

Quantitative Risk Management

Case for Week 2

John A. Dodson

September 14, 2016

1 Primer: Entropy

The entropy of a (univariate) random variable is a measure of how much information you know about the values it could take. This may seem like an absurd notion, but it turns out to be very important in the setting of fitting models to data. In fact, one might describe the goal of fitting a model to data as finding the version of the model that minimizes the average entropy of each observation.

$$H_X \triangleq -\mathbb{E} \log f_X$$

The entropy of a non-random (degenerate) variable is zero. The entropy of a fair coin toss is $\log 2 \approx 0.7$ “nats” (or 1 “bit” using the base-two logarithm). The entropy of a random “byte” (a discrete random variable which can take on any of 256 equally-probable values) is 8 bits or about 5.5 nats.

You might intuit that, for a non-discrete random variable, the value of the location parameter probably does not matter to the entropy, but the value of the scale parameter probably does. After all, the notion of variance is evocative of uncertainty.

Let’s explore this intuition. For $X \sim \mathcal{N}(\mu, \sigma^2)$, we have

$$-\log f_X(X) = \log \sqrt{2\pi\sigma^2} + \frac{(X - \mu)^2}{2\sigma^2}$$

so

$$H_X = \log \sqrt{2\pi e\sigma^2}$$

which is an increasing function in the variance parameter σ^2 . The entropy of a standard normal is $\log \sqrt{2\pi e} \approx 1.4$ nats or just over two bits.

What about a more general case? Say you have a univariate random variable X with a known entropy H_X , and you are interested in the entropy of a new random variable $Y = aX + b$ for some fixed numbers a and b . We call this an “affine transformation”, and it encompasses any arbitrary shift or change of units of the mean or variance (if these quantities exist).

Since the probability density of Y is

$$f_{aX+b}(y) = \frac{1}{|a|} f_X\left(\frac{y-b}{a}\right)$$

in terms of the probability density of X ,

$$H_Y = - \int \log\left(\frac{1}{|a|} f_X\left(\frac{y-b}{a}\right)\right) \frac{1}{|a|} f_X\left(\frac{y-b}{a}\right) dy$$

Changing variables and expanding the logarithm,

$$\begin{aligned} &= - \int (\log f_X(x) - \log |a|) f_X(x) dx \\ &= H_X + \log |a| \end{aligned}$$

so we see that indeed entropy is unchanged under a location transformation and increasing under a scale transformation.

It might be useful to think of entropy as a generalization of variance.

2 Exercise: Bivariate Normal

The bivariate normal has five parameters: two means, two standard deviations, and a correlation. The standard version has only one parameter, $-1 \leq \rho \leq 1$. Its density is

$$f_{(X,Y)}(x,y) = \frac{1}{2\pi} e^{-\frac{x^2 - 2\rho xy + y^2}{2(1-\rho^2)}} \frac{1}{\sqrt{1-\rho^2}}$$

1. Evaluate the marginal density of Y
2. Calculate its entropy
3. Evaluate the (conditional) density of $Y|X$ for the event $P\{X = 1\} = 1$
4. Show that conditioning had reduced its entropy

3 Solution: Joint Density

The density for a multivariate normal random variable is

$$f_{\vec{X}}(x) = (2\pi)^{-n/2} e^{-\frac{1}{2}(\vec{x}-\vec{\mu})'\Sigma^{-1}(\vec{x}-\vec{\mu})} \sqrt{|\Sigma^{-1}|}$$

for positive-definite covariance matrix Σ of dimension n and mean vector $\vec{\mu}$.

For the exercise, we have $n = 2$, $\mu = 0$, and

$$\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$$

with correlation parameter $-1 < \rho < 1$. Hence

$$\Sigma^{-1} = \frac{1}{1-\rho^2} \begin{pmatrix} 1 & -\rho \\ -\rho & 1 \end{pmatrix}$$

and the density simplifies to

$$f_{(X,Y)}(x,y) = \frac{1}{2\pi} e^{-\frac{x^2 - 2\rho xy + y^2}{2(1-\rho^2)}} \frac{1}{\sqrt{1-\rho^2}}$$

4 Solution: Conditional Density

According to the exercise, we are interested in the event $P\{X = 1\} = 1$ and what it tells us about Y . Prior to the observation, we only have the marginal density for Y ,

$$\begin{aligned} f_Y(y) &= \int_{-\infty}^{\infty} f_{(X,Y)}(x, y) dx \\ &= \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}y^2} \end{aligned}$$

which has entropy

$$H_Y = E[-\log f_Y(Y)] = \log \sqrt{2\pi e}$$

Once we have the observation on X , we can work with the conditional density,

$$\begin{aligned} f_{Y|X}(y) &= \frac{f_{(X,Y)}(1, y)}{\int_{-\infty}^{\infty} f_{(X,Y)}(1, y) dy} \\ &= \frac{1}{\sqrt{2\pi}} e^{-\frac{(y-\rho)^2}{2(1-\rho^2)}} \frac{1}{\sqrt{1-\rho^2}} \end{aligned}$$

which has entropy

$$H_{Y|X} = E[-\log f_{Y|X}(Y) | X] = \log \sqrt{2\pi e (1-\rho^2)}$$

Clearly $H_{Y|X} \leq H_Y$. To the extent that there is correlation, the observation acts to lower the entropy of the r.v. we are interested in describing.

5 Extension: Regression

We have demonstrated that (in the bivariate normal case at least) conditioning conveys information. We can see this in a more familiar light by considering the general bivariate result with four additional parameters and an arbitrary conditioning event, $P\{X = x\} = 1$.

You can confirm that the general conditional density works out to be

$$f_{Y|X}(y) = \frac{1}{\sqrt{2\pi}} e^{-\frac{(y-\mu_Y - \rho \frac{\sigma_Y}{\sigma_X}(x-\mu_X))^2}{2\sigma_Y^2(1-\rho^2)}} \frac{1}{\sigma_Y \sqrt{1-\rho^2}}$$

In other words,

$$Y|X \sim \mathcal{N}\left(\mu_Y + \rho \frac{\sigma_Y}{\sigma_X}(x - \mu_X), \sigma_Y^2(1-\rho^2)\right)$$

suggesting the transformation

$$Y|X \triangleq \alpha + \beta x + \epsilon \tag{1}$$

where

$$\beta = \frac{\text{cov}(X, Y)}{\text{var } X} \tag{2a}$$

$$\alpha = EY - \beta EX \tag{2b}$$

$$\epsilon \sim \mathcal{N}(0, \text{var } Y - \beta^2 \text{var } X) \tag{2c}$$

You may recognize this as the “population” result from classical statistics for ordinary least squares regression.

This observation is useful because it gives us guidance not only on grounding classical regression in modern estimation theory, but also on how we might adapt to situations where relationships are not well described by linear correlations.