Information Theory, Pattern Recognition and Neural Networks

Reading associated with the current lectures – Chapters 3, 21 (especially sec 21.2), and 22 (especially sec 22.1), and 27.

Homework recommendations – Examples 22.1-4 (p. 300) and exercise 22.8.

Ex 3.10 (p57) (children); 8.10, black and white cards; **9.19** TWOS; **9.20**, birthday problem; 15.5, 15.6, (233) magic trick; 8.3 (140), 8.7; **22.11** sailor.

Ex 22.5.

Supervision 5. Below, please find the main assignment for supervision number 5. Overleaf is another activity that we'll discuss in supervision 5 or 6. Please look at both of these before supervision 5.

An old exam question

A spy would like you to write a computer program that recognises, given a small number of consecutive characters from the middle of a computer file, whether the file is an English-language document. Assuming that the two alternative hypotheses are that the file is an English-language document ($\mathcal{H}_E$), or that it is a random string of characters drawn from the same alphabet ($\mathcal{H}_R$), describe how you would solve this problem.

Estimate how many characters your method would need in order to work reasonably well.

Further questions

Maybe you would enjoy writing a program that implements your method?

How would your answers differ if instead the task were to distinguish

(a) English from German?

(b) English from Hsilgne (backwards English)?

[When I say German, let's assume German with no accent characters.]

(Some facts about English and German are supplied below.)

Letter frequencies of English and German

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
English (**e**)	.07	.01	.03	.03	.10	.02	.02	.05	.06	.001	.006	.03	.02	.06
German (**g**)	.06	.02	.03	.04	.15	.01	.03	.04	.07	.002	.01	.03	.02	.08

	O	P	Q	R	S	T	U	V	W	X	Y	Z	–
English (**e**)	.06	.02	.0009	.05	.05	.08	.02	.008	.02	.002	.01	.0008	.17
German (**g**)	.02	.007	.0002	.06	.06	.05	.04	.006	.02	.0003	.0003	.01	.14

The entropies of these two distributions are $H(\mathbf{e}) = 4.1$ bits; $H(\mathbf{g}) = 4.1$ bits; and the relative entropies between them are $D_{\mathrm{KL}}(\mathbf{e}||\mathbf{g}) = 0.16$ bits and $D_{\mathrm{KL}}(\mathbf{g}||\mathbf{e}) = 0.12$ bits. The relative entropies between the uniform distribution $\mathbf{u}$ and the English distribution $\mathbf{e}$ are $D_{\mathrm{KL}}(\mathbf{e}||\mathbf{u}) \simeq 0.6$ bits and $D_{\mathrm{KL}}(\mathbf{u}||\mathbf{e}) \simeq 1$ bits.

How well calibrated are your estimates of uncertainty?

According to BBC News, the Earth was almost put on impact alert by some astronomers who, on 13 January 2004, reckoned a 30m object, later designated 2004 AS1, had a one-in-four chance of hitting the planet within 36 hours. 2004 AS1 in fact passed the Earth at a distance of about 12 million km – roughly 1000 Earth–diameters – posing no danger to us whatsoever. Clearly the correct assessment of the probability of earth being hit should have been more like one in 10^6.

This sheet is intended to give insight into the accuracy of our own probability estimates.

Give a 94% confidence interval for the following quantities. Give the tightest* interval you can, while remaining 94% sure that the true value is in the interval. Don't look up answers before you have written down your interval – the aim of this exercise is to get a feel for how well calibrated your intervals are.

	Quantity	Lower bound	Guess	Upper bound	Ratio	Score
1	Mass of the textbook (g)					
2	Population of Britain (census, 2001)					
3	Population of Turkey (July 2004)					
4	Population of Luxembourg (July 2004)					
5	Number of British MEPs					
6	Starting pay of University Lecturer (Aug 2004)					
7	Parliamentary salary of MP (1/4/2005)					
8	Council tax, South Cambs. (£/house/yr) (band D, 2005-6)					
9	Fraction of central government expenditure that goes to 'Defence' (2004)					
10	UK prison population (as fraction of whole) (March 2005)					
11	Number of USA nuclear warheads (Feb 2003)					
12	Distance to sun (miles) (on 10 March 2005)					
13	Mean radius of earth (km)					
14	Speed of light ($\mathrm{m\,s^{-1}}$)					
15	Density of Gold ($\mathrm{g\,cm^{-3}}$)					
16	The ratio $\dfrac{\text{Density of Uranium}^{238}}{\text{Density of Gold}}$					

$\star$ My definition of 'tightest' is that an interval (l, u) gets tighter if the ratio u/l is reduced. If you want, you can score how tight your intervals are using $\log_2(\log_2 u/l)$ – the smaller, the better.

Ratio	$\log_2(\log_2 \text{Ratio})$
1.01	-6
1.02	-5
1.04	-4
1.09	-3
1.19	-2
1.41	-1
2	0
4	1
16	2
250	3
100,000	4