Probabilistic Models for Machine Learning (CS5560)

Summary Notes

January 4, 2022

Contents

Chapter 1

Overview of Probabilistic Models

The three main ingredients of Machine Learning pipeline are:

1. **Model:** This is mainly the way one chooses to represent/denote/describe concepts/knowledge relevant to the machine learning problem at hand. For example, so-called "probabilistic models", which are the models studied in this course, choose to represent concepts/knowledge via some appropriate likehihood function.

   Definition 1. *A likehood function is nothing but the probability mass function (pmf) in case of discrete random variables and the probability density function (pdf) in case of continuous random variables. It is often generalized to the case of multivariate random variables with some continuous and some discrete random variables using the Bayes rule. For example, consider the Bayes rule:* $p(x/y)p(y) = p(y/x)p(x)$*, where* $p(x/y), p(x)$ *are pmfs and* $p(y/x), p(y)$ *are pdfs. Then this common value is defined as the "joint likelihood",* $p(x, y) \equiv p(x/y)p(y) = p(y/x)p(x)$*. Further, the likelihoods* $p(y/x), p(x/y)$ *are qualified as "conditional liklihoods"; whereas* $p(x), p(y)$ *are qualified as "marginal likelihoods".* —[Important definition]

   (a) **Untrained Model:** This is a qualifier used to emphasize that "learning/training" has not yet been performed. In other words, untrained model denotes that model with no access/exposure (yet) to training data. Typically, untrained models are designed/hand-crafted using background/domain knowledge available apriori to accessing any training data.

   (b) **Trained Model:** This is a qualifier used to emphasize that "learning/training" has been performed and is ready for deployment in the application.

2. **Training Algorithm:** This is the key algorithm that uses the training data provided and "transforms" the untrained model into a trained model.

3. **Inference Algorithm:** This is the algorithm that uses the concepts encoded in the trained model and makes inferences useful for the application in hand. For example, in case of supervised learning, given a (new) input datapoint compute the corresponding label etc.

1.1 Types of Probabilistic Models

Based on what kind of a likelihood function is modelled, one can broadly categorize probabilistic models into two types:

1. **Generative Models:** Here, a likelihood function that can (re)generate training data is modeled. In other words, samples from the likelihood function modeled and those in the training data are expected to be `id` (identically distributed). For example, model $p(x)$ in case of unsupervised learning problems, or model $p(x, y)$ in case of supervised learning problems etc. Hence-forth, x is reserved for input-related variables and y is reserved for output/target/label related variables.

2. **Discriminative Models:** Here, the conditional of the target/output/label given the input is modelled. i.e., $p(y/x)$ is modelled. Needless to say, such models are relevant only for Supervised learning problems (and variants).

In order to enable any learning (from training data), untrained models must be flexible (choicy) enough. Based on the nature of flexibility/choice in the untrained models, we have the following categorization:

1. **Parametric Models:** Here, the untrained model is a collection/set of candidate likelihoods being modelled. The various candidates are indexed via, so-called, parameters. Which exact set of candidates are included is often designed using background/domain knowledge or convenience.

2. **Non-parametric Models:** Here, the untrained model is an/collection of abstract (yet valid) likelihood function(s) that is/are defined in terms of (not yet known) training data. Which abstract functional form is employed is often designed using background/domain knowledge or convenience.

Further, sometimes the choice/flexibility in the untrained model is specified as a soft-choice: for example, sometimes a prior likelihood that evaluates how

likely is each candidate likelihood function to be the "correct" likelihood function is provided. Again, background/domain knowledge is used to design such a prior likelihood. In other words, these are models where every candidate has some likelihood of being "correct". Such modelling is known as the Bayesian modelling.

At times, especially in so-called structured prediction problems, one models the likelihood function (either joint/marginal/conditional) in an indirect way via factorization. For example, instead of modeling $p(x, y)$ directly, one can (equivalently) model $p(x/y)$ and $p(y)$ individually. Such a indirect modeling has several advantages that may be clear later and is often performed in so-called Probabilistic Graphical Models.

1.2 Examples of Probabilistic Models

We begin with examples of Parametric Generative Models.

1.2.1 Parametric Generative Models

Standard (named) likelihoods provide our first examples: Bernoulli/Binomial, Multinoulli/Multinomial, Poisson, Uniform, Gaussian/Normal, Laplace, Gamma, Beta, Pareto, Dirichlet, Student T (refer sections 2.3-2.5 in Murphy [2012] for details) etc. For example, the (untrained) model that is collection of (univariate) Gaussian likelihoods with various means, variances is called as the (univariate) Gaussian Model. The parameters of these likelihoods are also well-known. For example, the mean, variance are the parameters for the Gaussian Model.

The most popular parametric generative models in machine learning are those from the so-called Exponential family:

Definition 2. *Let $\mathcal{X}$ be a given (arbitrary) set/domain/support. Let $\phi : \mathcal{X} \mapsto \mathbb{R}^n$ be a given (arbitrary) function, known as the feature map. Let $h : \mathcal{X} \mapsto \mathbb{R}^{++}$ be a given (arbitrary) function, known as the auxiliary likelihood. Consider the following likelihood, parametrized by w:* —[Important definition]

$$p_w(x) = \frac{1}{Z(w)} h(x) e^{w^\top \phi(x)}, \ x \in \mathcal{X}, \tag{1.1}$$

where $Z(w) \equiv \int_{\mathcal{X}} h(x) e^{w^\top \phi(x)} \mathrm{d}x$. Z is known as the Partition function. The domain of Z, i.e., $\mathcal{W} \equiv \{w \mid Z(w) < \infty\}$, is known as the Natural Parameter Space. Further, the parameters $w \in \mathcal{W}$ are sometimes qualified as Natural Parameters. The (untrained) model that consists of all likelihoods in (1.1) for

various parameters $w \in \mathcal{W}$ is said to be a model in the Exponential Family or is said to be a Maximum Entropy (Max-Ent) model. If the (natural) Parameter space is open[1] and convex[2], then the model is qualified as a linear (regular) exponential model. Various choices of h, ϕ correspond to the various models in this family. Typically, in machine learning, ϕ is designed using background/domain knowledge[3]. ϕ-statistic turns out be Sufficient[4] and hence ϕ is also popularly called as Sufficient Statistic.

It is not a difficult exercise to show that Bernoulli/Binomial, Multinoulli/Multinomial, Poisson, Gaussian/Normal, Laplace, Gamma, Beta, Pareto, Dirichlet (and many more) models are Max-Ent models[5] (refer sec.9.2.2. in Murphy [2012]). Infact, interested students may refer to the Pitman–Koopman–Darmois theorem to know how gigantic the Exponential family is. However, there are many useful models that do not belong to this family, e.g., Cauchy, hypergeometric, Student T, Gaussian mixture models etc.

While it is clear that the above can be employed in unsupervised learning, in case of supervised learning, generative models can also be employed using a joint feature map η: $p_w(x,y) = \frac{e^{w^\top \eta(x,y)}}{Z(w)}$. However, one needs to compute the corresponding conditional the $p_w(y/x)$ from $p_w(x,y)$ in order to perform inference i.e., $p_w(y/x) \propto p_w(x,y)$. However, whenever η does not factorize into x and y factors, $p_w(y/x)$ will not even belong to exponential family (for a fixed x). Hence, for convenience, η is often chosen as a function that factorizes: i.e., $\eta(x,y) \equiv \phi(x) \otimes \psi(y)$. Here $\otimes$ denotes tensor (outer) product. In this case it is easy to see that $p_w(y/x) \propto p_w(x,y) = e^{w^\top \phi(x) \otimes \psi(y)}$ belongs to exponential family with feature map as ψ and j^{th} parameter as $\sum_i w_{ij}\phi_i(x)$.

1.2.2 Parametric Discriminative Models

Once parametric generative models are understood, it is pretty easy to extend them to discriminative models. Recall that in such models, one needs to model a likelihood $p(y/x)$ (at every $x \in \mathcal{X}$). The most straight-forward way is to model the likelihood over $Y \in \mathcal{Y}$ at every x by the same generative model, for e.g., using an exponential family model with sufficient statistic $\phi : \mathcal{Y} \mapsto \mathbb{R}^n$. Now, denoting the

[1] Model is called regular if the natural parameter space is open.

[2] Natural paramater space is always convex.

[3] Unless it's a model selection problem, h does not matter because it is constant across likelihoods in a model.

[4] See `https://en.wikipedia.org/wiki/Sufficient_statistic`. Interested students please refer section 9.2.6 in Murphy [2012] to know why ϕ is a sufficient statistic.

[5] See `https://en.wikipedia.org/wiki/Category:Exponential_family_distributions`.

parameter at x by $\theta(x) \equiv W^\top \psi(x)$, for some appropriate feature map $\psi : \mathcal{X} \mapsto \mathbb{R}^d$, W is an $d \times n$ matrix, we have the parametric discriminative (exponential family) model specified by the conditional likelihood:

$$p_W(y/x) = \frac{1}{Z\left(W^\top \psi(x)\right)} h(y) e^{\psi(x)^\top W \phi(y)}, \ y \in \mathcal{Y}, \ x \in \mathcal{X}. \tag{1.2}$$

Note that the parameter for this model is $W_{d \times n}$ and by changing h, ψ, ϕ one can obtain various discriminative models in this family.

For example, in case $\mathcal{Y} = \mathbb{R}$, by fixing ϕ to be the sufficient statistic for the Gaussian model, and by defining $W^\top \psi(x) \equiv \begin{bmatrix} v^\top \eta(x)/a^2 \\ -1/\left(2a^2\right) \end{bmatrix}$, where $v, a > 0$ are parameters, one obtains that the likelihood $p_W(y/x)$ is a infact a Gaussian with mean as $v^\top \eta(x)$ and variance as a^2. This discriminative model, which belongs to the family defined by (1.2), is popular by the name Linear Regression Model. Please refer sections 7.1,7.2,7.3.2,7.3.3 in Murphy [2012] for details. In the case $\mathcal{X} = \mathbb{R}^n$, and joint feature map η is that corresponding to (multivariate) Gaussian, the same conditional (albeit with reparametrization) as in linear regression model can be obtained with a multivariate Gaussian generative model over $\mathcal{X} \times \mathcal{Y}$. Please refer theorem 4.31 in Murphy [2012] for more details.

Another popular example, in case $\mathcal{Y}$ is discrete, is obtained by fixing ϕ to be sufficient statistic for the Multinoulli, is known as the Logistic Regression Model. Please refer sections 8.1,8.2,8.3.7 in Murphy [2012] for details.

1.2.3 Parametric Bayesian Models

We critically observe that often in ML applications, the parameters have physical meaning and hence for the same reasons as why probabilistic modeling of X, Y (inputs/labels) is desirable, it is also desirable to probabilistically model the (uncertainties in the) parameters too! This is the core ideology in Bayesian methods.

Accordingly, in Bayesian models either $p(\theta, x, y)$ (generative) or $p(\theta, y/x)$ (discriminative) is specified. Here, θ is the parameter. In the following expressions are given for generative models, but the analogous expressions for discriminative models straight-forwardly follow.

Now, $p(\theta, x)$ is factorized as: $p(\theta, x) = p(x/\theta)p(\theta)$. The first factor is given by standard likelihood models (like those in exponential family). For example, $p(x/\theta)$ is Gaussian likelihood with parameter θ etc. The second factor is an appropriate "prior" likelihood over the parameters, which encodes the background information already available regarding the parameters. Overall, $p(\theta, x)$, captures the apriori beliefs/uncertainties regarding the involved parameters/labels/inputs.

In Bayesian framework, this specification of $p(\theta, x)$ is the so-called "untrained (prior) model". Note that the concept of untrained model is very different from that in non-Bayesian models in section 1.2. This is because parametric non-Bayesian models, are essentially collections of likelihoods with an unknown parameter; whereas in Bayesian models, there is only one likelihood, $p(\theta, x)$, which is completely known!

Examples of parametric Bayesian models are given in sections 3.3-3.3.2,3.4-3.4.2,3.5.1.2,7.6,7.6.3,8.4 in Murphy [2012]. In general, one can employ any exponential family likelihood and define a prior of the (natural) parameters to obtain a corresponding Bayesian model.

1.2.4 Half/Semi Bayesian (parametric) Hybrid Models

These methods are a mix of Bayesian and non-Bayesian methods. Consider a half-Bayesian (untrained) model, where a set of parametrized likelihoods are given (say an exponential family model); however, motivated by Bayesian methods, we assume a prior over the parameters is provided. Unlike Bayesian methods, the parameter is to be "learnt/estimated" using the training dataset as well as the prior.

1.3 Graphical Models

Probabilistic graphical models [Koller and Friedman, 2009] offer a very convenient way for encoding rich background/prior knowledge about dependencies among the relevant random variables. Such encoding eventually leads to improved training and inference. While Bayes Nets (section 1.3.1) model directional/causal dependencies, the MRFs (section 1.3.2) model symmetric/undirectional dependencies.

1.3.1 Bayesian Networks

Bayes nets or Bayesian networks or Directed Graphical Models or Belief/Decision networks intuitively model "causal" inter-dependencies among the relevant random variables using Directed Acyclic Graphs (DAGs). The random variables form the nodes of the DAG and directed edges intuitively encode causal relationships e.g., $X \to Y$ implies X is a direct cause of Y etc.

More formally, the DAG in the Bayes Net defines a set of conditional independence conditions, known as the local Markov in-dependencies: For each

random variable X_i: we have $X_i \perp nd(X_i) \backslash pa(X_i) \backslash \{X_i\} \mid pa(X_i)$. Refer definition 3.1 in Koller and Friedman [2009]. Here, $pa(X), nd(X)$ denote the parents and non-descendants of X respectively.

We say a DAG, $\mathcal{G}$, is an I-map for likelihood function p iff all the local Markov in-dependencies are true in p. Interestingly, given a topological ordering of nodes in $\mathcal{G}$, we can show that p factorizes as: $p(x_1, \ldots, x_n) = \prod_{i=1}^{n} p(x_i \mid pa(x_i))$, which is known as the chain rule for the Bayes Net.

1.3.2 Markov Random Fields

1.3.3 Conditional Random Fields

Chapter 2

Training Algorithms

2.1 Parametric Models

In case of parametric models, the problem of learning can be posed as that of estimating model parameters from the given training data (assuming training is nothing but a set of iid samples from the estimated likelihood). We discuss below two classical methods for parameter estimation.

2.1.1 Method of (matching) Moments

Motivated by law of large numbers, here, the population/true expectation and the sample/empirical mean of a sufficient statistic ϕ are matched to "learn" the parameters of the generative model:

$$\min_{w \in \mathcal{W}} \left\| \mathbb{E}_{X \sim p_w} [\phi(X)] - \frac{1}{m} \sum_{i=1}^{m} \phi(x_i) \right\|^2, \tag{2.1}$$

where $\mathcal{D} = \{x_1, \ldots, x_m\}$ is the training set. This is known as method of moments or method of matching moments. An interesting special case is when the optimal objective of (2.1) is zero, which leads to the popular moment matching equation:

$$\mathbb{E}_{X \sim p_w} [\phi(X)] = \frac{1}{m} \sum_{i=1}^{m} \phi(x_i) \tag{2.2}$$

2.1.2 MLE or M-projection

Motivated by law of large numbers, here, the likelihood, p_w, and the sample/empirical likelihood[1], $p_{\mathcal{D}}$, are matched to "learn" the parameters of the generative model:

$$\min_{w \in \mathcal{W}} \; KL(p_{\mathcal{D}} \parallel p_w), \tag{2.3}$$

where $\mathcal{D} = \{x_1, \ldots, x_m\}$ is the training set, and $KL(p \parallel q) \equiv \mathbb{E}_{X \sim p}\left[\log\left(\frac{p(X)}{q(X)}\right)\right]$, is the KL-divergence between likelihoods p, q. Due to the resemblence of (2.3) to the definition of projections, the parameter esitmation via (2.3) is known as M-projection estimation. Interested students may refer sections 8.5.1-8.5.3 for details.

Note the similar forms of (2.1) and (2.3). Infact, theorem 8.6 in Koller and Friedman [2009] shows that if optimal value of (2.1) is zero, then the optimal solutions of (2.1,2.3) are the same.

A straight-forward simplification shows that (2.3) is equivalent to:

$$\max_{w \in \mathcal{W}} \; p_w(\mathcal{D}), \tag{2.4}$$

where[2] $p_w(\mathcal{D}) \equiv p_w(x_1) \ldots p_w(x_m)$ is the likelihood of the training data. For obvious reasons, the parameter estimated via (2.4) (or equivalently via (2.3)) is known as maximum likelihood estimation. For exponential family models, (2.3) simplifies as:

$$\min_{w \in \mathcal{W}} \; \log\left(Z(w)\right) - w^\top \left(\frac{1}{m}\sum_{i=1}^{m} \phi(x_i)\right), \tag{2.5}$$

From the Holder's inequality it follows that the parameter space, $\mathcal{W}$, is always convex (e.g., refer Theorem 1 in `https://people.eecs.berkeley.edu/~jordan/courses/260-spring10/other-readings/chapter8.pdf`). From this it is straight-forward to show that the cummulant generating function is also convex. So problem (2.5) is a convex program. Additionally if we assume $\mathcal{W}$ is an open-set[3], then, standard results in convex optimization provide that w^* is an optimal solution of (2.5) if and only if gradient of the objective in (2.5) (wrt. w) at w^* is zero. Derivation in section 9.2.3 in Murphy [2012] shows that this gradient is: $\mathbb{E}_{X \sim p_w}[\phi(X)] - \frac{1}{m}\sum_{i=1}^{m} \phi(x_i)$. This result implies that the optimal solutions of MLE/M-projection are exactly same as those of MM (Method of Moments), whenever optimal value of (2.1) is zero. Also, in case (2.5) is not solvable, minimizing the norm of this gradient, (equivalent to (2.1)), may be desirable.

[1] Sample/Empirical likelihood is simply a uniform likelihood over the given samples.

[2] Abuse of notation!

[3] i.e., if we have a linear exponential model.

In case of discriminative models, methodologies like matching moments and MLE (M-projection) do not have a trivial generalization mainly because there are no multiple samples for the label y for a given input x in order to estimate $p(y/x)$. In the following sections, clever generalizations of these methodologies are presented.

2.1.3 Method of Partial/Conditional Moments

Recall that for exponential family generative models with factorized feature maps, $\eta(x,y) = \phi(x) \otimes \psi(y)$, the method of moments equation is: $\mathbb{E}_{(X,Y)\sim p_w}\left[\phi(X) \otimes \psi(Y)\right] = \frac{1}{m}\sum_{i=1}^{m} \phi(x_i) \otimes \psi(y_i)$. Also, recall that in case of discriminative models, the joint $p(x,y)$ is not modeled and only the conditional, $p(y/x)$, is modeled. Hence one cannot apply the above equation of matching moments directly. A way out is to construct a joint $q_w(x,y)$ from the modeled conditional $p_w(y/x)$ and the empirical marginal over inputs: $\hat{p}(x) = \frac{1}{m}\delta_{x_i}(x)$, i.e., $q(x,y) \equiv p_w(y/x)\hat{p}(x)$. Now, from total expectation rule: $\mathbb{E}_{(X,Y)\sim q_w}\left[\phi(X) \otimes \psi(Y)\right] = \mathbb{E}_{X\sim\hat{p}}\left[\mathbb{E}_{Y/X\sim p_w}\left[\phi(X) \otimes \psi(Y)/X\right]\right] = \frac{1}{m}\sum_{i=1}^{m} \mathbb{E}_{Y/X=x_i\sim p_w(y/x_i)}\left[\phi(x_i) \otimes \psi(Y)/x_i\right]$. This leads to the following method of partial/conditional moments:

$$\min_{w\in\mathcal{W}} \left\| \sum_{i=1}^{m} \phi(x_i) \otimes \mathbb{E}_{Y/X\sim p_w}\left[\psi(Y)/x_i\right] - \sum_{i=1}^{m} \phi(x_i) \otimes \psi(y_i) \right\|^2, \tag{2.6}$$

and the corresponding partial/conditional moment matching equation is:

$$\sum_{i=1}^{m} \phi(x_i) \otimes \mathbb{E}_{Y/X\sim p_w}\left[\psi(Y)/x_i\right] = \sum_{i=1}^{m} \phi(x_i) \otimes \psi(y_i) \tag{2.7}$$

2.1.4 MCLE

Let's begin by extending the notion of M-projection to discriminative models. Naive way is to find M-projection of the empirical conditional $\hat{p}_{\mathcal{D}}(y/x)$ onto the (discriminative exponential family) model. However, it is non-trivial to estimate $\hat{p}_{\mathcal{D}}(y/x)$ unless multiple samples of y are given for every x, which is not the case in ML.

Hence, here we insist on a relaxed M-projection:

$$\min_{w\in\mathcal{W}} \mathbb{E}_{X\sim\hat{p}(x)}\left[KL(\hat{p}_{\mathcal{D}}(\cdot/X) \parallel p_w(\cdot/X))\right], \tag{2.8}$$

Using the definition of KL and the total expectation rule, the above simplifies to[4]:

$$\max_{w \in \mathcal{W}} \ \Pi_{i=1}^{m} p_w(y_i/x_i), \tag{2.9}$$

which is nothing but maximizing the conditional likelihood of training data, hence the name Maximum Conditional Likelihood Estimation (MCLE).

For exponential family discriminative models, the optimality conditions for (2.9) turn out to be same as (2.7). Refer `https://1drv.ms/b/s!Au6Zdrbq2x4ph7RqM1dNyd1wu4qu3A?e=UKwxIQ` for details. Section 7.3.1 in Murphy [2012] simplifies this equation for the special case of linear regression model.

2.1.5 Computing gradient - Sampling Techniques

A simple technique for solving (2.5) or (2.1) or (2.7) is (projected/conditional/stochastic) gradient descent. For example, refer section 8.3.2 in Murphy [2012].

However computing the gradient requires to compute $\mathbb{E}_{X \sim p_w}[\phi(X)]$ or $\mathbb{E}_{Y/x_i \sim p_w}[\psi(Y)/x_i]$, at every iterate $w^{(k)}$, which may be computationally challenging and/or may require clever estimation. Also, the likelihood $p_w = \frac{\tilde{p}_w}{Z(w)}$, is known only up-to a constant i.e., the un-normalized likelihood $\tilde{p}_w$ is known.

One way to estimate this expectation is importance sampling explained in section 23.4 in Murphy [2012]. Such methods work well with low-dimensional feature maps. The idea used is very generic and has applications elsewhere in ML.

A more popular sampling technique is the Metropolis Hasting algorithm explained in sections 24.3 in Murphy [2012]. Gibbs sampling is a special case of it detailed in section 24.2 in the same book. Such methods, known as Markov Chain Monte Carlo (MCMC) methods, essentially sample from a Markov chain whose limiting stationary likelihood is p_w.

2.2 Bayesian Learning

Recall that in (untrained) Bayesian models, there is no "unknown" to be "learnt". This is because the (apriori) joint $p(\theta, x)$ is completely specified. So the natural question is what constitutes learning in Bayesian models? In Bayesian framework, learning is simply updating the prior beliefs encoded by $p(\theta, x)$ based on the training data. In other words, the goal is Bayesian learning is to find $p(\theta, x/\mathcal{D})$, where $\mathcal{D}$ is the training set. We shall refer to this $p(\theta, x/\mathcal{D})$ as the trained or

[4]Refer `https://www.iith.ac.in/~saketha/teaching/cs5560Scribe.pdf` for details.

posterior Bayesian model. As detailed below, this can be found by simple rules of probability:

$$
\begin{aligned}
p(\theta, x/\mathcal{D}) &= p(x/\theta, \mathcal{D})p(\theta/\mathcal{D}) \\
&= p(x/\theta)p(\theta/\mathcal{D}) && (\because Assumption1) \\
&\propto p(x/\theta)p(\mathcal{D}/\theta)p(\theta) && (\because BayesRule) \\
&= p(x/\theta)\Pi_{i=1}^{m} p(x_i/\theta)p(\theta) && (\because Assumption1)
\end{aligned} \tag{2.10}
$$

The above is key equations that relate the untrained and trained Bayesian models. Here, Assumption1 is that the samples (training or otherwise) are (mutually) independent given the parameter. Note that this assumption is very different from the assumption in non-Bayesian methods that the samples are (mutually, unconditionally) independent. Note that the last RHS in 2.10 is completely known given the untrained model. The term $p(\theta/\mathcal{D})$ is referred to as the posterior of parameters. Refer sections 3.3.3,3.4.3,4.6.1-4.6.3,7.6.1 in Murphy [2012] for details of posterior with various Bayesian models.

2.3 MAP Estimation

In case of the half/Semi Bayesian models presented in section 1.2.4, one may use the so-called MAP estimation procedure for learning the parameters:

$$\theta_{MAP} \equiv \arg\max_{\theta \in \Theta} p(\theta/\mathcal{D}), \tag{2.11}$$

where Θ is the parameter space. Note the similarity of (2.11) with (2.4); only the prior factor is included in the former. Refer sections 5.2.1,3.3.3.1,3.4.3,4.6.2.1,4.6.3.4,7.5,8.3.6 in Murphy [2012] for MAP estimates with few models.

Chapter 3

Inference Algorithms

Once the model is trained, then the trained model can be employed for various purposes. Depending on the application, various inference queries might be fired on the trained model. Here we study algorithms for some common inference queries.

3.1 Supervised Learning

In case of discriminative models trained with MLE/MAP, the posterior, $p_{w*}(y/x)$ is readily available. If the model is generative, then obtaining the posterior is not difficult as $p_{w*}(y/x) \propto p_{w*}(x, y)$. In case of Bayesian models one needs to perform Bayesian Model Averaging to obtain this posterior as described in the section below.

3.1.1 Bayesian Averaging

In case of Bayesian models, the posterior needs to be computed and is not readily available after training. The posterior predictive is computed as the so-called Bayesian average: $p(y/x, \mathcal{D}) = \int_\Theta p(y/x, \theta) p(\theta/\mathcal{D}) \, \mathrm{d}\theta$.

3.1.2 Loss-based inference

Given an input x, though the posterior provides a likelihood over y, the classification problem requires one to come up with a single label, $\hat{y}$. In other words, the inference problem is that of finding a label $\hat{y}$ that is "best" representative of the

posterior likelihood, $p(y/x)$. Here, $p(y/x)$ is $p_{w*}(y/x)$, in case of plug-in estimates like MLE/MAP and $p(y/x)$ is $p(y/x, \mathcal{D})$ for Bayesian models. Now, we pose this inference as the following optimization problem:

$$\hat{y} \equiv \arg\min_{y \in \mathcal{Y}} \mathbb{E}_{Y/x \sim p}[l(y, Y)/x], \tag{3.1}$$

where $l : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}$ is a loss function that captures the loss incurred by replacing the arguments with each other. In other words, we wish $\hat{y}$ to be that label that incurs least expected loss when it replaces a random label drawn from the trained posterior. We discuss algorithms for solving (3.1) in classification and regression below. Refer section 5.7 in Murphy [2012] for details.

Classification

In case of classification, the label space is discrete and hence the so-called 0-1 loss is well-suited: $l(y, y^{'}) = 1 \iff y \neq y^{'}$, and zero otherwise. With this loss, (3.1) turns out to be that of finding the mode of the posterior, $p_{w^*}(y/x)$. For exponential family discriminative models this is same as:

$$\hat{y} \equiv \arg\max_{y \in \mathcal{Y}} \phi(x)^{\top} W^* \psi(y). \tag{3.2}$$

Note that the inference is same as above even for the generative exponential family models with factored joint feature maps. If p, q are dimensionalities of input,label respectively, and c is the number of classes, then the computational complexity of finding posterior mode is $O(pq + cq)$. Refer section 5.7.1.1 in Murphy [2012] for details.

Regression

When the label space is R or it's uncountable subset, the natural loss function is the so-called squared-loss: $l_2(y, y^{'}) \equiv (y - y^{'})^2$. With this loss, (3.1) turns out to be same as computing the posterior mean, $\mathbb{E}_{Y/x \sim p_{w^*}}[Y/x]$. This can be estimated using any MCMC technique. Refer section 5.7.1.3 in Murphy [2012] for details.

3.2 Unsupervised Learning

In case of MLE/MAP, the generative $p_{w^*}(x)$ is readily available, whereas in Bayesian models, this generative is given by Bayesian averaging: $p(x/\mathcal{D}) = \int_{\Theta} p(x/\theta) p(\theta/\mathcal{D}) \, \mathrm{d}\theta$. Often computing this integral may be computationally costly and we may have

to resort to some MCMC based estimation. In the discussion below, $p(x)$ denotes either $p_{w^*}(x)$ or $p(x/\mathcal{D})$.

3.2.1 Novelty Detection

In novelty detection applications, given an input, x_0, one needs to infer whether it is a "novel" input (as compared to those seen in the training data) or not. One way to infer this is to fix some (low enough) threshold, ϵ, and infer that x_0 is novel iff $p(x_0) < \epsilon$.

3.2.2 Clustering

Assuming clusters are defined as contiguous regions of high density. Below is one inference algorithm [Ben-Hur et al., 2001] for finding clusters whenever the domain is a vector space.

1. fix a high enough threshold, h, and find highly density training datapoints: $H \equiv \{x \in \mathcal{D} \mid p(x) \geq h\}$.
2. build a graph with nodes as training datapoints and edges present iff the points in the line segment joining the nodes of the edge are all high density points i.e., their density is greater than h. In practice, one only checks at a fixed number of (equidistant) points in the line segement.
3. The connected components in this graph are declared as clusters.
4. In case the clustered information has to be retrieved for a new datapoint, x_0:
   (a) If $p(x_0 < h)$, then it is declared as any "outlier (novel point)" not belonging to any cluster.
   (b) Else, graph is modified (re-constructed efficiently) to include x_0 appropriately and then cluster information is again inferred based on the modified connected components.

Chapter 4

Model Selection

Given the great abundance of models described in chapter 1, in typical ML applications, one has to "choose" among multiple promising models. Here we discuss algorithms for "choosing" the "right" model.

4.1 Maximum Marginal Likelihood Estimation

Here we consider the special setting where a collection of Bayesian models each with a different hyper-parameter is available i.e., $p_\alpha(x, \theta)\ \alpha \in \mathcal{A}$. Here, $\mathcal{A}$ is the hyper-parameter space. Then model selection problem is nothing but hyper-parameter estimation problem.

Now, consider the corresponding Bayesian averaged models, or in other words, the marginal likelihoods: $p_\alpha(x) = \int_\Theta p_\alpha(x, \theta)\ \mathrm{d}\theta = \int_\Theta p(x/\theta) p_\alpha(\theta)\ \mathrm{d}\theta$. One way of performing hyper-parameter selection is to "pick" the best among these averaged models given some training data, which is a familiar problem discussed earlier (e.g., see section 2.1.2). For example, one can employ the MLE, which reads as the following in this case:

$$\alpha_{MMLE} \equiv \arg\max_{\alpha \in \mathcal{A}} p_\alpha(\mathcal{D}) = \arg\max_{\alpha \in \mathcal{A}} \frac{Z(alpha)}{N(alpha, \mathcal{D})}, \tag{4.1}$$

where Z is the normalization constant of the prior, and N is that of the posterior. Refer sections 5.3.2.1-5.3.2.3 in Murphy [2012] for some examples. This avatar of MLE is known as the Maximum Marginal Likelihood Estimation as the likelihoods are actually marginal likelihoods! Read section 5.3.1 in Murphy [2012] for an intuitive explanation of MMLE.

4.2 Hierarchical Bayesian Methods

Continuing the ideas in the previous section, after performing hyper-parameter estimation using MMLE, one can either perform MLE/MAP/BayesianAveraging within this "best" model. Further, one can define a prior over the hyper-parameters and may then perform MAP/BayesianAveraging over the hyper-parameters. The parameters of the hyper-parameter prior may then be called as hyper-hyper-parameters :) Such models, with two or more levels of priors defined are knwon as hierarchical Bayesian models. The fisrt table on page 173 lists some of the named combinations.

Hierarchical Bayesian models have far reaching applications beyong model selection. Some examples are given in sections 5.5.1,5.6.2.1 in Murphy [2012].

4.3 Validation

Refer section 6.5.3 in Murphy [2012] for details.

Bibliography

A. Ben-Hur, D. Horn, H. T. Siegelmann, and V. .Vapnik. Support Vector Clustering. *Journal of Machine Learning Research*, 2:125–137, 2001.

D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.

Kevin P. Murphy. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012. ISBN 0262018020, 9780262018029.