Sequential learning for a sustainable electrical system

Rencontres MathTech

SORBONNE UNIVERSITÉ

LPSM
Laboratoire de probabilités statistique & modélisation

Margaux Brégère

IHES - January, 23rd 2025

EDF

As electricity is hard to store, balance between production and demand must be strictly maintained

Adapt production and Forecast demand

Optimization **Statistics**

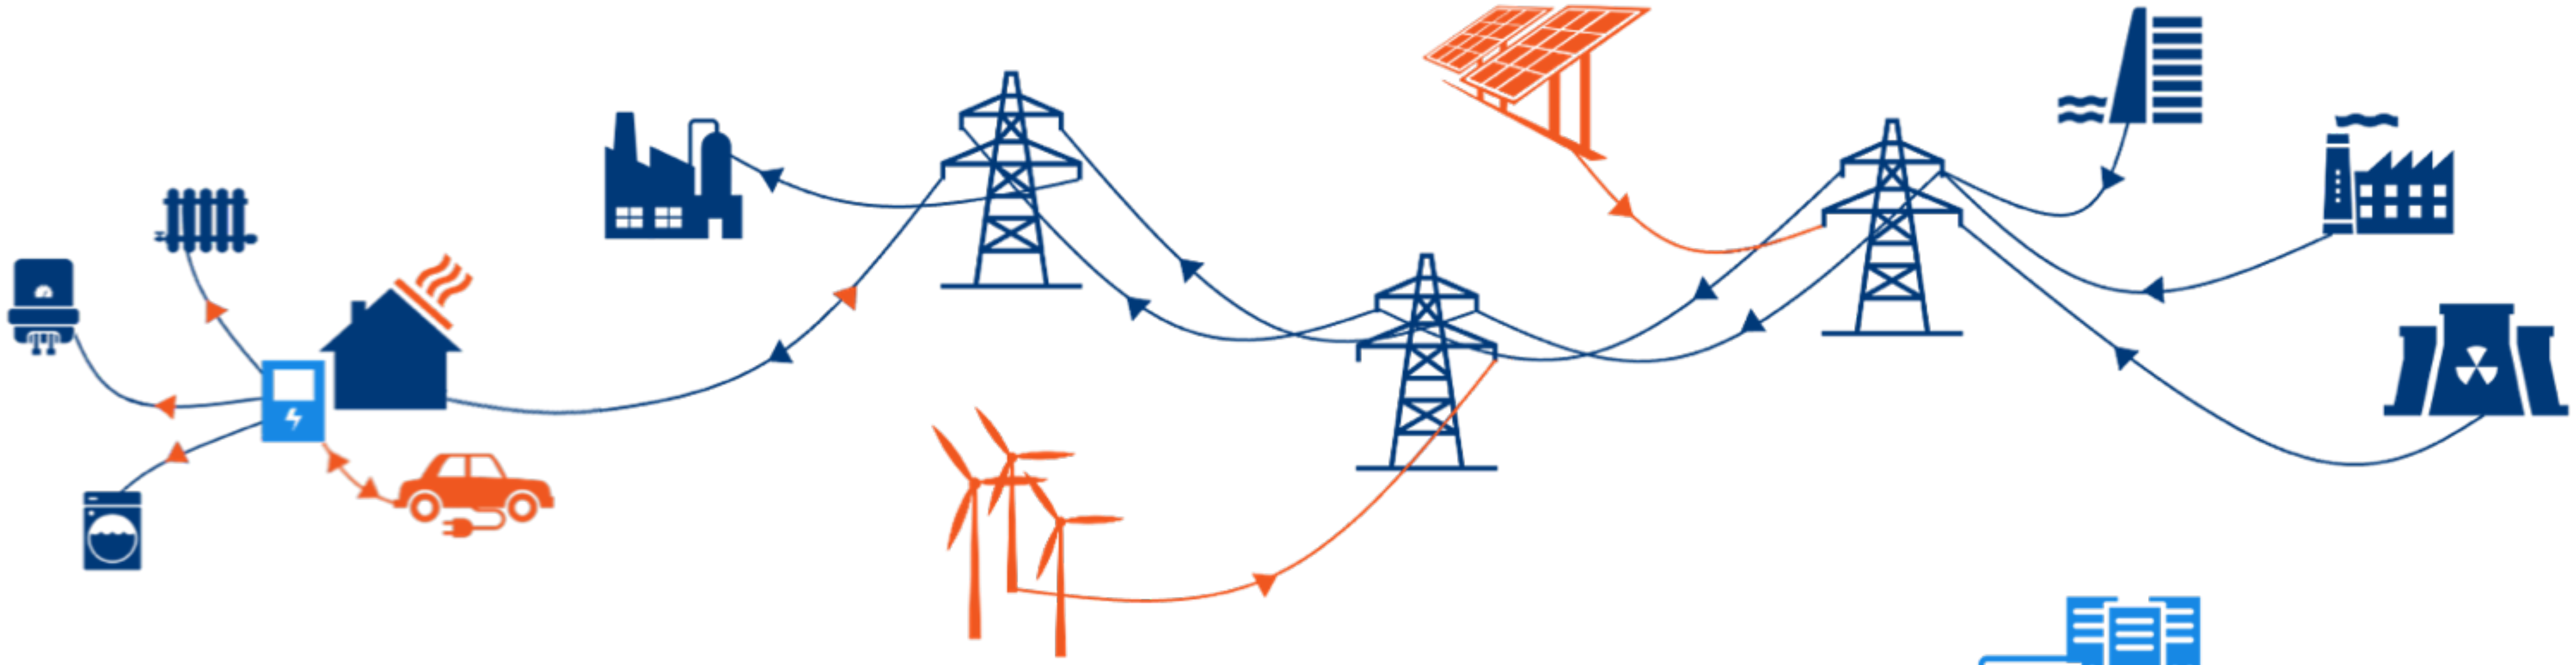

The energy and digital transition context

- New uses of electricity and electrification of numerous applications
- Massive development of intermittent renewables
- Increasingly rapid availability of data, smart meters and high-performance computing resources

Raises new challenges

- Changes in electricity demand (energy crisis, sobriety, self-consumption, electric vehicles, increase from the current 450 TWh to 645 TWh according to « Energy Futures 2050 »...)
- Need for electrical flexibilities (from 13 to 17 GW in 2050)
- Explosion of artificial intelligence (increasingly complex and costly models)

As electricity is hard to store, balance between production and demand must be strictly maintained

Forecast demand and renewables

and

Adapt production

Manage electrical flexibilities

1. Online learning for electrical system forecasting

→ reconciliation of regional forecasts

Forecast demand and renewables

and

Adapt production

Manage electrical flexibilities

2. Reinforcement learning for demand side management

→ algorithms for thermostatically controlled loads

3. Automated machine learning and explainability

→ application to electrical demand forecasting models

1. Online reconciliation of electricity demand forecasts

Malo Huard, Milvue

Raffaele Mattera, University of Campania Luigi Vanvitelli

[1] Online Hierarchical Forecasting for Power Consumption Data, Margaux Brégère and Malo Huard, International Journal of Forecasting, 2022, IIF-Tao Hong Award

[2] Spatio-temporal Clustering and Reconciliation for Regional Electricity Demand Forecasting, Margaux Brégère and Raffaele Mattera, Submitted, 2024

Motivation

- Forecasts needed at various aggregated levels
  - France: managing the overall balance and planning cross-border exchanges
  - Consumer type: designing offers
  - Regions: dispatching electricity at network junctions

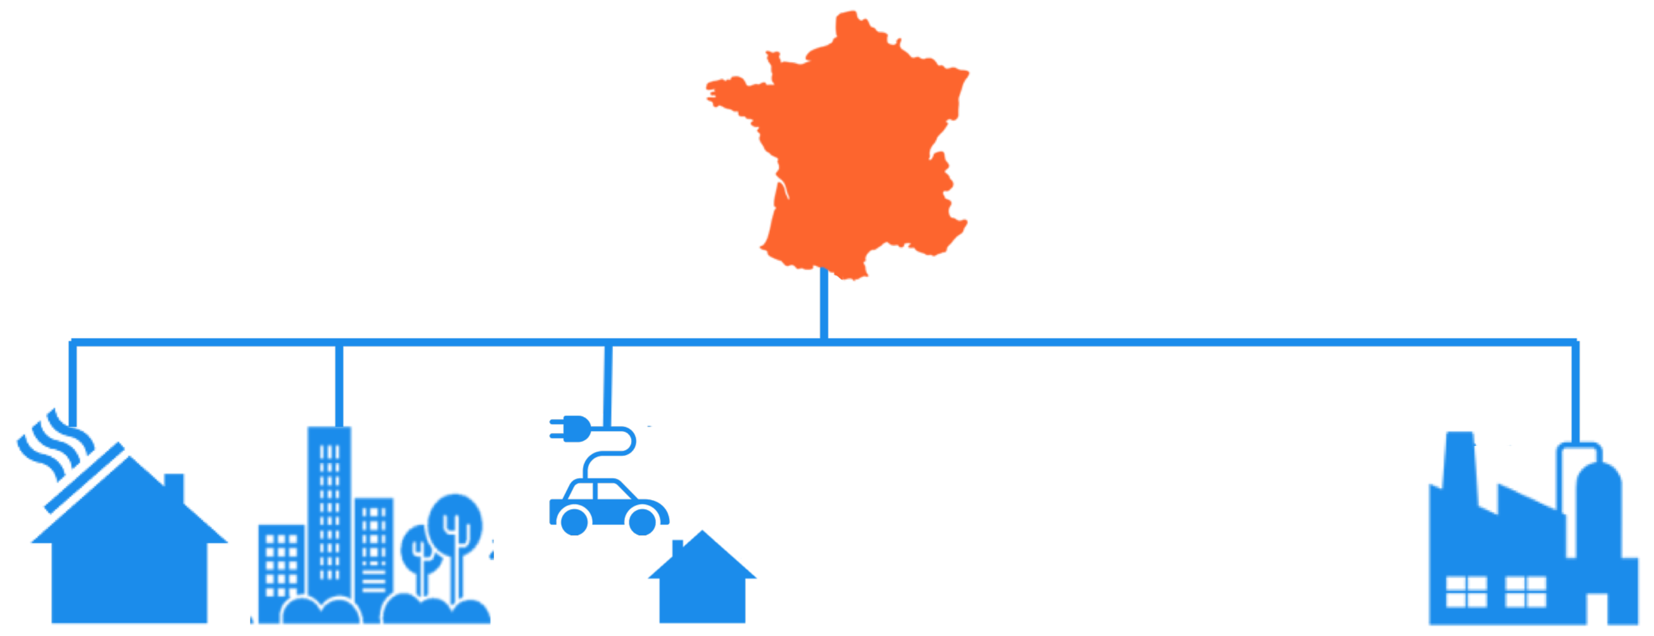

- Benchmark forecasts at each aggregated levels
  - France: easier to forecast (smoother)
  - Consumer type: same behavior
  - Regions: local weather

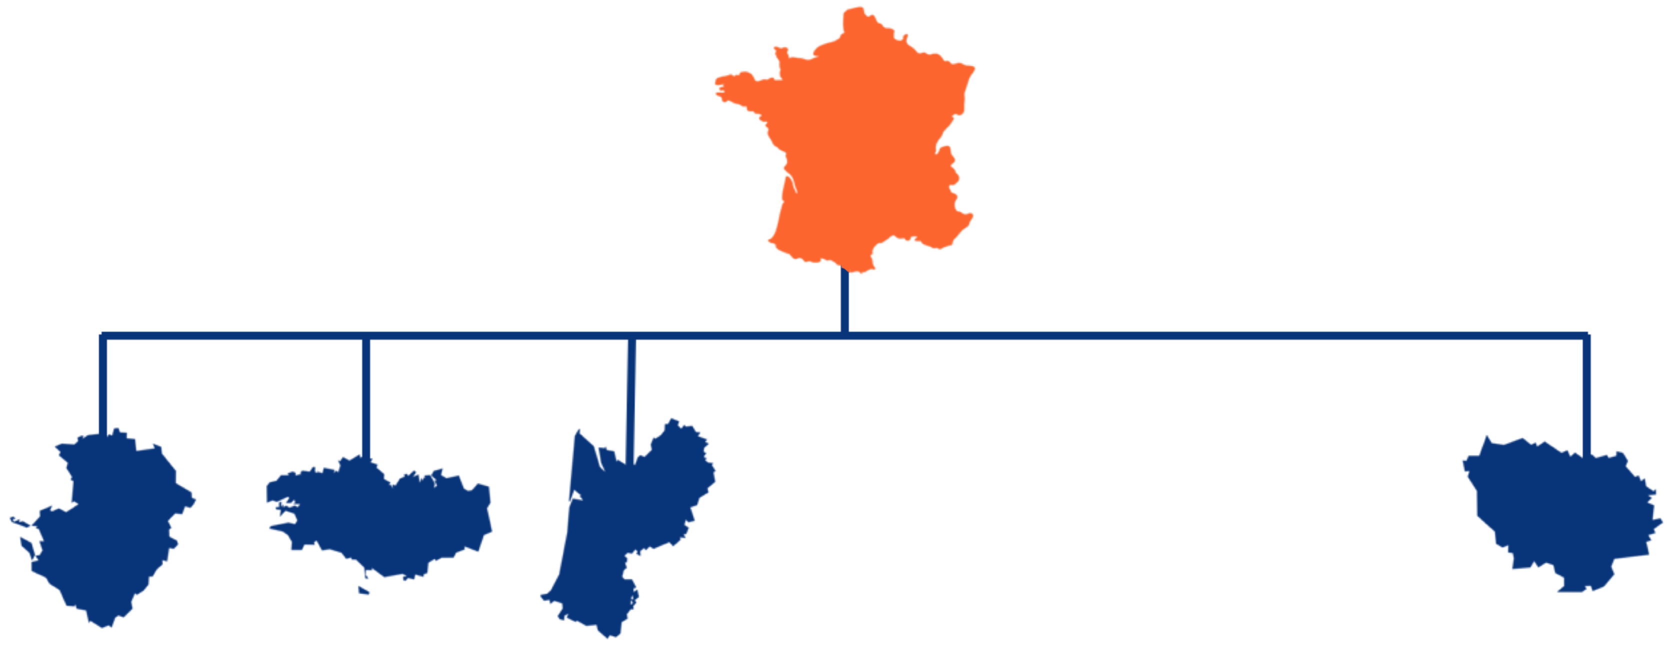

Correlated and connected times series through summation constraints

→ Reconciliation

Modeling

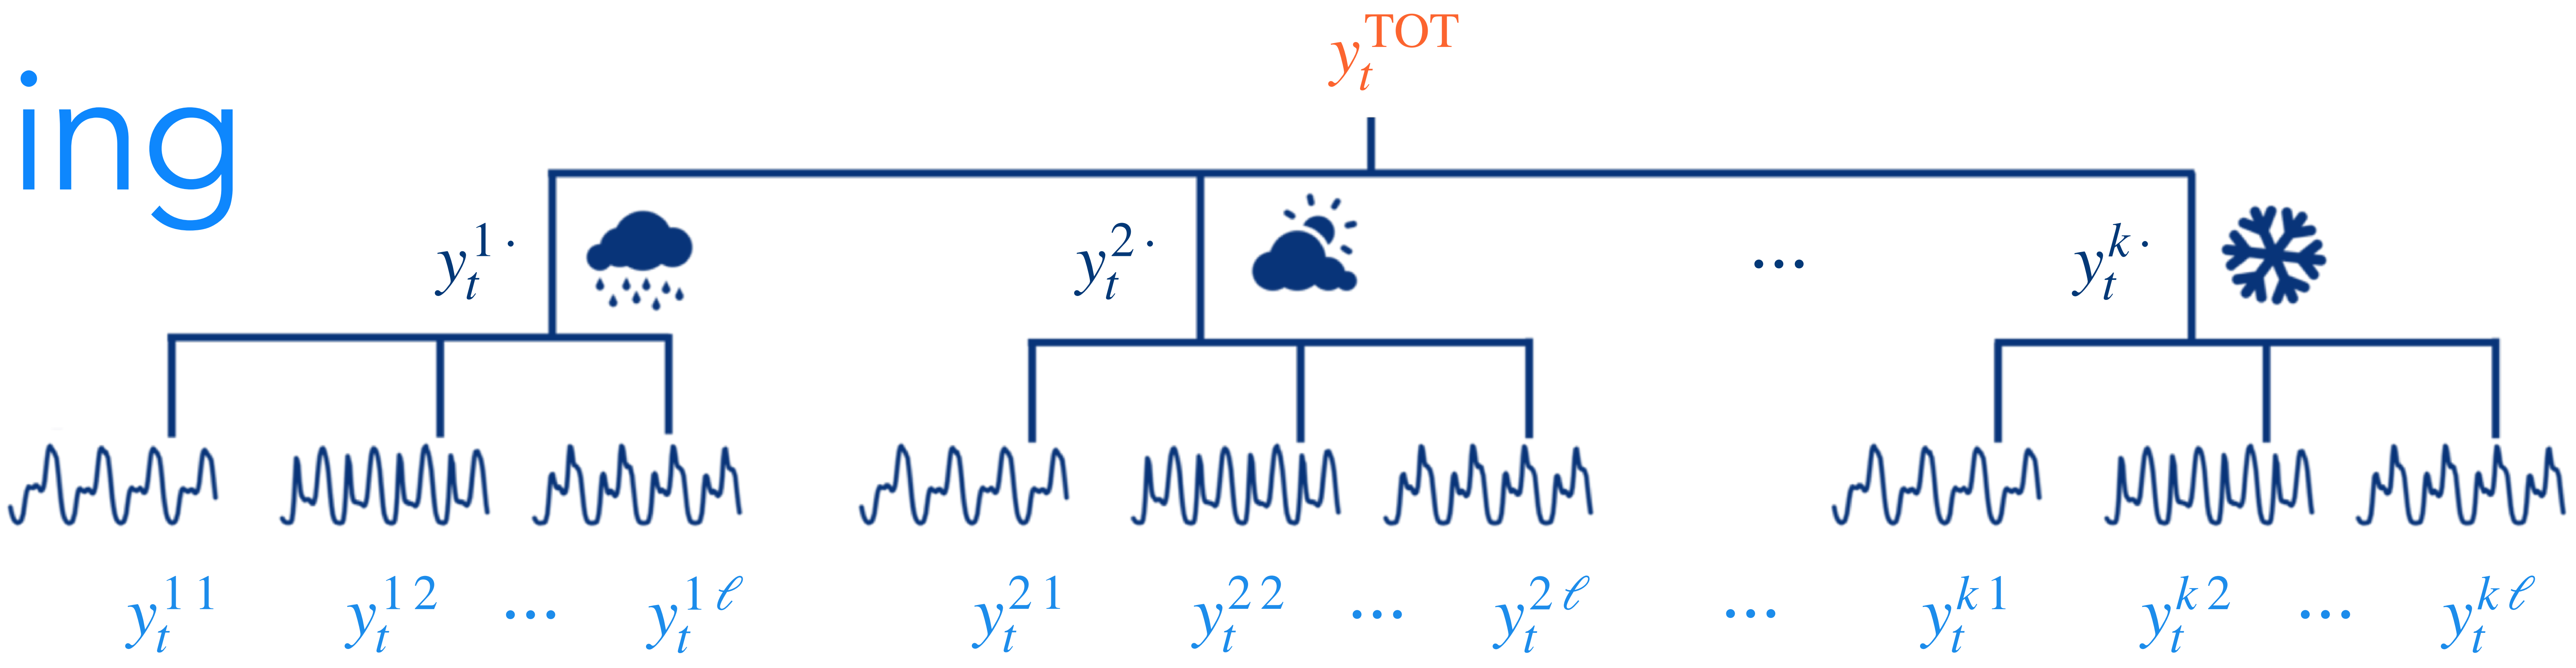

$$S = \begin{array}{c} \begin{array}{ccc} \overbrace{\qquad\qquad\qquad}^{\text{Region 1}} & \overbrace{\qquad\qquad\qquad}^{\text{Region 2}} & \qquad \overbrace{\qquad\qquad\qquad}^{\text{Region } k} \end{array} \\ \left(\begin{array}{cccccccccccccc} 1 & 1 & \dots & 1 & 1 & 1 & \dots & 1 & \dots\dots & 1 & 1 & \dots & 1 \\ 1 & 1 & \dots & 1 & 0 & 0 & \dots & 0 & \dots\dots & 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 & 1 & 1 & \dots & 1 & \dots\dots & 0 & 0 & \dots & 0 \\ & & & & & & \vdots & & & & & & \\ 0 & 0 & \dots & 0 & 0 & 0 & \dots & 0 & \dots\dots & 1 & 1 & \dots & 1 \\ & & & & & & I_n & & & & & & \end{array}\right) \end{array} \begin{array}{l} \}\ \text{All} \\ \\ \left.\begin{array}{l} \\ \\ \\ \end{array}\right\}\ \text{Regions} \\ \}\ \text{clusters x regions} \end{array}$$

$$y_t = \begin{bmatrix} y_t^{TOT} \\ y_t^{1\cdot} \\ \vdots \\ y_t^{k\cdot} \\ y_t^{11} \\ y_t^{12} \\ \vdots \\ y_t^{21} \\ \vdots \\ y_t^{k\ell} \end{bmatrix}$$

With the most disaggregated level time series $b_t = \left(y_t^{11}, y_t^{12}, \dots, y_t^{1\ell}, y_t^{21}, \dots, y_t^{k\ell}\right)^T$, $y_t = Sb_t$

From base forecasts $\hat{y}_t$ to reconciled ones: $\tilde{y}_t = S\tilde{b}_t$

All linear methods can be written as $\tilde{y}_t = SP\hat{y}_t$:

- P projects base forecasts into bottom level disaggregated forecasts: $\tilde{b}_t = P\hat{y}_t$
- S sums them: $\tilde{y}_t = S\tilde{b}_t$

Minimum trace reconciliation (MinT - Wickramasuriya et al. 2019):

assuming base forecast errors are stationary conditionally to data observed, the optimal reconciliation (which minimizes the variance of the reconciled forecast errors) is obtained for

$$P = \left(S^{\mathrm{T}}\Sigma^{\dagger}S\right)^{-1}S^{\mathrm{T}}\Sigma^{\dagger} \text{ with } \Sigma = \mathbb{E}\left[\left(y_t - \hat{y}_t\right)\left(y_t - \hat{y}_t\right)^{\mathrm{T}}\right]$$

Remarks:

- Stationarity implies unbiased base forecast
- Reconciled forecasts will also be unbiased $\Leftrightarrow SPS = S$
- Challenge: estimating the variance-covariance matrix of the base forecast errors Σ

Application: French electricity consumption

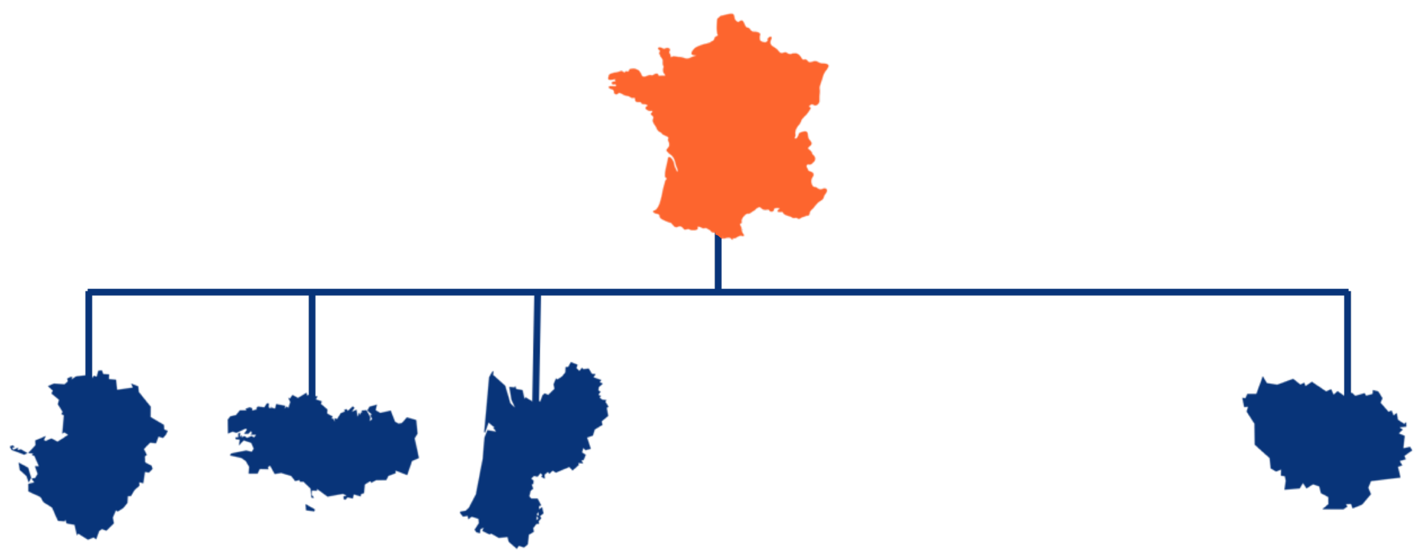

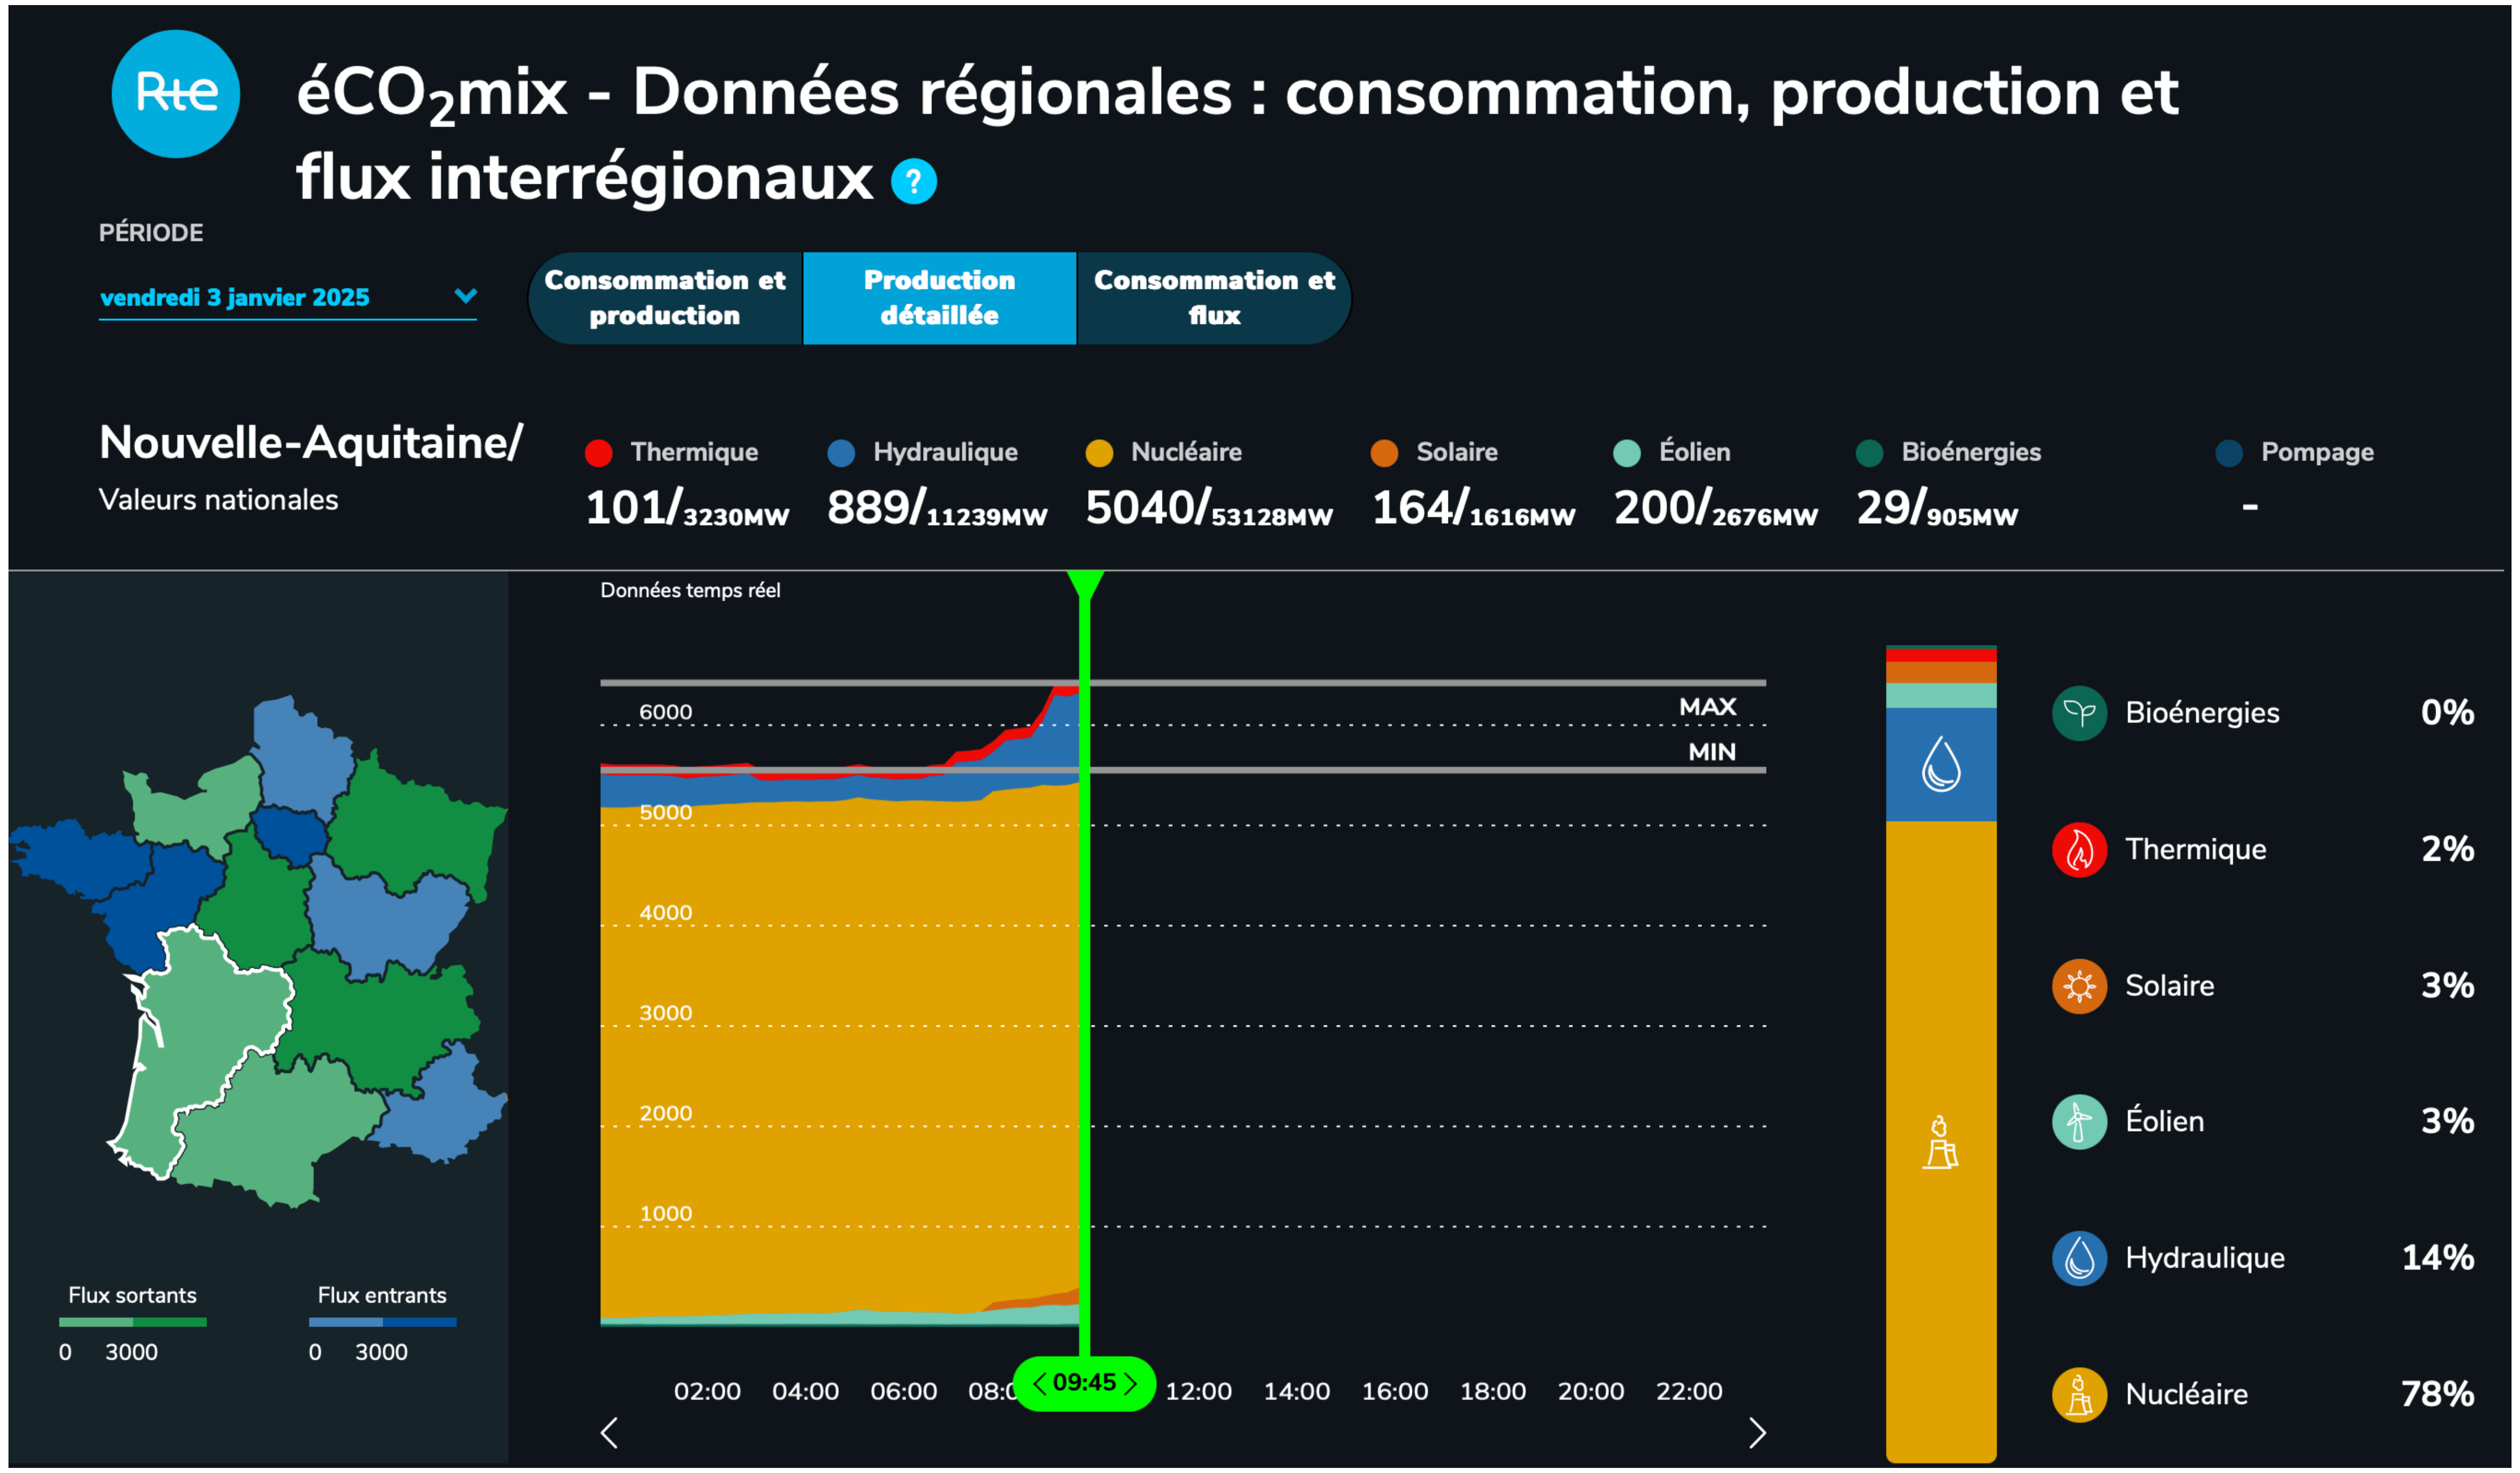

Base forecasts generation using online learning

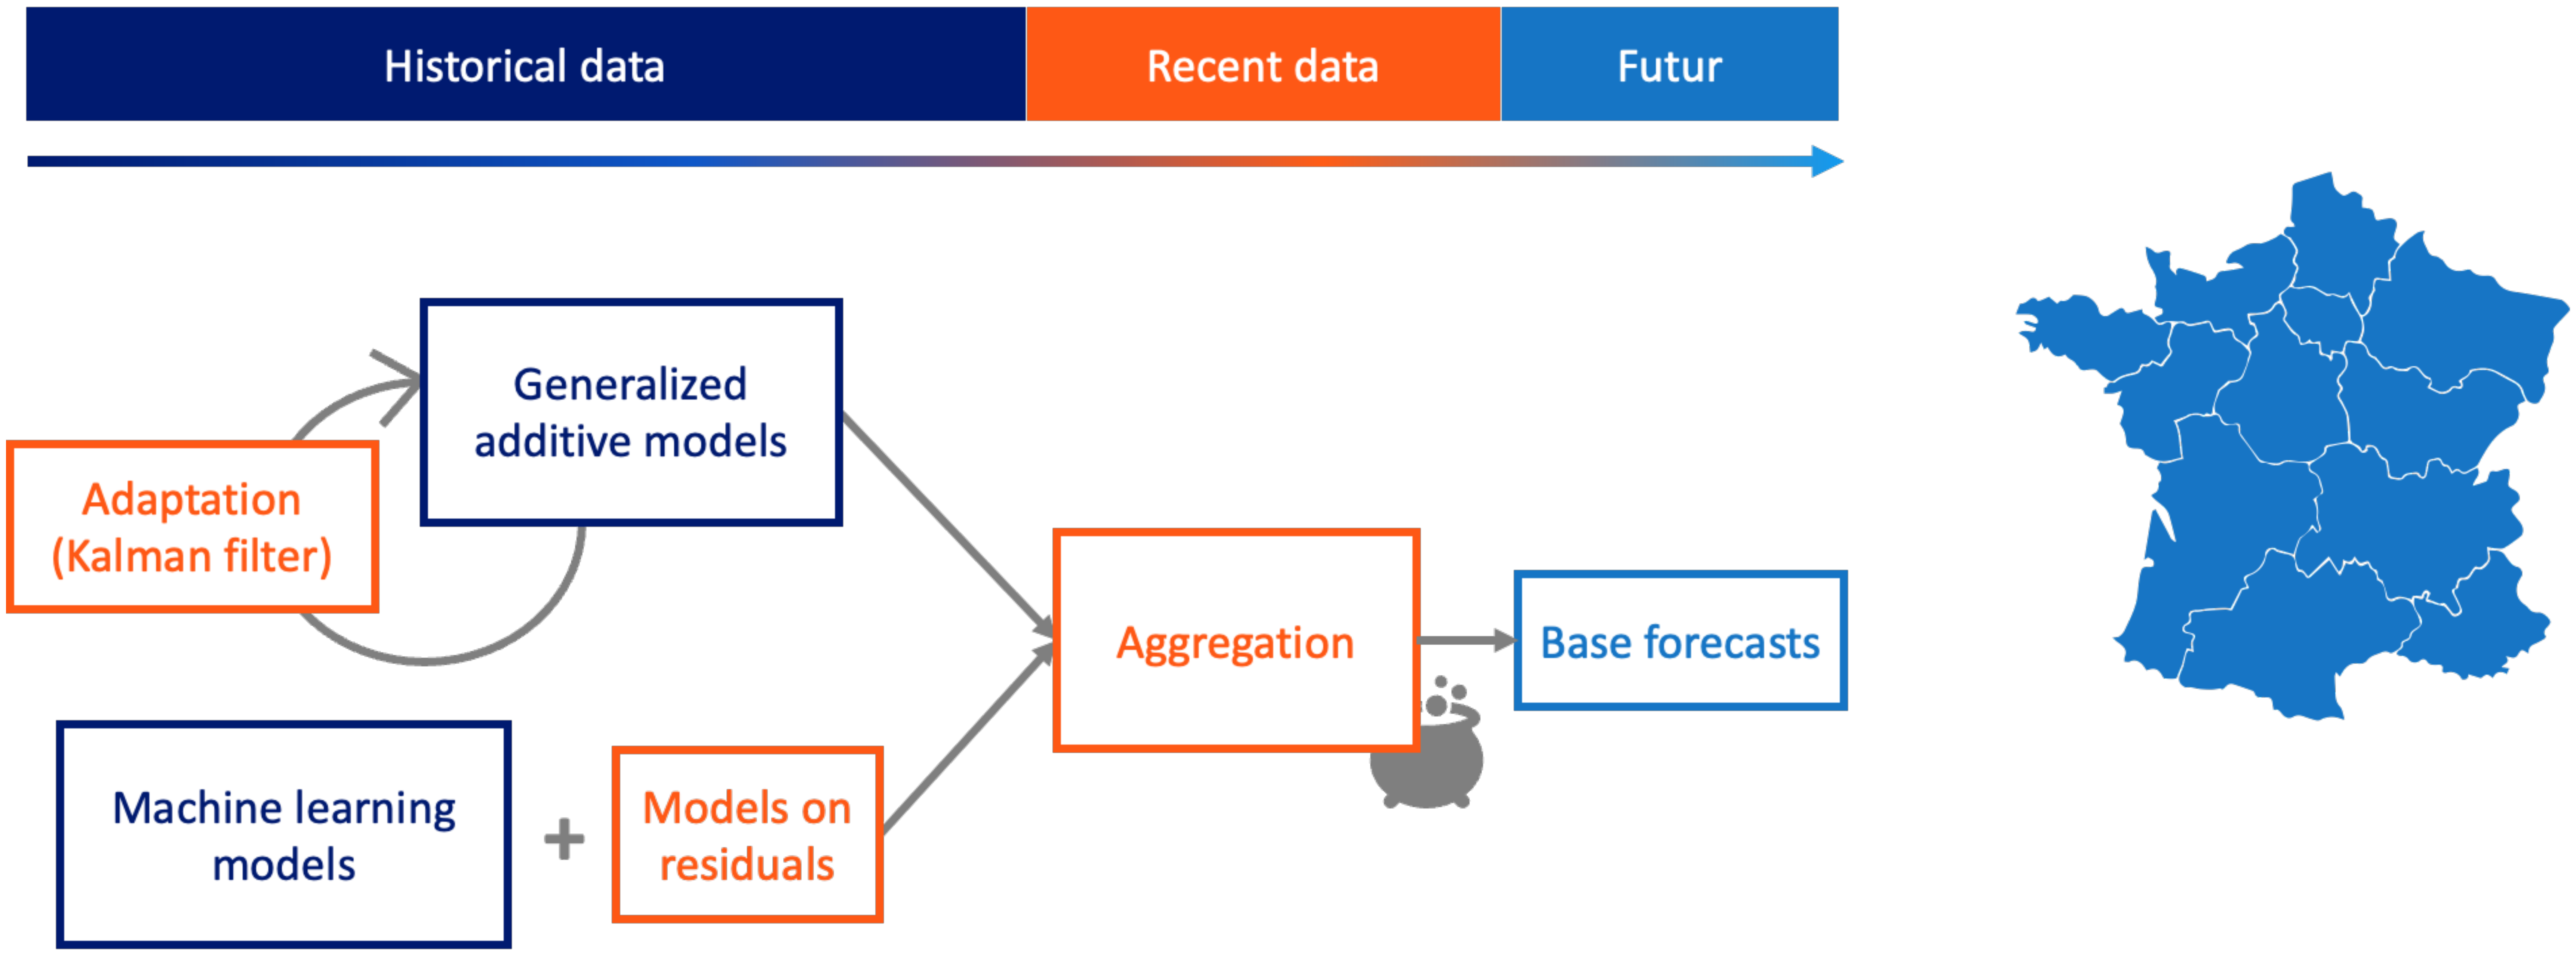

→ 13 base forecasts with $\hat{y}_t^{\text{France}} \neq \hat{y}_t^{\text{Nouvelle Aquitaine}} + \hat{y}_t^{\text{Bretagne}} + \ldots + \hat{y}_t^{\text{Île-de-France}}$

Online MinT

Input

- d: delay in data reception
- τ: window for the variance-covariance matrix of the base forecast errors estimation

For $t = 1,\ldots,T$

- For each level $\gamma \in \{\text{France}, \text{Auvergne-Rhône-Alpes}, \ldots, \text{Provence Alpes Côte d'Azur}\}$

  Generate online base forecast $\hat{y}_t = \left(\hat{y}_t^\gamma\right)_\gamma$
- Compute online empirical covariance matrix

$$\left(\hat{\Sigma}_t\right)_{\gamma\gamma'} = \frac{1}{\tau}\sum_{s=t-d-\tau}^{t-d}\left(e_s^\gamma - \bar{e}_s^\gamma\right)\left(e_s^{\gamma'} - \bar{e}_s^{\gamma'}\right) \quad \text{with} \quad e_s^\gamma = y_s^\gamma - \hat{y}_s^\gamma \text{ and } \bar{e}_s^\gamma = \frac{1}{\tau}\sum_{s=t-d-\tau}^{t-d} e_s^\gamma$$

- Reconcile base forecasts: $\tilde{y}_t = SP_t\hat{y}_t$ with $P_t = \left(S^{\mathrm{T}}\widehat{\Sigma}_t^\dagger S\right)^{-1}S^{\mathrm{T}}\widehat{\Sigma}_t^\dagger$

Results - 01.11.24 - 31.12.24

- Gam France: generalized additive model + Kalman filter on model effects
- Gam Regions: Bottom up approaches based on 13 (one for each french region) generalized additive model + online linear regression on models effects
- Gam MinT: Online MinT on using Gam Kalman (for France) and the 13 models (of the regions) of the bottom up approach as base forecasts
- Best model: online aggregation of many models

Model	RMSE (MW)	MAPE (%)	Mean bias (MW)
Gam France	1 318	1.83	90
Gam Regions (Bottom-Up)	1 220	1.63	-77
Gam MinT	**1 156**	**1.56**	**-31**
Best model	1 128	1.51	-23
RTE D-2	1455	2.04	339

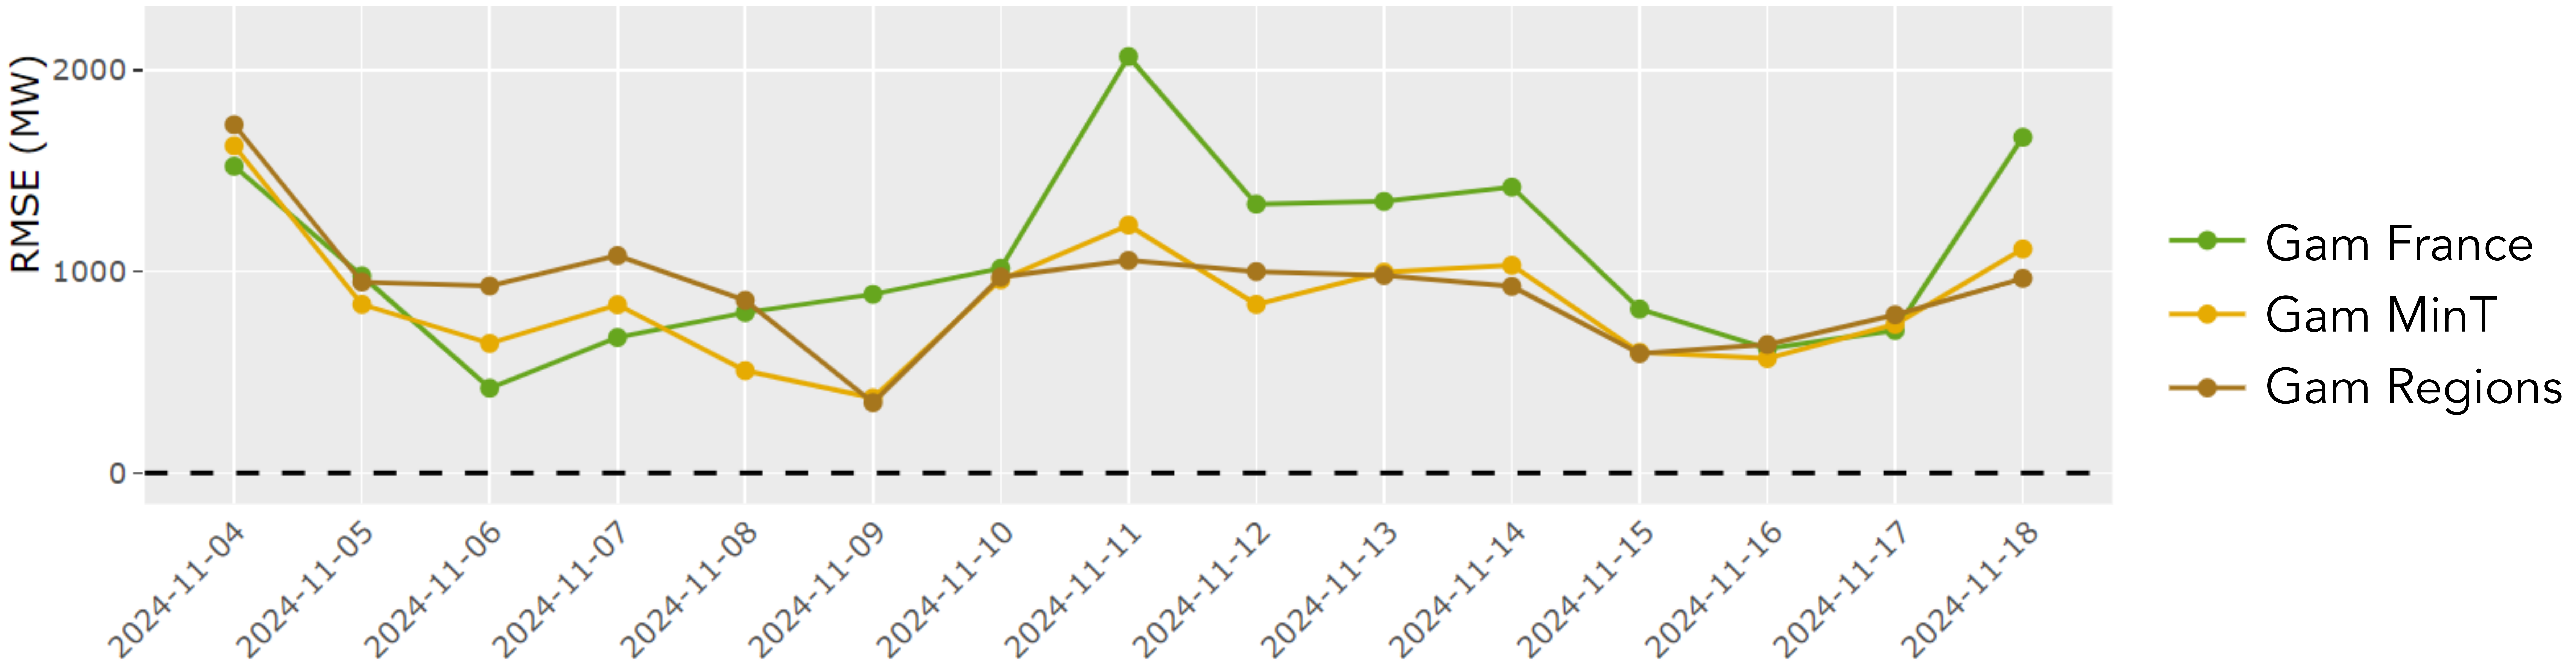

Prospects

- Use city data (bi-level hierarchy)
- Temporal and spatio-temporal reconciliation
- Combine clustering with reconciliation
- Reconciliation for probabilistic forecasts

2. Reinforcement learning for demand side management

- Demand Response: Send incentive signals → Bandits[3]
- Demand Despatch: Control flexible devices

Bianca Marin Moreno

Nadia Oudjane, EDF

Pierre Gaillard, Inria

[3] Stochastic Bandit Algorithms for Demand Side Management, University of Paris Saclay (LMO, EDF R&D, Inria Paris), under the supervision of G. Stoltz, Y. Goude and P. Gaillard, 2020

Mean Field Approach[4]

Control of M water heaters with same characteristics without compromising service quality

For water heater j, day t and time of the day n:

State: $x_{j,t}^{n} = \left(\text{Temperature}_{j,t}^{n}, \text{ON/OFF}_{j,t}^{n}\right)$

Action: $a_{j,t}^{n} = \left(\text{Turn/Keep ON/OFF}_{j,t}^{n}\right)$

$$\left.\begin{array}{l}\text{New state depends on:} \\ \quad\text{Temperature evolution (deterministic PDE)} \\ \quad\quad\text{+ Eventuel water drains (probabilistic law)} \\ \quad\quad\text{+ Action to turn/keep ON/OFF (service quality)}\end{array}\right\} \text{Markov Decision Process (MDP) } p$$

Mean Field assumption ($M \to \infty$): Control the state-action distribution $\mu^{\pi,p}$ induces by a policy π

[4] (Online) Convex Optimization for Demand-Side Management: Application to Thermostatically Controlled Loads, Bianca Marin Moreno, Margaux Brégère, Pierre Gaillard and Nadia Oudjane, 2024

Control with Mean Field Approach

At each day $t = 1, \ldots, T$

For each water-heater $j = 1, \ldots, M$

Initialization: $(x^0_{j,t}, a^0_{j,t}) \sim \mu_0$

For each instant of the day $n = 1, \ldots, N$

Send to all water heaters action $a^n_{j,t} \sim \pi^n_t(\cdot \mid x^n_{j,t})$

Observe new states $(x^n_{j,t})$ for all n and j

Loss function $F_t(\mu^{\pi_t, p})$ is exposed

Compute $\pi_{t+1} = (\pi^1_{t+1}, \ldots \pi^N_{t+1})$

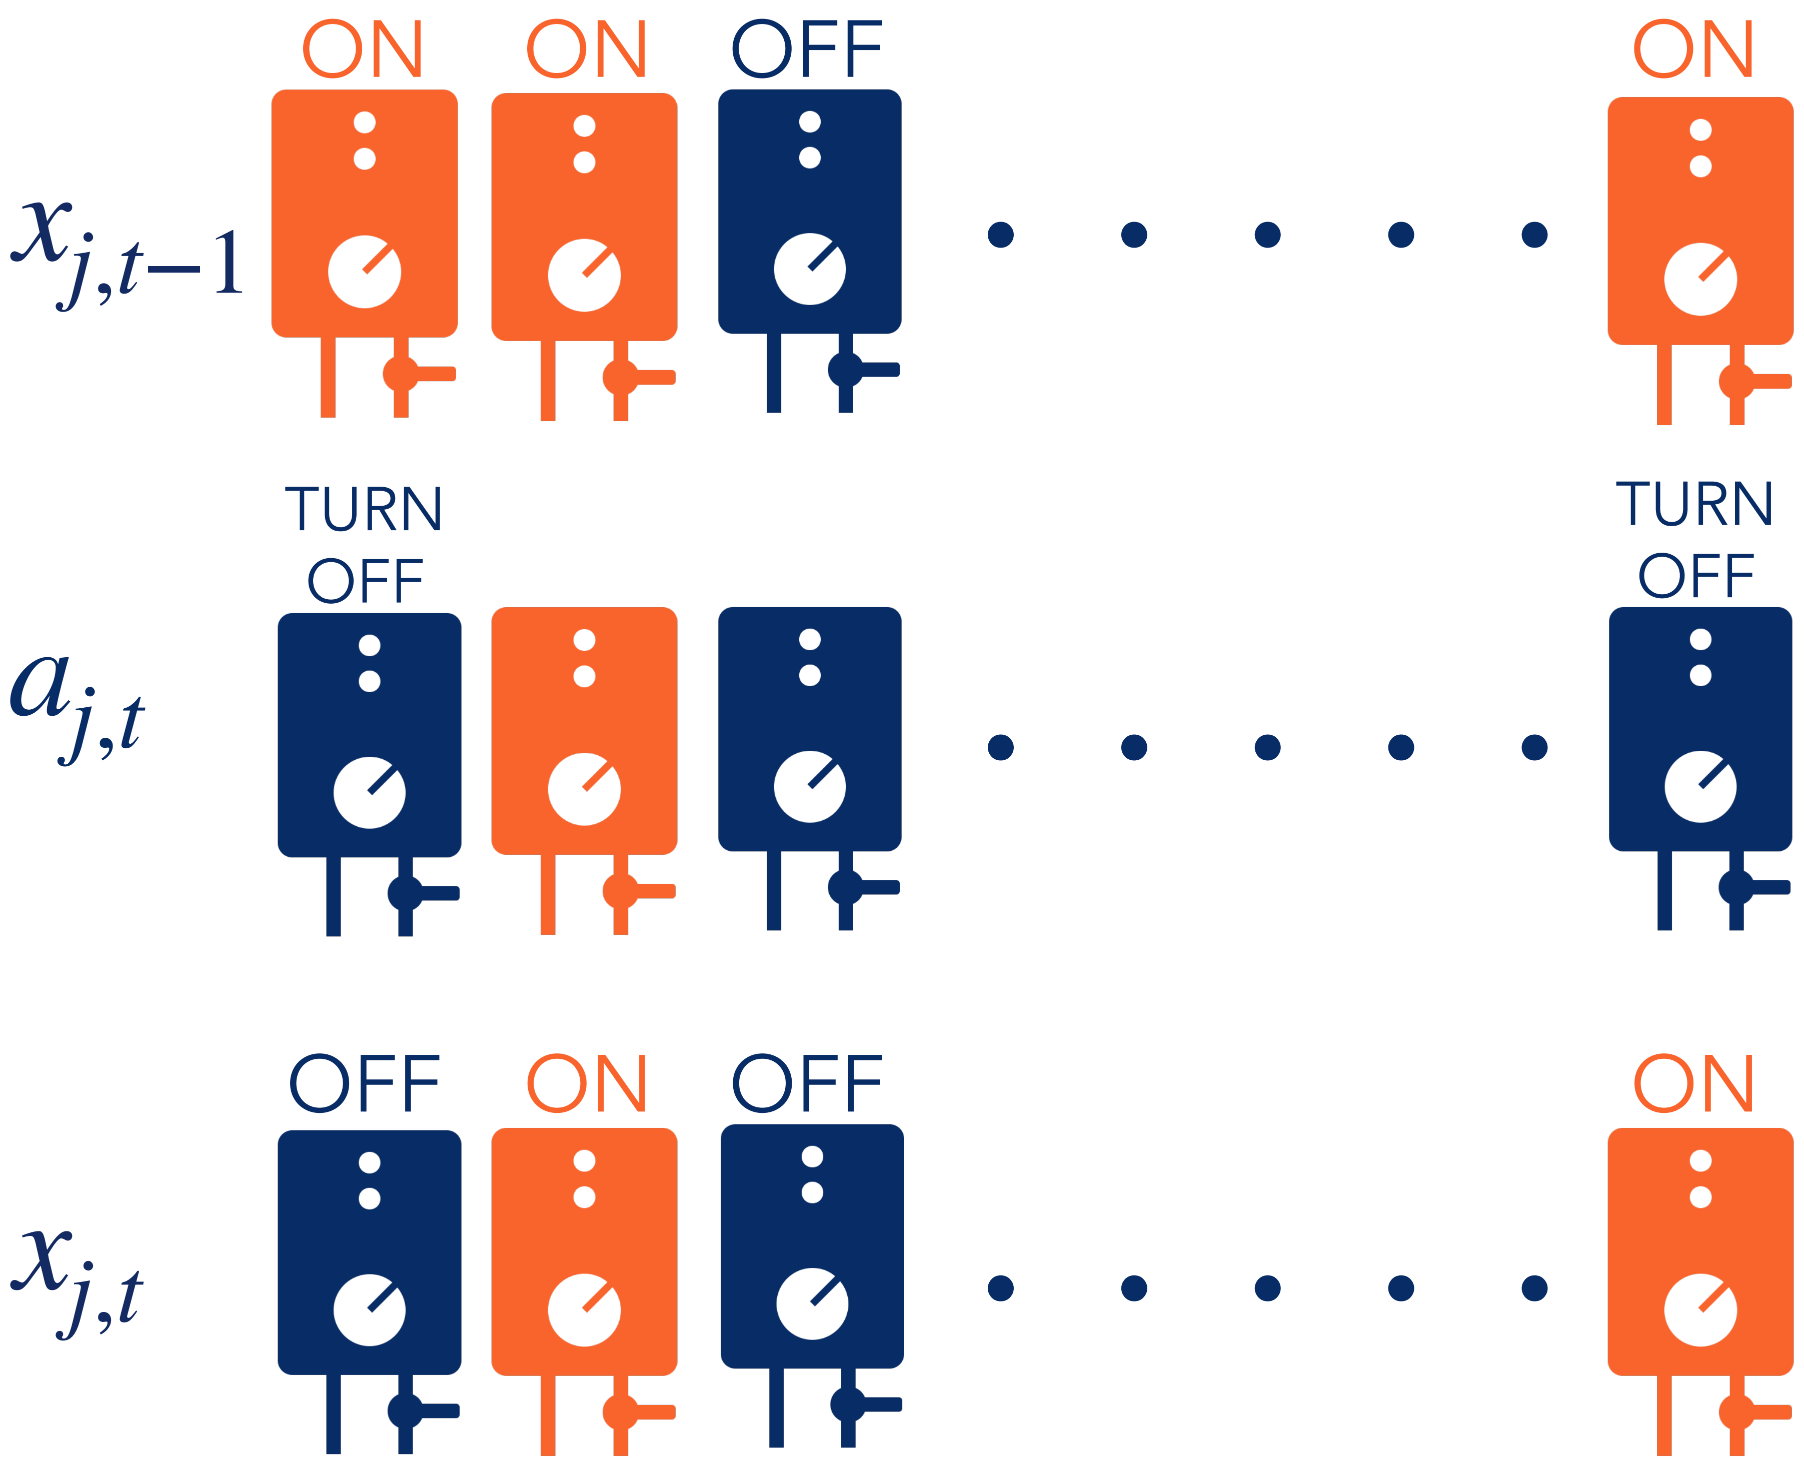

Aim: Find $\pi^\star \in \operatorname{argmin}_\pi \sum_{t=1}^{T} F_t\left(\mu^{\pi, p}\right)$

with F_t the quadratic difference between the consumption for all water-heaters and the target at t

CURL in online learning scenario[5]

- Mirror-Descent approach for CURL (convex reinforcement learning) when p and $F_t = F$ are known: $\pi^{\mathrm{MD}}(F, p)$

Initialization: policy π_1 (nominal = without control)

At each day $t = 1, \ldots, T$

…

Update the estimation of the MDP using the new observations: $\hat{p}_{t+1}$

Act if $F_{t+1} = F_t$ and compute $\pi_{t+1} = \pi^{\mathrm{MD}}(F_t, \hat{p}_{t+1})$

[5] Efficient Model-Based Concave Utility Reinforcement Learning through Greedy Mirror Descent, Bianca Marin Moreno, Margaux Brégère, Pierre Gaillard and Nadia Oudjane, AISTAT 2024

Extension: non-stationary MDP[6]

At each day $t = 1,\ldots,T$

- Restart previous algorithm $\mathscr{E}$ from the beginning: $\mathscr{E}^t$
- Define a new policy by aggregating the t policies $\pi_t = \sum_{s=1}^{t} \omega_{s,t}\pi_t^s$

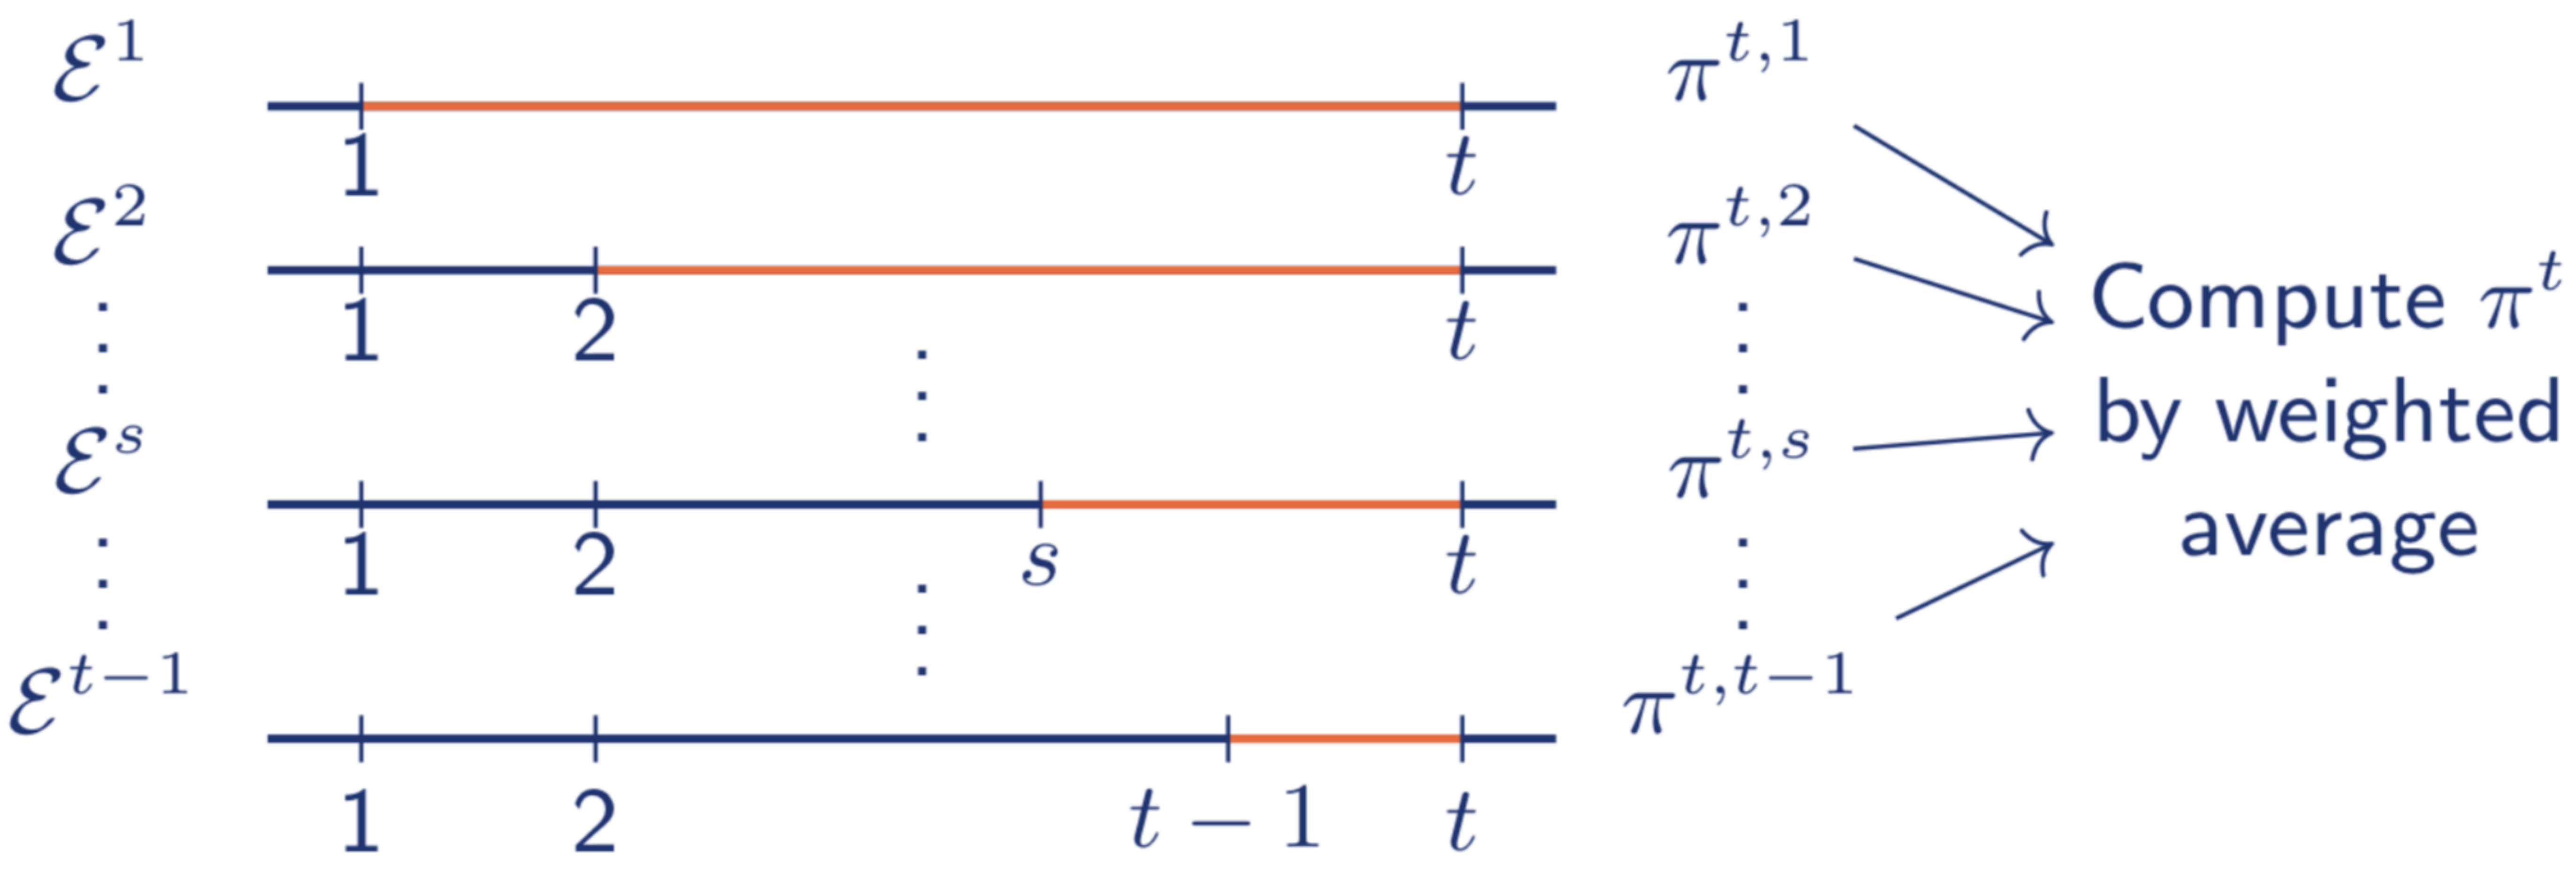

Prospects

- Constrained MDPs: $\mu_0^{\pi,p} = \mu_N^{\pi,p}$ (work in progress)
- Impact of the number of devices M on the control
- Extension of the current work to smart charging

[6] MetaCURL: Non-stationary Concave Utility Reinforcement Learning, Bianca Marin Moreno, Margaux Brégère, Pierre Gaillard and Nadia Oudjane, NeurIPS, 2024

3. Automated machine learning and Explainability

DRAGON

DiRected Acyclic Graphs OptimizatioN

Julie Keisler

XPC

eXplainability through Positive Contributions

Gaspard Berthelier

Sequential learning for AutoML

Train a neural network is expensive and time-consuming

Aim: for a search space Λ (set of possible architectures) and a budget T, find the best neural network:

$$\arg\min_{\lambda \in \Lambda} \ell\Big(f_\lambda\big(\mathscr{D}_{\text{TEST}}\big)\Big)$$

At each round $t = 1,\ldots,T$

- Choose hyper-parameters $\lambda_t \in \Lambda$
- Train network f_{λ_t} on $\mathscr{D}_{\text{TRAIN}}$
- Observe the forecast error $\ell_t = \ell\Big(f_{\lambda_t}\big(\mathscr{D}_{\text{VALID}}\big)\Big)$

Output (best arm identification): $\arg\min_{f_{\lambda_t}} \ell\Big(f_{\lambda_t}\big(\mathscr{D}_{\text{VALID}}\big)\Big)$

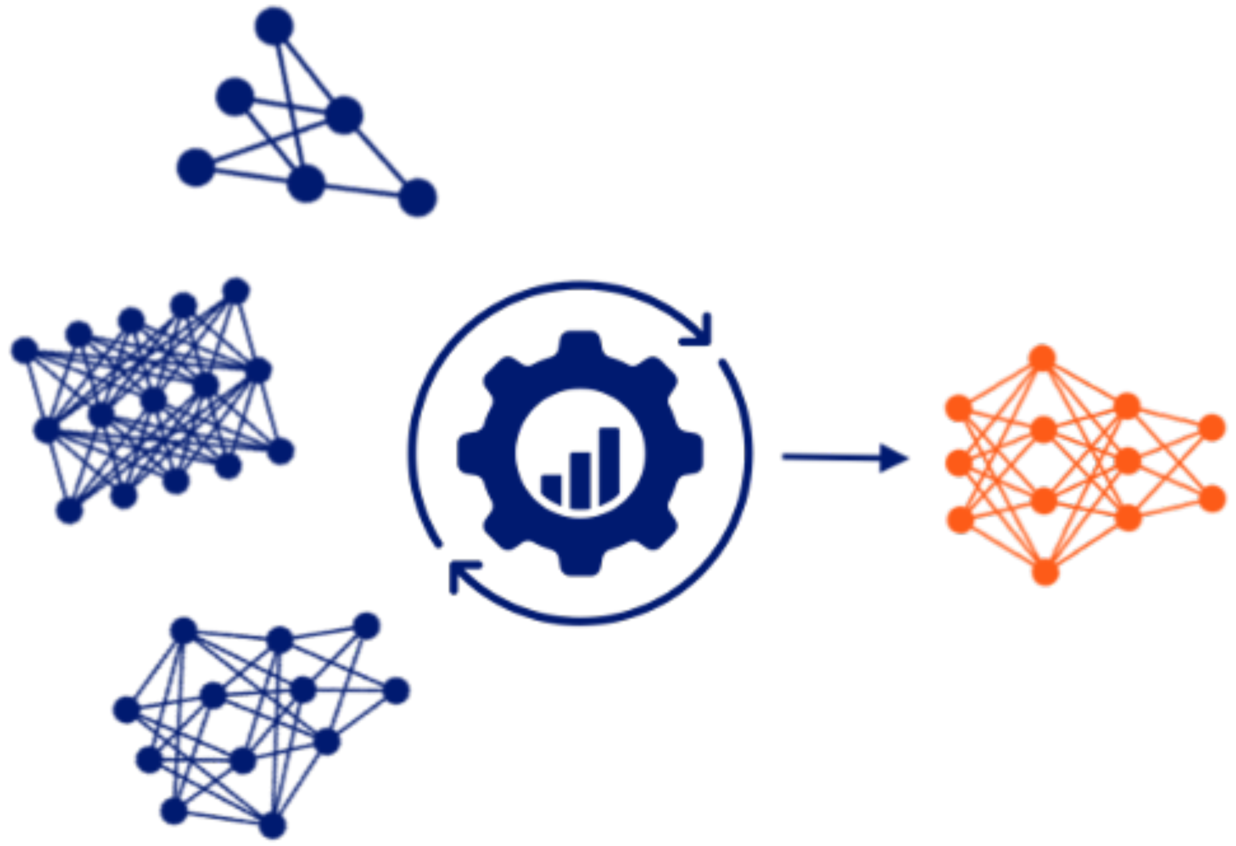

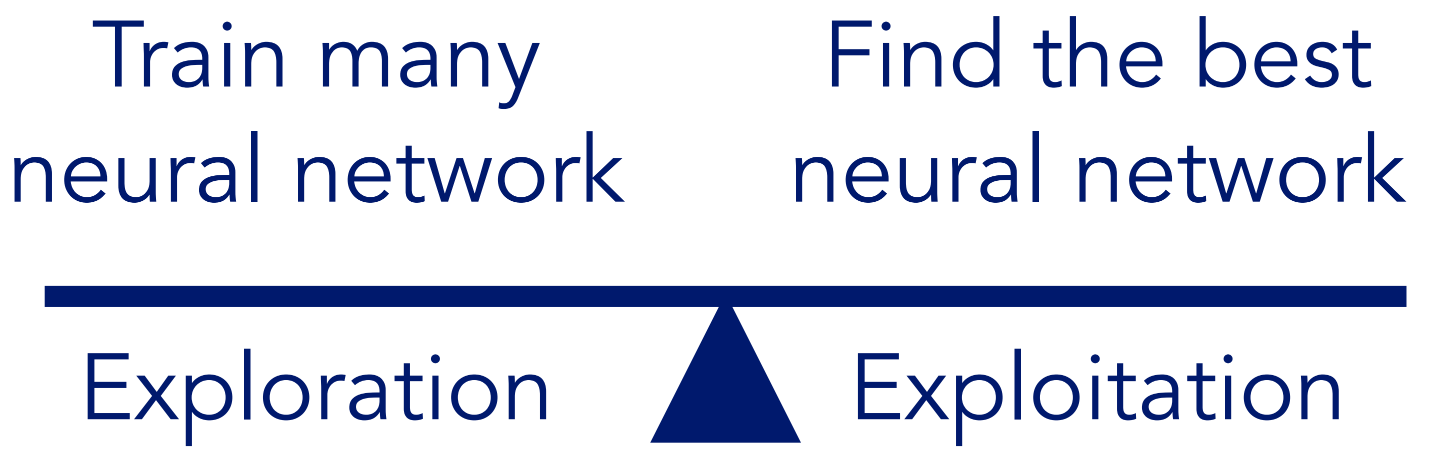

[7] A bandit approach with evolutionary operators for model selection: Application to neural architecture optimization for image classification, M. Brégère and J. Keisler, Submitted, 2024

Explainability of electricity demand forecasting models

Shapley value approach for positive component decomposition

Each feature value is a "player" in a collaborative game where the prediction is the payoff

4 115

+ 32 205

+ 11 503

=

47 823 MW

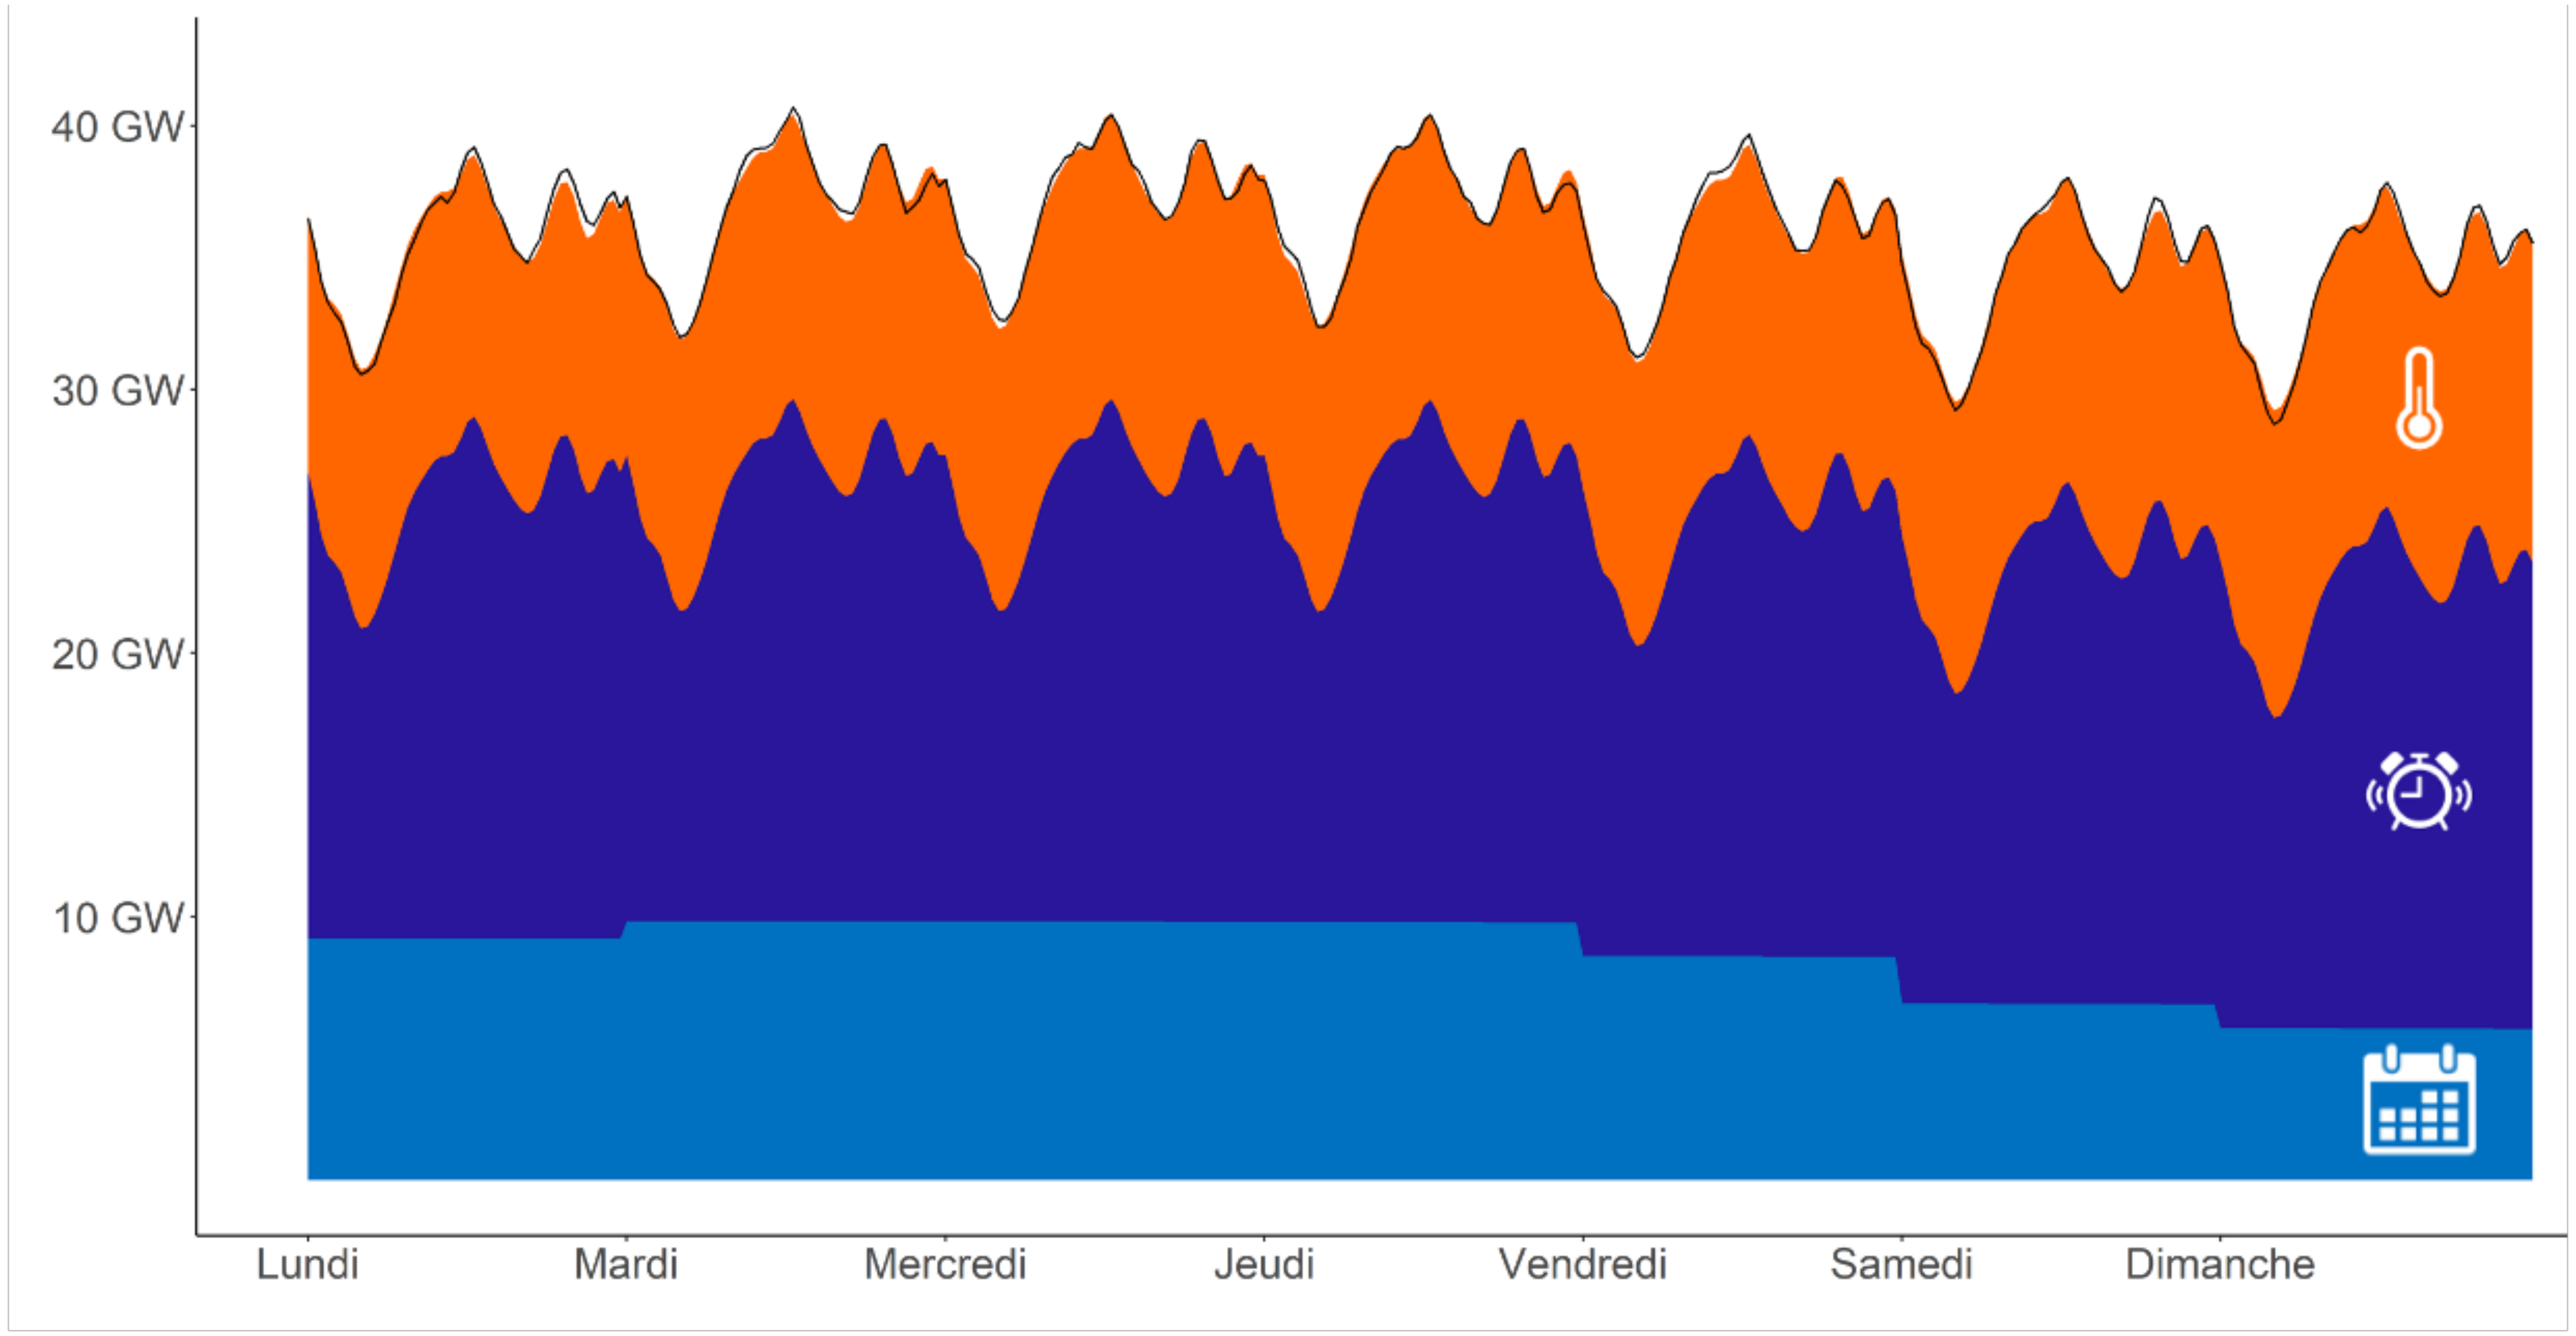

Thank you!