CVPR 2020

Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem

Matthias De Lange[1] Xu Jia[2] Sarah Parisot[2,3]
Ales Leonardis[2] Gregory Slabaugh[2] Tinne Tuytelaars[1]

[1]KU Leuven [2]Huawei, Noah's Ark Lab [3]Mila

{firstname.lastname}@kuleuven.be {firstname.lastname}@huawei.com

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

Roadmap

- **Model Personalization: What and how?**
- Novel Benchmarks
- Adaptation on the server
- Adapting locally
- Conclusion

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Model Personalization

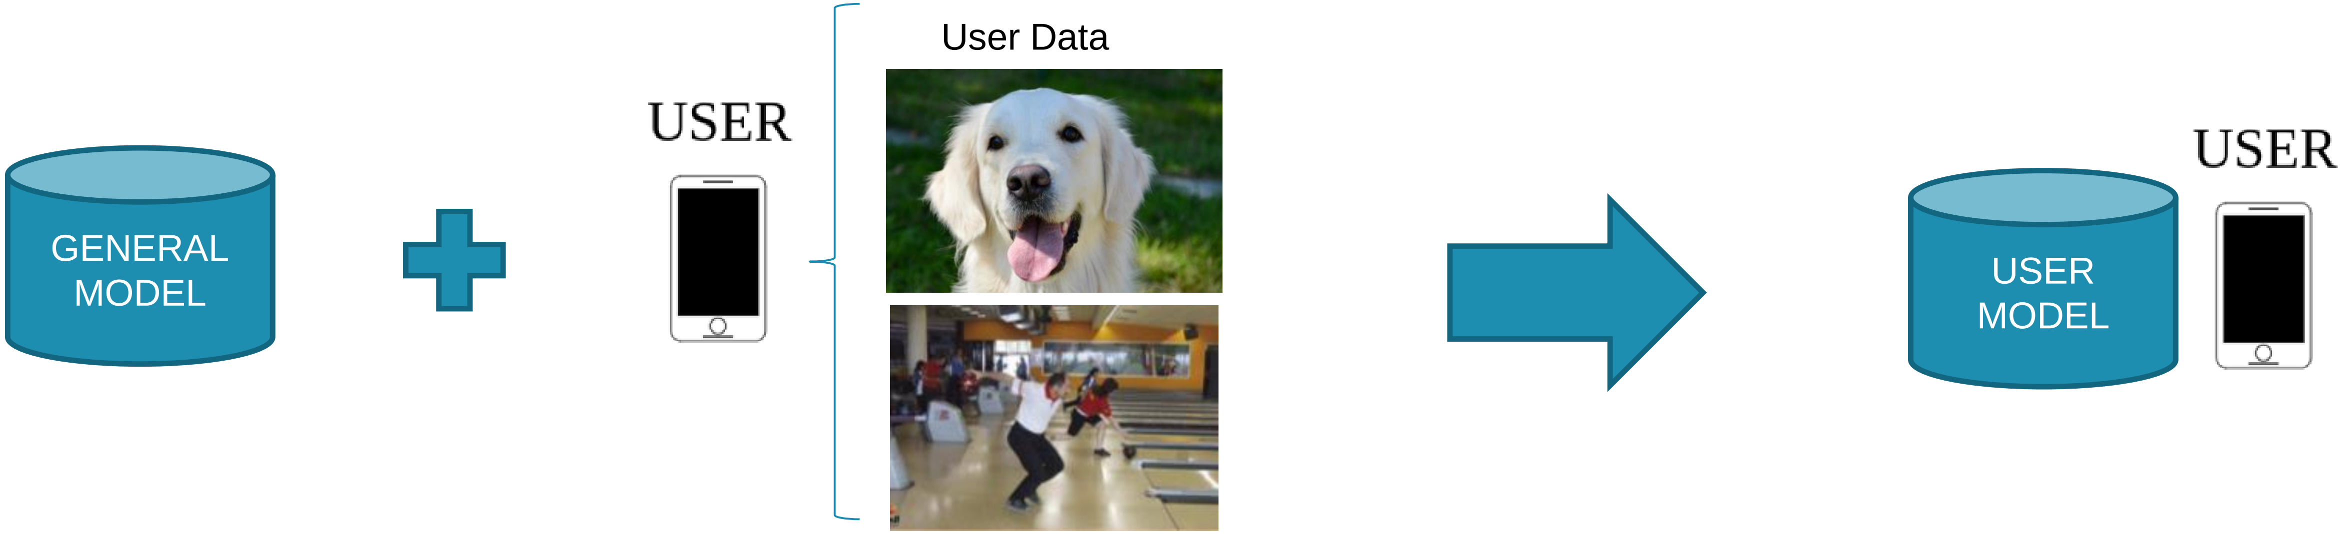

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

Model Personalization
2 Ways

- User-Adaptation **on the server**:
  + High Capacity

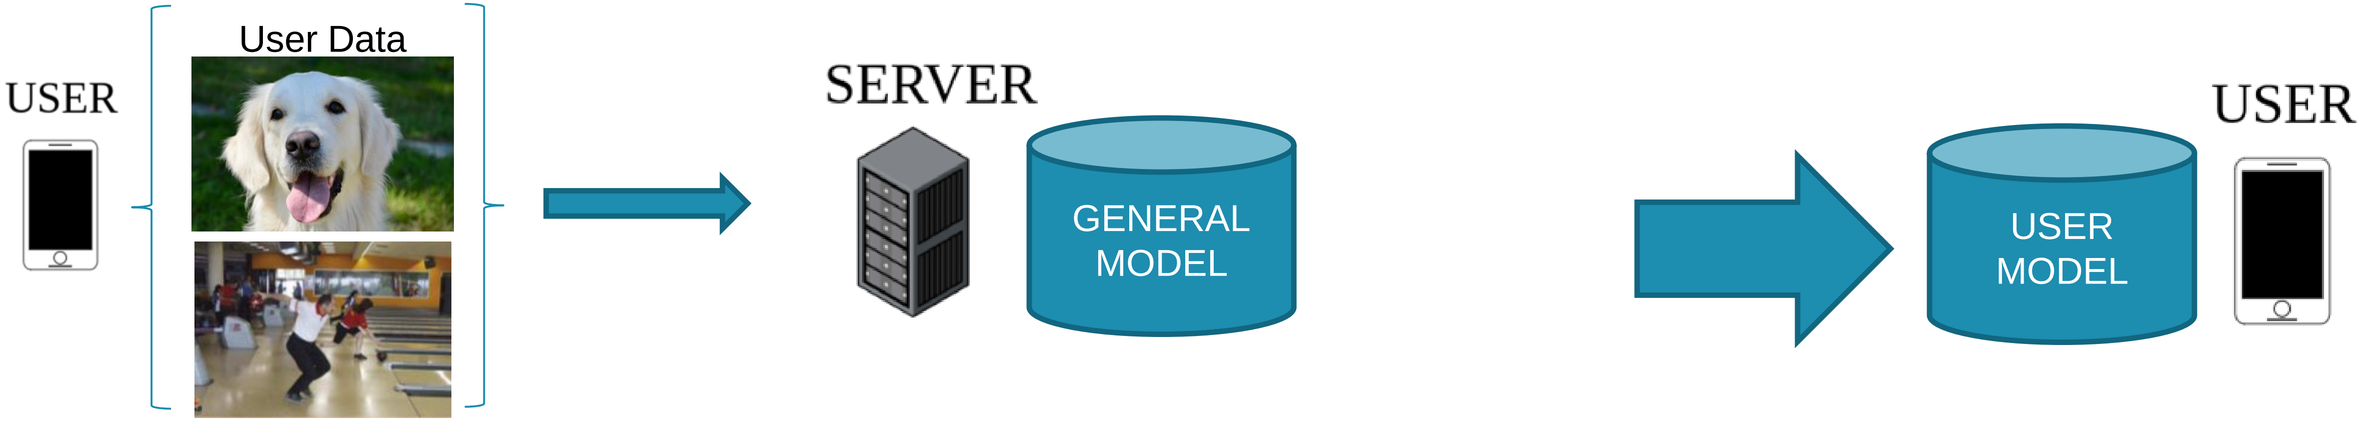

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Privacy

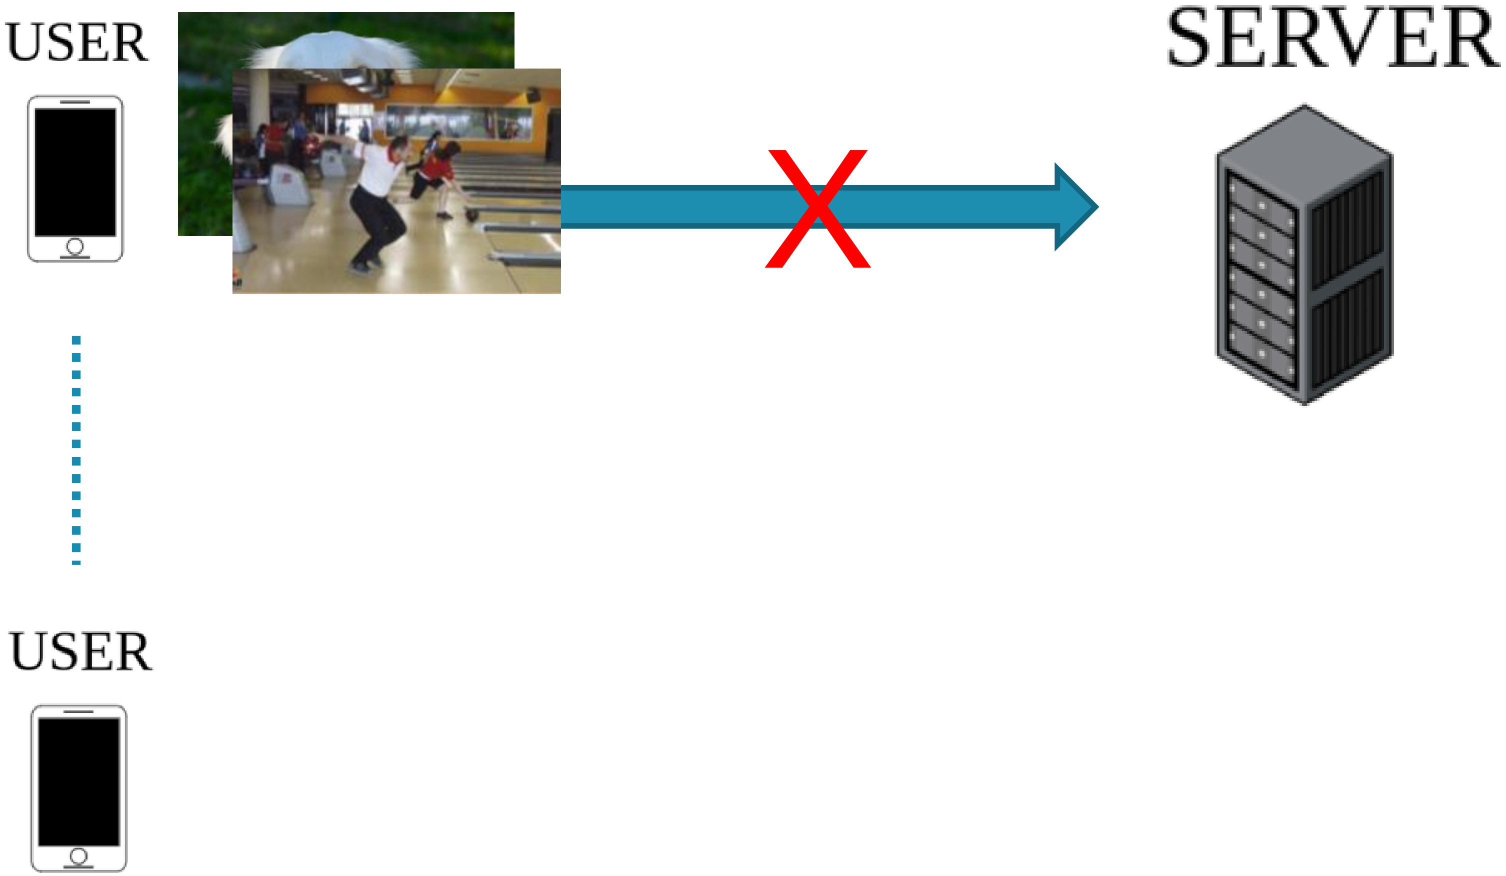

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Scalability

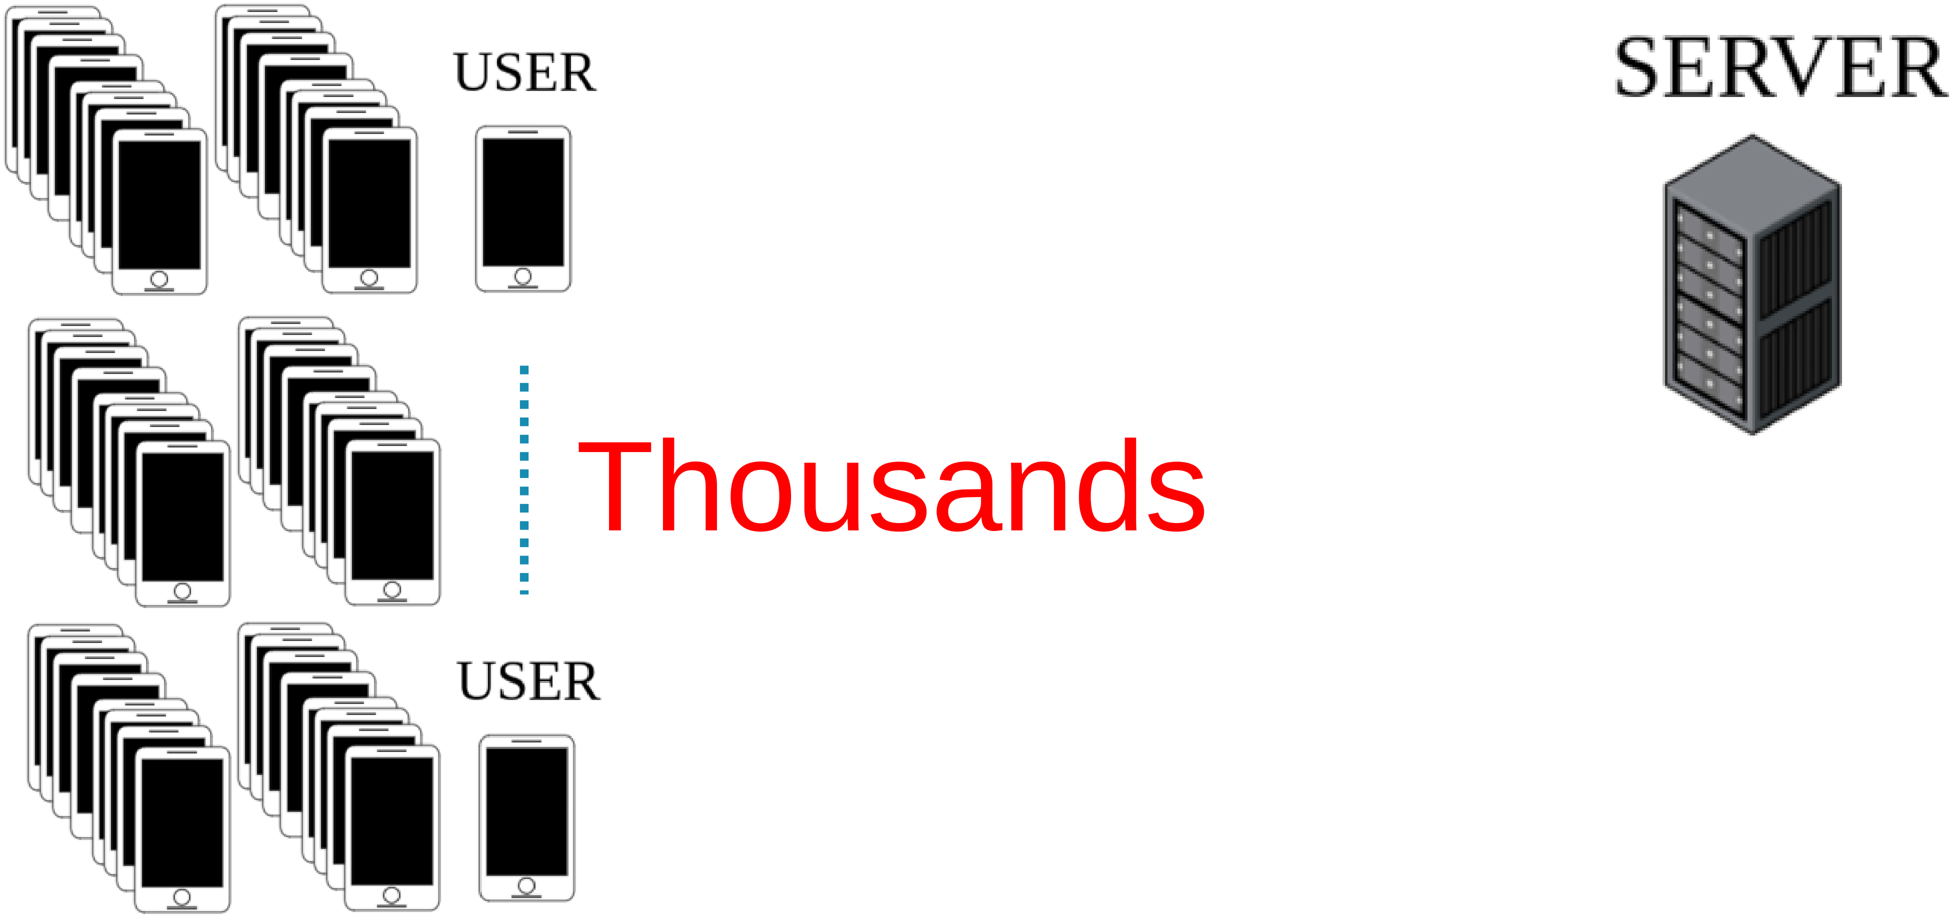

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Model Personalization
2 Ways

- User-Adaptation local **on user device**:
  - Low capacity
  + No privacy issues
  + No scalability issues

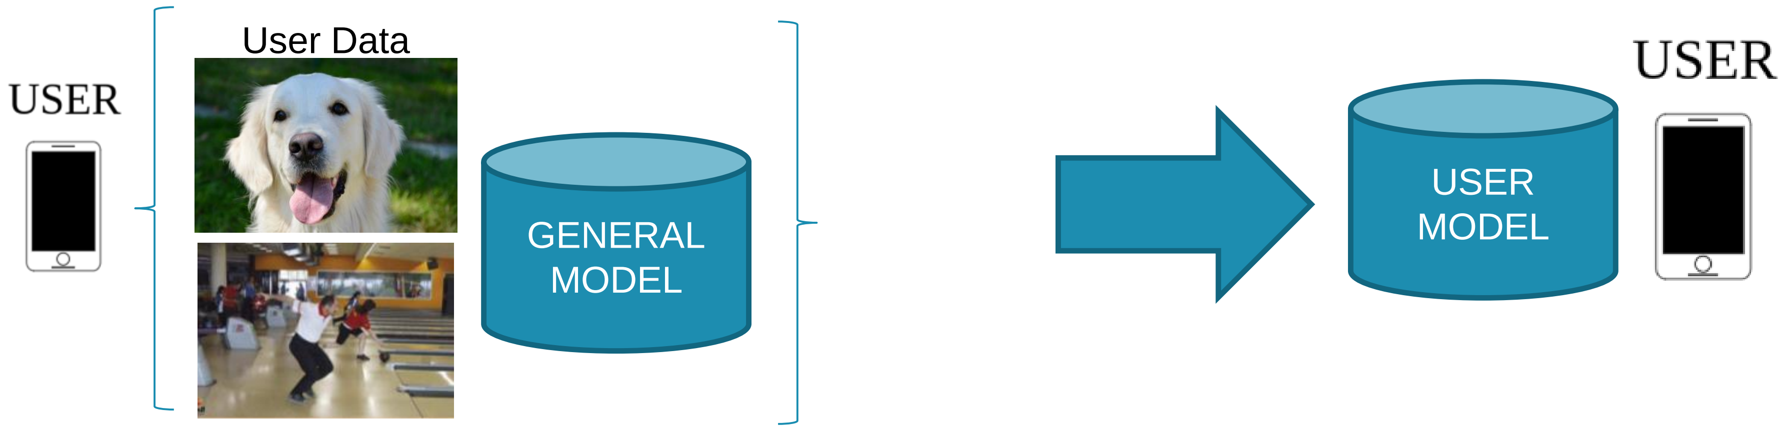

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Supervision

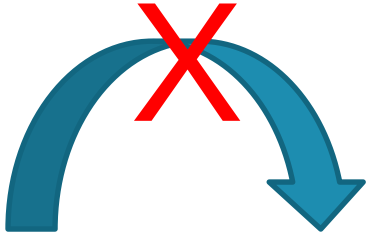

USER

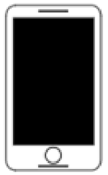

NO LABELS

USER

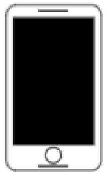

SERVER

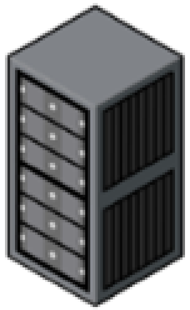

HUAWEI Mila

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

PSI KU LEUVEN

Model Personalization
2 Ways

On the server:

+ High Capacity
- Privacy
- Scalability
- Unlabeled user-data

Locally:

- Low Capacity
+ Privacy
+ Scalability
- Unlabeled user-data

Questions? matthias.delange@kuleuven.be

Model Personalization
2 Ways

On the server:

+ High Capacity
- Privacy
- Scalability
- Unlabeled user-data

Locally:

- Low Capacity
+ Privacy
+ Scalability
- Unlabeled user-data

One Framework

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

Model Personalization
2 Ways

On the server:

+ High Capacity
- Privacy
- Scalability
- Unlabeled user-data

Locally:

- Low Capacity
+ Privacy
+ Scalability
- Unlabeled user-data

One Framework

2x adaptation

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

Dual User-Adaptation framework (DUA)

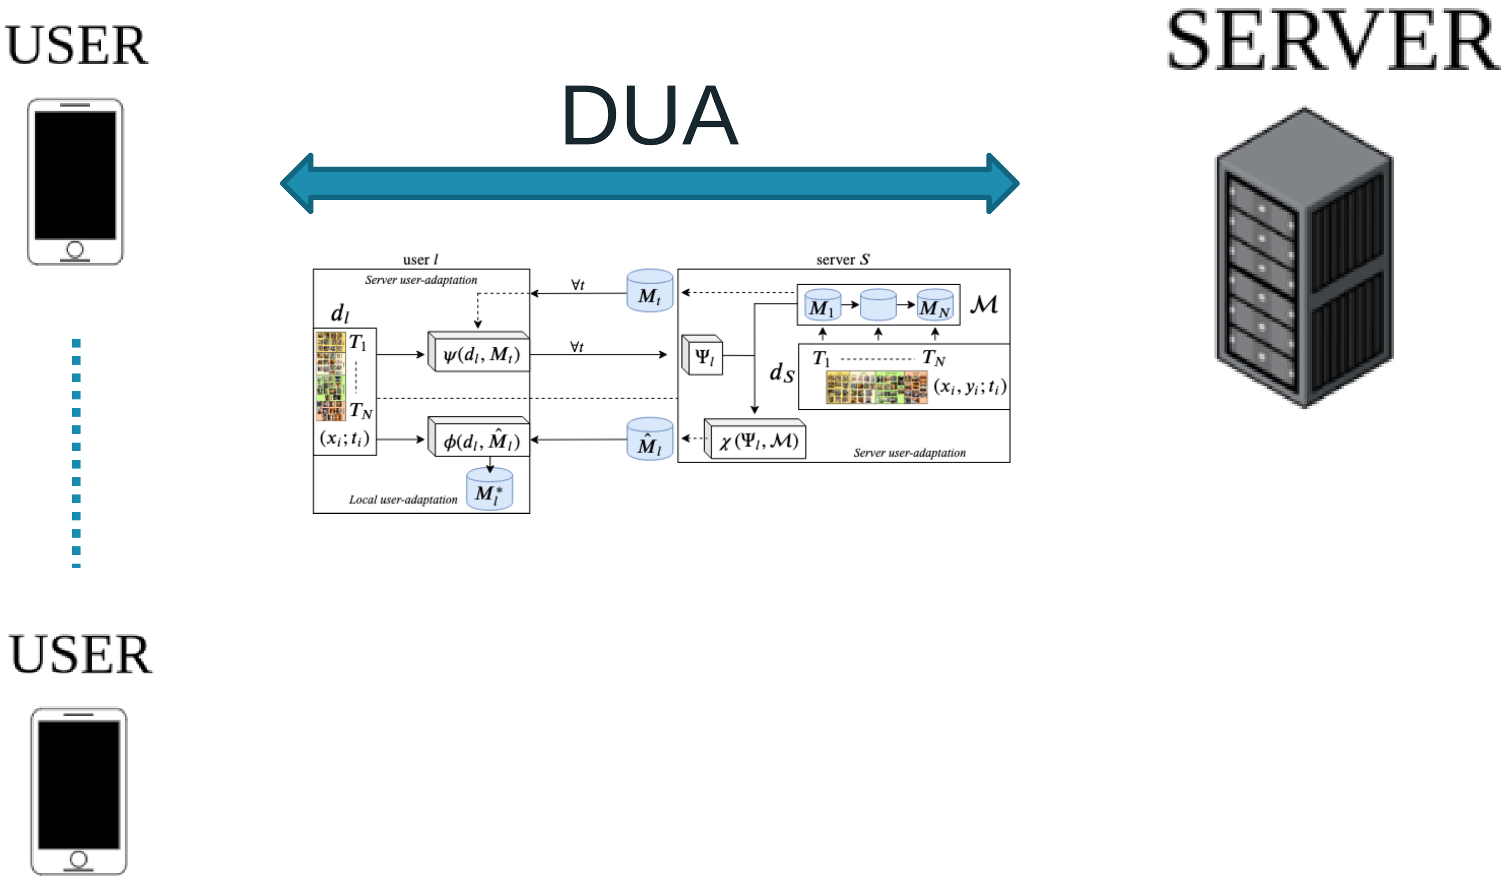

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Roadmap

- Model Personalization: What and how?
- **Novel Benchmarks**
- Adaptation on the server
- Adapting locally
- Conclusion

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

[1] De Lange, Matthias, et al. "A continual learning survey: Defying forgetting in classification tasks." *arXiv preprint arXiv:1909.08383* (2020).

Benchmarks

- 2 components
  - Users with different preferences (prior) → Validation/ Evaluation sets
  - Server with large dataset → Training set

- Task incremental continual learning, see [1]
  - Divide into sequence of tasks

Questions? matthias.delange@kuleuven.be

Benchmarks: Numbers

5 tasks of 2 subsequent numbers

User 1

Test data

MNIST

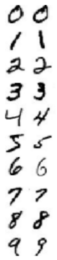

User 2

Test data

SVHN

SERVER

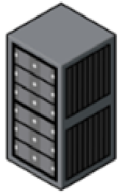

Train data

MNIST

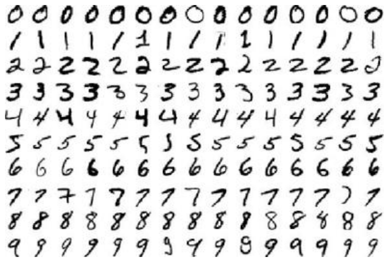

+

Train data

SVHN

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Benchmarks: CatPrior

Each user has 3 local category (scene) preferences per Task (Supercategory)

User 1

. . .

User 5

3 local preferences

3 local preferences

3 local preferences

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

HUAWEI Mila PSI KU LEUVEN

Benchmarks: TransPrior

10 users, Each user has 1 transformation

Spatters, elastic transformation, saturation, defocus blur, Gaussian noise, brightness, Gaussian blur, jpeg compression, contrast and impulse noise.

Roadmap

- Model Personalization: What and how?
- Novel Benchmarks
- **Adaptation on the server**
- Adapting locally
- Conclusion

Questions? matthias.delange@kuleuven.be

Server adaptation in a

fashion?

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

Server adaptation in a

Scalable

Privacy Preserving

Unsupervised

fashion?

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

[2] Lee, Sang-Woo, et al. "Overcoming catastrophic forgetting by incremental moment matching." *Advances in neural information processing systems*. 2017.

Continual Learning

- Continual learning major focus on *Catastrophic Forgetting*
- Many of its properties suit our setting:

Local scalability → limit user resources to model learning multiple tasks

Distributed scalability → limit server resources to many personalized models e.g. task incremental with IMM [2]

Questions? matthias.delange@kuleuven.be

Continual Learning

Distributed scalability with IMM

1. Learn task-specific models
2. Get model importance weights
3. Merge using importance weights

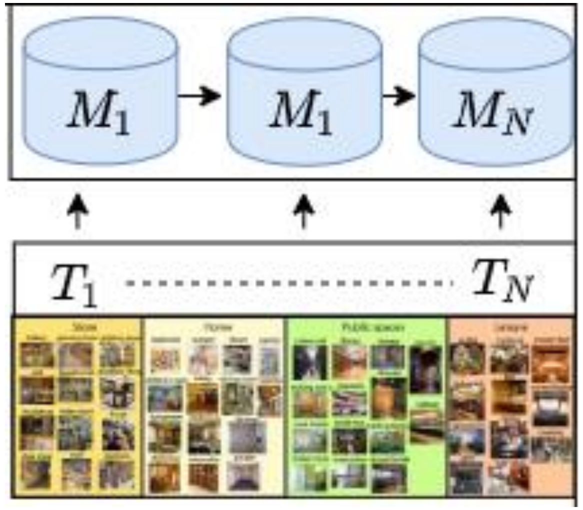

→ #models = #tasks

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

Dual User-Adaptation framework (DUA)

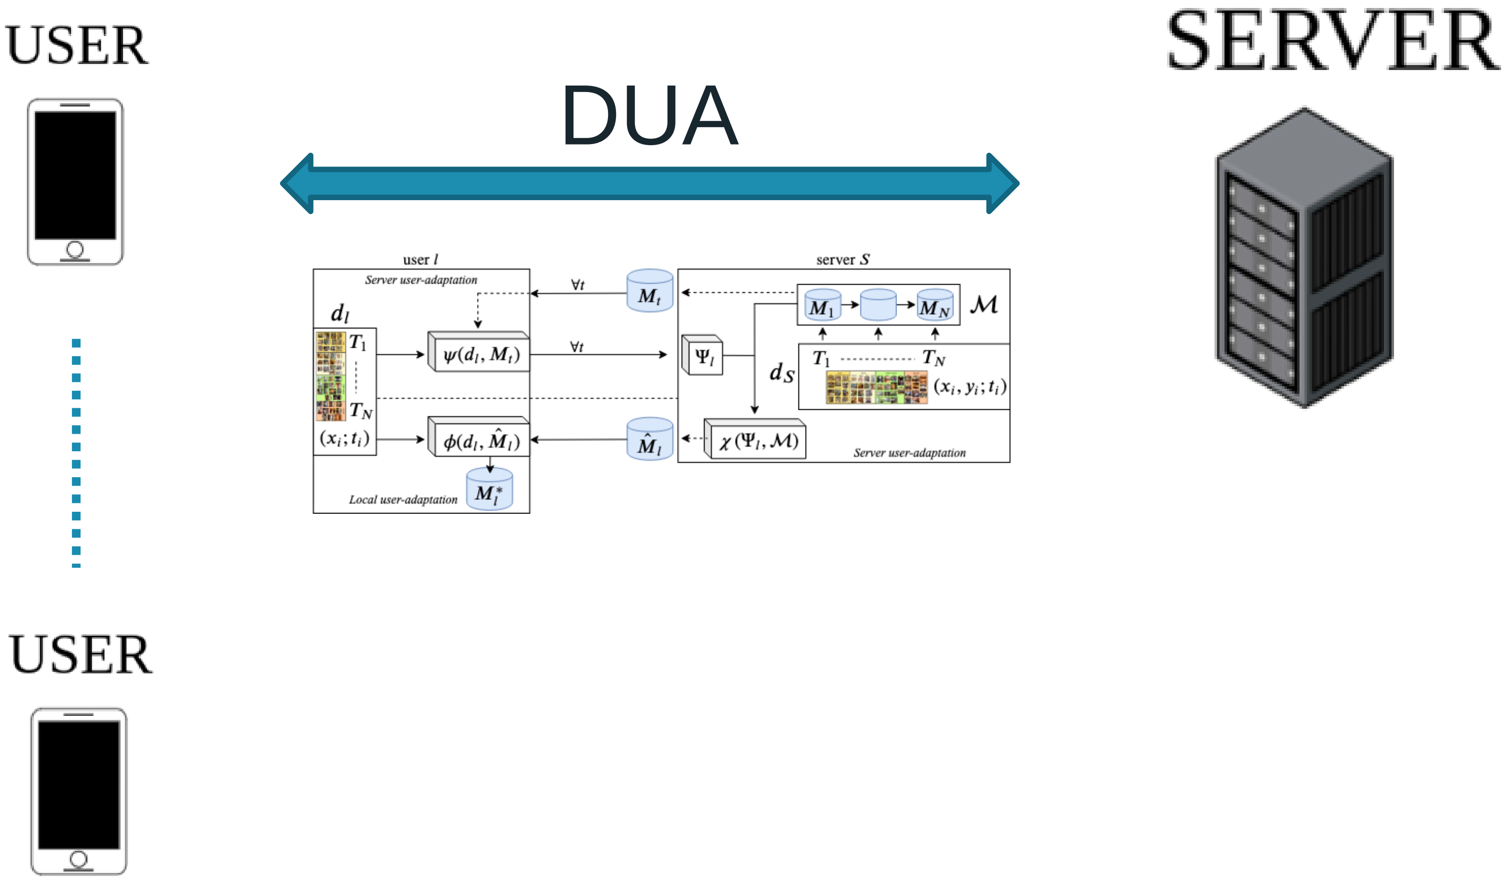

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

(1) Serverside Adaptation

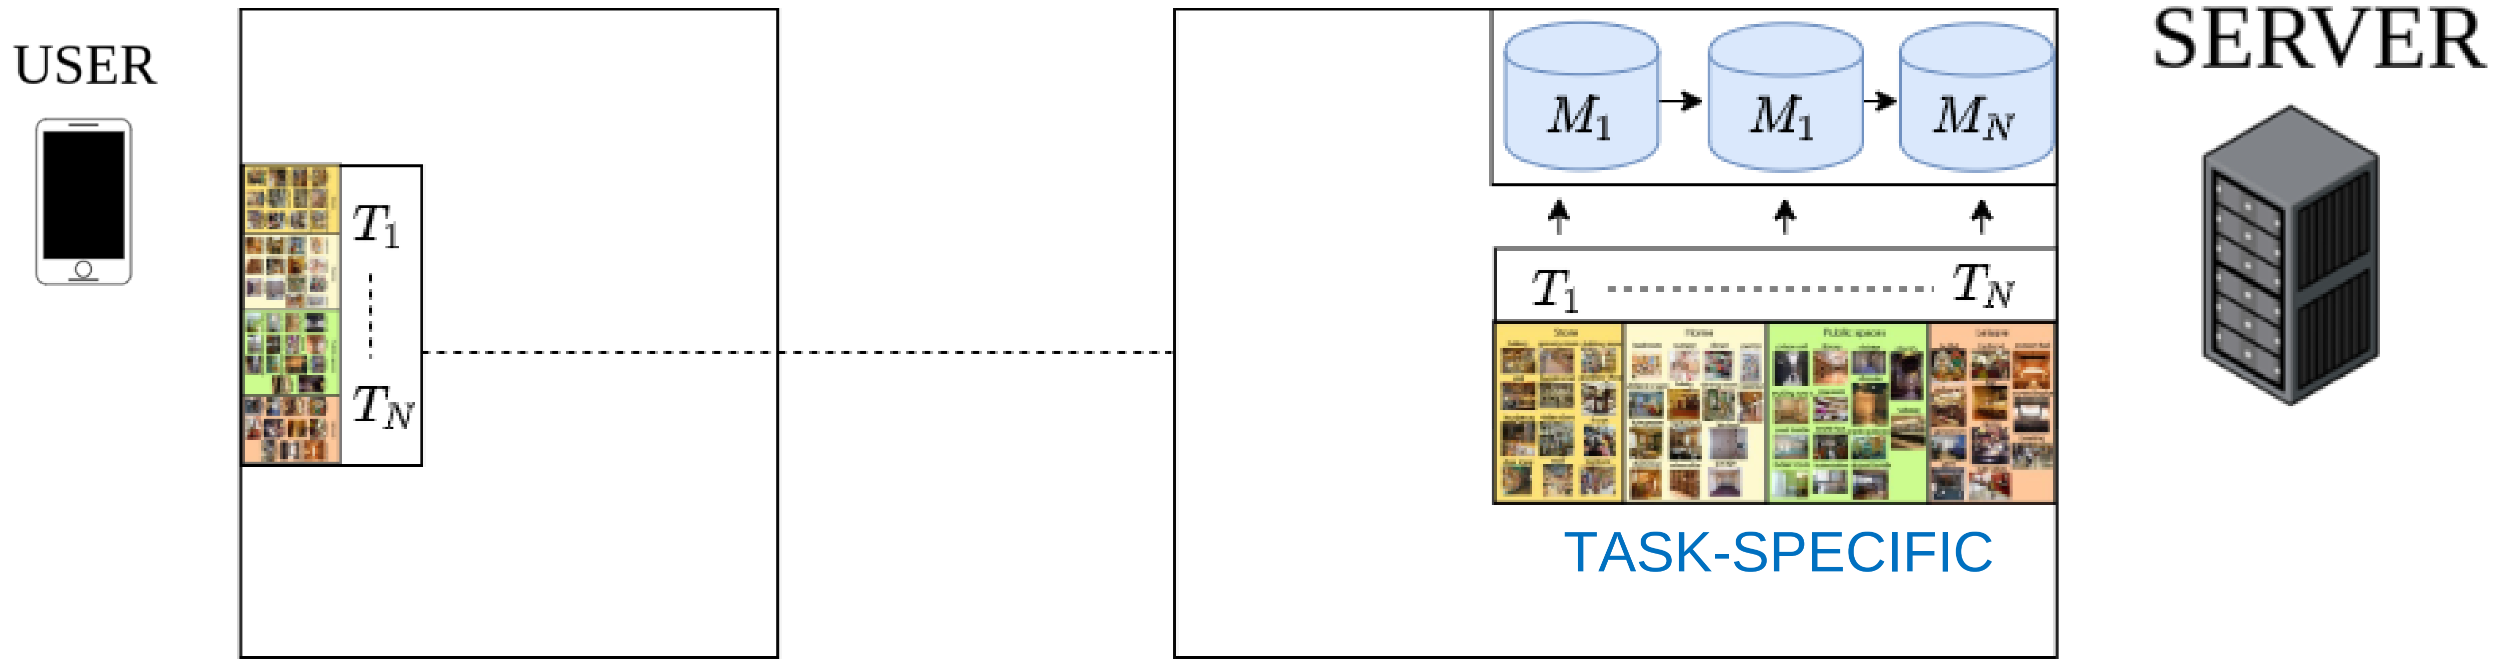

✓ Scalable

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

(1) Serverside Adaptation

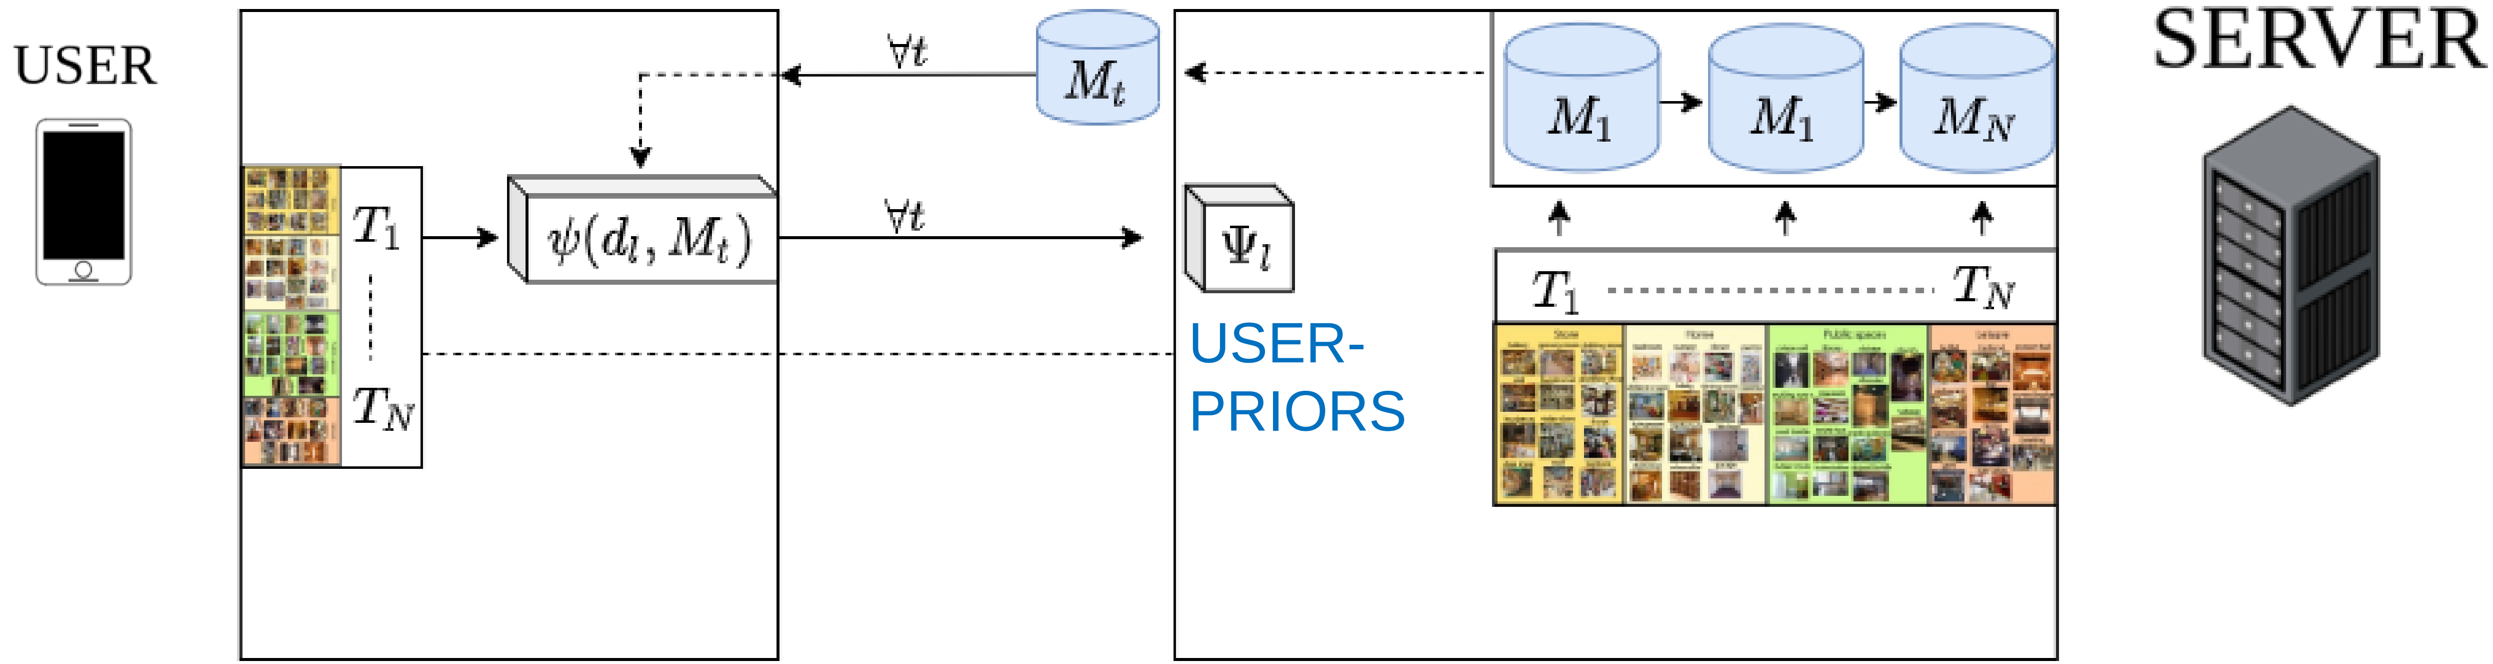

- ✓ Scalable
- ✓ Privacy

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

(1) Serverside Adaptation

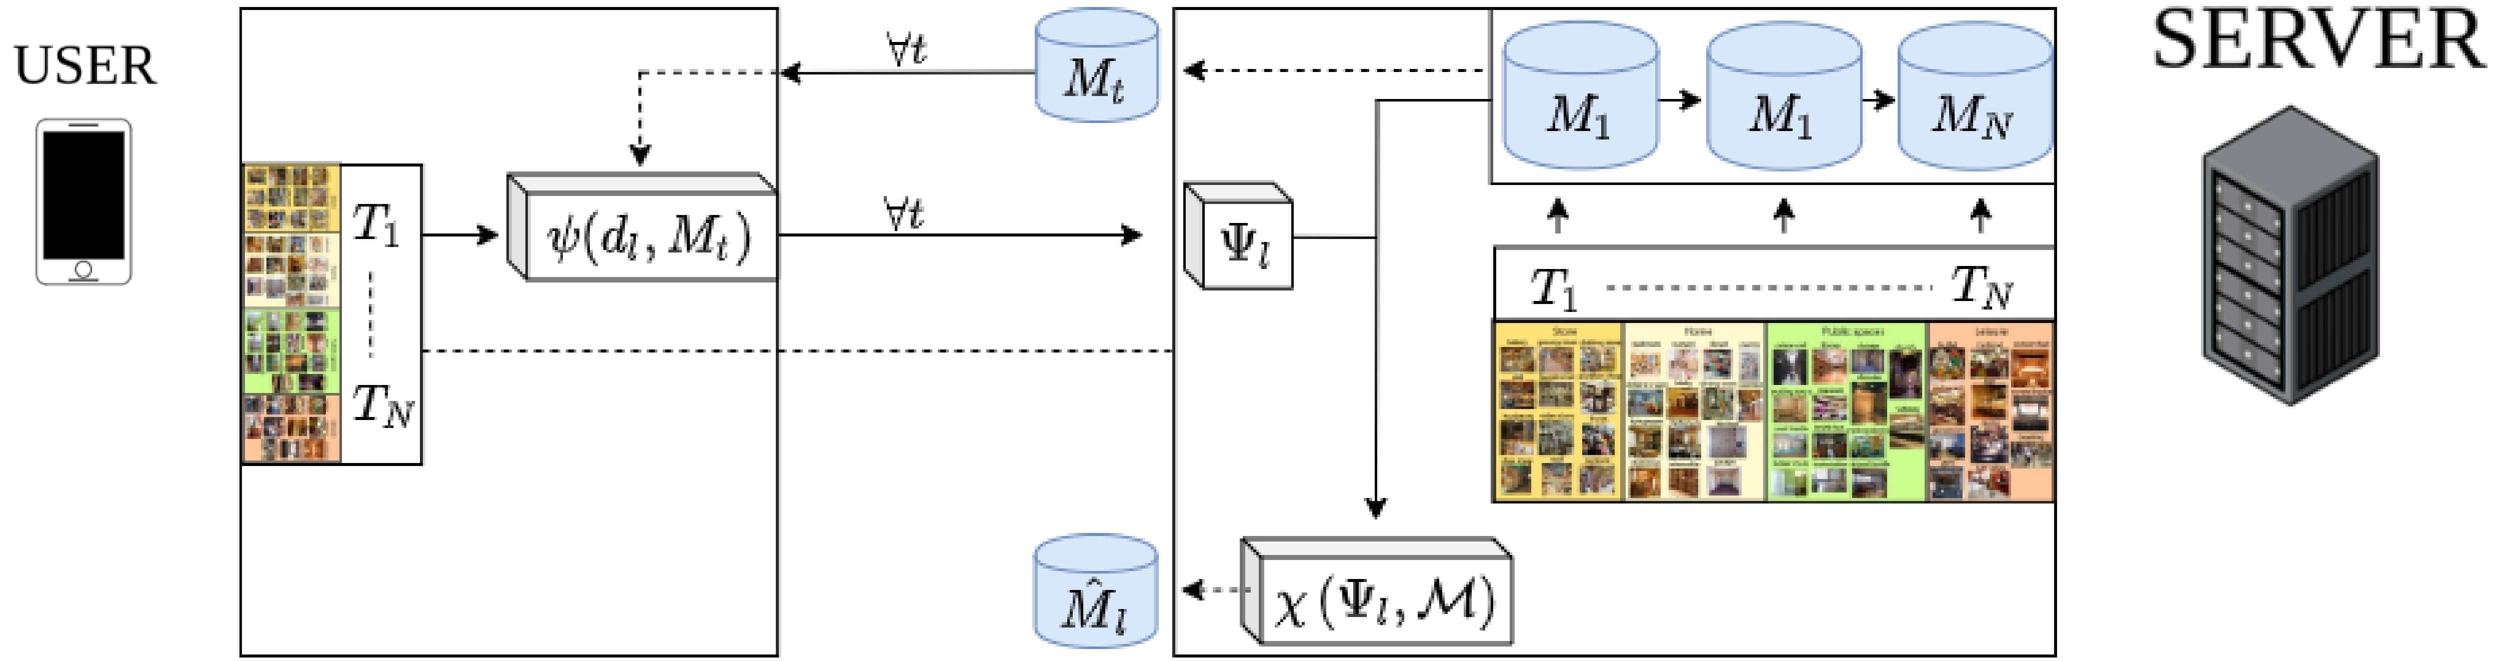

AGGREGATE

✓ Scalable
✓ Privacy

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

Supervision

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

[3] Aljundi, Rahaf, et al. "Memory aware synapses: Learning what (not) to forget." *(ECCV)*. 2018.

Merging IMM-models

- Originally importance weights diagonal Fisher Information Matrix (FIM)
  - Loss-based → Requires labels

- Instead unsupervised MAS [3] importance weights, based on output function?

Data Setup	Model	MAS-IMM	FIM-IMM
CatPrior	AlexNet	67.39 (0.73)	67.42 (0.23)
	VGG11	76.77 (0.30)	76.29 (0.43)
TransPrior	AlexNet	46.51 (-0.14)	46.68 (-0.35)
	VGG11	53.49 (-0.17)	53.14 (0.07)
Numbers	MLP	84.36 (-0.40)	87.68 (0.07)

Remote Adaptive Continual Learning (RACL)

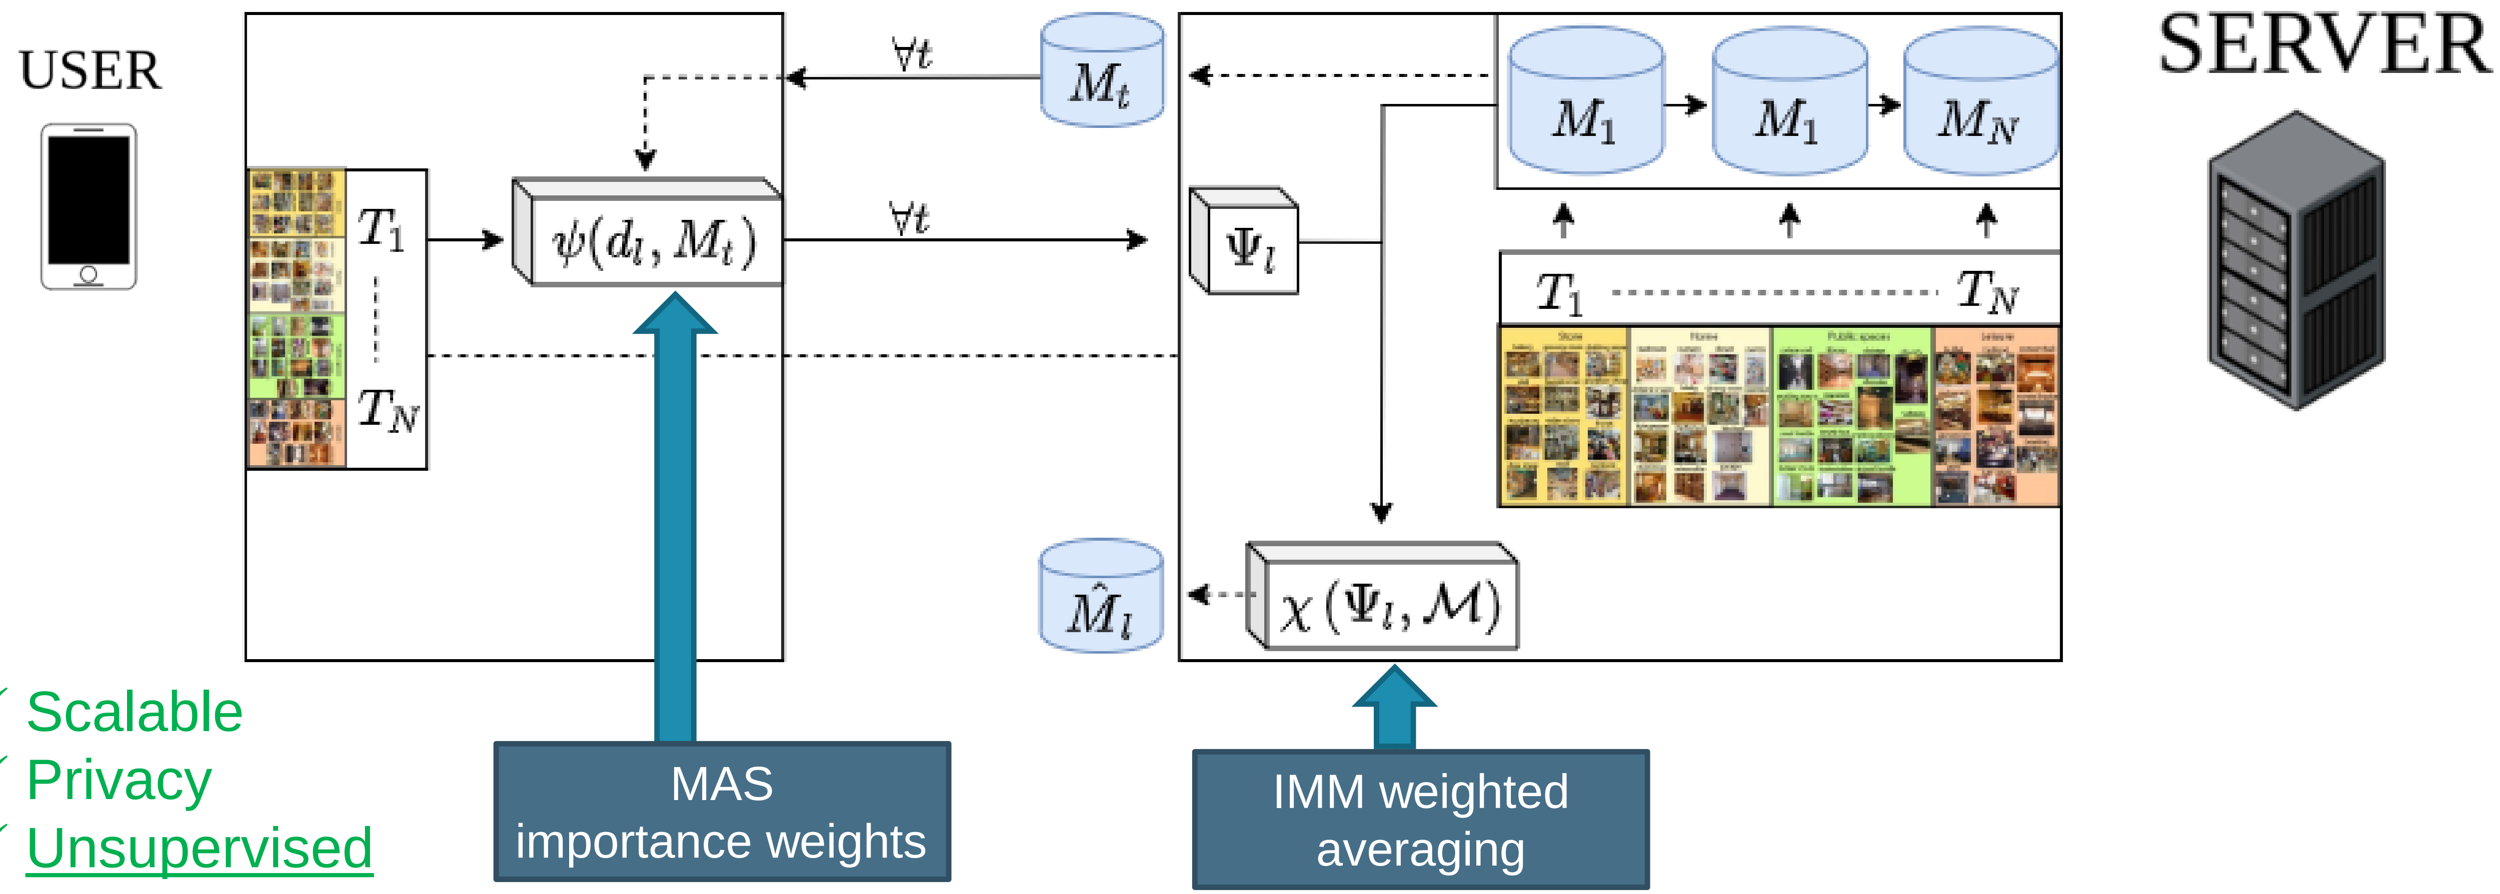

- ✓ Scalable
- ✓ Privacy
- ✓ Unsupervised

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

User-Adaptation

- RACL/IMM → User-specific/General model
- MAS/FIM → Unsupervised/supervised importance weights

Method	**Alexnet**		**VGG11**		**MLP**	**Adapt.**	**Unsup.**
	CatPrior	*TransPrior*	*CatPrior*	*TransPrior*	*Numbers*		
MAS-RACL	66.97 (0.88)	**47.04 (-0.27)**	**77.32 (0.77)**	**53.59 (-0.14)**	84.01 (-0.22)	✓	✓
MAS-IMM	67.39 (0.73)	46.51 (-0.14)	76.77 (0.30)	53.49 (-0.17)	84.36 (-0.40)	✗	✓
FIM-RACL	67.20 (0.73)	**47.32 (-0.51)**	**76.53 (0.68)**	**53.73 (-0.13)**	**87.83 (0.30)**	✓	✗
FIM-IMM	67.42 (0.23)	46.68 (-0.35)	76.29 (0.43)	53.14 (0.07)	87.68 (0.07)	✗	✗

Improvements insignificant → Why?

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Analyzing Importance Weights

We found consistently:

1. High correlations for different datasets on a same model
2. Low correlation for same data on different models

Importance weights indicate parameter importance for the specific model, rather than the data they are estimated from!

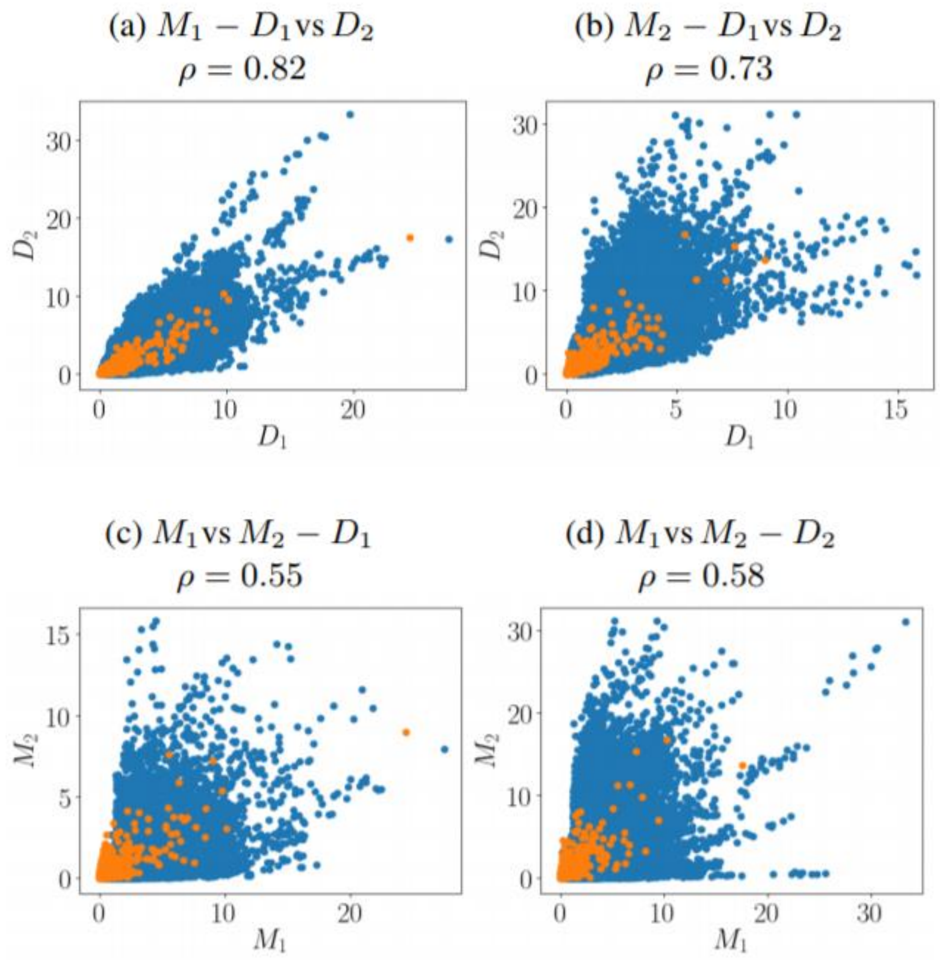

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Roadmap

- Model Personalization: What and how?
- Novel Benchmarks
- Adaptation on the server
- **Adapting locally**
- Conclusion

Questions? matthias.delange@kuleuven.be

(1) Serverside Adaptation

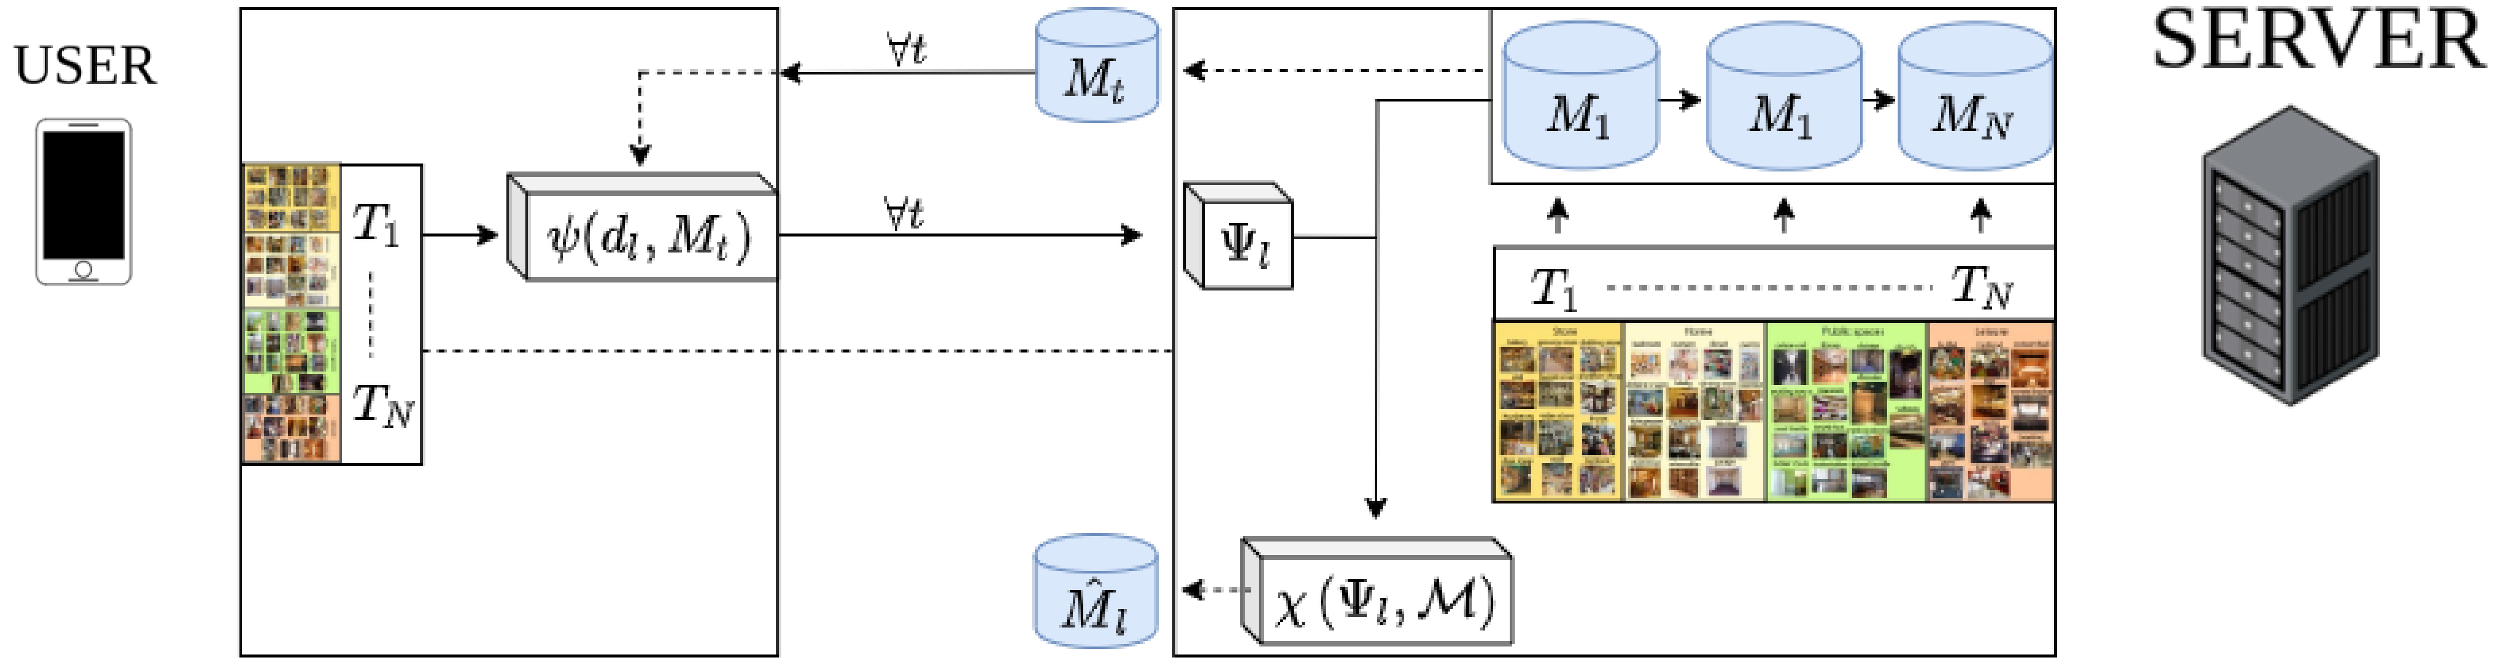

- ✓ Scalable
- ✓ Privacy
- ✓ Unsupervised

HUAWEI Mila

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

PSI KU LEUVEN

(2) Local Adaptation

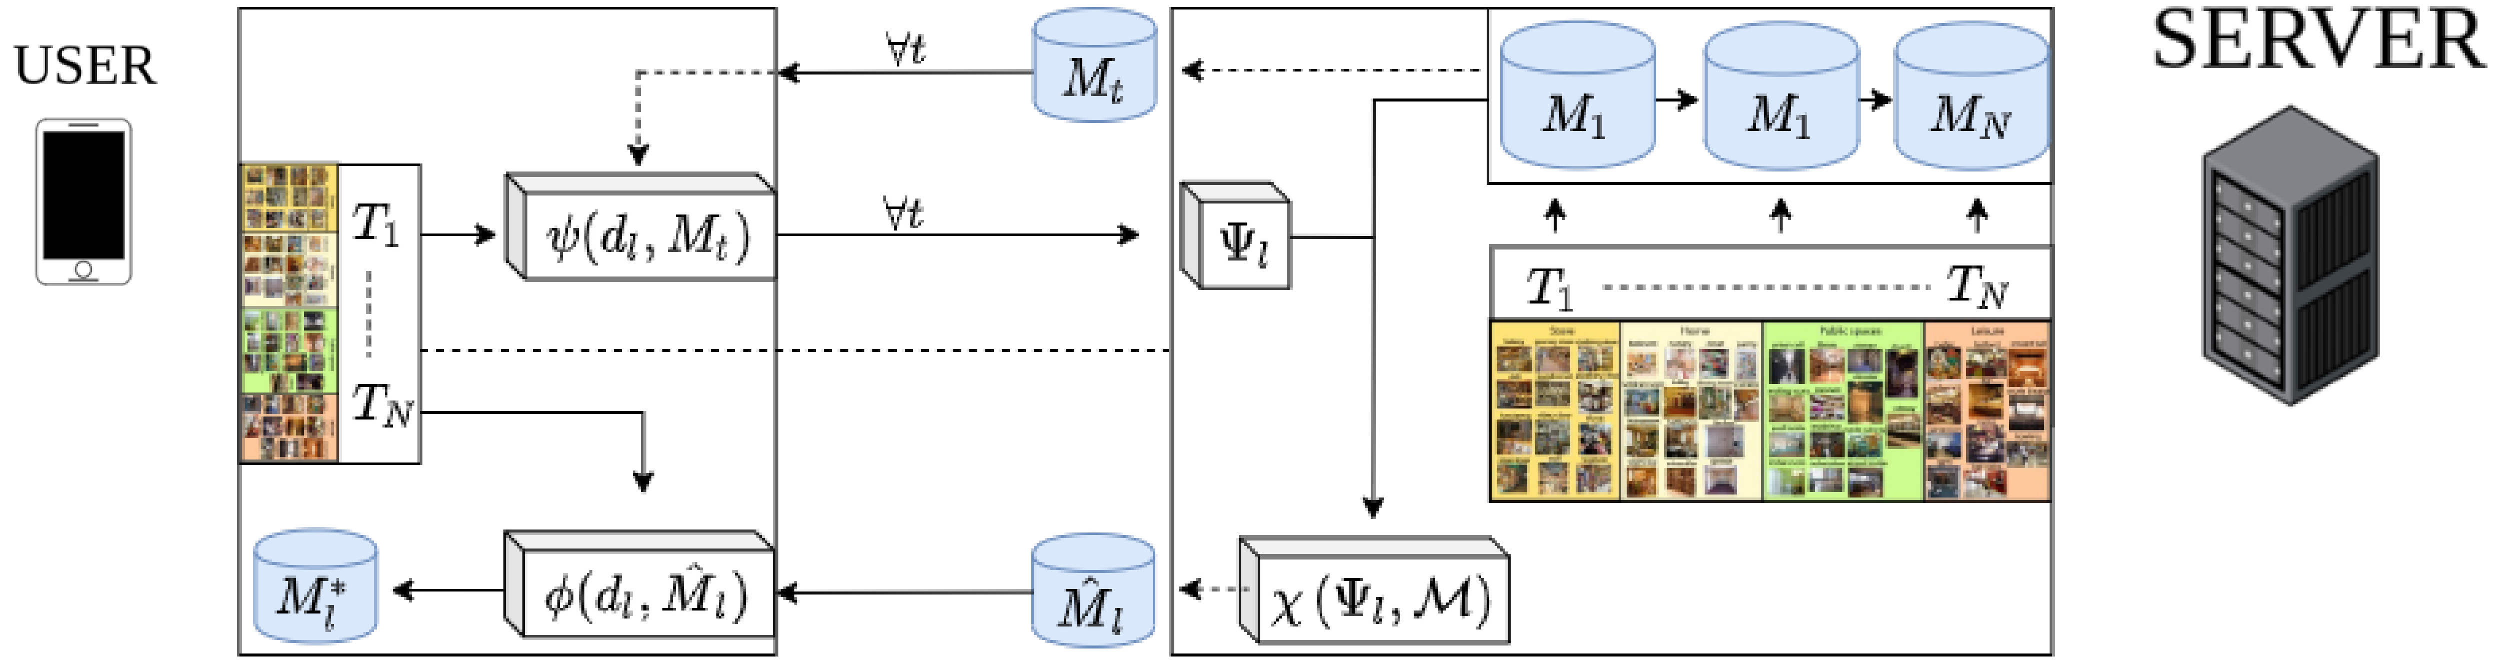

LOCAL ADAPTATION

- ✓ Scalable
- ✓ Privacy
- ✓ Unsupervised
- ✓ 2 x Adaptation

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

Local Adaptation

- Adapt from general Server domain → User domain
- With limited user resources
  - Adapt Batch Normalization stats (AdaBN) → Unsupervised
  - Train few BN parameters (AdaBN-S) → Supervised

Method	**CatPrior**			**TransPrior**		
	BN	*AdaBN*	*AdaBN-S*	*BN*	*AdaBN*	*AdaBN-S*
MAS-RACL	58.05 (2.74)	58.30 (2.34)	60.68 (2.67)	30.14 (2.69)	30.19 (2.50)	32.82 (3.25)
FIM-RACL	59.58 (2.14)	59.71 (1.61)	62.43 (1.84)	32.15 (1.53)	32.04 (1.33)	34.80 (2.13)
Task Experts	80.78 (5.61)	n/a	n/a	68.22 (11.35)	n/a	n/a
MAS-IMM	55.55 (2.69)	55.89 (2.69)	58.87 (2.81)	29.36 (2.63)	29.15 (2.45)	31.73 (3.22)
FIM-IMM	61.50 (-0.03)	61.35 (-0.46)	63.99 (-0.16)	32.08 (1.32)	31.86 (1.21)	34.48 (2.05)
MAS	65.58 (3.96)	64.15 (4.04)	67.10 (4.66)	37.32 (2.64)	35.64 (2.88)	40.51 (2.69)
EWC	66.20 (2.88)	64.03 (3.43)	67.54 (3.90)	37.16 (2.85)	35.44 (3.12)	40.05 (3.18)
LWF	70.76 (0.73)	70.37 (0.43)	72.73 (1.03)	40.22 (0.43)	39.51 (0.12)	43.07 (0.52)
Joint	75.75 (n/a)	72.13 (n/a)	76.39 (n/a)	46.53 (n/a)	41.18 (n/a)	48.50 (n/a)

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Local Adaptation

- Still open problem lightweight, unsupervised domain adaptation
- Relaxing unsupervised user training → Consistent improvements ≈ 3%

Method	CatPrior			TransPrior		
	BN	*AdaBN*	*AdaBN-S*	*BN*	*AdaBN*	*AdaBN-S*
MAS-RACL	58.05 (2.74)	58.30 (2.34)	60.68 (2.67)	30.14 (2.69)	30.19 (2.50)	32.82 (3.25)
FIM-RACL	59.58 (2.14)	59.71 (1.61)	62.43 (1.84)	32.15 (1.53)	32.04 (1.33)	34.80 (2.13)
Task Experts	80.78 (5.61)	n/a	n/a	68.22 (11.35)	n/a	n/a
MAS-IMM	55.55 (2.69)	55.89 (2.69)	58.87 (2.81)	29.36 (2.63)	29.15 (2.45)	31.73 (3.22)
FIM-IMM	61.50 (-0.03)	61.35 (-0.46)	63.99 (-0.16)	32.08 (1.32)	31.86 (1.21)	34.48 (2.05)
MAS	65.58 (3.96)	64.15 (4.04)	67.10 (4.66)	37.32 (2.64)	35.64 (2.88)	40.51 (2.69)
EWC	66.20 (2.88)	64.03 (3.43)	67.54 (3.90)	37.16 (2.85)	35.44 (3.12)	40.05 (3.18)
LWF	70.76 (0.73)	70.37 (0.43)	72.73 (1.03)	40.22 (0.43)	39.51 (0.12)	43.07 (0.52)
Joint	75.75 (n/a)	72.13 (n/a)	76.39 (n/a)	46.53 (n/a)	41.18 (n/a)	48.50 (n/a)

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

Roadmap

- Model Personalization: What and how?
- Novel Benchmarks
- Adaptation on the server
- Adapting locally
- **Conclusion**

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

Dual User-Adaptation

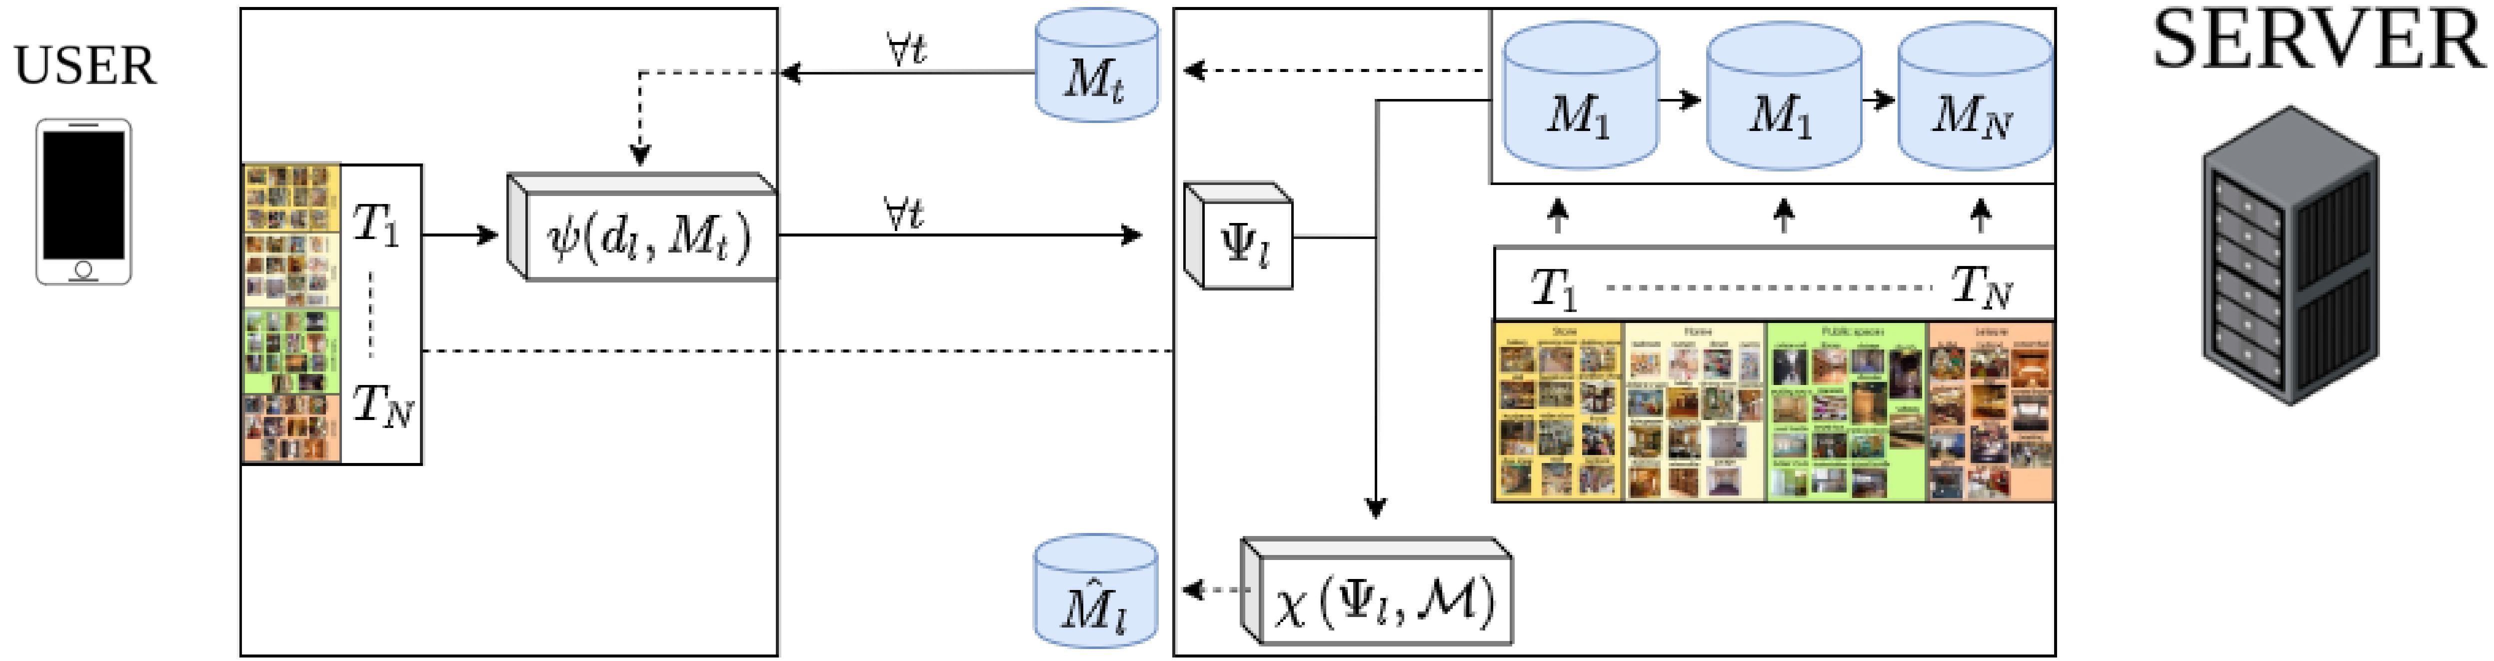

- ✓ Scalable
- ✓ Privacy
- ✓ Unsupervised

(2) Local Adaptation

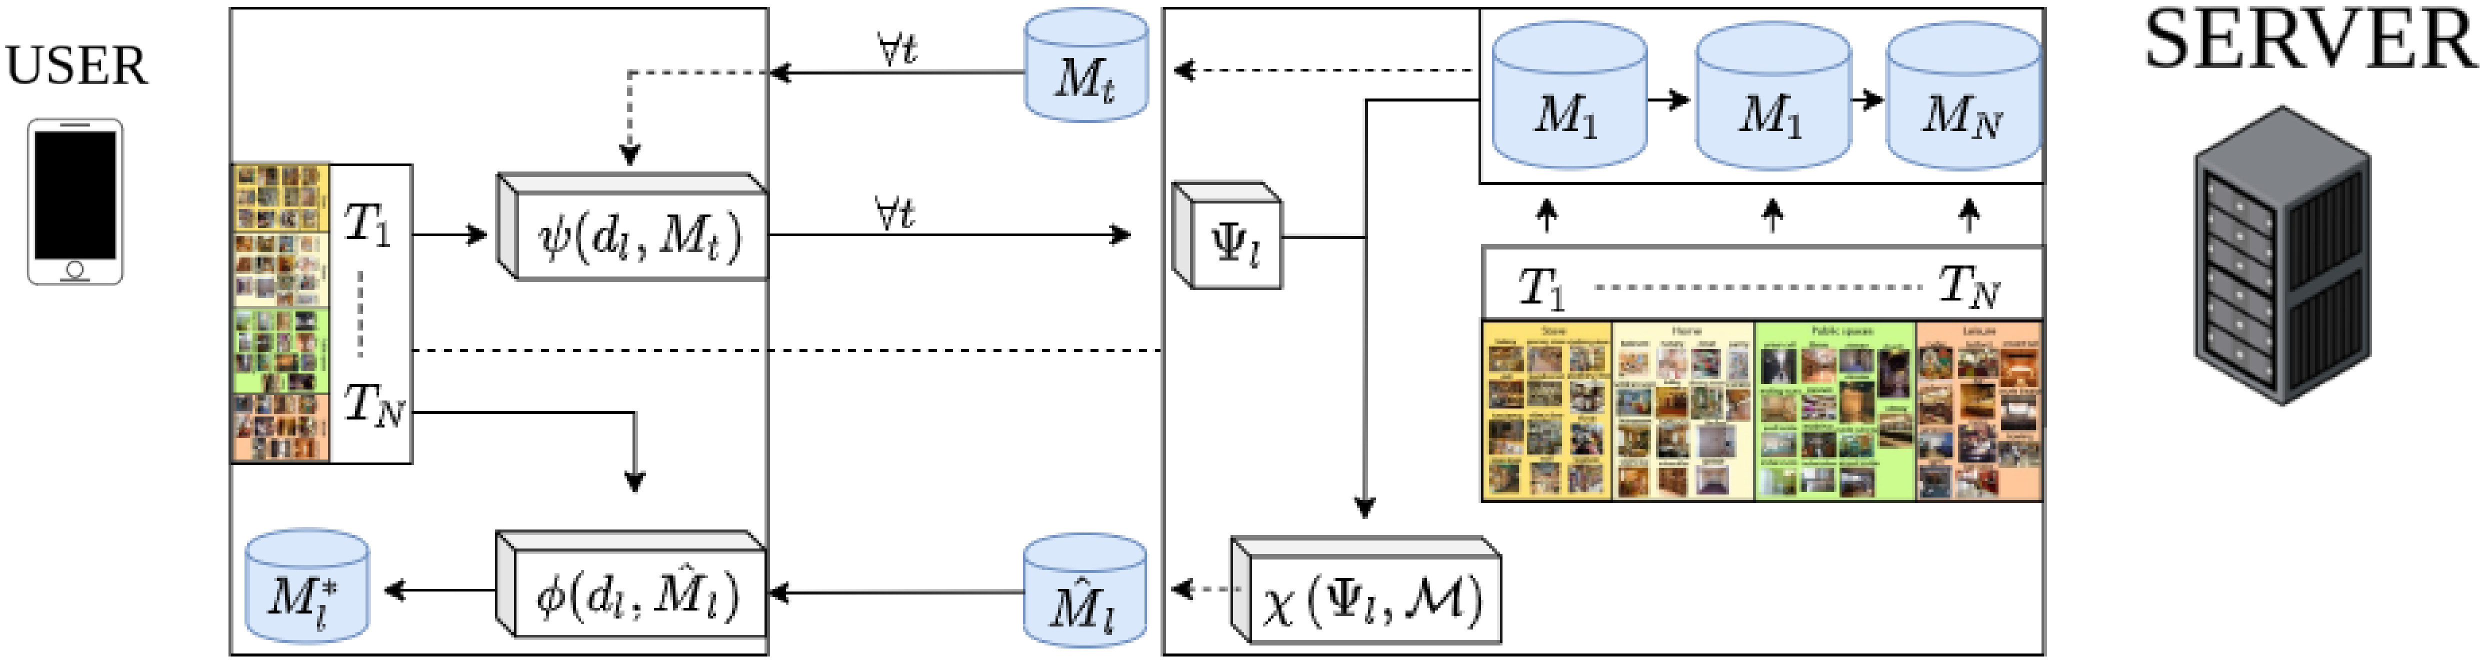

- ✓ Scalable
- ✓ Privacy
- ✓ Unsupervised
- ✓ 2 x Adaptation

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.“, CVPR (2020).

Questions? matthias.delange@kuleuven.be

(2) Local Adaptation

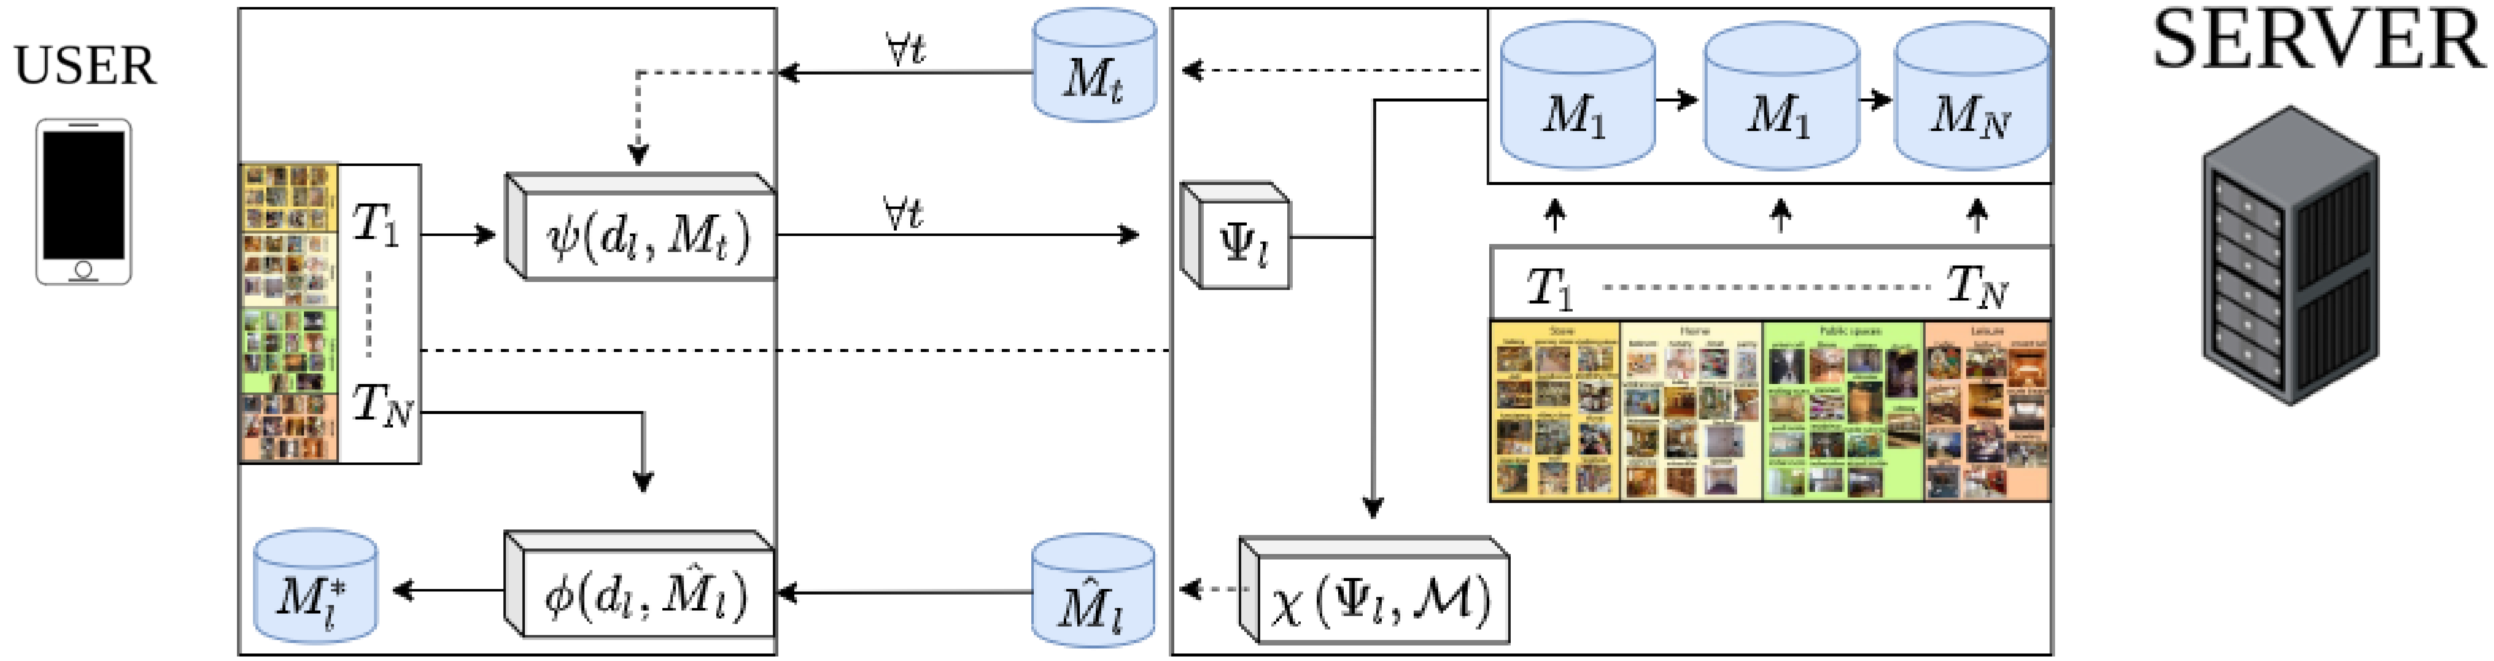

✓ Scalable
✓ Privacy
✓ Unsupervised
✓ 2 x Adaptation

Open Problem:
- Data-dependent importance weights
- Domain Adaptation

De Lange, Matthias, et al. "Unsupervised Model Personalization while Preserving Privacy and Scalability: An Open Problem.", CVPR (2020).

Questions? matthias.delange@kuleuven.be

Code

https://github.com/mattdl/DUA

Questions?

matthias.delange@kuleuven.be