Learning with Decision-Dependent Uncertainty

Tito Homem-de-Mello

Department of Mechanical and Industrial Engineering
University of Illinois at Chicago

NSF Workshop, University of Maryland, May 2010

Learning and deciding

Many decision-making problems fit the following framework:

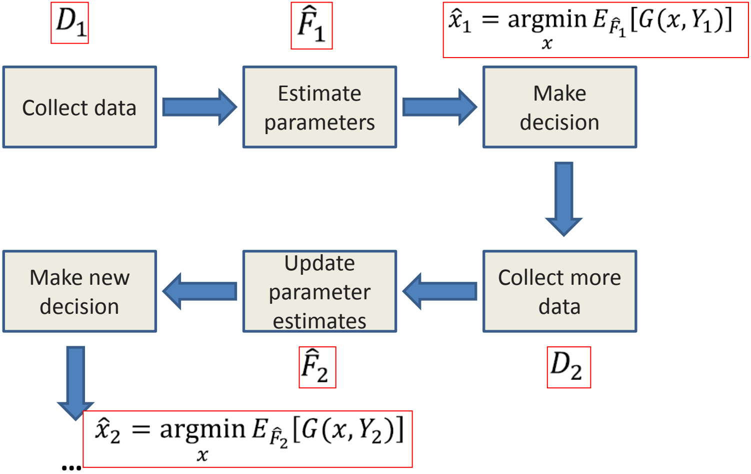

Convergence

Does this process converge?

- This is clearly the case when the observed data are i.i.d. observations from a given distribution F.
  - In that case, $\hat{F}_n$ is the empirical distribution corresponding to samples from F, so under mild conditions $\hat{x}_n \to x^*$.
- But what if the demand depends on the decisions?
  - Many cases (e.g. some revenue management problems) fall into that framework.
  - In that case, the true distribution is $F = F(x, \cdot)$.
  - If the dependence is ignored, the forecast/decision process converges to x° satisfying $x^\circ = \text{argmin } E_{F(x^\circ,\cdot)}[G(x, Y)]$, which can be suboptimal. (Cooper+HM+Kleywegt 2006, Lee+HM+Kleywegt 2009)

Convergence

Does this process converge?

- This is clearly the case when the observed data are i.i.d. observations from a given distribution F.
  - In that case, $\hat{F}_n$ is the empirical distribution corresponding to samples from F, so under mild conditions $\hat{x}_n \to x^*$.
- But what if the demand depends on the decisions?
  - Many cases (e.g. some revenue management problems) fall into that framework.
  - In that case, the true distribution is $F = F(x, \cdot)$.
  - If the dependence is ignored, the forecast/decision process converges to x° satisfying $x^\circ = \operatorname{argmin} E_{F(x^\circ, \cdot)}[G(x, Y)]$, which can be suboptimal. (Cooper+HM+Kleywegt 2006, Lee+HM+Kleywegt 2009)

Convergence

Does this process converge?

- This is clearly the case when the observed data are i.i.d. observations from a given distribution F.
  - In that case, $\hat{F}_n$ is the empirical distribution corresponding to samples from F, so under mild conditions $\hat{x}_n \to x^*$.
- But what if the demand depends on the decisions?
  - Many cases (e.g. some revenue management problems) fall into that framework.
  - In that case, the true distribution is $F = F(x, \cdot)$.
  - If the dependence is ignored, the forecast/decision process converges to x° satisfying $x^\circ = \operatorname{argmin} E_{F(x^\circ,\cdot)}[G(x, Y)]$, which can be suboptimal. (Cooper+HM+Kleywegt 2006, Lee+HM+Kleywegt 2009)

Convergence

Does this process converge?

- This is clearly the case when the observed data are i.i.d. observations from a given distribution F.
    - In that case, $\hat{F}_n$ is the empirical distribution corresponding to samples from F, so under mild conditions $\hat{x}_n \to x^*$.
- But what if the demand depends on the decisions?
    - Many cases (e.g. some revenue management problems) fall into that framework.
    - In that case, the true distribution is $F = F(x, \cdot)$.
    - If the dependence is ignored, the forecast/decision process converges to x° satisfying $x^\circ = \text{argmin } E_{F(x^\circ,\cdot)}[G(x, Y)]$, which can be suboptimal. (Cooper+HM+Kleywegt 2006, Lee+HM+Kleywegt 2009)

Learning the dependence

Can we learn the structure of the dependence of F on x, and simultaneously optimize the objective function?

- For example, consider the simple case where
  - $Y_n = \beta_0 + \beta_1 x + \varepsilon_n$, where β_0 and β_1 are unknown constants and $\{\varepsilon_n\}$ is i.i.d.
  - $G(x, Y) = xY$.
- At iteration n,
  - Estimate β_0 and β_1 from $\hat{x}_1, \ldots, \hat{x}_n$ using least-squares;
  - Compute $\hat{x}_{n+1} = \dfrac{-\hat{\beta}_0^n}{2\hat{\beta}_1^n}$.
- This kind of procedure is called certainty equivalent control in the control and econometrics literature.
- Very hard to show convergence!

Learning the dependence

Can we learn the structure of the dependence of F on x, and simultaneously optimize the objective function?

- For example, consider the simple case where
  - $Y_n = \beta_0 + \beta_1 x + \varepsilon_n$, where β_0 and β_1 are unknown constants and $\{\varepsilon_n\}$ is i.i.d.
  - $G(x, Y) = xY$.
- At iteration n,
  - Estimate β_0 and β_1 from $\hat{x}_1, \ldots, \hat{x}_n$ using least-squares;
  - Compute $\hat{x}_{n+1} = \dfrac{-\hat{\beta}_0^n}{2\hat{\beta}_1^n}$.
- This kind of procedure is called certainty equivalent control in the control and econometrics literature.
- Very hard to show convergence!

Learning the dependence

Can we learn the structure of the dependence of F on x, and simultaneously optimize the objective function?

- For example, consider the simple case where
  - $Y_n = \beta_0 + \beta_1 x + \varepsilon_n$, where β_0 and β_1 are unknown constants and $\{\varepsilon_n\}$ is i.i.d.
  - $G(x, Y) = xY$.
- At iteration n,
  - Estimate β_0 and β_1 from $\hat{x}_1, \ldots, \hat{x}_n$ using least-squares;
  - Compute $\hat{x}_{n+1} = \dfrac{-\hat{\beta}_0^n}{2\hat{\beta}_1^n}$.
- This kind of procedure is called certainty equivalent control in the control and econometrics literature.
- Very hard to show convergence!

Learning the dependence

Can we learn the structure of the dependence of F on x, and simultaneously optimize the objective function?

- For example, consider the simple case where
  - $Y_n = \beta_0 + \beta_1 x + \varepsilon_n$, where β_0 and β_1 are unknown constants and $\{\varepsilon_n\}$ is i.i.d.
  - $G(x, Y) = xY$.
- At iteration n,
  - Estimate β_0 and β_1 from $\hat{x}_1, \ldots, \hat{x}_n$ using least-squares;
  - Compute $\hat{x}_{n+1} = \dfrac{-\hat{\beta}_0^n}{2\hat{\beta}_1^n}$.
- This kind of procedure is called certainty equivalent control in the control and econometrics literature.
- Very hard to show convergence!

Can we use stochastic approximation?

- We could apply an SA-type procedure to the function $g(x) = E[G(x, Y)]$.
- We cannot get derivatives estimators if we don't know how Y depends on x.
- Even if we know the dependence form, it may be hard to prove convergence.
  - Again, suppose $Y = \beta_0 + \beta_1 x + \varepsilon$, $G(x, Y) = xY$. Then, $g'(x) = \beta_0 + 2\beta_1 x$.
  - We can estimate β_0 and β_1 from regression, but does it converge?
- We could use finite differences, but convergence is slow.

Can we use stochastic approximation?

- We could apply an SA-type procedure to the function $g(x) = E[G(x, Y)]$.
- We cannot get derivatives estimators if we don't know how Y depends on x.
- Even if we know the dependence form, it may be hard to prove convergence.
  - Again, suppose $Y = \beta_0 + \beta_1 x + \varepsilon$, $G(x, Y) = xY$. Then, $g'(x) = \beta_0 + 2\beta_1 x$.
  - We can estimate β_0 and β_1 from regression, but does it converge?
- We could use finite differences, but convergence is slow.

Can we use stochastic approximation?

- We could apply an SA-type procedure to the function $g(x) = E[G(x, Y)]$.
- We cannot get derivatives estimators if we don't know how Y depends on x.
- Even if we know the dependence form, it may be hard to prove convergence.
  - Again, suppose $Y = \beta_0 + \beta_1 x + \varepsilon$, $G(x, Y) = xY$. Then, $g'(x) = \beta_0 + 2\beta_1 x$.
  - We can estimate β_0 and β_1 from regression, but does it converge?
- We could use finite differences, but convergence is slow.

Can we use stochastic approximation?

- We could apply an SA-type procedure to the function $g(x) = E[G(x, Y)]$.
- We cannot get derivatives estimators if we don't know how Y depends on x.
- Even if we know the dependence form, it may be hard to prove convergence.
  - Again, suppose $Y = \beta_0 + \beta_1 x + \varepsilon$, $G(x, Y) = xY$. Then, $g'(x) = \beta_0 + 2\beta_1 x$.
  - We can estimate β_0 and β_1 from regression, but does it converge?
- We could use finite differences, but convergence is slow.

Comparing the methods

$\beta_0 = 5,\ \beta_1 = 1,\ \varepsilon =$Normal(0,1)

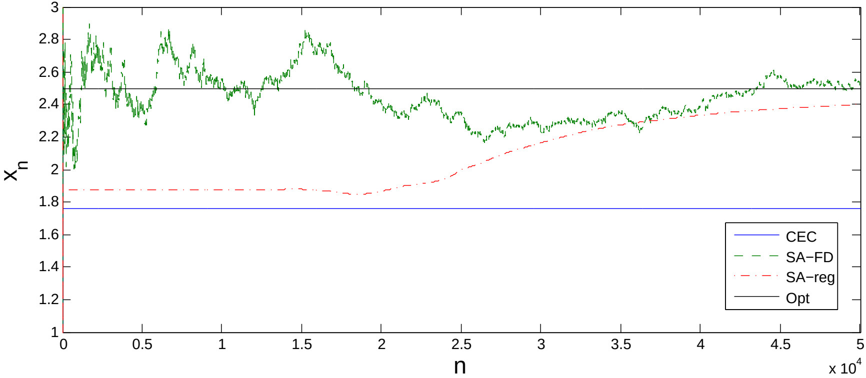

Issues

- Is it better to first learn the dependence structure, and then optimize? (e.g., Besbes and Zeevi 2010 in the context of demand functions)
- Alternatively, we can view the problem as a "black-box" system where we only observe outputs (in this case, revenues) in terms of the inputs (the decisions).
- We want to optimize — *however*, in the context we are interested in, function evaluations are expensive.
- We are looking at methods that choose decision points randomly according to some distribution, which is updated from iteration to iteration.

Issues

- Is it better to first learn the dependence structure, and then optimize? (e.g., Besbes and Zeevi 2010 in the context of demand functions)
- Alternatively, we can view the problem as a "black-box" system where we only observe outputs (in this case, revenues) in terms of the inputs (the decisions).
- We want to optimize — *however*, in the context we are interested in, function evaluations are expensive.
- We are looking at methods that choose decision points randomly according to some distribution, which is updated from iteration to iteration.

Issues

- Is it better to first learn the dependence structure, and then optimize? (e.g., Besbes and Zeevi 2010 in the context of demand functions)
- Alternatively, we can view the problem as a "black-box" system where we only observe outputs (in this case, revenues) in terms of the inputs (the decisions).
- We want to optimize — *however*, in the context we are interested in, function evaluations are expensive.
- We are looking at methods that choose decision points randomly according to some distribution, which is updated from iteration to iteration.

Issues

- Is it better to first learn the dependence structure, and then optimize? (e.g., Besbes and Zeevi 2010 in the context of demand functions)
- Alternatively, we can view the problem as a "black-box" system where we only observe outputs (in this case, revenues) in terms of the inputs (the decisions).
- We want to optimize — *however*, in the context we are interested in, function evaluations are expensive.
- We are looking at methods that choose decision points randomly according to some distribution, which is updated from iteration to iteration.