Online Moment Constrained Optimal Transport applied to Electric Vehicle Charging

Thomas Le Corre, Ana Bušić

Integrating more renewable energies in the electric mix

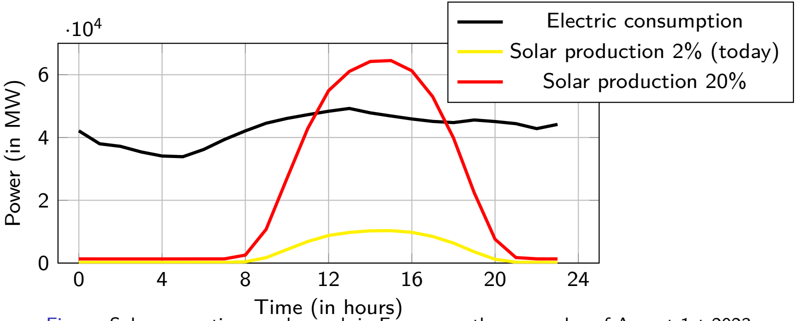

Figure: Solar generation vs. demand, in France, on the sunny day of August 1st 2023

How to do it ?

Use the flexibility of the demand

An example of flexibility: Charging a fleet of Electric Vehicles

A central planner wants to charge optimally a huge fleet of EVs over a finite time horizon. Different constraints must be taken into account:

- Satisfy EV owner requirements.
- Exploit EVs flexibility as much as possible.

Figure: An EV Charging station

Example of a Mean Field Control Problem

$$\min_{\mu}\left\{\int_{\mathcal{X}} c(x)d\mu(x) \ : \ \int_{\mathcal{X}} f(x)d\mu(x) \leq 0 \ \text{ and } \ \mu_1 = \nu\right\} \quad (1)$$

where μ is the distribution of $X = (S, W)$, and ν is the first marginal of μ.

Objectives

- Coordination of an ensemble of agents to achieve a desired goal.
- Enforcement of physical constraints, both spatial and dynamics.
- Enforcement of strict constraints on the distribution of exogenous variables.

Our Problem: *Moment Constrained Optimal Transport for Control* With given μ_1, μ_2, functions $\{f^{(m)}\}$, and $\varepsilon > 0$,

$$\min_{\pi,\mu}\Big\{\langle \pi, c\rangle + \varepsilon KL(\pi\|\mu_1 \otimes \mu_2) : \pi \in K(\mu_1, \mu),\ \mu \in \mathcal{P}_f\Big\}$$

.

Our Problem: *Moment Constrained Optimal Transport for Control* With given μ_1, μ_2, functions $\{f^{(m)}\}$, and $\varepsilon > 0$,

$$\min_{\pi,\mu}\left\{\langle \boldsymbol{\pi}, \boldsymbol{c}\rangle + \varepsilon KL(\pi \| \mu_1 \otimes \mu_2) : \pi \in K(\mu_1, \mu),\ \mu \in \mathcal{P}_f\right\}$$

- $\langle \pi, c\rangle = \int_{\mathcal{X}\times\mathcal{X}} c(x,y) d\pi(x,y)$ penalizes the deviation from a certain policy (Plug when Arrive for EVs).

.

Our Problem: *Moment Constrained Optimal Transport for Control* With given μ_1, μ_2, functions $\{f^{(m)}\}$, and $\varepsilon > 0$,

$$\min_{\pi,\mu}\Big\{\langle \pi, c\rangle + \boldsymbol{\varepsilon KL(\pi \| \mu_1 \otimes \mu_2)} : \pi \in K(\mu_1, \mu),\ \mu \in \mathcal{P}_f\Big\}$$

- $\langle \pi, c\rangle = \int_{\mathcal{X}\times\mathcal{X}} c(x,y)d\pi(x,y)$ penalizes the deviation from a certain policy (Plug when Arrive for EVs).
- $KL(\pi \| \mu_1 \otimes \mu_2)$ is an entropic term for regularization, ease of computation and incorporating hard constraints. ε is chosen small.

.

Our Problem: *Moment Constrained Optimal Transport for Control* With given μ_1, μ_2, functions $\{f^{(m)}\}$, and $\varepsilon > 0$,

$$\min_{\pi,\mu}\left\{\langle \pi, c\rangle + \varepsilon KL(\pi \| \mu_1 \otimes \mu_2) : \boldsymbol{\pi} \in \boldsymbol{K(\mu_1, \mu)},\ \mu \in \mathcal{P}_f\right\}$$

- $\langle \pi, c\rangle = \int_{\mathcal{X}\times\mathcal{X}} c(x,y)d\pi(x,y)$ penalizes the deviation from a certain policy (Plug when Arrive for EVs).
- $KL(\pi \| \mu_1 \otimes \mu_2)$ is an entropic term for regularization, ease of computation and incorporating hard constraints. ε is chosen small.
- $K(\mu_1, \mu)$ is the set of all bivariate distributions π with first and second marginals: $\pi_1 = \mu_1$ and $\pi_2 = \mu$ and preserving distribution of exogenous variable (arrival time for example).

.

Our Problem: *Moment Constrained Optimal Transport for Control* With given μ_1, μ_2, functions $\{f^{(m)}\}$, and $\varepsilon > 0$,

$$\min_{\pi,\mu}\left\{\langle \pi, c\rangle + \varepsilon KL(\pi \| \mu_1 \otimes \mu_2) : \pi \in K(\mu_1, \mu),\ \boldsymbol{\mu} \in \boldsymbol{\mathcal{P}_f}\right\}$$

- $\langle \pi, c\rangle = \int_{\mathcal{X}\times\mathcal{X}} c(x,y) d\pi(x,y)$ penalizes the deviation from a certain policy (Plug when Arrive for EVs).
- $KL(\pi \| \mu_1 \otimes \mu_2)$ is an entropic term for regularization, ease of computation and incorporating hard constraints. ε is chosen small.
- $K(\mu_1, \mu)$ is the set of all bivariate distributions π with first and second marginals: $\pi_1 = \mu_1$ and $\pi_2 = \mu$ and preserving distribution of exogenous variable (arrival time for example).
- $\mathcal{P}_f = \{\mu \in \mathcal{B}(\mathcal{X}) : \langle \mu, f_m\rangle \leq 0 \ : \ 1 \leq m \leq M\}$ is the set of distributions μ respecting moment constraints (here constraints on global consumption).

A large fleet of EVs arrives to a charging station between 9am and $10:30$am and must be fully charged by 5pm.
A central planner wants to maintain constraints for the aggregate power consumption. The goal is power tracking: total power consumption should follow a reference signal (r_t) over a time period $[t_1, t_2]$.

A large fleet of EVs arrives to a charging station between 9am and $10:30$am and must be fully charged by 5pm.
A central planner wants to maintain constraints for the aggregate power consumption. The goal is power tracking: total power consumption should follow a reference signal (r_t) over a time period $[t_1, t_2]$.

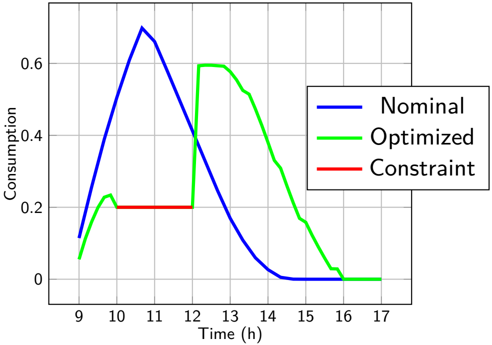

Figure: Optimized consumption compared to the nominal consumption

It is possible to add a slope control constraint to avoid a peak when the signal constraint is released.

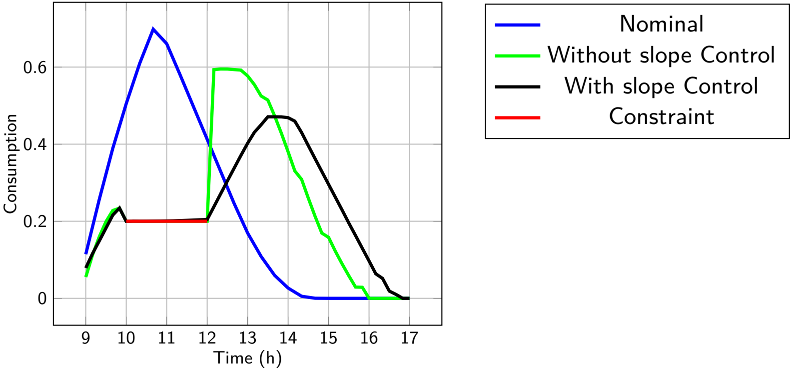

Figure: Optimized consumption with and without controlling the slope

Absolute continuity:

$$KL(\pi \| \mu_1 \otimes \mu_2) = +\infty \text{ if } \exists x \in \mathcal{X}, \mu_\lambda(x) = \pi_2(x) > 0 \text{ and } \mu_2(x) = 0$$

Thus, $\forall x \in \mathcal{X}, \mu_2(x) = 0 \Rightarrow \mu_\lambda(x) = 0$

Absolute continuity:

$$KL(\pi \| \mu_1 \otimes \mu_2) = +\infty \text{ if } \exists x \in \mathcal{X}, \mu_\lambda(x) = \pi_2(x) > 0 \text{ and } \mu_2(x) = 0$$

Thus, $\forall x \in \mathcal{X}, \mu_2(x) = 0 \Rightarrow \mu_\lambda(x) = 0$

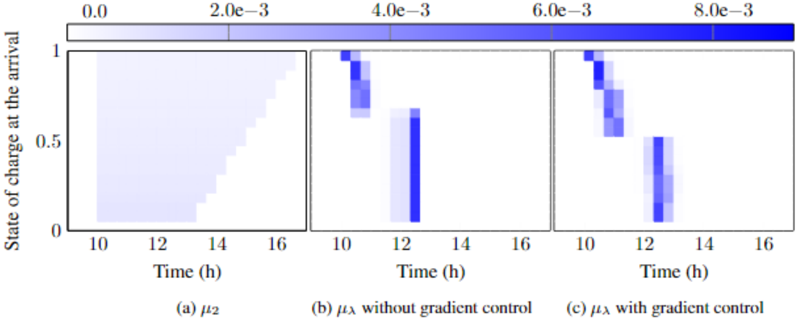

Figure 1: For vehicles arriving at 10am : (a) μ_2 designed to encode physical and quality of service constraints; (b) optimized μ without gradient control; (c) optimized μ with gradient control.

At each time step $t \in [0, 24]$, this sequence of actions take place:

1. New EVs arrive at the charging station and are added to the list of non-charging EVs.

At each time step $t \in [0, 24]$, this sequence of actions take place:

1. New EVs arrive at the charging station and are added to the list of non-charging EVs.

2. $\mu_1^{(t)}$ is defined by the "Plug when Arrive" strategy and $\mu_2^{(t)}$ is defined as "Plug with a uniform distribution" strategy, designed to incorporate the strong constraint of respecting the quality of service through the absolute continuity of μ with μ_2 (due to the KL term).

At each time step $t \in [0, 24]$, this sequence of actions take place:

1. New EVs arrive at the charging station and are added to the list of non-charging EVs.

2. $\mu_1^{(t)}$ is defined by the "Plug when Arrive" strategy and $\mu_2^{(t)}$ is defined as "Plug with a uniform distribution" strategy, designed to incorporate the strong constraint of respecting the quality of service through the absolute continuity of μ with μ_2 (due to the KL term).

3. The central planner will minimize the MCOT-C problem to obtain an optimal policy.

At each time step $t \in [0, 24]$, this sequence of actions take place:

1. New EVs arrive at the charging station and are added to the list of non-charging EVs.

2. $\mu_1^{(t)}$ is defined by the "Plug when Arrive" strategy and $\mu_2^{(t)}$ is defined as "Plug with a uniform distribution" strategy, designed to incorporate the strong constraint of respecting the quality of service through the absolute continuity of μ with μ_2 (due to the KL term).

3. The central planner will minimize the MCOT-C problem to obtain an optimal policy.

4. Each vehicle's plugging time is randomly chosen according to the policy.

At each time step $t \in [0, 24]$, this sequence of actions take place:

1. New EVs arrive at the charging station and are added to the list of non-charging EVs.
2. $\mu_1^{(t)}$ is defined by the "Plug when Arrive" strategy and $\mu_2^{(t)}$ is defined as "Plug with a uniform distribution" strategy, designed to incorporate the strong constraint of respecting the quality of service through the absolute continuity of μ with μ_2 (due to the KL term).
3. The central planner will minimize the MCOT-C problem to obtain an optimal policy.
4. Each vehicle's plugging time is randomly chosen according to the policy.

The previous policy is used as the regularizer for the next step

10 000 EVs transactions in the Netherlands during 2019:

Elaad OpenDataset. Available online:
https://platform.elaad.io/analyses/ElaadNL_opendata.php

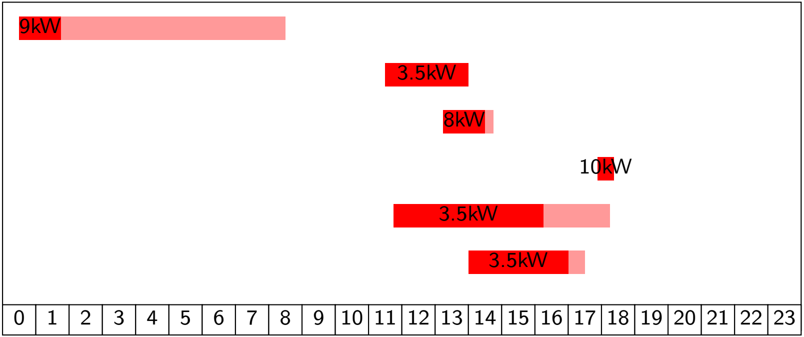

Figure: Sample of vehicles arriving on the 1st of January

We select the 7253 transactions happening during weekdays:

- Training set: 90% of transactions. The predicted distribution ν is computed on this set and the expected number of EV arriving during the day is $N = \frac{6540}{9} \simeq 727$.

- Test day: the remaining 10% of weekdays.

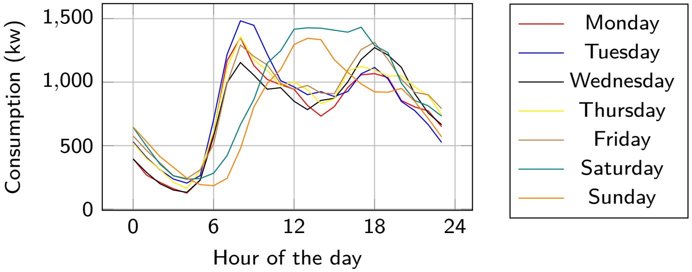

Figure: Aggregated Power consumption depending on the weekday

It is possible to implement constraints on the maximum power and to the gradient of the power consumption.

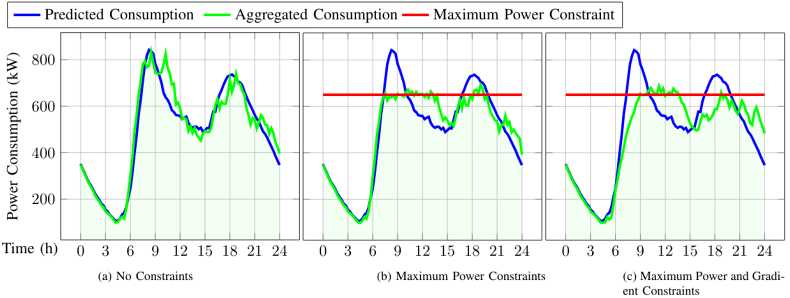

Fig. 2: (a) Consumptions for the "Plug When Arrive" μ_1 strategy with the arrival of EV predicted with ν and with the real distribution of EV; (b) Optimized Consumption for a constraint of 650kW for the aggregated consumption; (c) Optimized consumption for the same maximum power constraint and a constraint of 120kW/h for the gradient of the aggregated consumption.

Compliance with the constraint seems robust to the prediction ν.

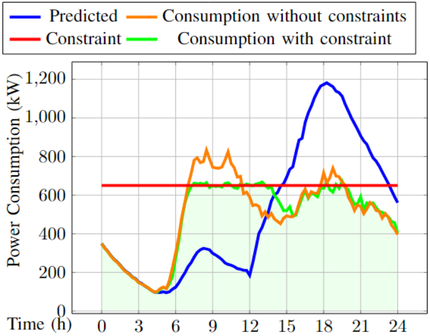

Fig. 3: When the prediction ν differs greatly from the reality

Compliance with the constraint seems robust to the prediction ν.

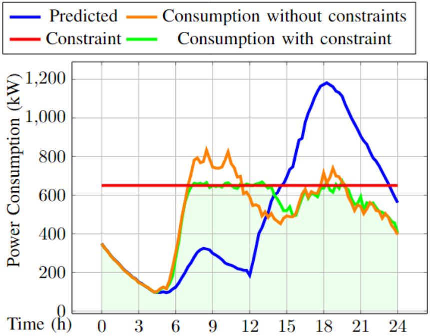

Fig. 3: When the prediction ν differs greatly from the reality

But the average time between arrival time t_a and connection time t_c increases from 11 to 12 minutes.

Conclusions:

- MCOT-C is an MFC problem with an Optimal Transport penalization of the deviation from a certain policy.
- The KL term allows incorporating hard physical constraints "for free".
- Our algorithm shows good results on this real dataset.

What to do next ?

- Compare with the literature
- MCOT-C is particularly interesting when the state space is large (or infinite). We would like to explore this aspect from a theoretical point of view.