18.175: Lecture 13

More large deviations

Scott Sheffield

MIT

Outline

Outline

Legendre transform

Large deviations

Legendre transform

- Define **Legendre transform** (or Legendre dual) of a function $\Lambda : \mathbb{R}^d \to \mathbb{R}$ by

$$\Lambda^*(x) = \sup_{\lambda \in \mathbb{R}^d} \{(\lambda, x) - \Lambda(\lambda)\}.$$

- Let's describe the Legendre dual geometrically if $d = 1$: $\Lambda^*(x)$ is where tangent line to Λ of slope x intersects the real axis. We can "roll" this tangent line around the convex hull of the graph of Λ, to get all Λ^* values.
- Is the Legendre dual always convex?
- What is the Legendre dual of x^2? Of the function equal to 0 at 0 and ∞ everywhere else?
- How are derivatives of Λ and Λ^* related?
- What is the Legendre dual of the Legendre dual of a convex function?
- What's the higher dimensional analog of rolling the tangent line?

Outline

Legendre transform

Large deviations

Outline

Legendre transform

Large deviations

Recall: moment generating functions

- Let X be a random variable.
- The **moment generating function** of X is defined by $M(t) = M_X(t) := E[e^{tX}]$.
- When X is discrete, can write $M(t) = \sum_x e^{tx} p_X(x)$. So $M(t)$ is a weighted average of countably many exponential functions.
- When X is continuous, can write $M(t) = \int_{-\infty}^{\infty} e^{tx} f(x) dx$. So $M(t)$ is a weighted average of a continuum of exponential functions.
- We always have $M(0) = 1$.
- If $b > 0$ and $t > 0$ then $E[e^{tX}] \geq E[e^{t \min\{X,b\}}] \geq P\{X \geq b\} e^{tb}$.
- If X takes both positive and negative values with positive probability then $M(t)$ grows at least exponentially fast in $|t|$ as $|t| \to \infty$.

Recall: moment generating functions for i.i.d. sums

- We showed that if $Z = X + Y$ and X and Y are independent, then $M_Z(t) = M_X(t)M_Y(t)$
- If $X_1 \ldots X_n$ are i.i.d. copies of X and $Z = X_1 + \ldots + X_n$ then what is M_Z?
- Answer: M_X^n.

Large deviations

- Consider i.i.d. random variables X_i. Can we show that $P(S_n \geq na) \to 0$ exponentially fast when $a > E[X_i]$?
- Kind of a quantitative form of the weak law of large numbers. The empirical average A_n is *very* unlikely to ϵ away from its expected value (where "very" means with probability less than some exponentially decaying function of n).

General large deviation principle

- More general framework: a *large deviation principle* describes limiting behavior as $n \to \infty$ of family $\{\mu_n\}$ of measures on measure space $(\mathcal{X}, \mathcal{B})$ in terms of a *rate function* I.
- The **rate function** is a lower-semicontinuous map $I : \mathcal{X} \to [0, \infty]$. (The sets $\{x : I(x) \le a\}$ are closed — rate function called "good" if these sets are compact.)
- **DEFINITION:** $\{\mu_n\}$ satisfy LDP with rate function I and speed n if for all $\Gamma \in \mathcal{B}$,

$$-\inf_{x \in \Gamma^0} I(x) \le \liminf_{n\to\infty} \frac{1}{n} \log \mu_n(\Gamma) \le \limsup_{n\to\infty} \frac{1}{n} \log \mu_n(\Gamma) \le -\inf_{x \in \overline{\Gamma}} I(x).$$

- **INTUITION:** when "near x" the probability density function for μ_n is tending to zero like $e^{-I(x)n}$, as $n \to \infty$.
- **Simple case:** I is continuous, Γ is closure of its interior.
- **Question:** How would I change if we replaced the measures μ_n by weighted measures $e^{(\lambda n, \cdot)}\mu_n$?
- Replace $I(x)$ by $I(x) - (\lambda, x)$? What is $\inf_x I(x) - (\lambda, x)$?

Cramer's theorem

- Let μ_n be law of empirical mean $A_n = \frac{1}{n}\sum_{j=1}^n X_j$ for i.i.d. vectors $X_1, X_2, \ldots, X_n$ in $\mathbb{R}^d$ with same law as X.
- Define **log moment generating function** of X by

$$\Lambda(\lambda) = \Lambda_X(\lambda) = \log M_X(\lambda) = \log \mathbb{E}e^{(\lambda, X)},$$

  where $(\cdot,\cdot)$ is inner product on $\mathbb{R}^d$.
- Define **Legendre transform** of Λ by

$$\Lambda^*(x) = \sup_{\lambda \in \mathbb{R}^d} \{(\lambda, x) - \Lambda(\lambda)\}.$$

- **CRAMER'S THEOREM:** μ_n satisfy LDP with convex rate function Λ^*.

Thinking about Cramer's theorem

- Let μ_n be law of empirical mean $A_n = \frac{1}{n}\sum_{j=1}^n X_j$.
- **CRAMER'S THEOREM:** μ_n satisfy LDP with convex rate function

$$I(x) = \Lambda^*(x) = \sup_{\lambda \in \mathbb{R}^d} \{(\lambda, x) - \Lambda(\lambda)\},$$

where $\Lambda(\lambda) = \log M(\lambda) = \mathbb{E}e^{(\lambda, X_1)}$.
- This means that for all $\Gamma \in \mathcal{B}$ we have this **asymptotic lower bound** on probabilities $\mu_n(\Gamma)$

$$-\inf_{x \in \Gamma^0} I(x) \le \liminf_{n\to\infty} \frac{1}{n} \log \mu_n(\Gamma),$$

so (up to sub-exponential error) $\mu_n(\Gamma) \ge e^{-n \inf_{x\in\Gamma^0} I(x)}$.
- and this **asymptotic upper bound** on the probabilities $\mu_n(\Gamma)$

$$\limsup_{n\to\infty} \frac{1}{n} \log \mu_n(\Gamma) \le -\inf_{x\in\overline{\Gamma}} I(x),$$

which says (up to subexponential error) $\mu_n(\Gamma) \le e^{-n \inf_{x\in\overline{\Gamma}} I(x)}$.

Proving Cramer upper bound

- Recall that $I(x) = \Lambda^*(x) = \sup_{\lambda \in \mathbb{R}^d}\{(\lambda, x) - \Lambda(\lambda)\}$.
- For simplicity, assume that Λ is defined for all x (which implies that X has moments of all orders and Λ and Λ^* are strictly convex, and the derivatives of Λ and Λ' are inverses of each other). It is also enough to consider the case X has mean zero, which implies that $\Lambda(0) = 0$ is a minimum of Λ, and $\Lambda^*(0) = 0$ is a minimum of Λ^*.
- We aim to show (up to subexponential error) that $\mu_n(\Gamma) \leq e^{-n \inf_{x \in \overline{\Gamma}} I(x)}$.
- If Γ were singleton set $\{x\}$ we could find the λ corresponding to x, so $\Lambda^*(x) = (x, \lambda) - \Lambda(\lambda)$. Note then that
$$\mathbb{E}e^{(n\lambda, A_n)} = \mathbb{E}e^{(\lambda, S_n)} = M_X^n(\lambda) = e^{n\Lambda(\lambda)},$$
and also $\mathbb{E}e^{(n\lambda, A_n)} \geq e^{n(\lambda, x)}\mu_n\{x\}$. Taking logs and dividing by n gives $\Lambda(\lambda) \geq \frac{1}{n}\log \mu_n + (\lambda, x)$, so that $\frac{1}{n}\log \mu_n(\Gamma) \leq -\Lambda^*(x)$, as desired.
- General Γ: cut into finitely many pieces, bound each piece?

Proving Cramer lower bound

- Recall that $I(x) = \Lambda^*(x) = \sup_{\lambda \in \mathbb{R}^d}\{(\lambda, x) - \Lambda(\lambda)\}$.
- We aim to show that asymptotically $\mu_n(\Gamma) \geq e^{-n \inf_{x \in \Gamma^0} I(x)}$.
- It's enough to show that for each given $x \in \Gamma^0$, we have that asymptotically $\mu_n(\Gamma) \geq e^{-n \inf_{x \in \Gamma^0} I(x)}$.
- Idea is to weight the law of X by $e^{(\lambda, x)}$ for some λ and normalize to get a new measure whose expectation is this point x. In this new measure, A_n is "typically" in Γ for large Γ, so the probability is of order 1.
- But by how much did we have to modify the measure to make this typical? Not more than by factor $e^{-n \inf_{x \in \Gamma^0} I(x)}$.

MIT OpenCourseWare
http://ocw.mit.edu

18.175 Theory of Probability
Spring 2014

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.