

Correcting Predictions for Approximate Bayesian Inference

Tomasz Kuśmierczyk, Joseph Sakaya, Arto Klami

Helsinki Institute for Information Technology HIIT

Department of Computer Science, University of Helsinki

{tomasz.kusmierczyk, joseph.sakaya, arto.klami}@helsinki.fi

Abstract

Bayesian models quantify uncertainty and facilitate optimal decision-making in downstream applications. For most models, however, practitioners are forced to use approximate inference techniques that lead to sub-optimal decisions due to incorrect posterior predictive distributions. We present a novel approach that corrects for inaccuracies in posterior inference by altering the decision-making process. We train a separate model to make optimal decisions under the approximate posterior, combining interpretable Bayesian modeling with optimization of direct predictive accuracy in a principled fashion. The solution is generally applicable as a plug-in module for predictive decision-making for arbitrary probabilistic programs, irrespective of the posterior inference strategy. We demonstrate the approach empirically in several problems, confirming its potential.

1 Introduction

Bayesian inference provides the fundamental basis for modeling uncertainty. The posterior distribution $p(\theta|\mathcal{D})$ provides a complete summary of what is known about the parameters θ of a model $p(y, \theta)$ given some observed data $\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N$. In particular, the posterior is necessary and sufficient information for making optimal decisions under uncertainty (Berger 1985).

For predictive models $p(y|x)$ (conditioning on \mathcal{D} is assumed implicitly), such as regression or classification, the decisions correspond to committing to individual decision h that minimizes the risk \mathcal{R} , the expectation of a loss $\ell(y, h)$ over the uncertainty in the model parameters. For example, to minimize the squared loss we should commit to the mean of the predictive distribution $p(y|x) = \int p(y|x, \theta)p(\theta|\mathcal{D})d\theta$, whereas for a loss that penalizes for overestimation of the true value the optimal decision is a suitably chosen lower quantile of the predictive distribution. For many losses the Bayes optimal decision is a simple statistic of the predictive distribution (Table 1).

Unfortunately, the decisions obtained by minimizing the risk are optimal only if we have access to the true posterior $p(\theta|\mathcal{D})$. For most models we can only obtain an approximation $q(\theta) \approx p(\theta|\mathcal{D})$. The almost universally adopted practice is to plug in the approximation in lieu of the posterior in

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

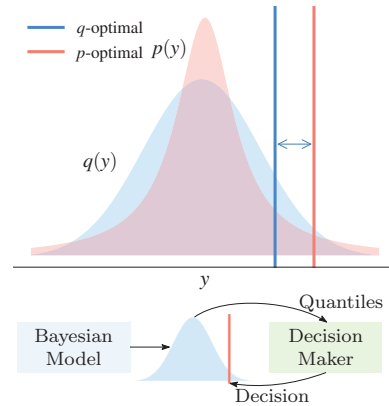


Figure 1: Bayesian decision theory gives a closed-form analytic rule for making optimal decisions for a given loss, based on the posterior predictive distribution (red). If the same decision rule is used for approximate predictive distribution (blue) based on approximate inference strategy, the numerical decision may be badly off as indicated by the gap ϵ . Our approach learns to correct for this mistake, by replacing the theory-based decisions with parametric decision-making module – a neural network – that maps quantiles of the incorrect predictive distribution into optimal decisions.

the predictive distribution and proceed as though the result was the true predictive distribution. Even though this works relatively well for good approximations (e.g., well-converged Markov Chain Monte Carlo sampler), it may fail miserably for bad approximations. Furthermore, even for good approximations it would make sense to account for common biases; e.g., mean-field variational approximations are known to underestimate the posterior variation and hence also the variation in the predictive distribution (Blei, Kucukelbir, and McAuliffe 2017). Consequently, using decision rules optimal for the true posterior will have systematic bias as well.

Even though the problem persists in virtually all use of approximate Bayesian inference in decision-making problems, the literature addressing the issue is limited. The best examples build on the concept of *loss-calibrated inference*. Originally proposed in the context of variational approximation by Lacoste-Julien, Huszár, and Ghahramani (2011), the core idea is to account for the eventual decision problem (charac-

terized by some loss) during inference by altering the learning objective for approximate inference. Recently Cobb, Roberts, and Gal (2018) and Kuśmierczyk, Sakaya, and Klami (2019) have presented practical algorithms building on this principle, but despite the elegant theory the procedure runs into several practical difficulties: It is computationally heavy, requires knowing the decision loss during inference, and the empirical improvements are typically meagre. A similar approach has been proposed also for MCMC (Abbasnejad, Domke, and Sanner 2015), with similar drawbacks.

We propose an alternative, novel approach for making optimal decisions based on a probabilistic model characterized by an approximate posterior. Instead of modifying the *inference procedure* as in loss-calibrated inference, we modify the *decision making phase* by replacing the analytic Bayes optimal decisions with a parametric function tuned to minimize the risk for the approximate posterior at hand. The proposed approach is conceptually simple, can be applied on top of any approximate posterior irrespective of the manner it was obtained, does not require re-computing the posterior approximation for new losses, and it is demonstrated empirically to clearly outperform the loss-calibrated inference approach (Kuśmierczyk, Sakaya, and Klami 2019) in some scenarios. When applied on the true posterior the procedure reverts back to the standard Bayes optimal decision.

Our key idea is to couple the Bayesian model with a decision-making module, in practice a neural network, which takes as input a characterization of the predictive distribution $p(y|x)$ and outputs the decision h . It mimics optimal Bayesian decision rules by using the same input sufficient for optimal decisions, but is optimized to minimize a particular risk. The principle is illustrated in Figure 1.

Training a flexible mapping from the predictive distribution to the decisions by direct optimization of the empirical risk can, naturally, overfit to the specific data collection. Furthermore, a flexible mapping may, under some conditions, overrule the underlying model, by making predictions not supported by the assumed model. We overcome these issues by presenting a generalized Bayesian inference strategy, building on Bissiri, Holmes, and Walker (2016), over the decision-making modules. We provide a justified prior distribution that controls the decision-making module, by regularizing it towards the standard Bayes optimal decision, and demonstrate automatic prior specification using a light bootstrap procedure (Efron and Tibshirani 1994).

Our result is a plug-in tool that is applicable for every Bayesian model designed for predictive tasks. One can use any modeling framework or inference strategy, as long as we have access to samples from the predictive distribution. We demonstrate the approach for standard Bayesian models ranging from matrix factorization to sparse regression, and show improved decisions for poor approximations.

2 Background

To facilitate understanding the rest of the paper, we briefly summarize Bayesian decision theory, the concept of approximate posterior inference, and the standard approach for incorporating approximate posteriors into Bayesian decision making.

Table 1: Example losses and their p -optimal decisions.

LOSS	EXPRESSION	DECISION RULE
squared	$(h - y)^2$	$\text{mean}_p[y]$
absolute	$ h - y $	$\text{median}_p[y]$
imbalanced absolute	$\begin{cases} a \cdot h - y & y \geq h \\ b \cdot h - y & y < h \end{cases}$	$\frac{a}{a+b}$ -quantile $_p[y]$
tilted	$\begin{cases} t \cdot h - y & y \geq h \\ (1 - t) \cdot h - y & y < h \end{cases}$	t -quantile $_p[y]$

2.1 Bayesian Decision Theory

Bayesian decision theory (Berger 1985) defines a rigorous framework for decision-making under uncertainty. Given a loss function $\ell(\theta, h)$ and a posterior distribution $p(\theta|\mathcal{D})$ of a parametric model conditioned on the dataset \mathcal{D} , the Bayes optimal decision h_p minimizes expected posterior loss $\mathcal{R}_p(h) = \int \ell(\theta, h)p(\theta|\mathcal{D}) d\theta$.

In supervised settings, decisions are made for pairs (x, y) , and the risk $\mathcal{R} = \mathbb{E}_{(x,y) \sim p_{\mathcal{D}}} \ell(y, h(x))$ is an expectation over unknown data generating distribution $p_{\mathcal{D}}$. For setups where we do not explicitly model $p(x)$ but do assume a model for $p(y|x)$, we can define conditional risk for x using

$$\mathcal{R}_p(h|x) = \int \ell(y, h)p(y|\mathcal{D}, x)dy. \quad (1)$$

where $p(y|\mathcal{D}, x) = \int p(y|\theta, x)p(\theta|\mathcal{D})d\theta$. The optimal decisions for individual data points are then $h_p = \arg \min_{h \in \mathcal{H}} \mathcal{R}_p(h|x)$, denoted as p -optimal to indicate they are optimal for the true predictive distribution. Similarly, we call (1) the p -risk or simply risk. Table 1 presents p -optimal decisions for several example losses.

In practice the risk needs to be approximated by *empirical risk* $\mathcal{E}\mathcal{R} = \sum_{n=1}^N \ell(y_n, h(x_n))$, based on some data collection, where the distribution of x and y is purely empirical. More detailed discussion of Bayesian decision theory for supervised learning can be found in Lacoste-Julien, Huszár, and Ghahramani (2011).

2.2 Approximate Inference

Typically, the posterior distribution $p(\theta|\mathcal{D})$ is analytically intractable, and we need to resort to approximate inference techniques to replace it with a computationally tractable proxy $q(\theta)$. A wide array of inference tools are available with different trade-offs between computational time and efficiency. Monte Carlo Markov Chain (MCMC) methods approximate $p(\theta|\mathcal{D})$ with samples from the posterior, with state-of-the-art algorithms using gradient-based information to improve convergence and speed (Hoffman and Gelman 2014). Distributional approximations, in turn, assume a parametric approximation family $q(\theta; \alpha)$ and optimize for its parameters α to minimize some discrepancy measure between the approximation and the true posterior. This can be done in multitude of ways, the most common strategies being variational approximation (Blei, Kucukelbir, and McAuliffe 2017) and expectation propagation (Gelman et al. 2017; Minka 2001). Distributional approximations often have computational advantage over MCMC, but are biased or inaccurate especially when using poor approximation family (Yao et

al. 2018). For example, variational approximation minimizes a Kullback-Leibler divergence between the approximation and the true posterior, and hence underestimates posterior variation (Blei, Kucukelbir, and McAuliffe 2017).

Probabilistic programming tools, such as Stan (Carpenter et al. 2017) and Edward (Tran et al. 2016) have recently made practical Bayesian modeling easier, by coupling automatic inference engines based on Hamiltonian Monte Carlo (Hoffman and Gelman 2014) or gradient-based variational approximation (Titsias and Lázaro-Gredilla 2014; Kucukelbir et al. 2017) with easy model specification language. The typical mode of operation today is to specify the model in such a language, and the practitioner may not even care about the approximation strategy being used. This has made Bayesian modeling possible for wider audience, which implies more research is needed on how the models are being used in downstream applications.

2.3 Decision-making for Approximate Inference

We often do not have access to the true posterior $p(\theta|\mathcal{D})$ to evaluate the p -risk in Eq. 1. We instead use the approximation $q(\theta)$ as a proxy for the posterior. Following the nomenclature of Lacoste-Julien, Huszár, and Ghahramani (2011), we define the q -posterior predictive distribution $p_q(y|\mathcal{D}, x) = \int p(y|\theta, x)q(\theta) d\theta$ and the q -risk

$$\mathcal{R}_q(h|x) = \int \ell(y, h)p_q(y|\mathcal{D}, x)dy.$$

The q -optimal decisions $h_q = \arg \min_{h \in \mathcal{H}} \mathcal{R}_q(h|x)$ minimize the q -risk. Despite the name, the q -optimal decisions are not optimal with respect to the p -risk. Instead, we typically have $\mathcal{R}_p(h_q|x) > \mathcal{R}_p(h_p|x)$. Section 3 describes how the discrepancy between h_q and h_p can be mitigated, when only having access to the approximation $q(\theta)$.

3 Predictive Decision Theory

Our goal is to make good – possibly even optimal – decisions, as measured by the risk \mathcal{R} under a probabilistic model $p(y, \theta)$, when only having access to an approximation $q(\theta)$ of the true posterior $p(\theta|\mathcal{D})$. We do this by introducing a parametric decision-making module $f(p_q(y), \omega)$, a function that takes as input a characterization of the predictive distribution $p_q(y)$ under the approximate posterior and produces as an output the decision h . This module can be applied on top of arbitrary probabilistic models.

In the following, we describe the necessary technical elements required for implementing the proposed Predictive Decision Theory (PDT). We first formulate the problem as generalized Bayesian inference over the decision-makers in Section 3.1, and then explain how practical decision-making modules can be implemented in Section 3.2.

3.1 Decision Belief Distributions

To provide justified uncertainty quantification for the decisions, we build on the generalized Bayesian inference framework by Bissiri, Holmes, and Walker (2016). They present a coherent procedure for updating beliefs for scenarios where the parameter of interest is connected to observations via

arbitrary loss functions, instead of likelihood functions. Their key result is that (using generic notation to avoid confusion with the symbols used in main derivations of our results)

$$p(\eta|y) \propto e^{-\ell(\eta, y)}p(\eta)^{\lambda_1}$$

is a valid posterior distribution for any loss $\ell(\eta, y)$ that depends on the parameters η and the observation y , and any constant $\lambda_1 > 0$ allowing for calibration of relative impact of the prior and the loss. In other words, we can use exponentiated negative losses in place of likelihoods and still characterize the uncertainty on η in a coherent manner.

The framework provides direct basis for decision-problems involving the model parameters θ itself. Instead, we use it for predictive problems characterized by loss $\ell(h, y)$ defined in terms of predicted quantities, for which the decision h is made based on the predictive distribution $p(y)$ that is obtained via marginalization of model parameters θ . In particular, we consider *decision belief distributions*

$$p(h|y) \propto e^{-\ell(h, y)}p(h)^{\lambda_1}$$

and tie decisions for different y with parametric *decision-making module* (DM) $h = f(p_q(y), \omega)$. DM makes decisions on the observed data y , is parameterized by ω , and takes as input the predictive distribution $p(y)$ (or an approximation $p_q(y)$) from already trained model. This allows us to define the Bayesian update rule for belief distribution of a decision-making module

$$p(\omega|y) \propto e^{-\ell'(\omega, y)}p(\omega)^{\lambda_1}, \quad (2)$$

where $\ell'(\omega, y) := \ell(f(p_q(y), \omega), y)$ and $p(\omega)$ is some prior defined over the parameters of the DM. The constant $\lambda_1 > 0$ allows tuning the compromise, but importantly all values result in valid posterior.

3.2 Decision Maker

We replace the standard decision-making procedure explained in Section 2.3 with parametric function (decision maker) $f(p_q(y), \omega)$, for which posterior inference is characterized by Eq. (2). For a collection of N data instances the log-posterior is

$$\begin{aligned} \log p(\omega|\mathcal{D}) &= -\sum_{n=1}^N \ell'(\omega, y_n) + \lambda_1 \log p(\omega) + C_1 \\ &\propto -\underbrace{\frac{1}{N} \sum_{n=1}^N \ell'(\omega, y_n)}_{\mathcal{ER}} + \lambda \log p(\omega) + C, \quad (3) \end{aligned}$$

where C is a normalization constant that has no effect on inference and $\lambda := \frac{\lambda_1}{N}$. Empirical risk \mathcal{ER} appears naturally as a part of the objective. Suitable priors, however, help us to address the issue of generalization, to make sure the decisions are good for the true risk instead of just the empirical one. The formulation defines the basis for full posterior inference, but we focus on finding a single good decision-maker and hence resort to MAP estimation of Eq. (3) that is less computationally expensive.

Regularization with q -optimal Decisions We start by specifying the prior $p(\omega)$ used for controlling the flexibility of the DM, by building on the assumption that $q(\theta)$ is reasonably good approximation of $p(\theta|\mathcal{D})$. When they are identical the q -optimal decision h_q corresponds exactly to the p -optimal decision h_p , and for small approximation errors we would still expect to find the optimal decision in the neighborhood of the theoretical optimum.

Building on the above principle we define an implicit prior. Instead of placing a prior on the parameters ω directly, we place a prior on the decisions $h = f(p_q(y), \omega)$, which induces a prior on ω . For the decisions we assume the prior

$$h^n \sim \mathcal{N}(h_q^n, 1), \quad (4)$$

where h_q^n is q -optimal decision for n th data point and the variance is set to unit value as it can be subsumed to the λ parameter in Eq. (3). Small λ allows DM to deviate arbitrary far from h_q , whereas large λ forces DM to emulate q -optimal decisions. A good value can be selected, e.g., by cross-validation.

A more flexible prior is obtained by allowing point-specific variation with

$$h^n \sim \mathcal{N}(\mu_h^n, \sigma_h^n).$$

This prior requires an automatic procedure for setting μ_h^n and σ_h^n , which we do based on the expected variation caused by the underlying distribution $p(x)$. The decisions are made conditional on the covariates and the total risk averages over the data generating distribution $p_{\mathcal{D}}(x, y) \approx p(y|x)p(x)$, as briefly explained in Section 2.1. The natural variation in samples produced by the generating distribution provides a reasonable basis for estimating the allowed variation in decisions as well, and to estimate this we apply a bootstrap procedure (Efron and Tibshirani 1994).

With bootstrapping, we draw several new datasets \mathcal{D}_l ($l = 1 \dots L$), using $L = 5$ in our empirical experiments. For each dataset, we obtain the posterior approximation $q(\theta)$ and the q -optimal decisions h_{ql} , and estimate the parameters for each data point using $\hat{\mu}_h^n = \frac{1}{L} \sum_{l=1}^L h_{ql}^n$, $\hat{\sigma}_h^n = \sqrt{\frac{1}{L} \sum_{l=1}^L (h_{ql}^n - \hat{\mu}_h^n)^2}$. Finally, under independence assumption on h^n , we can define the prior for ω as

$$\log p(\omega) = \sum_{n=1}^N \log N(f(p_q(y_n), \omega) | \mu_h^n, \sigma_h^n), \quad (5)$$

where for a simpler case $(\mu_h^n, \sigma_h^n) := (h_q^n, 1)$ and for bootstrap-based prior $(\mu_h^n, \sigma_h^n) := (\hat{\mu}_h^n, \hat{\sigma}_h^n)$. The overall strength of the prior is in both cases controlled by λ .

Representation of the Predictive Distribution Bayesian decision theory states that the predictive distribution $p(y)$ is necessary and sufficient information for making optimal decisions. We retain this assumption when switching to parametric decision-makers, by using the predictive distribution as the sole input for the decision-maker. However, it is rarely available as analytic expression.

A practical and model-independent representation is obtained by sampling a collection of S data points y_s from the model $p(y, \theta)$ under the approximate posterior $q(\theta)$. Such a

collection can be summarized with suitable finite statistic that can then be passed as input for the decision-maker; we use B evenly-spaced empirical quantiles, but for example a histogram would work as well. This representation only requires the ability to sample from the predictive distribution, covering essentially all probabilistic models of interest. It could even generalize for simulator-based models that lack closed-form expression for likelihood but still enable sampling from the model (Gutmann and Corander 2016).

The parameters B and S influence the richness and accuracy of the input representation, but as illustrated empirically in Section 5.1 the procedure is not very sensitive to the choices. Since computing the representation is cheap we can safely use large S , and small B is enough because the decision-makers can implicitly interpolate between the available quantiles if needing more granularity.

Parametric Decision Makers Most Bayesian optimal decisions are simple summary statistics of the predictive distributions (Table 1), and simplified illustrations of the procedure like Figure 1 intuitively suggest that it may be enough to modify the specific threshold, e.g., by lowering or increasing the quantile for tilted loss to account for the discrepancy between the predictive distributions. Such a decision-maker is easy to implement and has limited room for overfitting, but we will later show empirically (Section 5.1) that it is not sufficient for notably improving the predictions. Instead, we need more flexible functions.

A natural choice for a flexible model-independent decision maker is to use a neural network, interpreted as arbitrary mapping from the predictive distribution $p(y)$ to the decision h . The decision-maker is parameterized by ω , which denotes collectively the set of all network weights. Such modules are easy to implement in modern machine learning platforms, allowing for flexible choice of network architectures. In our experiments, we use a simple feed-forward network with 3 hidden layers (with 20, 20 and 10 nodes) with ReLU activation and Adam optimizer with learning rate = 0.01, but note that the specific network details are not important. Instead, we simply need a network that is able to relatively flexibly convert a fixed-dimensional quantile representation of the predictive distribution into a scalar decision.

4 Related Work

Our work lies in the intersection of Bayesian modeling and learning theory. The general principle of minimizing the regularized risk for making decisions that generalize for the data-generating distribution is the cornerstone of standard learning theory (see, e.g., Hastie, Tibshirani, and Friedman (2009)). The current work does not utilize or require results in recent learning theory literature, but we note that there may be interesting connections useful for future work. Our method also loosely relates to work on quantile regression (Koenker and Hallock 2001) and the training of decision-makers to generate prediction intervals (Pearce et al. 2018). The main difference is that we build on existing Bayesian models and merely use a neural network to correct for approximation error, whereas they use black-box methods for the predictive distribution itself.




Figure 2: Matrix Factorization of *last.fm* data: (Left:) Predictive performance on test data, showing the proposed method clearly outperforms LCVI for tilted losses. (Middle:) Regularization effects on test and training data for squared and tilted ($t = 0.5$) loss, showing how the regularization completely removes overfitting. (Right:) Effects of number of samples and predictive distribution representation (number of quantiles) on test data performance for tilted ($t = 0.5$) loss.

The most closely related work is on direct loss-calibration of posterior approximations (Lacoste-Julien, Huszár, and Ghahramani 2011; Abbasnejad, Domke, and Sanner 2015). We propose an alternative for that approach and will empirically compare the proposed solution against the latest loss-calibration method (Kuśmierczyk, Sakaya, and Klami 2019), and hence describe here their approach briefly.

The core idea of loss-calibration is to replace the original learning objective, in case of variational approximation, a lower bound to the marginal likelihood $\log p(\mathcal{D})$, with an augmented objective that involves a separate term accounting for the loss

$$\underbrace{\mathbb{E}_{q(\theta; \alpha)} [\log p(\mathcal{D}, \theta) - \log q(\theta; \alpha)]}_{\text{lower bound to } \log p(\mathcal{D})} + \underbrace{\mathbb{E}_{q(\theta; \alpha)} [\log u(\theta, h)]}_{\text{utility term}}.$$

The function $u(\theta, h)$ is the loss $\ell(\theta, h)$ transformed to utility (Berger 1985). This augmented objective is maximized, typically with an alternating algorithm, with respect to both the parameters α of the approximation and the decisions h . The decisions influence the approximation and the approximation influence the decisions, and hence the optimization is tightly coupled. This implies that the procedure needs to be carried out separately for all losses of interest and the algorithmic details need to be derived for each approximation strategy separately. Our approach is applied on top of existing approximation and hence does not suffer from these limitations: Once the approximation is available, we can easily optimize for any loss without needing to re-iterate the inference.

5 Experiments

To illustrate and validate the procedure we conduct a series of experiments, using variational approximation as $q(\theta)$. We first compare the method against the alternative of calibrating the posterior inference to account for the loss for a matrix factorization model, and then demonstrate improved decisions for a sparse regression model and a multilevel model for cases with approximations. We also study practical decisions regarding the input representation and regularization. The code to reproduce our experiments is available online.¹

¹https://github.com/tkusmierczyk/correcting_approximate_bayesian_predictions

5.1 Predictive Performance

We compare our approach to Loss-calibrated Variational Inference (LCVI) by Kuśmierczyk, Sakaya, and Klami (2019), the closest alternative to our method, using the same simplified probabilistic matrix factorization model

$$Y \sim \mathcal{N}(ZW, \sigma_y), \quad W_{ik} \sim \mathcal{N}(0, \sigma_W), \quad Z_{kj} \sim \mathcal{N}(0, \sigma_z)$$

they used for demonstrating loss-calibration improves predictions over standard variational inference on a subset of *last.fm* data (Bertin-Mahieux et al. 2011). We use their code² to precisely reproduce the experiment, matching all of their modeling and approximation choices ($K = 20$, $\sigma = 10$, mean-field Gaussian approximation, and log-transformation for the count data).

Figure 2 (left) compares relative reduction of empirical risk ($\mathcal{R} = \frac{\mathcal{E}\mathcal{R}_{\text{VI}} - \mathcal{E}\mathcal{R}}{\mathcal{E}\mathcal{R}_{\text{VI}}}$, where $\mathcal{E}\mathcal{R}_{\text{VI}}$ is the risk of q -optimal decision on standard approximation) on the test data for the same four losses they used. For tilted losses the proposed approach using neural network as decision-maker dramatically outperforms LCVI, reaching 7 – 24% risk reduction depending on the t parameter of the loss. This demonstrates the proposed approach has clear practical value over the alternative of fine-tuning the approximation itself. The simpler decision-maker that only optimizes for the quantile is unable to improve the decisions, demonstrating that it is preferable to learn a more flexible mapping.

For squared loss even the neural decision-maker performs poorly, due to overfitting. Figure 2 (middle) illustrates how regularizing the decision-maker towards the q -optimal decisions using the prior (5) elegantly removes the overfitting. The method never achieves risk reduction, but instead converges to the standard q -optimal estimator. Consequently, there are scenarios for which the computationally more expensive LCVI is able to improve the predictions, whereas the proposed method is not.

5.2 Posterior Representation

In practical computation the posterior predictive distribution is presented using empirical quantiles of samples drawn from the predictive distribution (Section 3.2). The choice of the

²<https://github.com/tkusmierczyk/lcvi>

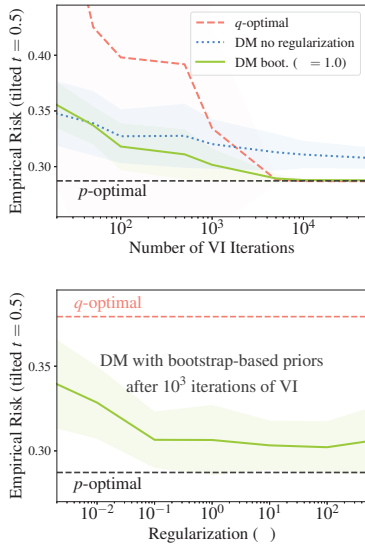


Figure 3: (Top:) Risk for tilted loss ($t = 0.5$) on test subset of *radon* data. The proposed method with regularization (solid green) is able to significantly reduce the risk over the q -optimal baseline (dashed red) for poor approximations. (Bottom:) The method is very insensitive to the regularization parameter λ ; all choices improve over the q -optimal baseline.

sample size S and the number of quantiles B naturally influences the representation accuracy. We study the effect of these parameters on the same data and model as above.

Figure 2 (right) compares test data empirical risk averaged over 10 different random initializations for various representations. Since it is beneficial to use large sample size, and the computational cost of drawing the samples is negligible compared to other stages of the procedure (fitting the approximation, training the decision-maker) we use $S = 1000$ in other experiments. The risk improves also when increasing the representation size, and for the other experiments we use $B = 20$ quantiles. However, it is worth noting that already very few quantiles are sufficient for improving the risk compared to q -optimal decisions (q -risk = 0.46).

5.3 Handling VI Failures

The main use for the method is in correcting for inaccurate posterior approximations. Hence, we conduct experiments on two separate regression models to validate this. At the same time, we illustrate the bootstrap-based regularization strategy.

Multilevel Models with Poor Approximation Since variational approximations are trained with iterative algorithms, we can construct a controlled experiment by terminating the inference algorithm early and varying the termination point. In general, approximations trained for shorter time are further away from the final approximation and the true posterior.

We use the *radon* data and the multi-level model

$$y_i = \alpha_{j[i]} + \beta x_i + \epsilon_i, \quad \alpha_j = \gamma_0 + \gamma_1 u_j + \zeta_j, \quad \zeta_j \sim \mathcal{N}(0, \sigma_\alpha^2)$$

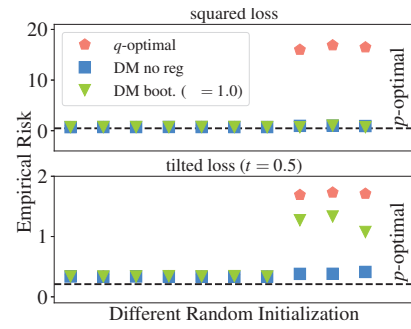


Figure 4: Risk for squared and tilted ($t = 0.5$) loss on test subset of *corn* data. For three initializations the approximation fails to converge and q -optimal decisions have very high risk, but the decision-maker (blue) is able to correct for the failure.

for modeling it (Gelman and Hill 2006). The model is implemented using the publicly available Stan code³, using the Minnesota subset of the data and conforms to its details. The inference is carried out using automatic differentiation variational inference (Kucukelbir et al. 2017), and we split the data randomly into equally sized training and test set.

Figure 3 (top) compares performance of different decision-making strategies (vertical axis) w.r.t quality of the approximation fit controlled by number of training iterations (horizontal axis). Hamiltonian Monte Carlo (Hoffman and Gelman 2014) provides here sufficiently accurate posterior to act as the baseline, and we see that variational approximation converges sufficiently close to the true posterior around 10^4 iterations, reaching the p -risk. The proposed method without regularization achieves relatively good risk already for very poor posterior approximations, but never converges to the optimal decision due to overfitting. The variant regularized with the bootstrap-based prior in Eq. (5) achieves the best of both worlds, reducing the risk for poor approximations but converging to the p -optimal decisions when the approximation becomes good. Figure 3 (bottom) plots the risk at 10^3 iterations (still incorrect approximation) as a function of the regularization parameter λ , showing that the method is robust to the choice of λ ; all values within the range improve compared to the q -optimal baseline, and for $\lambda \in [10^{-1}, 10^2]$ the risk is virtually identical.

Sparse Models with Failure of Convergence Variational approximations have trouble with multimodal or otherwise complicated posteriors. Yao et al. (2018) use models with horseshoe priors (Carvalho, Polson, and Scott 2010) for demonstrating this, and we use the same model family to demonstrate how we can still recover good decisions even when this happens. We use the sparse regression model

$$y|\beta \sim \mathcal{N}(X\beta, \sigma), \quad \beta_j|\tau, \lambda, c \sim \mathcal{N}(0, \tau^2 \tilde{\lambda}_j^2), \\ \lambda_j \sim C^+(0, 1), \quad \tau \sim C^+(0, \tau_0), \quad c^2 \sim \text{Inv-Gamma}(2, 8),$$

³<http://mc-stan.org/users/documentation/case-studies/radon.html>

following closely a classification model proposed by Piironen and Vehtari (2017) with public Stan code⁴, but change to Gaussian likelihood suitable for regression problems. We apply the model on the *corn* data (Chen and Martin 2009), that we randomly split into equally sized training and test subsets.

For this model the quality of the variational approximation is sensitive to random initialization and the stochastic variation during the optimization, so that occasionally the posterior is reasonable whereas for some runs it converges to a very bad solution. Figure 4 illustrates the risk for 10 independent runs with different random seeds, showing that for 7 runs all decision-making strategies are sufficient. However, for 3 of the runs the q -optimal decision fails miserably, yet the parametric decision-maker is able to recover essentially p -optimal decisions. In other words, we show that the proposed strategy is able to correct for the mistake in the posterior.

5.4 Model Faithfulness

The main goal of our work is to improve the decisions made given a specific probabilistic model, which means the decisions should remain faithful to the underlying model: The decision-maker should correct for mistakes in the posterior approximation but not for model misspecification, since the users needs to be able to rely on interpretations of the model. Next we show the regularization strategy proposed in Section 3.2 achieves this.

Figure 5 shows the predictive distribution of intentionally incorrect model, a linear model for highly non-linear data. Even though the decision-maker only takes as input the predictive distribution (red shaded area), a flexible enough DM learns to ignore the model and returns the upper quantile of the data distribution that minimizes the risk. This can be done because there exists a mapping from the predictive quantiles back to the covariates, and hence the neural network can implicitly use the covariates themselves for predictions. In other words, it can learn to directly map the non-linear data, without conforming to the linear model.

Regularizing with (4) smoothly transitions the decisions towards the q -optimal decision, here a linear function of the covariates. Sufficiently strong regularization prevents the flexible DM module from overriding the model.

6 Conclusion

Even though Bayesian decision theory (Berger 1985) elegantly separates inference from decisions, this separation is constantly being misused by applying decision rules derived for the true posterior to approximations. This may dramatically affect the decisions, typically in a manner that is difficult to notice since there is no access to the true p -risk.

Some effort has been made to remedy this, by directly modifying the inference process to account for the decision loss (Lacoste-Julien, Huszár, and Ghahramani 2011; Cobb, Roberts, and Gal 2018; Kuśmierczyk, Sakaya, and Klami 2019; Abbasnejad, Domke, and Sanner 2015), but

⁴https://github.com/yao-yl/Evaluating-Variational-Inference/blob/master/R_code/glm_bernoulli_rhs.stan




Figure 5: Without regularization the decision-maker can ignore the underlying linear model (red), but the proposed regularization strategy prevents this and provides predictions that are compatible with the model and only account for approximation error. Here the q -optimal decisions are also p -optimal, because there is no approximation error.

these works are computationally costly and have demonstrated only marginal improvements. We proposed the alternative of retaining the original approximation and changing the decision-making process instead, by replacing the analytic decision with parametric decision-making module. This has the advantage of being agnostic to the inference strategy used, making the approach applicable to, e.g., improving predictive accuracy of models for which inference has already been carried out. To demonstrate the generality of our method, we ran several experiments on top of posterior approximations computed by probabilistic programming system Stan (Carpenter et al. 2017), demonstrating improved decisions for cases where the posterior approximation is inaccurate.

While the notion of directly optimizing for empirical risk seems unorthodox in a Bayesian sense, we presented several arguments supporting the validity of the approach. By building on generalized Bayesian inference (Bissiri, Holmes, and Walker 2016) we reformulate reasoning over the decision-makers as Bayesian inference, and we provide a prior distribution that allows controlling for how faithful we remain to the underlying model, with automatic procedure for selecting parameters that does not allow significant deviation. While some questions of interest still remain, we have demonstrated that the quantitative value of predictive Bayesian models can be improved by incorporating flexible decision-makers borrowed from the deep learning literature without compromising the interpretability of Bayesian modeling.

Acknowledgements

This work was supported by the Academy of Finland, grant 313125, and the Flagship programme: Finnish Center for Artificial Intelligence, FCAI. We thank the Finnish Grid and Cloud Infrastructure (urn:nbn:fi:research-infras-2016072533) for computational resources.

References

- Abbasnejad, E.; Domke, J.; and Sanner, S. 2015. Loss-calibrated Monte Carlo Action Selection. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- Berger, J. O. 1985. *Statistical Decision Theory and Bayesian Analysis; 2nd edition*. Springer Series in Statistics. New York: Springer.
- Bertin-Mahieux, T.; Ellis, D. P.; Whitman, B.; and Lamere, P. 2011. The Million Song Dataset. In *Proceedings of the 12th International Conference on Music Information Retrieval*.
- Bissiri, P. G.; Holmes, C. C.; and Walker, S. G. 2016. A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 78(5):1103–1130.
- Blei, D. M.; Kucukelbir, A.; and McAuliffe, J. D. 2017. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association* 112(518).
- Carpenter, B.; Gelman, A.; Hoffman, M. D.; Lee, D.; Goodrich, B.; Betancourt, M.; Brubaker, M.; Guo, J.; Li, P.; and Riddell, A. 2017. Stan: A Probabilistic Programming Language. *Journal of Statistical Software* 76(1).
- Carvalho, C. M.; Polson, N. G.; and Scott, J. G. 2010. The horseshoe estimator for sparse signals. *Biometrika* 97:465–480.
- Chen, T., and Martin, E. 2009. Bayesian linear regression and variable selection for spectroscopic calibration. *Analytica chimica acta* 631(1):13–21.
- Cobb, A. D.; Roberts, S. J.; and Gal, Y. 2018. Loss-Calibrated Approximate Inference in Bayesian Neural Networks. In *Theory of Deep Learning workshop, ICML*.
- Efron, B., and Tibshirani, R. J. 1994. *An introduction to the bootstrap*. CRC press.
- Gelman, A., and Hill, J. 2006. *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press.
- Gelman, A.; Vehtari, A.; Jylänki, P.; Sivula, T.; Tran, D.; Sahai, S.; Blomstedt, P.; Cunningham, J. P.; Schiminovich, D.; and Robert, C. 2017. Expectation propagation as a way of life: A framework for Bayesian inference on partitioned data. *arXiv preprint arXiv:1412.4869*.
- Gutmann, M. U., and Corander, J. 2016. Bayesian optimization for likelihood-free inference of simulator-based statistical models. *Journal of Machine Learning Research* 17(125):1–47.
- Hastie, T.; Tibshirani, R.; and Friedman, J. 2009. *The Elements of Statistical Learning: Data mining, inference and prediction*. Springer, 2nd edition edition.
- Hoffman, M. D., and Gelman, A. 2014. The no-u-turn sampler: Adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research* 15:1593–1623.
- Koenker, R., and Hallock, K. F. 2001. Quantile regression. *Journal of Economic Perspectives* 15(4):143–156.
- Kucukelbir, A.; Tran, D.; Ranganath, R.; Gelman, A.; and Blei, D. M. 2017. Automatic Differentiation Variational Inference. *The Journal of Machine Learning Research* 18(1).
- Kuśmierczyk, T.; Sakaya, J.; and Klami, A. 2019. Variational Bayesian Decision-making for Continuous Utilities. *arXiv preprint arXiv:1902.00792*.
- Lacoste-Julien, S.; Huszár, F.; and Ghahramani, Z. 2011. Approximate inference for the loss-calibrated Bayesian. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*.
- Minka, T. P. 2001. Expectation Propagation for Approximate Bayesian Inference. In *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*.
- Pearce, T.; Brintrup, A.; Zaki, M.; and Neely, A. 2018. High-quality prediction intervals for deep learning: A distribution-free, ensembled approach. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 4075–4084.
- Piironen, J., and Vehtari, A. 2017. On the hyperprior choice for the global shrinkage parameter in the horseshoe prior. In *Artificial Intelligence and Statistics*, 905–913.
- Titsias, M., and Lázaro-Gredilla, M. 2014. Doubly Stochastic Variational Bayes for non-Conjugate Inference. In *Proceedings of 31st International Conference on Machine Learning*.
- Tran, D.; Kucukelbir, A.; Dieng, A. B.; Rudolph, M.; Liang, D.; and Blei, D. M. 2016. Edward: A Library for Probabilistic Modeling, Inference, and Criticism. *arXiv:1610.09787*.
- Yao, Y.; Vehtari, A.; Simpson, D.; and Gelman, A. 2018. Yes, but did it work?: Evaluating variational inference. In *Proceedings of the 35th International Conference on Machine Learning*.