Online Submodular Maximization via Online Convex Optimization

Alireza Kazemipour

CMPUT 676: Optimization and Decision-Making under Uncertainty

December 9th, 2024

Structure of the presentation

- **What?**
- **How?**
- **Why?**
- **Proposal/What I've been doing.**

(What) What does the paper try to do?

Maximizing $\underbrace{\text{monotone}}_{I}$ $\underbrace{\text{submodular}}_{II}$ functions under general $\underbrace{\text{matroid}}_{III}$ constraints via online convex optimization.

(What) What does the paper try to do?

~~Maximizing~~ $\underbrace{\textbf{monotone}}_{I}$ $\underbrace{\text{\sout{submodular}}}_{II}$ **functions** ~~under general~~ $\underbrace{\text{\sout{matroid constraints}}}_{III}$ ~~via online convex optimization~~.

- Let $V \triangleq [n], n \in \mathbb{N}$.
- A **set function** $f : 2^V \to \mathbb{R}$ is monotone if:

$$f(A) \leq f(B), \forall A, B \text{ that } A \subseteq B \subseteq 2^V.$$

(What) What does the paper try to do?

~~Maximizing~~ $\underbrace{\text{~~monotone~~}}_{I}$ $\underbrace{\textbf{submodular}}_{II}$ **functions** ~~under general~~ $\underbrace{\text{~~matroid~~}}_{III}$ ~~constraints via online convex optimization~~.

- Let $V \triangleq [n], n \in \mathbb{N}$.
- A **set function** $f : 2^V \to \mathbb{R}$ is submodular if:

$$f(B \cup \{v\}) - f(B) \leq f(A \cup \{v\}) - f(A),$$
$$\forall A, B \text{ that } A \subseteq B \subseteq 2^V \text{ and } v \in V \setminus B.$$

(What) What does the paper try to do?

~~Maximizing~~ $\underbrace{\text{~~monotone~~}}_{I}$ $\underbrace{\textbf{submodular}}_{II}$ **functions** ~~under general~~ $\underbrace{\text{~~matroid~~}}_{III}$ ~~constraints via online convex optimization~~.

- Let $V \triangleq [n], n \in \mathbb{N}$.
- A **set function** $f : 2^V \to \mathbb{R}$ is submodular if:

$$f(B \cup \{v\}) - f(B) \leq f(A \cup \{v\}) - f(A),$$
$$\forall A, B \text{ that } A \subseteq B \subseteq 2^V \text{ and } v \in V \setminus B.$$

- *It kinda resembles the notion of convexity/concavity in set functions.*

(What) What does the paper try to do?

~~Maximizing $\underset{I}{\underline{\text{monotone}}}$ $\underset{II}{\underline{\text{submodular}}}$ functions under general~~
$\underset{III}{\underline{\textbf{matroid constraints}}}$ ~~via online convex optimization~~.

- Let $V \triangleq [n], n \in \mathbb{N}$, and $\mathcal{I} \subseteq 2^V$.
- A **matroid** is a pair $\mathcal{M} = (V, \mathcal{I})$ such that:
  - ◇ If $B \in \mathcal{I}$ and $A \subseteq B$ then $A \in \mathcal{I}$.
  - ◇ If $A, B \in \mathcal{I}$ and $|A| < |B|$ then there exists a $b \in B$ such that $A \cup \{b\} \in \mathcal{I}$

(What) What does the paper try to do?

~~Maximizing~~ $\underbrace{\text{monotone}}_{I}$ $\underbrace{\text{submodular}}_{II}$ ~~functions under general~~ $\underbrace{\textbf{matroid constraints}}_{III}$ ~~via online convex optimization~~.

- Let $V \triangleq [n], n \in \mathbb{N}$, and $\mathcal{I} \subseteq 2^V$.
- A **matroid** is a pair $\mathcal{M} = (V, \mathcal{I})$ such that:
  - ◇ If $B \in \mathcal{I}$ and $A \subseteq B$ then $A \in \mathcal{I}$.
  - ◇ If $A, B \in \mathcal{I}$ and $|A| < |B|$ then there exists a $b \in B$ such that $A \cup \{b\} \in \mathcal{I}$
- *It generalizes the concept of linear independence in vector spaces to sets.*

(What) Let's put them in an example.

Facility Location Problem (FLP): We want to choose a subset of potential locations $S \subseteq L = \{l_1, \cdots, l_n\}$ to open our facilities (e.g., warehouses) to minimize the total building cost of the facility, and minimizing the distance d between the locations and clients $C = \{c_1, \cdots c_m\}$ that should be served. So the objective is:

$$\min \left[\sum_{i \in S} \text{cost}_{\text{build}}(l_i) + \sum_{j \in C} \min_{l \in S} d(l, c_j) \right] =$$

$$\max \left[- \sum_{i \in S} \text{cost}_{\text{build}}(l_i) - \sum_{j \in C} \min_{l \in S} d(l, c_j) \right]$$

(What) Let's put them in an example.

$$f(B \cup \{v\}) - f(B) \leq f(A \cup \{v\}) - f(A), A \subseteq B$$

- FLP is a submodular maximization (minimization) problem, because:
  - **Diminishing returns**. At a certain point, the benefit of opening a new facility is less valuable than the initial phase, because the existing facilities have already done a lot of the heavy lifting.

(What) Let's put them in an example.

$$f(B \cup \{v\}) - f(B) \leq f(A \cup \{v\}) - f(A), A \subseteq B$$

- FLP is a submodular maximization (minimization) problem, because:
  - **Diminishing returns**. At a certain point, the benefit of opening a new facility is less valuable than the initial phase, because the existing facilities have already done a lot of the heavy lifting.
- The decisions are of the form $\{0, 1\}^n$.

It generalizes the linear independence in vectors to sets.

(What) FLP

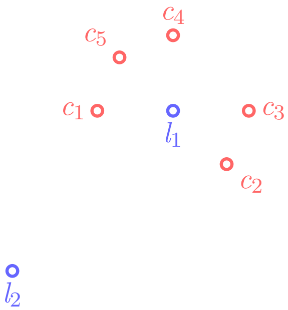

(What) FLP

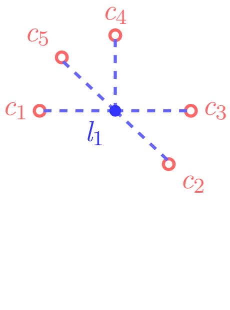

(What) FLP

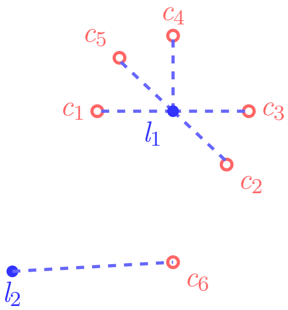

(What) Submodular Maximization is NP-Hard

- In general the problem of maximizing submodular functions, even in offline setting, is NP-hard[1] as there are reductions to the Traveling Salesman Problem.

[1] Krause and Golovin, "Submodular function maximization."

(What) Submodular Maximization is NP-Hard

- In general the problem of maximizing submodular functions, even in offline setting, is NP-hard[1] as there are reductions to the Traveling Salesman Problem.
- The best approximation ratio by the greedy algorithm is:

$$\alpha = 1 - \frac{1}{e}$$

[1] Krause and Golovin, "Submodular function maximization."

(What) So the paper does:

In the online setting, the reward (cost) is revealed one by one. Hence:

- Let $\mathcal{X}$ be the decision space forming a matroid.
- Let $f_t : \mathcal{X} \to \mathbb{R}_{\geq 0}$ be a reward function selected by adversary among the set of submodular functions $\mathcal{F}$.

The paper's algorithm tries to find a policy π_x to minimize:

(What) So the paper does:

In the online setting, the reward (cost) is revealed one by one. Hence:

- Let $\mathcal{X}$ be the decision space forming a matroid.
- Let $f_t : \mathcal{X} \to \mathbb{R}_{\geq 0}$ be a reward function selected by adversary among the set of submodular functions $\mathcal{F}$.

The paper's algorithm tries to find a policy π_x to minimize:

- **Static regret**:

$$SR_T(\pi_{\mathbf{x}}, \alpha) := \sup_{(f_t)_{t=1}^T \in \mathcal{F}^T} \left\{ \alpha \max_{\mathbf{u} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{u}) - \mathbb{E}\left[\sum_{t=1}^T f_t(\mathbf{x}_t) \right] \right\}$$

(What) So the paper does:

In the online setting, the reward (cost) is revealed one by one. Hence:

- Let $\mathcal{X}$ be the decision space forming a matroid.
- Let $f_t : \mathcal{X} \to \mathbb{R}_{\geq 0}$ be a reward function selected by adversary among the set of submodular functions $\mathcal{F}$.

The paper's algorithm tries to find a policy π_x to minimize:

- **Static regret**:

$$SR_T(\pi_{\mathbf{x}}, \alpha) := \sup_{(f_t)_{t=1}^{T} \in \mathcal{F}^T} \left\{ \alpha \max_{\mathbf{u} \in \mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{u}) - \mathbb{E}\left[\sum_{t=1}^{T} f_t(\mathbf{x}_t) \right] \right\}$$

- **Dynamic regret**:

$$DR_T(\pi_{\mathbf{x}}, P_T, \alpha) := \sup_{(f_t)_{t=1}^{T} \in \mathcal{F}^T} \left\{ \alpha \sum_{t=1}^{T} \max_{\mathbf{u}_t \in \mathcal{X}} f_t(\mathbf{u}_t) - \mathbb{E}\left[\sum_{t=1}^{T} f_t(\mathbf{x}_t) \right] \right\}$$

where $\sum_{t=1}^{T-1} \|\mathbf{u}_{t+1} - \mathbf{u}_t\| \leq P_T$

(What) So the paper does:

$$SR_T(\pi_{\mathbf{x}}, \alpha) := \sup_{(f_t)_{t=1}^T \in \mathcal{F}^T} \left\{ \alpha \max_{\mathbf{u} \in \mathcal{X}} \sum_{t=1}^T f_t(\mathbf{u}) - \mathbb{E}\left[\sum_{t=1}^T f_t(\mathbf{x}_t) \right] \right\}$$

$$DR_T(\pi_{\mathbf{x}}, P_T, \alpha) := \sup_{(f_t)_{t=1}^T \in \mathcal{F}^T} \left\{ \alpha \sum_{t=1}^T \max_{\mathbf{u}_t \in \mathcal{X}} f_t(\mathbf{u}_t) - \mathbb{E}\left[\sum_{t=1}^T f_t(\mathbf{x}_t) \right] \right\}$$

- **Full-Information**
- **Bandit**
- **Optimistic**: Some predictions about f_t are available to the learner.

(How) Concave Relaxation

Let $\mathcal{Y} \triangleq \text{conv}(\mathcal{X})$, $\Xi : \mathcal{Y} \to \mathcal{X}$ a randomized rounding, and let $\tilde{f} : \mathcal{Y} \to \mathbb{R}$ be a concave L-Lipschitz function such that:

$$\tilde{f}(\mathbf{x}) \geq f(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X}$$

and

$$\mathbb{E}_{\Xi}[f(\Xi(\mathbf{y}))] \geq \alpha \cdot \tilde{f}(\mathbf{y}), \forall \mathbf{y} \in \mathcal{Y}$$

Then since $\mathcal{Y}$ is convex and compact, and $\tilde{f}$ is concave and L-Lipschitz, we can run an OCO method (e.g., Online mirror Ascent (Descent)) and get a regret which it holds that:

$$SR_T(\pi_{\mathbf{x}}, \alpha) \leq \alpha \cdot SR_T(\pi_{\mathbf{y}})$$

and

$$DR_T(\pi_{\mathbf{x}}, P_T, \alpha) \leq \alpha \cdot DR_T(\pi_{\mathbf{y}}, P_T)$$

(How) Weighted Threshold Potential Functions

One such functions f where they are also equal to $\tilde{f}$, and can represent a wide variety of problems, are Weighted Threshold Potential (WTP) functions:

$$f(\mathbf{x}) \triangleq \sum_{i \in I} c_i \Psi(\mathbf{x}, S_i, \mathbf{w}_i, b_i), \forall \mathbf{x} \in \{0,1\}^n$$

- I: An arbitrary index set.
- $c_i \in \mathbb{R}_{\geq 0}$ for $i \in I$.
- $S_i \subseteq [n]$
- $b_i \in \mathbb{R}_{\geq 0} \cup \{\infty\}$: A threshold.
- $\mathbf{w}_i \in [0, b]^{|S_i|}$
- $\Psi(\mathbf{x}, S, \mathbf{w}, b) = \min \left\{ b, \sum_{j \in S} x_j w_j \right\}$

(How) Randomized Swap Rounding

- We found f and $\tilde{f}$; if we find $\Xi : \mathcal{Y} \to \mathcal{X}$, then we are done.

$$\tilde{f}(\mathbf{x}) \geq f(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X}$$

$$\mathbb{E}_{\Xi}[f(\Xi(\mathbf{y}))] \geq \alpha \cdot \tilde{f}(\mathbf{y}), \forall \mathbf{y} \in \mathcal{Y}$$

(How) Randomized Swap Rounding

- We found f and $\tilde{f}$; if we find $\Xi : \mathcal{Y} \to \mathcal{X}$, then we are done.
- Randomized Swap Rounding[2] (RSR) on (WTP) would result in:

$$\alpha = 1 - \frac{1}{e}$$

$$\tilde{f}(\mathbf{x}) \geq f(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X}$$

$$\mathbb{E}_{\Xi}[f(\Xi(\mathbf{y}))] \geq \alpha \cdot \tilde{f}(\mathbf{y}), \forall \mathbf{y} \in \mathcal{Y}$$

[2] Chekuri, Vondrák, and Zenklusen, "Dependent randomized rounding via exchange properties of combinatorial structures".

(How) Summary

1. Construct the convex hull of $\mathcal{X} = \{0, 1\}^n$.
   - $\mathcal{Y} \triangleq \text{conv}(\mathcal{X})$
2. Run an OCO method on $\mathcal{Y}$. *(Yet to be discussed)*
3. Convert the fractional solution obtained on $\mathcal{Y}$ to an integral solution $\in \mathcal{X}$ via RSR with $\alpha = 1 - \frac{1}{e}$ approximation.
4. Enjoy the regret bounds.

$$SR_T(\pi_{\mathbf{x}}, \alpha) \leq \alpha \cdot SR_T(\pi_{\mathbf{y}})$$

$$DR_T(\pi_{\mathbf{x}}, P_T, \alpha) \leq \alpha \cdot DR_T(\pi_{\mathbf{y}}, P_T)$$

(How) Optimistic Online Mirror Ascent

Algorithm 1: OOMA: Two-Update-Per-Step

Require: $\eta \in \mathbb{R}_{\geq 0}$ `/* learning rate */`

$\Phi : \mathcal{Y} \to \mathbb{R}$ `/* mirror map */`

$M_2, M_3, \cdots, M_{T+1}$ `/* Sequence of predictions */`

1: Let $\mathbf{y}_1 = \mathbf{z}_1 = \arg\max_{\mathbf{y} \in \mathcal{Y}} \Phi(\mathbf{y})$

2: **for** $t = 1$ to T **do**

3: Play $\mathbf{y}_t$.

4: Observe the reward function $\tilde{f}_t$ and let $\nabla_t = \nabla \tilde{f}_t(\mathbf{y}_t)$

5: $\mathbf{z}_{t+1} = \arg\max_{\mathbf{z} \in \mathcal{Y}} \langle \mathbf{z}, \nabla_t \rangle - \frac{B_\Phi(\mathbf{z}; \mathbf{z}_t)}{\eta}$ `/* Adapt the secondary decision */`

6: $\mathbf{y}_{t+1} = \arg\max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, M_{t+1} \rangle - \frac{B_\Phi(\mathbf{y}; \mathbf{z}_{t+1})}{\eta}$ `/* Adapt the primary decision */`

7: **end for**

(Why) Because OSM via OCO is the best :)

Paper	Prob. Class	$(1-1/e)$-regret (Full Information)									Time
		Static			**Dynamic**			**Optimistic**			
		Uni.	**Part.**	**Gen.**	**Uni.**	**Part.**	**Gen.**	**Uni.**	**Part.**	**Gen.**	
[4]	GS	$r\sqrt{\log(n)\,T}$	×	×	×	×	×	×	×	×	$T^4 O_b$
[2]	GS	$\sqrt{r\log\left(\frac{n}{r}\right)T}$			×	×	×	×	×	×	$nr^2 + O_m \cdot n^4/\epsilon^3$ $\log(n^3T/\epsilon)$
[3]	GS	$r\sqrt{r\log(nT)T}$	×	×	$\sqrt{r(r\log(nT)+P_T)T}$	×	×	×	×	×	nr
[1]	GS	$r^{\frac{3}{2}}\sqrt{\log(n)T}$		×	×	×	×	×	×	×	$n^2 c_p$
[30]	DR-S	$\sqrt{rnT}$			×	×	×	×	×	×	$\sqrt{T}O_{oco} \cdot O_m + nr^2$
[31]	DR-S	$T^{\frac{4}{5}}$			×	×	×	×	×	×	$T^{\frac{3}{5}}O_{oco} \cdot O_m + nr^2$
[32]	DR-S	$\sqrt{rnT}$			×	×	×	×	×	×	$O_{oco} \cdot O_m + nr^2$
[33]	LWD	$n(\alpha+2)\sqrt{T}$			×	×	×	×	×	×	TO_α
This work	WTP	$r\sqrt{\log(\frac{n}{r})T}$			$\sqrt{r(r\log(\frac{n}{r})+\log(n)P_T)T}$			$\sqrt{r(r\log(\frac{n}{r})+\log(n)P_T} \cdot \sqrt{\sum_{t=1}^{T}\lVert\mathbf{g}_t-\mathbf{g}_t^{\pi}\rVert_\infty^2}$			nr^2

Comparison to previous results

(Proposal 1) OOMA used in the paper was old!

Algorithm 2: OOMA: Two-Update-Per-Step

Require: $\eta \in \mathbb{R}_{\geq 0}$ `/* learning rate */`

$\Phi : \mathcal{Y} \to \mathbb{R}$ `/* mirror map */`

$M_2, M_3, \cdots, M_{T+1}$ `/* Sequence of predictions */`

Let $\mathbf{y}_1 = \mathbf{z}_1 = \arg\max_{\mathbf{y} \in \mathcal{Y}} \Phi(\mathbf{y})$

for $t = 1$ to T **do**

  Play $\mathbf{y}_t$.

  Observe the reward function $\tilde{f}_t$ and let $\nabla_t = \nabla \tilde{f}_t(\mathbf{y}_t)$

  $\mathbf{z}_{t+1} = \arg\max_{\mathbf{z} \in \mathcal{Y}} \langle \mathbf{z}, \nabla_t \rangle - \frac{B_\Phi(\mathbf{z}; \mathbf{z}_t)}{\eta}$ `/* Adapt the secondary decision */`

  $\mathbf{y}_{t+1} = \arg\max_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, M_{t+1} \rangle - \frac{B_\Phi(\mathbf{y}; \mathbf{z}_{t+1})}{\eta}$ `/* Adapt the primary decision */`

end for

(Proposal 1) OOMA was based on this paper

Online Learning with Predictable Sequences

Alexander Rakhlin
University of Pennsylvania

Karthik Sridharan
University of Pennsylvania

May 27, 2014

(Proposal 1) OOMA was improved later!

Proceedings of Machine Learning Research 76:1–40, 2017 Algorithmic Learning Theory 2017

A Modular Analysis of Adaptive (Non-)Convex Optimization: Optimism, Composite Objectives, and Variational Bounds

Pooria Joulani POORIA@UALBERTA.CA
Department of Computing Science
University of Alberta
Edmonton, Alberta, Canada

András György A.GYORGY@IMPERIAL.AC.UK
Department of Electrical and Electronic Engineering
Imperial College London
London, UK

Csaba Szepesvári SZEPESVA@UALBERTA.CA
Department of Computing Science
University of Alberta
Edmonton, Alberta, Canada

Algorithm 3: OOMA: One-Update-Per-Step

Require: $\eta \in \mathbb{R}_{\geq 0}$ `/* learning rate */`
$\Phi : \mathcal{Y} \to \mathbb{R}$ `/* mirror map */`
$M_2, M_3, \cdots, M_{T+1}$ `/* Sequence of predictions */`

1: Let $\mathbf{y}_1 = \mathbf{z}_1 = \arg\max_{\mathbf{y}\in\mathcal{Y}} \Phi(\mathbf{y})$
2: **for** $t = 1$ to T **do**
3: Play $\mathbf{y}_t$.
4: Observe the reward function $\tilde{f}_t$ and let $\nabla_t = \nabla \tilde{f}_t(\mathbf{y}_t)$
5: $\mathbf{y}_{t+1} = \arg\max_{\mathbf{y}\in\mathcal{Y}} \langle \mathbf{y}, \nabla_t - M_t + M_{t+1} \rangle - \frac{B_\Phi(\mathbf{y};\mathbf{y}_t)}{\eta}$
6: **end for**

(Proposal 1) I proved the new OOMA's dynamic regret

Matching the new algorithm's regret bounds with the one used in the paper.

$$DR_T(\pi_{\mathbf{y}}, P_T) \leq \frac{\eta}{2\rho} \sum_{t=1}^{T} \|\nabla_t - M_t\|_*^2 + \frac{(D_\Phi^2 + 2G_\Phi P_T)}{\eta}$$

And when the optimal learning rate $\eta^* = \sqrt{\frac{2\rho(D_\Phi^2 + 2G_\Phi P_T)}{\sum_{t=1}^{T} \|\nabla_t - M_t\|_*^2}}$ is selected:

$$DR_T(\pi_{\mathbf{y}}, P_T) \leq \sqrt{\frac{2(D_\Phi^2 + 2G_\Phi P_T)}{\rho} \sum_{t=1}^{T} \|\nabla_t - M_t\|_*^2}$$

(Proposal 2) Contradiction of the optimal learning rate

- The optimal learning rate $\eta^* = \sqrt{\frac{2\rho(D_\Phi^2 + 2G_\Phi P_T)}{\sum_{t=1}^{T} \|\nabla_t - M_t\|_*^2}}$ is dependent on the knowledge of $\sum_{t=1}^{T-1} \|\mathbf{u}_{t+1} - \mathbf{u}_t\| \leq P_T$.
- This is a violation of the **adversarial assumption**, as the learner knows how much budget the adversary has to change $\mathbf{u_t}$s!

Adaptive Online Learning in Dynamic Environments

Lijun Zhang, Shiyin Lu, Zhi-Hua Zhou
National Key Laboratory for Novel Software Technology
Nanjing University, Nanjing 210023, China
{zhanglj, lusy, zhouzh}@lamda.nju.edu.cn

(Proposal 2) Extension to OOMA

Design an algorithm for Optimistic Online Mirror Ascent without the knowledge of the path length and match the previous regret bounds up to polylogarithmic factors.

(Proposal 2) Extension to OOMA

Algorithm 4: Optimistic Online Mirror Ascent: One-Update-Per-Step with Adaptive learning Rate

Require: $\alpha \in \mathbb{R}_{\geq 0}$ `/* learning rate */`
$\Phi : \mathcal{Y} \to \mathbb{R}$ `/* mirror map */`
$M_2, M_3, \cdots, M_{T+1}$ `/* Sequence of predictions */`

Let $N = 1 + \lceil \log_2 \sqrt{4 + 8T} \rceil$
Let $\eta^{(i)} = \frac{\sqrt{2\rho} D_\Phi 2^{i-1}}{2L\sqrt{T}}, i = 1, \cdots, N$
Let $\mathbf{p}_1 = [\frac{1}{N}, \cdots, \frac{1}{N}] \in [0, 1]^N$ `/* Uniform prior over experts */`
Let $\mathbf{y}_1 = \text{argmax}_{\mathbf{y} \in \mathcal{Y}} \Phi(\mathbf{y})$
for $t = 1$ to T **do**
  Play $\mathbf{y}_t$.
  Observe the reward function $\tilde{f}_t$ and let $\nabla_t = \nabla \tilde{f}_t(\mathbf{y}_t)$
  Update each expert: $\mathbf{y}_{t+1}^{(i)} = \text{argmax}_{\mathbf{y} \in \mathcal{Y}} \langle \mathbf{y}, \nabla_t - M_t + M_{t+1} \rangle - \frac{B_\Phi(\mathbf{y}; \mathbf{y}_t)}{\eta^{(i)}}, i = 1, \cdots, N$
  Update each weight: $\mathbf{p}_{t+1}^{(i)} = \frac{\mathbf{p}_t^{(i)} \exp\left(\alpha \tilde{f}_t\left(\mathbf{y}_t^{(i)}\right)\right)}{\sum_{j=1}^N \mathbf{p}_t^{(j)} \exp\left(\alpha \tilde{f}_t\left(\mathbf{y}_t^{(j)}\right)\right)}, i = 1, \cdots, N$
  $\mathbf{y}_{t+1} = \sum_{i=1}^N \mathbf{p}_{t+1}^{(i)} \mathbf{y}_{t+1}^{(i)}$
end for

Thank You! *:)*