



# Privacy-Preserving Probabilistic Forecasting in Local Energy Communities

Jean-François Toubreau, from 'Energy Systems Integration & Modeling' - ESIM

[Jean-francois.toubreau@kuleuven.be](mailto:Jean-francois.toubreau@kuleuven.be)

# Content of the presentation

- Context:  
Why do we need forecasts in local energy communities ?
- Methodology:  
How can we provide privacy by-design forecasting models ?
  - Federated learning → All data are kept local
  - Horizontal (cross-series) learning → deal with users with different history
  - Differential privacy → to prevent inference of raw data from the trained model
- Case study for analyzing the value of collaborative learning
- Conclusions (take home messages) & perspectives

# Local energy communities



Context: The integration of distributed energy resources (DER), such as photovoltaics and electric vehicles, complicates the operation states of the power distribution grids.

Global problem: nodal voltage may violate frequently (**issues for equipment**).

Targeted solution: To proactively manage local energy exchanges, e.g., through local energy communities (LECs).


However, to ensure optimal coordination between resources, the community needs to be informed with accurate predictions of the future system state.

# Goal of the work

Develop a new framework for the short-term probabilistic forecasting of nodal voltage magnitudes, exploiting the information from smart metering devices.

Nodal voltages are governed by intricate:

- time correlations
- space dependencies arising from network constraints (i.e., neighboring buses are likely to exhibit similar voltage patterns).



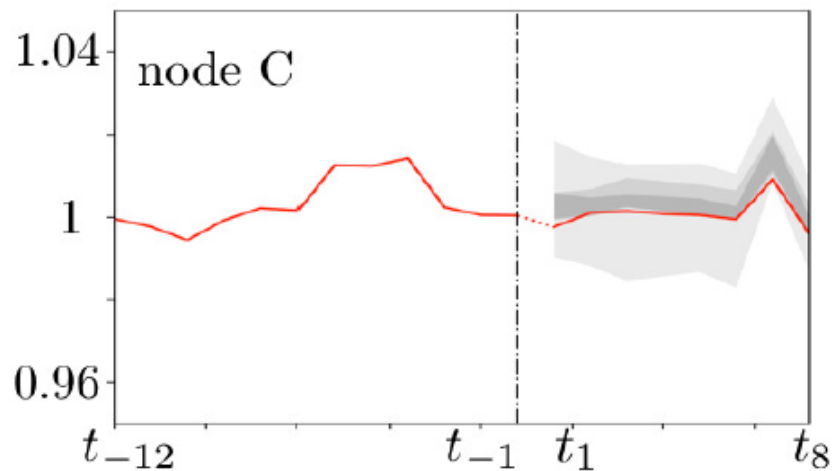
# Quantifying the value of collaboration

## Case study:


IEEE European Low Voltage Test Feeder

Multi-horizon of 8 intervals of 30 minutes (i.e., 4 hour ahead)

A look-back window of  $k = 12$  intervals (i.e., 6 hours) is selected to capture past dynamics



Legend: — real value  
— quantiles [0.01, 0.99]  
— quantiles [0.25, 0.75]



## Input features:

Past data $\mathbf{x}_{:t_0,n}^{(p)}$	past (local) load, past (aggregated) PV generation, past (aggregated) imports-exports, calendar information (weekday, hour of the day)
Known future data $\mathbf{x}_{t_1:T,n}^{(f)}$	forecasted (aggregated) imports-exports, calendar information (weekday, hour of the day)
Static data $x_n^{(s)}$	node location (feeder to which the node is connected, and distance to the origin node)

# Quantifying the value of collaboration

- 1) (Non-private) centralized models relying on complete information, i.e., an ideal (but unrealistic) case where each end-user shares its private data to form a single database, from which the community manager trains a single model.
- 2) Fully private models, where a different model is trained by each end-user based only on its own private data.

PERFORMANCE AND TRAINING TIME OF DIFFERENT MODELS.

Model	Centralized		Local (private)	
	$QL^{\text{tot}}$ [pu]	time [min] (epochs)	$QL^{\text{tot}}$ [pu]	time [min] (epochs)
Prob-Persist	-	-	0.210	-
Prob-Avg	0.319	-	0.192	-
QRF	0.048	148	0.156	172
QGBDT	0.041	64	0.151	76
DFNN	0.043	10 (33)	0.144	16 (27)
LSTM	0.034	52 (18)	0.186	63 (4)
BLSTM	0.032	88 (16)	0.172	105 (4)
A-BLSTM	0.192	54 (11)	-	-

**Quantile loss** over a test set of 96 days for 57 end-users:

- High differences in performance between centralized and local models, which reflects the importance of capturing nodal dependencies between voltage levels
- RNNs fail with data scarcity

# Towards privacy-preserving forecasting



Specific problem: Traditional algorithms assume that private data can be freely accessed from a centralized location.

The nodal (e.g., smart meter) data are owned by end-users, who may be reluctant to share this information (as it may reveal private aspects such as home occupancy, routines, and usage of specific appliances).

➔ It is important to implement new collaborative learning strategies, leveraging the raw data of all end-users (capture dependencies among stakeholders), while complying with privacy requirements.

# Vertical versus Horizontal learning






Client A

$x_1^A$   $x_2^A$   $y^A$

Sample 1			
Sample 2			
Sample 3			

Client B

$x_1^B$   $x_2^B$   $y^B$



















		
		
		

$x_1$  Consumption, PV generation, past voltage levels, node location, etc.  
 $x_2$   
 $y$  Voltage level





















Horizontal learning

$x_1$   $x_2$   $y$

Sample 1			
Sample 2			
Sample 3			
Sample 4			
Sample 5			
Sample 6			

Vertical learning



















$x_1^A$   $x_2^A$   $x_1^B$   $x_2^B$   $y^A$   $y^B$

Sample 1						
Sample 2						
Sample 3						



# Vertical versus Horizontal learning

## Horizontal learning



















Sample 1			
Sample 2			
Sample 3			
Sample 4			
Sample 5			
Sample 6			

Data efficient → more samples to train the model

Can handle cold-start forecasting (for clients with very little history)

Global model → Cross-series learning for space dependencies

## Vertical learning

	$x_1^A$	$x_2^A$	$x_1^B$	$x_2^B$		$y^A$	$y^B$
Sample 1							
Sample 2							
Sample 3							

Intuitively more suited for capturing dependencies between clients

More collaborative by design → less prone to trust issues

All clients need to have the same history (database is limited to the client with the shortest history)


# Horizontal Federated Learning (FL)

Federated learning is a *distributed approach where a community of clients is coordinated by a central server to learn a global model, without sharing any raw client data.*

- In FL, private local data stays local, and only the trained model parameters are shared.
  - ➔ FL provides the privacy of local data
  - ➔ FL avoids expensive data transfer
- FL support an arbitrary (even dynamic) number of client nodes, which are coordinated through one central server
- FL naturally hedge against local data scarcity

# Federated Averaging algorithm

[1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In Proc. of AISTATS, pages 1273– 1282. PMLR, 2017.



(1) the server initializes the parameters of the model

(2) the server broadcasts the global model to the users

(3) each selected user performs local computations on its private dataset, and computes the new (locally) optimized model

(4) The local updates are then uploaded to the server

(5) The new global model is then computed, and the procedure is iterated until convergence

# Limitations of Federated Learning

- During the federated learning, the model can be accessed by adversaries to infer raw local information.
- It has been shown that **trained models may be reverse-engineered** to extract detailed input information from the end-users involved in the training phase
- Encrypted computations can be used to protect the training procedure, but
  - possibility to break current cryptographic functions
  - encryption schemes can be computationally expensive

# Differential Privacy (DP)

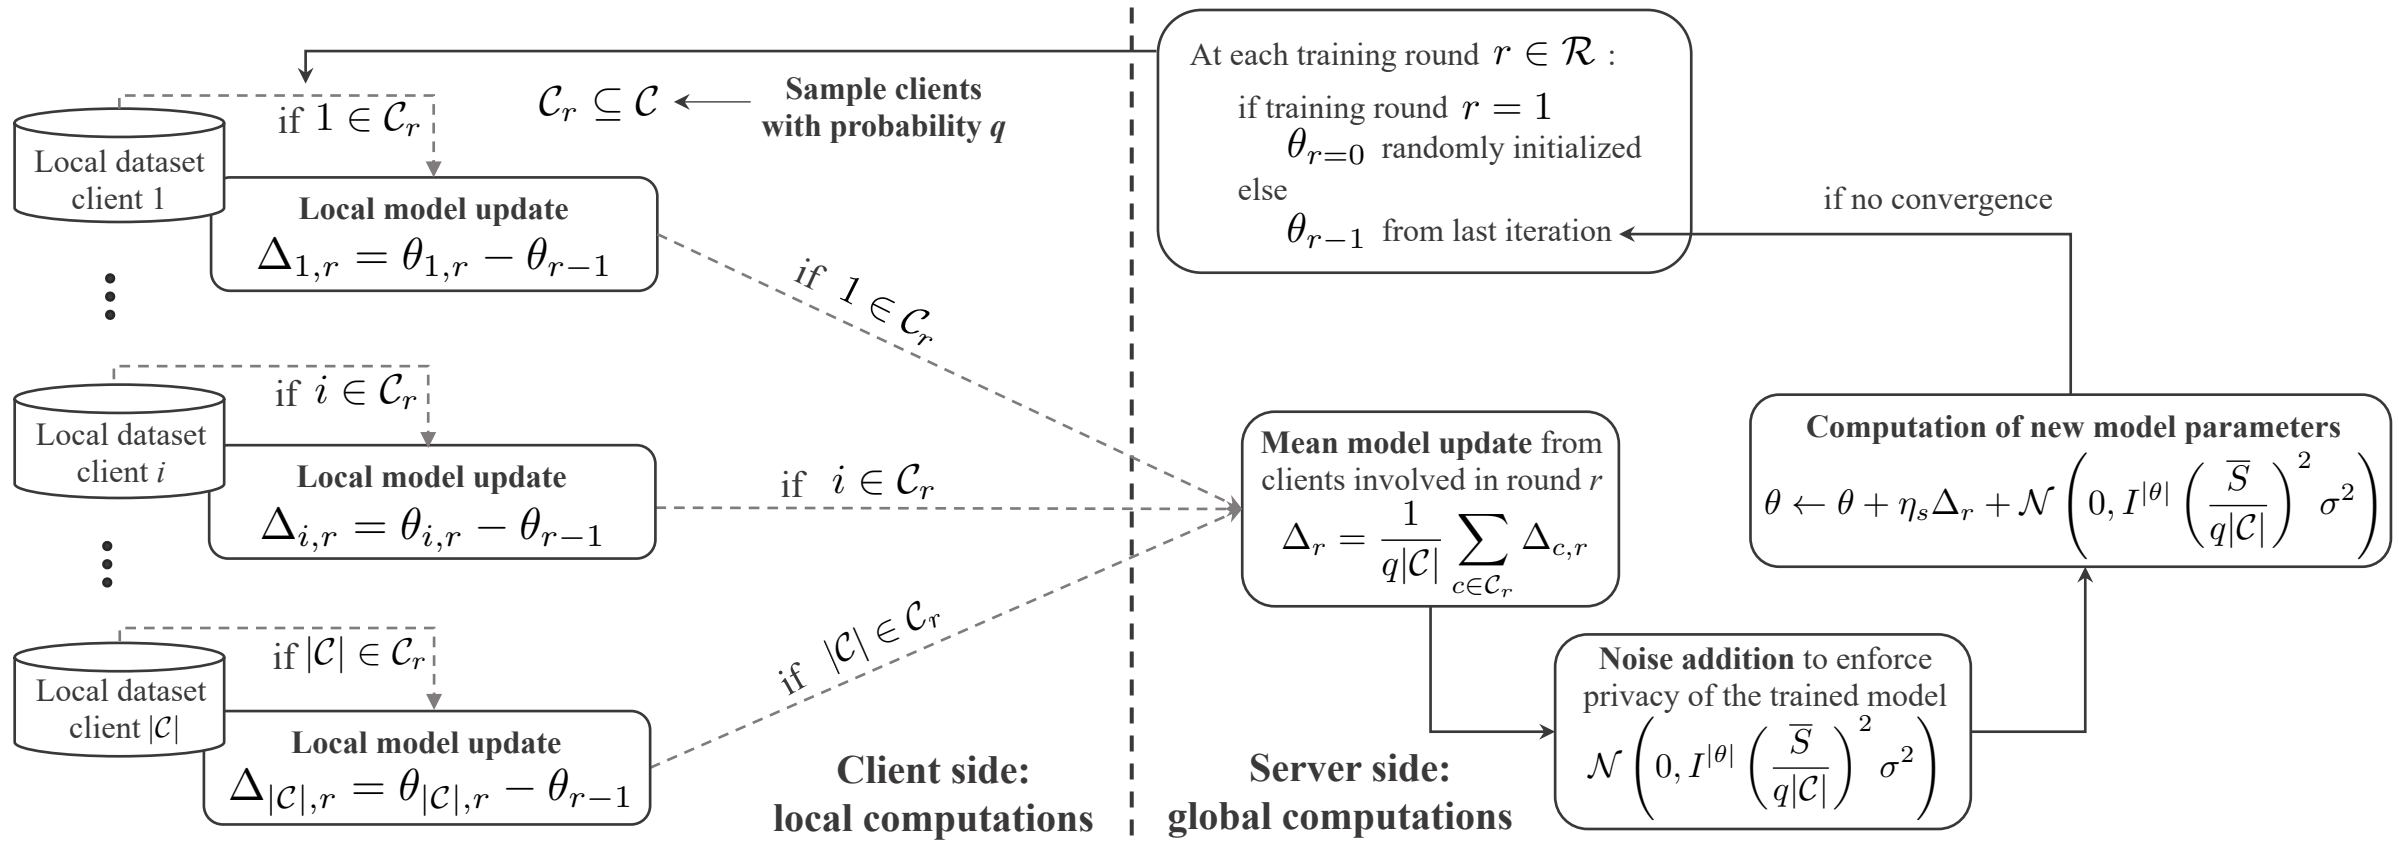
- DP enables to bound and quantify the privacy leakage of sensitive information when performing learning tasks.
- DP is based on the notion of adjacent databases  $D$  and  $D'$ , which differ by the addition or removal of a single user.

$(\epsilon, \delta)$ -DP:

$$\Pr(\underbrace{\mathcal{M}(D)}_{=\theta_1} \in S) \leq e^\epsilon \Pr(\underbrace{\mathcal{M}(D')}_{=\theta_2} \in S) + \delta$$

- $\epsilon$  is the privacy loss yielding an upper bound of how much the probability of converging to a particular set of weights is affected by including (or removing) a single client during training
- no bound on  $\epsilon$  with probability  $\delta$

# Proposed framework




# Quantifying the value of collaboration

## Case study:

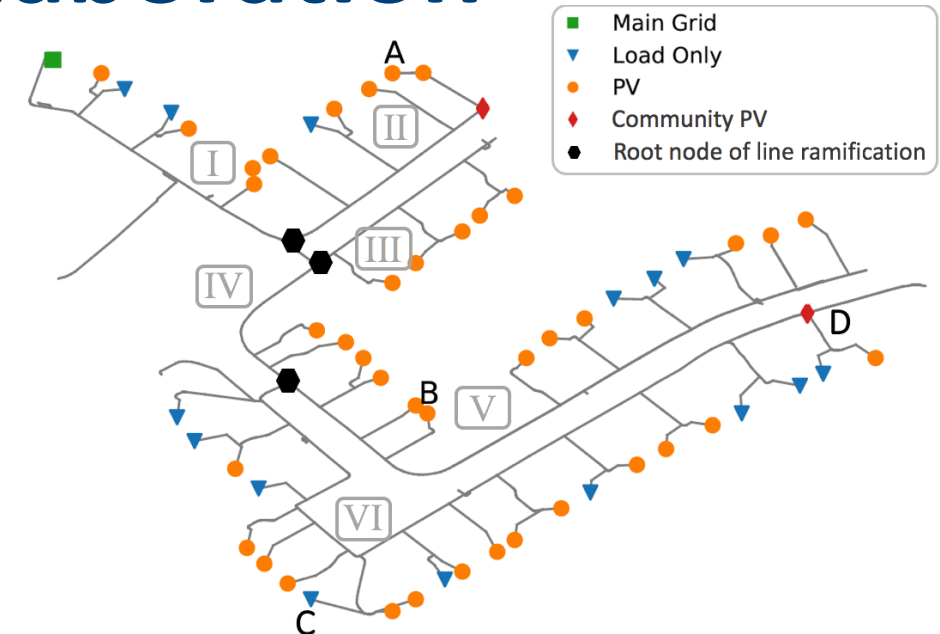
IEEE European Low Voltage Test Feeder

Multi-horizon of 8 intervals of 30 minutes (i.e., 4 hour ahead)

A look-back window of  $k = 12$  intervals (i.e., 6 hours) is selected to capture past dynamics



Legend: — real value  
 — quantiles [0.01, 0.99]  
 — quantiles [0.25, 0.75]




## Input features:

Past data $\mathbf{x}_{:t_0,n}^{(p)}$	past (local) load, past (aggregated) PV generation, past (aggregated) imports-exports, calendar information (weekday, hour of the day)
Known future data $\mathbf{x}_{t_1:T,n}^{(f)}$	forecasted (aggregated) imports-exports, calendar information (weekday, hour of the day)
Static data $x_n^{(s)}$	node location (feeder to which the node is connected, and distance to the origin node)

# A glimpse into the RNN model

The multi-horizon time-series forecasting problem can be naturally treated as a sequence-to-sequence task.

Both encoder and decoder blocks are modeled by recurrent neural networks due to their ability to represent inter-temporal dependencies.



Trained with quantile regression,  
for  $q = 1, 10, 25, 50, 75, 90$  and  $99\%$ .


→ the output corresponding to each  
time step is thus 7-dimensional.



# Utility privacy trade-off

There is an inherent trade-off between the performance of the prediction model and the privacy, i.e.,  $(\epsilon, \delta)$ -values:

- privacy breaches may occur at each round of the training
- privacy can be improved by reducing the expected number of end-users used at each round
- privacy can be controlled by the noise multiplier, which plays on the variance of the noise injected during the weight update



# Main conclusions

Privacy-preserving forecasting model:

- to distribute the computations among the parties
- to derive a tradeoff between utility and privacy by embedding the learning procedure into a differentially private mechanism.

Outcomes show that compact recurrent models are inherently more robust to noise, which makes them natural candidates for the development of privacy-enhancing techniques in renewable-dominated smart grids.

# Perspectives

- Tracking the privacy spent by each client, which is highly challenging since the set of clients participating in each round is private.
- Developing state-of-the-art attack frameworks to have an empirical evaluation of how much information an adversary can actually infer from trained models.
- Using split learning, which is more tailored to perform vertical learning
- Connect with data markets, for data pricing or valuation approaches

**Thank you for your attention**



[Follow ESIM on LinkedIn](#)