TECHNICAL DERIVATIONS FROM CS5560 (Probabilistic Models for ML)

① MLE from first principles.

Let $p^*(x)$ denote the unknown likelihood being modeled.

Let $p_\theta(x)$ denote the ~~dash~~ likelihood corresponding to model parameter θ.

Our goal is to minimize KL divergence between them.

Accordingly, $\theta^{MLE} \equiv \underset{\theta \in \Theta}{\operatorname{argmin}} \; KL\left(p^*(x) \,||\, p_\theta(x)\right)$, where Θ is set of model parameters.

$$= \underset{\theta \in \Theta}{\operatorname{argmin}} \; \mathbb{E}_{x \sim p^*(x)}\left[\log\left(\frac{p^*(x)}{p_\theta(x)}\right)\right] \qquad (\because \text{defn. of KL})$$

$$= \underset{\theta \in \Theta}{\operatorname{argmin}} - \mathbb{E}_{x \sim p^*(x)}\left[\log p_\theta(x)\right] \qquad (\because \text{prop. of log \& } p^* \text{ not dependent on } \theta \text{ \& linearity of } \mathbb{E})$$

$$= \underset{\theta \in \Theta}{\operatorname{argmax}} \; \mathbb{E}_{x \sim p^*(x)}\left[\log p_\theta(x)\right] \qquad (\because \min_x -f(x) = -\max_x f(x))$$

$$\approx \underset{\theta \in \Theta}{\operatorname{argmax}} \; \frac{\sum_{i=1}^{m} \log p_\theta(x_i)}{m} \qquad (\because \text{Using sample mean approximation})$$

$$= \underset{\theta \in \Theta}{\operatorname{argmax}} \; \prod_{i=1}^{m} p_\theta(x_i) \qquad (\because \text{prop. of log, exp, monotonicity})$$

$$= \underset{\theta \in \Theta}{\operatorname{argmax}} \; p_\theta(\mathcal{D}) \qquad (\because \text{Samples are independent})$$

Here, argmin denotes ~~the~~ optimal value of the variable at which min. is attained.

|||ly argmax " " " " " max. " " .

② Derivation of MCLE from first principles.

Let $p^*(y/x)$ be the unknown posterior being modeled.

Let $p_\theta(y/x)$ be the posterior likelihood at parameter θ.

Our goal is to minimize the expected KL divergence between these two.

Accordingly, $\theta^{MCLE} \equiv \underset{\theta \in \Theta}{\operatorname{argmin}} \; E_{x \sim p^*(x)}\left[KL\left(p^*(y/x) \,||\, p_\theta(y/x)\right)\right]$ (here, Θ is set of parameters) ($p^*(x)$ is the corresponding unknown marginal.)

$$= \underset{\theta \in \Theta}{\operatorname{argmin}} \; E_{x \sim p^*(x)}\left[E_{y/x \sim p^*(y/x)}\left[\log \frac{p^*(y/x)}{p_\theta(y/x)}\right]\right] \quad (\because \text{defn. of } KL)$$

$$= \underset{\theta \in \Theta}{\operatorname{argmin}} \; E_{x,y \sim p^*(x,y)}\left[\log \frac{p^*(y/x)}{p_\theta(y/x)}\right] \quad (\because \text{Total expectation rule})$$

$$= \underset{\theta \in \Theta}{\operatorname{argmin}} \; -E_{x,y \sim p^*(x,y)}\left[\log p_\theta(y/x)\right] \quad (\because \text{linearity of } E \text{ \& prop. of log; } p^* \text{ is independent of } \theta)$$

$$= \underset{\theta \in \Theta}{\operatorname{argmax}} \; E_{x,y \sim p^*(x,y)}\left[\log p_\theta(y/x)\right] \quad (\because \text{minimizer of } f \text{ is same as maximizer of } -f)$$

$$\approx \underset{\theta \in \Theta}{\operatorname{argmax}} \; \frac{\sum_{i=1}^{m} \log p_\theta(y_i/x_i)}{m} \quad (\because \text{sample mean approximation})$$

$$= \underset{\theta \in \Theta}{\operatorname{argmax}} \; \underbrace{\prod_{i=1}^{m} p_\theta(y_i/x_i)}$$

→ defined as conditional likelihood of training data (at θ).

Scanned by CamScanner

Theorem 1: Let $f: \mathcal{X} \to \mathbb{R}$ be a differentiable function, where $\mathcal{X}$ is open $\subseteq \mathbb{R}$.

$$x^* \text{ is local minima/maxima} \implies f'(x^*) = 0$$

Comment:
Hence, a valid strategy for finding maximizers/minimizers for f would be to look at all points where slope is zero, and then find at which of those "critical" points (in $\mathcal{X}$), the function is maximized/minimized!.

③ Derivation of MLE for Bernoulli model.

$\cancel{p(\mathcal{D})}$ $p_\theta(\mathcal{D}) = \theta^{m_1}(1-\theta)^{m_2}$, where m_1 is the number of samples with first outcome.
m_2 " " " " second "

Hence. $\theta^{MLE} \equiv \underset{\theta \in [0,1]}{\operatorname{argmax}}\ \theta^{m_1}(1-\theta)^{m_2}$.

$$= \begin{cases} 1 & \text{if } m_2 = 0 \\ 0 & \text{if } m_1 = 0 \\ \underset{\theta \in (0,1)}{\operatorname{argmax}}\ \theta^{m_1}(1-\theta)^{m_2} & \text{else} \end{cases} \quad \text{I} \quad (\because \text{ objective is non-zero in } (0,1))$$

critical points θ^* $\Rightarrow m_1 \theta^{*\,m_1-1}(1-\theta^*)^{m_2} - \theta^{*\,m_1} m_2 (1-\theta^*)^{m_2-1} = 0$

$$\Leftrightarrow \theta^* = \frac{m_1}{m_1+m_2}.$$

~~Since there is only one critical point~~

Since this is unique and at end points the function $\to 0$, this must be a global maximizer.

Alternatively, one can check the sign of double derivative to see it is < 0 at θ^*.

$$\therefore \theta^{MLE} = \frac{m_1}{m_1+m_2}$$

$(\because$ it covers all the three cases in I$)$

(4) Derivation of MLE for Multinoulli model.

For multinoulli, $p_\theta(\mathcal{D}) = \theta_1^{m_1}\theta_2^{m_2}\cdots\theta_{c-1}^{m_{c-1}}\theta_c^{m_c}$, where m_i is no. samples with i^{th} outcome

$\theta_i \geq 0 \ \forall i, \ \sum_{i=1}^{c}\theta_i = 1.$

Let's eliminate θ_c using the equality $\sum_{i=1}^{c}\theta_i = 1$

$$\Rightarrow \theta_c = 1 - \sum_{i=1}^{c-1}\theta_i$$

But we need to still contrain $\theta_c \geq 0 \iff 1-\sum_{i=1}^{c-1}\theta_i \geq 0.$

Hence, $p_\theta(\mathcal{D}) = \theta_1^{m_1}\theta_2^{m_2}\cdots\theta_{c-2}^{m_{c-2}}\theta_{c-1}^{m_{c-1}}\left(1-\sum_{i=1}^{c-1}\theta_i\right)^{m_c}$, (I)

$\theta_1 \geq 0, \ldots, \theta_{c-1} \geq 0, 1-\sum_{i=1}^{c-1}\theta_i \geq 0.$

Now,

$$\max_{\theta_1\geq 0,\ldots,\theta_{c-1}\geq 0, 1-\sum_{i=1}^{c-1}\theta_i \geq 0} p_\theta(\mathcal{D})$$

$$= \max_{\theta_1 \geq 0,\ldots,\theta_{c-2}\geq 0} \ \max_{\substack{\theta_{c-1}\geq 0, \\ \theta_{c-1} \leq 1-\sum_{i=1}^{c-2}\theta_i}} p_\theta(\mathcal{D})$$

$\left(\because \text{order of variables does not matter in optimization}\right)$

$$= \max_{\theta_1\geq 0,\ldots,\theta_{c-2}\geq 0} \theta_1^{m_1}\theta_2^{m_2}\cdots\theta_{c-2}^{m_{c-2}} \cdot \max_{\theta_{c-1}\in[0,\alpha]} \theta_{c-1}^{m_{c-1}}\left(\alpha-\theta_{c-1}\right)^{m_c},$$

where $\alpha \equiv 1-\sum_{i=1}^{c-2}\theta_i$

$$= \max_{\theta_1\geq 0,\ldots\theta_{c-2}\geq 0} \theta_1^{m_1}\theta_2^{m_2}\cdots\theta_{c-2}^{m_{c-2}} \max_{\bar\theta\in[0,1]} \alpha^{m_{c-1}+m_c}\bar\theta^{m_{c-1}}(1-\bar\theta)^{m_c},$$

where $\bar\theta = \frac{\theta_{c-1}}{\alpha}.$

But, from (3) we have $\operatorname{argmax}_{\bar\theta\in[0,1]} \bar\theta^{m_{c-1}}(1-\bar\theta)^{m_c} = \frac{m_{c-1}}{m_{c-1}+m_c}.$

$$\iff \theta_{c-1}^* = \frac{\left(1-\sum_{i=1}^{c-2}\theta_i\right)m_{c-1}}{m_{c-1}+m_c} \quad (\because \ )$$

Substituting this θ_{c-1}^* in the objective, we obtain

$$\underset{\theta \in \Delta}{\arg\max}\ p_\theta(\mathcal{D}) = \underset{\theta_1 \geq 0 \ldots \theta_{c-2} \geq 0}{\arg\max}\ \theta_1^{m_1} \cdots \cdots \theta_{c-2}^{m_{c-2}} \left(1 - \sum_{i=1}^{c-2} \theta_i\right)^{m_{c-1}+m_c}$$

($\because$ ignoring terms not involving the variables)

But, this problem has the same form as (I).

$$\text{So, } \theta_{c-2}^* = \frac{\left(1 - \sum_{i=1}^{c-3} \theta_i\right) m_{c-2}}{m_{c-2} + m_{c-1} + m_c}.$$

$\vdots$

$$\theta_2^* = \frac{(1-\theta_1)\, m_2}{m_2 + m_3 + \cdots + m_{c-1} + m_c}$$

$$\theta_1^* = \frac{1 \cdot m_1}{m_1 + m_2 + \cdots + m_c}$$

Substituting values bottom-up, we obtain $\underline{\underline{\theta_i^* = \frac{m_i}{\sum_{j=1}^{c} m_j}}}$

Date ___/___/___
Page ______
SCHOOL MATE

Theorem 2: Let $f : \mathcal{X} \subseteq \mathbb{R}^n \to \mathbb{R}$ be a twice differentiable function. Then Let $\mathcal{X}$ be an open set.

Then,

(I) x^* is local max./min $\Longrightarrow \nabla f(x^*) = 0$

&,

(II) $\nabla f(x^*) = 0, \ \underbrace{\nabla^2 f(x^*) \succ 0}_{\text{Hessian of } f \text{ at } x^* \text{ is positive definite.}} \Rightarrow x^*$ is local min.

&, (III) ~~Theorem 3~~: ~~If~~ $\nabla^2 f(x) \succeq 0 \ \forall x \in \mathcal{X} \Rightarrow f$ is convex.
$\mathcal{X}$ is convex set

~~Theorem 3~~ **Comments:**

(1) The assumptions are critical for truthfulness of the theorem.

For eg.

(a) if $\mathcal{X}$ is not open, then (I) need not be true!

(b) if $\nabla^2 f(x^*) \succeq 0, \nabla f(x^*) = 0$, then (II) need not "

(c) if $\mathcal{X}$ is not open & $\mathcal{X}$ is not convex then (III) " "

Date ___/___/___
Page ______
SCHOOL MATE

Theorem 3: If f is convex, then every local min. is a global min.

e.g. From (8.7) in Murphy's book, it is easy to verify that the negative log-likelihood in logistic regression is convex.

Theorem 4: Let $\{x^{(n)}\}$ be the sequence of points obtained with gradient descent, with stepsizes chosen using Armijo's rule. Then every limit point of $\{x^{(n)}\}$ is a stationary point.

Comments

① The theorem does not guarantee convergence, But it is as good because: for eg. let the iterates alternate between two points. Then both these will be stationary points!

② If f is convex, then convergence is guaranteed!

Scanned by CamScanner