Last lecture, we introduced the method of moments for estimating one or more parameters $\theta$ in a parametric model. This lecture, we discuss a different method called maximum likelihood estimation. The focus of this lecture will be on how to compute this estimate; subsequent lectures will study its statistical properties.

### 13.1 Maximum likelihood estimation

Consider data $X_{1}, \ldots, X_{n} \stackrel{I I D}{\sim} f(x \mid \theta)$, for a parametric model $\{f(x \mid \theta): \theta \in \Omega\}$. Given the observed values $X_{1}, \ldots, X_{n}$ of the data, the function

$$
\operatorname{lik}(\theta)=f\left(X_{1} \mid \theta\right) \times \ldots \times f\left(X_{n} \mid \theta\right)
$$

of the parameter $\theta$ is called the likelihood function. If $f(x \mid \theta)$ is the PMF of a discrete distribution, then $\operatorname{lik}(\theta)$ is simply the probability of observing the values $X_{1}, \ldots, X_{n}$ if the true parameter were $\theta$. The maximum likelihood estimator (MLE) of $\theta$ is the value of $\theta \in \Omega$ that maximizes $\operatorname{lik}(\theta)$. Intuitively, it is the value of $\theta$ that makes the observed data "most probable" or "most likely".

The idea of maximum likelihood is related to the use of the likelihood ratio statistic in the Neyman-Pearson lemma. Recall that for testing

$$
\begin{aligned}
& H_{0}:\left(X_{1}, \ldots, X_{n}\right) \sim g \\
& H_{1}:\left(X_{1}, \ldots, X_{n}\right) \sim h
\end{aligned}
$$

where $g$ and $h$ are joint PDFs or PMFs for $n$ random variables, the most powerful test rejects for small values of the likelihood ratio

$$
L\left(X_{1}, \ldots, X_{n}\right)=\frac{g\left(X_{1}, \ldots, X_{n}\right)}{h\left(X_{1}, \ldots, X_{n}\right)}
$$

In the context of a parametric model, we may consider testing $H_{0}: X_{1}, \ldots, X_{n} \stackrel{I I D}{\sim} f\left(x \mid \theta_{0}\right)$ versus $H_{1}: X_{1}, \ldots, X_{n} \stackrel{I I D}{\sim} f\left(x \mid \theta_{1}\right)$, for two different parameter values $\theta_{0}, \theta_{1} \in \Omega$. Then

$$
\begin{aligned}
& g\left(X_{1}, \ldots, X_{n}\right)=f\left(X_{1} \mid \theta_{0}\right) \times \ldots \times f\left(X_{n} \mid \theta_{0}\right), \\
& h\left(X_{1}, \ldots, X_{n}\right)=f\left(X_{1} \mid \theta_{1}\right) \times \ldots \times f\left(X_{n} \mid \theta_{1}\right),
\end{aligned}
$$

so the likelihood ratio is exactly $\operatorname{lik}\left(\theta_{0}\right) / \operatorname{lik}\left(\theta_{1}\right)$. The MLE (if it exists and is unique) is the value of $\theta \in \Omega$ for which $\operatorname{lik}(\theta) / \operatorname{lik}\left(\theta^{\prime}\right)>1$ for any other value $\theta^{\prime} \in \Omega$.

### 13.2 Examples

Computing the MLE is an optimization problem. Maximizing $\operatorname{lik}(\theta)$ is equivalent to maximizing its (natural) logarithm

$$
l(\theta)=\log (\operatorname{lik}(\theta))=\sum_{i=1}^{n} \log f\left(X_{i} \mid \theta\right)
$$

which in many examples is easier to work with as it involves a sum rather than a product. Let's work through several examples:
Example 13.1. Let $X_{1}, \ldots, X_{n} \stackrel{\text { IID }}{\sim} \operatorname{Poisson}(\lambda)$. Then

$$
\begin{aligned}
l(\lambda) & =\sum_{i=1}^{n} \log \frac{\lambda^{X_{i}} e^{-\lambda}}{X_{i}!} \\
& =\sum_{i=1}^{n}\left(X_{i} \log \lambda-\lambda-\log \left(X_{i}!\right)\right) \\
& =(\log \lambda) \sum_{i=1}^{n} X_{i}-n \lambda-\sum_{i=1}^{n} \log \left(X_{i}!\right) .
\end{aligned}
$$

This is differentiable in $\lambda$, so we maximize $l(\lambda)$ by setting its first derivative equal to 0 :

$$
0=l^{\prime}(\lambda)=\frac{1}{\lambda} \sum_{i=1}^{n} X_{i}-n
$$

Solving for $\lambda$ yields the estimate $\hat{\lambda}=\bar{X}$. Since $l(\lambda) \rightarrow-\infty$ as $\lambda \rightarrow 0$ or $\lambda \rightarrow \infty$, and since $\hat{\lambda}=\bar{X}$ is the unique value for which $0=l^{\prime}(\lambda)$, this must be the maximum of $l$. In this example, $\hat{\lambda}$ is the same as the method-of-moments estimate.
Example 13.2. Let $X_{1}, \ldots, X_{n} \stackrel{\text { IID }}{\sim} \mathcal{N}\left(\mu, \sigma^{2}\right)$. Then

$$
\begin{aligned}
l\left(\mu, \sigma^{2}\right) & =\sum_{i=1}^{n} \log \left(\frac{1}{\sqrt{2 \pi \sigma^{2}}} e^{-\frac{\left(X_{i}-\mu\right)^{2}}{2 \sigma^{2}}}\right) \\
& =\sum_{i=1}^{n}\left(-\frac{1}{2} \log \left(2 \pi \sigma^{2}\right)-\frac{\left(X_{i}-\mu\right)^{2}}{2 \sigma^{2}}\right) \\
& =-\frac{n}{2} \log (2 \pi)-\frac{n}{2} \log \left(\sigma^{2}\right)-\frac{1}{2 \sigma^{2}} \sum_{i=1}^{n}\left(X_{i}-\mu\right)^{2} .
\end{aligned}
$$

Considering $\sigma^{2}$ (rather than $\sigma$ ) as the parameter, we maximize $l(\lambda)$ by settings its partial derivatives with respect to $\mu$ and $\sigma^{2}$ equal to 0 :

$$
\begin{aligned}
& 0=\frac{\partial l}{\partial \mu}=\frac{1}{\sigma^{2}} \sum_{i=1}^{n}\left(X_{i}-\mu\right) \\
& 0=\frac{\partial l}{\partial \sigma^{2}}=-\frac{n}{2 \sigma^{2}}+\frac{1}{2 \sigma^{4}} \sum_{i=1}^{n}\left(X_{i}-\mu\right)^{2}
\end{aligned}
$$

Solving the first equation yields $\hat{\mu}=\bar{X}$, and substituting this into the second equation yields $\hat{\sigma}^{2}=\frac{1}{n} \sum_{i=1}^{n}\left(X_{i}-\bar{X}\right)^{2}$. Since $l\left(\mu, \sigma^{2}\right) \rightarrow-\infty$ as $\mu \rightarrow-\infty, \mu \rightarrow \infty, \sigma^{2} \rightarrow 0$, or $\sigma^{2} \rightarrow \infty$, and as $\left(\hat{\mu}, \hat{\sigma}^{2}\right)$ is the unique value for which $0=\frac{\partial l}{\partial \mu}$ and $0=\frac{\partial l}{\partial \sigma^{2}}$, this must be the maximum of $l$. Again, the MLEs are the same as the method-of-moments estimates.
Example 13.3. Let $X_{1}, \ldots, X_{n} \stackrel{\text { IID }}{\sim} \operatorname{Gamma}(\alpha, \beta)$. Then

$$
\begin{aligned}
l(\alpha, \beta) & =\sum_{i=1}^{n} \log \left(\frac{\beta^{\alpha}}{\Gamma(\alpha)} X_{i}^{\alpha-1} e^{-\beta X_{i}}\right) \\
& =\sum_{i=1}^{n}\left(\alpha \log \beta-\log \Gamma(\alpha)+(\alpha-1) \log X_{i}-\beta X_{i}\right) \\
& =n \alpha \log \beta-n \log \Gamma(\alpha)+(\alpha-1) \sum_{i=1}^{n} \log X_{i}-\beta \sum_{i=1}^{n} X_{i} .
\end{aligned}
$$

To maximize $l(\alpha, \beta)$, we set its partial derivatives equal to 0 :

$$
\begin{aligned}
& 0=\frac{\partial l}{\partial \alpha}=n \log \beta-\frac{n \Gamma^{\prime}(\alpha)}{\Gamma(\alpha)}+\sum_{i=1}^{n} \log X_{i}, \\
& 0=\frac{\partial l}{\partial \beta}=\frac{n \alpha}{\beta}-\sum_{i=1}^{n} X_{i} .
\end{aligned}
$$

The second equation implies that the MLEs $\hat{\alpha}$ and $\hat{\beta}$ satisfy $\hat{\beta}=\hat{\alpha} / \bar{X}$. Substituting into the first equation and dividing by $n, \hat{\alpha}$ satisfies

$$
\begin{equation*}
0=\log \hat{\alpha}-\frac{\Gamma^{\prime}(\hat{\alpha})}{\Gamma(\hat{\alpha})}-\log \bar{X}+\frac{1}{n} \sum_{i=1}^{n} \log X_{i} . \tag{13.1}
\end{equation*}
$$

The function $f(\alpha)=\log \alpha-\frac{\Gamma^{\prime}(\alpha)}{\Gamma(\alpha)}$ decreases from $\infty$ to 0 as $\alpha$ increases from 0 to $\infty$, and the value $-\log \bar{X}+\frac{1}{n} \sum_{i=1}^{n} \log X_{i}$ is always negative (by Jensen's inequality) -hence (13.1) always has a single unique root $\hat{\alpha}$, which is the MLE for $\alpha$. The MLE for $\beta$ is then $\hat{\beta}=\hat{\alpha} / \bar{X}$.

Unfortunately there is no closed-form expression for this root $\hat{\alpha}$. (In particular, the MLE $\hat{\alpha}$ is not the method-of-moments estimator for $\alpha$.) We may compute the root numerically using the Newton-Raphson method: We start with an initial guess $\alpha^{(0)}$, which (for example) may be the method-of-moments estimator

$$
\alpha^{(0)}=\frac{\bar{X}^{2}}{\frac{1}{n} \sum_{i=1}^{n}\left(X_{i}-\bar{X}\right)^{2}} .
$$

Having computed $\alpha^{(t)}$ for any $t=0,1,2, \ldots$, we compute the next iteration $\alpha^{(t+1)}$ by approximating the equation (13.1) with a linear equation using a first-order Taylor expansion around $\hat{\alpha}=\alpha^{(t)}$, and set $\alpha^{(t+1)}$ as the value of $\hat{\alpha}$ that solves this linear equation. In detail, let $f(\alpha)=\log \alpha-\frac{\Gamma^{\prime}(\alpha)}{\Gamma(\alpha)}$. A first-order Taylor expansion around $\hat{\alpha}=\alpha^{(t)}$ in (13.1) yields the linear approximation

$$
0 \approx f\left(\alpha^{(t)}\right)+\left(\hat{\alpha}-\alpha^{(t)}\right) f^{\prime}\left(\alpha^{(t)}\right)-\log \bar{X}+\frac{1}{n} \sum_{i=1}^{n} \log X_{i}
$$

and we set $\alpha^{(t+1)}$ to be the value of $\hat{\alpha}$ solving this linear equation, i.e. ${ }^{1}$

$$
\alpha^{(t+1)}=\alpha^{(t)}+\frac{-f\left(\alpha^{(t)}\right)+\log \bar{X}-\frac{1}{n} \sum_{i=1}^{n} \log X_{i}}{f^{\prime}\left(\alpha^{(t)}\right)}
$$

The iterations $\alpha^{(0)}, \alpha^{(1)}, \alpha^{(2)}, \ldots$ converge to the MLE $\hat{\alpha}$.
Example 13.4. Let $\left(X_{1}, \ldots, X_{k}\right) \sim \operatorname{Multinomial}\left(n,\left(p_{1}, \ldots, p_{k}\right)\right)$. (This is not quite the setting of $n$ IID observations from a parametric model, as we have been considering, although you can think of $\left(X_{1}, \ldots, X_{k}\right)$ as a summary of $n$ such observations $Y_{1}, \ldots, Y_{n}$ from the parametric model Multinomial $\left(1,\left(p_{1}, \ldots, p_{k}\right)\right)$, where $Y_{i}$ indicates which of $k$ possible outcomes occurred for the $i$ th observation.) The log-likelihood is given by

$$
l\left(p_{1}, \ldots, p_{k}\right)=\log \left(\binom{n}{X_{1}, \ldots, X_{k}} p_{1}^{X_{1}} \ldots p_{k}^{X_{k}}\right)=\log \binom{n}{X_{1}, \ldots, X_{k}}+\sum_{i=1}^{k} X_{i} \log p_{i}
$$

and the parameter space is

$$
\Omega=\left\{\left(p_{1}, \ldots, p_{k}\right): 0 \leq p_{i} \leq 1 \text { for all } i \text { and } p_{1}+\ldots+p_{k}=1\right\} .
$$

To maximize $l\left(p_{1}, \ldots, p_{k}\right)$ subject to the linear constraint $p_{1}+\ldots+p_{k}=1$, we may use the method of Lagrange multipliers: Consider the Lagrangian

$$
L\left(p_{1}, \ldots, p_{k}, \lambda\right)=\log \binom{n}{X_{1}, \ldots, X_{k}}+\sum_{i=1}^{k} X_{i} \log p_{i}+\lambda\left(p_{1}+\ldots+p_{k}-1\right)
$$

for a constant $\lambda$ to be chosen later. Clearly, subject to $p_{1}+\ldots+p_{k}=1$, maximizing $l\left(p_{1}, \ldots, p_{k}\right)$ is the same as maximizing $L\left(p_{1}, \ldots, p_{k}, \lambda\right)$. Ignoring momentarily the constraint $p_{1}+\ldots+p_{k}=1$, the unconstrained maximizer of $L$ is obtained by setting for each $i=1, \ldots, k$

$$
0=\frac{\partial L}{\partial p_{i}}=\frac{X_{i}}{p_{i}}+\lambda
$$

which yields $\hat{p}_{i}=-X_{i} / \lambda$. For the specific choice of constant $\lambda=-n$, we obtain $\hat{p}_{i}=X_{i} / n$ and $\sum_{i=1}^{n} \hat{p}_{i}=\sum_{i=1}^{n} X_{i} / n=1$, so the constraint is satisfied. As $\hat{p}_{i}=X_{i} / n$ is the unconstrained maximizer of $L\left(p_{1}, \ldots, p_{k},-n\right)$, this implies that it must also be the constrained maximizer of $L\left(p_{1}, \ldots, p_{k},-n\right)$, so it is the constrained maximizer of $l\left(p_{1}, \ldots, p_{k}\right)$. So the MLE is given by $\hat{p}_{i}=X_{i} / n$ for $i=1, \ldots, k$.

[^0]
[^0]:    ${ }^{1}$ If this update yields $\alpha^{(t+1)} \leq 0$, we may reset $\alpha^{(t+1)}$ to be a very small positive value.

