

## Lecture 5: Review on Statistical Inference

*Instructor: Yen-Chi Chen*

These notes are partially based on those of Mathias Drton.

## 5.1 Introduction

In this lecture, we will review the basic concept of statistical inference. In particular, we will focus on the so-called parametric model. We will assume that the data is from certain parametric family (e.g., a Gaussian) and the goal is to make inference about the underlying parameter (e.g., the mean and variance of a Gaussian) using the data.

Statistical inference is about drawing conclusions from the data. This process often involves with estimating some parameters of interest. In a parametric model, the parameter of interest are often the parameters of the corresponding parametric family. For a parametric model, there are three common tasks in statistical inference – estimating the underlying parameter, providing an interval inference about the underlying parameter, and testing if the underlying parameter satisfy certain conditions.

In statistics, there are two major paradigm for making inference – the frequentist paradigm and the Bayesian paradigm<sup>1</sup>. We will talk about their principles of estimation, interval inference, and testing.

Note that there is no right or wrong about each paradigm – they are just different ways of making arguments. Each paradigm is a self-consistent way to making logical arguments and has its own advantages and limitations.

## 5.2 Frequentist Approach

The Frequentist approach is the paradigm we learn from Statistics 101. It interprets probability as the long term frequency. In Frequentist approach, the parameter of interest is a fixed and unknown number.

### 5.2.1 Estimation

In a parametric model, we often estimate the parameter of interest using the so-called **maximum likelihood estimator (MLE)**. The idea is very simple. Suppose we observe only one observation  $X$  from a PDF/PMF  $p(x)$ . The parametric model assumes that such a PDF/PMF can be written as  $p(x) = p(x; \theta)$ , where  $\theta$  is the parameter of the model ( $\theta$  is often the parameter of interest) inside a parameter space  $\Theta$  ( $\theta \in \Theta$ ). The idea of MLE is to ask the following question: given the observation  $X$ , which  $\theta$  is the *most likely* parameter that generates  $X$ ? To answer this question, we can vary  $\theta$  and examine the value of  $p(X; \theta)$ .

Because we are treating  $X$  as fixed and  $\theta$  being something that we want to optimize, we can view the problem as finding the best  $\theta$  such that the **likelihood function**  $L(\theta|X) = p(X; \theta)$  is maximized. The MLE uses

<sup>1</sup>There are also other paradigms such as the fiducial inference ([https://en.wikipedia.org/wiki/Fiducial\\_inference](https://en.wikipedia.org/wiki/Fiducial_inference)) but the Frequentists and Bayesian are the two major paradigm.

the  $\theta$  that maximizes the likelihood value. Namely,

$$\hat{\theta}_{MLE} = \operatorname{argmax}_{\theta} L(\theta|X).$$

When we have multiple observations  $X_1, \dots, X_n$ , the likelihood function can be defined in a similar way – we use the joint PDF/PMF to define the likelihood function. Let  $p(x_1, \dots, x_n; \theta)$  be the joint PDF/PMF. Then the likelihood function is

$$L_n(\theta) = L(\theta|X_1, \dots, X_n) = p(X_1, \dots, X_n; \theta).$$

Note that when we assume IID observations,

$$L_n(\theta) = \prod_{i=1}^n L(\theta|X_i) = \prod_{i=1}^n p(X_i; \theta).$$

In many cases, instead of using the likelihood function, we often work with the **log-likelihood function**

$$\ell_n(\theta) = \log L_n(\theta).$$

Because taking the logarithmic does not change the maximizer of a function, the maximizer of the log-likelihood function is the same as the maximizer of the likelihood function. There are both computational and mathematical advantages of using a log-likelihood function over likelihood function. To see this, we consider the case of IID sample. Computationally, the likelihood function often has a very small value due to the product form of PDF/PMFs. So it is very likely that the number is too small, making the computation very challenging. Mathematically, when we take log of the likelihood function, the product of PDF/PMFs becomes an additive form

$$\ell_n(\theta) = \log L_n(\theta) = \sum_{i=1}^n \log p(X_i; \theta).$$

Under IID assumption, each  $\log p(X_i; \theta)$  is an IID random variable so the central limit theorem and the law of large number can be applied to the average, making it possible to analyze its asymptotic behavior.

Since under the IID assumptions, we have many advantages, we will assume IID from now on. Because MLE finds the maximum of  $\ell_n(\theta)$ , a common trick to find MLE is to study the gradient of the log-likelihood function, which is also known as the **score function**:

$$S_n(\theta) = \frac{\partial}{\partial \theta} \ell_n(\theta) = \sum_{i=1}^n s(\theta|X_i),$$

where  $s(\theta|X_i) = \frac{\partial}{\partial \theta} \ell(\theta|X_i) = \frac{\partial}{\partial \theta} \log p(X_i; \theta)$ . Under suitable conditions, the MLE satisfies the *score equation*:

$$S_n(\hat{\theta}_{MLE}) = 0.$$

Note that if there are more than one parameter, say  $\theta \in \mathbb{R}^p$ , the score equation will be a system of  $p$  equations.

Because the MLE is at the maximal point of the likelihood function, the curvature of the likelihood function around the maximal will determine its stability. To measure the curvature, we use the **Fisher's information matrix**:

$$I_1(\theta_0) = \mathbb{E}(S_1(\theta_0)S_1^T(\theta_0)).$$

When the data is indeed generated from a PDF/PMF  $p(x; \theta_0)$  and some other regularity conditions are met, the Fisher's information matrix measures the curvature of the model

$$\mathbb{E} \left[ \frac{\partial^2}{\partial \theta \partial \theta^T} \ell_n(\theta) \right] = -nI_1(\theta_0).$$

So sometimes we can just utilize the second-order derivative to compute Fisher's information matrix. Moreover,

$$\sqrt{n} \left( \hat{\theta}_{MLE} - \theta_0 \right) \xrightarrow{D} N(0, I_1^{-1}(\theta_0)).$$

Namely, the MLE is asymptotically normally distributed around the true parameter  $\theta_0$  and the covariance is determined by the Fisher's information matrix. Note that the asymptotic normality also implies that  $\hat{\theta}_{MLE} - \theta_0 \xrightarrow{P} 0$ .

**Example 5.1 (Binomial distribution)** Assume that we obtain a single observation  $Y \sim \text{Bin}(n, p)$ , and we assume that  $n$  is known. The goal is to estimate  $p$ . The log-likelihood function is

$$\ell(p) = Y \log p + (n - Y) \log(1 - p) + C_n(Y),$$

where  $C_n(Y) = \log \binom{n}{Y}$  is independent of  $p$ . The score function is

$$S(p) = \frac{Y}{p} - \frac{n - Y}{1 - p}$$

so solving the score equation gives us  $\hat{p}_{MLE} = \frac{Y}{n}$ . Moreover, the Fisher's information is

$$I(p) = \mathbb{E} \left\{ \frac{\partial}{\partial p} S(p) \right\} = -\frac{\mathbb{E}(Y)}{p^2} - \frac{n - \mathbb{E}(Y)}{(1 - p)^2} = \frac{n}{p(1 - p)}.$$

**Example 5.2 (Multinomial distribution)** Let  $X_1, \dots, X_n$  be IID from a multinomial distribution such that  $P(X_1 = j) = p_j$  for  $j = 1, \dots, s$  and  $\sum_{j=1}^s p_j = 1$ . Note that the parameter space is  $\Theta = \{(p_1, \dots, p_s) : 0 \leq p_j, \sum_{j=1}^s p_j = 1\}$ . By setting  $N_j = \sum_{i=1}^n I(X_i = j)$  for each  $j = 1, \dots, s$ , we obtain the random vector  $(N_1, \dots, N_s) \sim \text{Multinomial}(n, p)$ , where  $p = (p_1, \dots, p_s)$ . The parameters of interest are  $p_1, \dots, p_s$ . In this case, the likelihood function is

$$L_n(p_1, \dots, p_s) = \frac{n!}{N_1! \dots N_s!} p_1^{N_1} \dots p_s^{N_s}$$

and the log-likelihood function is

$$\ell_n(p_1, \dots, p_s) = \sum_{j=1}^s N_j \log p_j + C_n,$$

where  $C_n$  is independent of  $p$ . Note that naively computing the score function and set it to be 0 will not grant us a solution (think about why) because we do not use the constraint of the parameter space – the parameters are summed to 1. To use this constraint in our analysis, we consider adding the Lagrange multipliers and optimize it:

$$F(p, \lambda) = \sum_{j=1}^s N_j \log p_j + \lambda \left( 1 - \sum_{j=1}^s p_j \right).$$

Differentiating this function with respect to  $p_1, \dots, p_s$ , and  $\lambda$  and set it to be 0 gives

$$\frac{\partial F}{\partial p_j} = \frac{N_j}{p_j} - \lambda = 0 \Rightarrow N_j = \lambda \hat{p}_{MLE, j}$$

and  $1 - \sum_{j=1}^s p_j = 0$ . Thus,  $n = \sum_{j=1}^s N_j = \lambda \sum_{j=1}^s p_j = \lambda$  so  $\hat{p}_{MLE, j} = \frac{N_j}{n}$ .

## 5.2.2 Confidence Intervals

In some analysis, we not only want to have just a point estimate of the parameter of interest, but also want to use an interval to infer the parameter of interest. And we also want to assign a level to this interval to describe how ‘accurate’ this interval is. Note that here the concept of accuracy is not well-defined – we will talk about it later. Ideally, given an accuracy level, we want an interval as small as possible.

The Frequentist and the Bayesian defines the accuracy differently so their construction of intervals are also different. In short, the Frequentists defines the accuracy as the *long term frequency coverage* of the underlying true parameter of interest whereas the Bayesian defines the accuracy in terms of covering the posterior probability. In this section, we will talk about the Frequentist approach and the interval is known as the **confidence interval**. The accuracy that Frequentists are using is called the *confidence level*.

Formally, given a confidence level  $1 - \alpha$ , a confidence interval of  $\theta_0$  is a random interval  $C_{n,\alpha}$  that can be constructed solely from the data (i.e., can be constructed using  $X_1, \dots, X_n$ ) such that

$$P(\theta \in C_{n,\alpha}) \geq 1 - \alpha + o(1).$$

Beware, what is random is not  $\theta$  but the interval  $C_{n,\alpha}$ . The quantity  $P(\theta \in C_{n,\alpha})$  is also called the (Frequentist) coverage. Note that we allow the coverage to be *asymptotically*  $1 - \alpha$ ; when there is no  $o(1)$  term, we will say that the confidence interval has a finite sample coverage. A confidence interval with the above property is also called a (asymptotically) valid confidence interval.

**Normal confidence interval.** A traditional approach to constructing a confidence interval of  $\theta_0$  is based on the asymptotic normality of the MLE:

$$\sqrt{n}(\hat{\theta}_{MLE} - \theta_0) \xrightarrow{D} N(0, I_1^{-1}(\theta_0)).$$

When the dimension of the parameter is 1, a simple confidence interval is

$$\hat{\theta}_{MLE} \pm z_{1-\alpha/2} \cdot \sigma_{\theta_0},$$

where  $\sigma_{\theta_0}^2 = \frac{1}{n} I_1^{-1}(\theta_0)$ . Such interval is not a confidence interval because  $\theta_0$  is unknown. We can modify it using a plug-in estimate of the Fisher’s information:

$$C_{n,\alpha} = [\hat{\theta}_{MLE} - z_{1-\alpha/2} \cdot \sigma_{\hat{\theta}_{MLE}}, \hat{\theta}_{MLE} + z_{1-\alpha/2} \cdot \sigma_{\hat{\theta}_{MLE}}],$$

where  $z_\beta$  the  $\beta$ -percentile of the standard normal distribution. Using the Slutsky’s theorem, you can easily show that this confidence interval has the asymptotic coverage.

When the dimension of the parameter is greater than 1, there are multiple ways we can construct a confidence interval. Note that in this case, the set  $C_{n,\alpha}$  is no longer an interval but a region/set so it is often called a confidence region/set. A simple approach of constructing a confidence set is via an ellipse. Note that the asymptotic normality also implies that (using continuous mapping theorem)

$$n(\hat{\theta}_{MLE} - \theta_0)^T I_1(\hat{\theta}_{MLE})(\hat{\theta}_{MLE} - \theta_0) \xrightarrow{D} \chi_p^2,$$

where  $\chi_p^2$  denotes the  $\chi^2$  distribution with a degree of freedom  $p$ . So we construct the confidence set using

$$C_{n,\alpha} = \left\{ \theta : n(\hat{\theta}_{MLE} - \theta)^T I_1(\hat{\theta}_{MLE})(\hat{\theta}_{MLE} - \theta) \leq \chi_{p,1-\alpha}^2 \right\},$$

where  $\chi_{p,\beta}^2$  is the  $\beta$ -percentile of the  $\chi^2$  distribution with a degree of freedom  $p$ .

**Bootstrap confidence interval.** Bootstrap approach is an Monte Carlo method for assessing the uncertainty of an estimator. It can be used to compute the variance of an estimator (not necessarily the MLE) and

construct a confidence interval. In the case of likelihood inference, the bootstrap approach has an advantage that we do not need to know the closed-form of  $I_1(\theta)$  to construct the confidence interval or to approximate the variance of the MLE.

While there are many variant of bootstrap methods, we introduce the simplest one – the empirical bootstrap. For simplicity, we assume that the dimension of  $\theta$  is 1 (the bootstrap works for higher dimensions as well). Let  $X_1, \dots, X_n$  be the original sample. We then *sample with replacement* from the original sample to obtain a new sample of each size  $X_1^*, \dots, X_n^*$ . This new sample is called a bootstrap sample. We find the MLE using the bootstrap sample and let  $\hat{\theta}_{MLE}^*$  denote the bootstrap MLE. Now we repeat the bootstrap process  $B$  times, leading to  $B$  bootstrap MLEs

$$\hat{\theta}_{MLE}^{*(1)}, \dots, \hat{\theta}_{MLE}^{*(B)}.$$

Let  $t_\beta$  denotes the  $\beta$ -percentile of these  $B$  values, i.e.,

$$\hat{t}_\beta = \hat{G}^{-1}(\beta), \quad \hat{G}(t) = \frac{1}{B} \sum_{b=1}^B I(\hat{\theta}_{MLE}^{*(b)} \leq t).$$

Then the bootstrap confidence interval of  $\theta_0$  is

$$C_{n,\alpha} = [\hat{t}_{\alpha/2}, \hat{t}_{1-\alpha/2}].$$

One can prove that under very mild conditions, the bootstrap confidence interval has asymptotic coverage.

The power of the bootstrap method is that *we do not use anything about the Fisher's information!* As long as we can compute the estimator, we can construct an asymptotically valid confidence interval. Note that if we do know the Fisher's information, the bootstrap method can be modified using the bootstrap  $t$ -distribution method, which provides a better asymptotic coverage (namely, the  $o(1)$  decays faster to 0 than the above method and the normal confidence interval)<sup>2</sup>.

### 5.2.3 Test of Significance

Statistical test is about how to design a procedure that allows us to make scientific discovery. Such a procedure has to be able to handle the uncertain nature of our data. In statistics, we model the data as random variables so the testing procedure needs to account for the randomness.

Let  $\mathcal{D}_n = \{X_1, \dots, X_n\}$  denotes our data. The testing procedure involves two competing hypotheses:

- *Null hypothesis  $H_0$* : the hypothesis that we want to challenge. It is often related to the current scientific knowledge.
- *Alternative hypothesis  $H_a$* : the hypothesis that complements to the null hypothesis. It is the hypothesis we would like to prove to be plausible using our data.

The goal is to see if we have strong enough evidence (from  $\mathcal{D}_n$ ) that we can argue the alternative hypothesis is more reasonable than the null hypothesis. If we do have enough evidence, then we will reject the null hypothesis. When the null hypothesis reflects the scenarios that can be explained by the current scientific knowledge, rejecting the null hypothesis means that we have discovered something new.

To design a testing procedure, we need to quantify the notion of evidence. The Frequentists and the Bayesian use different ways to measure the evidence. The Frequentist approach is the *p-value* whereas the Bayesian approach is the *Bayes factor*. We will talk about Bayes factor later so here we focus on the *p-value*.

<sup>2</sup>see, e.g., Chapter 2 of *All of nonparametric statistics* by Larry Wasserman and Chapter 3.5 of *The bootstrap and Edgeworth expansion* by Peter Hall.

Here is a summary on Frequentist approach of hypothesis test.

1. Based on the model and null hypothesis, design a test statistic.
2. Compute the distribution of the test statistics under the null hypothesis.
3. Plug-in the data into the test statistic, compute the probability of observing a more extreme data against the null hypothesis. Such a probability is the p-value.
4. Compare p-value to the significance level. If p-value is less than the significance level, we reject the null hypothesis.

The central idea of hypothesis test is to control the *type-1 error*, the probability of falsely rejecting  $H_0$  when  $H_0$  is correct. Essentially, the p-value can be interpreted as *if we reject the null hypothesis (under this p-value), then our type-1 error is the same as the p-value*. The significance level reflects the amount of type-1 error we can tolerate so when p-value is less than the significance level, we can reject  $H_0$ . Due to the construction of p-value, a small p-value means that the null hypothesis does not fit to the data very well (so we are seeing an extreme event if  $H_0$  is true). Thus, small p-value or rejecting  $H_0$  under a small significance level means that we have more evidence against  $H_0$ .

Note that there is another quantity called *type-2 error*, the probability of not rejecting  $H_0$  when  $H_0$  is false. Namely, type-2 error is concerned with the case that we fail to reject  $H_0$  when we should.

In statistics, we often control type-1 error first and the hope that the type-2 error is also small. When do we put more emphasis on type-1 error? This has something to do with the philosophy of scientific research. The scientific approach is a systematic way to acquire reliable knowledge. Thus, every discovery we made should be accompanied with sufficient evidences. In Frequentist approach, the measure of evidence against  $H_0$  is the p-value – the smaller p-value, the more evidence. Thus, controlling type-1 error means that we put requirements on the amount of evidence we need to claim a scientific discovery.

While there are many possible ways to construct a test statistic, here we consider two common approaches: Wald test and the likelihood ratio test.

**Wald test.** Assume in a simple case where we model the data as IID from a parametric model  $p(x; \theta)$  with a 1D parameter. Suppose that we are comparing two hypotheses

$$H_0 : \theta = \theta_1, \quad H_a : \theta \neq \theta_1.$$

Since the MLE is a good estimate of  $\theta$ , we can design our test statistic using the MLE. Because we know that the MLE has asymptotic normality,

$$\sqrt{n}I_1(\theta)^{-1/2}(\hat{\theta}_n - \theta) \xrightarrow{D} N(0, 1),$$

we can then use

$$T_n(\theta) = \sqrt{n}I_1(\theta)^{-1/2}(\hat{\theta}_n - \theta)$$

as our test statistic. If  $H_0$  is true,  $T_n(\theta_1)$  should behaves like a standard normal distribution. When we observe the actual data, we can then compare the observed value of  $T_n(\theta_1)$  against 0 (since  $T_n(\theta_1) = 0$  means a perfect match with  $H_0$ ). We reject  $H_0$  if  $T_n(\theta_1)$  is either too large or too small. The p-value is then

$$p(\mathcal{D}_n) = 1 - 2 \cdot \Phi^{-1}(|T_n(\theta_1)|),$$

where  $\Phi(t) = P(Z \leq t)$  is the CDF of a standard normal distribution. Note that the above method is also known as the **Wald test**.

**Likelihood ratio test (LRT).** LRT is another popular way to conduct hypothesis test under a parametric model. It can be easily applied to multivariate parameters so now we assume that there are  $p$  parameters, i.e.,  $\theta \in \Theta \subset \mathbb{R}^p$ . The null hypothesis and the alternative hypothesis are

$$H_0 : \theta \in \Theta_0, \quad H_a : \theta \in \Theta \setminus \Theta_0.$$

Let  $\ell_n(\theta)$  be the log-likelihood function and let

$$\hat{\theta}_{MLE} = \operatorname{argmax}_{\theta \in \Theta} \ell_n(\theta), \quad \hat{\theta}_{MLE,0} = \operatorname{argmax}_{\theta \in \Theta_0} \ell_n(\theta)$$

be the global MLE and the constrained MLE (in the null hypothesis space). The LRT starts with a test statistic

$$\text{LRT} = \frac{\sup_{\theta \in \Theta_0} L_n(\theta)}{\sup_{\theta \in \Theta} L_n(\theta)} = \frac{L_n(\hat{\theta}_{MLE,0})}{L_n(\hat{\theta}_{MLE})}$$

and reject  $H_0$  if LRT is too small. To get the p-value of the LRT, we often use the following fact:

$$-2 \log \text{LRT} = 2 \left( \ell_n(\hat{\theta}_{MLE}) - \ell_n(\hat{\theta}_{MLE,0}) \right) \xrightarrow{D} \chi_r^2,$$

where  $r = \dim(\Theta) - \dim(\Theta_0)$  is the difference of the dimensions/degree-of-freedom between the full parameter space and constrained parameter space. The above result is also known as the likelihood ratio approximation or Wilk's theorem. The LRT proposes to use the  $T_n = 2 \left( \ell_n(\hat{\theta}_{MLE}) - \ell_n(\hat{\theta}_{MLE,0}) \right)$  as our test statistic and compared it with the CDF of  $\chi_r^2$  to obtain the p-value.

**Remark 5.3** *Here are some remarks about LRT:*

- Why  $r$  degrees of freedom? *Some of you may be wondering why we are getting a  $\chi^2$  distribution with a degree of freedom  $r$ . Here is a simple explanation using geometry. Recall that the MLE behaves like a normal distribution around the true parameter  $\theta_0$ . If  $H_0$  is true,  $\theta_0 \in \Theta_0$  so  $\hat{\theta}_{MLE}$  will be asymptotically normally distributed around  $\theta_0$ . Since  $\hat{\theta}_{MLE}$  is the maximizer over  $\Theta$  so it has  $p$  degrees of freedom (it can move in each of the  $p$  dimensions). The constrained MLE  $\hat{\theta}_{MLE,0}$  is the maximizer under  $\Theta_0$ , which has  $r$  constraints. Thus,  $\Theta_0$  has  $p - r$  degrees of freedom, which implies that its maximizer  $\hat{\theta}_{MLE,0}$  uses  $p - r$  degree of freedom. The remaining degrees of freedom is  $p - (p - r) = r$  so this is why we are obtaining a  $\chi^2$  distribution with  $r$  degrees of freedom. Note that one can replace the  $r$  constraints to be saying that  $\Theta_0$  is a  $p - r$  dimensional manifold.*
- Equivalence between likelihood ratio test and Wald test. *Asymptotically, one can show that the likelihood ratio test and the Wald test are the same under classical assumptions on the MLE. Also, there is another test that is closely related to them called the score test, which is based on the value of score function as a test statistic. Again, asymptotically the score test and other two tests are equivalent. However, when the likelihood function is more complex, such as having multiple local maxima. These three tests may not be the same (they can be quite different)<sup>3</sup>.*
- Relation to confidence interval. *Hypothesis test can be used to construct a confidence interval. Consider testing the null hypothesis:  $H_0 : \theta = \theta_1$  for a specific value  $\theta_1$  under a significance level  $\alpha$ . For each  $\theta_1 \in \Theta$ , we can do a hypothesis test. Some parameters will be rejected whereas the others will not. Let  $\hat{A}_{n,\alpha}$  be the collection of parameters that the null hypothesis will not be rejected. Then you can show that  $\hat{A}_{n,\alpha}$  is a confidence interval of  $\theta$ . This approach is also known as the confidence interval from inverting a hypothesis test. Note that if we are given a  $1 - \alpha$  confidence interval of the parameter of interest  $\theta$ , we can use it to test the null hypothesis  $H_0 : \theta \in \Theta_0$ . If the confidence interval intersects*

<sup>3</sup>see <https://arxiv.org/abs/1807.04431>

with  $\Theta_0$ , we cannot reject  $H_0$ . If they are disjoint, then we can reject  $H_0$  under a significance level  $\alpha$ . Thus, hypothesis test problem and confidence intervals are highly related to each other. However, testing a given hypothesis is often an easier problem than constructing a confidence interval because confidence interval requires a procedure that is valid regardless of the null hypothesis being correct or not. On the other hand, hypothesis test problem only requires a procedure that works when  $H_0$  is true.

There are many misconceptions about p-values. To clarify what p-value stands for and what it does NOT stand for, I obtain the following 6 principles from American Statistical Association's website<sup>4</sup>:

1. P-values can indicate how incompatible the data are with a specified statistical model.
2. P-values do not measure the probability that the studied hypothesis is true, or the probability that the data were produced by random chance alone.
3. Scientific conclusions and business or policy decisions should not be based only on whether a p-value passes a specific threshold.
4. Proper inference requires full reporting and transparency.
5. A p-value, or statistical significance, does not measure the size of an effect or the importance of a result.
6. By itself, a p-value does not provide a good measure of evidence regarding a model or hypothesis.

### 5.2.4 Model Mis-specification (♠♠♠)

Many theory about the MLE assumes that the population distribution function belongs to our parametric family. However, this is a very strong assumption in reality. It is very likely that the population distribution function does not belong to our parametric family (e.g., the population PDF is not Gaussian but we fit a Gaussian to it). What will happen in this case for our MLE? Will it still converge to something? If so, what will be the quantity that it is converging?

Model mis-specification studies the situation like this – we assume a wrong model for the population distribution function. Let  $p_0(x)$  be the population PDF and we assume that the population PDF can be written as  $p(x; \theta)$ . However,  $p_0 \neq p(x; \theta)$  for every  $\theta \in \Theta$ . It turns out that the MLE  $\hat{\theta}_{MLE}$  still converges under mild assumptions to a quantity  $\theta^*$  in probability. Moreover, the corresponding PDF/PMF  $p(x; \theta^*)$  has an interesting relation with  $p_0(x)$ . Assume that the RV  $X$  has a PDF/PMF  $p_0$ . Then

$$\mathbb{E} \left\{ \log \left( \frac{p_0(X)}{p(X; \theta^*)} \right) \right\} = \inf_{\theta \in \Theta} \mathbb{E} \left\{ \log \left( \frac{p_0(X)}{p(X; \theta)} \right) \right\} = \inf_{\theta \in \Theta} \text{KL}(p_0, p_\theta),$$

where KL is also known as the *Kullback-Liebler (KL) divergence* and  $p_\theta(x) = p(x; \theta)$ . Namely, the MLE corresponds to the parametric distribution in the specified family that minimizes the KL divergence to the population distribution.

In model mis-specification case, the MLE still satisfies the score equation (under appropriate assumption) but the Fisher's information may not reflect the actual curvature of the likelihood function around  $\theta^*$ . The asymptotic covariance (related to the curvature) of  $\hat{\theta}_{MLE}$  will be  $\Sigma = I_1^{-1}(\theta^*) \mathbb{E}(S_1(\theta^*) S_1^T(\theta^*)) I_1^{-1}(\theta^*)$  and we still have

$$\sqrt{n}(\hat{\theta}_{MLE} - \theta^*) \xrightarrow{D} N(0, \Sigma).$$

<sup>4</sup><https://www.amstat.org/asa/files/pdfs/P-ValueStatement.pdf>

### 5.3 Bayesian Approach

The Bayesian inference is an alternative statistical paradigm to the Frequentist approach. The Bayesian approach interprets the probability in a broader sense that include subjective probability, which allows us to assign probability to almost every quantity in our model (including the parameter of interest and even a statistical model). The Bayesian inference relies on a simple decision theoretic rule – if we are competing two or more choices, we always choose the one with higher probability. This simple rule allows us to design an estimator, construct an interval, and perform hypothesis test.

In the Bayesian analysis, we assign a probability to every parameter in our model. For a parametric model  $p(x; \theta)$ , the parameter of interest  $\theta$  is given a **prior distribution**  $\pi(\theta)$  that reflects our belief about the value of  $\theta$ . In a sense, the prior distribution quantifies our subjective belief about the parameter  $\theta$ . The higher value of  $\pi(\theta)$  indicates that we believe that  $\theta$  is a more likely value of it.

Using the prior distribution, the Bayesian probability model can be written as follows:

$$\begin{aligned} X_1, \dots, X_n | \theta &\stackrel{IID}{\sim} p(x|\theta) \\ \theta &\sim \pi. \end{aligned}$$

The Bayesian inference focuses on the distribution of  $\theta$  after observing  $X_1, \dots, X_n$ :

$$\pi(\theta | X_1, \dots, X_n) = \frac{p(X_1, \dots, X_n, \theta)}{p(X_1, \dots, X_n)} \propto \underbrace{p(X_1, \dots, X_n | \theta)}_{\text{likelihood}} \times \underbrace{\pi(\theta)}_{\text{prior}}.$$

This distribution is also known as the **posterior distribution**.

The posterior distribution informs us about how our prior belief is updated after seeing the data. It is the central quantity in Bayesian inference – all our decisions will be related to it. In Bayesian's point of view, probability models are just mathematical tools for analyzing data. We do not assume that the data is generated from a probability distribution. We just *view* the data as generated from  $p(x; \theta)$ . Given that we do not assume the probability model to be the *true* model, there is NO true parameter so we cannot talk about conventional statistical errors. However, Bayesian does have another way to expressing the error in our inference – the posterior distribution. The posterior distribution reflects our belief about the parameter after seeing the data, we can use it as a measure of *uncertainty* about  $\theta$ . If the posterior distribution is more spread out, then the uncertainty in our inference is larger. On the other hand, if the posterior distribution is very concentrated, then there is very little (Bayesian) uncertainty.

**Remark 5.4 (Interpretation of a prior distribution)** *The prior distribution is often interpreted as as our subject belief about the likelihood of the parameter  $\theta$ , which can be viewed as an optimal choice under a stochastic decision process. To simplify the problem we assume that  $\Theta = \{0, 1, 2\}$ . Even without any data at hand, we can ask ourselves about our belief about each parameter value. Some people may think that 1 is the most likely one; some may think that 2 is the most likely one. To make our belief more precise, we use probability to work on it. Let  $\pi(j)$  be the number that reflects our belief about  $\theta = j$ . We interpret the numerical value of  $\pi(j)$  as follows. We are forced to guess the answer of  $\theta = j$  versus  $\theta \neq j$ . If the answer is  $\theta = j$  and we indeed guess it correctly, we will be rewarded  $\delta$  dollar. If the answer is  $\theta \neq j$  and we get it correct, we will be rewarded 1 dollar. If we get it wrong, we do not lose anything. Our principle is to maximize our expected reward. Now assume that the true value of  $\theta$  has equal probability of being  $j$  or not  $j$ . Then what should we choose?  $\theta = j$  or  $\theta \neq j$ ? Now we think about this problem by varying  $\delta$  from 0 to infinity. When  $\delta$  is small, unless we have very strong belief on  $\theta = j$ , we will not bid on it. When increasing  $\delta$ , at certain threshold we will switch our decision from bidding on  $\theta \neq j$  to  $\theta = j$ . Let this threshold be  $\eta_j$ .  $\eta_j$  is a number that reflects our belief about  $\theta = j$  and we associate it with our prior*

$$\pi(j) = \frac{1}{1 + \eta_j} \Leftrightarrow \eta_j = \frac{1 - \pi(j)}{\pi(j)} \quad (\text{odds of } \theta = j).$$

Here, you see that we only use one simple decision rule – bidding on the one with a higher expected outcome. This allows us to quantify our belief.

### 5.3.1 Bayesian Estimation

There are two common estimator in Bayesian inference: the posterior mean and the maximum a posteriori estimation (MAP).

**Posterior mean.** Just like we often use the sample mean as an estimator of the population mean, the mean of the posterior distribution is a common quantity that was used as an estimator of  $\theta$ :

$$\hat{\theta}_\pi = \mathbb{E}(\theta|X_1, \dots, X_n) = \int \theta \cdot \pi(\theta|X_1, \dots, X_n) d\theta.$$

It represents the average location of our belief about the parameter after seeing the data.

**Maximum a posteriori estimation (MAP).** Another common estimator of  $\theta$  is the MAP; it relies on the similar principle as the MLE – we choose the one that is the most likely. Here the ‘likely’ is interpreted as our posterior belief about the parameter of interest  $\theta$ . Formally, MAP is defined as

$$\hat{\theta}_{MAP} = \operatorname{argmax}_\theta \pi(\theta|X_1, \dots, X_n).$$

**Example 5.5 (Bayesian Binomial distribution)** Assume that we have an observation  $Y \sim \operatorname{Bin}(N, \theta)$  where  $N$  is known and the parameter of interest is  $\theta$ :

$$P(Y = y|\theta) = \binom{N}{y} \theta^y (1 - \theta)^{N-y}.$$

We use a Beta distribution with parameters  $(\alpha, \beta)$  as our prior distribution for  $\theta$ . Namely,

$$\pi(\theta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1} (1 - \theta)^{\beta-1},$$

where  $\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dt$  is the Gamma function and  $\alpha, \beta > 0$ . Note that  $(\alpha, \beta)$  are called the hyperparameters and are known quantities (because we know our belief about the data). For a Beta distribution with parameter  $\alpha, \beta$ , the mean is  $\frac{\alpha}{\alpha + \beta}$ .

The posterior distribution is

$$\begin{aligned} \pi(\theta|Y) &= \frac{\binom{N}{Y} \theta^Y (1 - \theta)^{N-Y} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1} (1 - \theta)^{\beta-1}}{\int \binom{N}{Y} \theta^Y (1 - \theta)^{N-Y} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1} (1 - \theta)^{\beta-1} d\theta} \\ &\propto \theta^{Y + \alpha - 1} (1 - \theta)^{N - Y + \beta - 1} \end{aligned}$$

so it is a Beta distribution with parameters  $(Y + \alpha, N - Y + \beta)$ . Then the posterior mean and MAP are

$$\hat{\theta}_\pi = \frac{Y + \alpha}{N + \alpha + \beta}, \quad \hat{\theta}_{MAP} = \frac{Y + \alpha - 1}{N + \alpha + \beta - 2}$$

(these are the mean and the mode of a Beta distribution).

Note that in this problem, the MLE is  $\hat{\theta}_{MLE} = \frac{Y}{N}$ . Thus, the posterior mean has an interesting decomposition:

$$\begin{aligned}\hat{\theta}_\pi &= \frac{Y + \alpha}{N + \alpha + \beta} \\ &= \hat{\theta}_\pi = \frac{Y}{N + \alpha + \beta} + \frac{\alpha}{N + \alpha + \beta} \\ &= \frac{Y}{N} \times \frac{N}{N + \alpha + \beta} + \frac{\alpha}{\alpha + \beta} \times \frac{\alpha + \beta}{N + \alpha + \beta} \\ &= \hat{\theta}_{MLE} \times W + \text{Prior mean} \times (1 - W),\end{aligned}$$

where  $W = \frac{N}{N + \alpha + \beta}$  is a weight that is tending to 1 when  $N \rightarrow \infty$ . This phenomenon – the posterior mean can be written as the weighted average of the MLE and the prior mean – occurs in several scenarios. Moreover, the fact that the weights  $W \rightarrow 1$  as the sample size  $N \rightarrow \infty$  means that when we have more and more data, the prior distribution seems to be irrelevant. Thus, the posterior mean would have a similar asymptotic property as the sample mean. However, this is not a general phenomenon; often only certain combination of prior and likelihood models will have this feature.

**Example 5.6 (Multinomial-Dirichlet model)** Suppose  $\theta \in \mathbb{R}^K$  satisfying  $\theta_j \geq 0$  and  $\sum_{j=1}^K \theta_j = 1$  (this is called a  $K$ -simplex). Given  $\theta$ , the random vector  $X \in \{0, 1, \dots, N\}^K$  is from a multinomial distribution  $\text{Mult}(N, \theta)$  if  $\sum_{j=1}^K X_j = N$  and its joint PMF is

$$p(x_1, \dots, x_K | \theta) = \frac{N!}{x_1! \dots x_K!} \prod_{j=1}^K \theta_j^{x_j}.$$

You can easily see that the Binomial distribution is a special case of the multinomial distribution when  $K = 2$ .

To place a prior on feasible values of  $\theta$  ( $K$ -simplex), we consider the Dirichlet distribution:  $\theta \sim \text{Dir}(\alpha)$ , where  $\alpha \in \mathbb{R}_{>0}^K$  and the PDF of  $\theta$  is

$$\pi(\theta) = \frac{1}{D(\alpha)} \prod_{j=1}^K \theta_j^{\alpha_j - 1},$$

when  $\theta_j \geq 0$  and  $\sum_{j=1}^K \theta_j = 1$  and the density is 0 outside this feasible range. The Dirichlet distribution can be viewed as a generalized Beta distribution—when  $K = 2$ , the Dirichlet distribution reduces to Beta distribution.

In this case, the posterior distribution of  $\theta$  given  $X$  is

$$\pi(\theta | x) \propto p(x | \theta) \pi(\theta) \propto \theta_j^{X_j + \alpha_j - 1},$$

so it is Dirichlet distribution with parameter  $\alpha + X$ . You can show that the posterior mean again can be decomposed into a convex combination between the MLE and the mean of the prior.

**Example 5.7 (Normal-normal model)** Suppose we have random variable  $X | \theta \sim N(\theta, \sigma^2)$  and we place a prior distribution on  $\theta$  as  $\theta \sim N(\mu, \tau^2)$ . We assume that  $\sigma^2, \mu, \tau^2$  are known. In this case, the posterior distribution of  $\theta$  given  $X$  is

$$\pi(\theta | x) \propto p(x | \theta) \pi(\theta) \propto e^{-\frac{1}{2\sigma^2}(x-\theta)^2} \cdot e^{-\frac{1}{2\tau^2}(\theta-\mu)^2},$$

which is a normal distribution

$$\theta | X \sim N\left(\frac{1/\tau^2}{1/\tau^2 + 1/\sigma^2} \mu + \frac{1/\sigma^2}{1/\tau^2 + 1/\sigma^2} X, \frac{1}{1/\tau^2 + 1/\sigma^2}\right). \quad (5.1)$$

Equation (5.1) clearly shows how the posterior mean  $\mathbb{E}(\theta|X) = \frac{1/\tau^2}{1/\tau^2+1/\sigma^2}\mu + \frac{1/\sigma^2}{1/\tau^2+1/\sigma^2}X$  can be written as a convex combination between the sample mean  $X$  and the prior's mean  $\mu$ .

**Remark 5.8** Some remarks about the Bayesian inference.

- Choice of prior and conjugate prior. The choice of prior reflects our belief about the parameter before seeing any data. Sometimes people want to choose a prior distribution such that the posterior distribution is in the same family as the prior distribution, just like what we have observed in the above example. If a prior distribution and a likelihood function leads to a posterior that belongs to the same family as the prior, we call this prior **conjugate prior**. In all the previous three examples, we are using a conjugate prior for the problem. There are several conjugate priors know to date, see [https://en.wikipedia.org/wiki/Conjugate\\_prior](https://en.wikipedia.org/wiki/Conjugate_prior) for an incomplete list of cases.

Another common choice of prior is called the **Jeffreys prior**<sup>5</sup>, which chooses a prior  $\pi(\theta) \propto \sqrt{\det(I_1(\theta))}$ , where  $I_1(\theta)$  is the Fisher's information matrix. One can view the Jeffreys prior as the prior that we do not have any prior belief about  $\theta$ ; or more formally, an uninformative prior.

- Challenge of computing the posterior. In general, if we do not choose a conjugate prior, the posterior distribution could be difficult to compute. The challenge often comes from the normalization quantity  $p(X_1, \dots, X_n)$  in the denominator of the posterior  $\pi(\theta|X_1, \dots, X_n)$  (the numerator is just the prior times the likelihood). In practice we will use Monte Carlo method to compute the posterior – we generate points from  $\pi(\theta|X_1, \dots, X_n)$  and as we generate enough points, these points should approximate the true posterior distribution well. We will talk more about this later in the lecture of MCMC (Monte Carlo Markov Chain).
- Consistency. In pure Bayesian's point of view, statistical consistency is not an important property because probability model is a working model to describe the data and we do not need to assume that there exists an actual parameter that generates the data. Thus, the posterior distribution is the quantity that we really need to make our inference. However, sometimes Bayesian estimators, such as the posterior mean or MAP, does have statistical consistency. Namely,  $\hat{\theta}_\pi \xrightarrow{P} \theta_0$  and  $\hat{\theta}_{MAP} \xrightarrow{P} \theta_0$ , where the data  $X_1, \dots, X_n \stackrel{IID}{\sim} p(x; \theta_0)$ . This is often related to the Bernstein-von Mises theorem<sup>6</sup>. Although statistical consistency was not an important property in Bayesian paradigm (because Bayesian does not assume the data is indeed from a probability model; probability models are just a mathematical model to help us analyze the data), still many researchers would prove consistency when proposing a Bayesian approach.

### 5.3.2 Bayesian Interval: Credible Interval

The Bayesian's interval is very straight forward – since the posterior contains our belief about the parameter after seeing the data, we just construct the interval using the posterior distribution. Given a credible level  $1 - \alpha$ , a credible interval  $D_{n,\alpha}$  satisfies

$$1 - \alpha = \int_{D_{n,\alpha}} \pi(\theta|X_1, \dots, X_n) d\theta.$$

Often we will choose the credible interval such that it is shortest. It turns out that (actually it is not hard to prove) such a credible interval is related to the upper level set of the posterior. Given a level  $\lambda$ , we define

$$L_\lambda = \{\theta \in \Theta : \pi(\theta|X_1, \dots, X_n) \geq \lambda\}.$$

<sup>5</sup>see [https://en.wikipedia.org/wiki/Jeffreys\\_prior](https://en.wikipedia.org/wiki/Jeffreys_prior) for more details.

<sup>6</sup>[https://en.wikipedia.org/wiki/Bernstein%E2%80%93von\\_Mises\\_theorem](https://en.wikipedia.org/wiki/Bernstein%E2%80%93von_Mises_theorem)

Define  $V(\lambda) = \int_{L_\lambda} \pi(\theta|X_1, \dots, X_n) d\theta$ . When  $\lambda$  is very large,  $V(\lambda)$  is very small. When we decrease  $\lambda$ ,  $V(\lambda)$  will increase since we are including more regions. At a critical level  $\lambda_\alpha$ , we will have exactly

$$V(\lambda_\alpha) = 1 - \alpha.$$

We will then use the level set  $L_{\lambda_\alpha}$  as the credible interval. You can show that under appropriate conditions ( $\pi(\theta|X_1, \dots, X_n)$  has not flat region and is Lipschitz),  $L_{\lambda_\alpha}$  is a shortest credible interval with a credible level  $1 - \alpha$ .

Because the above (shortest) credible interval is very straight forward, when people are talking about credible intervals, they are often referring to this interval.

### 5.3.3 Bayesian Testing: Bayes Factor

Bayesian hypothesis testing is also very straight forward – as you can guess, it is based on the posterior distribution. Instead of just putting priors on parameter, using our decision theoretic way, we can also *put priors on the hypotheses*.

Recall that  $H_0$  and  $H_a$  are the null and alternative hypotheses. Let  $\pi(H_0)$  and  $\pi(H_a) = 1 - \pi(H_0)$  denote the prior distribution on the two hypotheses and let  $\pi(H_0|X_1, \dots, X_n)$  and  $\pi(H_a|X_1, \dots, X_n)$  be the posterior distribution given the data.

The testing procedure is very simple – we reject  $H_0$  if

$$\frac{\pi(H_0|X_1, \dots, X_n)}{\pi(H_a|X_1, \dots, X_n)} < 1 \Leftrightarrow \frac{p(X_1, \dots, X_n|H_0)}{p(X_1, \dots, X_n|H_a)} \cdot \frac{\pi(H_0)}{\pi(H_a)} < 1. \quad (5.2)$$

Namely, we reject  $H_0$  if our posterior belief about  $H_0$  is less than that of  $H_a$ .

In many scenarios, the hypotheses will not directly give us a probability model related to the data. In a parametric model, they often put some constraints on the parameter. Thus, we can then rewrite the posterior  $p(X_1, \dots, X_n|H_0)$  as

$$p(X_1, \dots, X_n|H_0) = \int p(X_1, \dots, X_n, \theta|H_0) d\theta = \int p(X_1, \dots, X_n|\theta, H_0) \pi(\theta|H_0) d\theta.$$

Using the above equality, we define the **Bayes factor** as

$$\text{BF}(X_1, \dots, X_n) = \frac{p(X_1, \dots, X_n|H_0)}{p(X_1, \dots, X_n|H_a)} = \frac{\int p(X_1, \dots, X_n|\theta, H_0) \pi(\theta|H_0) d\theta}{\int p(X_1, \dots, X_n|\theta, H_a) \pi(\theta|H_a) d\theta}. \quad (5.3)$$

Using the Bayes factor, we can rewrite the decision rule as: we reject  $H_0$  if

$$\text{BF}(X_1, \dots, X_n) \cdot \underbrace{\frac{\pi(H_0)}{\pi(H_a)}}_{\text{odds}} < 1.$$

Namely, if the Bayes factor is less than the inverse of the odds of  $H_0$ , we reject  $H_0$ .

The Bayes factor can be viewed as a *Bayesian version of p-value* – the smaller, the less favor for  $H_0$ . To see how the Bayes factor is like the p-value, note that the Frequentist way of rejecting  $H_0$  is if the p-value is less than a pre-specified significance level  $\alpha$ . The Bayesian's threshold is given by the odds of  $H_0$  from the prior distribution. If the inverse of the Bayes factor is greater than the odds of  $H_0$ , we reject  $H_0$ . In Bayesian, the threshold in testing (significance level) has a simple interpretation – the odds of our prior belief about the null hypothesis.

**Example 5.9 (Bayesian Binomial distribution: revisited)** Now we come back to the Binomial sampling example where we have an observation  $Y \sim \text{Bin}(N, \theta)$  where  $N$  is known and the parameter of interest is  $\theta$ :

$$P(Y = y|\theta) = \binom{N}{y} \theta^y (1 - \theta)^{N-y}.$$

We use a Beta distribution with parameters  $(\alpha, \beta)$  as our prior distribution for  $\theta$ . Now we consider testing

$$H_0 : \theta = 0.5, \quad H_a : \theta \neq 0.5.$$

Our goal is to compute the Bayes factor. First, we compute the numerator:

$$p(Y|H_0) = \int_0^1 p(Y|\theta, H_0) \pi(\theta|H_0) d\theta = \int_0^1 p(Y|\theta) \delta(\theta = 0.5) d\theta = p(Y|\theta = 0.5) = \binom{N}{Y} 0.5^Y 0.5^{N-Y} = \binom{N}{Y} 0.5^N.$$

The denominator of the Bayes factor will be

$$\begin{aligned} p(Y|H_a) &= \int_0^1 p(Y|\theta, H_a) \pi(\theta|H_a) d\theta \\ &= \int_0^1 p(Y|\theta, H_a) \pi(\theta) d\theta \quad (\text{Note that } \theta = 0.5 \text{ is just a single point so it does not affect the integral}) \\ &= \int_0^1 \binom{N}{Y} \theta^Y (1 - \theta)^{N-Y} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1} (1 - \theta)^{\beta-1} d\theta \\ &= \binom{N}{Y} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \frac{\Gamma(Y + a)\Gamma(N - Y + b)}{\Gamma(N + a + b)}. \end{aligned}$$

Therefore, the Bayes factor is

$$\text{BF}(Y) = \frac{0.5^N}{\frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \frac{\Gamma(Y + a)\Gamma(N - Y + b)}{\Gamma(N + a + b)}}.$$

As you can see, the computation of Bayes factor does not involve the prior on  $H_0$ . However, it does involve the prior of the parameter  $\theta$  (and this is why the value of it depends on the hyperparameters  $\alpha, \beta$ ).

There is an interesting relation between the Bayes factor and the likelihood ratio test. In the above Binomial example, if we use the likelihood ratio test, we obtain a test statistic:

$$\text{LRT} = \frac{p(Y|\theta = 0.5)}{p(Y|\theta = \hat{\theta}_{MLE})} = \frac{\binom{N}{Y} 0.5^N}{\sup_{\theta \in [0,1]} \binom{N}{Y} \theta^Y (1 - \theta)^{N-Y}} = \frac{0.5^N}{(Y/N)^Y (1 - Y/N)^{N-Y}}.$$

The numerator is the same while the denominator is slightly different – the LRT only uses the maximum likelihood value whereas the Bayes factor uses the average value within  $H_a$ ! An interesting fact: if the alternative hypothesis is a simple hypothesis  $H_a : \theta = \theta_a$  for a fixed quantity  $\theta_a$ , then the Bayes factor and the LRT statistic coincides.

### 5.3.4 Bayesian nonparametric (♠♠♠)

Bayesian approach can also be applied to nonparametric estimation. Nonparametric estimation refers to the case where we do not assume the parameter of interest is a vector or number. One common problem is density estimation. If our goal is to estimate the underlying PDF or CDF, we can do it without assuming the data being from a parametric family such as a Gaussian. For instance, histogram can be used as a density estimator without assuming the data is from a Gaussian or any other parametric family (there are

many more advanced techniques such as the kernel density estimator, orthonormal basis approach, wavelet approach, ...etc).

The challenge of Bayesian inference in nonparametric problem is that we need to put a prior on ‘function space’ (a collection of functions). Often a function space does not admit a density function. However, there are some tricks that we can put a prior on function space.

- *Dirichlet process.* The Dirichlet process<sup>7</sup> is a stochastic process that generates random probability distributions. It also has several other interesting names such as the Chinese buffet process and stick-breaking process. The appealing feature of Dirichlet process is that it is constructed from an algorithmic scheme so there is a simple way to sampling points from a random distribution generated from this process.
- *Sequential mixture model.* Another way to assign priors on distributions is to use a sequential mixture model. For instance, a Gaussian mixture model with parameters drawn from some distribution. When we allow the number of mixture to increase with sample size, we can approximate many distributions in the function space. Note that when we allow the number of mixture to increase, this implies that our prior distribution is changing with respect to the sample size.
- *Basis approach.* A smooth density function with a compact support may be written as  $p_0(x) = \sum_{k=1}^{\infty} \theta_k \phi_k(x)$ , where  $\{\phi_k(x) : k = 1, 2, \dots\}$  is a basis and  $\{\theta_k : k = 1, 2, \dots\}$  are the coefficients. For instance, one may choose the cosine basis or spline basis. Putting priors on  $\theta_k$  leads to a prior distribution over functions. Note that in practice we often had to truncate the basis at certain level, i.e., we use only  $\{\phi_k(x) : k = 1, 2, \dots, N\}$  for some  $N = N_n$ . Often we allow  $N = N_n$  to increase with respect to the sample size so that our prior covers most part of the function space.
- *$\epsilon$ -cover/bracketing.* There is another approach to assign priors over function space using the covering/bracketing of function space. An  $\epsilon$ -cover is a collection of functions in a function space such that every function inside the space has a distance at most  $\epsilon$  to the nearest element inside the cover (bracketing is a generalization of this concept). There are many  $\epsilon$ -covers but we often use those with the minimal number of elements. Now consider a sequence of  $\epsilon$ -cover with  $\epsilon = \epsilon_n \rightarrow 0$  as  $n \rightarrow \infty$ . Let  $\Pi_n$  be a uniform distribution over each element of  $\epsilon_n$ -cover. Then the mixture distribution  $\Pi = \sum_{n=1}^{\infty} \beta_n \Pi_n$ , with an appropriate choice of  $\{\beta_n\}$  forms a prior distribution that can well-approximate almost every function in the function space (when the function space is sufficiently smooth). This idea is a powerful tool in constructing a prior with amazing theoretical properties<sup>8</sup> although it is often hard to numerically compute this prior.

## 5.4 Prediction

Prediction is about the question that given observations  $X_1, \dots, X_n$ , what will be the possible range of the next observation  $X_{n+1}$ ? Using probability model, we know that the next observation  $X_{n+1}$  will have a distribution. So the prediction problem is about finding the distribution  $p(x_{n+1}|X_1, \dots, X_n)$ . Here we will show how the Frequentist prediction method and the Bayesian prediction method are different in a parametric model.

The Frequentist approach is very simple – given that we have already assigned a parametric model  $p(x; \theta)$  for the observations, we should just use it as our predictive distribution. To use this model, we need to choose

<sup>7</sup>[https://en.wikipedia.org/wiki/Dirichlet\\_process](https://en.wikipedia.org/wiki/Dirichlet_process)

<sup>8</sup>see, e.g., the famous paper Ghosal et al. (2000) “Convergence rates of posterior distributions”: <https://projecteuclid.org/euclid.aos/1016218228>

the parameter  $\theta$ . A simple choice is the MLE  $\hat{\theta}_{MLE}$ . Thus, we will predict the distribution of  $X_{n+1}$  as

$$p(x_{n+1}|X_1, \dots, X_n) = p(x_{n+1}; \theta = \hat{\theta}_{MLE}).$$

The Bayesian approach again relies on the posterior. The predictive distribution  $p(x_{n+1}|X_1, \dots, X_n)$  can be written as

$$\begin{aligned} p(x_{n+1}|X_1, \dots, X_n) &= \int p(x_{n+1}, \theta|X_1, \dots, X_n) d\theta \\ &= \int p(x_{n+1}|\theta, X_1, \dots, X_n) \pi(\theta|X_1, \dots, X_n) d\theta \\ &= \int p(x_{n+1}|\theta) \underbrace{\pi(\theta|X_1, \dots, X_n)}_{\text{posterior distribution}} d\theta. \end{aligned}$$

Thus, the predictive distribution is the *averaged* distribution of  $p(x|\theta)$  where we average  $\theta$  over the posterior distribution.

Here, as you can see, the two paradigms make prediction using different principles – the Frequentists use only the *most likely* model to make predictions whereas the Bayesians use the *averaged* model over the posterior distribution to make predictions.