Learning To Generate Image Embeddings With User-Level Differential Privacy

Zheng Xu*, Maxwell Collins*, Yuxiao Wang, Liviu Panait, Sewoong Oh, Sean Augenstein, Ting Liu, Florian Schroff, H. Brendan McMahan

Google Research

Poster 368, Tue PM

JUNE 18-22, 2023
CVPR
VANCOUVER, CANADA

How can we protect privacy when training an image embedding model from user data?

- **User-level DP:** mathematical guarantees that a model won't memorize user data; successfully applied in small on-device language models in production.
- **DP-FedEmb**: a new algorithm to train large image-to-embedding feature extractors specifically designed for scalability to achieve strong privacy-utility trade-offs
  - Virtual clients, partial aggregation, private local fine-tuning, and public pretraining
- Superior utility under same privacy budget on benchmark datasets DigiFace, EMNIST, GLD and iNaturalist for faces, landmarks and natural species.
- It is possible to achieve strong user-level DP guarantees of single-digit epsilon while controlling the utility drop within 5%, when millions of users can participate in training .

Key algorithm design choices

- construction of virtual clients
- selection of what information is shared among users

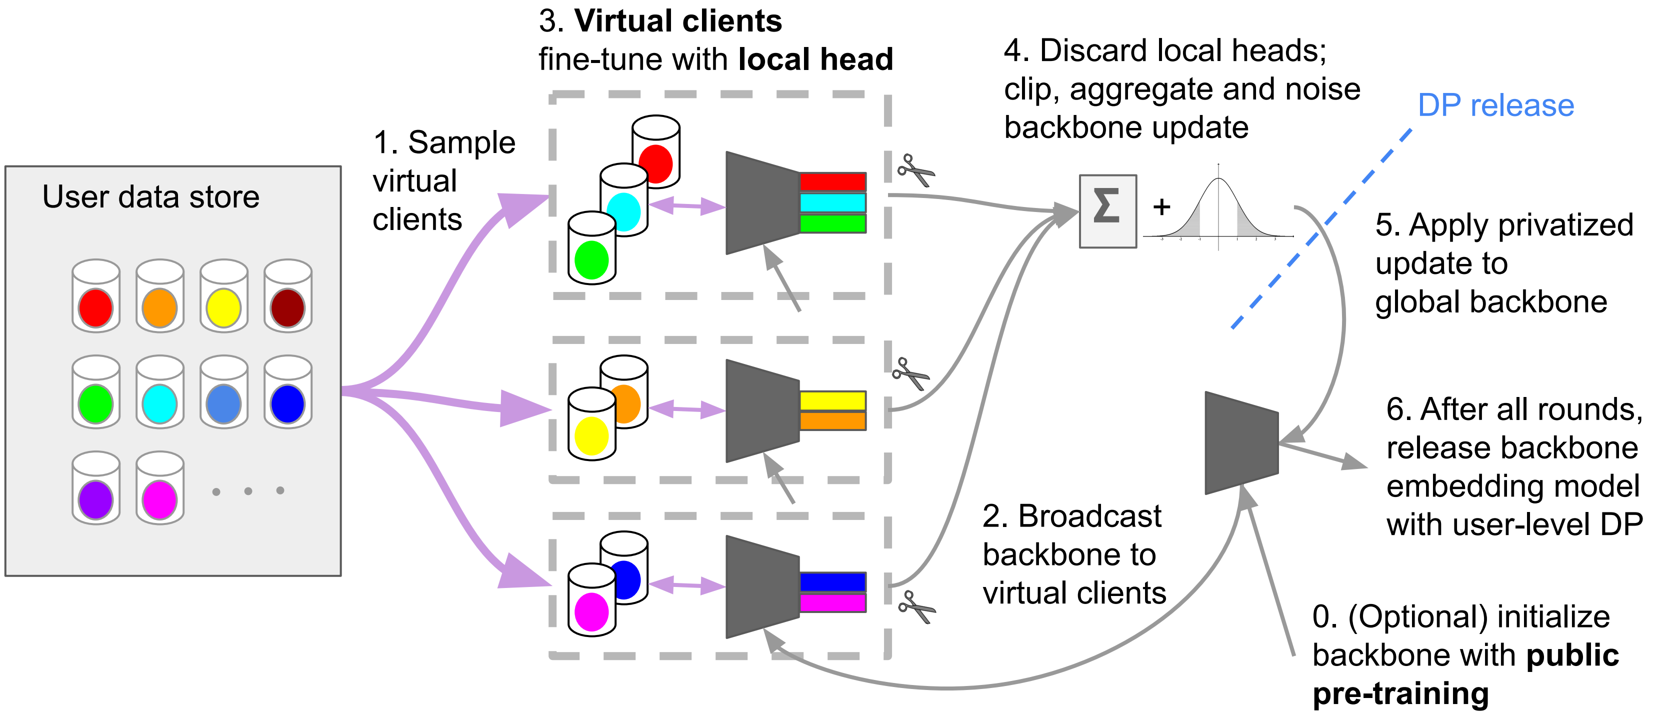

User Data (in Datacenter)

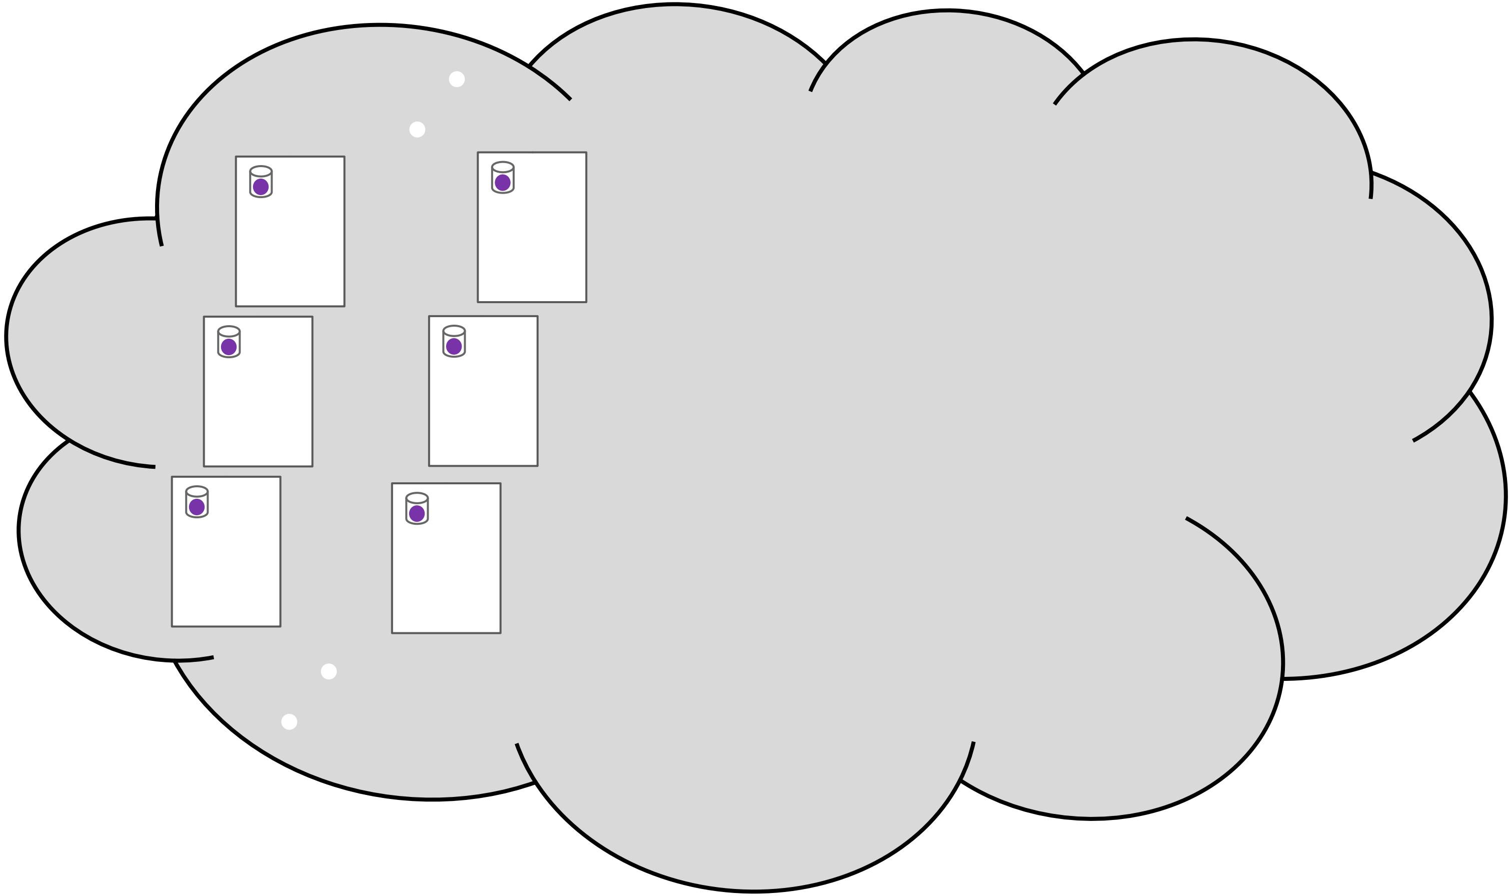

Machine Learning from User Data

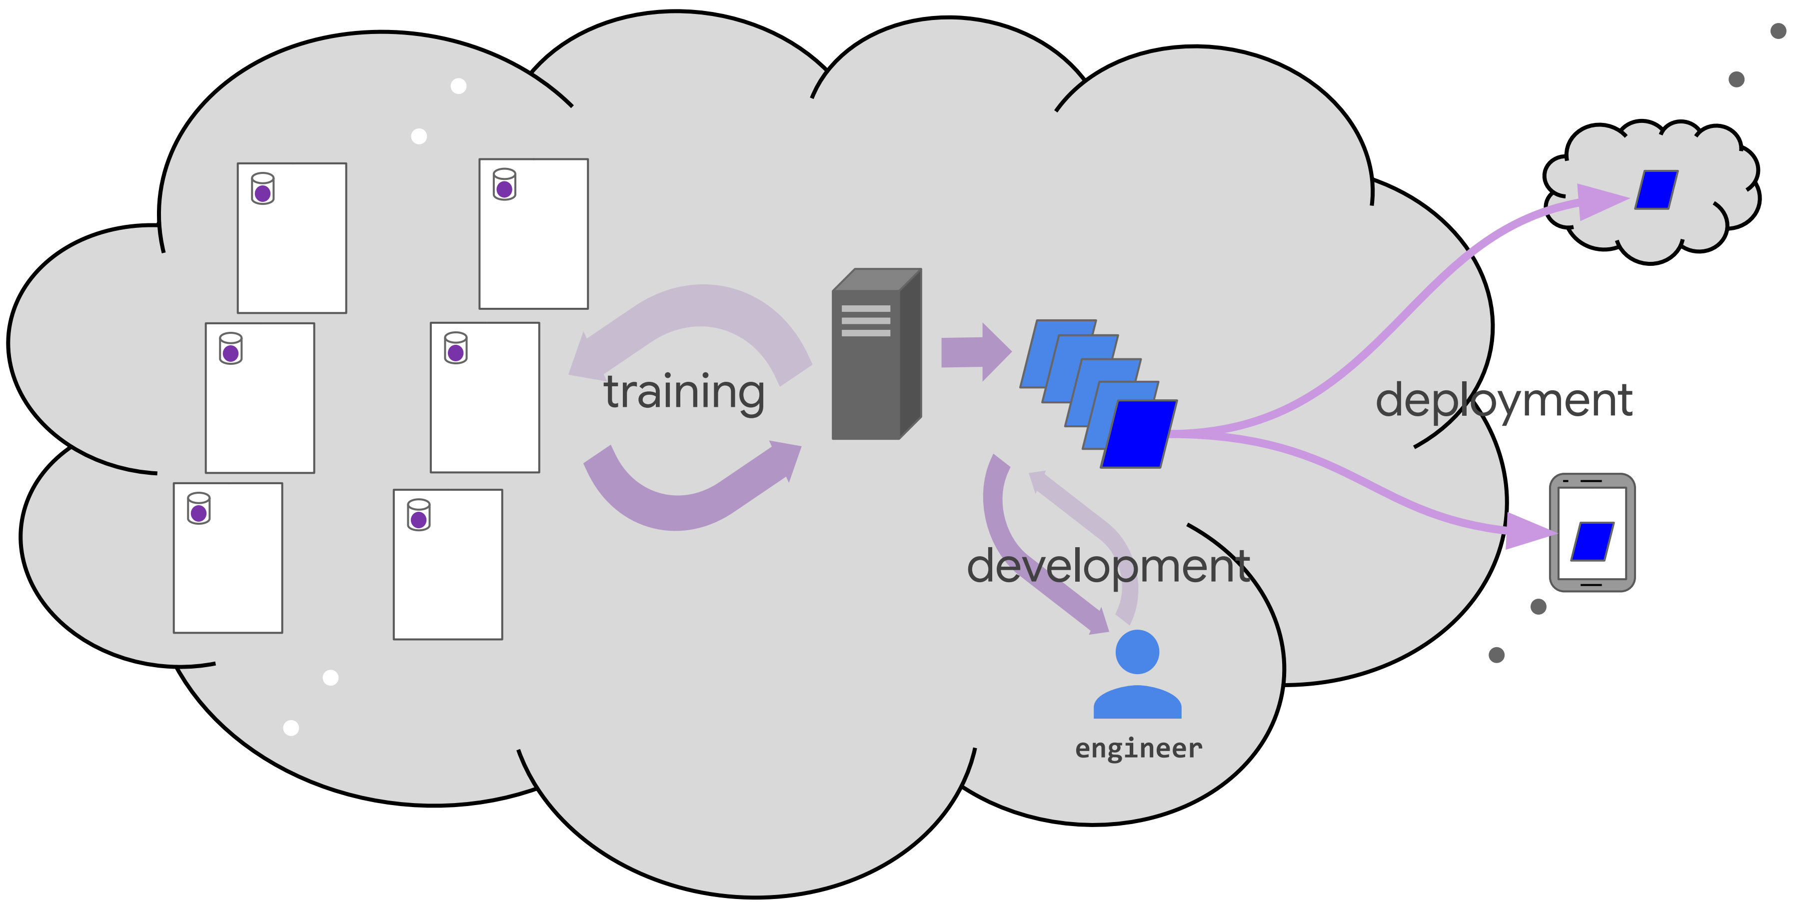

User-level Differential Privacy

Data anonymization: model won't memorize individual user's data

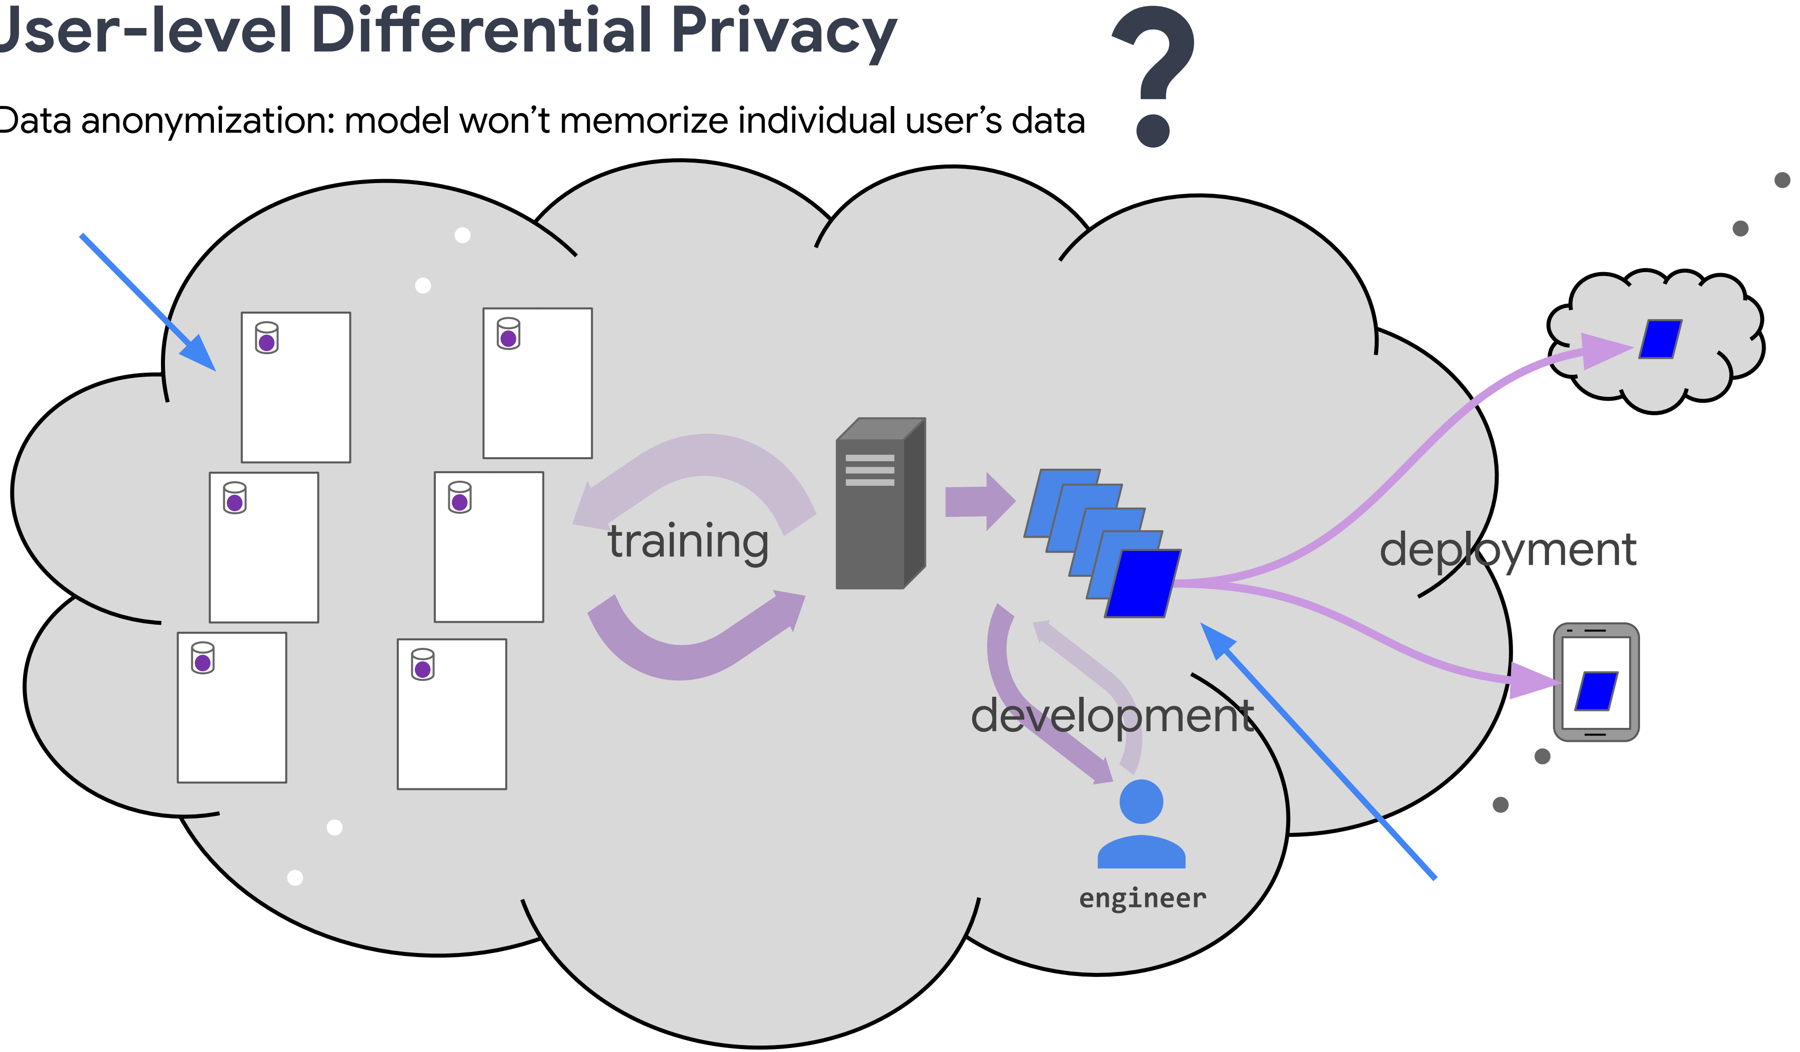

User-level Differential Privacy by “Federated” Algorithm

Data anonymization: model won’t memorize individual user’s data

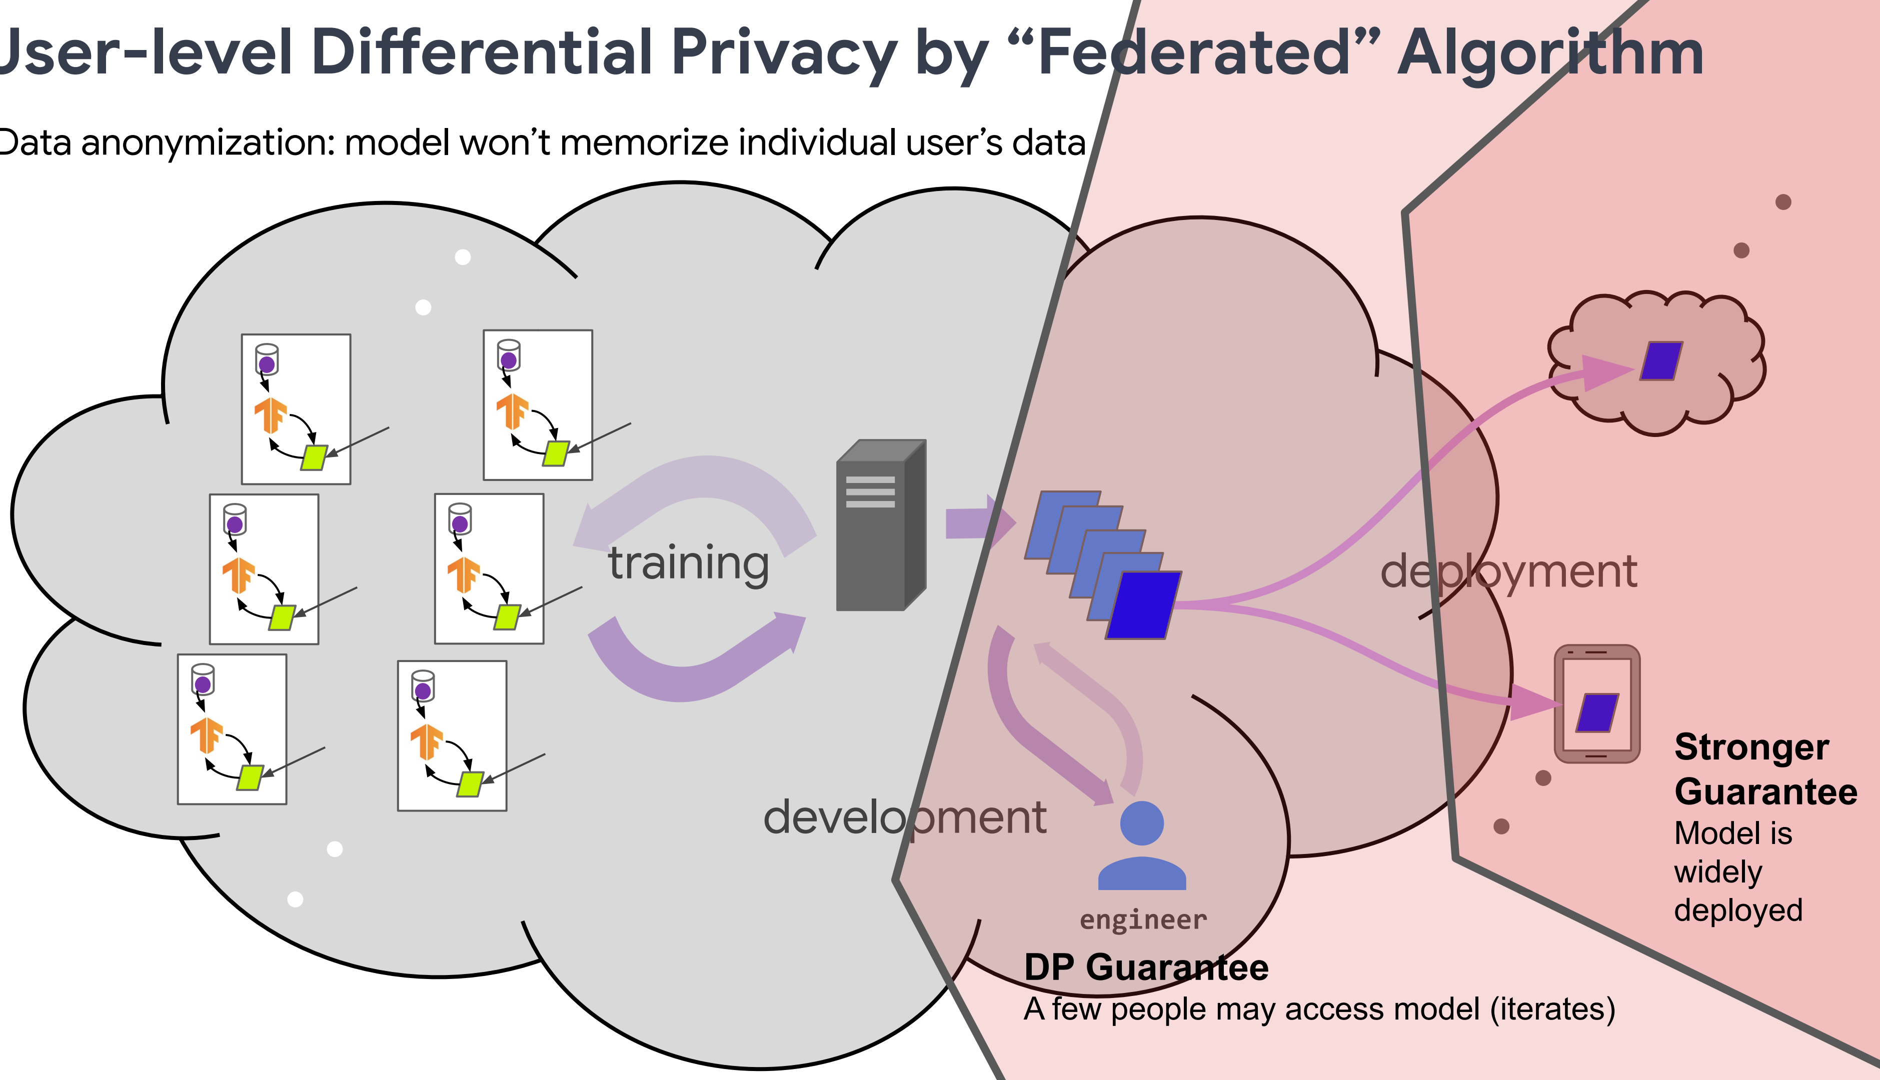

User-level Differential Privacy by “Federated” Algorithm

“Natural” fit

- Data granularity by users
- Infrequent aggregation and model release

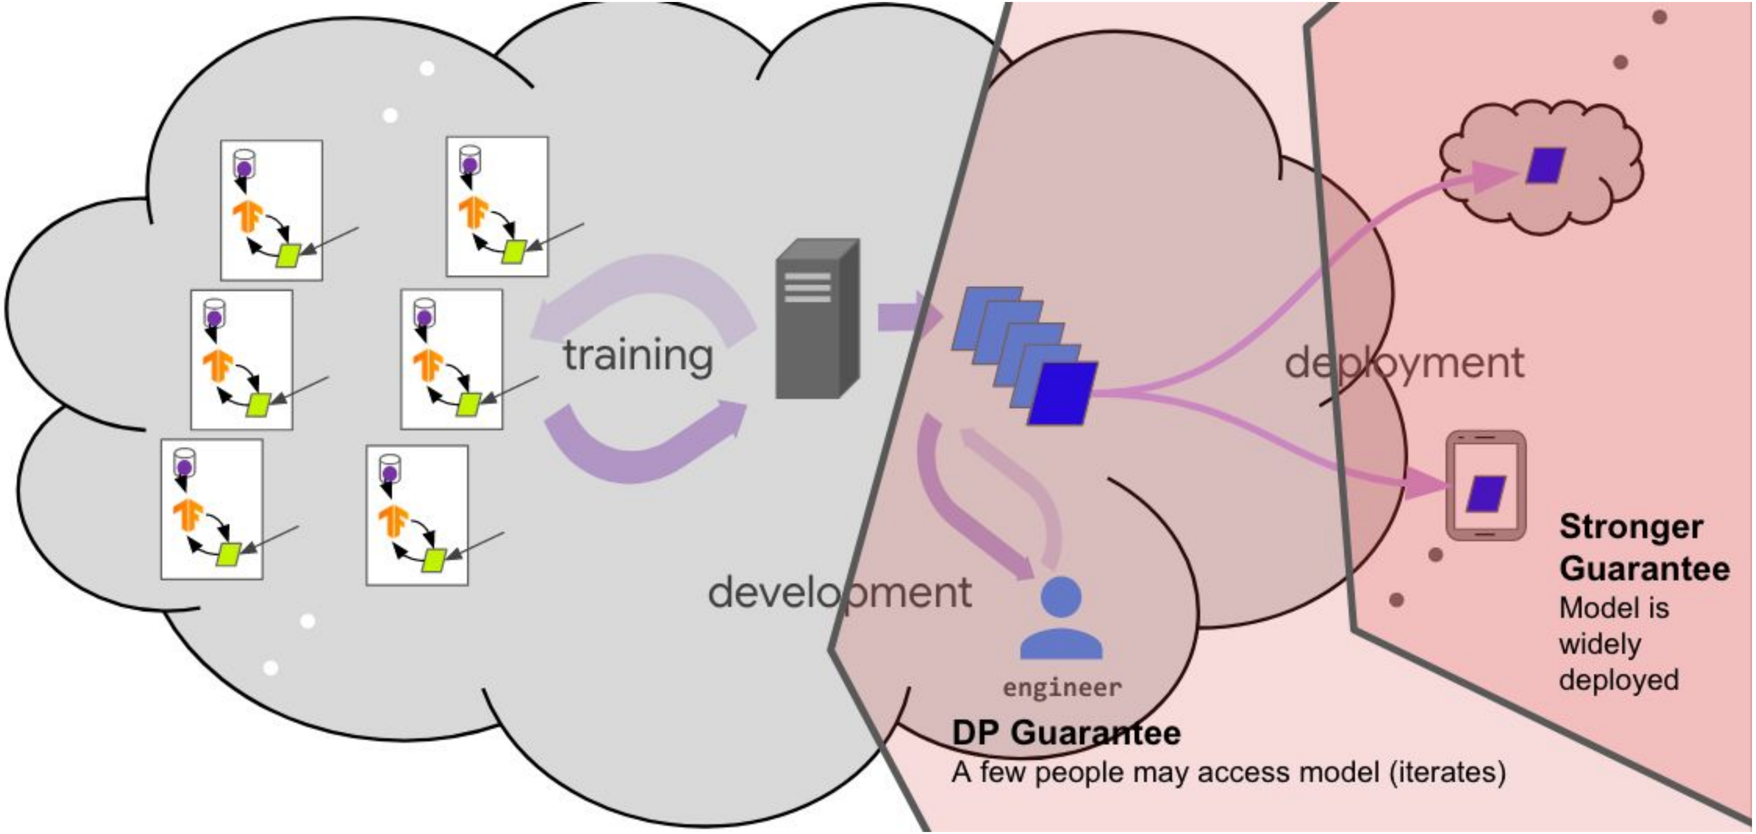

Differentially Private Federated Averaging (DP-FedAvg)

User-siloed data

- Conceptual broadcast and aggregation

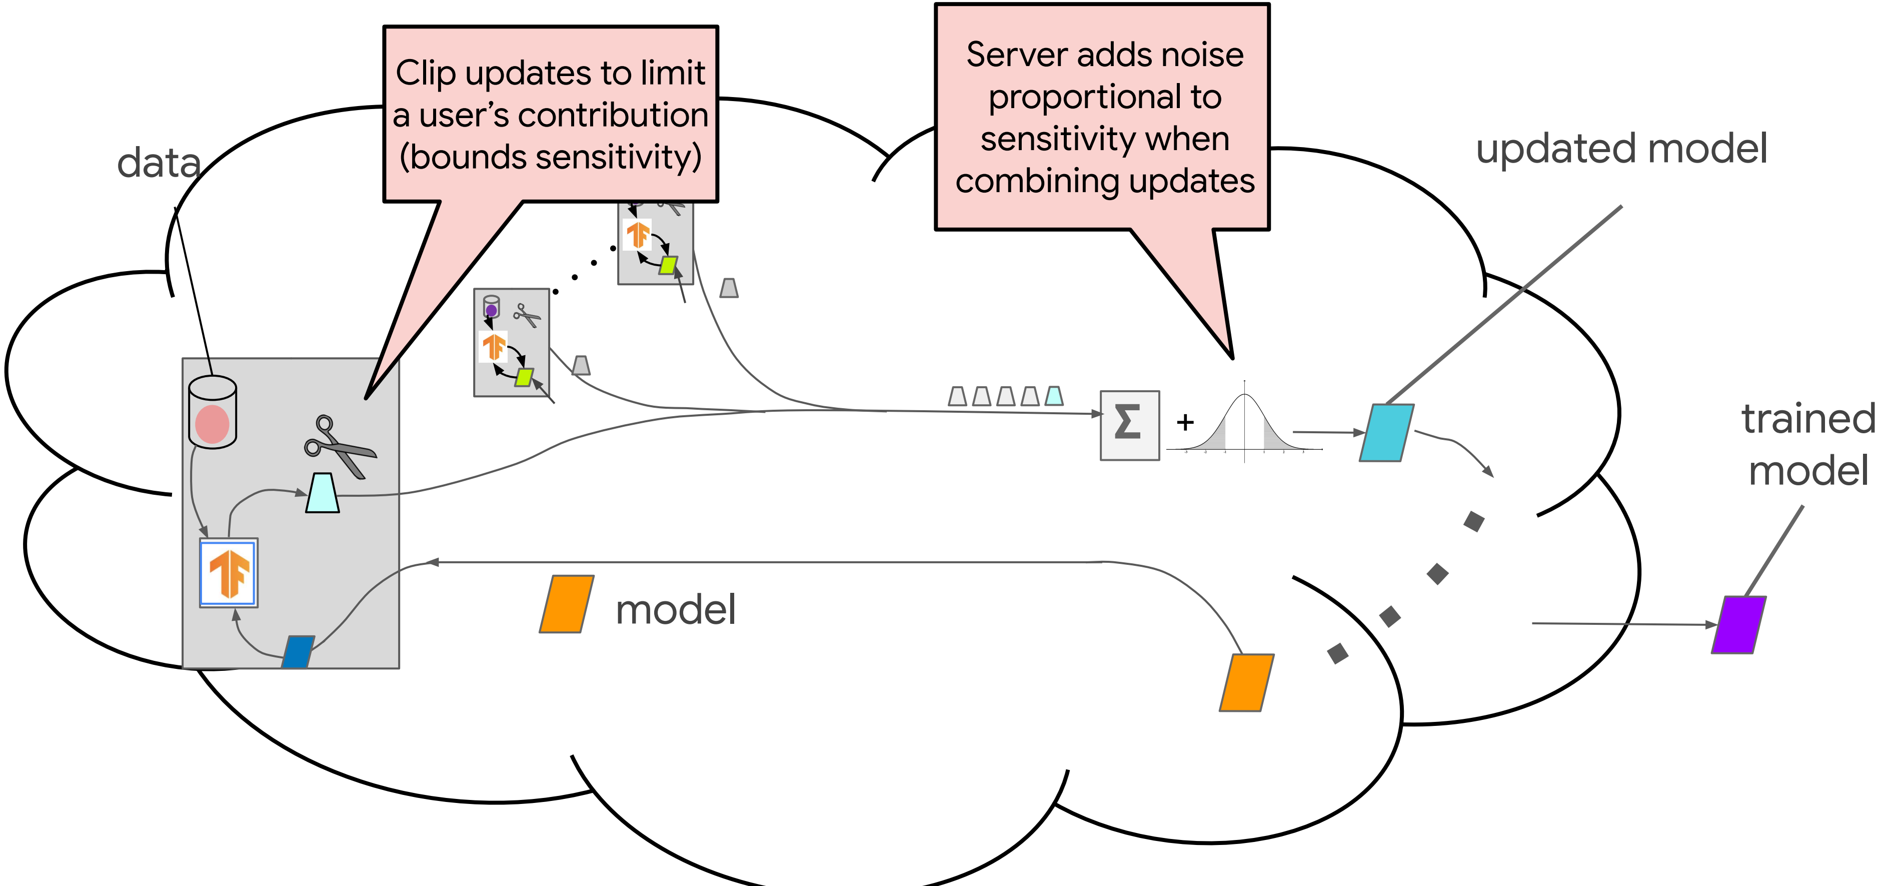

Image Embedding Models

Image Embedding Models

Challenges with DP-FedAvg

- Heterogeneity: contrastive samples in user-partitioned data
- High-dimensionality that cannot be solved by sampled softmax

Backbone parameters

- ResNet-50 ~23.77M
- MobileNetV2 ~3.49M

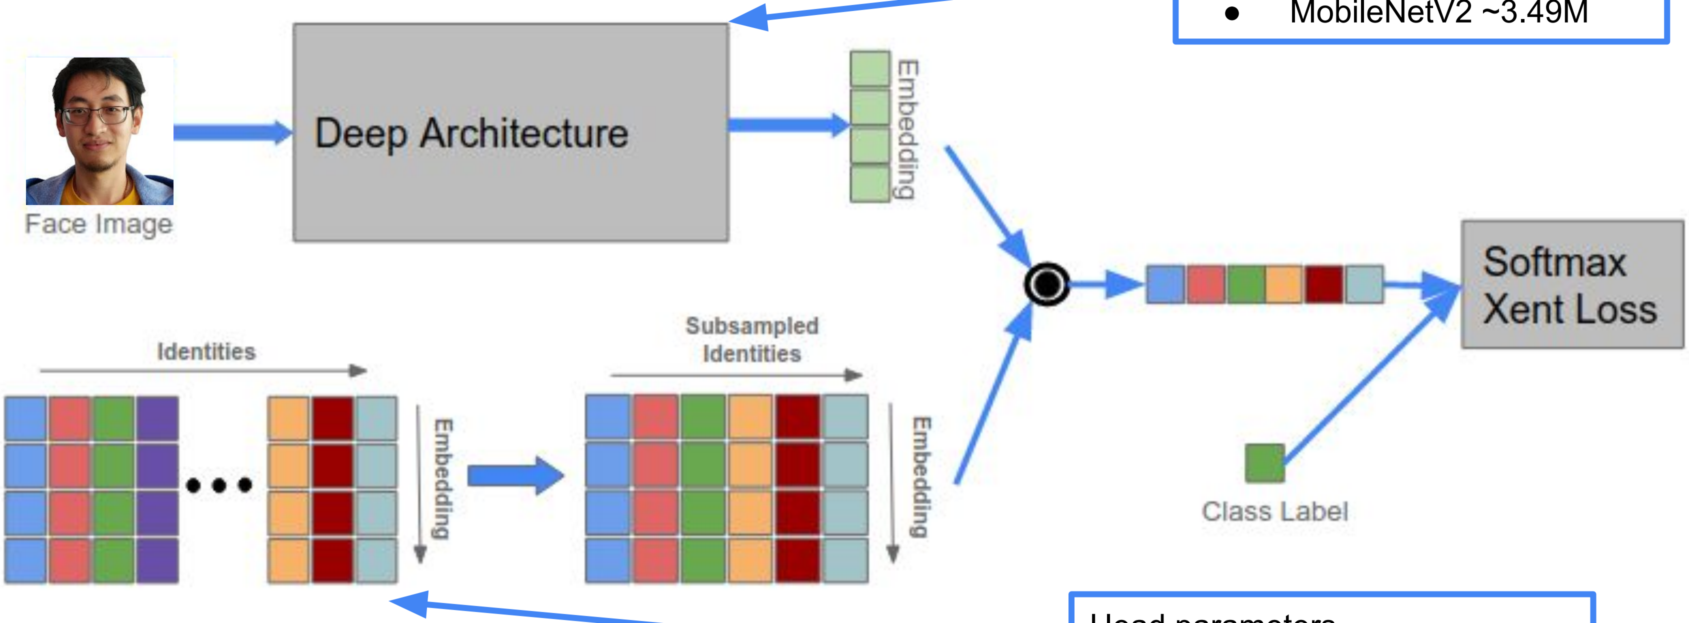

Head parameters

- 10K identities ~1.28M
- 10M identities ~1280M

Embedding Models with User-level DP

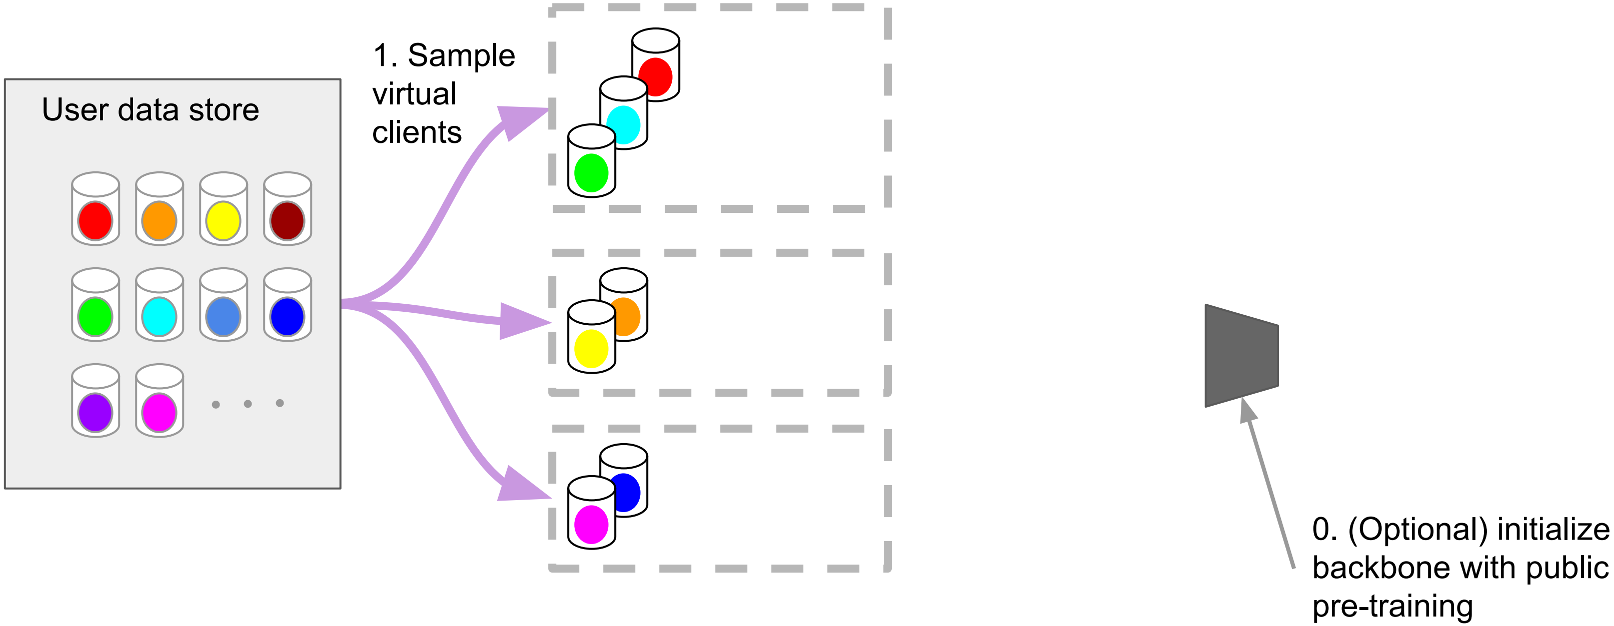

Embedding Models with User-level DP

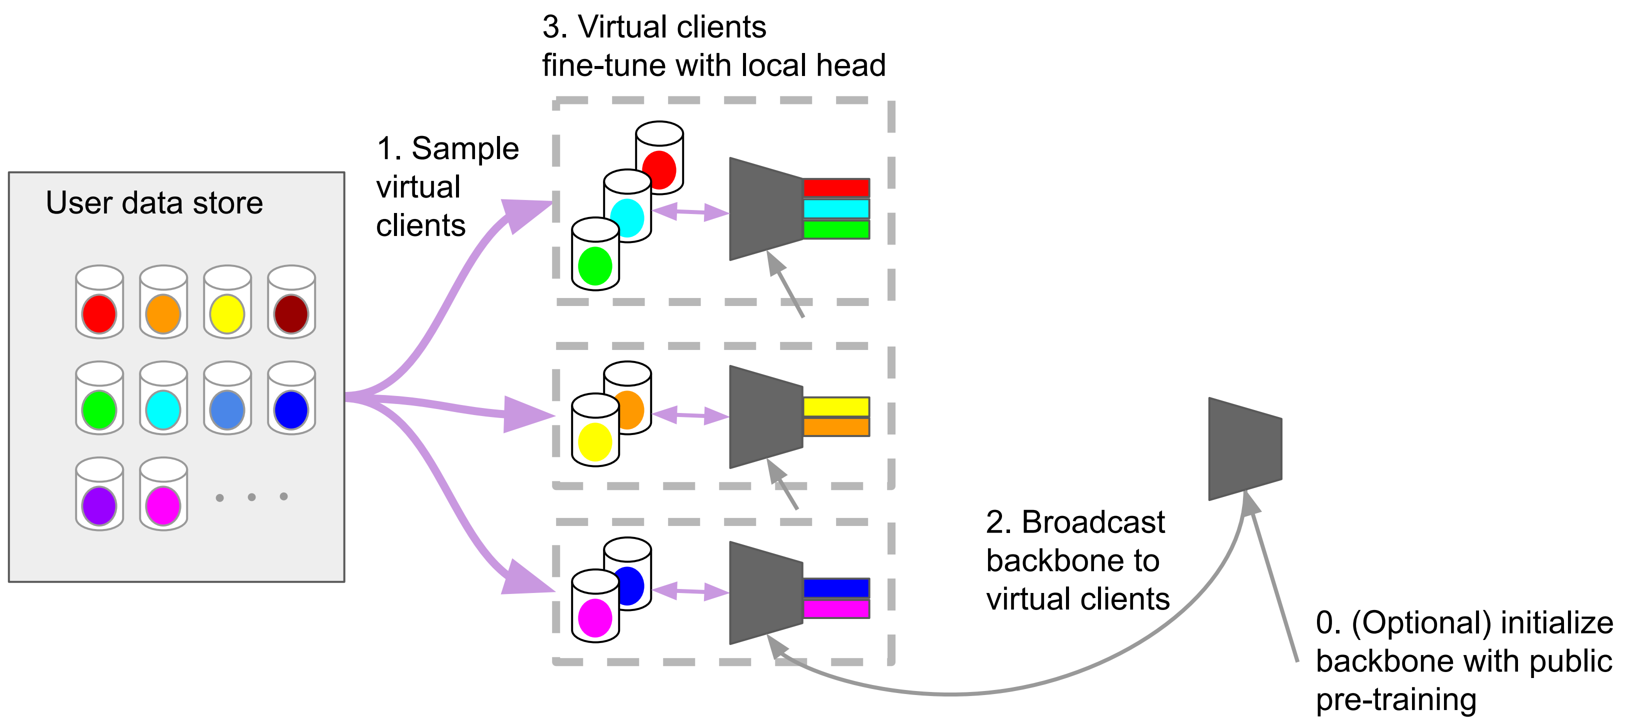

Embedding Models with User-level DP

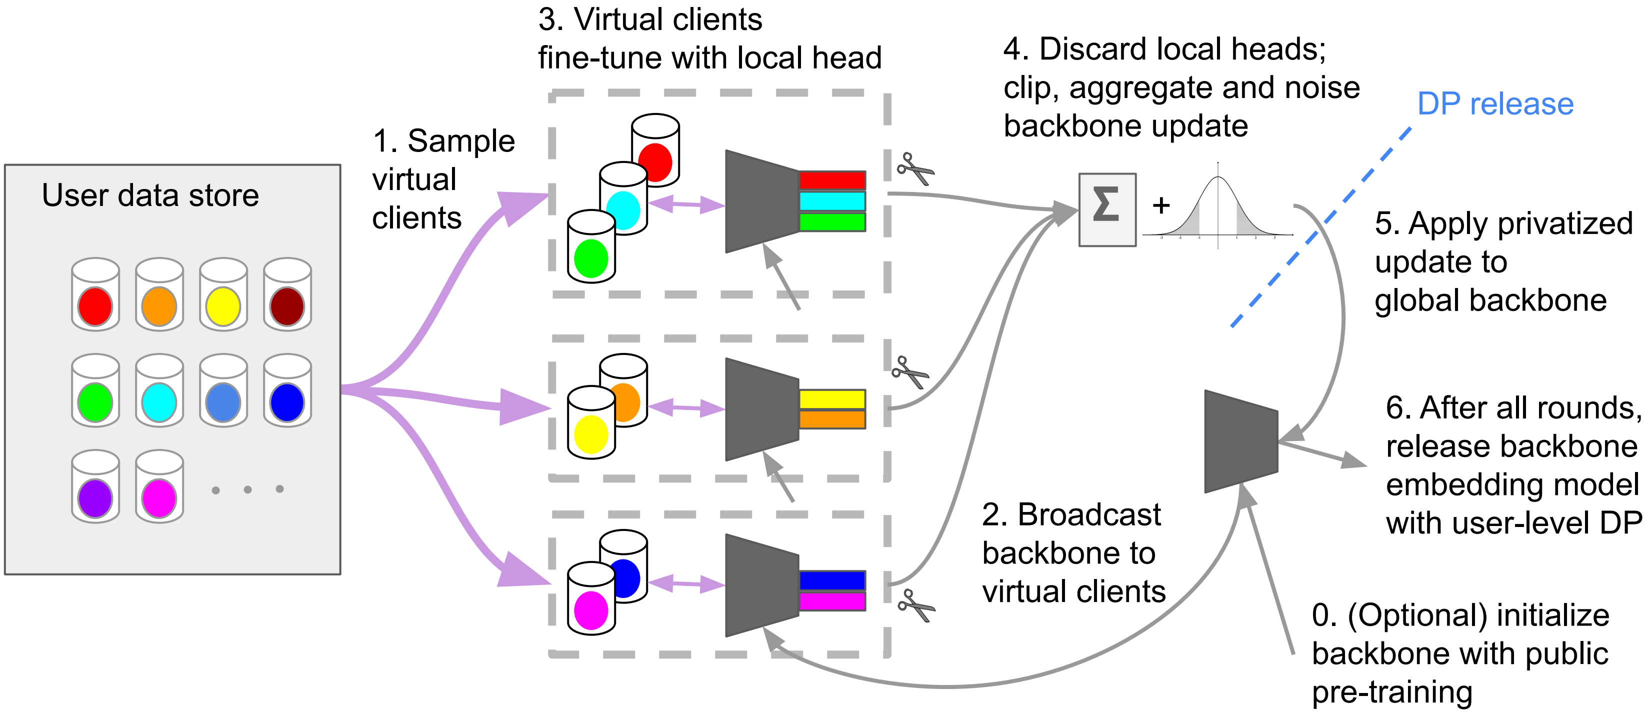

Embedding Models with User-level DP

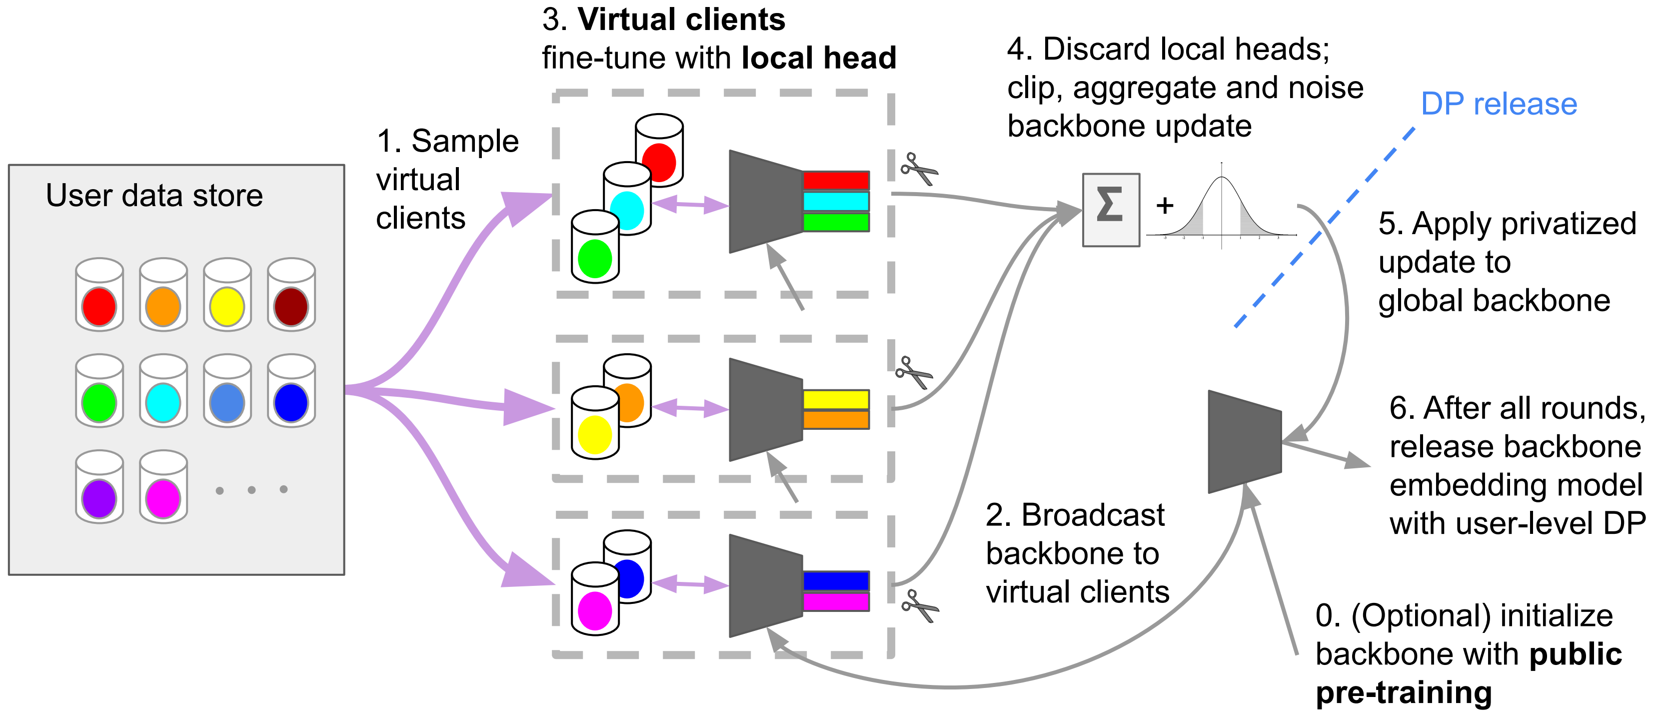

Results on DigiFace

Privacy guarantees with less than 5% utility drop

Algorithm	Hyperparameters		Privacy (10M users)		Recall@FAR=1e−3	
	Noise	SerLR	RDP-ϵ	zCDP	Validation	Test
Centralized	0	0.05	∞	∞	75.55 ± 0.05	75.53 ± 0.12
DP-FedAvg	0.015×64	0.5	5.62	-	72.57 ± 0.12	72.37 ± 0.09
DP-FedEmb	0.02×64	0.2	3.90	-	72.63 ± 0.05	72.37 ± 0.09
DP-FTRL-FedEmb	0.26×64	0.2	9.67	1.28	72.2 ± 0.29	71.87 ± 0.26

- Centralized baseline is a suboptimal repro removing tricks like data augmentation that are not currently implemented in federated training yet
- Formal privacy guarantees are based on extrapolation
  - More users are available in a practical setting
  - For sufficiently large data, the utility accuracy will not drop if noise multiplier and clients per round proportionally increase; 32*8 GPUs can be used for 8 days
- Verified that the conclusions on DigiFace are very similar to conclusions generated from experiments on natural facial images

Takeaways

- Differential privacy guarantees are achievable in practice
  - Scale is the key: large amount of data and computation resources
  - Improving privacy-utility trade-off by public data, new algorithms, DP mechanism and accounting
- Privacy is not “free”
  - Computation and infra support
  - Common understanding of the techniques: verifiable, auditing
  - Engineering efforts / migration cost

Learning To Generate Image Embeddings With User-Level Differential Privacy

Zheng Xu*, Maxwell Collins*, Yuxiao Wang, Liviu Panait, Sewoong Oh, Sean Augenstein, Ting Liu, Florian Schroff, H. Brendan McMahan

Google Research

Poster 368, Tue PM

JUNE 18-22, 2023

CVPR

VANCOUVER, CANADA