CIS 620 — Advanced Topics in AI
Profs. M. Kearns and L. Saul
Problem Set 2
Distributed: Monday, January 28, 2002
Due: Wednesday, February 6, 2002 (start of class)

1. *Bernoulli distribution*

   (a) *Left-sided bound on large deviations*
   Consider N i.i.d. Bernoulli random variables x_i ($i = 1 \ldots N$) with mean μ. Let $\mu' = \mu - \varepsilon$, where $\varepsilon > 0$ and $\mu' > 0$. Show that:

   $$\Pr\left[\frac{1}{N}\sum_i x_i \;\leq\; \mu'\right] \;\leq\; e^{-N d_{\mathrm{KL}}(\mu',\mu)},$$

   where $d_{\mathrm{KL}}(\mu', \mu)$ is the KL distance

   $$d_{\mathrm{KL}}(\mu',\mu) \;=\; \mu' \log\left(\frac{\mu'}{\mu}\right) + (1-\mu')\log\left(\frac{1-\mu'}{1-\mu}\right).$$

   (b) *KL distance*
   Let μ and μ' denote the means of Bernoulli random variables. Show that

   $$\frac{\partial^2}{\partial\mu^2}\left[d_{\mathrm{KL}}(\mu,\mu')\right] \;\geq\; 4 \quad \text{for all } \mu.$$

   Use this inequality to derive the lower bound:

   $$d_{\mathrm{KL}}(\mu,\mu') \;\geq\; 2(\mu-\mu')^2.$$

   (c) *Hoeffding bound*
   Consider N i.i.d. Bernoulli random variables x_i ($i = 1 \ldots N$) with mean μ. Assuming the results in parts (a) and (b), derive the simplified bound:

   $$\Pr\left[\frac{1}{N}\sum_i x_i \;\leq\; \mu - \varepsilon\right] \;\leq\; e^{-2N\varepsilon^2}$$

2. *Gaussian distribution*

   (a) *Generating function*
   Compute the generating function $\mathrm{E}[e^{kx}]$ for a Gaussian random variable with mean μ and variance σ^2:

   $$\mathrm{E}[e^{kx}] \;=\; \int_{-\infty}^{\infty} dx\, p(x) e^{kx} \quad \text{where} \quad p(x) \;=\; \frac{1}{\sqrt{2\pi\sigma^2}}\, e^{-\frac{(x-\mu)^2}{2\sigma^2}}.$$

   You may assume without proof that the distribution is properly normalized: $\int_{-\infty}^{\infty} dx\, p(x) = 1$.

   (b) *KL distance*
   Evaluate the KL distance

   $$\mathrm{KL}(p_1,p_2) \;=\; \int dx\, p_1(x) \log\left[\frac{p_1(x)}{p_2(x)}\right]$$

   between two Gaussian distributions $p_1(x)$ and $p_2(x)$ with means μ_1 and μ_2 and variances σ_1^2 and σ_2^2.

(c) *Large deviations*
Consider N i.i.d. Gaussian random variables x_i ($i = 1 \ldots N$) with mean μ and variance σ^2. Show that:

$$\Pr\left[\frac{1}{N}\sum_i x_i \;\geq\; \mu+\varepsilon\right] \;\leq\; e^{-N\varepsilon^2/(2\sigma^2)}.$$

3. *Heavy-tailed distribution* (extra credit)
The *Cauchy* distribution with mean zero and *width* α is given by:

$$p(x) \;=\; \frac{\alpha}{\pi}\left(\frac{1}{x^2+\alpha^2}\right).$$

(a) *Width and tails*
Show that $\Pr[|x| \leq \alpha] = \frac{1}{2}$ and that $\mathrm{E}[x^2] = \infty$.

(b) *Stability*
The sum of N i.i.d. Cauchy random variables with mean zero and width α is itself Cauchy distributed with mean zero and width $N\alpha$. (You are not asked to prove this.) Clearly, this process does not converge to a Gaussian distribution as $N \rightarrow \infty$. What assumption of the Central Limit Theorem is violated in this case?

4. *MATLAB by example*
Type these commands into MATLAB and use the *help* facility to understand the syntax. You will need to program in MATLAB for later problem sets.

```
% GAUSSIAN DISTRIBUTION
x = [-4:0.01:4];
figure(1); clf;
subplot(2,1,1); plot(x,exp(-x.*x/2)/sqrt(2*pi));
subplot(2,1,2); hist(randn(10000,1),32);

% KL DISTANCE FOR BERNOULLI
u = [0.001:0.001:0.999];
v = 0.5;
kl = u.*log(u./v) + (1-u).*log((1-u)./(1-v));
figure(3); clf;
plot(u,kl,'b-',u,2*(u-v).^2,'g-');
set(gca,'FontSize',18);
legend('KL distance','lower bound');
```

5. *Lower bound on planning from a generative model.*
Let A be any algorithm that uses a generative model for an MDP M as a subroutine, takes an arbitrary state $\vec{x}$ and an arbitrarily small value $\epsilon > 0$ as inputs, and outputs an action $a = A(\vec{x})$. (Note that the output of A may be stochastic due to sampling from the generative model.) Let the policy determined by A for any fixed $\epsilon > 0$ satisfy

$$V^*(\vec{x}) - V^A(\vec{x}) \leq \epsilon$$

simultaneously for all $\vec{x}$. Thus, the policy computed by A is near-optimal. (The sparse sampling algorithm described in class is an example of such an algorithm.) Construct an MDP M that gives the strongest lower bound you can on the number of calls A must make to the generative model as a function of the ϵ-horizon time $H_\epsilon = (1/(1-\gamma))\log(1/((1-\gamma)\epsilon))$.