

Selected problems in Machine Learning: Warm-up test, October 4th, 2017

- Plot the following functions: (a) $f(x) = x \exp(x)$, (b) $f(x) = \ln \frac{x^2-1}{x^2+1}$, (c) $f(x) = \left| \frac{x-1}{1-2x} \right|$, (d) $f(x) = \sqrt{1 - \exp(-x^2)}$
- Define convex function.
- Define convex region in R^2 .
- What is gradient?
- Find a growing function which maps R to $(0, 1)$.
- Given joint distribution $p(A, B)$, how we call $p(A)$ and $p(A|B)$?
- Explain “curse of dimensionality”.
- Explain the main difference between the frequentist and Bayesian interpretation of probability.
- Derive Bayes’ theorem.
- Define independence (independent random variables X and Y).
- Define conditional independence (variable X independent of Y given Z).
- What does it mean that a collection of random variables (e.g. a sequence) is i.i.d. (independent and identically distributed)?
- What is the relation between a probability density function and associated cumulative distribution function?
- Explain the difference between the terms probability and likelihood.
- Why we use probabilistic methods?
- Where does uncertainty in NLP tasks come from?
- What is correlation?
- What is variance?
- What is covariance matrix?
- Let’s suppose that the sequence (1, 3, 4, 4, 8) is drawn from $\mathcal{N}(\mu, \sigma^2)$. What are the values of μ and σ^2 (according to Maximum Likelihood)?
- What can you say about a multidimensional Gaussian with covariance matrix having zeros everywhere outside its diagonal?
- Why is Gaussian distribution so special?
- Which types of random variables cannot be modeled by Gaussian distributions?
- Could you sketch a histogram for length distribution of dependency relations? (distinguish orientation)
- Could you sketch a histogram for length distribution of anaphoric relations? (distinguish orientation)
- What are Monte Carlo methods used for?
- How can you generate samples from a uniform distribution in 3D unit ball?
- How can you generate samples from 1D Gaussian distribution?
- If you have a generator of numbers from the uniform distribution on $[0, 1]$, how would you generate samples with probability density $p(X)$, $x \in R$.
- Suppose you have a sampler of x-y pairs from a joint distribution $p(x, y)$. How can you generate samples from the associated marginal distribution $p(x)$?
- Explain how Gibbs sampling works.
- For which distributions you can *not* use Gibbs sampler?
- What is entropy?
- Which distribution on discrete values has the lowest entropy?
- Which distribution on continuous values has the lowest entropy?

36. What is mutual information?
37. What is KL divergence? Why it is not a distance (in the mathematical sense)?
38. Explain the difference between generative models and discriminative models.
39. Explain the difference between classification and regression.
40. What is separation boundary (in classifiers)?
41. What does it mean that two sets of points are linearly separable.
42. Explain how entropy maximization can be used in classification.
43. Explain how entropy minimization can be used in classification.
44. Explain how Naive Bayes classifier works.
45. Plot the sigmoid function used in logistic regression ($f(z) = \frac{1}{1+e^{-z}}$).
46. Design feature transformation functions which make the following sets linearly separable: (a) $A = (2, 0)$ and $B = (1, 0), (3, 0), (2, 1), (2, -1)$, (b) $A = (0, 0), (2, 2)$ and $B = (1, 1), (3, 3)$.
47. Explain how k-means clustering works.
48. Explain how EM works.
49. What is cost function (loss function)?
50. What is bias-variance trade-off.
51. Illustrate underfit/overfit problems in regression by fitting a polynomial function to a sequence of points in 2D.
52. Illustrate underfit/overfit problems in classification by modeling two (partially overlapping) classes of points in 2D.
53. Plot a typical dependence of training and test error rates (vertical axis) on training data size (horizontal axis).
54. Plot a typical dependence of training and test error rates (vertical axis) on model complexity (horizontal axis).
55. What problem is typically signalled by test error being much higher than training error?
56. What problem is typically signalled by test and training errors being stabilized after a limited portion of training data?
57. What is regularization used for?
58. Explain the difference between parametric and non-parametric methods.
59. Name some quantities that can be used for feature selection.