

**Exercise 2.1 Entropy as a measure of uncertainty**

Let me try to introduce you to entropy measures in an intuitive way. For a precise and neat formulation you should read pages 13-17 of the script.

The goal here is to find a way of quantifying the ignorance one has about a given phenomenon. The first step is to express our knowledge with probability distributions. As usual we start with a rather ridiculous example: my father forgot where he left his glasses again, and gave my brother, my sister and me ten minutes to try to find them. Our house has 6 rooms: kitchen, parents bedroom, children bedroom, dinning room, living room and bathroom. Each of us searched for a while and tried to figure where the glasses could be. Our knowledge after ten minutes is represented in Fig. 1.



Figure 1: Our knowledge on the whereabouts of the glasses: I (first graph) searched in three of the rooms and found nothing so I think it is either in my dad’s room (the most likely option, given our previous experience), in the kitchen or in the bathroom. My brother (second graph) also searched in the bathroom, narrowing his options to parents room and kitchen. My sister (third graph) is almost in the same conditions except that she did not search very thoroughly so assumes they may have escaped her in one of the other rooms. The cat (fourth graph) has no idea.

These probability distributions have some things in common (like the height of the peak probability for the three siblings, or the size of the support, for my sister and the cat), but also some differences — it is clear that we have different degrees of certainty about where the glasses may be.

So, looking at those probability distributions, how can we quantify our knowledge, or lack of it? The answer is: it depends. It depends on how you want to use your knowledge.

If your father was feeling lucky and just asked us what was the one room where he was more likely to find the glasses, i.e. if we had to guess where the glasses were in a single shot, than the relevant quantity would be the peak probability, that gives us the probability of guessing correctly. In this case one could use the min-entropy to quantify our ignorance,

$$H_{\min}(X)_P = -\log \max_{x \in \mathcal{X}} P_X(x). \tag{1}$$

In this case my guess is as good as my the one by my sister — even though our state of knowledge is clearly different. This is how the min-entropy sees these probability distributions:



Figure 2: The min-entropy only cares about the peak probability, that expresses the probability of taking a correct guess in a single shot approach, ignoring how the other the probabilities are spread.

On the other hand, if our father wanted to be absolutely sure he would find the glasses and asked us to give him a list of the rooms where the glasses could still be, even if unlikely cases, the peak probabilities would be irrelevant and the only thing to matter would be the support of the distribution.



Figure 3: The max-entropy cares about the support of the distribution and is indifferent to the way the probabilities are spread. This makes it see similar distributions as if they were very different (for instance graphs 2 and 3), while two distributions that are quite distinct (like 3 and 4) may have the same max-entropy.

The max-entropy gives us the size (in bits) of the memory that is required to store all the possibilities, so it would be suited to quantify our ignorance in this case.

$$H_{\max}(X)_P = \log |P_X|. \quad (2)$$

You may think that it is somehow *unfair* that my sister, who knows much more than the cat about where the glasses may be, is considered as ignorant as it according to this entropy measure. In particular, you may find it a bit of a waste of time (or, in any another example, memory) that she sends our father to look for the glasses in the four rooms where she is *almost* certain they will not be. If she takes a very small chance of being wrong — a small error tolerance — she may dismiss the very unlikely event that the glasses are in one of those four rooms and tell the father to search only in the first two rooms.

Fortunately *someone* in the information theory community thought about this question before my sister and introduced the smooth max-entropy,

$$H_{\max}^\epsilon(X)_P = \min_{Q_X \in \mathcal{B}^\epsilon(P_X)} H_{\max}(X)_Q, \quad (3)$$

where the minimum is taken over all probability distributions  $Q_X$  that are  $\epsilon$ -close to  $P_X$  according to the trace distance. In practice the smooth max-entropy takes the tail of the distribution and wipes unlikely events until it reaches a maximum weight of  $\epsilon$ .



Figure 4: For the smooth max-entropy these two distributions look the same if the total probability of the four unlikely events in the tail of the second distribution sum up to less than the error tolerance  $\epsilon$ .

There is also a smooth min-entropy,

$$H_{\min}^\epsilon(X)_P = \max_{Q_X \in \mathcal{B}^\epsilon(P_X)} H_{\min}(X)_Q, \quad (4)$$

which does not have an immediate one-shot meaning. We will see next week that in the iid limit the smooth min- and max-entropies converge.

Small remark before we continue: one question you may ask is “why are all these entropies logarithmic?”. The short answer is that we want them to be additive, i.e. we want the entropy of two uncorrelated events to be the sum of the individual entropies, or in other words “what we do not know about two things that have nothing to do with each other is just the sum of our ignorance on each of the things” (what I do not know about oranges and war tanks is my ignorance about oranges plus... I am sure you got it by now). More in page 13 of the script.

My intuition for the Shannon entropy is not as clear as for the previous two (I am hoping these two were clear!). Let me try anyway. Sticking to the example, if our father had lost the glasses many times and we had reached the same probability distributions all those times, then the Shannon entropy would tell us how *surprised*, in average, we would be with the actual whereabouts of the glasses.

It makes sense that the smaller the probability of an event, the more surprised we would be if that event happened. Since we everything is logarithmic here, we define the surprise content, or *surprisal* of an event  $E$  as  $h(E)_P = -\log P_X(E)$  (again, check page 13 of the script for more convincing arguments). The Shannon entropy is just the expectation value of the surprisal, i.e. its average over all possible outcomes,

$$H(X)_P = - \sum_{x \in \mathcal{X}} P_X(x) \log P_X(x). \quad (5)$$



Figure 5: All probability distributions are different for the Shannon entropy. This may seem good but has a drawback: it has no operational meaning for a single shot experiment. In particular, very similar distributions, like those of graphs 2 and 3, may have very different Shannon entropies.

More remarks: In general  $H_{\min}(X)_P \leq H(X)_P \leq H_{\max}(X)_P$ . In the worst case scenario, i.e. the uniform distribution, they are all the same. In the i.i.d. limit their smooth versions are all the same too (but we will cover that later).

## Exercise 2.2 Mutual Information

Two new things in this exercise: conditional entropies and mutual information.

Conditional entropy quantifies our ignorance about something given our knowledge about a (hopefully) related event — for instance our uncertainty about the weather tomorrow after listening to the radio forecast.

The Shannon conditional entropy of  $X$  given  $Y$  is defined as the expectation value of the surprisal of  $x$  knowing  $Y = y$ ,

$$\begin{aligned} H(X|Y) &= \langle h(x|Y=y)_{P_{xy}} \rangle_{xy} \\ &= \langle -\log P_{X|Y=y}(x) \rangle_{xy} \\ &= -\sum_{x,y} P_{XY}(x,y) \log P_{X|Y=y}(x). \end{aligned} \quad (6)$$

As we have  $P_{X|Y=y}(x) = P_{XY}(x,y)/P_Y(y)$ , so comes

$$H(X|Y) = H(XY) - H(Y). \quad (7)$$

Conditional min- and max-entropies are given on page 16 of the script.

The mutual information tells us how correlated two experiments (read random variables) are. If they are maximally correlated (like a very accurate forecast and the actual weather) then you can determine each one of them from the other. If they are uncorrelated (like solar flares and the price of gold) then knowing one of them does not help you at all to guess the other.

The mutual information between  $X$  and  $Y$  is defined in a natural way as “what we have learned about  $X$  by knowing  $Y$ ”, or “what we know about  $X$  now that we know  $Y$  minus what we knew about  $X$  before knowing  $Y$ ”, or, to make it even more bizarre, “what we *did not know* about  $X$  before minus what we *do not know* about  $X$  now that we know  $Y$ ”, i.e. the entropy of  $X$  minus the conditional entropy of  $X$  given  $Y$ ,

$$I(X : Y) = H(X) - H(X|Y). \quad (8)$$

Notice that the mutual information is symmetric,  $I(X : Y) = H(X) + H(Y) - H(XY) = I(Y : X)$ .

Back to the exercise, you have to apply this concept to calculate the mutual information between your guess and the actual weather, and also between the guess of your grandfather and the weather.

The conditional and marginal probabilities for the radio forecast case are represented in the figure below. Check solution sheet 1 for details. Remember that in the case of a sunny forecast any strategy was equally bad — so for simplicity you may assume you trust the radio report and say it will not rain.




Figure 6: The radio forecast and the actual weather: marginal and conditional probabilities. Naturally, you can get the joint probabilities using  $P_{X|Y=y}(x) = P_{XY}(x,y)/P_Y(y)$ .