Moment Constrained Optimal Transport for Energy Demand Management of Heterogeneous Loads[1]

Julien Cardinal, **Thomas Le Corre**, Ana Bušić

November 18, 2025

[1]This work was carried out in the framework of the AI-NRGY project, funded by France 2030 (Grant No: ANR-22-PETA-0004).

Outline

Motivation

Context:

- High penetration of solar PV.
- A fixed daily demand profile.
- Lack of flexible demand or storage to absorb surplus.

Implications:

- Need for flexible generation.
- Grid stability challenges.
- Motivation for demand response and storage strategies.

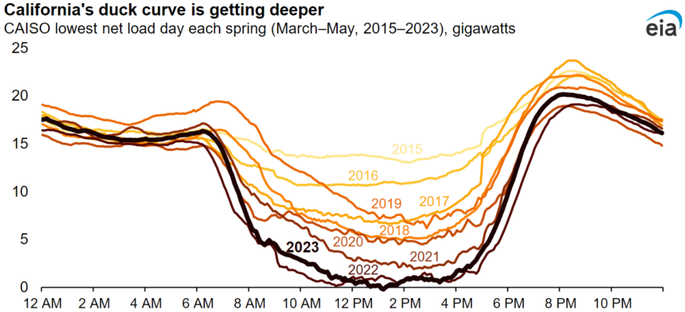

Challenges

- **Diversity of agents:** There are a large number of flexible loads that we want to consider:

Electric vehicles in a car park

Electric water heaters

...

- **Scalability:** We need a large number of agents to achieve flexibility.
- **Real time adaptation:** The behavior of each flexible agent is influenced by external factors that cannot be perfectly predicted in advance (e.g., water draws for water heaters, arrival and departure of electric vehicles, etc.).

Agents

Population Setup:

- N agents partitioned into H classes: $\mathcal{I}^{(h)}$ agents in class h, $N^{(h)} = |\mathcal{I}^{(h)}|$, $\alpha^{(h)} = N^{(h)}/N$.
- Each agent has state $s^i \in \mathcal{S}^{(h)}$ and control $w^i \in \mathcal{W}^{(h)}$.
- Controlling an agent of class h induces a cost $c^{(h)}(s,w)$.

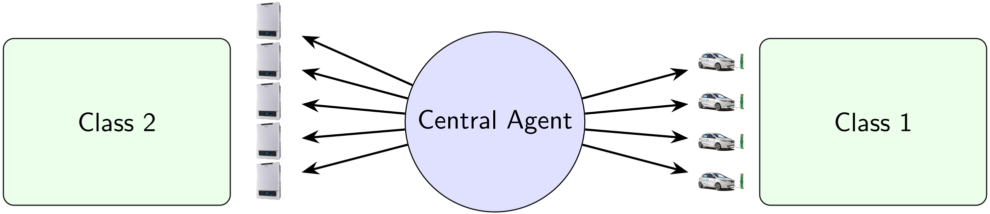

Global Constraints

Aggregate Constraints:

$$\forall a \in \{1, \ldots, A\}, \quad \sum_{h=1}^{H} \sum_{i \in \mathcal{I}^{(h)}} f_a^{(h)}(X^i) \leq 0$$

Examples:

- Total power consumption cap.
- Signal tracking (modeled as pairs of inequalities).
- Slope constraint (Consumption should not increase or decrease too abruptly).
- Constraints specific to a class (total power cap for a car park,...).

Objective Function

Minimization Problem We consider that a central agent (such as an electricity provider) wants to minimize :

$$\min_{\{w^i\}_{i=1}^N} \sum_{h=1}^{H} \sum_{i \in \mathcal{I}^{(h)}} c^{(h)}(s^i, w^i) \quad \text{s.t. } \forall a \in \{1, \dots, A\}, \quad \sum_{h=1}^{H} \sum_{i \in \mathcal{I}^{(h)}} f_a^{(h)}(X^i) \leq 0$$

The complexity of solving exactly this problem grows with the number of agents.

Mean-Field intuition

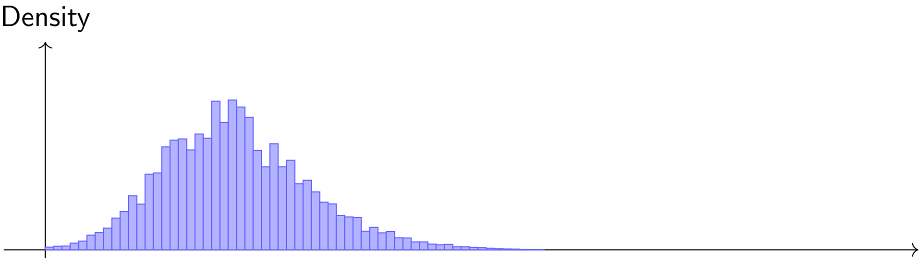

Mean-Field intuition

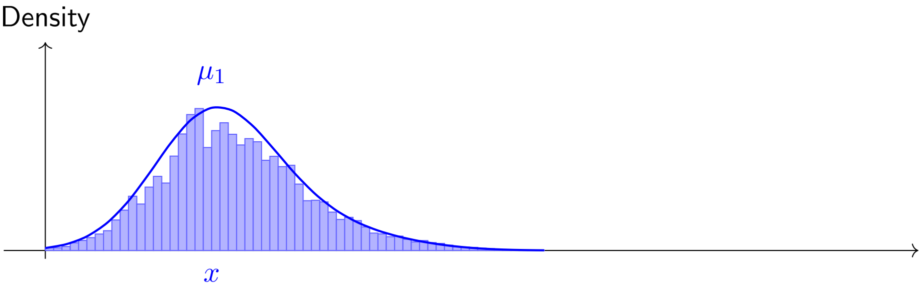

Mean-Field intuition

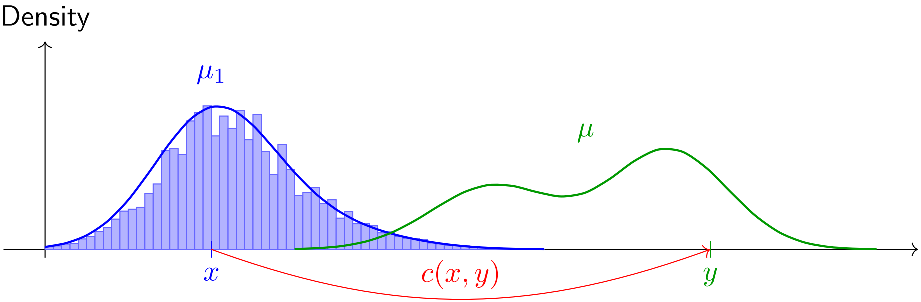

This type of problem is well known in **Optimal Transport** for a known μ, but here we know only:

$$\mu \in \mathcal{P}_f = \left\{ \{\mu^{(h)}\}_{h=1}^{H} : \sum_{h=1}^{H} \alpha^{(h)} \langle \mu^{(h)}, f_a^{(h)} \rangle \leq 0, \forall a \right\}$$

Mean-Field Formulation

Class-wise Distributions:

- Each class h is described by its distribution $\mu_1^{(h)} \in \mathcal{B}(\mathcal{X}^{(h)})$.
- The aggregate system: $\{\mu_1^{(h)}\}_{h=1}^{H}$.

Controlled Coupling:

- Joint distributions $\pi^{(h)} \in \mathcal{B}(\mathcal{X}^{(h)} \times \mathcal{X}^{(h)})$ represent the transport of the initial to the controlled distribution.
- First marginal fixed: $\pi_1^{(h)} = \mu_1^{(h)}$, for the control to be feasible.

$$\pi^{(h)} \in \mathcal{U}(\mu_1^{(h)}) = \left\{ \pi^{(h)} : \pi_1^{(h)} = \mu_1^{(h)},\ \pi^{(h)}((s_1,.),(s_2,.)) = 0 \ \forall s_1 \neq s_2 \in \mathcal{S} \right\}$$

- Second marginal must respect the **Global Moment Constraints:**

$$\{\pi_2^{(h)}\}_{h=1}^{H} \in \mathcal{P}_f = \left\{ \{\mu^{(h)}\}_{h=1}^{H} : \sum_{h=1}^{H} \alpha^{(h)} \langle \mu^{(h)}, f_a^{(h)} \rangle \leq 0, \forall a \right\}$$

Extension to Heterogeneous Loads

Multi-Class MCOT-C Problem:

$$\min_{\{\pi^{(h)}\}_{h=1}^{H}} \sum_{h=1}^{H} \alpha^{(h)} \Big(\langle \pi^{(h)}, c^{(h)} \rangle + \varepsilon D_{\mathsf{KL}}(\pi^{(h)} \| \mu_1^{(h)} \otimes \mu_2^{(h)}) \Big)$$

$$\text{s.t. } \pi^{(h)} \in \mathcal{U}(\mu_1^{(h)}), \quad \{\pi_2^{(h)}\}_{h=1}^{H} \in \mathcal{P}_f$$

where **Entropic regularization** ensures computational feasibility and incorporate hard constraints:

$$D_{\mathsf{KL}}(\pi^{(h)} \| \mu_1^{(h)} \otimes \mu_2^{(h)}) = \int \log \frac{\pi^{(h)}(x,y)}{\mu_1^{(h)}(x)\mu_2^{(h)}(y)} \pi^{(h)}(x,y) dx dy$$

Dual Problem and Gradient

Heterogeneous MCOT Dual Problem:

$$\varphi(\lambda) = \min_{\{\pi^{(h)}\}} \sum_{h=1}^{H} \alpha^{(h)} \left(\varepsilon D_{\mathsf{KL}}(\pi^{(h)} \| \mu_1^{(h)} \otimes \mu_2^{(h)}) - \langle \pi^{(h)}, \ell_\lambda^{(h)} \rangle \right) \quad \text{s.t. } \pi^{(h)} \in \mathcal{U}(\mu_1^{(h)})$$

Remark: The minimization *decouples* across classes h:

$$\varphi(\lambda) = \sum_{h=1}^{H} \alpha^{(h)} \underbrace{\min_{\pi^{(h)}} \left(\varepsilon D_{\mathsf{KL}}(\pi^{(h)} \| \mu_1^{(h)} \otimes \mu_2^{(h)}) - \langle \pi^{(h)}, \ell_\lambda^{(h)} \rangle \right)}_{\text{class-wise subproblem}}$$

This problem admits a **closed-form solution:** and the **gradient of the dual** can be expressed as:

$$\nabla\varphi(\lambda) = \sum_{h=1}^{H} \alpha^{(h)} \mathbb{E}_\lambda^{(h)}[f^{(h)}] \quad \Rightarrow \mathbb{E}_\lambda^{(h)} \text{can be computed from class-wise subproblems.}$$

Gradient Descent Algorithm for MCOT

Algorithm Gradient Descent to solve MCOT

1: Initialize λ
2: **for** $k = 1, \ldots, K$ **do**
3: $G \leftarrow \sum_{h=1}^{H} \alpha^{(h)} \mathbb{E}_{\lambda}^{(h)}[f^{(h)}]$
4: $\lambda \leftarrow \lambda - \rho_k G$
5: **end for**

- $\mathbb{E}_{\lambda}^{(h)}[f^{(h)}]$ computed:
  - Directly for small/finite state spaces (e.g., EVs)
  - Monte Carlo for large/infinite state spaces (e.g., WHs)
- K = number of iterations, $\{\rho_k\}$ = step sizes (empirically chosen)
- Objective: update dual variable λ to satisfy global constraints across heterogeneous classes.

Real time adaptation

Objective: Satisfy constraints (tracking, max consumption) *without knowing future load dynamics*.

Unknowns:

- Water Heaters: future water drains
- Electric Vehicles: number of EVs and charging needs

Algorithm Adaptation in real time

1: **for** $t = 1, \ldots, T$ **do**
2: Observe the unknowns and adapt the prediction
3: Run the previous algorithm to obtain a policy
4: Apply the policy
5: **end for**

Electric Vehicles: Dataset and State Space

Dataset: Elaad OpenDataset, composed of 10.000 transactions in the Netherlands, in 2019.

EV description: $\mathcal{S}^{\text{EVs}}$ is composed of:

- Arriving time t_a
- Leaving time t_l (deadline to be fully charged)
- Charging need Δt_n (duration to full charge)

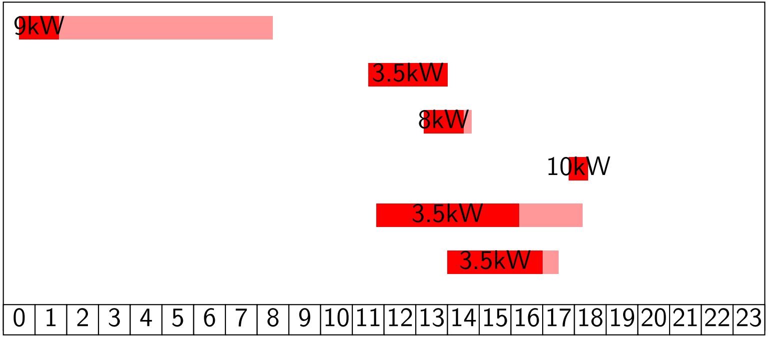

Figure: Sample of vehicles arriving on the 1st of January

EVs: Nominal Behavior and Control

Controlled state space: charging start t_c within

$$\mathcal{W}_t^{\text{EVs}} = \{t, \ldots, T\} \cap [t_a, t_l - \Delta t_n]$$

$$\mathcal{X}_t^{\text{EVs}} = \mathcal{S}^{\text{EVs}} \times \mathcal{W}_t^{\text{EVs}}$$

Nominal behavior: Plug when Arrive: $\mu_{1,t}(s, t_c) = \nu_t(s)\delta(t_c - t_a)$

Local cost function: Moving from this nominal behavior induce a quadratic penalty for late charging start

$$c(x, (t_a, t_l, \Delta t_n), t_c) = (t_c - t_a)^2$$

Water Heaters: Model and Dynamics

State variables for each WH:

- Mean temperature $\theta^i(t) \in \Theta$
- Power mode $m^i(t) \in \{0, 1\}$ (Off/On)

Differential equation: The temperature θ is following the following differential equation:

$$\frac{d\theta(t)}{dt} = -\underbrace{\rho(\theta(t) - \theta_{amb})}_{\text{heat loss}} + \underbrace{\sigma m(t) p}_{\text{Joule effect}} - \underbrace{\sigma \epsilon(t)}_{\text{water drain effect}}$$

Dataset: Drains averaged each 10 minutes from data generated by the simulator SMACH (Multi-agent Simulation of Human Activity in the Household).

Water Heaters: Controlled Switching

Nominal policy μ_1: Switch when reaching θ_{min} or θ_{max}.

Controlled policy: allow up to 2 additional switches t_1, t_2

Local cost: number of added switches:

$$c(x, s, \text{switches}) = \#\text{switches}$$

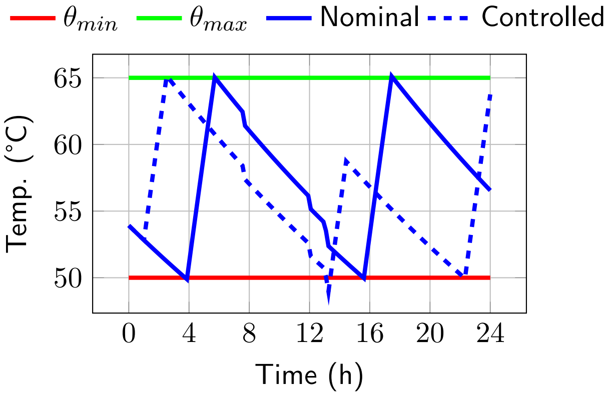

Figure: Nominal trajectory and controlled trajectory (t_1 =1:00 and at t_2 =14:20)

Results for Water Heaters

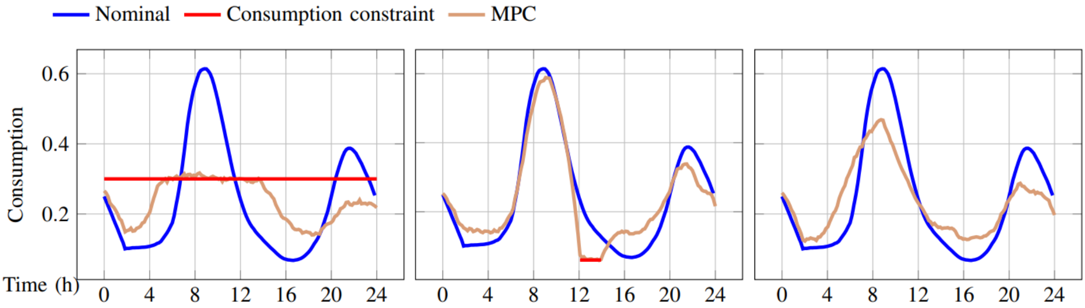

(a) Whole day constraint (b) Constraint between 12:00 and 14:00 (c) Constraint with gradient control

Fig. 8: (a) Optimized consumption compared to the nominal to limit the global consumption under 30% of the maximum consumption; (b) Limit of 6% between 12:00 and 14:00; (c) Gradient control for the whole day.

Results: Total Power Consumption

Setup: $N^{\text{WHs}} = 1000$, $N^{\text{EVs}} = 1000$

Constraints considered:

- **Tracking:** $P^{\text{Tot}}_{t'} = r_{t'}$, moves and reduces peak consumption.
- **Consumption cap:** $P^{\text{Tot}}_{t'} \leq u_{t'}$, e.g., limit 0.7 MW.

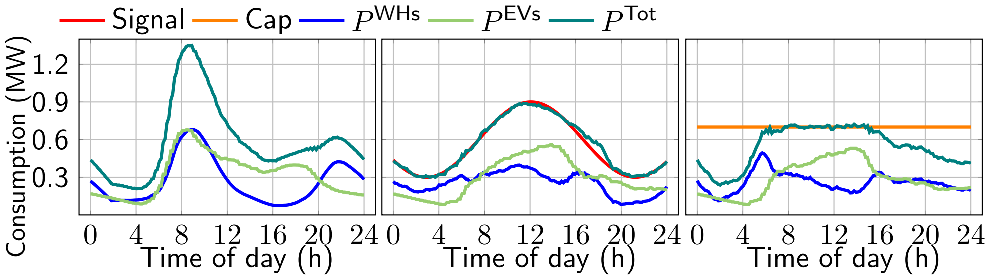

(a) Nominal (b) Tracking (c) Capped consumption

Summary of Contributions

- Extension of the MCOT-C framework to **heterogeneous populations** of controllable electrical loads.
- Coordination of multiple agent classes with **distinct dynamics and cost structures**, under global operational constraints.
- **Mean-field formulation** with class-wise distributions and entropic regularization $\Rightarrow$ tractable dual problem.
- Development of a **gradient-based algorithm** and integration into a **Model Predictive Control (MPC)** scheme for online adaptation.
- Validation on **realistic datasets** (EVs, WHs) $\Rightarrow$ effective constraint enforcement and interpretable decentralized policies.

Future Research Directions

- **Robustness analysis:** study the sensitivity of the MPC strategy to modeling errors and stochastic disturbances.
- **Dynamic constraints:** extend the framework to time-coupled constraints for more realistic scenarios.
- **Learning-based methods:** integrate online learning or reinforcement learning to update forecasts and estimate unknown agent models, improving long-term control under uncertainty.

Questions?