Machine Learning	Prof. Dr. Bernt Schiele
	Dr. Paul Swoboda & Dr. Gerard Pons-Moll
Exercise Sheet 1 - 22.10.2018	due: 29.10.2018, 14:15

Exercise 1 - Correlation, Independence and all the rest

a. **(1 Point)** Let X be a $\mathbb{R}$-valued random variable with symmetric density function ($p(x) = p(-x)$, $\forall x \in \mathbb{R}$) and define $Y = X^2$. Compute the correlation of X and Y. What does this result tell you about the dependence of X and Y ? Are they independent ?

b. **(2 Points)** Suppose you flip a biased coin which shows head with probability $0 \leq p \leq 1$. You predict head with probability $0 \leq q \leq 1$. What is the probability that your prediction is correct under the assumption that the coin flip and your prediction are independent ? How has q to be chosen in order to maximize this probability (for fixed p) ?

Exercise 2 - Transformation law

a. **(3 Points)** Let U_1, U_2 be two independent real-valued random variables with uniform distribution on [0, 1]. Compute the density of the variables $X = g(U)$, where

$$X_1 = \sqrt{-2\ln U_2}\sin(2\pi U_1),$$
$$X_2 = \sqrt{-2\ln U_2}\cos(2\pi U_1).$$

The result is $p(x_1, x_2) = \frac{1}{2\pi}e^{-\frac{x_1^2+x_2^2}{2}}$ (a centered Gaussian in $\mathbb{R}^2$ with covariance matrix $C = \mathbb{1}$, where $\mathbb{1}$ is the identity matrix in $\mathbb{R}^2$). This is the so-called **Box-Muller**-transform which is used to generate samples of a centered Gaussian with unit variance (this is what you get if you use the command `randn` in Matlab).

Hints:

a. Derive the inverse function g^{-1}, that is derive U_1, U_2 as a function of X_1, X_2 (you have to invert the given nonlinear system of equations). Here, $\sin^2\alpha + \cos^2\alpha = 1$ is quite helpful.

b. Then apply the transformation law from the lecture. First, derive the Jacobian of g^{-1} and then compute its determinant.

Exercise 3 - Phenomena in high dimensions and sampling

This exercise shows that the geometric intuition we have from living in three dimensions can usually not be transferred to higher dimensional spaces. However, often one encounters learning problems where one has a lot of features und thus one is working in a high-dimensional space. Therefore it is important to get some idea what can go wrong in high dimensions.

The volume $\mathrm{vol}(B_d(r))$ of the d-dimensional ball $B_d(r)$ ($B_d(r) = \{x \in \mathbb{R}^d \mid \|x\|_2 \leq r\}$) of radius r in $\mathbb{R}^d$ is given as

$$\mathrm{vol}(B_d(r)) = \frac{\pi^{\frac{d}{2}} r^d}{\Gamma(\frac{d}{2}+1)},$$

where Γ is the Gamma function and one has

$$\Gamma\left(\frac{d}{2}+1\right) = \begin{cases} \left(\frac{d}{2}\right)! & \text{if } d \text{ even,} \\ \sqrt{\pi}\frac{d!!}{2^{\frac{d+1}{2}}} & \text{if } d \text{ is odd} \end{cases}.$$

Note, that $d!!$ is the **double factorial** defined as $d!! = d(d-2)\ldots 1$ where $1!! = 0!! = 1$.

a. **(1 Point)** Derive the limit of the volume of $B_d(1)$ as $d \to \infty$.

b. **(3 Points)** Suppose you want to sample from the uniform measure on $B_d(1)$ (not just the surface - the whole volume).

   - What is the density of the uniform probability measure on $B_d(1)$?
   - Somebody suggests *rejection sampling* to get samples from the uniform measure on $B_d(1)$. First one draws a point X uniformly from the d-dimensional cube $[-1, 1]^d$ ($\mathtt{X} = 2 * \mathtt{rand}(\mathtt{d}, 1) - 1$ in Matlab). Then one selects X as a sample if $\|X\|_2 \leq 1$ otherwise one discards it. What is the probability that a sample is discarded ? How many samples do you have to draw on average in the first step in order to get 1000 samples from the d-dimensional unit ball $B_d(1)$ for $d = 3$ and $d = 20$?

c. **(2 Points)** Clearly, the method based on rejection sampling does not work in high dimensions. Describe a more efficient way of sampling from the d-dimensional unit ball using Gaussian samples where no sample is discarded ? Prove that the method you suggest generates samples from the uniform distribution on $B_d(1)$.

Hints:

a. The probability in b) should be calculated and not be determined by numerical experiments (for $d = 20$ it will anyway not work :)).

b. There might be several ways in part c). The hint points in one particular direction.

   - In spherical coordinates $(r, \theta_1, \ldots, \theta_{d-1})$ on $\mathbb{R}^d$ one has for every $S \subset \mathbb{R}^d$,

   $$\int_S f(x_1, \ldots, x_d)\, dx_1 \ \ldots \ dx_d = \int_S f(r, \theta_1, \ldots, \theta_d)\, r^{d-1}\, dr\, d\Omega,$$

   where $d\Omega$ is the surface element on the sphere (up to a factor this is the uniform measure on the surface of $B_d(1)$). You may use that the surface S_{d-1} of the unit sphere in $\mathbb{R}^d$ has $(d-1)$-dimensional volume $\mathrm{vol}(S_{d-1}) = d\,\mathrm{vol}(B_d(1))$.
   - This suggests a two-step procedure. First one has to sample a point uniformly on the surface of $B_d(1)$. This will yield the directional part and second one has to modify the norm of this point in a way that it behaves as the radial part of the uniform measure on $B_d(1)$.

Submission instructions

- We accept both handwritten and electronic submissions. So you can choose what is more convenient for you. In any case, you should specify full names and immatriculation IDs of all team members. Obviously, programming tasks you can submit only electronically.
- Handwritten submissions should be submitted in the lecture hall of Monday's lecture (before the lecture starts).
- Electronic submissions should be emailed to the corresponding tutor:

  a. Apratim Bhattacharyya (Wednesday 8-10): abhattac@mpi-inf.mpg.de
  b. Maksym Andriushchenko (Thursday 8-10): s8mmandr@stud.uni-saarland.de
  c. Max Losch (Friday 16-18): mlosch@mpi-inf.mpg.de

  If not all 3 students belong to the same tutorial group, then you should email your submission to **only** one tutor (e.g. to the tutor of the first author of your homework), so please do not put other tutors in copy of the email.

  The email subject must have the following form: "[ML18/19 Exercise] Sheet X", where X is the number of the current exercise sheet. Then please specify in the email full names and immatriculation IDs of all team members. Then please attach all your files as a single zip archive, which consists of your immatriculation IDs, e.g. "2561234_2561235_2561236.zip".

- Reminder: you should submit in groups of 3. Otherwise, we will later on merge the groups smaller than 3 students.