Fisher-Rao Gradient Descent for Stochastic Control Problems

Lukasz Szpruch
University of Edinburgh, The Alan Turing Institute, London

joint work with B.Kerimkulov and D.Siska (UoE)

Outline

- Stochastic Control and Pontryagin's Maximum Principle.
- Methods of successive approximation.
- Optimisation on $\mathscr{P}(\mathbb{R}^d)$.
- Convergence of Gradient Descent for regularised stochastic control problem.

Stochastic Control

Stochastic Control

For the initial state $x \in \mathbb{R}^d$ and the admissible control $\pi \in \mathscr{A}$ controlled dynamics

$$dX_s(\pi) = b_s(X_s(\pi), \pi_s)\, ds + \sigma_s(X_s(\pi), \pi_s)\, dW_s\,,\ s \in [0, T]\,,\ X_0(\pi) = x\,.$$

Objective:

$$J^\tau(\pi) := \mathbb{E}\left[\int_0^T [f_s(X_s(\pi), \pi_s) + \tau\, h(\pi_s)]\, ds + g(X_T(\pi))\right]$$

where $h : \mathscr{P}(\mathbb{R}^p) \to \mathbb{R} \cup \{\infty\}$ is a convex function and

$$\mathscr{A} := \Big\{ \pi : \Omega \to \mathscr{M}(E) : \pi_t \in \mathscr{P}(\mathbb{R}^p),\, \pi_t \text{ is } \mathscr{F}_t\text{-measurable } \forall t \in [0, T]\,,$$
$$\text{and } \pi(da, dt) = \pi_t(a)\, da\, dt \text{ for a.a. } t \in [0, T] \Big\}\,.$$

Stochastic Control

For the initial state $x \in \mathbb{R}^d$ and the admissible control $\pi \in \mathscr{A}$ controlled dynamics

$$dX_s(\pi) = b_s(X_s(\pi), \pi_s)\, ds + \sigma_s(X_s(\pi), \pi_s)\, dW_s\,,\ s \in [0, T]\,,\ X_0(\pi) = x\,.$$

Objective:

$$J^\tau(\pi) := \mathbb{E}\left[\int_0^T [f_s(X_s(\pi), \pi_s) + \tau\, h(\pi_s)]\, ds + g(X_T(\pi))\right]$$

where $h : \mathscr{P}(\mathbb{R}^p) \to \mathbb{R} \cup \{\infty\}$ is a convex function and

$$\mathscr{A} := \Big\{ \pi : \Omega \to \mathscr{M}(E) : \pi_t \in \mathscr{P}(\mathbb{R}^p),\, \pi_t \text{ is } \mathscr{F}_t\text{-measurable } \forall t \in [0, T]\,,$$
$$\text{and } \pi(da, dt) = \pi_t(a)\, da\, dt \text{ for a.a. } t \in [0, T] \Big\}\,.$$

Aim: minimise J over $\mathscr{A}$ (open loop controls)

Stochastic Control

For the initial state $x \in \mathbb{R}^d$ and the admissible control $\pi \in \mathscr{A}$ controlled dynamics

$$dX_s(\pi) = b_s(X_s(\pi), \pi_s)\, ds + \sigma_s(X_s(\pi), \pi_s)\, dW_s\,,\ s \in [0,T]\,,\ X_0(\pi) = x\,.$$

Objective:

$$J^\tau(\pi) := \mathbb{E}\left[\int_0^T [f_s(X_s(\pi), \pi_s) + \tau\, h(\pi_s)]\, ds + g(X_T(\pi))\right]$$

where $h : \mathscr{P}(\mathbb{R}^p) \to \mathbb{R} \cup \{\infty\}$ is a convex function and

$$\mathscr{A} := \Big\{ \pi : \Omega \to \mathscr{M}(E) : \pi_t \in \mathscr{P}(\mathbb{R}^p),\ \pi_t \text{ is } \mathscr{F}_t\text{-measurable } \forall t \in [0,T],$$
$$\text{and } \pi(da, dt) = \pi_t(a)\, da\, dt \text{ for a.a. } t \in [0,T] \Big\}.$$

Aim: minimise J over $\mathscr{A}$ (open loop controls)

Example 1 (Relaxed Control)

$$b_t(x,\pi) = \int \bar{b}_t(x,a)\pi(da)\,,\ \text{and}\ \ \sigma_t(x,\pi)(\sigma_t(x,\pi))^\top = \int \bar{\sigma}_t(x,a)\bar{\sigma}_t(x,a)^\top \pi(da)$$

Pontryagin's Maximum Principle

Define Hamiltonian

$$H_t^\tau(x,y,z,m) := b_t(x,m)\cdot y + \mathrm{tr}(\sigma_t^\top(x,m)z) + f_t(x,m) + \tau\, h(m)\,.$$

Pontryagin's Maximum Principle

Define Hamiltonian

$$H_t^\tau(x,y,z,m) := b_t(x,m)\cdot y + \mathrm{tr}(\sigma_t^\top(x,m)z) + f_t(x,m) + \tau\, h(m)\,.$$

For any $\pi,\pi' \in \mathscr{A}$ we have that

$$\frac{d}{d\varepsilon} J^\tau(\pi + \varepsilon(\pi' - \pi))\big|_{\varepsilon=0} = \mathbb{E}\int_0^T \left[\int \frac{\delta H_s^\tau}{\delta m}(X_s(\pi), Y_s(\pi), Z_s(\pi), \pi_s, a)(\pi_s' - \pi_s)(a)\, da\right] ds\,.$$

Pontryagin's Maximum Principle

Define Hamiltonian

$$H_t^\tau(x,y,z,m) := b_t(x,m)\cdot y + \mathrm{tr}(\sigma_t^\top(x,m)z) + f_t(x,m) + \tau\, h(m)\,.$$

For any $\pi, \pi' \in \mathscr{A}$ we have that

$$\frac{d}{d\varepsilon} J^\tau(\pi + \varepsilon(\pi' - \pi))\big|_{\varepsilon=0} = \mathbb{E}\int_0^T \left[\int \frac{\delta H_s^\tau}{\delta m}(X_s(\pi), Y_s(\pi), Z_s(\pi), \pi_s, a)(\pi_s' - \pi_s)(a)\, da\right] ds\,.$$

Necessary condition: Let $\pi^* \in \mathscr{A}$ be an (locally) optimal control and $(X(\pi^*), Y(\pi^*), Z(\pi^*))$ be the associated forward-backward system

$$\begin{aligned}
&dX_s(\pi^*) = b_s(X_s(\pi^*), \pi_s^*)\, ds + \sigma_s(X_s(\pi^*), \pi_s^*)\, dW_s\,,\ X_0(\pi^*) = x\,.\\
&dY_s(\pi^*) = -(D_x H_s^0)(X_s(\pi^*), Y_s(\pi^*), Z_s(\pi^*), \pi_s^*)\, ds + Z_s(\pi^*)\, dW_s\,,\\
&Y_T(\pi^*) = (D_x g)(X_T(\pi^*)),\ .
\end{aligned}$$

then

$$\pi_t^* \in \operatorname*{argmin}_{\pi\in\mathscr{A}} H_t^\tau(X_t(\pi^*), Y_t(\pi^*), Z_t(\pi^*), \pi)\,, \quad t \in [0,T]\,,\ \text{pointwise optimisation!}$$

Pontryagin's Maximum Principle

Define Hamiltonian

$$H^\tau_t(x,y,z,m) := b_t(x,m)\cdot y + \mathrm{tr}(\sigma_t^\top(x,m)z) + f_t(x,m) + \tau\, h(m)\,.$$

For any $\pi, \pi' \in \mathscr{A}$ we have that

$$\frac{d}{d\varepsilon} J^\tau(\pi + \varepsilon(\pi' - \pi))\big|_{\varepsilon=0} = \mathbb{E}\int_0^T \left[\int \frac{\delta H^\tau_s}{\delta m}(X_s(\pi), Y_s(\pi), Z_s(\pi), \pi_s, a)(\pi'_s - \pi_s)(a)\, da\right] ds\,.$$

Necessary condition: Let $\pi^* \in \mathscr{A}$ be an (locally) optimal control and $(X(\pi^*), Y(\pi^*), Z(\pi^*)$ be the associated forward-backward system

$$dX_s(\pi^*) = b_s(X_s(\pi^*), \pi^*_s)\, ds + \sigma_s(X_s(\pi^*), \pi^*_s)\, dW_s\,,\ X_0(\pi^*) = x\,.$$
$$dY_s(\pi^*) = -(D_x H^0_s)(X_s(\pi^*), Y_s(\pi^*), Z_s(\pi^*), \pi^*_s)\, ds + Z_s(\pi^*)\, dW_s\,,$$
$$Y_T(\pi^*) = (D_x g)(X_T(\pi^*)),\ .$$

then

$$\pi^*_t \in \operatorname*{argmin}_{\pi\in\mathscr{A}} H^\tau_t(X_t(\pi^*), Y_t(\pi^*), Z_t(\pi^*), \pi)\,, \quad t \in [0,T]\,,$$ pointwise optimisation!

Sufficient condition: If g is convex and H^τ is convex in (x,m) and π^* and $(X(\pi^*), Y(\pi^*), Z(\pi^*))$ are as above then $J^\tau(\pi^\star) = \inf_{\pi\in\mathscr{A}} J^\tau(\pi)$

Solving Stochastic Control

- DPP and HJB equation
  - Solve nonlinear PDE or corresponding (2)BSDE
  - 'Linearise' with policy or value iteration and solve linear PDE [Kerimkulov et al., 2020, Gobet and Labart, 2010]

Solving Stochastic Control

- DPP and HJB equation
  - Solve nonlinear PDE or corresponding (2)BSDE
  - 'Linearise' with policy or value iteration and solve linear PDE [Kerimkulov et al., 2020, Gobet and Labart, 2010]
- Maximum principle
  - Solve the corresponding (non-Markov) FBSDE
  - 'Linearise' with policy and solve a sequence of decoupled FBSDEs: Modified Successive Approximation (MSA) [Kerimkulov et al., 2021]

Solving Stochastic Control

- DPP and HJB equation
  - Solve nonlinear PDE or corresponding (2)BSDE
  - 'Linearise' with policy or value iteration and solve linear PDE [Kerimkulov et al., 2020, Gobet and Labart, 2010]
- Maximum principle
  - Solve the corresponding (non-Markov) FBSDE
  - 'Linearise' with policy and solve a sequence of decoupled FBSDEs: Modified Successive Approximation (MSA) [Kerimkulov et al., 2021]

Other works on MSA

- [Krylov and Chernousko, 1963]: divergent without regularisation, convergence for controlled ODEs, rate in special cases
- [Li et al., 2018]: connection to training of deep recurrent neural nets
- [Reisinger et al., 2021]: MSA-like algorithm but BSDE solved via PDE-based method giving Markov controls

MSA I: Bregman Proximal Descent

Initialization: a guess $\pi^0 = (\pi^0_s)_{s\in[0,T]}$ and fix $\lambda > 0$.
while difference between $J^0(\pi^{n+1})$ and $J^0(\pi^n)$ is large **do**
Given a control $\pi^n = (\pi^n_s)_{s\in[0,T]}$ solve the following forward SDE, then solve backward SDE $\Theta_s(\pi) := (X_s(\pi), Y_s(\pi), Z_s(\pi))$ for $s \in [0,T]$:

$$dX_s(\pi^n) = b_s(X_s(\pi^n), \pi^n_s)\,ds + \sigma_s(X_s(\pi^n), \pi^n_s)\,dW_s\,,\;\; X_0(\pi^n) = x\,,$$
$$dY_s(\pi^n) = -(D_x H^0_s)(\Theta_s(\pi^n), \pi^n_s)\,ds + Z_s(\pi^n)\,dW_s\,,\; Y_T(\pi^n) = (D_x g)(X_T(\pi^n))\,.$$

Update the control $\forall s \in [0,T]$ as

$$\pi^{n+1}_s \in \arg\min_{\pi\in\mathscr{P}(\mathbb{R}^p)} H^\tau_s(\Theta_s(\pi^n), \pi) + \lambda\, D_h(\pi|\pi^n_s)\,,$$

end while
return π^{n+1}.

MSA II: Mirror Descent

Initialization: a guess $\pi^0 = (\pi^0_s)_{s\in[0,T]}$ and fix $\lambda > 0$.

while difference between $J^0(\pi^{n+1})$ and $J^0(\pi^n)$ is large **do**

Given a control $\pi^n = (\pi^n_s)_{s\in[0,T]}$ solve the following forward SDE, then solve backward SDE $\Theta_s(\pi) := (X_s(\pi), Y_s(\pi), Z_s(\pi))$ for $s \in [0,T]$:

$$dX_s(\pi^n) = b_s(X_s(\pi^n), \pi^n_s)\,ds + \sigma_s(X_s(\pi^n), \pi^n_s)\,dW_s\,,\;\; X_0(\pi^n) = x\,,$$
$$dY_s(\pi^n) = -(D_x H^0_s)(\Theta_s(\pi^n), \pi^n_s)\,ds + Z_s(\pi^n)\,dW_s\,,\; Y_T(\pi^n) = (D_x g)(X_T(\pi^n))\,.$$

Update the control $\forall s \in [0,T]$ as

$$\pi^{n+1}_s \in \arg\min_{\pi\in\mathscr{P}(\mathbb{R}^p)} \int \frac{\delta H^\tau_s}{\delta m}(\Theta_s(\pi^n), \pi^n_s, a)(\pi - \pi^n_s)(da) + \lambda\, D_h(\pi|\pi^n_s)\,.$$

end while

return π^{n+1}.

Optimisation on $\mathscr{P}(\mathscr{X})$

Preliminaries

Let $h : \mathscr{P}(\mathscr{X}) \to \mathbb{R}$. Assume that $h \in \mathscr{C}^1$ (i.e differentiable in a sense of linear functional derivative). We say that $\mu \mapsto h(\mu)$ is convex if for all $\mu, \mu' \in \mathscr{P}$

$$h(\mu') - h(\mu) - \int \frac{\delta h}{\delta \mu}(\mu, y)(\mu' - \mu)(dy) \geq 0 .$$

Preliminaries

Let $h : \mathscr{P}(\mathscr{X}) \to \mathbb{R}$. Assume that $h \in \mathscr{C}^1$ (i.e differentiable in a sense of linear functional derivative). We say that $\mu \mapsto h(\mu)$ is convex if for all $\mu, \mu' \in \mathscr{P}$

$$h(\mu') - h(\mu) - \int \frac{\delta h}{\delta \mu}(\mu, y)(\mu' - \mu)(dy) \geq 0 .$$

Let $h \in \mathscr{C}^1$ be convex. For all $\mu, \mu' \in \mathscr{P}$ define h-Bregman divergence to be

$$D_h(\mu', \mu) = h(\mu') - h(\mu) - \int \frac{\delta h}{\delta \mu}(\mu, y)(\mu' - \mu)(dy) .$$

Properties of h-Bregman Divergence

P.1. $D_h(\mu, \mu') \geq 0$ for all $\mu, \mu' \in \mathscr{P}(\mathscr{X})$. When h is strictly convex $D_h(\mu, \mu') = 0 \implies \mu' = \mu$.

Properties of h-Bregman Divergence

P.1. $D_h(\mu,\mu') \geq 0$ for all $\mu,\mu' \in \mathscr{P}(\mathscr{X})$. When h is strictly convex $D_h(\mu,\mu') = 0 \implies \mu' = \mu$.

P.2. For any $\mu' \in \mathscr{P}(\mathscr{X})$ we have that $\mu \mapsto D_h(\mu,\mu')$ and

$$\frac{\delta D_h}{\delta \mu}(\mu,\nu,y) = \frac{\delta h}{\delta \mu}(\mu,y) - \frac{\delta h}{\delta \mu}(\nu,y).$$

Properties of h-Bregman Divergence

P.1. $D_h(\mu,\mu') \geq 0$ for all $\mu,\mu' \in \mathscr{P}(\mathscr{X})$. When h is strictly convex $D_h(\mu,\mu') = 0 \implies \mu' = \mu$.

P.2. For any $\mu' \in \mathscr{P}(\mathscr{X})$ we have that $\mu \mapsto D_h(\mu,\mu')$ and

$$\frac{\delta D_h}{\delta \mu}(\mu,\nu,y) = \frac{\delta h}{\delta \mu}(\mu,y) - \frac{\delta h}{\delta \mu}(\nu,y).$$

P.3. h convex (strictly convex, strongly convex) $\implies$ $\mu \mapsto D_h(\mu,\nu)$ convex (strictly convex, strongly convex)

Properties of h-Bregman Divergence

P.1. $D_h(\mu,\mu') \geq 0$ for all $\mu,\mu' \in \mathscr{P}(\mathscr{X})$. When h is strictly convex $D_h(\mu,\mu') = 0 \implies \mu' = \mu$.

P.2. For any $\mu' \in \mathscr{P}(\mathscr{X})$ we have that $\mu \mapsto D_h(\mu,\mu')$ and

$$\frac{\delta D_h}{\delta \mu}(\mu,\nu,y) = \frac{\delta h}{\delta \mu}(\mu,y) - \frac{\delta h}{\delta \mu}(\nu,y).$$

P.3. h convex (strictly convex, strongly convex) $\implies \mu \mapsto D_h(\mu,\nu)$ convex (strictly convex, strongly convex)

P.4. Let $h,g \in \mathscr{C}^1$ be convex and $\alpha,\beta \in \mathscr{R}$. Then

$$D_{\alpha h+\beta g}(\mu,\nu) = \alpha D_h(\mu,\nu) + \beta D_g(\mu,\nu).$$

Setup

Let $F : \mathscr{P}(\mathscr{X}) \to \mathbb{R}$. Consider the following optimisation problem

$$\min_{\mu} F(\mu)$$

under the assumption that F is lower bounded i.e. $\min_\mu F(\mu) > -\infty$.

Setup

Let $F : \mathscr{P}(\mathscr{X}) \to \mathbb{R}$. Consider the following optimisation problem

$$\min_{\mu} F(\mu)$$

under the assumption that F is lower bounded i.e. $\min_{\mu} F(\mu) > -\infty$.

We say that function F is *L-relatively smooth with respect to h* if

$$D_F(\mu,\mu') = F(\mu) - F(\mu') - \int \frac{\delta F}{\delta \mu}(\mu', y)(\mu - \mu')(dy) \leq L\, D_h(\mu,\mu').$$

Setup

Let $F : \mathscr{P}(\mathscr{X}) \to \mathbb{R}$. Consider the following optimisation problem

$$\min_{\mu} F(\mu)$$

under the assumption that F is lower bounded i.e. $\min_{\mu} F(\mu) > -\infty$.

We say that function F is L-relatively smooth with respect to h if

$$D_F(\mu,\mu') = F(\mu) - F(\mu') - \int \frac{\delta F}{\delta \mu}(\mu',y)(\mu-\mu')(dy) \leq L\, D_h(\mu,\mu').$$

We say that function F is l-strongly convex relative to h if

$$D_F(\mu,\mu') = F(\mu) - F(\mu') - \int \frac{\delta F}{\delta \mu}(\mu',y)(\mu-\mu')(dy) \geq l\, D_h(\mu,\mu').$$

Setup

Let $F : \mathscr{P}(\mathscr{X}) \to \mathbb{R}$. Consider the following optimisation problem

$$\min_{\mu} F(\mu)$$

under the assumption that F is lower bounded i.e. $\min_{\mu} F(\mu) > -\infty$.

We say that function F is L-relatively smooth with respect to h if

$$D_F(\mu,\mu') = F(\mu) - F(\mu') - \int \frac{\delta F}{\delta \mu}(\mu',y)(\mu-\mu')(dy) \leq L\, D_h(\mu,\mu').$$

We say that function F is l-strongly convex relative to h if

$$D_F(\mu,\mu') = F(\mu) - F(\mu') - \int \frac{\delta F}{\delta \mu}(\mu',y)(\mu-\mu')(dy) \geq l\, D_h(\mu,\mu').$$

- [Aubin-Frankowski et al.,]: mirror descent of measure space with relative smoothness and convexity
- [Lu et al., 2018]: mirror descent of $\mathbb{R}^d$ with relative smoothness and convexity

Mirror Descent

Mirror Descent is defined as

$$\mu^{n+1} \in \operatorname*{argmin}_{\mu} \left\{ \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu - \mu^n)(dy) + L D_h(\mu, \mu^n) \right\}.$$

When h is strictly convex, then the minimiser is unique.

Mirror Descent

Mirror Descent is defined as

$$\mu^{n+1} \in \underset{\mu}{\operatorname{argmin}} \left\{ \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu - \mu^n)(dy) + L D_h(\mu, \mu^n) \right\}.$$

When h is strictly convex, then the minimiser is unique.

Remark: One can replace L with any $\lambda \geq L$.

Towards gradient flow - a digression

The first-order condition yields that

$$\frac{\delta F}{\delta \mu}(\mu^n, y) + L\frac{\delta D_h}{\delta \mu'}(\mu^{n+1}, \mu^n, y) = \text{constant}.$$

Hence the flow is (implicitly) defined us

$$\frac{\delta h}{\delta \mu}(\mu^{n+1}, y) - \frac{\delta h}{\delta \mu}(\mu^n, y) = -\frac{1}{L}\left(\frac{\delta F}{\delta \mu}(\mu^n, y) + constant\right)$$

Towards gradient flow - a digression

The first-order condition yields that

$$\frac{\delta F}{\delta \mu}(\mu^n, y) + L\frac{\delta D_h}{\delta \mu'}(\mu^{n+1}, \mu^n, y) = \text{constant}.$$

Hence the flow is (implicitly) defined us

$$\frac{\delta h}{\delta \mu}(\mu^{n+1}, y) - \frac{\delta h}{\delta \mu}(\mu^n, y) = -\frac{1}{L}\left(\frac{\delta F}{\delta \mu}(\mu^n, y) + constant\right)$$

Formally passing to the limit $L \to \infty$ gives (up to a constant)

$$\frac{d}{dt}\frac{\delta h}{\delta \mu}(\mu_t, y) = -\frac{\delta F}{\delta \mu}(\mu_t, y).$$

Towards gradient flow - a digression

The first-order condition yields that

$$\frac{\delta F}{\delta \mu}(\mu^n, y) + L\frac{\delta D_h}{\delta \mu'}(\mu^{n+1}, \mu^n, y) = \text{constant}.$$

Hence the flow is (implicitly) defined us

$$\frac{\delta h}{\delta \mu}(\mu^{n+1}, y) - \frac{\delta h}{\delta \mu}(\mu^n, y) = -\frac{1}{L}\left(\frac{\delta F}{\delta \mu}(\mu^n, y) + constant\right)$$

Formally passing to the limit $L \to \infty$ gives (up to a constant)

$$\frac{d}{dt}\frac{\delta h}{\delta \mu}(\mu_t, y) = -\frac{\delta F}{\delta \mu}(\mu_t, y).$$

Taking $h(\mu) = \int \log \mu(x)\mu(dx)$, formally $D_h(\mu, \nu) = \mathrm{KL}(\mu, \nu)$, the flow becomes

$$\frac{d}{dt}\log \mu_t(y) = -\frac{\delta F}{\delta \mu}(\mu_t, y).$$

or equivalently

$$\frac{d}{dt}\mu_t(y) = -\frac{\delta F}{\delta \mu}(\mu_t, y)\mu_t(y),$$

which is a birth-death (or Fisher-Rao) gradient flow.

Towards gradient flow - a digression

The first-order condition yields that

$$\frac{\delta F}{\delta \mu}(\mu^n, y) + L\frac{\delta D_h}{\delta \mu'}(\mu^{n+1}, \mu^n, y) = \text{constant}.$$

Hence the flow is (implicitly) defined us

$$\frac{\delta h}{\delta \mu}(\mu^{n+1}, y) - \frac{\delta h}{\delta \mu}(\mu^n, y) = -\frac{1}{L}\left(\frac{\delta F}{\delta \mu}(\mu^n, y) + constant\right)$$

Formally passing to the limit $L \to \infty$ gives (up to a constant)

$$\frac{d}{dt}\frac{\delta h}{\delta \mu}(\mu_t, y) = -\frac{\delta F}{\delta \mu}(\mu_t, y).$$

Taking $h(\mu) = \int \log \mu(x)\mu(dx)$, formally $D_h(\mu, \nu) = \mathrm{KL}(\mu, \nu)$, the flow becomes

$$\frac{d}{dt}\log \mu_t(y) = -\frac{\delta F}{\delta \mu}(\mu_t, y).$$

or equivalently

$$\frac{d}{dt}\mu_t(y) = -\frac{\delta F}{\delta \mu}(\mu_t, y)\mu_t(y),$$

which is a birth-death (or Fisher-Rao) gradient flow.

- [Liu et al., 2023]: established exponential convergence via Polyak-Łojasiewicz condition
- [Liero et al., 2022], [Chizat et al., 2018]: General Theory

Proximal Bregman Lemma

Lemma 2 (Three point lemma / Bregman proximal inequality)

{lem

Let $G : \mathscr{P} \to \mathbb{R}$ be convex. Define, for all $\mu \in \mathscr{P}$

$$\bar{\nu} = \operatorname*{argmin}_{\nu} \{ G(\nu) + D_h(\nu, \mu) \} .$$

Then

$$G(\nu) + D_h(\nu, \mu) \geq G(\bar{\nu}) + D_h(\bar{\nu}, \mu) + D_h(\nu, \bar{\nu}) .$$

Proximal Bregman Lemma

Lemma 2 (Three point lemma / Bregman proximal inequality)

Let $G : \mathscr{P} \to \mathbb{R}$ be convex. Define, for all $\mu \in \mathscr{P}$

$$\bar{\nu} = \underset{\nu}{\operatorname{argmin}} \{ G(\nu) + D_h(\nu, \mu) \} .$$

Then

$$G(\nu) + D_h(\nu, \mu) \geq G(\bar{\nu}) + D_h(\bar{\nu}, \mu) + D_h(\nu, \bar{\nu}) .$$

Consider a convex function

$$\nu \mapsto G(\nu, \mu^n) := \frac{1}{L} \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\nu - \mu^n)(dy) .$$

Lemma 2 with $\mu = \mu^n$ and $\bar{\nu} = \mu^{n+1}$ (due definition of $(\mu^n)_n$), for any ν

$$\begin{aligned} &D_h(\nu, \mu^{n+1}) - D_h(\nu, \mu^n) \\ &\quad \leq \frac{1}{L} \left(\int \frac{\delta F}{\delta \mu}(\mu^n, y)(\nu - \mu^n)(dy) - \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu^{n+1} - \mu^n)(dy) \right) - D_h(\mu^{n+1}, \mu^n) . \end{aligned}$$

Lemma 3 (Energy dissipation)

Let $F \in \mathscr{C}^1$ be L-relatively smooth with respect to h. Then

$$F(\mu^{n+1}) \leq F(\mu^n) \text{ and } \lim_{n\to\infty} D_h(\mu^{n+1}, \mu^n) = 0.$$

Proof. From L-relatively smoothness and the definition of $(\mu^n)_n$, for any $\mu \in \mathscr{P}$

$$F(\mu^{n+1}) \leq F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu^{n+1} - \mu^n)(dy) + L\, D_h(\mu^{n+1}, \mu^n)$$

$$\leq F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu - \mu^n)(dy) + L D_h(\mu, \mu^n).$$

Taking $\mu = \mu^n$ completes the proof of $F(\mu^n) \searrow$.

Lemma 3 (Energy dissipation)

Let $F \in \mathscr{C}^1$ be L-relatively smooth with respect to h. Then

$$F(\mu^{n+1}) \le F(\mu^n) \text{ and } \lim_{n\to\infty} D_h(\mu^{n+1},\mu^n) = 0.$$

Proof. From L-relatively smoothness and the definition of $(\mu^n)_n$, for any $\mu \in \mathscr{P}$

$$F(\mu^{n+1}) \le F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu^{n+1} - \mu^n)(dy) + L\, D_h(\mu^{n+1},\mu^n)$$

$$\le F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu - \mu^n)(dy) + LD_h(\mu,\mu^n).$$

Taking $\mu = \mu^n$ completes the proof of $F(\mu^n) \searrow$. Using $L-$smoothness and Proximal Lemma

$$F(\mu^{n+1}) \le F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu^{n+1} - \mu^n)(dy) + L\, D_h(\mu^{n+1},\mu^n)$$

$$\le F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\nu - \mu^n)(dy) + LD_h(\nu,\mu^n) - LD_h(\nu,\mu^{n+1}).$$

Take $\nu = \mu^n$ so

$$D_h(\mu^{n+1},\mu^n) \le \frac{1}{L}\left(F(\mu^n) - F(\mu^{n+1})\right).$$

Lemma 3 (Energy dissipation)

Let $F \in \mathscr{C}^1$ be L-relatively smooth with respect to h. Then

$$F(\mu^{n+1}) \leq F(\mu^n) \text{ and } \lim_{n\to\infty} D_h(\mu^{n+1},\mu^n) = 0.$$

Proof. From L-relatively smoothness and the definition of $(\mu^n)_n$, for any $\mu \in \mathscr{P}$

$$F(\mu^{n+1}) \leq F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu^{n+1} - \mu^n)(dy) + L\,D_h(\mu^{n+1},\mu^n)$$

$$\leq F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu - \mu^n)(dy) + LD_h(\mu,\mu^n).$$

Taking $\mu = \mu^n$ completes the proof of $F(\mu^n) \searrow$. Using $L-$smoothness and Proximal Lemma

$$F(\mu^{n+1}) \leq F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu^{n+1} - \mu^n)(dy) + L\,D_h(\mu^{n+1},\mu^n)$$

$$\leq F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\nu - \mu^n)(dy) + LD_h(\nu,\mu^n) - LD_h(\nu,\mu^{n+1}).$$

Take $\nu = \mu^n$ so

$$D_h(\mu^{n+1},\mu^n) \leq \frac{1}{L}\left(F(\mu^n) - F(\mu^{n+1})\right).$$

Since $F(\mu^{n+1}) \leq F(\mu^n)$

$$\sum_{n=1}^{m} D_h(\mu^{n+1},\mu^n) \leq \frac{1}{L}\left(F(\mu^0) - F(\mu^{m+1})\right) \leq \frac{1}{L}\left(F(\mu^0) - \min_{\mu} F(\mu)\right).$$

Theorem 4

Let $F \in \mathscr{C}^1$ be L-relatively smooth with respect to h. If $\mu \mapsto F(\mu)$ is convex, then

$$F(\mu^n) - F(\mu^*) \leq \frac{L}{n} D_h(\mu^*, \mu^0) \quad \textit{where} \quad \mu^* \in \operatorname*{argmin}_{\mu} F(\mu).$$

If F is l-strongly convex relative to h then

$$F(\mu^n) - F(\mu^*) \leq D_h(\mu^n, \mu^*) \leq L\, D_h(\mu^0, \mu^*) e^{-\frac{l}{L} n}.$$

Theorem 4

Let $F \in \mathscr{C}^1$ be L-relatively smooth with respect to h. If $\mu \mapsto F(\mu)$ is convex, then

$$F(\mu^n) - F(\mu^*) \leq \frac{L}{n} D_h(\mu^*, \mu^0) \quad \textit{where} \quad \mu^* \in \operatorname*{argmin}_{\mu} F(\mu)\,.$$

If F is l-strongly convex relative to h then

$$F(\mu^n) - F(\mu^*) \leq D_h(\mu^n, \mu^*) \leq L\, D_h(\mu^0, \mu^*) e^{-\frac{l}{L} n}\,.$$

Proof. Proximal lemma and $L-$smoothness implies

$$\begin{aligned}
& D_h(\nu, \mu^{n+1}) - D_h(\nu, \mu^n) \\
& \leq \frac{1}{L} \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\nu - \mu^n)(dy) - \frac{1}{L} \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu^{n+1} - \mu^n)(dy) - D_h(\mu^{n+1}, \mu^n) \\
& \leq \frac{1}{L} \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\nu - \mu^n)(dy) + \frac{1}{L} \left(F(\mu^n) - F(\mu^{n+1})\right)\,.
\end{aligned}$$

Theorem 4

Let $F \in \mathscr{C}^1$ be L-relatively smooth with respect to h. If $\mu \mapsto F(\mu)$ is convex, then

$$F(\mu^n) - F(\mu^*) \leq \frac{L}{n} D_h(\mu^*, \mu^0) \quad \textit{where} \quad \mu^* \in \operatorname*{argmin}_{\mu} F(\mu).$$

If F is l-strongly convex relative to h then

$$F(\mu^n) - F(\mu^*) \leq D_h(\mu^n, \mu^*) \leq L\, D_h(\mu^0, \mu^*) e^{-\frac{l}{L} n}.$$

Proof. Proximal lemma and $L-$smoothness implies

$$\begin{aligned}
&D_h(\nu, \mu^{n+1}) - D_h(\nu, \mu^n) \\
&\leq \frac{1}{L} \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\nu - \mu^n)(dy) - \frac{1}{L} \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu^{n+1} - \mu^n)(dy) - D_h(\mu^{n+1}, \mu^n) \\
&\leq \frac{1}{L} \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\nu - \mu^n)(dy) + \frac{1}{L} \left(F(\mu^n) - F(\mu^{n+1}) \right).
\end{aligned}$$

When F is convex

$$D_h(\nu, \mu^{n+1}) - D_h(\nu, \mu^n) \leq \frac{1}{L} \left(F(\nu) - F(\mu^n) + F(\mu^n) - F(\mu^{n+1}) \right).$$

Since $F(\mu^n)_n$ is decreasing

$$m(F(\mu^m) - F(\nu)) \leq \sum_{n=0}^{m-1} \left(F(\mu^{n+1}) - F(\nu) \right) \leq L\, D_h(\nu, \mu^0) - L\, D_h(\nu, \mu^m),$$

Proof cont.

If in addition F is l-strongly convex relative to h i.e

$$\int \frac{\delta F}{\delta \mu}(\mu^n, y)(\nu - \mu^n)(dy) \leq F(\nu) - F(\mu^n) - l\, D_h(\nu, \mu^n).$$

then

$$D_h(\nu, \mu^{n+1}) - D_h(\nu, \mu^n) \leq \frac{1}{L}\left(F(\nu) - F(\mu^{n+1})\right) - \frac{l}{L} D_h(\nu, \mu^n).$$

Since $F(\mu^*) - F(\mu^n) \leq 0$ conclusion follows from discrete Gronwall's lemma and the fact that when F is L-smooth relative to h we have

$$\begin{aligned} F(\mu^n) - F(\mu^*) &\leq \int \frac{\delta F}{\delta \mu}(\mu^*, y)(\mu^n - \mu^*)(dy) + L\, D_h(\mu^n, \mu^*) \\ &= L D_h(\mu^n, \mu^*). \end{aligned}$$

Bregman proximal descent

$$\mu^{n+1} \in \operatorname*{argmin}_{\mu} \{F(\mu) + LD_h(\mu, \mu^n)\} .$$

[1] Observed by Pierre-Cyril Aubin-Frankowski

Bregman proximal descent

$$\mu^{n+1} \in \underset{\mu}{\operatorname{argmin}} \{F(\mu) + LD_h(\mu, \mu^n)\} .$$

First order condition yileds

$$\frac{\delta h}{\delta \mu}(\mu^{n+1}, y) - \frac{\delta h}{\delta \mu}(\mu^n, y) = -\frac{1}{L}\left(\frac{\delta F}{\delta \mu}(\mu^{n+1}, y) + constant\right)$$

[1] Observed by Pierre-Cyril Aubin-Frankowski

Bregman proximal descent

$$\mu^{n+1} \in \underset{\mu}{\operatorname{argmin}} \{F(\mu) + LD_h(\mu, \mu^n)\} .$$

First order condition yileds

$$\frac{\delta h}{\delta \mu}(\mu^{n+1}, y) - \frac{\delta h}{\delta \mu}(\mu^n, y) = -\frac{1}{L}\left(\frac{\delta F}{\delta \mu}(\mu^{n+1}, y) + constant\right)$$

Let $g(\mu) = \frac{1}{L}F(\mu) + h(\mu)$ (so that $Lh = Lg - F$) for all $\mu \in \mathscr{P}(\mathscr{X})$. Note that by linearity of $h \mapsto D_h(\mu, \nu)$

$$\begin{aligned} F(\mu) + L\, D_h(\mu, \mu^n) &= F(\mu) - D_F(\mu, \mu^n) + L\, D_g(\mu, \mu^n) \\ &= F(\mu^n) + \int \frac{\delta F}{\delta \mu}(\mu^n, y)(\mu - \mu^n)(dy) + L\, D_g(\mu, \mu^n) . \end{aligned}$$

Proximal descent with h-Bregman divergence is equivalent to mirror descent with g-Bregman divergence![1]

---

[1] Observed by Pierre-Cyril Aubin-Frankowski

Bregman proximal descent

Need to check that relative smoothness and convexity assumption hold with g-Bregman divergence.

Bregman proximal descent

Need to check that relative smoothness and convexity assumption hold with g-Bregman divergence.

L-smoothness relative to g: Since $D_h(\mu, \mu') \geq 0$

$$D_F(\mu, \mu') = LD_g(\mu, \mu') - LD_h(\mu, \mu') \leq LD_g(\mu, \mu').$$

Bregman proximal descent

Need to check that relative smoothness and convexity assumption hold with g-Bregman divergence.

L-smoothness relative to g: Since $D_h(\mu,\mu') \geq 0$

$$D_F(\mu,\mu') = LD_g(\mu,\mu') - LD_h(\mu,\mu') \leq LD_g(\mu,\mu').$$

$\frac{Ll}{L-l}$-convexity relative to g: We know that

$$D_F(\mu,\mu') \geq lD_h(\mu,\mu') = l(D_g(\mu,\mu') - \frac{1}{L} D_F(\mu,\mu')).$$

Hence

$$D_F(\mu,\mu') \geq \frac{Ll}{L-l} D_g(\mu,\mu').$$

Bregman proximal descent

Need to check that relative smoothness and convexity assumption hold with g-Bregman divergence.

L-smoothness relative to g: Since $D_h(\mu, \mu') \geq 0$

$$D_F(\mu, \mu') = LD_g(\mu, \mu') - LD_h(\mu, \mu') \leq LD_g(\mu, \mu').$$

$\frac{Ll}{L-l}$-convexity relative to g: We know that

$$D_F(\mu, \mu') \geq lD_h(\mu, \mu') = l(D_g(\mu, \mu') - \frac{1}{L} D_F(\mu, \mu')).$$

Hence

$$D_F(\mu, \mu') \geq \frac{Ll}{L-l} D_g(\mu, \mu').$$

Bregman proximal does not require smoothness assumption to converge (due the scheme being implicit) at the cost of harder optimisation problem

Convergence Analysis

Key Challenges

$$dX_s(\pi) = b_s(X_s(\pi), \pi_s)\, ds + \sigma_s(X_s(\pi), \pi_s)\, dW_s\,,\ s \in [0, T],\ X_0(\pi) = x\,.$$

$$J^\tau(\pi) := \mathbb{E}\left[\int_0^T [f_s(X_s(\pi), \pi_s) + \tau\, h(\pi_s)]\, ds + g(X_T(\pi))\right]$$

Objective: Prove the convergence of the MSA II with

$$\pi_s^{n+1} \in \arg\min_{\pi \in \mathscr{P}(\mathbb{R}^p)} \int \frac{\delta H_s^\tau}{\delta m}(\Theta_s(\pi^n), \pi_s^n, a)(\pi - \pi_s^n)(da) + \lambda\, D_h(\pi | \pi_s^n)\,.$$

Key Challenges

$$dX_s(\pi) = b_s(X_s(\pi),\pi_s)\,ds + \sigma_s(X_s(\pi),\pi_s)\,dW_s\,,\ s\in[0,T],\ X_0(\pi)=x\,.$$

$$J^\tau(\pi) := \mathbb{E}\left[\int_0^T [f_s(X_s(\pi),\pi_s)+\tau\, h(\pi_s)]\,ds + g(X_T(\pi))\right]$$

Objective: Prove the convergence of the MSA II with

$$\pi_s^{n+1} \in \arg\min_{\pi\in\mathscr{P}(\mathbb{R}^p)} \int \frac{\delta H_s^\tau}{\delta m}(\Theta_s(\pi^n),\pi_s^n,a)(\pi-\pi_s^n)(da) + \lambda\, D_h(\pi|\pi_s^n)\,.$$

we need

- Relative smoothness of $\pi \mapsto J^\tau(\pi)$

Key Challenges

$$dX_s(\pi) = b_s(X_s(\pi),\pi_s)\,ds + \sigma_s(X_s(\pi),\pi_s)\,dW_s\,,\ s\in[0,T],\ X_0(\pi)=x\,.$$

$$J^\tau(\pi) := \mathbb{E}\left[\int_0^T [f_s(X_s(\pi),\pi_s)+\tau\,h(\pi_s)]\,ds + g(X_T(\pi))\right]$$

Objective: Prove the convergence of the MSA II with

$$\pi_s^{n+1} \in \arg\min_{\pi\in\mathscr{P}(\mathbb{R}^p)} \int \frac{\delta H_s^\tau}{\delta m}(\Theta_s(\pi^n),\pi_s^n,a)(\pi-\pi_s^n)(da) + \lambda\, D_h(\pi|\pi_s^n)\,.$$

we need

- Relative smoothness of $\pi \mapsto J^\tau(\pi)$ - these are obtained by utilising the theory of Bounded Mean Oscillation (BMO) martingales required for estimates on the adjoint Backward Stochastic Differential Equation (BSDE).

Key Challenges

$$dX_s(\pi) = b_s(X_s(\pi),\pi_s)\,ds + \sigma_s(X_s(\pi),\pi_s)\,dW_s\,,\; s\in[0,T],\; X_0(\pi)=x\,.$$

$$J^\tau(\pi) := \mathbb{E}\left[\int_0^T [f_s(X_s(\pi),\pi_s) + \tau\,h(\pi_s)]\,ds + g(X_T(\pi))\right]$$

Objective: Prove the convergence of the MSA II with

$$\pi_s^{n+1} \in \arg\min_{\pi\in\mathscr{P}(\mathbb{R}^p)} \int \frac{\delta H_s^\tau}{\delta m}(\Theta_s(\pi^n),\pi_s^n,a)(\pi-\pi_s^n)(da) + \lambda\,D_h(\pi|\pi_s^n)\,.$$

we need

- Relative smoothness of $\pi \mapsto J^\tau(\pi)$ - these are obtained by utilising the theory of Bounded Mean Oscillation (BMO) martingales required for estimates on the adjoint Backward Stochastic Differential Equation (BSDE).
- Convexity (to obtain linear convergence) or relative convexity (to obtain exponential convergence) of $\pi \mapsto J^\tau(\pi)$

Key Challenges

$$dX_s(\pi) = b_s(X_s(\pi),\pi_s)\,ds + \sigma_s(X_s(\pi),\pi_s)\,dW_s\,,\ s\in[0,T]\,,\ X_0(\pi)=x\,.$$

$$J^\tau(\pi) := \mathbb{E}\left[\int_0^T [f_s(X_s(\pi),\pi_s)+\tau\, h(\pi_s)]\,ds + g(X_T(\pi))\right]$$

Objective: Prove the convergence of the MSA II with

$$\pi_s^{n+1} \in \arg\min_{\pi\in\mathscr{P}(\mathbb{R}^p)} \int \frac{\delta H_s^\tau}{\delta m}(\Theta_s(\pi^n),\pi_s^n,a)(\pi-\pi_s^n)(da) + \lambda\, D_h(\pi|\pi_s^n)\,.$$

we need

- Relative smoothness of $\pi\mapsto J^\tau(\pi)$ - these are obtained by utilising the theory of Bounded Mean Oscillation (BMO) martingales required for estimates on the adjoint Backward Stochastic Differential Equation (BSDE).
- Convexity (to obtain linear convergence) or relative convexity (to obtain exponential convergence) of $\pi\mapsto J^\tau(\pi)$ - but $\pi\mapsto J^\tau(\pi)$ is not convex...

Assumption 1 (Assumptions on b and σ)

The functions b and σ are jointly continuous in t and $\mathscr{C}^1$.
Moreover, $\forall x, x' \in \mathbb{R}^d, \forall m, m' \in \mathscr{P}(\mathbb{R}^p), \forall t \in [0,T]$, $\Phi_t = (b_t, \sigma_t)$

$$|(D_x\Phi_t)(x,m)| \leq K$$

$$|\Phi_t(x,m) - \Phi_t(x,m')|^2 \leq K\, D_h(m|m')$$

$$|(D_x\Phi_t)(x,m) - (D_x\Phi_t)(x',m')|^2 \leq K|x-x'|^2 + K\, D_h(m|m'),$$

$$\left|\int \left(D_x \frac{\delta \Phi_t}{\delta m}\right)(x,m,a)(\mu-\mu')(da)\right|^2 + \left|\int \frac{\delta \Phi_t}{\delta m}(x,m,a)(\mu-\mu')(da)\right|^2 \leq K\, D_h(\mu|\mu')$$

$$\left|\int\int \frac{\delta^2 b_t}{\delta m^2}(x,m,a,a')(\mu-\mu')(da')(\mu-\mu')(da)\right| \leq K D_h(\mu|\mu'), \text{ and } \frac{\delta^2 \sigma_t}{\delta m^2}(x,m,a,a') = 0.$$

Assumption 2 (Assumptions on f and g)

The function f is jointly continuous in t, f and g are in $\mathscr{C}^1$. Moreover, there is a constant $K \geq 0$ such that $\forall x, x' \in \mathbb{R}^d, \forall m, m' \in \mathscr{P}_2(\mathbb{R}^p), \forall a \in \mathbb{R}^p, \forall t \in [0,T]$,

$$|(D_x g)(x)| + |(D_x f_t)(x,m)| \leq K$$

$$|(D_x f_t)(x,m) - (D_x f_t)(x',m')|^2 \leq K|x-x'|^2 + K\, D_h(m|m'),$$

$$|(D_x g)(x) - (D_x g)(x')| \leq K|x-x'|$$

$$\left| \int \left(D_x \frac{\delta f_t}{\delta m} \right)(x,m,a)(\mu-\mu')(da) \right|^2 + \left| \int \frac{\delta f_t}{\delta m}(x,m,a)(\mu-\mu')(da) \right|^2 \leq K D_h(\mu|\mu')$$

$$\left| \int \int \frac{\delta^2 f_t}{\delta m^2}(x,m,a,a')(\mu-\mu')(da')(\mu-\mu')(da) \right| \leq K D_h(\mu|\mu'),.$$

Relative Smoothness

Lemma 5

Let $\Theta_s(\pi) := (X_s(\pi), Y_s(\pi), Z_s(\pi))$ *for* $s \in [0,T]$. *For any admissible controls* $\pi, \pi' \in \mathscr{A}$ *there exists a constant* $C > 0$ *such that*

$$J^\tau(\pi') - J^\tau(\pi) \leq \mathbb{E}\int_0^T \left[\int \frac{\delta H_s^\tau}{\delta m}(\Theta_s(\pi), \pi_s, a)(\pi_s' - \pi_s)(a)\, da\right] ds + C\mathbb{E}\int_0^T D_h(\pi_s' | \pi_s)\, ds\,.$$

Relative Smoothness

Lemma 5

Let $\Theta_s(\pi) := (X_s(\pi), Y_s(\pi), Z_s(\pi))$ *for* $s \in [0, T]$. *For any admissible controls* $\pi, \pi' \in \mathscr{A}$ *there exists a constant* $C > 0$ *such that*

$$J^\tau(\pi') - J^\tau(\pi) \le \mathbb{E}\int_0^T \left[\int \frac{\delta H_s^\tau}{\delta m}(\Theta_s(\pi), \pi_s, a)(\pi'_s - \pi_s)(a)\, da\right] ds + C\mathbb{E}\int_0^T D_h(\pi'_s | \pi_s)\, ds\,.$$

Note that

$$J^\tau(\pi^{n+1}) \le J^\tau(\pi^n) + \mathbb{E}\int_0^T \left[\int \frac{\delta H_s^\tau}{\delta m}(\Theta_s(\pi^n), \pi_s^n, a)(\pi_s^{n+1} - \pi_s^n)(a)\, da\right] ds + C\mathbb{E}\int_0^T D_h(\pi_s^{n+1} | \pi_s^n)\, ds\,.$$

and by definition of the update for any $\pi \in \mathscr{A}$

$$\int \frac{\delta H_s^\tau}{\delta m}(\Theta_s(\pi^n), \pi_s^n, a)(\pi_s^{n+1} - \pi_s^n)(a)\, da + \lambda\, D_h(\pi_s^{n+1} | \pi_s^n) \le 0$$

Convexity

Assumption 3 (Assumptions on convexity)

- *the map* $x \mapsto g(x)$ *is convex*
- *the map* $(x,m) \mapsto H_t^0(x,y,z,m)$ *is convex for all* (t,y,z), *in the sense that for all* $x,x' \in \mathbb{R}^d$ *and all* $m,m' \in \mathscr{P}(\mathbb{R}^p)$ *it holds that*

$$H_t^0(x',y,z,m') - H_t^0(x,y,z,m) - (D_x H_t^0)(x,y,z,m)(x-x') - \int \frac{\delta H_t^0}{\delta m}(x,y,z,m,a)(m-m')(da) \geq 0\,.$$

Convexity

Assumption 3 (Assumptions on convexity)

- *the map* $x \mapsto g(x)$ *is convex*
- *the map* $(x,m) \mapsto H_t^0(x,y,z,m)$ *is convex for all* (t,y,z)*, in the sense that for all* $x,x' \in \mathbb{R}^d$ *and all* $m,m' \in \mathscr{P}(\mathbb{R}^p)$ *it holds that*

$$H_t^0(x',y,z,m') - H_t^0(x,y,z,m) \\ - (D_x H_t^0)(x,y,z,m)(x-x') - \int \frac{\delta H_t^0}{\delta m}(x,y,z,m,a)(m-m')(da) \geq 0\,.$$

Theorem 6

For any $\pi,\pi' \in \mathscr{A}$

$$J^\tau(\pi) - J^\tau(\pi') - \mathbb{E}\int_0^T \int \frac{\delta H_s^\tau}{\delta m}(\Theta_s(\pi'),\pi'_s,a)(\pi_s - \pi'_s)(da)\,ds \geq \tau\,\mathbb{E}\int_0^T \left[D_h(\pi|\pi')\right]\,ds\,.$$

Convergence of the MSA

Theorem 7

Let $\{\pi^n\}_{n\geq 0}$ *be the sequence given by the MSA II. Then,*

- *when* $\tau > 0$*:*

$$J^\tau(\pi^n) - J^\tau(\pi^*) \leq \lambda e^{-\frac{\tau}{\lambda}n}\mathbb{E}\int_0^T D_h(\pi_s^*|\pi_s^0)\,ds\,.$$

- *when* $\tau = 0$*:*

$$J^0(\pi^n) - J^0(\pi^*) \leq \frac{\lambda}{n}\mathbb{E}\int_0^T D_h(\pi_s^*|\pi_s^0)\,ds\,.$$

Bergman Proximal Descent aka MSA I

For MSA I the update is:

$$\pi_s^{n+1} \in \arg\min_{\pi \in \mathscr{P}(\mathbb{R}^p)} H_s^\tau(\Theta_s(\pi^n), \pi) + \lambda D_h(\pi | \pi_s^n),$$

Bergman Proximal Descent aka MSA I

For MSA I the update is:

$$\pi_s^{n+1} \in \arg\min_{\pi \in \mathscr{P}(\mathbb{R}^p)} H_s^\tau(\Theta_s(\pi^n), \pi) + \lambda D_h(\pi|\pi_s^n),$$

Lemma 8

Let $\Theta_s(\pi) := (X_s(\pi), Y_s(\pi), Z_s(\pi))$ for $s \in [0, T]$. For any admissible controls $\pi, \pi' \in \mathscr{A}$ there exists a constant $C > 0$ such that

$$\begin{aligned} J^\tau(\pi') - J^\tau(\pi) \leq \mathbb{E} \int_0^T [H_s^\tau(\Theta_s(\pi), \pi_s') - H_s^\tau(\Theta_s(\pi), \pi_s)]\, ds \\ + C\, \mathbb{E} \int_0^T D_h(\pi_s'|\pi_s)\, ds. \end{aligned}$$

Bergman Proximal Descent aka MSA I

For MSA I the update is:

$$\pi_s^{n+1} \in \arg \min_{\pi \in \mathscr{P}(\mathbb{R}^p)} H_s^\tau(\Theta_s(\pi^n), \pi) + \lambda D_h(\pi | \pi_s^n),$$

Lemma 8

Let $\Theta_s(\pi) := (X_s(\pi), Y_s(\pi), Z_s(\pi))$ for $s \in [0, T]$. For any admissible controls $\pi, \pi' \in \mathscr{A}$ there exists a constant $C > 0$ such that

$$J^\tau(\pi') - J^\tau(\pi) \leq \mathbb{E} \int_0^T [H_s^\tau(\Theta_s(\pi), \pi_s') - H_s^\tau(\Theta_s(\pi), \pi_s)] \, ds + C \, \mathbb{E} \int_0^T D_h(\pi_s' | \pi_s) \, ds .$$

Note that

$$J^\tau(\pi^{n+1}) \leq J^\tau(\pi^n) + \mathbb{E} \int_0^T [H_s^\tau(\Theta_s(\pi^n), \pi_s^{n+1}) - H_s^\tau(\Theta_s(\pi^n), \pi_s^n)] \, ds + C \, \mathbb{E} \int_0^T D_h(\pi_s^{n+1} | \pi_s^n) \, ds .$$

and

$$H_s^\tau(\Theta_s(\pi^n), \pi^{n+1} - H_s^\tau(\Theta_s(\pi^n), \pi^n)) + \lambda D_h(\pi^{n+1} | \pi_s^n) \leq 0$$

Bergman Proximal Descent aka MSA I

For MSA I the update is:

$$\pi_s^{n+1} \in \arg\min_{\pi \in \mathscr{P}(\mathbb{R}^p)} H_s^\tau(\Theta_s(\pi^n), \pi) + \lambda D_h(\pi | \pi_s^n),$$

Lemma 8

Let $\Theta_s(\pi) := (X_s(\pi), Y_s(\pi), Z_s(\pi))$ for $s \in [0, T]$. For any admissible controls $\pi, \pi' \in \mathscr{A}$ there exists a constant $C > 0$ such that

$$J^\tau(\pi') - J^\tau(\pi) \le \mathbb{E} \int_0^T [H_s^\tau(\Theta_s(\pi), \pi_s') - H_s^\tau(\Theta_s(\pi), \pi_s)]\, ds + C\, \mathbb{E} \int_0^T D_h(\pi_s' | \pi_s)\, ds.$$

Note that

$$J^\tau(\pi^{n+1}) \le J^\tau(\pi^n) + \mathbb{E} \int_0^T [H_s^\tau(\Theta_s(\pi^n), \pi_s^{n+1}) - H_s^\tau(\Theta_s(\pi^n), \pi_s^n)]\, ds + C\, \mathbb{E} \int_0^T D_h(\pi_s^{n+1} | \pi_s^n)\, ds.$$

and

$$H_s^\tau(\Theta_s(\pi^n), \pi^{n+1} - H_s^\tau(\Theta_s(\pi^n), \pi^n)) + \lambda D_h(\pi^{n+1} | \pi_s^n) \le 0$$

Convergence rates can be established as shown above.

Outlook

- Extension to cover policies of the feedback form: need regularity of controls
- Continuous time gradient flow: need existence&uniqueness of the flow (clear in KL setting, less clear for general D_h)
- Extension to games: Best Response and Fictitious play
- Study regret for gradient/descent flows based learning algorithms.

3+1 years Research positions available in London/Oxford/Edinburgh

References I

[Aubin-Frankowski et al.,] Aubin-Frankowski, P.-C., Korba, A., and Léger, F. Mirror descent with relative smoothness in measure spaces, with application to sinkhorn and em. In *Advances in Neural Information Processing Systems*.

[Chizat et al., 2018] Chizat, L., Peyré, G., Schmitzer, B., and Vialard, F.-X. (2018). An interpolating distance between optimal transport and fisher–rao metrics. *Foundations of Computational Mathematics*, 18:1–44.

[Gobet and Labart, 2010] Gobet, E. and Labart, C. (2010). Solving bsde with adaptive control variate. *SIAM Journal on Numerical Analysis*, 48(1):257–277.

[Kerimkulov et al., 2020] Kerimkulov, B., Šiška, D., and Szpruch, L. (2020). Exponential convergence and stability of howard's policy improvement algorithm for controlled diffusions. *SIAM Journal on Control and Optimization*, 58(3):1314–1340.

[Kerimkulov et al., 2021] Kerimkulov, B., Šiška, D., and Szpruch, L. (2021). A modified msa for stochastic control problems. *Applied Mathematics & Optimization*, pages 1–20.

[Krylov and Chernousko, 1963] Krylov, I. A. and Chernousko, F. L. (1963). On a method of successive approximations for the solution of optimal control problems. Technical report, FOREIGN TECHNOLOGY DIV WRIGHT-PATTERSON AFB OH.

[Li et al., 2018] Li, Q., Chen, L., Tai, C., and Weinan, E. (2018). Maximum principle based algorithms for deep learning. *Journal of Machine Learning Research*, 18(165):1–29.

[Liero et al., 2022] Liero, M., Mielke, A., and Savaré, G. (2022). Fine properties of geodesics and geodesic λ-convexity for the hellinger-kantorovich distance. *arXiv preprint arXiv:2208.14299*.

[Liu et al., 2023] Liu, L., Majka, M. B., and Szpruch, Ł. (2023). Polyak–łojasiewicz inequality on the space of measures and convergence of mean-field birth-death processes. *Applied Mathematics & Optimization*, 87(3):48.

[Lu et al., 2018] Lu, H., Freund, R. M., and Nesterov, Y. (2018). Relatively smooth convex optimization by first-order methods, and applications. *SIAM Journal on Optimization*, 28(1):333–354.

[Reisinger et al., 2021] Reisinger, C., Stockinger, W., and Zhang, Y. (2021). A fast iterative pde-based algorithm for feedback controls of nonsmooth mean-field control problems. *arXiv preprint arXiv:2108.06740*.