

Article

Optimal Belief Approximation

Reimar H. Leike ^{1,2,*} and Torsten A. Enßlin ^{1,2}

¹ Max-Planck-Institut für Astrophysik, Karl-Schwarzschildstr. 1, 85748 Garching, Germany; ensslin@mpa-garching.mpg.de

² Ludwig-Maximilians-Universität München, Geschwister-Scholl-Platz 1, 80539 Munich, Germany

* Correspondence: reimar@mpa-garching.mpg.de

Received: 18 April 2017; Accepted: 5 July 2017; Published: 4 August 2017

Abstract: In Bayesian statistics probability distributions express beliefs. However, for many problems the beliefs cannot be computed analytically and approximations of beliefs are needed. We seek a loss function that quantifies how “embarrassing” it is to communicate a given approximation. We reproduce and discuss an old proof showing that there is only one ranking under the requirements that (1) the best ranked approximation is the non-approximated belief and (2) that the ranking judges approximations only by their predictions for actual outcomes. The loss function that is obtained in the derivation is equal to the Kullback-Leibler divergence when normalized. This loss function is frequently used in the literature. However, there seems to be confusion about the correct order in which its functional arguments—the approximated and non-approximated beliefs—should be used. The correct order ensures that the recipient of a communication is only deprived of the minimal amount of information. We hope that the elementary derivation settles the apparent confusion. For example when approximating beliefs with Gaussian distributions the optimal approximation is given by moment matching. This is in contrast to many suggested computational schemes.

Keywords: information theory; Bayesian inference; loss function; axiomatic derivation; machine learning

1. Introduction

In Bayesian statistics, probabilities are interpreted as degrees of belief. For any set of mutually exclusive and exhaustive events, one expresses the state of knowledge as a probability distribution over that set. The probability of an event then describes the personal confidence that this event will happen or has happened. As a consequence, probabilities are subjective properties reflecting the amount of knowledge an observer has about the events; a different observer might know which event happened and assign different probabilities. If an observer gains information, she updates the probabilities she had assigned before.

If the set of possible mutually exclusive and exhaustive events is infinite, it is generally impossible to store all entries of the corresponding probability distribution on a computer or communicate it through a channel with finite bandwidth. One therefore needs to approximate the probability distribution which describes one’s belief. Given a limited set X of approximative beliefs $q(s)$ on a quantity s , what is the best belief to approximate the actual belief as expressed by the probability $p(s)$?

In the literature, it is sometimes claimed that the best approximation is given by the $q \in X$ that minimizes the Kullback–Leibler divergence (“approximation” KL) [1]

$$\text{KL}(p, q) = \sum_s p(s) \ln \left(\frac{p(s)}{q(s)} \right) \quad (1)$$

where q is the approximation and p is the real belief. We refer to this functional as “approximation” KL to emphasize its role in approximation, which will be derived in the course of this paper and to distinguish it from the same functional, with q being a prior belief and p being the posterior belief to

which this KL is minimized in inference. We will refer to the functional with q being the input and p obtained through minimization as “inference KL”. In Equation (1), minimization is done with respect to its second argument. The derivation of this particular functional form varies from field to field.

For example, in coding theory, one tries to minimize the amount of bandwidth needed to transmit a message. Given a prior q over the symbols that the message consists of, an optimal scheme can be derived. The approximation KL gives the expected amount of extra bits needed to transmit such a message if the symbols are actually drawn from the probability distribution p instead of q [2]. If we know that p is the real probability distribution, the best approximative probability distribution $q \in X$ to base a coding on is therefore the one minimizing the approximation KL. However, it is not clear that minimizing the amount of bits transferred is the best or even only measure expressing how good such an approximation is in general.

In machine learning and deep learning, neural networks are trained to understand abstract data d ; for example, to assign a label s to it. This task can be viewed as fitting an approximative probability distribution $q(s|d)$ to a true generating probability distribution $p(s|d)$. For this, the approximative probability distribution is parametrized (to a neural network) and then matched to the true, generating probability distribution using a loss function and samples. The most frequently used loss function is the cross entropy, which is equivalent to the approximation KL. The reason to use this form is often either inspired from coding theory, or by experimental experience [3].

Another argument for minimizing the approximation KL is given in Chapter 13 of Reference [4], where it is claimed that this yields the maximum likelihood estimation to $p(s)$ among the probability distributions in X and that it gives an unbiased and unique approximation. Interchanging the arguments of the Kullback–Leibler divergence (the inference KL used in variational Bayes) generally leads to a biased estimate and does not necessarily yield a unique result. These arguments undoubtedly give evidence for why minimizing the approximation KL gives a good estimate. However, this does not exclude all other methods. Having an unbiased estimate refers to getting the right mean. In our picture, this is a result of optimal approximation and not a requirement for optimality. Additionally, this result was derived with the help of information geometry, whose applicability to non-local problems is criticized, for example, in References [5,6].

Contrary to the evidence for minimizing the approximation KL, we find many examples where an approximation is made by minimizing other functionals; for example, minimizing the inference KL (e.g., [7–13]). For many but not all of them, this is because minimizing the approximation KL is not feasible in practice in their case due to the real distribution p not being accessible.

In this paper, we seek to bring together the different motivations and give a full and consistent picture. The proof we present here is not new; it goes back to [14], where there is an exact mathematical derivation for probability densities analogously to our derivation. However, there are earlier publications dealing with the discrete case [15–17]. Although this proof dates back at least 40 years, its implication on approximating beliefs seem to be quite unknown—especially in the community of physicists applying Bayesian methods. In this paper, we reproduce a slightly modified version of this proof, give the result a new interpretation and add further justification for the prerequisites used, laying emphasis on why one has to accept the axioms necessary for the derivation if one is a Bayesian. We provide argumentation for why believe approximation is an important and omnipresent topic.

We lay the emphasis of this paper more on interpretation of results and justification of prerequisites, and thus present an easy version of the proof where the loss function is assumed to be differentiable. The proof can however be extended to the general case of non-differentiable loss [18]. The argument we reproduce gives evidence that minimizing the approximation KL is the best approximation in theory. This argument does not rest on information geometry nor is it restricted to coding theory. By imposing two consistency requirements, one is able to exclude all ranking functions with the exception of one for ranking the approximative probability distributions $q \in X$. For this, one employs the principle of loss functions [19], also called cost functions, regret functions, (or with flipped sign, utility functions or score functions) and shows that the unique loss function for ranking approximated probability distributions

is the approximation KL. For us, a ranking is a total order indicating preference, whereas a loss is a map to \mathbb{R} , which induces a ranking but additionally gives an absolute scale to compare preferences. The presented axiomatic derivation does not give rise to any new method, but it enables a simple checking for whether a certain approximation is most optimally done through the approximation KL. There are many other examples of axiomatic derivations seeking to support information theory on a fundamental level. Some notable examples are Cox derivation [20] of Bayesian probability theory as a unique extension of Boolean algebra as well as a scientific discussion on the maximum entropy principle [21–23], establishing the inference KL as unique inference tool (which gave rise to the naming convention in this paper). Most of these arguments rely on page-long proofs to arrive at the Kullback–Leibler divergence. The proof that is sketched in this paper is only a few lines long, but nonetheless standard literature for axiomatic derivation in Bayesianism does not cite this “easy” derivation (e.g., the influential Reference [21]). As already discussed, approximation is an important and unavoidable part of information theory, and with the axiomatic derivation presented here we seek to provide orientation to scientists searching for a way to approximate probability distributions.

In Section 2, we introduce the concept of loss functions, which is used in Section 3 to define an optimal scheme for approximating probability distributions that express beliefs. We briefly discuss the relevance of our derivations for the scientific community in Section 4. We conclude in Section 5.

2. Loss Functions

The idea to evaluate predictions based on loss functions dates back 70 years, and was first introduced by Brier [24]. We explain loss functions by the means of parameter estimation. Imagine that one would like to give an estimate of a parameter s that is not known, which value of s should be taken as an estimate? One way to answer this question is by using loss functions. For this note that $p(s)$ is now formally a probability measure, however we choose to write $\int ds p(s)$ instead of $\int dp(s)$ as if $p(s)$ would be a probability density. A loss function in the setting of parameter estimation is a function that takes an estimate σ for s and quantifies how “embarrassing” this estimate is if $s = s_0$ turns out to be the case:

$$\mathcal{L}(\sigma, s_0)$$

The expected embarrassment can be computed by using the knowledge $p(s)$ about s :

$$\langle \mathcal{L}(\sigma, s_0) \rangle_{p(s_0)} = \int ds_0 \mathcal{L}(\sigma, s_0) p(s_0)$$

The next step is to take the estimate σ that minimizes the expected embarrassment; that is, the expectation value of the loss function. For different loss functions, one arrives at different recipes for how to extract an estimate σ from the belief $p(s)$; for example, for $s \in \mathbb{R}$:

$$\mathcal{L}(\sigma, s_0) = \begin{cases} -\delta(\sigma - s_0) & \Rightarrow \text{Take } \sigma \text{ such that} \\ & p(s)|_{s=\sigma} \text{ is maximal} \\ |\sigma - s_0| & \Rightarrow \text{Take } \sigma \text{ to be the median} \\ (\sigma - s_0)^2 & \Rightarrow \text{Take } \sigma \text{ to be the mean} \end{cases} \quad (2)$$

In the context of parameter estimation, there is no general loss function that one should take. In many scientific applications, the third option is favored, but different situations might enforce different loss functions. In the context of probability distributions, one has a mathematical structure available to guide the choice. In this context, one can restrict the possibilities by requiring consistent loss functions.

3. The Unique Loss Function

How embarrassing is it to approximate a probability distribution by $q(s)$ even though it is actually $p(s)$? We quantify the embarrassment in a loss function

$$\mathcal{L}\left(\frac{q}{m}, s_0\right) \quad (3)$$

which says how embarrassing it is to tell someone $q(s)$ is one's belief about s in the event that later s is measured to be s_0 . Here m is introduced as reference measure to make \mathcal{L} coordinate independent. For a finite set coordinate independence is trivially fulfilled and it might seem that having a reference measure m is superficial. Note however, that it is a sensible additional requirement to have the quantification be invariant under splitting of events, i.e., mapping to a bigger set where two now distinguishable events represent one former large event. The quotient $\frac{q}{m}$ is invariant under such splitting of events, whereas q itself is not. The reference measure m can be any measure such that q is absolutely continuous with respect to m .

Note further that we restrict ourselves to the case that we get to know the exact value of s . This does not make our approach less general; imagining that we would instead take a more general loss $L(q, \tilde{q}(s))$ where \tilde{q} is the knowledge about s at some later point, then we may define $\mathcal{L}\left(\frac{q}{m}, s_0\right) = L(q, \delta_{ss_0})$ with δ denoting the Kronecker or Dirac delta function, and thus restrict ourselves again to the case of exact knowledge. This line of reasoning was spelled out in detail by John Skilling [25]:

“If there are general theories, then they must apply to special cases”.

To decide which belief to tell someone, we look at the expected loss

$$\left\langle \mathcal{L}\left(\frac{q}{m}, s_0\right) \right\rangle_{p(s_0)} = \int ds_0 \mathcal{L}\left(\frac{q}{m}, s_0\right) p(s_0) \quad (4)$$

and try to find a $q \in X$ that minimizes this expected loss. To sum up, if we are given a loss function, we have a recipe for how to optimally approximate the belief. Which loss functions are sensible, though? We enforce two criteria that a good loss function should satisfy.

Criterion 1. (Locality) If $s = s_0$ turned out to be the case, \mathcal{L} only depends on the prediction q actually makes about s_0 :

$$\mathcal{L}\left(\frac{q}{m}, s_0\right) = \mathcal{L}\left(\frac{q(s_0)}{m(s_0)}\right) \quad (5)$$

Note that we make an abuse of notation here, denoting the function on both sides of the equation by the same symbol. The criterion is called locality because it demands that the functional of $\frac{q}{m}$ should be evaluated locally for every s_0 . It also forbids a direct dependence of the loss \mathcal{L} on s_0 which excludes losses that are a priori biased towards certain outcomes s_0 .

This form of locality is an intrinsically Bayesian property. Consider a situation where one wants to decide which of two rivaling hypotheses to believe. In order to distinguish them, some data d are measured. To update the prior using Bayes theorem, one only needs to know how probable the measured data d are given each hypothesis, not how probable other possible data $\tilde{d} \neq d$ that were not measured are. This might seem intuitive, but there exist hypothesis decision methods (not necessarily based on loss functions) that do not fulfill this property. For example, the non-Bayesian p -value depends mostly on data that were not measured (all the data that are at least as “extreme” as the measured data). Thus, it is a property of Bayesian reasoning to judge predictions only by using what was predicted about things that were measured.

The second criterion is even more natural. If one is not restricted in what can be told to others, then the best thing should be to tell them the actual belief p .

Criterion 2. (Optimality of the actual belief, properness.) Let X be the set of all probability distributions over s . For all p and all m , the probability distribution $q \in X$ with minimal expected loss is $q = p$:

$$0 = \left(\frac{\partial}{\partial q(s)} \left\langle \mathcal{L} \left(\frac{q}{m}, s_0 \right) \right\rangle_p \right)_{q=p} \tag{6}$$

The last criterion is also referred to as a proper loss (score) function in the literature; see Reference [26] for a mathematical overview of different proper scoring rules. Our version of this property is slightly modified to the version that is found in the literature as we demand this optimum to be obtained independently of a reference measure m . There is a fundamental Bayesian desiderata stating that “If there are multiple ways to arrive at a solutions, then they must agree.” We’d like to justify why this is a property that is absolutely important. If one uses statistics as a tool to answer some question, then if that answer would depend on how statistics is applied, then this statistic itself is inconsistent. In our case, where the defined loss function is dependent on an arbitrary reference measure m , the result is thus forced to be independent of that m .

Note furthermore that although intuitively we want the global optimum to be at the actual belief p (referred to as strictly proper in the literature), mathematically we only need it to be an extreme value for our derivation.

Having these two consistency requirements fixed, we derive which kind of consistent loss functions are possible. We insert Equation (5) into Equation (6), expand the domain of the loss function to not necessarily normalized positive vectors $q(s)$ but introduce λ as a Lagrange multiplier to account for the fact that we minimize under the constraint of normalization. We compute

$$\begin{aligned} 0 &= \left(\frac{\partial}{\partial q(s)} \int ds_0 \left(\mathcal{L} \left(\frac{q(s_0)}{m(s_0)} \right) p(s_0) + \lambda q(s_0) \right) \right)_{q=p} \\ &= \int ds_0 \partial \mathcal{L} \left(\frac{p(s_0)}{m(s_0)} \right) \frac{\delta(s-s_0)}{m(s_0)} p(s_0) + \lambda \delta(s-s_0) \\ &= \partial \mathcal{L} \left(\frac{p(s)}{m(s)} \right) \frac{p(s)}{m(s)} + \lambda \\ &\Rightarrow \partial \mathcal{L} \left(\frac{p(s)}{m(s)} \right) = - \frac{\lambda m(s)}{p(s)} \end{aligned} \tag{7}$$

Here $\partial \mathcal{L}$ denotes the derivative of \mathcal{L} . In the next step we substitute $x := \frac{p(s)}{m(s)}$ for the quotient. Note that Equation (7) holds for all positive real values of $x \in \mathbb{R}^+$ since the requested measure independence of the resulting ranking permits to insert any measure m . We then obtain

$$\begin{aligned} \partial \mathcal{L}(x) &= - \frac{\lambda}{x} \\ \Rightarrow \mathcal{L}(x) &= -C \ln(x) + D \end{aligned} \tag{8}$$

where $C > 0$ and D are constants with respect to q . Note that through the two consistency requirements one is able to completely fix what the loss function is. In the original proof of [14], there arises an additional possibility for \mathcal{L} if the sample space consists of 2 elements. In that case, the locality axiom, as it is used in the literature, does not constrain \mathcal{L} at all. In our case, where we introduced m as a reference measure, we are able to exclude that possibility. Note that the constants C and D are irrelevant for determining the optimal approximation as they do not affect where the minimum of the loss function is.

To sum up our result if one is restricted to the closed set X of probability distributions, one should take that $q \in X$ that minimizes

$$\left\langle \mathcal{L} \left(\frac{q}{m}, s_0 \right) \right\rangle_{p(s_0)} = - \int ds_0 p(s_0) \ln \left(\frac{q(s_0)}{m(s_0)} \right) \tag{9}$$

in order to obtain the optimal approximation, where it is not important what m is used.

If one takes $m = 1$, this loss is the cross entropy

$$\langle -\ln(q(s_0)) \rangle_{p(s_0)} .$$

If one desires a rating of how good an approximation is, and not only a ranking which approximation is best, one could go one step further and enforce a third criterion:

Criterion 3. (zero loss of the actual belief) *For all p , the expected loss of the probability distribution p is 0:*

$$0 = \left\langle \mathcal{L} \left(\frac{p}{m}, s_0 \right) \right\rangle_p \quad (10)$$

This criterion trivially forces $m = p$ and makes the quantification unique while inducing the same ranking. Thus we arrive at the Kullback-Leibler divergence

$$\text{KL}(p, q) = \int ds_0 p(s_0) \ln \left(\frac{p(s_0)}{q(s_0)} \right)$$

as the optimal rating and ranking function.

To phrase the result in words, the optimal way to approximate the belief p is such that given the approximated belief q , the amount of information $\text{KL}(p, q)$ that has to be obtained for someone who believes q to arrive back at the actual belief p is minimal. We should make it as easy as possible for someone who got an approximation q to get to the correct belief p . This sounds like a trivial statement, explaining why the approximation KL is already widely used for exactly that task.

4. Discussion

We briefly discuss the implications of these results.

In comparison to Reference [2,4], we presented another more elementary line of argumentation for the claim that the approximation KL is the correct ranking function for approximating which holds in a more general setting.

Other works that base their results on minimizing the inference KL ($\text{KL}(q, p)$) for belief approximation are not optimal with respect to the ranking function we derived in Section 3. One reason for preferring the for this purpose non-optimal inference KL is that it is computationally feasible for many applications, in contrast to the optimal approximation. As long as the optimal scheme is not computationally accessible, this argument has its merits.

Another reason for minimizing the inference KL for approximation that is often cited (e.g., [27]) is that it gives a lower bound to the log-likelihood

$$\ln(p(d|s)) = \left\langle \ln \left(\frac{p(d, s)}{q(s)} \right) \right\rangle_{q(s)} + \text{KL}(q, p) \quad (11)$$

which for example gives rise to the expectation maximization (EM-) algorithm [28]. However, the method only gives rise to maximum a posteriori or maximum likelihood solutions, which corresponds to optimizing the δ -loss of Equation (2).

In Reference [11], it is claimed that minimizing the inference KL yields more desirable results since for multi-modal distributions, individual modes can be fitted with a mono-modal distribution such as a Gaussian distribution, whereas the resulting distribution has a very large variance when minimizing the approximation KL to account for all modes. In Figure 1 there is an example of this behavior. The true distribution of the quantity s is taken to be a mixture of two standard Gaussians with means ± 3 . It is approximated with one Gaussian distribution by using the approximation KL and the inference KL. When using the approximation KL, the resulting distribution has a large variance to cover both peaks. Minimizing the inference KL leads to a sharply peaked approximation around one

peak. A user of this method might be very confident that the value of s must be near 3, even though the result is heavily dependent on the initial condition of the minimization and could have become peaked around -3 just as well.

We find that fitting a multi-modal distribution with a mono-modal one will yield suboptimal results irrespective of the fitting scheme. An approximation should always have the goal to be close to the target that is being approximated. If it is already apparent that this goal cannot be achieved, it is recommended to rethink the set of approximative distributions and not dwell on the algorithm used for approximation.

In Reference [12] an approximative simulation scheme called information field dynamics is described. There, a Gaussian distribution q is matched to a time-evolved version $U(p)$ of a Gaussian distribution p . This matching is done by minimizing the inference KL. In this particular case (at least for information preserving dynamics), the matching can be made optimal without making the algorithm more complicated. Since for information preserving dynamics time evolution is just a change of coordinates and the Kullback–Leibler divergence is invariant under such transformations, one can instead match the Gaussian distribution p and $U^{-1}(q)$ by minimizing $KL(p, U^{-1}(q)) = KL(U(p), q)$, which is just as difficult in terms of computation.

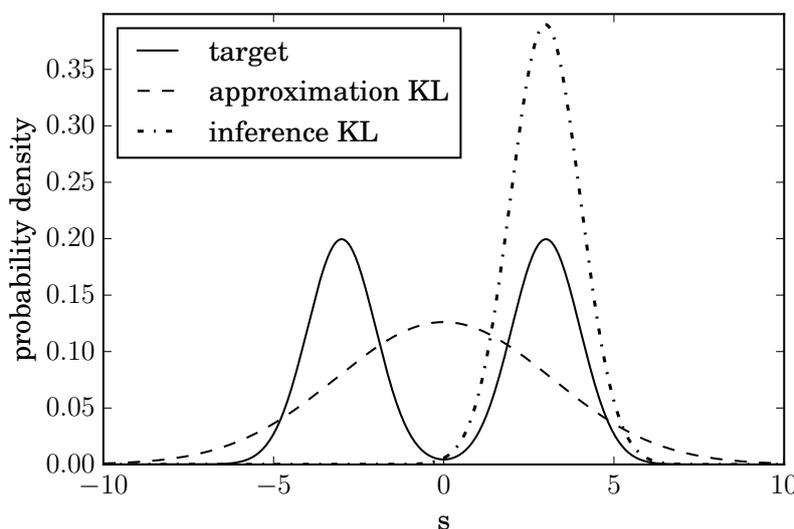


Figure 1. Results of approximating a target distribution in s with a Gaussian distribution. KL: Kullback–Leibler.

In Reference [13] it is claimed that the inference KL yields an optimal approximation scheme fulfilling certain axioms. This result is the exact opposite of our result. This disagreement is due to an assumed consistency of approximations. In Reference [13], further approximations are forced to be consistent with earlier approximations; i.e., if one does two approximations, one gets the same result as with one joint approximation. Due to this requirement, the derived functional cannot satisfy some of the axioms that we used. In our picture, it is better to do one large approximation instead of many small approximations. This is in accordance to the behavior of other approximations. For example, when step-wise rounding the real number 1.49, one gets 2 if it is first rounded to one decimal and then to integer precision compared to being rounded to integer precision directly where one gets 1. If information is lost due to approximation, it is natural for further approximations to be less precise than if one were to approximate in one go.

There also exist cases where we could not find any comments explaining why the arguments of the Kullback–Leibler divergence are in that particular order. In general, it would be desirable that authors provide a short argumentation for why they choose a particular order of the arguments of the KL divergence.

5. Conclusions

Using the two elementary consistency requirements on locality and optimality, as expressed by Equations (5) and (6), respectively, we have shown that there is only one ranking function that ranks how good an approximation of a belief is, analogously to Reference [14]. By minimizing $KL(p, q)$ with respect to its second argument $q \in X$, one gets the best approximation to p . This is claimed at several points in the literature. Nevertheless, we found sources where other functionals were minimized in order to obtain an approximation. This confusion is probably due to the fact that for the slightly different task of updating a belief q under new constraints, $KL(p, q)$ has to be minimized with respect to p , its first argument [29,30]. We do not claim that any of the direction of Kullback–Leibler divergence are wrong by themselves, but one should be careful of when to use which.

We hope that for the case of approximating a probability distribution p by another q we have given convincing and conclusive arguments for why this should be done by minimizing $KL(p, q)$ with respect to q , if this is feasible.

Acknowledgments: We would like to thank A. Caticha, J. Skilling, V. Böhm, J. Knollmüller, N. Porqueres, and M. Greiner and six anonymous referees for the discussions and their valuable comments on the manuscript.

Author Contributions: Reimar H. Leike conceived the idea and wrote the paper, Torsten A. Enßlin provided guidance and constructive feedback on the theory and the paper. Both authors have read and approved the final manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Kullback, S.; Leibler, R. On Information and Sufficiency. *Ann. Math. Stat.* **1951**, *22*, 79–86.
2. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*, 2nd ed.; John Wiley & Sons: Hoboken, NJ, USA, 2006.
3. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, MA, USA, 2016.
4. Opper, M.; Saad, D. *Advanced Mean Field Methods: Theory And Practice*; MIT Press: Cambridge, MA, USA, 2010.
5. Skilling, J. Bayesian Inference and maximum entropy methods in science and engineering. In Proceedings of the 33rd International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering (MaxEnt 2013), Canberra, ACT, Australia, 15–20 December 2013; Volume 1636, pp. 24–29.
6. Skilling, J. Bayesian inference and maximum entropy methods in science and engineering (MAXENT 2014). In Proceedings of the 34th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering (MAXENT 2014), Château Clos Lucé, Parc Leonardo Da Vinci, Amboise, France, 21–26 September 2014; Volume 1641, pp. 17–42.
7. Fox, C.W.; Roberts, S.J. A tutorial on variational Bayesian inference. *Artif. Intell. Rev.* **2012**, *38*, 85–95.
8. Enßlin, T.A.; Weig, C. Inference with minimal Gibbs free energy in information field theory. *Phys. Rev.* **2010**, *82*, 051112.
9. Blei, D.M.; Kucukelbir, A.; McAuliffe, J.D. Variational Inference: A Review for Statisticians. *arXiv* **2016**, arXiv:1601.00670.
10. Pinski, F.J.; Simpson, G.; Stuart, A.M.; Weber, H. Algorithms for Kullback-Leibler Approximation of Probability Measures in Infinite Dimensions. *arXiv* **2014**, arXiv:1408.1920.
11. Pinski, F.; Simpson, G.; Stuart, A.; Weber, H. Kullback-Leibler Approximation for Probability Measures on Infinite Dimensional Spaces. *arXiv* **2013**, arXiv:1310.7845.
12. Enßlin, T.A. Information field dynamics for simulation scheme construction. *Phys. Rev.* **2013**, *87*, 013308.
13. Tseng, C.-Y.; Caticha, A. Using Relative Entropy to Find Optimal Approximations: An Application to Simple Fluids. *Phys. Stat. Mech. Appl.* **2008**, *387*, 6759–6770.
14. Bernardo, J.M. Expected Information as Expected Utility. *Ann. Stat.* **1979**, *7*, 686–690.
15. Aczél, J.; Pfanzagl, J. Remarks on the Measurement of Subjective Probability and Information. *Metrika* **1967**, *11*, 91–105.
16. McCarthy, J. Measures of the value of information. *Proc. Natl. Acad. Sci. USA* **1956**, *42*, 654–655.
17. Good, I.J. Rational Decisions. *J. R. Stat. Soc. Ser. B* **1952**, *14*, 107–114.
18. Harremoës, P. Divergence and Sufficiency for Convex Optimization. *arXiv* **2017**, arXiv:1701.01010.

19. Cramér, H. *On the Mathematical Theory of Risk*; Centraltryckeriet: Alingsås, Sweden, 1930.
20. Cox, R.T. Probability, Frequency and Reasonable Expectation. *Am. J. Phys.* **1946**, *14*, 1.
21. Jaynes, E.T. *Probability Theory*; Larry Bretthorst, G., Ed.; Cambridge University Press: Cambridge, UK, 2003; p. 758, ISBN 0521592712.
22. Skilling, J. *Maximum-Entropy and Bayesian Methods in Science and Engineering*; Springer: Berlin/Heidelberg, Germany, 1988; pp. 173–187.
23. Caticha, A. *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*; Erickson, G.J., Zhai, Y., Eds.; American Institute of Physics Conference Series; American Institute of Physics: College Park, MD, USA, 2004; pp. 75–96.
24. Brier, G.W. Verification of forecasts expressed in terms of probability. *Mon. Weather Rev.* **1950**, *78*, 1–3.
25. Skilling, J. *Maximum Entropy and Bayesian Methods*; Springer: Berlin/Heidelberg, Germany, 1989; pp. 45–52.
26. Gneiting, T.; Raftery, A.E. Strictly Proper Scoring Rules, Prediction, and Estimation. *J. Am. Stat. Assoc.* **2007**, *102*, 359–378.
27. Bishop, C. *Pattern Recognition and Machine Learning (Information Science and Statistics)*; Springer: New York, NY, USA, 2007.
28. Dempster, A.P.; Laird, N.M.; Rubin, D.B. Maximum Likelihood from Incomplete Data via the EM Algorithm. *J. R. Stat. Soc. Ser. B* **1977**, *39*, 1–38.
29. Csiszar, I. Why Least Squares and Maximum Entropy? An Axiomatic Approach to Inference for Linear Inverse Problems. *Ann. Stat.* **1991**, *19*, 2032–2066.
30. Caticha, A. Towards an Informational Pragmatic Realism. *arXiv* **2014**, arXiv:1412.5644.



© 2017 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).