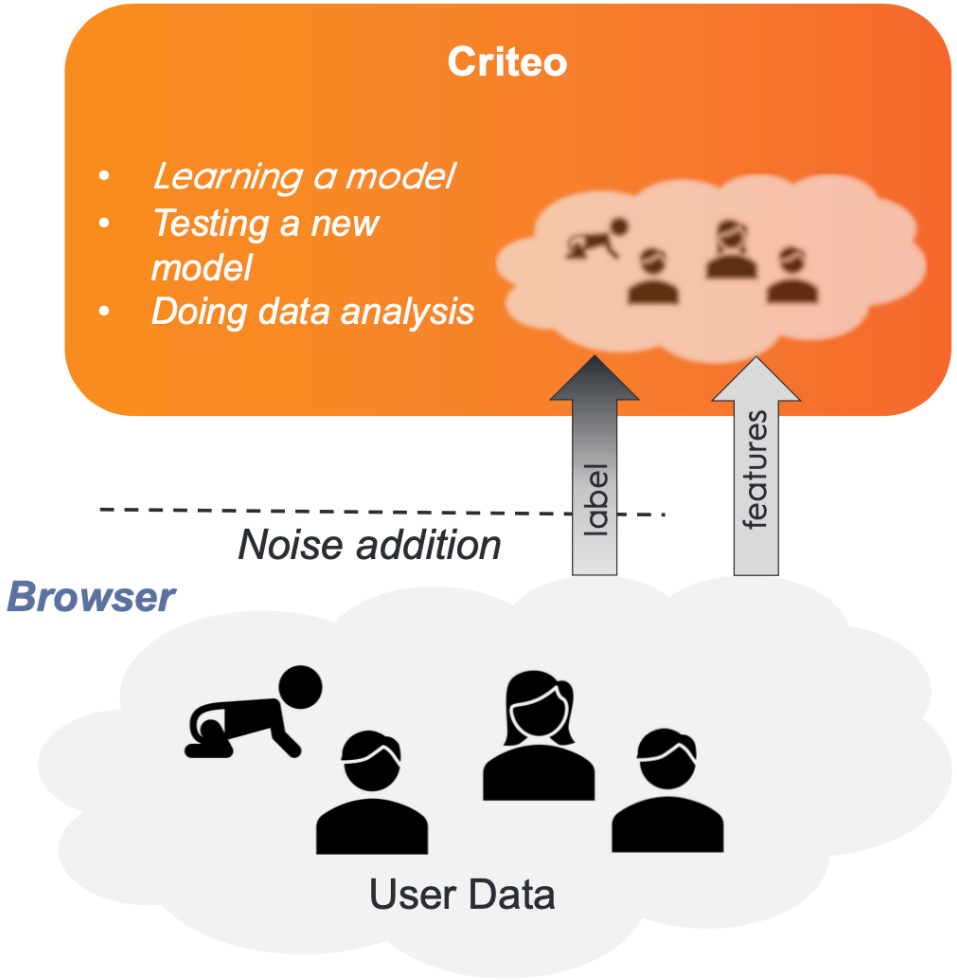

Position Paper: Open Research Challenges for Private Advertising Systems under Local Differential Privacy

Matilde Tullii [*2], Solenne Gaucher[*2], Hugo Richard[*1], Eustache Diemert[1], Vianney Perchet[1, 2], Alain Rakotomamonjy[1], Clément Calauzènes[1], and Maxime Vono[1]

[1]Criteo AI Lab, France
[2]ENSAE, Crest, France

February 5, 2024

Global DP Learning Paradigm – Aggregation API

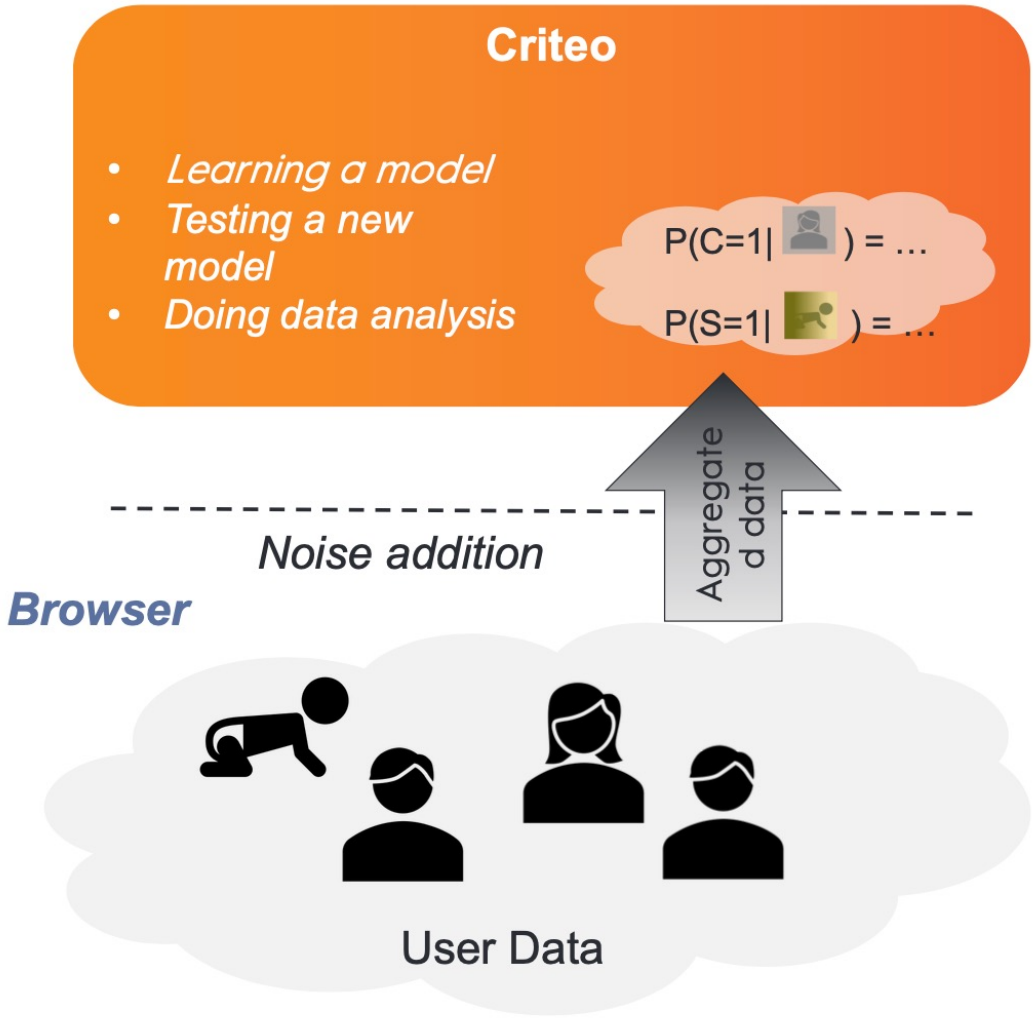

- **Pros:**
  - aggregate statistics are partially re-usable accross tasks
  - less noise than local DP
- **Cons:**
  - learning approaches could be envisioned leveraging GLMs structure
  - global full DP performance is not satisfying
  - local label DP is promising (WALR)

Global DP Learning Paradigm – Learning Trusted Server

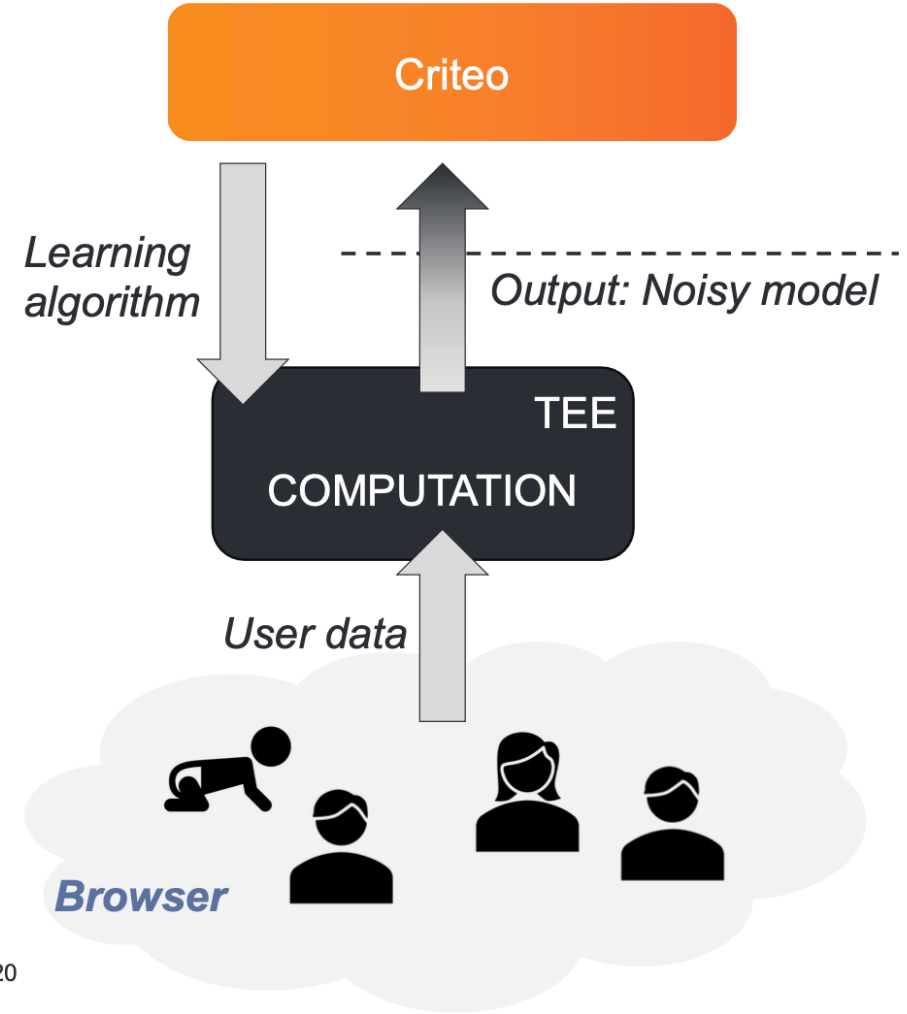

- **Pros:**
  - agnostic to the ML model
  - less noise than local DP
- **Cons:**
  - the more models/data analytics on the same data batch, the higher the noise
  - model monitoring/debugging?
  - distributed learning?

Impact of DP noise on offline performance

- **Rationale:**
  - general private learning problem is more complex (infra cost, availability of features at training/inference time,…)
  - we want to isolate the DP effect to have first insights for decision making
- **Base ML model:** ridge logistic regression
- **Privacy unit:** display (for the sake of simplicity)
- **Approches benchmarked:**
  - base model without DP
  - learning on event-level data (with full and label DP)
  - learning via gradients queried from an Aggregation API (with full and label DP)
  - learning inside a trusted server with global DP (DP-SGD)

Empirical Results - Dataset

Criteo Attribution Modeling for Bidding Dataset

https://ailab.criteo.com/criteo-attribution-modeling-bidding-dataset/

- Sample of 30 days of Criteo live traffic data.
- Each example corresponds to a click and contains:
  - **Features:** campaign ID, 9 contextual features, and the cost paid for the display.
  - **Label:** a 0/1 field indicating whether there was a conversion in the 30 days after the click and that is last-touch attributed to this click.
  - **User ID:** can be used to evaluate User x Time privacy unit.
- Number of rows is 5,947,563. Conversion rate (under last-touch attribution) is 6.74%.

Empirical Results – Performance (epsilon = 1)

Private Learning Approach	DP paradigm	DP type	Relative uplift (in LLH)
Baseline without DP	N/A	N/A	26%
Learning on event—level data	**Local DP**	**Full**	**< 0%**
		Label	**< 0%**
Learning via gradients queried from Aggregation API	Global DP	Full	21%
		Label	23%
Learning on aggregated data from Aggregation API	Global DP	Label	22%
DP-SGD	Global DP	Full	24%

Empirical Results – Performance (epsilon = 5)

Private Learning Approach	DP paradigm	DP type	Relative uplift (in LLH)
Baseline without DP	N/A	N/A	26%
Learning on event—level data	Local DP	Full	16%
		Label	25%
Learning via gradients queried from Aggregation API	Global DP	Full	23%
		Label	26%
Learning on aggregated data from Aggregation API	Global DP	Label	24%
DP-SGD	Global DP	Full	26%

Empirical Results – Performance (epsilon = 5)

Private Learning Approach	DP paradigm	DP type	Relative uplift (in LLH)
Baseline without DP	N/A	N/A	26%
Learning on event—level data	**Local DP**	**Full**	**16%**
		Label	**25%**
Learning via gradients queried from Aggregation API	Global DP	Full	23%
		Label	26%
Learning on aggregated data from Aggregation API	Global DP	Label	24%
DP-SGD	**Global DP**	**Full**	**26%**

Next Steps

- Present to PATCG
    - A benchmark of SOTA on learning from noisy event-level reports, esp. featuring user-level budgeting
    - Insights regarding learning inside a trusted server (focus on TEE-empowered trusted server)
- Happy to collaborate on how to extend WALR (or related ideas) beyond GLMs