

Variational Predictive Information Bottleneck

Alexander A. Alemi
Google Research
alemi@google.com

October 18, 2019

Abstract

In classic papers, Zellner [1988, 2002] demonstrated that Bayesian inference could be derived as the solution to an information theoretic functional. Below we *derive* a generalized form of this functional as a variational lower bound of a predictive information bottleneck objective. This generalized functional encompasses most modern inference procedures and suggests novel ones.

1 Introduction

Consider a data generating process $\phi \sim p(\phi)$ from which we have some N draws that constitute our training set, $\mathbf{x}_P = \{x_1, x_2, \dots, x_N\} \sim p(x|\phi)$. We can also imagine (potentially infinitely many) future draws from this same process $\mathbf{x}_F = \{x_{N+1}, \dots\} \sim p(x|\phi)$. The *predictive information* $I(\mathbf{x}_P; \mathbf{x}_F)$ ¹ gives a unique measure of the complexity of a data generating process [Bialek et al., 2001].




Figure 1: Graphical model under consideration.

The goal of *learning* is to capture this complexity. To perform learning, we form a global representation of the dataset $p(\theta|\mathbf{x}_P)$. This can be thought of as a learning algorithm, that, given a set of observations, produces a summary statistic of the dataset that we hope is useful for predicting future draws from the same process. This algorithm could be deterministic or more generally, stochastic.

¹We use $I(x; y)$ for the *mutual information* between two random variables: $I(x; y) \equiv \mathbb{E}_{p(x,y)} \left[\log \frac{p(x,y)}{p(x)p(y)} \right]$

For example, imagine training a neural network on some data with stochastic gradient descent. Here the training data would be \mathbf{x}_P , the test data \mathbf{x}_F and the neural network parameters would be θ . Our training procedure implicitly samples from the distribution $p(\theta|\mathbf{x}_P)$.

How do we judge the utility of this learned global representation? The mutual information $I(\theta; \mathbf{x}_F)$ quantifies the amount of information our representation captures about future draws.² To maximize learning we therefore aim to maximize this quantity.

This is, of course, only interesting if we constrain how expressive our global representation is, for otherwise we could simply retain the full dataset. The amount of information retained about the observed data: $I(\theta; \mathbf{x}_P)$ is a direct measure of our representation’s complexity. The bits a learner extracts from data provides upper bounds on generalization [Bassily et al., 2017].

2 Predictive Information Bottleneck

Combined, these motivate the *predictive information bottleneck objective*, a generalized *information bottleneck* [Bialek et al., 2001, Tishby et al., 2000]:

$$\max_{p(\theta|\mathbf{x}_P)} I(\theta; \mathbf{x}_F) \quad \text{s.t.} \quad I(\theta; \mathbf{x}_P) = I_0. \quad (1)$$

We can turn this into an unconstrained optimization problem with the use of a Lagrange multiplier β :

$$\max_{p(\theta|\mathbf{x}_P)} I(\theta; \mathbf{x}_F) - (1 - \beta)I(\theta; \mathbf{x}_P). \quad (2)$$

While this objective seems wholly out of reach, we can make progress by noting that our random variables satisfy the Markov chain: $\mathbf{x}_F \leftarrow \phi \rightarrow \mathbf{x}_P \rightarrow \theta$, in which θ and \mathbf{x}_F are conditionally independent given \mathbf{x}_P :

$$I(\theta; \mathbf{x}_F, \mathbf{x}_P) = I(\theta; \mathbf{x}_F) + I(\theta; \mathbf{x}_P|\mathbf{x}_F) = I(\theta; \mathbf{x}_P) + \underline{I(\theta; \mathbf{x}_F|\mathbf{x}_P)}. \quad (3)$$

This implies:

$$I(\theta; \mathbf{x}_F) = I(\theta; \mathbf{x}_P) - I(\theta; \mathbf{x}_P|\mathbf{x}_F). \quad (4)$$

and the equivalent unconstrained optimization problem:³

$$\min_{p(\theta|\mathbf{x}_P)} I(\theta; \mathbf{x}_P|\mathbf{x}_F) - \beta I(\theta; \mathbf{x}_P). \quad (5)$$

The first term here: $I(\theta; \mathbf{x}_P|\mathbf{x}_F)$ is the residual information between our global representation and the dataset after we condition on full knowledge of the data generating procedure. This is a direct measure of the inefficiency of our proposed representation.

²It is interesting to note that in the limit of an infinite number of future draws, $I(\theta; \mathbf{x}_F)$ approaches $I(\theta; \phi)$. Therefore, the amount of information we have about an infinite number of future draws from the process is the same as the amount of information we have about the nature and identity of the data generating process itself.

³A similar transformation for the (local) variational information bottleneck appeared in Fischer [2019].

3 Variational Predictive Information Bottleneck

Simple variational bounds [Poole et al., 2019] can be derived for this objective, just as was done for the (local) information bottleneck objective in Alemi et al. [2016]. First, we demonstrate a variational upper bound on $I(\theta; \mathbf{x}_P | \mathbf{x}_F)$:⁴

$$I(\theta; \mathbf{x}_P | \mathbf{x}_F) = \left\langle \log \frac{p(\theta | \mathbf{x}_P)}{p(\theta | \mathbf{x}_F)} \right\rangle \leq \left\langle \log \frac{p(\theta | \mathbf{x}_P)}{q(\theta)} \right\rangle. \quad (6)$$

Here we upper bound the residual information by using a variational approximation to $p(\theta | \mathbf{x}_F)$, the marginal of our global representation over all datasets drawn from the same data generating procedure. Any distribution $q(\theta)$ independent of \mathbf{x}_F suffices.

Next we variationally lower bound $I(\theta; \mathbf{x}_P)$ with:

$$I(\theta; \mathbf{x}_P) = \left\langle \log \frac{p(\mathbf{x}_P | \theta)}{p(\mathbf{x}_P)} \right\rangle \geq H(\mathbf{x}_P) + \sum_i \langle \log q(x_i | \theta) \rangle. \quad (7)$$

The entropy of the training data $H(\mathbf{x}_P)$ is a constant outside of our control that can be ignored. Here we variationally approximate the “posterior” of our global representation with a factorized “likelihood”: $\prod_i q(x_i | \theta) = q(\mathbf{x}_P | \theta) \sim p(\mathbf{x}_P | \theta)$. Notice that while $p(\mathbf{x}_P | \theta)$ will not factorize in general, we can certainly consider a family of variational approximations that do.

Combining these variational bounds, we generate the objective:

$$\min_{p(\theta | \mathbf{x}_P)} \left\langle \log \frac{p(\theta | \mathbf{x}_P)}{q(\theta)} - \beta \sum_i \log q(x_i | \theta) \right\rangle. \quad (8)$$

We have thus *derived*, as a variational lower bound on the predictive information bottleneck, the objective Zellner [1988] postulates (with $\beta = 1$) is satisfied for inference procedures that optimally process information. As Knoblauch et al. [2019] demonstrates, this encompasses a wide array of modern inference procedures, including Generalized Bayesian Inference [Bissiri et al., 2016] and a generalized Variational Inference, dubbed Gibbs VI [Alquier et al., 2016, Futami et al., 2017].⁵ Below we highlight some of these and other connections.

4 Connections

If, in Equation (8), we identity $q(\theta)$ with a *fixed prior* and $q(x | \theta)$ with a *fixed likelihood* of a generative model, optimizing this objective for $p(\theta | \mathbf{x}_P)$ in the space

⁴ $\langle \cdot \rangle$ is used to denote expectations, and unless denoted otherwise with respect to the full joint density $p(\theta | \mathbf{x}_P) p(\mathbf{x}_P | \phi) p(\phi) p(\mathbf{x}_F | \phi)$

⁵To incorporate the Generalized VI [Knoblauch et al., 2019] with divergence measures other than KL, we need only replace our mutual informations (which are KL based) with their corresponding generalizations.

of all probability densities gives the generalized Boltzmann distribution [Jaynes, 1957]:

$$p(\theta|\mathbf{x}_{\mathcal{P}}) = q(\theta) \left[\prod_i q(x_i|\theta) \right]^\beta / \mathcal{Z}, \quad (9)$$

where \mathcal{Z} is the *partition function*.⁶ This is a generalized form of Bayesian Inference called the *power likelihood* [Holmes and Walker, 2017, Royall and Tsou, 2003]. Here the inverse temperature β acts as a Lagrange multiplier controlling the trade-off between the amount of information we retain about our observed data ($I(\theta; \mathbf{x}_{\mathcal{P}})$) and how much predictive information we capture ($I(\theta; \mathbf{x}_{\mathcal{F}})$). As $\beta \rightarrow \infty$ (temperature goes to zero), we recover the maximum likelihood solution. At $\beta = 1$ (temperature = 1) we recover ordinary Bayesian inference. As $\beta \rightarrow 0$ (temperature goes to infinity), we recover just prior predictive inference that ignores the data entirely. These limits are summarized in Table 1.

Limit	Inference	Equivalent Objective	$p(\theta \mathbf{x}_{\mathcal{P}})$
β	Generalized Bayes	$\max I(\theta; \mathbf{x}_{\mathcal{F}}) - (1 - \beta)I(\theta; \mathbf{x}_{\mathcal{P}})$	$\propto q(\theta)q(\mathbf{x}_{\mathcal{P}} \theta)^\beta$
$\beta \rightarrow 0$	Prior Predictive	$\min I(\theta; \mathbf{x}_{\mathcal{P}} \mathbf{x}_{\mathcal{F}})$	$q(\theta)$
$\beta \rightarrow 1$	Bayesian	$\max I(\theta; \mathbf{x}_{\mathcal{F}})$	$\propto q(\theta)q(\mathbf{x}_{\mathcal{P}} \theta)$
$\beta \rightarrow \infty$	Maximum Likelihood	$\max I(\theta; \mathbf{x}_{\mathcal{P}})$	$\arg \min_{\theta} q(\mathbf{x}_{\mathcal{P}} \theta)$

Table 1: Power Bayes can be recovered as a variational lower bound on the predictive information bottleneck objective (Equation (5)).

More generally, notice that in Equation (8) the densities $q(x|\theta)$ and $q(\theta)$ are not literally the likelihood and prior of a generative model, they are variational approximations that we have complete freedom to specify. This allows us to describe other more generalized forms of Bayesian inference such as Divergence Bayes or the full Generalized Bayes [Knoblauch et al., 2019, Bissiri et al., 2016] provided we can interpret the chosen loss function as a conditional distribution.

If we limit the domain of $p(\theta|\mathbf{x}_{\mathcal{P}})$ to a restricted family of parametric distributions, we immediately recover not only standard variational inference, but a broad generalization known as Gibbs Variational Inference [Knoblauch et al., 2019, Alquier et al., 2016, Futami et al., 2017].

Furthermore, nothing prevents us from making $q(x|\theta)$ or $q(\theta)$ themselves parametric and simultaneously optimizing those. Optimizing the prior with a fixed likelihood, unconstrained $p(\theta|\mathbf{x}_{\mathcal{P}})$, and $\beta = 1$ the objective mirrors Empirical Bayesian [Maritz and Lwin, 2018] approaches, including the notion of *reference priors* [Mattingly et al., 2018, Berger et al., 2009]. Alternatively, optimizing a parametric likelihood with a parametric representation $p(\theta|\mathbf{x}_{\mathcal{P}})$, fixed prior, and $\beta = 1$ equates to a Neural Process [Garnelo et al., 2018].

Consider next *data augmentation*, where we have some stochastic process that modifies our data with implicit conditional density $t(x'|x)$. If the augmentation

⁶ $\mathcal{Z} \equiv \int d\theta q(\theta) \left[\prod_i q(x_i|\theta) \right]^\beta$

procedure is centered about zero so that $\langle x' \rangle_{t(x'|x)} = x$ and our chosen likelihood function is concave, then we have:

$$\log q(x|\theta) = \log q(\langle x' \rangle_{x' \sim t(x'|x)}|\theta) \geq \langle \log q(x'|\theta) \rangle_{x' \sim t(x'|x)}, \quad (10)$$

which maintains our bound. For example, for an exponential family likelihood and any centered augmentation procedure (like additive mean zero noise), doing generalized Bayesian inference on an augmented dataset is also a lower bound on the predictive information bottleneck objective.

5 Conclusion and Future Work

We have shown that a wide range of existing inference techniques are variational lower bounds on a single predictive information bottleneck objective.

This connection highlights the drawbacks of these traditional forms of inference. In all cases considered in the previous section, we made two choices that loosened our variational bounds. First, we approximated $p(\mathbf{x}_P|\theta)$, with a factorized approximation $q(\mathbf{x}_P|\theta) = \prod_i q(x_i|\theta)$. Second, we approximated the future conditional *marginal* $p(\theta|\mathbf{x}_F) = \int d\mathbf{x}_P p(\theta|\mathbf{x}_P)p(\mathbf{x}_P|\mathbf{x}_F)$ as an unconditional “prior”. Neither of these approximations is necessary.

For example, consider the following tighter “prior”:

$$q(\theta|\mathbf{x}_F) \sim \int d\mathbf{x}_P' p(\theta|\mathbf{x}_P')q(\mathbf{x}_P'|\mathbf{x}_F). \quad (11)$$

Here we reuse a tractable global representation $p(\theta|\mathbf{x}_P)$ and instead create a variational approximation to the density of alternative datasets drawn from the same process: $q(\mathbf{x}_P'|\mathbf{x}_F)$.

We believe this information-theoretic, representation-first perspective on learning has the potential to motivate new and better forms of inference.⁷

References

- Alexander A. Alemi and Ian Fischer. TherML: Thermodynamics of machine learning, 2018.
- Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.
- Pierre Alquier, James Ridgway, and Nicolas Chopin. On the properties of variational approximations of gibbs posteriors. *Journal of Machine Learning Research*, 17(236):1–41, 2016. URL <http://jmlr.org/papers/v17/15-290.html>.

⁷Extensions of this representation-first perspective beyond a single global representation towards supervised or unsupervised learning can be found in Alemi and Fischer [2018].

- Raef Bassily, Shay Moran, Ido Nachum, Jonathan Shafer, and Amir Yehudayoff. Learners that use little information. *arXiv preprint arXiv:1710.05233*, 2017.
- James O Berger, José M Bernardo, Dongchu Sun, et al. The formal definition of reference priors. *The Annals of Statistics*, 37(2):905–938, 2009.
- William Bialek, Ilya Nemenman, and Naftali Tishby. Predictability, complexity, and learning. *Neural computation*, 13(11):2409–2463, 2001.
- P. G. Bissiri, C. C. Holmes, and S. G. Walker. A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):1103–1130, 2016. doi: 10.1111/rssb.12158. URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/rssb.12158>.
- Ian Fischer. The conditional entropy bottleneck, 2019. URL <https://openreview.net/forum?id=rkVOXhAqY7>.
- Futoshi Futami, Issei Sato, and Masashi Sugiyama. Variational inference based on robust divergences. *arXiv preprint arXiv:1710.06595*, 2017.
- Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018.
- CC Holmes and SG Walker. Assigning a value to a power likelihood in a general bayesian model. *Biometrika*, 104(2):497–503, 2017.
- Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.
- Jeremias Knoblauch, Jack Jewson, and Theodoros Damoulas. Generalized variational inference. *arXiv preprint arXiv:1904.02063*, 2019.
- Johannes S Maritz and T Lwin. *Empirical bayes methods*. Routledge, 2018.
- Henry H Mattingly, Mark K Transtrum, Michael C Abbott, and Benjamin B Machta. Maximizing the information learned from finite data selects a simple model. *Proceedings of the National Academy of Sciences*, 115(8):1760–1765, 2018.
- Ben Poole, Sherjil Ozair, Aaron van den Oord, Alexander A Alemi, and George Tucker. On variational bounds of mutual information. *arXiv preprint arXiv:1905.06922*, 2019.
- Richard Royall and Tsung-Shan Tsou. Interpreting statistical evidence by using imperfect models: robust adjusted likelihood functions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):391–404, 2003.
- Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.

Arnold Zellner. Optimal information processing and bayes's theorem. *The American Statistician*, 42(4):278–280, 1988.

Arnold Zellner. Information processing and bayesian analysis. *Journal of Econometrics*, 107(1-2):41–50, 2002.