

Machine Learning for Computer Vision

January 23, 2019

Topic: Variational Inference

Exercise 1: Kullback-Leibler divergence

- a) What does the KL divergence describe? Is it a metric? Why?
- b) Compute the KL-divergence of two univariate normal distributions.
What if they have the same mean? What if they have the same variance?
- c) Consider a factorized variational distribution $q(Z)$. By using the technique of Lagrange multipliers, verify that minimization of $KL(p||q)$ with respect to one of the factors $q_i(Z_i)$ keeping all other factors fixed, leads to the solution:

$$q_j^*(Z_j) = \int p(Z) \prod_{i \neq j} dZ_i = p(Z_j)$$