Convergence of Dual Approximate Dynamic Programming

V. Leclère

March 28, 2012

Contents

An optimization problem

Let's consider the following dynamic optimization problem

$$\min_{\mathbf{X},\mathbf{U}} \quad \mathbb{E}\sum_{t=0}^{T} L_t\big(\mathbf{X}_t, \mathbf{U}_t, \mathbf{W}_t\big)$$

$$\mathbf{X}_{t+1} = f_t(\mathbf{X}_t, \mathbf{U}_t, \mathbf{W}_t)$$

$$\mathbf{X}_0 = x_0$$

$$\theta_t(\mathbf{U}_t) = \sum_{i=1}^{n} \mathbf{U}_t^i - \mathbf{D}_t = 0$$

Where $\mathbf{W}_t$ is a white noise, $\mathbf{X}_t$ is the state process, and $\mathbf{U}_t$ is the control process, both measurable with respect to $(\mathcal{F}_t) = \sigma(\mathbf{W}_0, \mathbf{D}_0, \cdots, \mathbf{W}_t, \mathbf{D}_t)$.
For example $\mathbf{U}_t = (\mathbf{U}_t^1, \cdots, \mathbf{U}_t^n)$ is the production of n power units and $\mathbf{D}_t$ is the demand.

Dualization

If we dualize the constraint on the control $\mathbf{U}_t$, we obtain

$$\max_{\boldsymbol{\lambda}} \quad \min_{\mathbf{X},\mathbf{U}} \quad \mathbb{E}\Big(\sum_{t=0}^{T} L_t(\mathbf{X}_t, \mathbf{U}_t, \mathbf{W}_t) + \boldsymbol{\lambda}_t \theta_t(\mathbf{U}_t) \Big)$$

$$\mathbf{X}_{t+1} = f_t(\mathbf{X}_t, \mathbf{U}_t, \mathbf{W}_t)$$

$$\mathbf{X}_0 = x_0$$

Where we can choose $\boldsymbol{\lambda}_t$ to be measurable with respect to $(\mathcal{F}_t)$.

Uzawa algorithm

At step k we want to have a process $\boldsymbol{\lambda}^{(k)}$ and solve

$$\min_{\mathbf{X},\mathbf{U}} \quad \mathbb{E}\sum_{t=0}^{T} L_t(\mathbf{X}_t,\mathbf{U}_t,\mathbf{W}_t) + \boldsymbol{\lambda}_t^{(k)}\theta_t(\mathbf{U}_t)$$
$$\mathbf{X}_{t+1} = f_t(\mathbf{X}_t,\mathbf{U}_t,\mathbf{W}_t)$$
$$\mathbf{X}_0 = x_0$$

we will then determine $\boldsymbol{\lambda}^{(k+1)}$ by a gradient step.
However solving this problem is quite difficult as $\boldsymbol{\lambda}^{(k)}$ is a non-markovian process, thus dynamic programming would have to be done on a state of dimension increasing with T, as we need to keep the whole past of $\boldsymbol{\lambda}^{(k)}$, which is numerically impossible.

Contents

DADP algorithm : general idea

The main idea of Dual Approximate Dynamic Programming is to replace $\boldsymbol{\lambda}_t$ by $\mathbb{E}(\boldsymbol{\lambda}_t \mid \mathbf{X}_t)$.

$$\min_{\mathbf{X},\mathbf{U}} \quad \mathbb{E}\sum_{t=0}^{T} L_t(\mathbf{X}_t, \mathbf{U}_t, \mathbf{W}_t) + \mathbb{E}(\boldsymbol{\lambda}_t^{(k)} \mid \mathbf{X}_t)\theta_t(\mathbf{U}_t)$$

$$\mathbf{X}_{t+1} = f_t(\mathbf{X}_t, \mathbf{U}_t, \mathbf{W}_t)$$

$$\mathbf{X}_0 = x_0$$

And solving this problem by Dynamic Programming, with the same state, is possible as $\mathbb{E}(\boldsymbol{\lambda}_t^{(k)} \mid \mathbf{X}_t)$ is a function of $\mathbf{X}_t$.

Generaly speaking we can introduce a “short memory” information process $\mathbf{Y}_t$, defining $\mathcal{B}_t = \sigma(\mathbf{Y}_t)$ and replace $\boldsymbol{\lambda}_t$ by $\mathbb{E}(\boldsymbol{\lambda}_t \mid \mathcal{B}_t)$.

DADP algorithm : as an approximation

It has been shown that this method is equivalent to solve

$$\min_{\mathbf{X},\mathbf{U}} \quad \mathbb{E}\sum_{t=0}^{T} L_t\big(\mathbf{X}_t, \mathbf{U}_t, \mathbf{W}_t\big)$$

$$\mathbf{X}_{t+1} = f_t(\mathbf{X}_t, \mathbf{U}_t, \mathbf{W}_t)$$

$$\mathbf{X}_0 = x_0$$

$$\mathbb{E}(\theta_t(\mathbf{U}_t) \mid \mathcal{B}_t) = 0$$

Contents

General Problem

We consider a stochastic optimization problem for a probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

$$(\mathcal{P}) \qquad \min_{\Theta(\mathbf{U}) \in -C} \quad J(\mathbf{U}) \;,$$

where the control is a random variable $\mathbf{U} \in \mathcal{U}$ on $(\Omega, \mathcal{F}, \mathbb{P})$ with value in a Banach $\mathbb{U}$, the criterion $J : \mathcal{U} \to \mathbb{R} \cup \{+\infty\}$ is an operator, and $\Theta : \mathcal{U} \to \mathcal{V}$ is the operator of constraints, with C a closed convex cone of $\mathcal{V}$.

Usual choice of criterion are :

- $J(\mathbf{U}) := \mathbb{E}\big(j(\mathbf{U})\big)$
- risk measures
- Worst-case scenario

Θ can take into account :

- almost sure constraints : $\Theta\big(\mathbf{U}\big)(\omega) := \theta\big(\mathbf{U}(\omega)\big)$, with $C = \{0\}$ and $\theta\big(\mathbf{U}\big) = 0$ is realized almost surely.
- measurability constraints : $\Theta\big(\mathbf{U}\big) := \mathbb{E}\big(\mathbf{U} \mid \mathcal{B}\big) - \mathbf{U}$, with $C = \{0\}$,
- risk constraint : $\Theta\big(\mathbf{U}\big) := \rho(\mathbf{U}) - a$, where ρ is a risk measure, and $C = \mathbb{R}^+$
- or probability constraint : $\Theta\big(\mathbf{U}\big) := \mathbb{P}\big(\mathbf{U} \in A\big) - p$, with $C = \mathbb{R}^+$, that is $\mathbb{P}\big(\mathbf{U} \in A\big) \leq p$.

We will consider a problem,

$$(\mathcal{P}) \qquad \min_{\mathbf{U}\in\mathcal{U}} \underbrace{J(\mathbf{U}) + \chi_{\mathbf{U}\in\mathcal{U}^{ad}}}_{:=\tilde{J}(\mathbf{U})},$$

with

$$\mathcal{U}^{ad} := \{\mathbf{U} \in \mathcal{U} | \Theta(\mathbf{U}) \in -C\}$$

and its approximation (for a subfield $\mathcal{F}_n$ of $\mathcal{F}$).

$$(\mathcal{P}_n) \qquad \min_{\mathbf{U}\in\mathcal{U}} \underbrace{J(\mathbf{U}) + \chi_{\mathbf{U}\in\mathcal{U}_n^{ad}}}_{:=\tilde{J}_n(\mathbf{U})},$$

with

$$\mathcal{U}_n^{ad} := \{\mathbf{U} \in \mathcal{U} | \mathbb{E}\big(\Theta(\mathbf{U}) \bigm| \mathcal{F}_n\big) \in -C\}$$

Contents

Epi-convergence

Let (A_n) a sequence of subset of a topological space (E, τ). We define

$$\limsup_n A_n = \{x \in E | \exists (x_{n_k}) \to_\tau x, \text{ with } \forall k, x_{n_k} \in A_{n_k}\}$$

$$\liminf_n A_n = \{x \in E | \exists (x_n) \to_\tau x, \text{ with } \forall n, x_n \in A_n\}$$

And we say that (A_n) converges to A, iff

$$\limsup_n A_n = \liminf_n A_n = A \ .$$

We say that a sequence of function (f_n) epi-converges to f, iff

$$\lim \mathrm{epi}\, f_n = \mathrm{epi}\, f \ .$$

Kudo-convergence

If $\mathcal{F}$ is a $\sigma-$algebra, and $(\mathcal{F}_n)$ a sequence of complete sub-$\sigma-$algebra of $\mathcal{F}$ we say that $(\mathcal{F}_n)$ Kudo converges to $\mathcal{F}_\infty$ if

$$\forall A \in \mathcal{F}, \quad \mathbb{P}(A|\mathcal{F}_n) \to_{\mathbb{P}} \mathbb{P}(A|\mathcal{F}_\infty)$$

or equivalently

$$\forall \mathbf{X} \in \mathcal{L}^\infty(\mathcal{F}), \quad \mathbb{E}|\mathbb{E}(\mathbf{X} \mid \mathcal{F}_n)| \to \mathbb{E}|\mathbb{E}(\mathbf{X} \mid \mathcal{F}_\infty)|$$

We begin by a lemma from Piccinini.

Theorem

Let $(\mathcal{F}_n)$ be a sequence of $\sigma-$algebra. The following statements are equivalent :

1. $\mathcal{F}_n \to \mathcal{F}_\infty$.
2. $\forall \mathbf{X} \in \mathcal{L}^p_E(\mathcal{F}), \quad \mathbb{E}\big(\mathbf{X} \mid \mathcal{F}_n\big) \to_{\mathcal{L}^p} \mathbb{E}\big(\mathbf{X} \mid \mathcal{F}_\infty\big)$.
3. $\forall \mathbf{X} \in \mathcal{L}^p_E(\mathcal{F}), \quad \mathbb{E}\big(\mathbf{X} \mid \mathcal{F}_n\big) \rightharpoonup_{\mathcal{L}^p} \mathbb{E}\big(\mathbf{X} \mid \mathcal{F}_\infty\big)$.

And we have the following corollary

Theorem

Let $(\mathcal{F}_n)$ be a sequence of $\sigma-$algebra, and $1 \leq p < \infty$. If $\mathcal{F}_n \to \mathcal{F}_\infty$, and $\mathbf{X}_n \to_{\mathcal{L}^p} \mathbf{X}$ (resp. $\mathbf{X}_n \rightharpoonup_{\mathcal{L}^p} \mathbf{X}$) then $\mathbb{E}\big(\mathbf{X}_n \mid \mathcal{F}_n\big) \to_{\mathcal{L}^p} \mathbb{E}\big(\mathbf{X} \mid \mathcal{F}_\infty\big)$ (resp. $\mathbb{E}\big(\mathbf{X}_n \mid \mathcal{F}_n\big) \rightharpoonup_{\mathcal{L}^p} \mathbb{E}\big(\mathbf{X} \mid \mathcal{F}_\infty\big)$).

Contents

Convergence result

Theorem

If $\mathcal{U}$ is endowed with a topology τ, and $\mathcal{V} = \mathcal{L}^p(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{V})$ endowed with the strong or weak topology (p being in $[1, \infty)$), and C is stable by $\mathbb{E}(\cdot \mid \mathcal{F}_n)$. If Θ and J are continuous, and $(\mathcal{F}_n)$ Kudo-converges to $\mathcal{F}$, then $\tilde{J}_n$ epi-converges to $\tilde{J}$.

Proof :

- It is sufficient to show that $\mathcal{U}_n^{ad} \to_{PK} \mathcal{U}^{ad}$, as it will imply that $\chi_{\mathcal{U}_n^{ad}} \to_e \chi_{\mathcal{U}^{ad}}$, and J being continuous we will have epi-convergence of $\tilde{J}_n$ to $\tilde{J}$.
- Stability of C imply that $\mathcal{U}^{ad} \subset \liminf_n \mathcal{U}_n^{ad}$.
- If $\mathbf{U} \in \limsup_n \mathcal{U}_n^{ad}$, we have $(\mathbf{U}_{n_k})_k \to_\tau \mathbf{U}$, such that for all $k \in \mathbb{N}$, $\mathbb{E}\big(\Theta(\mathbf{U}_{n_k}) | \mathcal{F}_{n_k}\big) \in -C$.
- Continuity of Θ, convergence of $\mathcal{F}_{n_k}$, preceding corollary, and closedness of $-C$ achieve the proof. ∎

Convergence result

Theorem

If $(\mathcal{F}_n) \to \mathcal{F}$, J and Θ are continuous, then we have $(\mathcal{P}_n) \to (\mathcal{P})$ in the following sense : If $(\mathbf{U}_n)$ is a sequence of control such that for all $n \in \mathbb{N}$,

$$\tilde{J}_n(\mathbf{U}_n) < \inf_{\mathbf{U} \in \mathcal{U}} \tilde{J}_n(\mathbf{U}) + \epsilon_n, \text{ where } \lim_n \epsilon_n = 0,$$

then, for every converging sub-sequence $\mathbf{U}_{n_k}$, we have

$$\tilde{J}(\lim_k \mathbf{U}_{n_k}) = \inf_{\mathbf{U} \in \mathcal{U}} \tilde{J}(\mathbf{U}) = \lim_k \tilde{J}_{n_k}(\mathbf{U}_{n_k})$$

Moreover if $(\mathcal{F}_n)$ is a filtration, then the convergence is monotonous.

Contents

Theorem

Let (x_n) be a sequence in a topological space. If from any subsequence (x_{n_k}) we can extract a sub-subsequence $(x_{\sigma(n_k)})$ converging to x^*, then (x_n) converges to x^*.

Let's note that $(\mathbf{U}_n)$ is such that from any subsequence there is a further subsequence converging almost surely to $\mathbf{U}$ is equivalent to $(\mathbf{U}_n) \to_{\mathbb{P}} \mathbf{U}$.

Theorem

Let $\Theta : \mathcal{U} \to \mathcal{V}$, where $\mathcal{U}$ is a set of random variable endowed with the topology of convergence in probability, and $(\mathcal{V}, \tau)$ is a topological space. If $\mathbf{U}_n \to \mathbf{U}$ almost surely imply $\Theta(\mathbf{U}_n) \to_\tau \Theta(\mathbf{U})$, then Θ is continuous.

Contents

Almost sure constraint

Theorem

If $\mathcal{U}$ is the set of random variable on $(\Omega, \mathcal{F}, \mathbb{P})$, with the topology of convergence in probability, and if θ is continuous and bounded, then $\Theta(\mathbf{U})(\omega) := \theta(\mathbf{U}(\omega))$ is continuous.

Proof : Suppose that $\mathbf{U}_n \to_{\text{a.s}} \mathbf{U}$, then by boundeness of θ we have that $\left(||\theta(\mathbf{U}_{\sigma(n_k)}) - \theta(\mathbf{U})||^p\right)_k$ is bounded, and thus by dominated convergence theorem we have that $\theta(\mathbf{U}_n) \to_{\mathcal{L}^p} \theta(\mathbf{U})$. ∎

Measurability constraint

Theorem

We set $\mathcal{U} = \mathcal{L}^p(\Omega, \mathcal{F}, \mathbb{P})$*,if* $\mathcal{B}$ *is a sub-*σ*-algebra of* $\mathcal{F}$*, then* $\Theta(\mathbf{U})(\omega) := \mathbb{E}(\mathbf{U} \mid \mathcal{B})(\omega) - \mathbf{U}(\omega)$*, is continuous.*

Proof : We have

$$
\begin{aligned}
||\Theta(\mathbf{U}_n) - \Theta(\mathbf{U})||_p &\leq ||\mathbf{U}_n - \mathbf{U}||_p + ||\mathbb{E}(\mathbf{U}_n - \mathbf{U} \mid \mathcal{B})||_p \\
&\leq 2||\mathbf{U}_n - \mathbf{U}||_p \to 0
\end{aligned}
$$

■

Risk constraints

Roughly speaking a convex risk measure is defined as

$$\rho(\mathbf{X}) = \max_{\mathbb{Q}\in\mathcal{Q}} \mathbb{E}_{\mathbb{Q}}\mathbf{X} - \rho^*(\mathbb{Q}) \ ,$$

where $\mathcal{Q}$ is a closed convex set of distributions.
We will here consider real valued control.

Theorem

If ρ is a convex risk function, such that $\mathcal{U} \subset int(dom(\rho))$, and $a \in \mathbb{R}$, then $\Theta\big(\mathbf{U}\big)(\omega) := \rho\big(\mathbf{U}\big) - a$ is continuous.

Proof : We define $\Theta : \mathcal{U} \to \mathbb{R}$ as $\Theta\big(\mathbf{U}\big)(\omega) := \rho\big(\mathbf{U}\big) - a$. This function is convex, and thus continuous on the interior of it's domain. ■

VaR constraint

Another risk measure that is widely used even if it has some serious drawback is the Value-at-Risk. If X is a real random variable its value at risk of level α can be defined as $VaR_\alpha(\mathbf{X}) := \inf\{F_{\mathbf{X}}^{-1}(\alpha)\}$ where $F_{\mathbf{X}}(x) := \mathbb{P}(\mathbf{X} \leq x)$.

Theorem

If $\mathcal{U}$ is such that every $\mathbf{U} \in \mathcal{U}$ have a positive density, then $\Theta(\mathbf{U}) := VaR_\alpha\big(\mathbf{U}\big)$ is continuous if we have endowed $\mathcal{U}$ with the topology of convergence in law.

Proof : By definition of convergence in law, if $\mathbf{U}_n \to \mathbf{U}$ in law, we have $\forall x \in \mathbb{R} \quad F_{\mathbf{U}_n}(x) \to F_{\mathbf{U}}(x)$, and $F_{\mathbf{U}}^{-1}$ is continuous which means that $\Theta(\mathbf{U}_n) \to \Theta(\mathbf{U})$. ∎

Attouch, H. (1984).
Variational convergence for functions and operators.
Pitman Advanced Pub. Program

Piccinini, L. (1998).
A new version of the multivalued Fatou lemma.
Journal of Applied Analysis

Kudo, H. (1974).
A Note on the Strong Convergence of σ-Algebras.
The Annals of Probability

Barty, K., Carpentier, P. and Girardeau, P. (2010)
Decomposition of large-scale stochastic optimal control problems.
RAIRO Operations Research, 2010, 44, 167-183.

Barty, K., Carpentier, P., Cohen, G. and Girardeau, P. (2010)
Price decomposition in large-scale stochastic optimal control.
arXiv, 2010, math.OC, 1012.2092.

The end

Thank you for your attention !