

DD2434 Machine Learning, Advanced Course

Assignment 1

Carl Henrik Ek

Deadline 12:00 (noon) (CET) November 19th, 2015

You will present the assignment will by a written report that you can mail to me at chek@kth.se before the deadline. From the report it should be clear what you have done and you need to support your claims with results. You are supposed to write down the answers to the specific questions detailed for each task. This report should clearly show how you have drawn the conclusions and come up with the derivations. Your assumptions, if any, should be stated clearly. For the practical part of the task you should not show any of your code but rather only show the results of your experiments using images and graphs together with your analysis.

Being able to communicate your results and conclusions is a key aspect of any scientific practitioner. It is up to you as a author to make sure that the report clearly shows what you have done. Based on this, and only this, we will decide if you pass the task. No detective work should be needed on our side. Therefore, neat and tidy reports please!

I very much recommend you to get used to \LaTeX to write your report. It is an amazing tool that you will find very useful in your further endeavors as a scientist.

The grading of the assignments will be as follows,

E Completed Task 2.1 and 2.2.

D E + Completed Task 2.3.

C D + Completed Task 2.4.

B C + Completed Task 2.5.

A B + Completed Task 2.6.

These grades are valid for review November 23rd, 2015. See the course [web page](#), HT 2015 - Assignments in the menu, for grading of delayed assignments.

Abstract

In this assignment we will examine several different aspects of building models of data. In the first task we will look at a supervised scenario where we are try and learn a model of a specific relationship between two different domains. This is a very common problem where we have observations in one domain, say an image of a face and then wish to infer the identity of the person. The second task we look at how we can perform unsupervised learning and learn a new representation of the data. This is related to finding hidden structures or patterns in the data which might contain important information. Finally we will end with a look at how we can approach model selection. This is very important as it gives us the tool to design different models and then choose the one that best represents our data. The important message that these exercises tries to convey is how we can integrate our beliefs with observations using a set of simple rules. The assignments are aimed at showing the key aspects of data modeling in a simple scenario such that our insights about the models are not “clouded” by the complexity data. It is left to you as a student to extend this knowledge to a realistic scenario with real data.

I The Prior $p(\mathbf{X})$, $p(\mathbf{W})$, $p(f)$

2.1 Theory

Regression is the task of estimating a continuous target variable \mathbf{Y} from an observed variate \mathbf{X} . The target and the observed variates are related to each other through a mapping,

$$f : \mathbf{X} \rightarrow \mathbf{Y}, \quad (1)$$

where f indicates the mapping. Given input output pairs $\{\mathbf{x}_i, \mathbf{y}_i\}_1^N$ our task is to estimate the mapping f such that we can infer the associated \mathbf{y}_i from previously unseen \mathbf{x}_i . In this task we will work with real vectorial data such that $\mathbf{x}_i \in \mathbf{X}$ where $\mathbf{x}_i \in \mathcal{R}^q$ and $\mathbf{y}_i \in \mathbf{Y}$ where $\mathbf{y}_i \in \mathcal{R}^D$. Being probabilistic means that we need to consider the uncertainty in both the observations as well as the relationship between the variates. Starting with the relationship between two *single* points \mathbf{x}_i^j and \mathbf{y}_i^j we can assume the following form of the likelihood,

$$p(\mathbf{y}_i^j | f, \mathbf{x}_i^j) \sim \mathcal{N}(f(\mathbf{x}_i), \sigma^2 \mathbf{I}). \quad (2)$$

Question 1: *Why is choosing a Gaussian likelihood a sensible thing to do? What does it mean that we have chosen a spherical covariance matrix for the likelihood?*

Assuming that each output point is independent given the input and the mapping we can write the likelihood of the data as follows,

$$p(\mathbf{Y} | f, \mathbf{X}) = \prod_i^N p(\mathbf{y}_i | f, \mathbf{x}_i). \quad (3)$$

Question 2: *If we do not assume that the data points are independent how would the likelihood look then? Remember that $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_N]$*

The task of regression means that we wish to infer \mathbf{y}_i from its corresponding variate \mathbf{x}_i . These two variates are related to each other by the mapping f so from a probabilistic view point we wish to find the mapping from the observed data. More specifically, taking uncertainty into account, what we wish to reach is the posterior distribution over the mapping given the observations,

$$p(f | \mathbf{X}, \mathbf{Y}). \quad (4)$$

2.1.1 Linear Regression

I recommend that you read Chapter 3 in Bishop 2006 before you start this part of the task. In order to proceed lets make an assumption about the mapping and model the relationship between the variates as a linear mapping. Further, lets make an assumption about the structure of the noise in the observations, assuming that they have been corrupted by additive Gaussian noise,

$$\mathbf{y}_i = \mathbf{W}\mathbf{x}_i + \epsilon, \quad (5)$$

where $\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$. From this we can formulate the likelihood of the data,

$$p(\mathbf{Y}|\mathbf{X}, \mathbf{W}) = \quad (6)$$

Question 3: *What is the specific form of the likelihood above, complete the right-hand side of the expression.*

The inference task we are interested in here is to learn the mapping, i.e. to infer \mathbf{W} from the data,

$$p(\mathbf{W}|\mathbf{X}, \mathbf{Y}) = \frac{1}{Z}p(\mathbf{Y}|\mathbf{X}, \mathbf{W})p(\mathbf{W}). \quad (7)$$

In the above equation we can see that we need to formulate our belief of the model parameters \mathbf{W} in a prior $p(\mathbf{W})$. We can make many different choices of priors, but a sensible choice would be to pick the conjugate prior i.e. a Gaussian prior over the parameters,

$$p(\mathbf{W}) = \mathcal{N}(\mathbf{W}_0, \tau^2\mathbf{I}). \quad (8)$$

Question 4: *Explain the concept of conjugate distributions. Why is this a motivated choice?*

The prior distribution provides us with a tool that tells us how likely or “how far” a parameter is from our belief.

Question 5: *What preferences in solution would these different priors encode?*

- $p(\mathbf{W}) \propto e^{-\|\mathbf{W}\|_{L_2}}$
- $p(\mathbf{W}) \propto e^{-\|\mathbf{W}\|_{L_1}}$

where L_1 and L_2 indicates the 1-norm and 2-norm.

The posterior is the object that integrates our prior beliefs with the data. In the next section we will see how this works in practice for the linear regression model that we have derived above.

Question 6: *Derive the posterior over the parameters. I recommend that you do these calculations by hand as it is very good practice. However, in order to pass the assignment you only need to outline the calculation and highlight the important steps.*

- *Why does it have the form that it does?*
- *What is the effect of the constant Z , are we interested in this?*

2.1.2 Non-parametric Regression

In the previous task we made the assumption that the relationship between the two variates were linear. This is quite a strong assumption that severely restricts the representative power of our model. The obvious way to proceed would be to add more parameters to f and use a higher-degree polynomial, but which one should we pick, degree 3 or 4, should we add a trigonometric function? These are very tricky questions that requires a lot of knowledge about the specific data that we are looking at, but we want to stay general, and the whole idea about Machine Learning is that we want the *data* to tell us this, not specify it before we start building our model.

Lets take a step back and think how a Bayesian would think in this situation. The above argument just says that we have a large uncertainty in, not only in the the parameters of the mapping, but also of the actual *form* of the mapping. Bayesian reasoning allows us to deal with this, its actually *exactly* this scenario that it was designed to deal with, we just need to formulate our uncertainty about the mapping in a prior over mappings and then use Bayes rule to reach the posterior, job done! The problem is just that we need to somehow formulate a prior over a the space of functions, which is quite a lot stranger mathematical object compared to the scalar valued parameters \mathbf{W} in the previous task. Before proceeding with this task I recommend that you read Bishop 2006, page 303 – 311

We know from the lecture that Gaussian Processes (\mathcal{GPs}) can be used to represent prior distributions over the space of functions. Rather than specifying a specific parametric form of the function \mathcal{GPs} is a non-parametric model.

Question 7: *What is a non-parametric model and what is the difference between non-parametrics and parametrics? In specific discuss these two aspects of non-parametrics,*

- *Representability?*
- *Interpretability?*

We will now proceed to look at the regression problem where we replace the linear assumption in the mapping to use a non-parametric prior over the space of functions. Lets make the same assumption about the observations as we did in the linear case,

$$\mathbf{y}_i = f(\mathbf{x}_i) + \epsilon. \quad (9)$$

This allows us to formulate the likelihood in the same manner as before. However, in the linear example we could easily formulate the relationship between \mathbf{x} , \mathbf{y} and f as the latter had a simple parametric form. Now we cannot do this anymore. To proceed, lets define the output of the function f as its own random variable,

$$\mathbf{y}_i = f_i + \epsilon, \quad (10)$$

where f_i is the *output* of the function at input location \mathbf{x}_i . The next step is to formulate the prior over the output of the function. This we can do using a \mathcal{GP} ,

$$p(f|\mathbf{X}, \boldsymbol{\theta}) = \mathcal{N}(\mathbf{0}, k(\mathbf{X}, \mathbf{X})) \quad (11)$$

where $k(\cdot, \cdot)$ is the covariance function and $\boldsymbol{\theta}$ is its parameters. We will refer to $\boldsymbol{\theta}$ as the *hyper-parameters* of the process. In this we have assumed that the data have been translated such as to have zero-mean such that we do not need to have a mean function in the prior.

Question 8: *Explain what this prior does? Why is it a sensible choice? Use images to show your reasoning. Clue: use the marginal distribution to explain the prior*

Given this formulation we can now formulate the full model.

Question 9: Formulate the joint likelihood of the full model that you have defined above,

$$p(\mathbf{Y}, \mathbf{X}, f, \boldsymbol{\theta})$$

Draw the graphical model to clearly show the assumptions that you have made.

Annoyingly we have added a new variable to our model which we are not really interested in. Specifically we have modeled the relationship between \mathbf{Y} and f and also f and \mathbf{X} but what we really are interested in is the relationship between \mathbf{Y} and \mathbf{X} . The motivation behind this is that we now have the possibility to have uncertainty in each of these stages, in our beliefs of the functions, and in how we believe the output of the function have generated the observed data. But again, we are not interested in f and therefore the variable should be marginalised out. Performing the marginalisation implies calculating the following integral,

$$p(\mathbf{Y}|\mathbf{X}, \boldsymbol{\theta}) = \int p(\mathbf{Y}|f)p(f|\mathbf{X}, \boldsymbol{\theta})df. \quad (12)$$

Question 10: Explain the marginalisation in Eq.12,

- Explain how this connects the prior and the data?
- How does the uncertainty “filter” through this?
- What does it imply that $\boldsymbol{\theta}$ is left on the left-hand side of the expression after marginalisation?

2.2 Practical

Now we will implement the approach we studied in the previous part. Remember to save images and figures to support your claims in part 1 as this will make the presentation much easier for me, and therefore also for you. There are a couple of packages that I found really useful when I implemented this code,

```
1 import pylab as pb
2 import numpy as np
3 from math import pi
4 from scipy.spatial.distance import cdist
5
6 # To sample from a multivariate Gaussian
7 f = np.random.multivariate_normal(mu,K);
8 # To compute a distance matrix between two sets of vectors
9 D = cdist(x1,x2)
10 # To compute the exponential of all elements in a matrix
11 E = np.exp(D)
```

2.2.1 Linear Regression

In this task we will implement the linear regression that we looked at in the previous task. We will look at both the prior and the posterior over the parameters \mathbf{W} and evaluate the effect this will have on the model. To do so we will need to have some data to experiment with. What we want to show is that the methodology that we have learned is capable of recovering the true underlying mapping from observed data, therefore lets generate some data and then simply throw the generating parameters

away.

$$y_i = w_0 x_i + w_1 + \epsilon \quad (13)$$

$$\mathbf{x} = [-1, -0.99, \dots, 0.99, 1] \quad (14)$$

$$\epsilon \sim \mathcal{N}(0, 0.3) \quad (15)$$

$$\mathbf{W} = [-1.3, 0.5] \quad (16)$$

Question 11:

1. Visualise the prior distribution over \mathbf{W} .
2. Pick a single data-point from the data and visualise the posterior distribution over \mathbf{W}
3. Sample from the posterior and show a couple of functions
4. Repeat 2 – 3 by adding additional data points

Describe the plots, and the behavior when adding more data? Is this a desirable behavior?

2.2.2 Non-parametric Regression

In this task we will implement and evaluate the effect of a \mathcal{GP} -prior. In specific we will look at the squared exponential covariance function,

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 e^{-\frac{(\mathbf{x}_i - \mathbf{x}_j)^T (\mathbf{x}_i - \mathbf{x}_j)}{l^2}} \quad (17)$$

You will need to first formulate the prior distribution and then the posterior. How to do this can be found in Bishop 2006, page 306 – 308. First we will look at the prior,

Question 12:

1. Create a \mathcal{GP} -prior with a squared exponential co-variance function.
2. Sample from this prior and visualise the samples
3. Show samples using different length-scale for the squared exponential

Explain the behavior of altering the length-scale of the covariance function.

This has been said many times before but it is something that cannot be stressed enough, priors are very important as they allow us to formulate our uncertainty in our beliefs in a principled manner. However, more important is that we can combine our beliefs with observations, this is what facilitates learning. The object that contains this is the posterior distribution. We will now perform a simple experiment on the posterior. The posterior and the prior are the same object if we do not have any observed data.

Question 13: *Explain the above statement, why is this?*

Lets generate some data that we know would not work particularly well using a linear model as in the previous task,

$$y_i = \sin(x_i) + \epsilon_i \quad (18)$$

$$\mathbf{x} = [-\pi, \dots, \pi]^T \quad (19)$$

$$\epsilon_i \sim \mathcal{N}(0, 0.5), \quad (20)$$

where the cardinality of \mathbf{x} is 7, i.e. we have 7 data-points. Now we have some observations of a noisy sin-wave which we can use together with our prior to see the posterior distribution over the functions.

Question 14:

1. Compute the predictive posterior distribution of the model
2. Sample from this posterior with points both close to the data and far away from the observed data.
3. Plot the data, the predictive mean and the predictive variance of the posterior from the data

Explain the behavior of the samples and compare the samples of the posterior with the ones from the prior. Is this behavior desirable? What would happen if you would add a diagonal co-variance matrix to the squared exponential?

II The Posterior $p(\mathbf{X}|\mathbf{Y})$

In the previous task we looked at learning a relationship between two variates \mathbf{X} and \mathbf{Y} such that we could infer one from the other. One way of thinking about this is that given the mapping f and the input \mathbf{X} we specify the outputs \mathbf{Y} , you can think of \mathbf{X} as a “representation” of \mathbf{Y} , i.e. that the former have generated the later. Actually this is exactly what we did in the linear regression task, we generated some data using a set of parameters \mathbf{w} then we threw them away and later recovered them back, but we retained \mathbf{x} in this process. In this task we will make things a bit more complicated by looking representation learning. This means we will only observe the outputs \mathbf{Y} and want to learn input \mathbf{X} that can represent \mathbf{Y} . Why would we ever want to do this? Lets take the example of an image. Images are very high-dimensional objects, a typical HD image \mathbf{y}_i concatenated into a vector will live in a space of $\mathbb{R}^{1920 \times 1080 \times 3}$. However, does the image actually have that many degrees-of-freedom? To simplify, given that you know the place the image was taken, the weather, the exact camera angle, the objects in the image wouldn't you be able to generate the pixel data? This is of course a massive simplification but ponder how many parameters you can come up with and I recon it will be less than the number of pixels in the image. Lets call all these factors that we came up with and refer to them as *generating parameters* just as we said that \mathbf{X} through f generated \mathbf{Y} in Task I. What representation learning is ideally about is to recover these generating parameters directly from the data. More specifically this relates to building a model of the data \mathbf{Y} and then looking at the posterior distribution over the input to the model \mathbf{X} .

The other new thing that we will introduce in this task is *learning*. This means that we specify a model like in Task I and then *fit* this model to the data. What this means is that the model have a set of parameters which we now will infer from the data.

2.3 Theory

The models that we will look like are much the same as in the first task, a linear model and a non-parametric Gaussian Process. The difference now is that the input locations \mathbf{X} are not known a-prior but rather we want to infer them from data. We will refer to the input locations as the *latent representation* of the observed data \mathbf{Y} . Think about how this relates to the latent space models that you worked on in the first part of the course, where you used discrete latent *states* to represent continuous data. This is very much the same thing, except for that we want to find the latent space from data and that rather than being a discrete variable it is continuous.

Lets start with the linear model,

$$p(\mathbf{Y}, \mathbf{X}, \mathbf{W}) = p(\mathbf{Y}|\mathbf{X}, \mathbf{W})p(\mathbf{X})p(\mathbf{W}). \quad (21)$$

The next step will take a bit of thinking, being fully Bayesian we would like to invert the model above and look at the conditional distribution over the variables that we want to infer. Think about this, does it make sense? Actually, there is a simple relationship between \mathbf{X} and \mathbf{W} , if I have one I have the other right? As an example, if I multiply each \mathbf{x}_i with a constant that is the same as dividing the \mathbf{W} with the same constant. The way to get away from this is to only look at one variable. But as our model contains both \mathbf{X} and \mathbf{W} how can we do this? We can specify a prior over the variable that we are not interested in an marginalise it out from the model, right?

What does this actually mean, previously we have been using prior distributions as a mean of encoding our beliefs about a variable before seeing data. Another equally valid explanation is as encoding our *preference* of a variable.

Question 15: *Elaborate on this, why can one view a prior as encoding a preference?*

Lets specify the prior over the latent variables as a spherical gaussian,

$$p(\mathbf{X}) = \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (22)$$

Question 16: *What type of “preference” does this prior encode?*

Now we can combine this prior with the likelihood, integrate out \mathbf{X} and reach the marginal distribution,

$$p(\mathbf{Y}|\mathbf{W}) = \int p(\mathbf{Y}|\mathbf{X}, \mathbf{W})d\mathbf{X}. \quad (23)$$

Question 17: *Perform the marginalisation in Eq. 23 and write down the expression. As previously, I do recommend that you do this by hand but to pass the assignment you only need to outline the calculations and show the approach that you would take.*

2.3.1 Learning

So far we have only created models and looked at the posterior. Now we will take one step further and learn the parameters of the model. A good background to what we will go through here can be found in Bishop 2006, page 9,23,26,30,165 & Section 1.2.4-1.2.6. The most straight forward manner to do learning in a probabilistic model is called Maximum-Likelihood or ML. This does exactly what it says on the box, we formulate the likelihood of the data and then we find the parameters that maximises the likelihood,

$$\hat{\mathbf{W}} = \operatorname{argmax}_{\mathbf{W}} p(\mathbf{Y}|\mathbf{X}, \mathbf{W}). \quad (24)$$

The next level up is to perform maximum-a-posteriori or *MAP* estimation. This means that we find the parameters that maximises the posterior distribution,

$$\hat{\mathbf{W}} = \operatorname{argmax}_{\mathbf{W}} \frac{p(\mathbf{Y}|\mathbf{X}, \mathbf{W})p(\mathbf{W})}{\int p(\mathbf{Y}|\mathbf{X}, \mathbf{W})p(\mathbf{W})d\mathbf{W}} = \operatorname{argmax}_{\mathbf{W}} p(\mathbf{Y}|\mathbf{X}, \mathbf{W})p(\mathbf{W}). \quad (25)$$

There is also an in-between stage which is often referred to as Type-II Maximum-Likelihood which implies maximisation of the marginal likelihood where you integrate out one parameter and then maximise over another,

$$\hat{\mathbf{W}} = \operatorname{argmax}_{\mathbf{W}} \int p(\mathbf{Y}|\mathbf{X}, \mathbf{W})p(\mathbf{X})d\mathbf{X}. \quad (26)$$

Question 18: *Compare they three estimation procedures above. (I recommend comparing them in log-space)*

- *How are the different?*
- *How are MAP and ML different when we observe more data?*
- *Why is the two last expressions of Eq. 25 equal?*

In the representation task we have a model with two variables \mathbf{W} and \mathbf{X} that interact. This means that a Type-II ML estimation is a sensible approach to learn the model.

Practical Optimisation In practice when performing optimisation on probabilistic models we often have to deal with exponentials. Exponentials are nice in many ways, they are for example infinitely differentiable, but they are a bit tricky to play with. Often we have the case that our parameters are actually in the exponents and then we can do a neat trick that makes life much easier. Rather than working directly on the exponent we perform all our learning in the log-space instead. The

reason that we can do that is because $\log(\cdot)$ is a monotonic function and therefore it will not alter the location of the extremes of the function. Further, most optimisation packages are designed to minimise a function rather than maximising it. This means that in practice we often formulate our optimisation problem as a minimisation of the *negative log* of a probability,

$$\hat{\theta} = \operatorname{argmax}_{\theta} p(\mathbf{Y}|\theta) = \operatorname{argmin}_{\theta} -\log(p(\mathbf{Y}|\theta)). \quad (27)$$

Now we will write down the objective function and its gradients which means we need to do some matrix algebra. Two really good references that I find incredibly useful when working with matrices is Petersen and Pedersen [2006, URL](#) and Magnus and Neudecker [1988, URL](#).

Question 19:

1. Write down the objective function $-\log(p(\mathbf{Y}|\mathbf{W})) = \mathcal{L}(\mathbf{W})$.
2. Write down the gradients of the objective with respect to the parameters $\frac{\delta \mathcal{L}}{\delta \mathbf{W}}$

In the practical section of this task we will perform the optimisation above for a real data-set.

2.3.2 Non-parametric

We will now move on to the non-linear case using \mathcal{GP} s to model the generative mapping from the latent space to the observed data. The procedure will be much the same as in the linear case where we compute the marginal likelihood of the model and then fit the parameters.

In the linear case we integrated out the latent locations \mathbf{X} and performed Type-II Maximum Likelihood estimation of the parameters of the mapping \mathbf{W} . It is not as straight forward in the non-parametric case. But as discussed before, rather than marginalising over \mathbf{X} we can marginalise out the mapping f ,

$$p(\mathbf{Y}|\mathbf{X}, \boldsymbol{\theta}) = \int p(\mathbf{Y}|f)p(f|\mathbf{X}, \boldsymbol{\theta})df \quad (28)$$

Question 20: Explain why it is simpler to marginalise out f than \mathbf{X} ?

Suggestion: Draw the graphical model and use this to motivate your explanation.

2.4 Practical

Now we will get our hands dirty, which for a machine learning scientist means including the data. Lets generate some data so that we know what we are looking to recover,

$$\mathbf{Y} = f_{\text{lin}}(f_{\text{non-lin}}(\mathbf{x})) \quad (29)$$

$$\mathbf{x} = [0, \dots, 4\pi] \quad (30)$$

$$|\mathbf{x}| = 100 \quad (31)$$

$$f_{\text{non-lin}}(x_i) = [x_i \sin(x_i), x_i \cos(x_i)] \quad (32)$$

$$f_{\text{lin}}(x') = \mathbf{A}^T \mathbf{x}' \quad (33)$$

$$\mathbf{A} = \mathbb{R}^{10 \times 2} \quad (34)$$

$$\mathbf{A}_{ij} \sim \mathcal{N}(0, 1). \quad (35)$$

The values in the linear mapping are draws from an independent Gaussian as we do not really care about the specific form of the mapping we only care about its **rank**.

Now we have generated a data-set $\mathbf{Y} \in \mathbb{R}^{100 \times 10}$ which have been generated from a one dimensional *generating parameter* $\mathbf{x} \in \mathbb{R}^{1 \times N}$. The aim is now to given only \mathbf{Y} recover \mathbf{x} , i.e a single line. This

is a very general and incredibly important task in machine learning, how to discover the parameters that have generated some observations and appears in many applications such as computer vision and computational biology just to name a few. It is important as many types of data are represented in very high-dimensional spaces which are very hard to interpret, providing the true generating parameters allows us to analyse the data and hopefully find the casual behavior in the data.

2.4.1 Linear Representation Learning

We have an objective function and we have the gradients with respect to the parameters that we want to learn. The actual optimisation we can do using gradient descent. This is well implemented in `scipy.optimize`. Have a look at [URL](#) for the different methods that are available. Below is the simple structure that you need to implement in order to get the `fmin` function working,

```

1  import numpy as np
2  import scipy as sp
3  import scipy.optimize as opt
4
5  def f(x, *args):
6      # return the value of the objective at x
7      return val
8
9  def dfx(x,*args):
10     # return the gradient of the objective at x
11     return val
12
13  x_star = opt.fmin_cg(f,x0,fprime=dfx, args=args)

```

Question 21: *Plot the representation that you have learned. Explain why it looks the way it does. Was this the result that you expected? Hint: Plot \mathbf{X} as a two-dimensional representation.*

2.4.2 OPTIONAL: Non-parametric Representation Learning

This task is just added for completion, you do not need to do it and the eventual completion of it will not in any way or form reflect your grade. However it is a very interesting example that shows the strength of non-parametric priors and can therefore be nice to at least look through. Exactly the techniques described below was used to produce the first machine learning algorithm that surpassed human performance in face recognition Lu and Tang 2014. These results have been highlighted by several general media outlets [Science,URL](#). These methods are part of a class of methods referred to as *Gaussian Process Latent Variable Models* or GP-LVM for short that were proposed in Lawrence 2005. If you are interested in these types of algorithms there is a very useful Python packages called GPy Hensman *et al.* 2014.

To perform this task you will need to compute the gradients of the marginal likelihood of the GP-regression model from Task I. In log-space, for a single point \mathbf{y} this becomes,

$$\mathcal{L} = -\frac{N}{2}\log(2\pi) - \frac{1}{2}\log|\mathbf{K}| - \frac{1}{2}\text{tr}(\mathbf{y}^T\mathbf{K}^{-1}\mathbf{y}), \quad (36)$$

where $|\cdot|$ is the determinant and tr is the trace of a matrix. In order to optimise the above with respect to \mathbf{X} we need to compute gradients of \mathcal{L} with respect to the latent variables. The simplest way to do this is to compute the gradients with respect to the kernel first and then use the chain rule to get the gradients of \mathbf{X} . Now we can proceed much in the same way as in the linear case to find the latent representation. The objective function is non-convex, this together with a gradient based optimisation means that we will need to have a descent initialisation to our objective. One approach would be to initialise the latent locations using the linear approach that we just did. In the following lets do exactly that and use the PPCA solution to initialise the non-linear model.




Figure 1: The above figure (Figure 3.13 in Bishop 2006) shows the idea of model complexity and Occam’s razor that was introduced in Mackay 1991. Occam’s Razor tells us that we should always choose the “simplest” model that explains all of our data. In the figure the data domain is ordered on the x-axis by increasing complexity and the evidence $p(\mathcal{D})$ is shown on the y-axis. Given that we want to model the data \mathcal{D}_0 which model should we choose? Model \mathcal{M}_1 places no probability mass over \mathcal{D}_0 so it is a bad choice. Both model \mathcal{M}_2 and \mathcal{M}_3 places probability mass over \mathcal{D}_0 but as \mathcal{M}_2 places more, due to it being less complex only modeling a part of the data domain, and should therefore be preferred according to Occam’s Razor.

III The Evidence $p(\mathbf{Y})$

One of the main arguments behind Bayesian reasoning is that it automatically implements Occam’s Razor, i.e. that automatically chooses the “correct” model complexity to perform a specific task. In this part of the assignment we will perform a study which shows that sometimes things are not as obvious as one might think. We will use the evidence of the data under the model as a means of measuring the complexity of the model.

2.5 Theory

In the practical part of the task we will perform the experiments outlined in the excellent paper by Murray and Ghahramani titled *A note on the evidence and Bayesian Occam’s razor*. I recommend that you read this paper and familiarise yourself with their discussion surrounding complexity. Do not expect any clear answers in this but you should grasp the discussion and be able to argue about what the results shows.

2.5.1 Data

Consider a very simple data domain $\mathcal{D} = \{y^i\}_{i=1}^9$ where $y^i \in \{-1, 1\}$. This data is structured according to a grid whos locations can be parametrised by $\mathcal{X} = \{\mathbf{x}^i\}_{i=1}^9$ where $\mathbf{x}^i = (\{-1, 0, +1\}, \{-1, 0, +1\})$. This means that our data domain \mathcal{D} contains $2^9 = 512$ different elements which is small enough for us to reason about but still complicated enough that it requires a sensible model.

2.5.2 Models

Given the data defined above we wish to create a model, i.e. something that will explain the statistical variations that are possible in \mathcal{D} . The simplest model that (I) can think of is something that simply takes all its probability mass and places it uniformly over the whole data space,

$$p(\mathcal{D}|M_0, \boldsymbol{\theta}_0) = \frac{1}{512}. \quad (37)$$

Question 22: *Why is this the simplest model, and what does it actually imply? Discuss its implications, why is this a bad model and why is it a good model.*

The first model Eq. 37 does not take any parameters at all which means it has no flexibility and uses no information about \mathcal{D} except for its cardinality. We can use what we know about the data in order to specify something slightly more representative. If we assume that each y^i are independent we can factorise the model into 9 separate models,

$$p(\mathcal{D}|M_1, \boldsymbol{\theta}_1) = \prod_{n=1}^9 p(y^n|M_1, \boldsymbol{\theta}_1), \quad (38)$$

where θ_i^j means the j :th element of the parameter vector for the t :th model. Each model can be expressed using an exponential function which relates the value y^i to its location \mathbf{x}^i ,

$$p(\mathcal{D}|M_1, \boldsymbol{\theta}_1) = \prod_{n=1}^9 \frac{1}{1 + e^{-y^n \theta_1^1 x_1^n}}, \quad (39)$$

Question 23: *Explain how the each separate model works? In what way is this model more or less flexible compared to M_0 ? How does this model spread its probability mass over \mathcal{D} ?*

We can continue to add more parameters and create further models,

$$p(\mathcal{D}|M_2, \boldsymbol{\theta}_2) = \prod_{n=1}^9 \frac{1}{1 + e^{-y^n (\theta_2^1 x_1^n + \theta_2^2 x_2^n)}} \quad (40)$$

$$p(\mathcal{D}|M_3, \boldsymbol{\theta}_3) = \prod_{n=1}^9 \frac{1}{1 + e^{-y^n (\theta_3^1 x_1^n + \theta_3^2 x_2^n + \theta_3^3)}}, \quad (41)$$

Question 24: *How have the choices we made above restricted the distribution of the model? What data sets are each model suited to model? What does this actually imply in terms of uncertainty? In what way are the different models more flexible and in what way are they more restrictive? Discuss and compare the models to each other?*

2.5.3 Evidence

The evidence of a model M_i is the distribution $p(\mathcal{D}|M_i)$. This distribution tells us how and where the model spreads its probability mass. Occam's razor can be interpreted in terms of the evidence such as we should choose a model which places most of its mass where we will see data and as little as possible elsewhere. In the previous section we have defined a small simple data domain \mathcal{D} and we will now evaluate where the different models defined above places their probability mass.

In order to "reach" the evidence of a model we need to first remove the dependency of the variable $\boldsymbol{\theta}$. This can be done by marginalising out the parameters from the model,

$$p(\mathcal{D}|M_i) = \int_{\forall \boldsymbol{\theta}} p(\mathcal{D}|M_i, \boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}. \quad (42)$$

Question 25: *Explain the process of marginalisation. Discuss its implications.*

The marginalisation above requires one more object that we haven't seen before $p(\boldsymbol{\theta}|M_i)$. This is the *prior* over the parameters of the model. Being Bayesian implies that you need to take uncertainty into account in all steps of your calculations this is true for the data but also true for the parameters. As we do not really know much at all about the parameters we would like to be very uncertain and allow for a large range of possible values of $\boldsymbol{\theta}$. One prior would be to choose a simple Gaussian with zero mean and a very large variance,

$$\begin{aligned}
 p(\boldsymbol{\theta}|M_i) &= \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \\
 \boldsymbol{\mu} &= \mathbf{0} \\
 \boldsymbol{\Sigma} &= \sigma^2 \mathbf{I} \\
 \sigma^2 &= 10^3
 \end{aligned}
 \tag{43}$$

Question 26: *What does this choice of prior imply? How does the choice of the parameters of the prior $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ effect the model?*

Now when we have defined the prior $p(\boldsymbol{\theta})$ we just need to perform the marginalisation in Eq. 42 to be able to evaluate the evidence. However, this integration is rather tricky to do analytically which means that we will here use an approximate integral using a naïve Monte Carlo approach,

$$p(\mathcal{D}|M_i) \approx \frac{1}{S} \sum_{s=1}^S p(\mathcal{D}|M_i, \boldsymbol{\theta}^s),
 \tag{44}$$

$$\boldsymbol{\theta}^s \sim p(\boldsymbol{\theta}|M_i)
 \tag{45}$$

where s indexes the samples from the prior of the parameters.

We will now proceed to implement the procedure explained above and see what implications the different models and the choices of prior distributions have in terms of the evidence.

2.6 Practical

Even though the functionality that we need to do the calculations described above are simple there is, as always with high-level languages, lots of useful packages available in Python that will make our life much easier. When I wrote the code for this I found the following libraries very useful,

```
1 import itertools as it
2 from math import exp, sqrt, pi
3 import scipy.stats
```

Below is just a suggestion of how and in what order to implement the code to be able to answer the questions for this task. I will only look at your plots and not your code so feel free to

1. First create the code that generates the dataset \mathcal{D} and the locations of the data \mathbf{x} , `itertools` will be very useful here. Write some simple functionality to visualise a single element of the data on the 3×3 grid defined by \mathbf{x} .
2. Create the code to represent each model M_0 to M_3
3. Create the code to sample from the prior $p(\boldsymbol{\theta}|M_i)$
4. Write the code to perform the Monte Carlo integration given a model, returning the evidence
5. Write the code to index the data-sets such that you can easily compare the models. A good suggestion is the one provided in the appendix of the paper.

Question 27: *For each model sum the evidence for the whole of \mathcal{D} what numbers do you get? Explain these numbers for all the models and relate them to each other.*

Question 28: *Plot the evidence over the whole data set for each model. The x-axis index the different instances in \mathcal{D} and each models evidence is on the y-axis. How do you interpret this, relate this to the parametrisation of each model.*

Question 29: *Find using `np.argmax` and `np.argmin` which part of the \mathcal{D} that is given most and least probability mass by each model. Plot the data-sets which are given the highest and lowest evidence for each model. Discuss these results, does it make sense?*

Question 30: *What is the effect of the prior $p(\boldsymbol{\theta})$.*

- *What happens if we change its parameters?*
- *What happens if we use a non-diagonal covariance matrix for the prior?*
- *Alter the prior to have a non-zero mean, such that $\mu = [5, 5]^T$?*
- *Redo evidence plot for these and explain the changes compared to using zero-mean.*

Final thoughts If you have made it this far and managed to do all the experiments outlined in this task you have done very well. As a last bit I recommend you to go back and read the abstract again, did performing this lab answer the main question stated there?

Good Luck!

References

- Christopher M Bishop. *Pattern recognition and machine learning*. 2006. URL: <http://www.library.wisc.edu/selectedtoocs/bg0137.pdf>.
- KB Petersen and MS Pedersen. “The matrix cookbook”. In: *Technical University of Denmark* (2006).
- Jan R Magnus and Heinz Neudecker. *Matrix differential calculus with applications in statistics and econometrics*. Wiley, 1988. URL: <http://www.narcis.nl/publication/RecordID/oai:wo.uvt.nl:153200>.
- Chaochao Lu and Xiaou Tang. “Surpassing Human-Level Face Verification Performance on LFW with GaussianFace”. In: *arXiv.org* (Apr. 2014). arXiv: [1404.3840v2](https://arxiv.org/abs/1404.3840v2) [cs.CV]. URL: <http://arxiv.org/abs/1404.3840v2>.
- Neil D Lawrence. “Probabilistic non-linear principal component analysis with Gaussian process latent variable models”. In: *The Journal of Machine Learning Research* 6 (2005), pp. 1783–1816. URL: <http://dl.acm.org/citation.cfm?id=1194904>.
- James Hensman *et al.* “GPY: A Gaussian process framework in python”. In: (2014). URL: <https://github.com/SheffieldML/GPY>.
- David J C Mackay. “Bayesian methods for adaptive models”. PhD thesis. California Institute of Technology, Dec. 1991. URL: <http://codatest2.library.caltech.edu/25/>.
- Iain Murray and Zoubin Ghahramani. *A note on the evidence and Bayesian Occam’s razor*. Tech. rep. GCNU-TR 2005-003. Aug. 2005. URL: <http://eprints.pascal-network.org/archive/00001165/>.