

Kullback–Leibler divergence

From Wikipedia, the free encyclopedia

(Redirected from Kullback-Leibler divergence)

In probability theory and information theory, the **Kullback–Leibler divergence**^{[1][2][3]} (also **information divergence**, **information gain**, **relative entropy**, or **KLIC**) is a non-symmetric measure of the difference between two probability distributions P and Q . KL measures the expected number of extra bits required to code samples from P when using a code based on Q , rather than using a code based on P . Typically P represents the "true" distribution of data, observations, or a precise calculated theoretical distribution. The measure Q typically represents a theory, model, description, or approximation of P .

Although it is often intuited as a distance metric, the KL divergence is not a true metric — for example, it's not symmetric: the KL from P to Q is not necessarily the same as the KL from Q to P .

KL divergence is a special case of a broader class of divergences called f -divergences. Originally introduced by Solomon Kullback and Richard Leibler in 1951 as the **directed divergence** between two distributions, it is not the same as a divergence in calculus. However, the KL divergence can be derived from the Bregman divergence.

Contents

- 1 Definition
- 2 Motivation
- 3 Properties
- 4 Relation to metrics
- 5 Relation to other quantities of information theory
- 6 KL divergence and Bayesian updating
 - 6.1 Bayesian experimental design
- 7 Discrimination information
 - 7.1 Principle of minimum discrimination information
- 8 Relationship to available work
- 9 Quantum information theory
- 10 Relationship between models and reality
- 11 Symmetrised divergence
- 12 Relationship to Hellinger distance
- 13 Other probability-distance measures
- 14 Data differencing
- 15 See also
- 16 References
- 17 External links

Definition

For probability distributions P and Q of a discrete random variable their K–L divergence is defined to be

$$D_{\text{KL}}(P\|Q) = \sum_i P(i) \log \frac{P(i)}{Q(i)}.$$

In words, it is the average of the logarithmic difference between the probabilities P and Q , where the average is taken using the probabilities P . The K-L divergence is only defined if P and Q both sum to 1 and if $Q(i) > 0$ for any i such that $P(i) > 0$. If the quantity $0 \log 0$ appears in the formula, it is interpreted as zero.

For distributions P and Q of a continuous random variable, KL-divergence is defined to be the integral:^[4]

$$D_{\text{KL}}(P\|Q) = \int_{-\infty}^{\infty} p(x) \log \frac{p(x)}{q(x)} dx,$$

where p and q denote the densities of P and Q .

More generally, if P and Q are probability measures over a set X , and Q is absolutely continuous with respect to P , then the Kullback–Leibler divergence from P to Q is defined as

$$D_{\text{KL}}(P\|Q) = - \int_X \log \frac{dQ}{dP} dP,$$

where $\frac{dQ}{dP}$ is the Radon–Nikodym derivative of Q with respect to P , and provided the expression on the right-hand side exists. Likewise, if P is absolutely continuous with respect to Q , then

$$D_{\text{KL}}(P\|Q) = \int_X \log \frac{dP}{dQ} dP = \int_X \frac{dP}{dQ} \log \frac{dP}{dQ} dQ,$$

which we recognize as the entropy of P relative to Q . Continuing in this case, if μ is any measure on X for which $p = \frac{dP}{d\mu}$ and $q = \frac{dQ}{d\mu}$ exist, then the Kullback–Leibler divergence from P to Q is given as

$$D_{\text{KL}}(P\|Q) = \int_X p \log \frac{p}{q} d\mu.$$

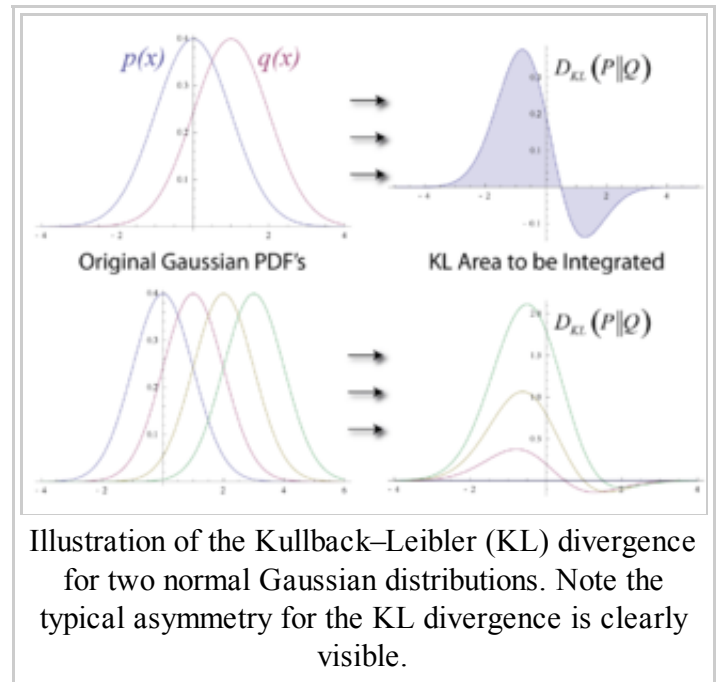
The logarithms in these formulae are taken to base 2 if information is measured in units of bits, or to base e if information is measured in nats. Most formulas involving the KL divergence hold irrespective of log base.

In this article, this will be referred to as the divergence from P to Q , although some authors call it the divergence "from Q to P " and others call it the divergence "between P and Q " (though note it is not symmetric as this latter terminology implies). Care must be taken due to the lack of standardization in terminology.

Motivation

In information theory, the Kraft–McMillan theorem establishes that any directly decodable coding scheme

for coding a message to identify one value x_i out of a set of possibilities X can be seen as representing an implicit probability distribution $q(x_i) = 2^{-l_i}$ over X , where l_i is the length of the code for x_i in bits. Therefore, KL divergence can be interpreted as the expected extra message-length per datum that must be communicated if a code that is optimal for a given (wrong) distribution Q is used, compared to using a code based on the true distribution P .



$$\begin{aligned} D_{\text{KL}}(P\|Q) &= -\sum_x p(x) \log q(x) + \sum_x p(x) \log p(x) \\ &= H(P, Q) - H(P) \end{aligned}$$

where $H(P, Q)$ is called the cross entropy of P and Q , and $H(P)$ is the entropy of P .

Properties

The Kullback–Leibler divergence is always non-negative,

$$D_{\text{KL}}(P\|Q) \geq 0,$$

a result known as Gibbs' inequality, with $D_{\text{KL}}(P\|Q)$ zero if and only if $P = Q$. The entropy $H(P)$ thus sets a minimum value for the cross-entropy $H(P, Q)$, the expected number of bits required when using a code based on Q rather than P ; and the KL divergence therefore represents the expected number of extra bits that must be transmitted to identify a value x drawn from X , if a code is used corresponding to the probability distribution Q , rather than the "true" distribution P .

The Kullback–Leibler divergence remains well-defined for continuous distributions, and furthermore is invariant under parameter transformations. It can therefore be seen as in some ways a more fundamental quantity than some other properties in information theory^[citation needed] (such as self-information or Shannon entropy), which can become undefined or negative for non-discrete probabilities.

The Kullback–Leibler divergence is additive for independent distributions in much the same way as Shannon entropy. If P_1, P_2 are independent distributions, with the joint distribution $P(x, y) = P_1(x)P_2(y)$, and Q, Q_1, Q_2 likewise, then

$$D_{\text{KL}}(P\|Q) = D_{\text{KL}}(P_1\|Q_1) + D_{\text{KL}}(P_2\|Q_2).$$

Relation to metrics

One might be tempted to call it a "distance metric" on the space of probability distributions, but this would

not be correct as the Kullback–Leibler divergence is not symmetric – that is, $D_{\text{KL}}(P \parallel Q) \neq D_{\text{KL}}(Q \parallel P)$, – nor does it satisfy the triangle inequality. Still, being a premetric, it generates a topology on the space of generalized probability distributions, of which probability distributions proper are a special case. More concretely, if $\{P_1, \dots, P_n\}$ is a sequence of distributions such that

$$\lim_{n \rightarrow \infty} D_{\text{KL}}(P_n \parallel Q) = 0$$

then it is said that $P_n \xrightarrow{D} Q$. Pinsker's inequality entails that $P_n \xrightarrow{D} P \Rightarrow P_n \xrightarrow{\text{TV}} P$, where the latter stands for the usual convergence in total variation.

Following Rényi (1970, 1961)^{[5][6]} the term is sometimes also called the **information gain** about X achieved if P can be used instead of Q . It is also called the **relative entropy**, for using Q instead of P .

Relation to other quantities of information theory

Many of the other quantities of information theory can be interpreted as applications of the KL divergence to specific cases.

The self-information,

$$I(m) = D_{\text{KL}}(\delta_{im} \parallel \{p_i\}),$$

is the KL divergence of the probability distribution $P(i)$ from a Kronecker delta representing certainty that $i=m$ — i.e. the number of extra bits that must be transmitted to identify i if only the probability distribution $P(i)$ is available to the receiver, not the fact that $i=m$.

The mutual information,

$$\begin{aligned} I(X; Y) &= D_{\text{KL}}(P(X, Y) \parallel P(X)P(Y)) \\ &= \mathbb{E}_X \{D_{\text{KL}}(P(Y|X) \parallel P(Y))\} \\ &= \mathbb{E}_Y \{D_{\text{KL}}(P(X|Y) \parallel P(X))\} \end{aligned}$$

is the KL divergence of the product $P(X)P(Y)$ of the two marginal probability distributions from the joint probability distribution $P(X, Y)$ — i.e. the expected number of extra bits that must be transmitted to identify X and Y if they are coded using only their marginal distributions instead of the joint distribution.

Equivalently, if the joint probability $P(X, Y)$ is known, it is the expected number of extra bits that must on average be sent to identify Y if the value of X is not already known to the receiver.

The Shannon entropy,

$$\begin{aligned} H(X) &= \text{(i)} \mathbb{E}_x \{I(x)\} \\ &= \text{(ii)} \log N - D_{\text{KL}}(P(X) \parallel P_U(X)) \end{aligned}$$

is the number of bits which would have to be transmitted to identify X from N equally likely possibilities, *less* the KL divergence of the uniform distribution $P_U(X)$ from the true distribution $P(X)$ — i.e. *less* the expected number of bits saved, which would have had to be sent if the value of X were coded according to

the uniform distribution $P_U(X)$ rather than the true distribution $P(X)$.

The conditional entropy,

$$\begin{aligned} H(X|Y) &= \log N - D_{\text{KL}}(P(X, Y) \| P_U(X)P(Y)) \\ &= \text{(i)} \log N - D_{\text{KL}}(P(X, Y) \| P(X)P(Y)) - D_{\text{KL}}(P(X) \| P_U(X)) \\ &= H(X) - I(X; Y) \\ &= \text{(ii)} \log N - \mathbb{E}_Y \{ D_{\text{KL}}(P(X|Y) \| P_U(X)) \} \end{aligned}$$

is the number of bits which would have to be transmitted to identify X from N equally likely possibilities, less the KL divergence of the product distribution $P_U(X)P(Y)$ from the true joint distribution $P(X, Y)$ — i.e. less the expected number of bits saved which would have had to be sent if the value of X were coded according to the uniform distribution $P_U(X)$ rather than the conditional distribution $P(X|Y)$ of X given Y .

The cross entropy between two probability distributions measures the average number of bits needed to identify an event from a set of possibilities, if a coding scheme is used based on a given probability distribution q , rather than the "true" distribution p . The cross entropy for two distributions p and q over the same probability space is thus defined as follows:

$$H(p, q) = \mathbb{E}_p[-\log q] = H(p) + D_{\text{KL}}(p \| q).$$

KL divergence and Bayesian updating

In Bayesian statistics the KL divergence can be used as a measure of the information gain in moving from a prior distribution to a posterior distribution. If some new fact $Y = y$ is discovered, it can be used to update the probability distribution for X from $p(x|I)$ to a new posterior probability distribution $p(x|y, I)$ using Bayes' theorem:

$$p(x|y, I) = \frac{p(y|x, I)p(x|I)}{p(y|I)}$$

This distribution has a new entropy

$$H(p(\cdot|y, I)) = \sum_x p(x|y, I) \log p(x|y, I),$$

which may be less than or greater than the original entropy $H(p(\cdot|I))$. However, from the standpoint of the new probability distribution one can estimate that to have used the original code based on $p(x|I)$ instead of a new code based on $p(x|y, I)$ would have added an expected number of bits

$$D_{\text{KL}}(p(\cdot|y, I) \| p(\cdot|I)) = \sum_x p(x|y, I) \log \frac{p(x|y, I)}{p(x|I)}$$

to the message length. This therefore represents the amount of useful information, or information gain, about X , that we can estimate has been learned by discovering $Y = y$.

If a further piece of data, $Y_2 = y_2$, subsequently comes in, the probability distribution for x can be updated

further, to give a new best guess $p(x|y_1, y_2, I)$. If one reinvestigates the information gain for using $p(x|y_1, I)$ rather than $p(x|I)$, it turns out that it may be either greater or less than previously estimated:

$$\sum_x p(x|y_1, y_2, I) \log \frac{p(x|y_1, y_2, I)}{p(x|I)} \text{ may be } \le \text{ or } > \text{ than } \sum_x p(x|y_1, I) \log \frac{p(x|y_1, I)}{p(x|I)}$$

and so the combined information gain does *not* obey the triangle inequality:

$$D_{\text{KL}}(p(\cdot|y_1, y_2, I) \| p(\cdot|I)) \text{ may be } <, = \text{ or } > \text{ than } \\ D_{\text{KL}}(p(\cdot|y_1, y_2, I) \| p(\cdot|y_1, I)) + D_{\text{KL}}(p(\cdot|y_1, I) \| p(x|I))$$

All one can say is that on *average*, averaging using $p(y_2|y_1, x, I)$, the two sides will average out.

Bayesian experimental design

A common goal in Bayesian experimental design is to maximise the expected KL divergence between the prior and the posterior^[citation needed]. When posteriors are approximated to be Gaussian distributions, a design maximising the expected KL divergence is called Bayes d-optimal.

Discrimination information

The Kullback–Leibler divergence $D_{\text{KL}}(p(x|H_1) \| p(x|H_0))$ can also be interpreted as the expected **discrimination information** for H_1 over H_0 : the mean information per sample for discriminating in favour of a hypothesis H_1 against a hypothesis H_0 , when hypothesis H_1 is true. Another name for this quantity, given to it by I.J. Good, is the expected weight of evidence for H_1 over H_0 to be expected from each sample.

The expected weight of evidence for H_1 over H_0 is **not** the same as the information gain expected per sample about the probability distribution $p(H)$ of the hypotheses,

$$D_{\text{KL}}(p(x|H_1) \| p(x|H_0)) \neq IG = D_{\text{KL}}(p(H|x) \| p(H|I)).$$

Either of the two quantities can be used as a utility function in Bayesian experimental design, to choose an optimal next question to investigate: but they will in general lead to rather different experimental strategies.

On the entropy scale of *information gain* there is very little difference between near certainty and absolute certainty—coding according to a near certainty requires hardly any more bits than coding according to an absolute certainty. On the other hand, on the logit scale implied by weight of evidence, the difference between the two is enormous – infinite perhaps; this might reflect the difference between being almost sure (on a probabilistic level) that, say, the Riemann hypothesis is correct, compared to being certain that it is correct because one has a mathematical proof. These two different scales of loss function for uncertainty are *both* useful, according to how well each reflects the particular circumstances of the problem in question.

Principle of minimum discrimination information

The idea of Kullback–Leibler divergence as discrimination information led Kullback to propose the Principle of **Minimum Discrimination Information** (MDI): given new facts, a new distribution f should

be chosen which is as hard to discriminate from the original distribution f_0 as possible; so that the new data produces as small an information gain $D_{\text{KL}}(f||f_0)$ as possible.

For example, if one had a prior distribution $p(x,a)$ over x and a , and subsequently learnt the true distribution of a was $u(a)$, the Kullback–Leibler divergence between the new joint distribution for x and a , $q(x|a)u(a)$, and the earlier prior distribution would be:

$$D_{\text{KL}}(q(x|a)u(a)||p(x,a)) = \mathbb{E}_{u(a)}\{D_{\text{KL}}(q(x|a)||p(x|a))\} + D_{\text{KL}}(u(a)||p(a)),$$

i.e. the sum of the KL divergence of $p(a)$ the prior distribution for a from the updated distribution $u(a)$, plus the expected value (using the probability distribution $u(a)$) of the KL divergence of the prior conditional distribution $p(x|a)$ from the new conditional distribution $q(x|a)$. This is minimised if $q(x|a) = p(x|a)$ over the whole support of $u(a)$; and we note that this result incorporates Bayes' theorem, if the new distribution $u(a)$ is in fact a δ function representing certainty that a has one particular value.

MDI can be seen as an extension of Laplace's Principle of Insufficient Reason, and the Principle of Maximum Entropy of E.T. Jaynes. In particular, it is the natural extension of the principle of maximum entropy from discrete to continuous distributions, for which Shannon entropy ceases to be so useful (see *differential entropy*), but the KL divergence continues to be just as relevant.

In the engineering literature, MDI is sometimes called the **Principle of Minimum Cross-Entropy** (MCE) or **Minxent** for short. This is not entirely helpful. Minimising the KL divergence of m from p with respect to m is equivalent to minimising the cross-entropy of p and m , since

$$H(p, m) = H(p) + D_{\text{KL}}(p||m),$$

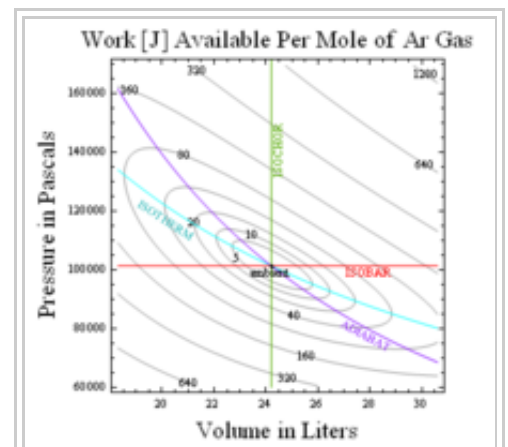
which is appropriate if one is trying to choose a least 'brain-damaged' approximation to p . However, this is just as often *not* the task one is trying to achieve. Instead, just as often it is m that is some fixed prior reference measure, and p that one is attempting to optimise by minimising $D_{\text{KL}}(p||m)$ subject to some constraint. This has led to some ambiguity in the literature, with some authors attempting to resolve the inconsistency by redefining cross-entropy to be $D_{\text{KL}}(p||m)$, rather than $H(p,m)$.

Relationship to available work

Surprisals^[7] add where probabilities multiply. The surprisal for an event of probability p is defined as $s \equiv k \ln[1/p]$. If k is $\{1, 1/\ln 2, 1.38 \times 10^{-23}\}$ then surprisal is in $\{\text{nats, bits, or } J/K\}$ so that, for instance, there are N bits of surprisal for landing all "heads" on a toss of N coins.

Best-guess states (e.g. for atoms in a gas) are inferred by maximizing the *average-surprisal* S (entropy) for a given set of control parameters (like pressure P or volume V). This constrained entropy maximization, both classically^[8] and quantum mechanically^[9], minimizes Gibbs availability in entropy units^[10] $A \equiv -k \ln Z$ where Z is a constrained multiplicity or partition function.

When temperature T is fixed, free-energy ($T \times A$) is also



Pressure versus volume plot of available work from a mole of Argon gas relative to ambient, calculated as T_0 times KL divergence.

minimized. Thus if T , V and number of molecules N are constant, the Helmholtz free energy $F \equiv U - TS$ (where U is energy) is minimized as a system "equilibrates." If T and P are held constant (say during processes in your body), the Gibbs free energy $G \equiv U + PV - TS$ is minimized instead. The change in free energy under these conditions is a measure of available work that might be done in the process. Thus available work for an ideal gas at constant temperature T_0 and pressure P_0 is $W = \Delta G = NkT_0 \Theta[V/V_0]$ where $V_0 = NkT_0/P_0$ and $\Theta[x] \equiv x - 1 - \ln x \geq 0$ (see also Gibbs inequality).

More generally^[11] the work available relative to some ambient is obtained by multiplying ambient temperature T_0 by KL-divergence or *net-surprisal* $\Delta I \geq 0$, defined as the average value of $k \ln[p/p_0]$ where p_0 is the probability of a given state under ambient conditions. For instance, the work available in equilibrating a monatomic ideal gas to ambient values of V_0 and T_0 is thus $W = T_0 \Delta I$, where KL-divergence $\Delta I = Nk(\Theta[V/V_0] + \frac{3}{2}\Theta[T/T_0])$. The resulting contours of constant KL-divergence, at right for a mole of Argon at standard temperature and pressure, for example put limits on the conversion of hot to cold as in flame-powered air-conditioning or in the unpowered device to convert boiling-water to ice-water discussed here.^[12] Thus KL-divergence measures thermodynamic availability in bits.

Quantum information theory

For density matrices P and Q on a Hilbert space the K–L divergence (or relative entropy as it is often called in this case) from P to Q is defined to be

$$D_{\text{KL}}(P\|Q) = \text{Tr}(P(\log(P) - \log(Q))).$$

In quantum information science it can also be used as a measure of entanglement in a state.

Relationship between models and reality

Just as KL-divergence of "ambient from actual" measures thermodynamic availability, KL-divergence of "model from reality" is also useful even if the only clues we have about reality are some experimental measurements. In the former case KL-divergence describes *distance to equilibrium* or (when multiplied by ambient temperature) the amount of *available work*, while in the latter case it tells you about surprises that reality has up its sleeve or, in other words, *how much the model has yet to learn*.

Although this tool for evaluating models against systems that are accessible experimentally may be applied in any field, its application to models in ecology via Akaike information criterion are particularly well described in papers^[13] and a book^[14] by Burnham and Anderson. In a nutshell the KL-divergence of a model from reality may be estimated, to within a constant additive term, by a function (like the squares summed) of the deviations observed between data and the model's predictions. Estimates of such divergence for models that share the same additive term can in turn be used to choose between models.

When trying to fit parametrized models to data there are various estimators which attempt to minimize Kullback–Leibler divergence, such as maximum likelihood and maximum spacing estimators.

Symmetrised divergence

Kullback and Leibler themselves actually defined the divergence as:

$$D_{\text{KL}}(P\|Q) + D_{\text{KL}}(Q\|P)$$

which is symmetric and nonnegative. This quantity has sometimes been used for feature selection in classification problems, where P and Q are the conditional pdfs of a feature under two different classes.

An alternative is given via the λ divergence,

$$D_{\lambda}(P\|Q) = \lambda D_{\text{KL}}(P\|\lambda P + (1 - \lambda)Q) + (1 - \lambda)D_{\text{KL}}(Q\|\lambda P + (1 - \lambda)Q),$$

which can be interpreted as the expected information gain about X from discovering which probability distribution X is drawn from, P or Q , if they currently have probabilities λ and $(1 - \lambda)$ respectively.

The value $\lambda = 0.5$ gives the Jensen–Shannon divergence, defined by

$$D_{\text{JS}} = \frac{1}{2}D_{\text{KL}}(P\|M) + \frac{1}{2}D_{\text{KL}}(Q\|M)$$

where M is the average of the two distributions,

$$M = \frac{1}{2}(P + Q).$$

D_{JS} can also be interpreted as the capacity of a noisy information channel with two inputs giving the output distributions p and q . The Jensen–Shannon divergence is the square of a metric that is equivalent to the Hellinger metric, and the Jensen–Shannon divergence is also equal to one-half the so-called *Jeffreys divergence* (Rubner et al., 2000; Jeffreys 1946).

Relationship to Hellinger distance

If P and Q are two probability measures, then the squared Hellinger distance is the quantity given by

$$H^2(P, Q) = \frac{1}{2} \int \left(\sqrt{\frac{dP}{d\lambda}} - \sqrt{\frac{dQ}{d\lambda}} \right)^2 d\lambda.$$

Noting that $x - 1 \geq \ln x$, so that in particular, $\sqrt{x} - 1 \geq \frac{1}{2} \ln(x)$, we see that

$$\sqrt{\frac{dP}{dQ}} - 1 \geq \frac{1}{2} \ln \frac{dP}{dQ}.$$

Taking expectations with respect to Q , we get

$$2H^2(P, Q) \leq E_Q \ln \frac{dQ}{dP} = \ln 2 \cdot D_{\text{KL}}(Q\|P).$$

Hence

$$D_{\text{KL}}(Q\|P) \geq \frac{2}{\ln 2} H^2(P, Q).$$

Other probability-distance measures

Other measures of probability distance are the *histogram intersection*, *Chi-square statistic*, *quadratic form distance*, *match distance*, *Kolmogorov–Smirnov distance*, and *earth mover's distance* ^[15].

Data differencing

Main article: Data differencing

Just as *absolute* entropy serves as theoretical background for data *compression*, *relative* entropy serves as theoretical background for data *differencing* – the absolute entropy of a set of data in this sense being the data required to reconstruct it (minimum compressed size), while the relative entropy of a target set of data, given a source set of data, is the data required to reconstruct the target *given* the source (minimum size of a patch).

See also

- Bregman divergence
- Deviance information criterion
- Bayesian information criterion
- Quantum relative entropy
- Information gain in decision trees
- Solomon Kullback and Richard Leibler
- Information theory and measure theory

References

- ¹ ^ Kullback, S.; Leibler, R.A. (1951). "On Information and Sufficiency". *Annals of Mathematical Statistics* **22** (1): 79–86. doi:10.1214/aoms/1177729694 (<http://dx.doi.org/10.1214%2Faoms%2F1177729694>) . MR39968 (<http://www.ams.org/mathscinet-getitem?mr=39968>) .
- ² ^ S. Kullback (1959) *Information theory and statistics* (John Wiley and Sons, NY).
- ³ ^ Kullback, S. (1987). "Letter to the Editor: The Kullback–Leibler distance". *The American Statistician* **41** (4): 340–341. JSTOR 2684769 (<http://www.jstor.org/stable/2684769>) .
- ⁴ ^ C. Bishop (2006). *Pattern Recognition and Machine Learning*. p. 55.
- ⁵ ^ A. Rényi (1970). *Probability Theory*. New York: Elsevier. Appendix, Sec.4. ISBN 0486458679.
- ⁶ ^ A. Rényi (1961). "On measures of information and entropy" (http://digitalassets.lib.berkeley.edu/math/ucb/text/math_s4_v1_article-27.pdf) . *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics and Probability 1960* (http://digitalassets.lib.berkeley.edu/math/ucb/text/math_s4_v1_article-27.pdf) . pp. 547–561. http://digitalassets.lib.berkeley.edu/math/ucb/text/math_s4_v1_article-27.pdf.
- ⁷ ^ Myron Tribus (1961) *Thermodynamics and thermostatics* (D. Van Nostrand, New York)
- ⁸ ^ E. T. Jaynes (1957) Information theory and statistical mechanics (<http://bayes.wustl.edu/etj/articles/theory.1.pdf>) , *Physical Review* **106**:620
- ⁹ ^ E. T. Jaynes (1957) Information theory and statistical mechanics II (<http://bayes.wustl.edu/etj/articles/theory.2.pdf>) , *Physical Review* **108**:171
- ¹⁰ ^ J.W. Gibbs (1873) A method of geometrical representation of thermodynamic properties of substances by means of surfaces, reprinted in *The Collected Works of J. W. Gibbs, Volume I Thermodynamics*, ed. W. R. Longley and R. G. Van Name (New York: Longmans, Green, 1931) footnote page 52.
- ¹¹ ^ M. Tribus and E. C. McIrvine (1971) Energy and information, *Scientific American* **224**:179–186.

12. ^ P. Fraundorf (2007) Thermal roots of correlation-based complexity (<http://www3.interscience.wiley.com/cgi-bin/abstract/117861985/ABSTRACT>) , *Complexity* **13**:3, 18–26
13. ^ Kenneth P. Burnham and David R. Anderson (2001) Kullback–Leibler information as a basis for strong inference in ecological studies (<http://www.publish.csiro.au/paper/WR99107.htm>) , *Wildlife Research* **28**:111–119.
14. ^ Burnham, K. P. and Anderson D. R. (2002) *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach, Second Edition* (Springer Science, New York) ISBN 978-0-387-95364-9.
15. ^ Rubner, Y., Tomasi, C., and Guibas, L. J., 2000. The Earth Mover's distance as a metric for image retrieval. *International Journal of Computer Vision*, **40**(2): 99–121.

External links

- Matlab code for calculating KL divergence (<http://www.mathworks.com/matlabcentral/fileexchange/loadFile.do?objectId=13089&objectType=file>)
- Sergio Verdú, Relative Entropy (http://videlectures.net/nips09_verdu_re/) , NIPS 2009. One-hour video lecture.

Retrieved from "http://en.wikipedia.org/wiki/Kullback%E2%80%93Leibler_divergence"

Categories: Statistical theory | Entropy and information | F-divergences | Thermodynamics

- This page was last modified on 11 December 2010 at 14:39.
 - Text is available under the Creative Commons Attribution-ShareAlike License; additional terms may apply. See Terms of Use for details.
- Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.