Privacy Preserving Demand Forecasting to Encourage Consumer Acceptance of Smart Energy Meters

Keele UNIVERSITY

Funded by UK Government

EUROPEAN UNION
European Regional Development Fund

Chris Briggs, Zhong Fan, Peter Andras

Tackling Climate Change with Machine Learning Workshop at NeurIPS 2020

Motivation

Climate Change

- Smart meter data contains an enormous amount of potential predictive power that will aid the transition away from fossil fuel technologies to cleaner and renewable technologies[1].

- Reliable forecasting will provide opportunity for more efficient optimisation of electricity grids to cope with varying energy demand and increasing contributions of renewables in the energy mix.

- Accurate forecasting is important here to understand how demand is evolving with consumer behaviour change (e.g. EV charging, electric heating and cooling).

UK smart meter rollout

- Smart Energy GB recently switched its advertising message
  - from: Helping you save energy in the home
  - to: Doing your bit to help upgrade the UK's energy system
- The focus here is to understand the demands on the energy system ahead of time to reduce wasted energy and plan energy infrastructure in the future

Motivation

Data Privacy

- Smart meter installation in most countries is an opt-in process and levels of adoption of smart meters remains low.
- Data privacy and security concerns are among the most cited reasons consumers give for rejecting a smart meter installation[2].
- High-resolution smart meter data is particularly sensitive as it can easily enable inference about household occupancy, lifestyle habits or even what and when specific appliances are being used in a household.
- The privacy concerns include government surveillance, energy suppliers selling data and illegal data acquisition/use.

Privacy issues with smart meters

What data do smart meters collect and transmit?

- Aggregate electricity consumption for a household:
  - How much electricity was consumed in a given time interval
  - What time it was consumed
  - + an identifier for the household (which can be linked to a customer account)

- Collected every 5, 15, 30 or 60 minutes depending on the type of smart meter. Up to 288 readings per day (vs ~1 meter reading per few months using an analog meter)

Privacy issues with smart meters

What can be inferred from this data?

- How many people occupy the household at particular times of day
- What times of day/week is the household unoccupied
- When specific appliances are being used and for how long
- Whether someone prefers to cook using an oven or a microwave
- Whether and how often exercise equipment is used (possibly indicating lifestyle changes)
- Whether an in-home alarm system was installed and in use
- Whether lighting is switched on at ”odd-hours” in the night (possibly indicating illness)

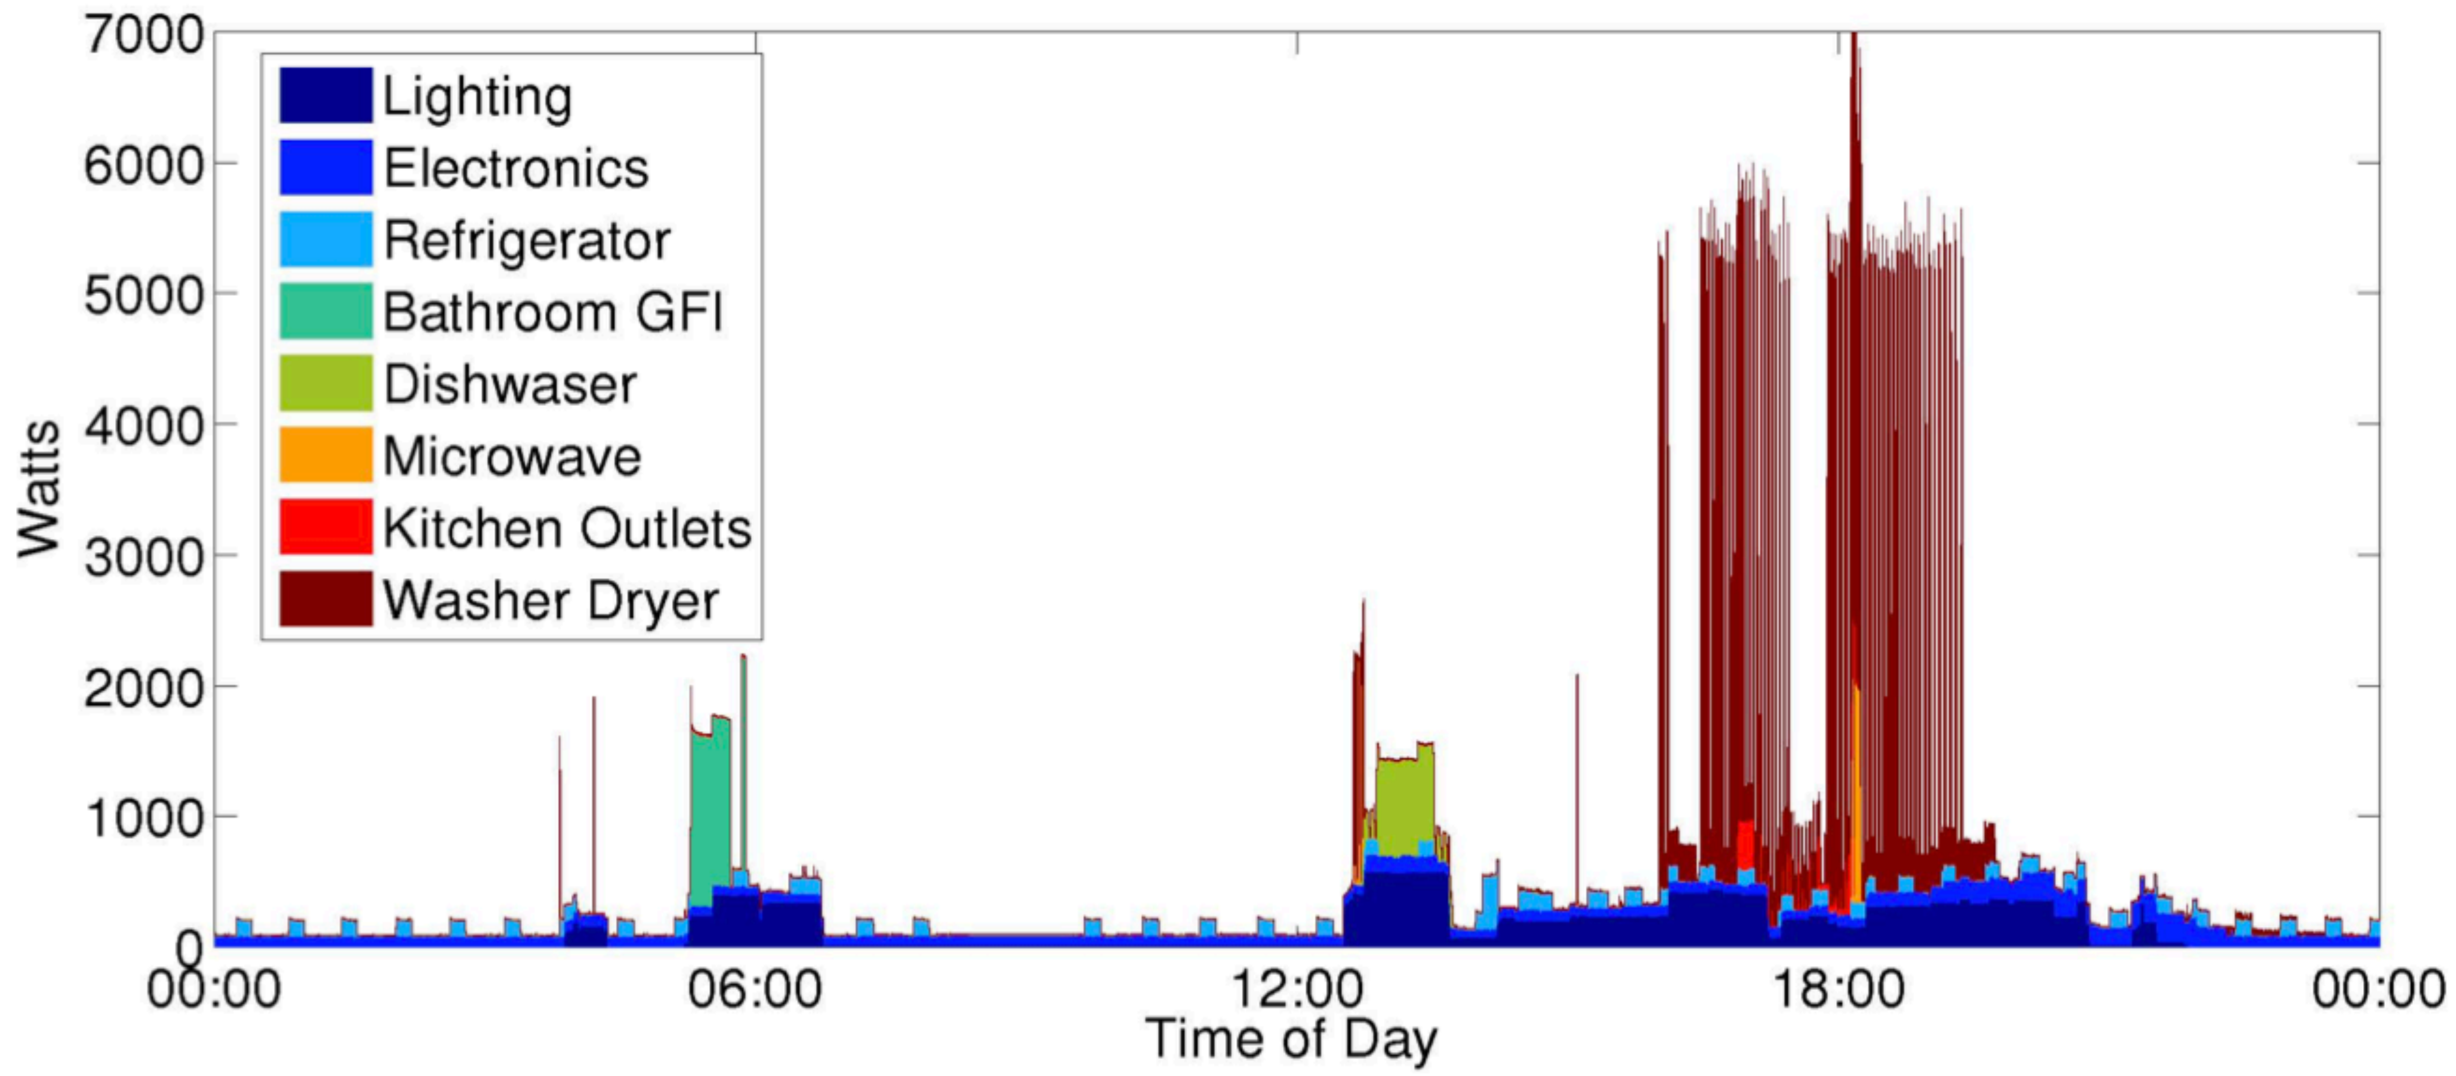

Appliance energy profiles inferred from aggregated smart meter data collected every 5 minutes

Proposed solution:

Federated Learning + Energy Demand Forecasting

Federated learning

What is it?

- Train a machine learning model on distributed data
- Users maintain control/privacy of their data
- Users contribute what their local model learns to a global learning objective
- Assumptions
  - Many participating users
  - Varying data distribution among users = challenging

Federated learning

How does it work?

Federated learning
How does it work?

Server

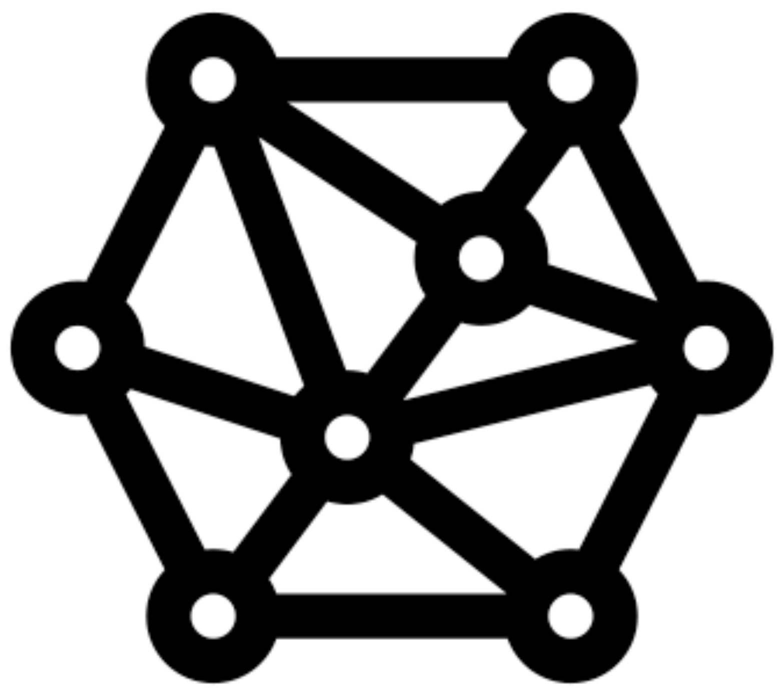

Global model

Clients

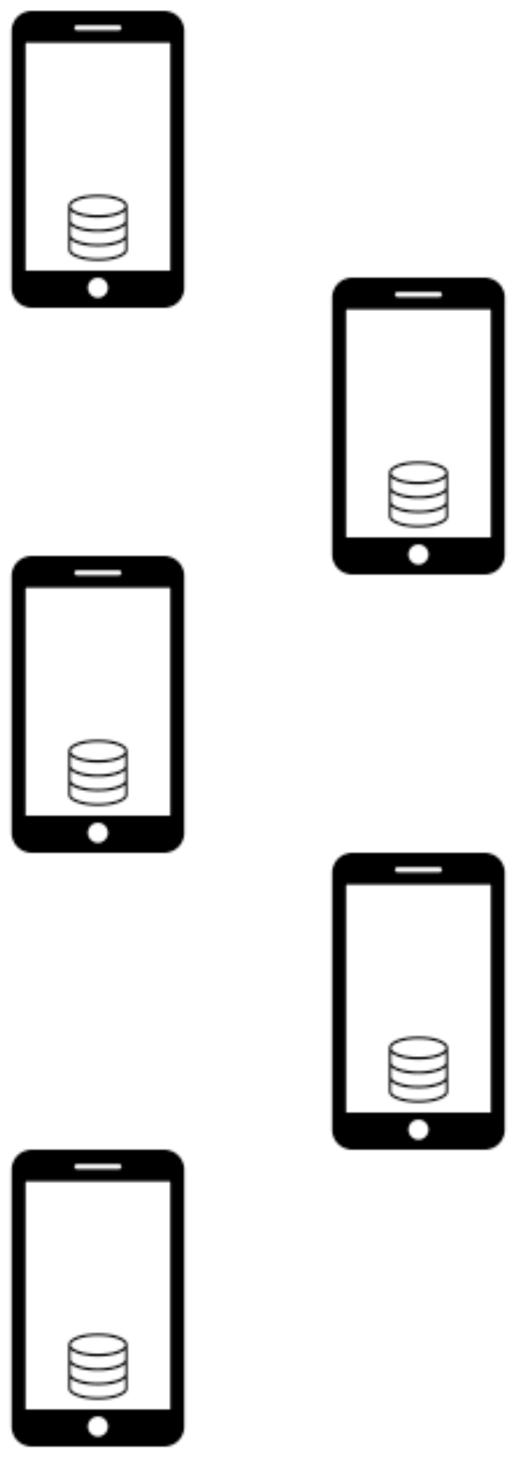

Federated learning

How does it work?

Federated learning

How does it work?

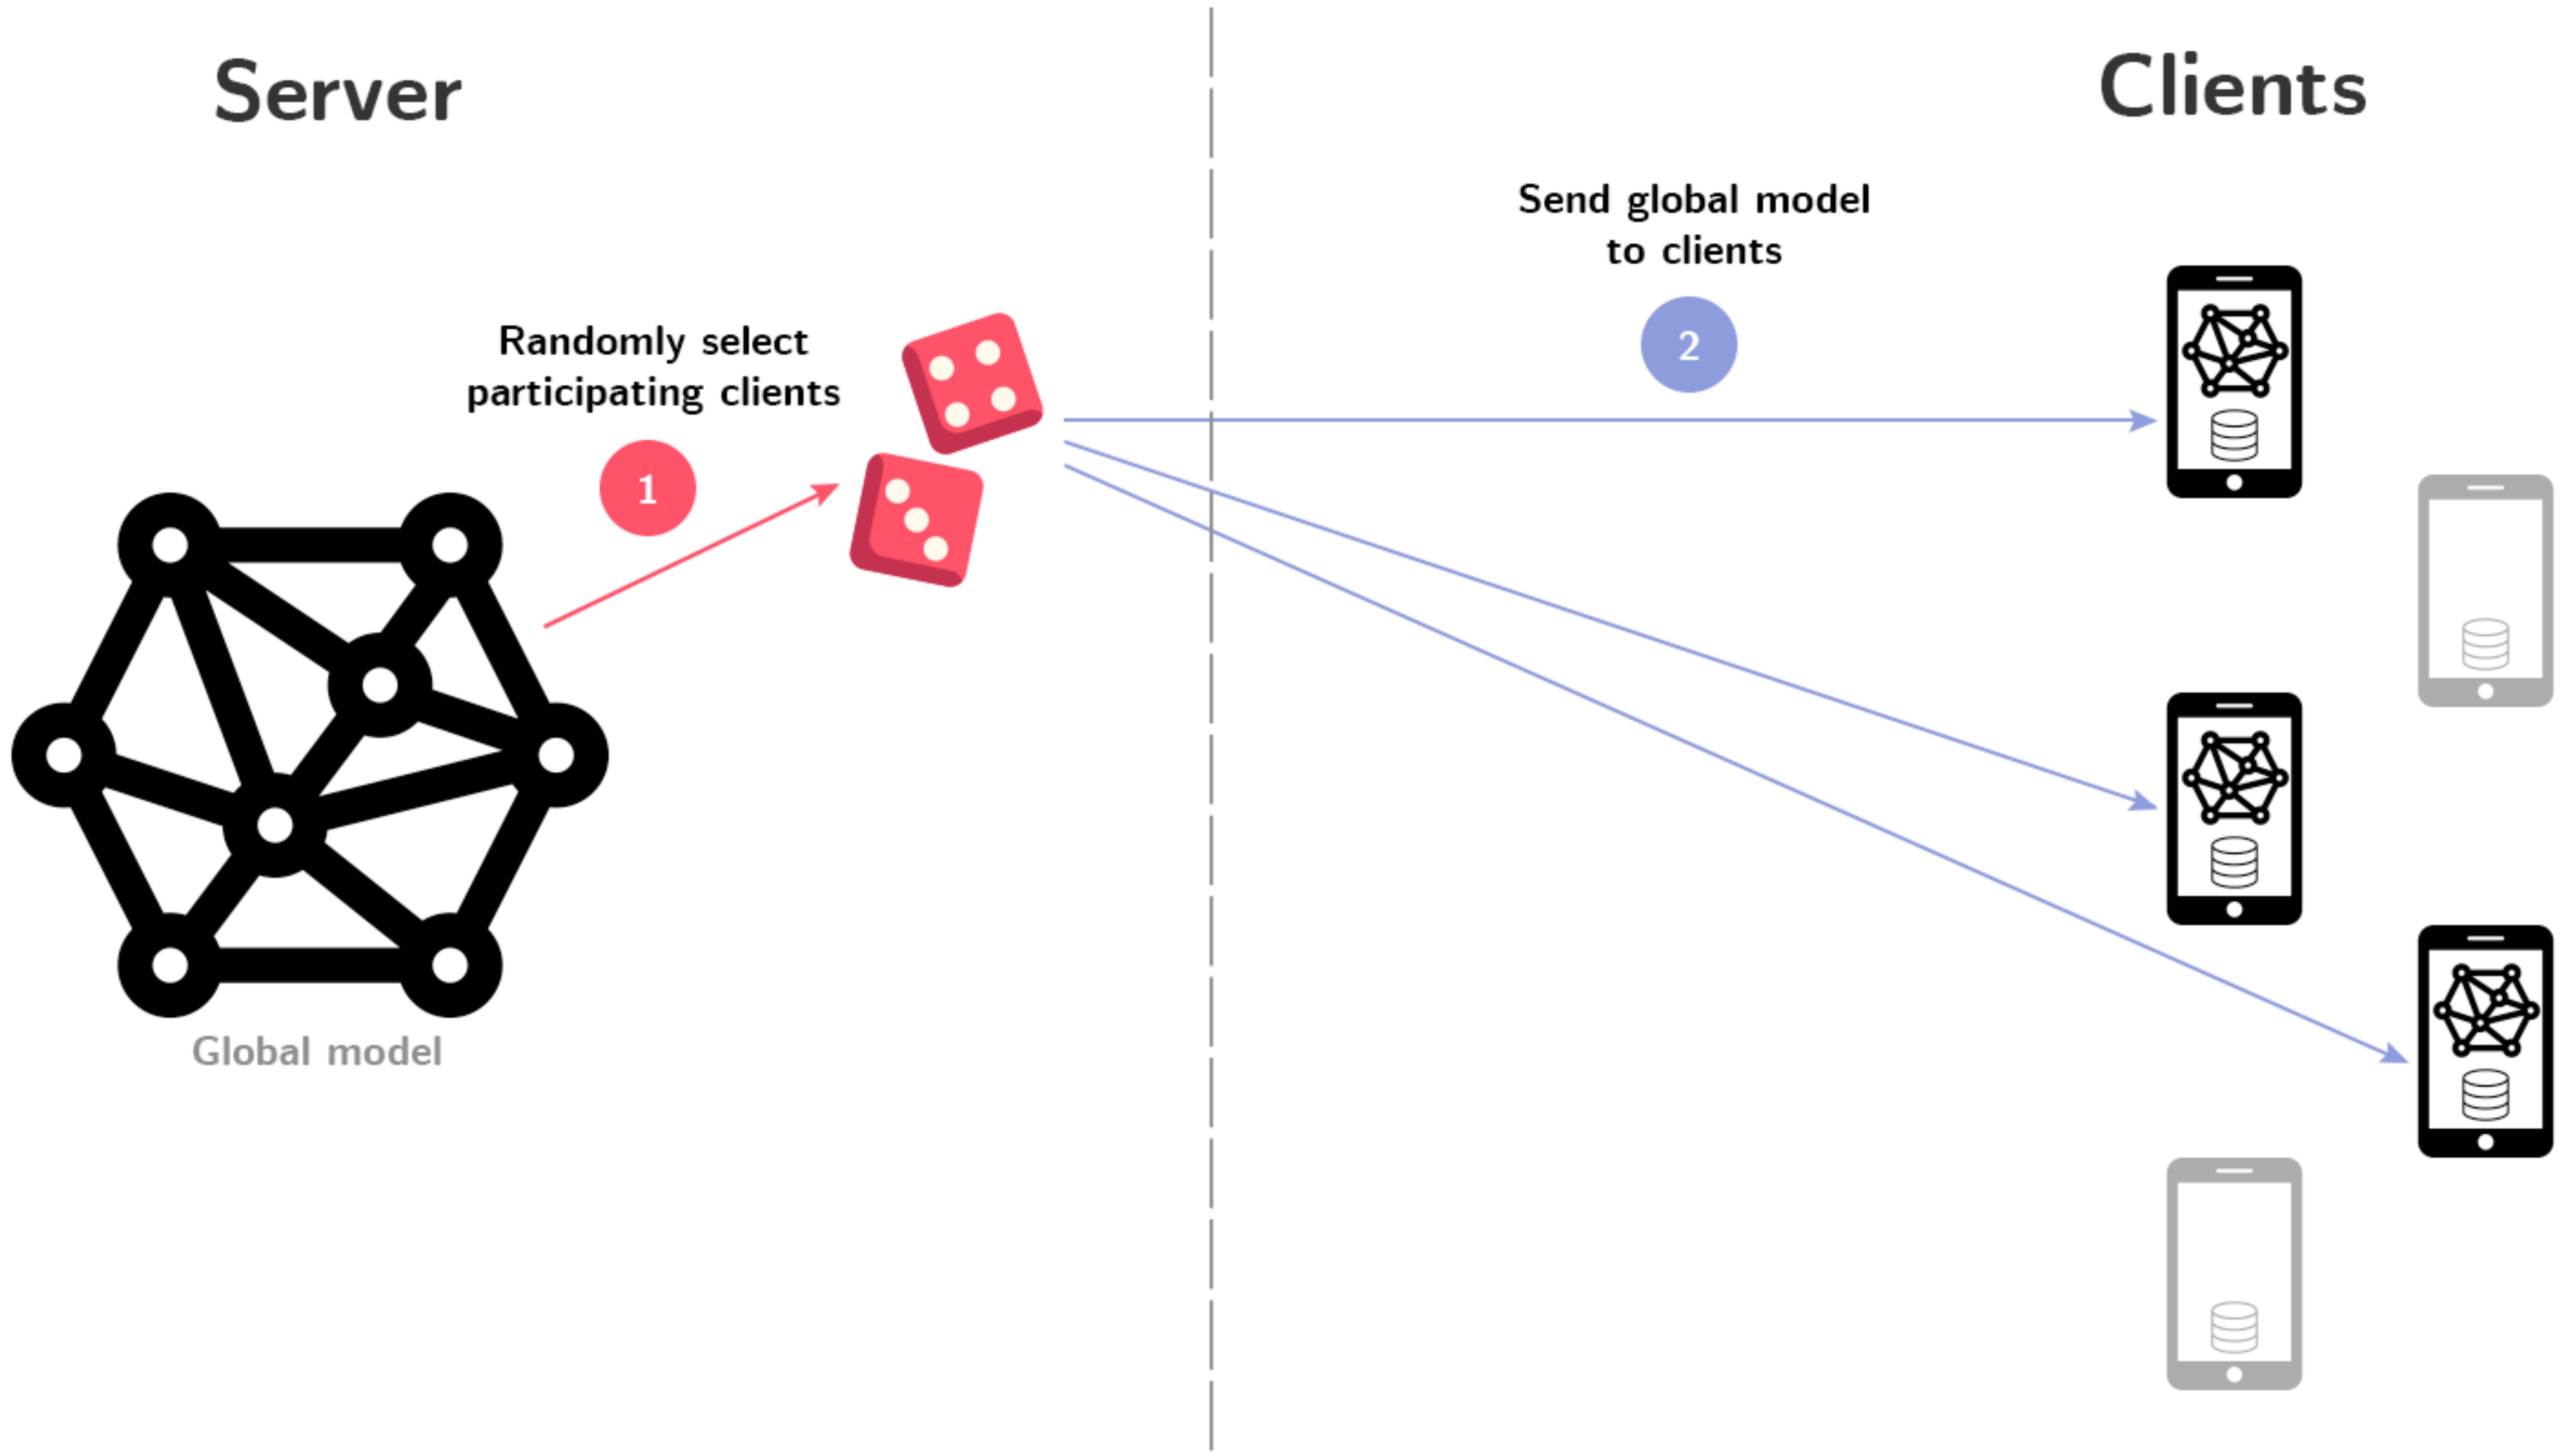

Federated learning

How does it work?

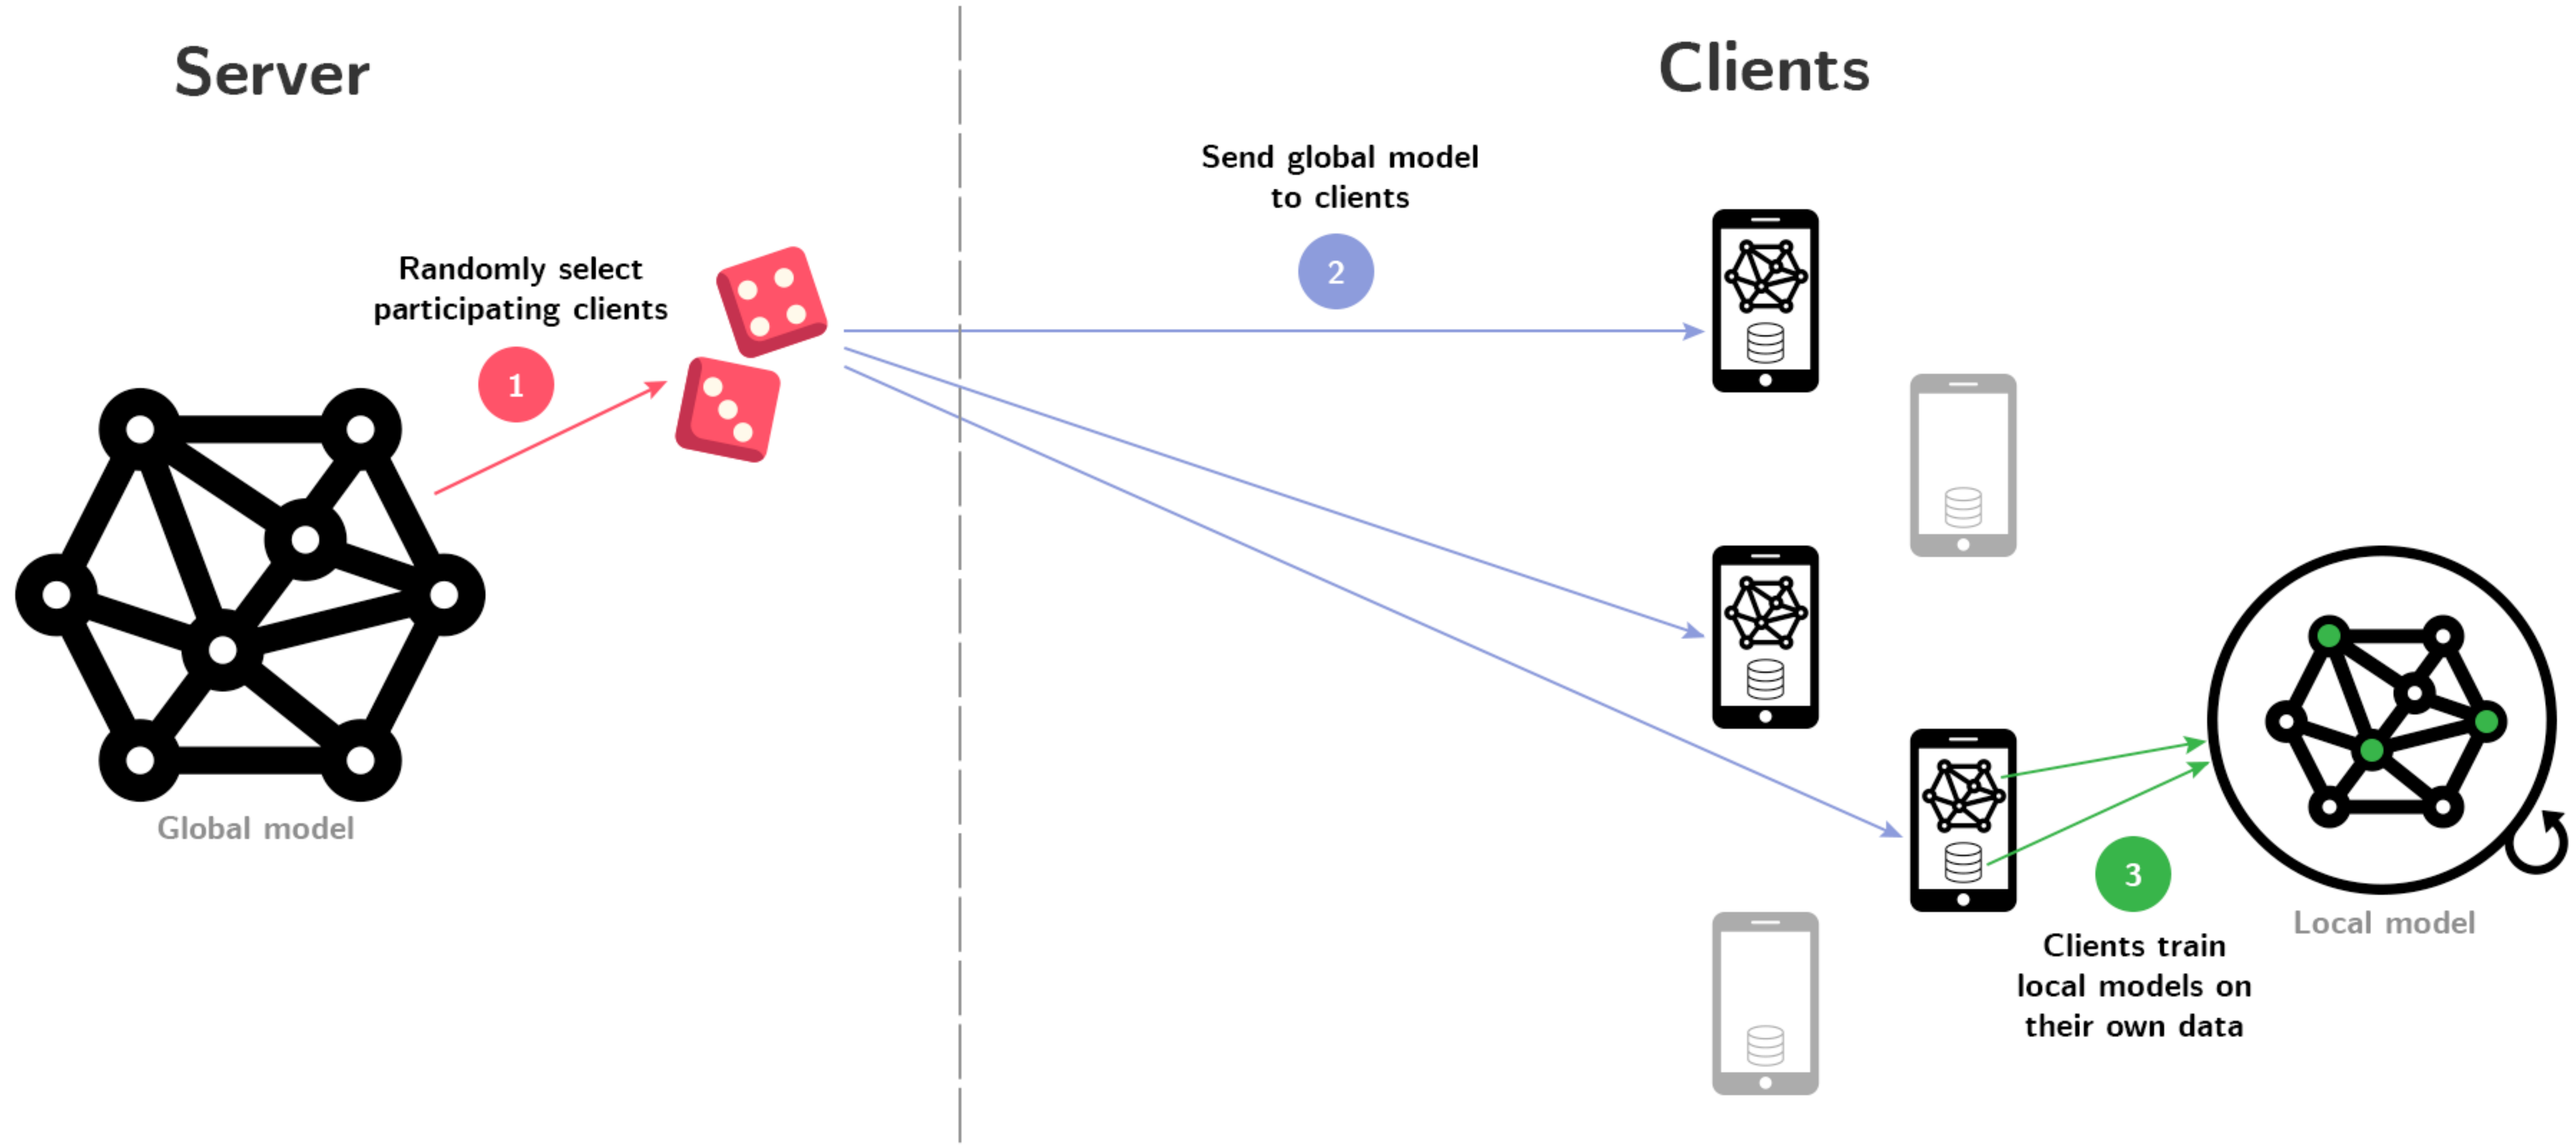

Federated learning

How does it work?

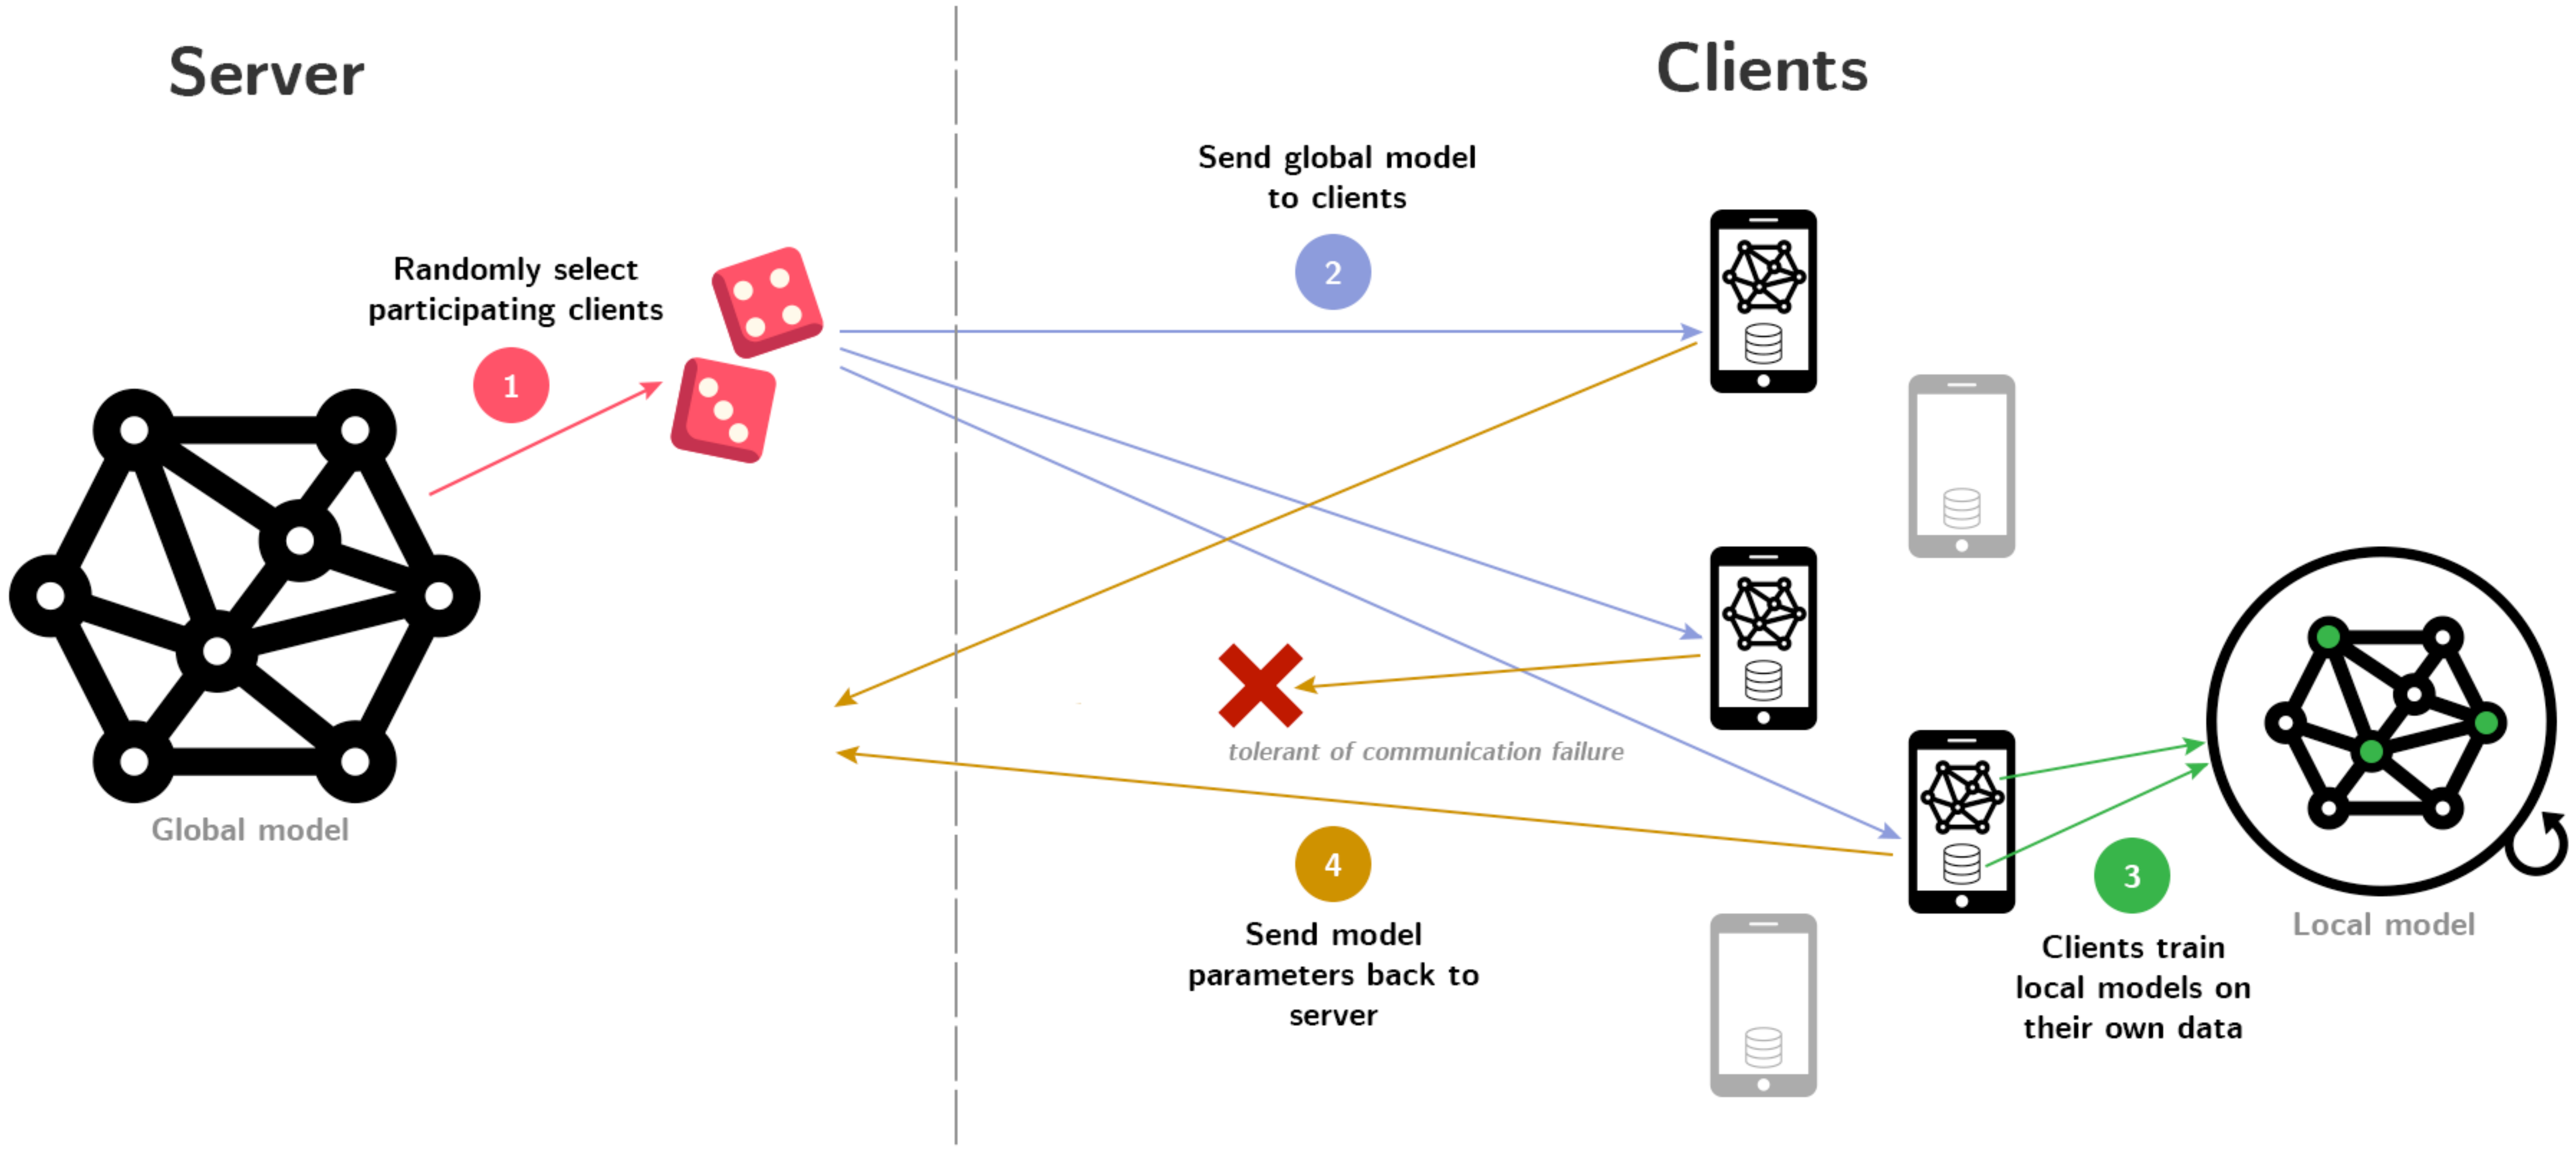

Federated learning

How does it work?

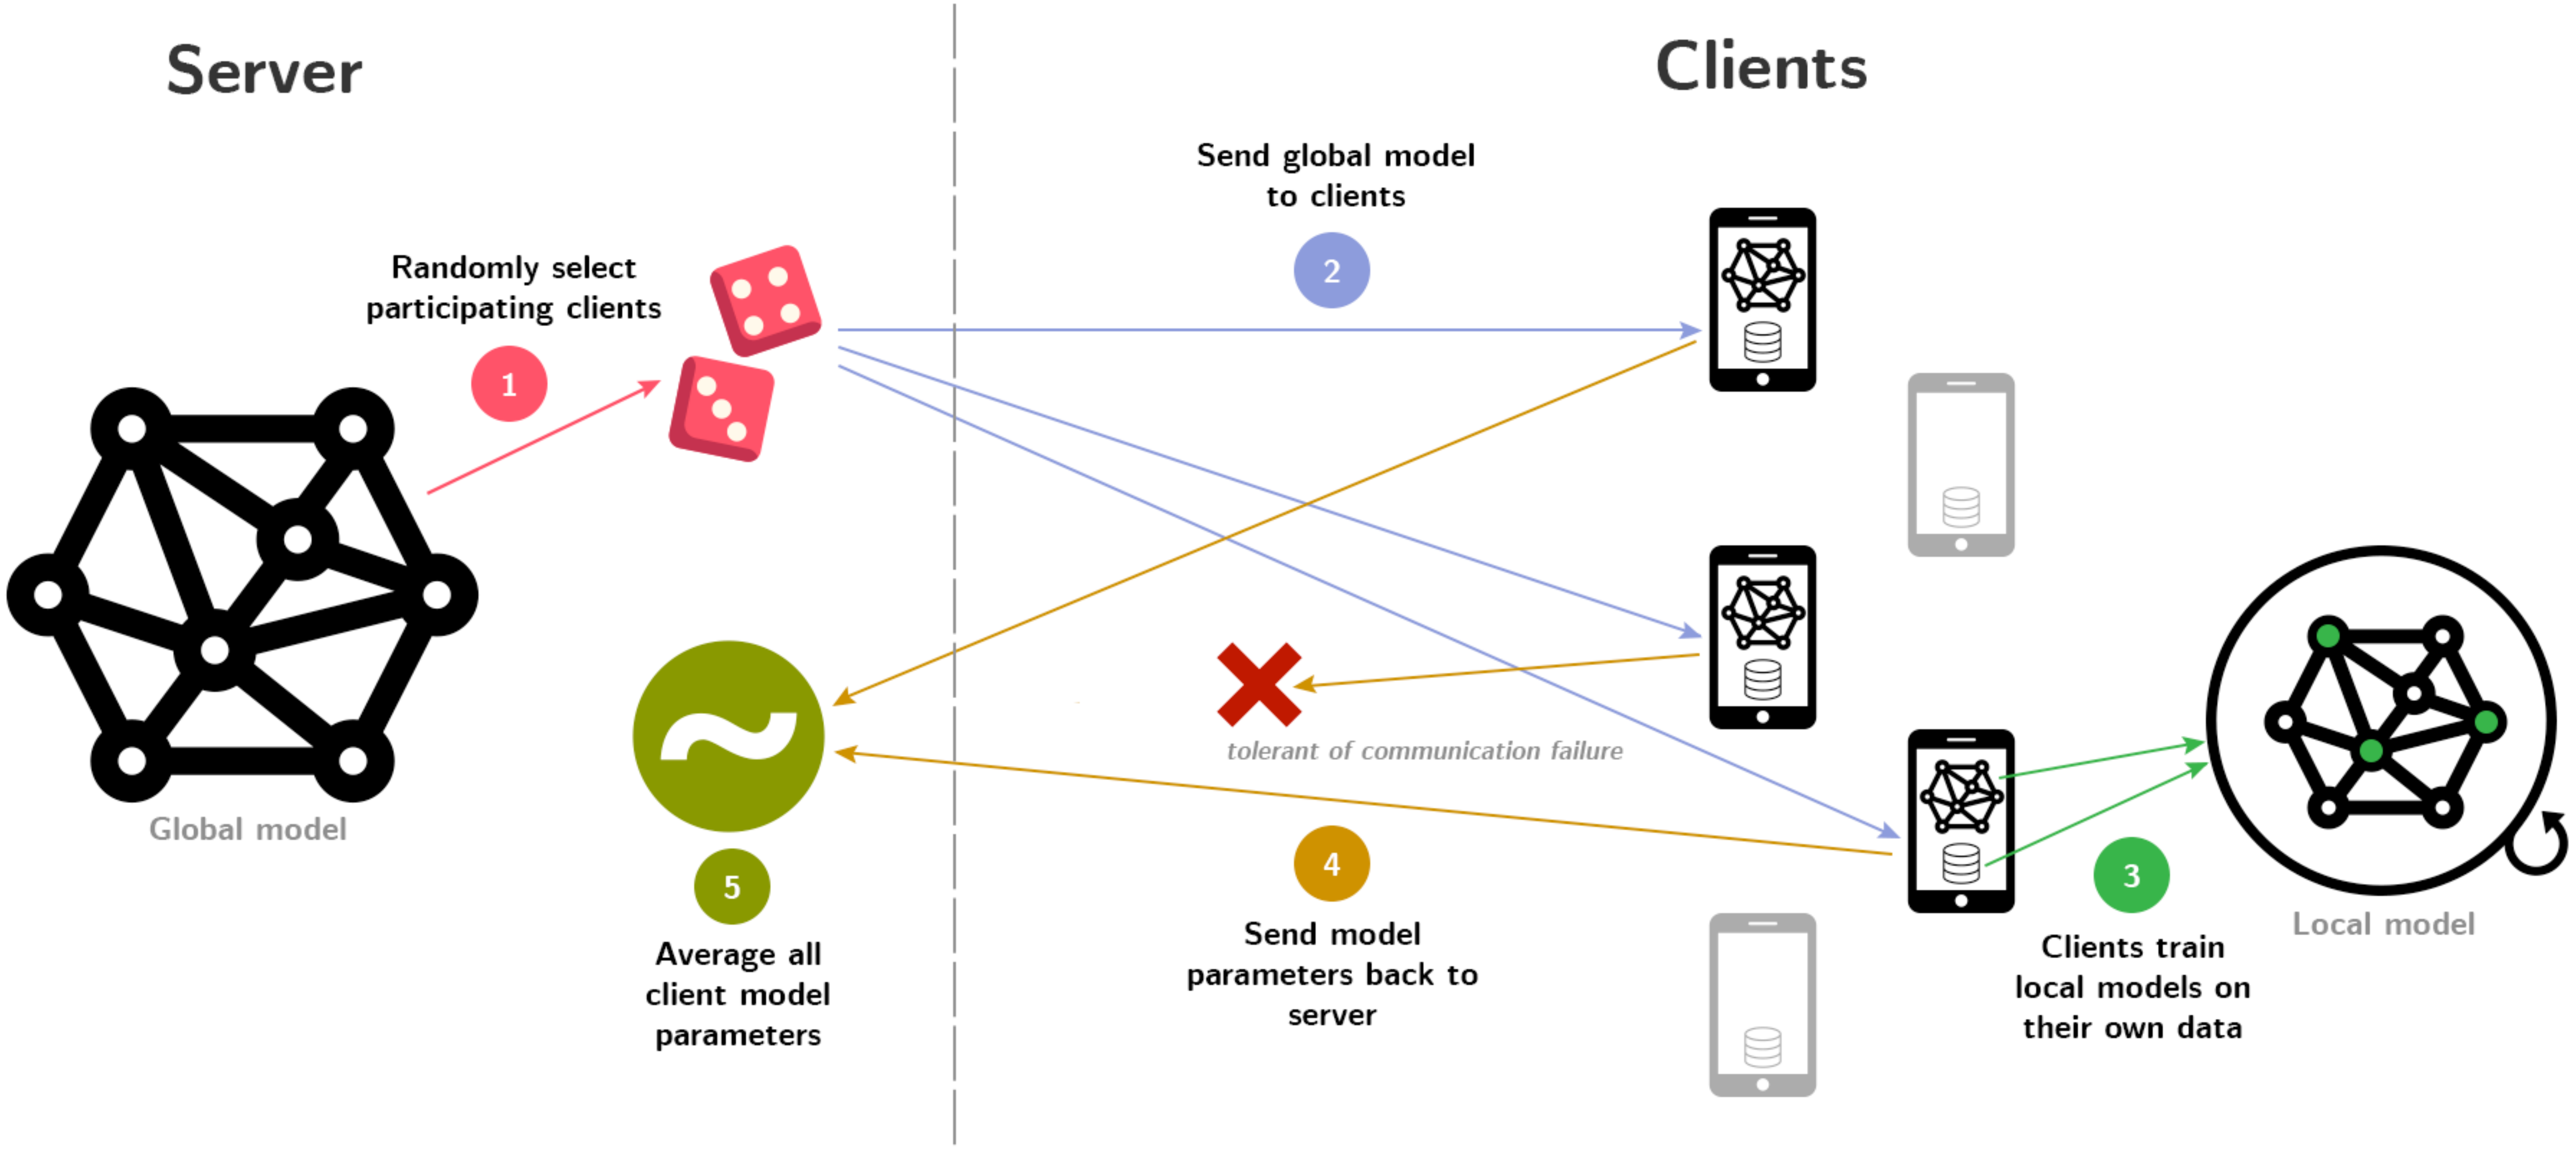

Federated learning

How does it work?

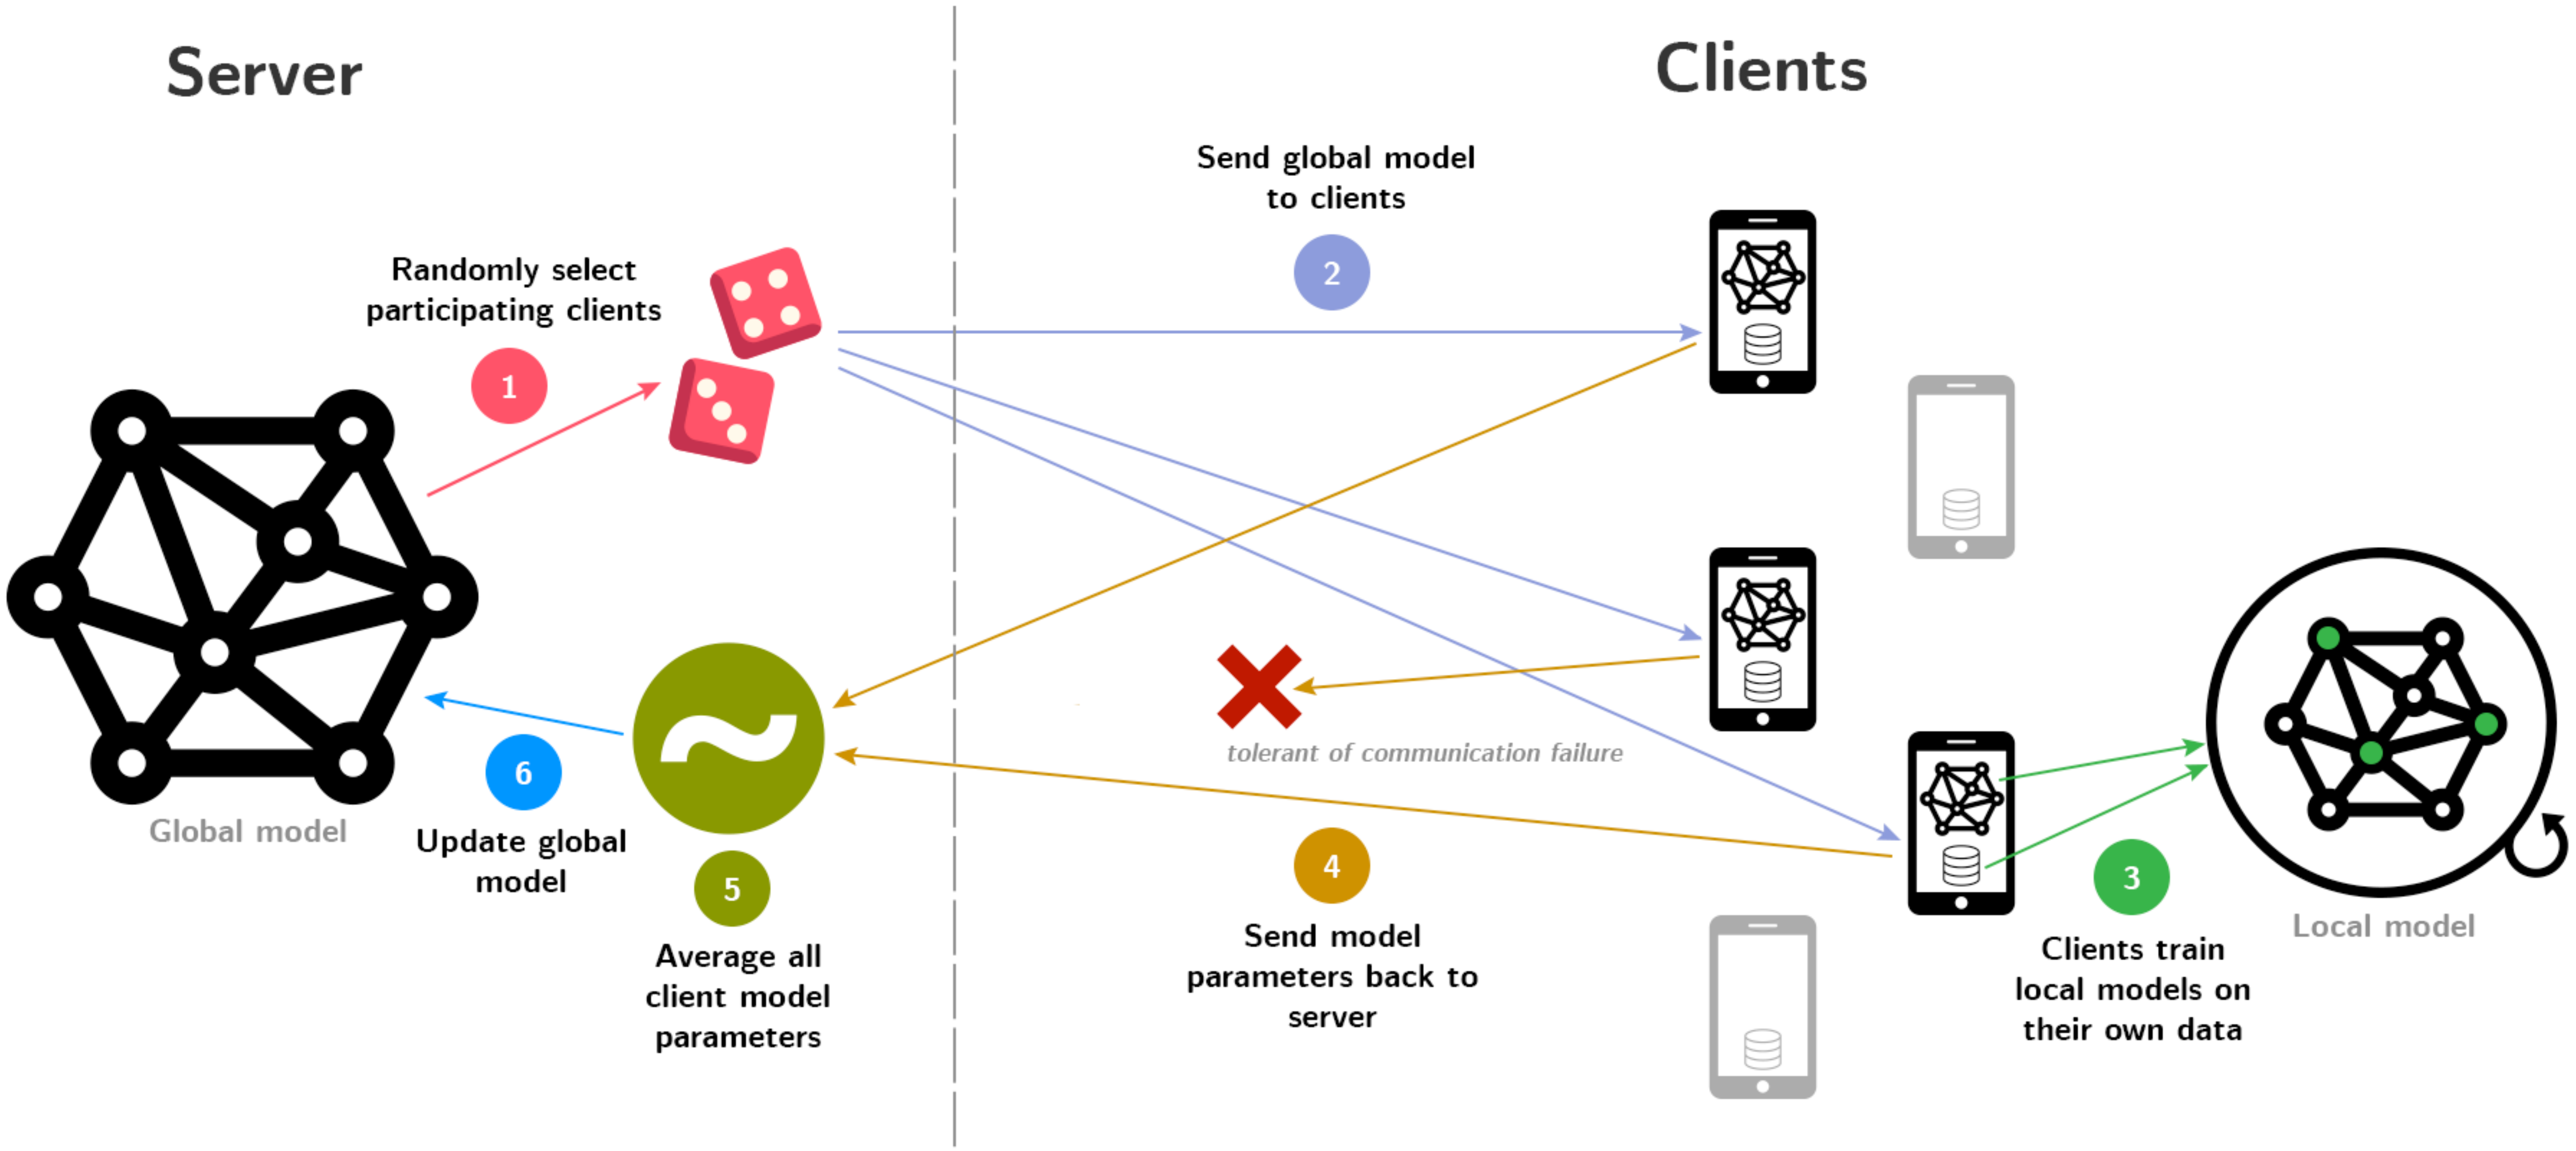

Federated learning

How does it work?

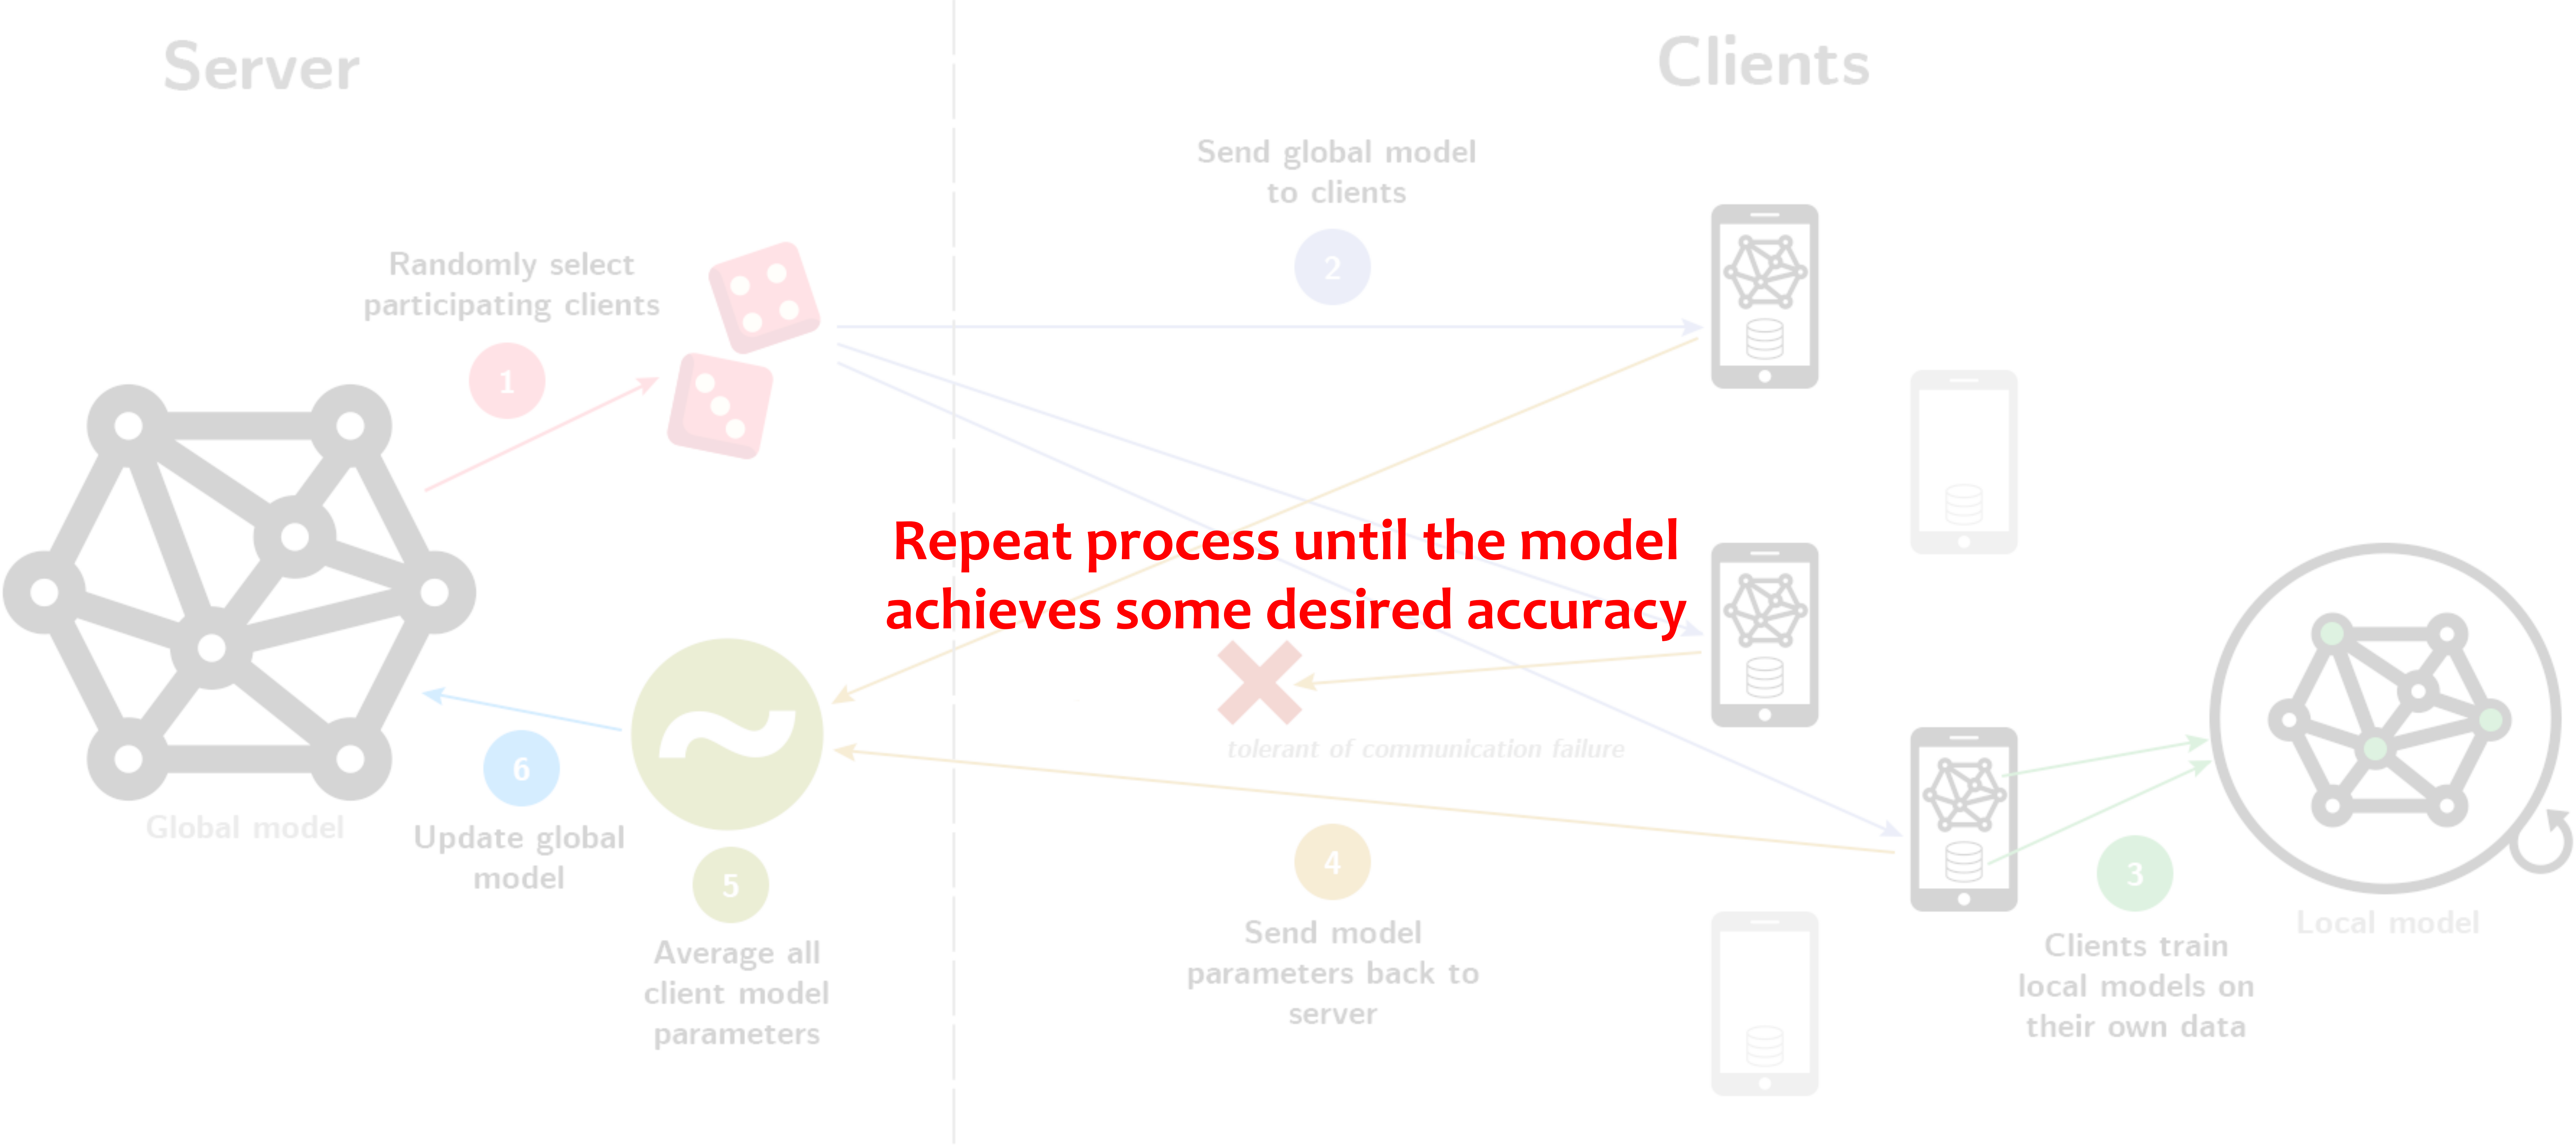

How does federated learning preserve privacy

- Private data is not collected/stored by a third party
  - The data never leaves the user's device
- Private data is not used for any other machine learning tasks that the user didn't agree to
- After a model is trained, new private data doesn't need to be sent to a third party to provide predictions
  - The final fully trained model is available to all users locally at the end of training

Energy demand dataset

- Low Carbon London project, led by UK Power Networks
- 5,567 smart meters installed in residential properties in London, UK
- Long baseline of energy consumption readings (2011 - 2014)
- Combined with temperature and humidity data from the Met Office UK
  - Highly correlated with energy demand - will help with predictive forecasting

Proposal of a Federated Learning Energy Demand Forecasting Application

- Privacy, data protection & security are the most cited concerns by users wary of upgrading to a smart meter.
- Federated learning could help alleviate these concerns - raw smart meter data is never shared.
- Demand prediction application:
    - Train a time-series/sequence model (e.g LSTM) using federated learning.
    - Compare with clustered version of FL to optimise for groups of households who use energy in a similar way.
    - Aggregate predictions using a privacy preserving method (further protecting individual user predictions)

Privacy Preserving Demand Forecasting to Encourage Consumer Acceptance of Smart Energy Meters

Christopher Briggs
School of Computing
& Mathematics
Keele University
Keele, Staffordshire, UK
c.briggs@keele.ac.uk

Zhong Fan
School of Computing
& Mathematics
Keele University
Keele, Staffordshire, UK
z.fan@keele.ac.uk

Peter Andras
School of Computing
& Mathematics
Keele University
Keele, Staffordshire, UK
p.andras@keele.ac.uk

Abstract

In this proposal paper we highlight the need for privacy preserving energy demand forecasting to allay a major concern consumers have about smart meter installations. High resolution smart meter data can expose many private aspects of a consumer's household such as occupancy, habits and individual appliance usage. Yet smart metering infrastructure has the potential to vastly reduce carbon emissions from the energy sector through improved operating efficiencies. We propose the application of a distributed machine learning setting known as federated learning for energy demand forecasting at various scales to make load prediction possible whilst retaining the privacy of consumers' raw energy consumption data.

1 Motivation & Impact

Smart meters are being deployed in many countries across the world for the purpose of optimising efficiency within electricity grids and giving consumers insight into their energy usage. The meters record energy consumption within a building directly from the electricity supply and periodically communicate this data to energy suppliers and other entities in the energy sector. Smart meter data contain an enormous amount of potential predictive power that will aid the transition from fossil fuel technologies to cleaner and renewable technologies [1]. However this high-resolution data is particular sensitive as it can easily enable inference about household occupancy, lifestyle habits or even what and when specific appliances are being used in a household.

A large contribution of renewables in the energy mix poses a significant challenge for balancing supply and demand. If peak demand coincides with low wind/solar inputs, energy must be provided by reliable backup generation, such as idling gas turbines. Such solutions are very costly, both economically and environmentally and serve to discourage the installation of large amounts of renewable energy production. Reliable forecasting will provide opportunity for more efficient optimisation of electricity grids to cope with varying energy demand.

With a rapidly evolving global climate and swinging extremes of temperature, large heating and cooling loads will become a more significant factor in energy demand profiles. Additionally, as other areas of our daily lives become decarbonised, loads from electric vehicles etc. will pose a strain on electricity grids. Accurate forecasting is important here to understand how demand is evolving with consumer behaviour change. Demand side response applications - those that actively or suggestively seek to reduce peak loads when necessary - can also benefit from accurate realtime forecasting.

Despite the benefits for promoting a greener energy sector, smart meter installation in most countries is an opt-in process and levels of adoption of smart meters remains low. Data privacy and security concerns are among the most cited reasons consumers give for rejecting a smart meter installation [2].

Preprint. Under review.

Thank you for listening

References

[1] Smart Energy GB, Smart meter benefits: Role of smart meters in responding to climate change. [Online]. Available: https://www.smartenergygb.org/en/-/media/SmartEnergy/essential-documents/press-resources/Documents/Smart-Energy-GB-report-2---Role-in-climate-change-mitigation-Final_updated-300819.ashx

[2] N. Balta-Ozkan, O. Amerighi, and B. Boteler, “A comparison of consumer perceptions to- wards smart homes in the UK, Germany and Italy: reflections for policy and future research,” Technology Analysis & Strategic Management, vol. 26, no. 10, pp. 1176–1195, Dec. 2014.

[3] UK Power Networks, SmartMeter Energy Consumption Data in London Households. [Online]. Available: https://data.london.gov.uk/dataset/smartmeter-energy-use-data-in-london-households

This work is partly supported by the SEND project (grant ref. 32R16P00706) funded by ERDF and BEIS. We would also like to thank Hildebrand Technology Limited for helpful discussions on the feasibility of deployment of a federated learning energy demand forecasting application.

Funded by UK Government

EUROPEAN UNION
European Regional Development Fund