
Variational Bayesian Unlearning

Quoc Phong Nguyen, Bryan Kian Hsiang Low, and Patrick Jaillet[†]

Dept. of Computer Science, National University of Singapore, Republic of Singapore

Dept. of Electrical Engineering and Computer Science, MIT, USA[†]

{qphong, lowkh}@comp.nus.edu.sg, jaillet@mit.edu[†]

Abstract

This paper studies the problem of approximately unlearning a Bayesian model from a small subset of the training data to be erased. We frame this problem as one of minimizing the Kullback-Leibler divergence between the approximate posterior belief of model parameters after directly unlearning from erased data vs. the exact posterior belief from retraining with remaining data. Using the *variational inference* (VI) framework, we show that it is equivalent to minimizing an evidence upper bound which trades off between fully unlearning from erased data vs. not entirely forgetting the posterior belief given the full data (i.e., including the remaining data); the latter prevents catastrophic unlearning that can render the model useless. In model training with VI, only an approximate (instead of exact) posterior belief given the full data can be obtained, which makes unlearning even more challenging. We propose two novel tricks to tackle this challenge. We empirically demonstrate our unlearning methods on Bayesian models such as sparse Gaussian process and logistic regression using synthetic and real-world datasets.

1 Introduction

Our interactions with *machine learning* (ML) applications have surged in recent years such that large quantities of users’ data are now deeply ingrained into the ML models being trained for these applications. This greatly complicates the regulation of access to each user’s data or implementation of *personal data ownership*, which are enforced by the General Data Protection Regulation in the European Union [24]. In particular, if a user would like to exercise her *right to be forgotten* [24] (e.g., when quitting an ML application), then it would be desirable to have the trained ML model “unlearn” from her data. Such a problem of *machine unlearning* [4] extends to the practical scenario where a small subset of data previously used for training is later identified as malicious (e.g., anomalies) [4, 9] and the trained ML model can perform well once again if it can unlearn from the malicious data.

A naive alternative to machine unlearning is to simply retrain an ML model from scratch with the data *remaining* after *erasing* that to be unlearned from. In practice, this is prohibitively expensive in terms of time and space costs since the remaining data is often large such as in the above scenarios. How then can a trained ML model *directly* and efficiently unlearn from a small subset of data to be erased to become (a) exactly and if not, (b) approximately close to that from retraining with the large remaining data? Unfortunately, (a) exact unlearning is only possible for selected ML models (e.g., naive Bayes classifier, linear regression, *k*-means clustering, and item-item collaborative filtering [4, 12, 30]). This motivates the need to consider (b) approximate unlearning as it is applicable to a broader family of ML models like neural networks [9, 13] but, depending on its choice of loss function, may suffer from *catastrophic unlearning*¹ that can render the model useless. For example, to mitigate this issue, the works of [9, 13] have to “patch up” their loss functions by additionally bounding the loss incurred

¹A trained ML model is said to experience *catastrophic unlearning* from the erased data when its resulting performance is considerably worse than that from retraining with the remaining data.

by erased data with a rectified linear unit and injecting a regularization term to retain information of the remaining data, respectively. This begs the question whether there exists a loss function that can *directly* quantify the approximation gap and *naturally* prevent catastrophic unlearning.

Our work here addresses the above question by focusing on the family of Bayesian models. Specifically, our proposed loss function measures the *Kullback-Leibler* (KL) divergence between the approximate posterior belief of model parameters by directly unlearning from erased data vs. the exact posterior belief from retraining with remaining data. Using the *variational inference* (VI) framework, we show that minimizing this KL divergence is equivalent to *minimizing* (instead of maximizing) a counterpart of the evidence lower bound called the *evidence upper bound* (EUBO) (Sec. 3.2). Interestingly, the EUBO lends itself to a natural interpretation of a trade-off between fully unlearning from erased data vs. not entirely forgetting the posterior belief given the *full* data (i.e., including the remaining data); the latter prevents catastrophic unlearning induced by the former.

Often, in model training, only an approximate (instead of exact) posterior belief of model parameters given the full data can be learned, say, also using VI. This makes unlearning even more challenging. To tackle this challenge, we analyse two sources of inaccuracy in the approximate posterior belief learned using VI, which lay the groundwork for proposing our first trick of an *adjusted likelihood* of erased data (Sec. 3.3.1): Our key idea is to curb unlearning in the region of model parameters with low approximate posterior belief where both sources of inaccuracy primarily occur. Additionally, to avoid the risk of incorrectly tuning the adjusted likelihood, we propose another trick of *reverse KL* (Sec. 3.3.2) which is naturally more protected from such inaccuracy without needing the adjusted likelihood. Nonetheless, our adjusted likelihood is general enough to be applied to reverse KL.

VI is a popular approximate Bayesian inference framework due to its scalability to massive datasets [15, 18] and its ability to model complex posterior beliefs using generative adversarial networks [33] and normalizing flows [21, 29]. Our work in this paper exploits VI to broaden the family of ML models that can be unlearned, which we empirically demonstrate using synthetic and real-world datasets on several Bayesian models such as sparse Gaussian process and logistic regression with the approximate posterior belief modeled by a normalizing flow (Sec. 4).

2 Variational Inference (VI)

In this section, we revisit the VI framework [2] for learning an approximate posterior belief of the parameters θ of a Bayesian model. Given a prior belief $p(\theta)$ of the unknown model parameters and a set D of training data, an approximate posterior belief $q(\theta|D)$ is being optimized by minimizing the KL divergence $\text{KL}[q(\theta|D) \parallel p(\theta)] = \int q(\theta|D) \log(q(\theta|D)/p(\theta)) d\theta$ or, equivalently, maximizing the *evidence lower bound* (ELBO) L [2]:

$$L = \int q(\theta|D) \log p(D|\theta) d\theta - \text{KL}[q(\theta|D) \parallel p(\theta)]. \quad (1)$$

Such an equivalence follows directly from $L = \log p(D) - \text{KL}[q(\theta|D) \parallel p(\theta)]$ where the log-marginal likelihood $\log p(D)$ is independent of $q(\theta|D)$. Since $\text{KL}[q(\theta|D) \parallel p(\theta)] \geq 0$, the ELBO L is a lower bound of $\log p(D)$. The ELBO L in (1) can be interpreted as a trade-off between attaining a higher likelihood of D (first term) vs. not entirely forgetting the prior belief $p(\theta)$ (second term).

When the ELBO L (1) cannot be evaluated in closed form, it can be maximized using *stochastic gradient ascent* (SGA) by approximating the expectation in

$$L = \mathbb{E}_{q(\theta|D)}[\log p(D|\theta) + \log(p(\theta)/q(\theta|D))] = \int q(\theta|D) (\log p(D|\theta) + \log(p(\theta)/q(\theta|D))) d\theta$$

with stochastic sampling in each iteration of SGA. The approximate posterior belief $q(\theta|D)$ can be represented by a simple distribution (e.g., in the exponential family) for computational ease or a complex distribution (e.g., using generative neural networks) for expressive power. Note that when the distribution of $q(\theta|D)$ is modeled by a generative neural network whose density cannot be evaluated, the ELBO can be maximized with adversarial training by alternating between estimating the log-density ratio $\log(p(\theta)/q(\theta|D))$ and maximizing the ELBO [33]. On the other hand, when the distribution of $q(\theta|D)$ is modeled by a normalizing flow (e.g., *inverse autoregressive flow* (IAF) [21]) whose density can be computed, the ELBO can be maximized with SGA.

3 Bayesian Unlearning

3.1 Exact Bayesian Unlearning

Let the (full) training data D be partitioned into a small subset D_e of data to be erased and a (large) set D_r of remaining data, i.e. $D = D_r \cup D_e$ and $D_r \cap D_e = \emptyset$. The problem of exact Bayesian unlearning involves recovering the exact posterior belief $p(\theta | D_r)$ of model parameters given remaining data D_r from that given full data D (i.e., $p(\theta | D)$ assumed to be available) by directly unlearning from erased data. Note that $p(\theta | D_r)$ can also be obtained from retraining with remaining data D_r , which is computationally costly, as discussed in Sec. 1. By using Bayes' rule and assuming conditional independence between D_r and D_e given θ ,

$$p(\theta | D_r) = p(\theta | D) p(D_e | D_r) / p(D_e) = p(\theta | D) p(D_e | \theta) / p(D_e) \quad (2)$$

When the model parameters are discrete-valued, $p(\theta | D_r)$ can be obtained from (2) directly. The use of a conjugate prior also makes unlearning relatively simple. We will investigate the more interesting case of a non-conjugate prior in the rest of Sec. 3.

3.2 Approximate Bayesian Unlearning with Exact Posterior Belief $p(\theta | D)$

The problem of approximate Bayesian unlearning differs from that of exact Bayesian unlearning (Sec. 3.1) in that only the approximate posterior belief $q_\mu(\theta | D_r)$ (instead of the exact one $p(\theta | D_r)$) can be recovered by directly unlearning from erased data. Since existing unlearning methods often use their model predictions to construct their loss functions [12, 14], we have initially considered doing likewise (albeit in the Bayesian context) by defining the loss function as the KL divergence between the approximate predictive distribution $q_\mu(y | D_r)$, $\sum_y p(y) q_\mu(\theta | D_r) d$ vs. the exact predictive distribution $p(y | D_r) = \sum_y p(y) p(\theta | D_r) d$ where the observation y (i.e., drawn from a model with parameters θ) is conditionally independent of D_r given θ . However, it may not be possible to evaluate these predictive distributions in closed form, hence making the optimization of this loss function computationally difficult. Fortunately, such a loss function can be bounded from above by the KL divergence between posterior beliefs $q_\mu(\theta | D_r)$ vs. $p(\theta | D_r)$, as proven in Appendix A:

$$\text{Proposition 1. } \text{KL}[q_\mu(y | D_r) \parallel p(y | D_r)] \leq \text{KL}[q_\mu(\theta | D_r) \parallel p(\theta | D_r)]^2$$

Proposition 1 reveals that reducing $\text{KL}[q_\mu(\theta | D_r) \parallel p(\theta | D_r)]$ decreases $\text{KL}[q_\mu(y | D_r) \parallel p(y | D_r)]$, thus motivating its use as the loss function instead. In particular, it follows immediately from our result below (i.e., proven in Appendix B) that minimizing $\text{KL}[q_\mu(\theta | D_r) \parallel p(\theta | D_r)]$ is equivalent to minimizing a counterpart of the ELBO called the evidence upper bound (EUBO) U :

Proposition 2. Define the EUBO as

$$U = \sum_{D_e} q_\mu(\theta | D_r) \log p(D_e | \theta) + \text{KL}[q_\mu(\theta | D_r) \parallel p(\theta | D)] \quad (3)$$

Then, $U = \log p(D_e | D_r) + \text{KL}[q_\mu(\theta | D_r) \parallel p(\theta | D_r)] - \log p(D_e | D_r)$ such that $p(D_e | D_r)$ is independent of $q_\mu(\theta | D_r)$.

From Proposition 2, minimizing EUBO (3) is equivalent to minimizing $\text{KL}[q_\mu(\theta | D_r) \parallel p(\theta | D_r)]$ which is precisely achieved using VI (i.e., by maximizing ELBO) from retraining with remaining data D_r . This is illustrated in Fig. 1a where unlearning from D_e by minimizing EUBO maximizes ELBO w.r.t. D_r ; in Fig. 1b, retraining with D_r by maximizing ELBO minimizes EUBO w.r.t. D_e .

The EUBO (3) can be interpreted as a trade-off between fully unlearning from erased data (1st term) vs. not entirely forgetting the exact posterior belief $p(\theta | D)$ given the full data D (i.e., including the remaining data D_r) (second term). The latter can be viewed as a regularization term to prevent catastrophic unlearning (i.e., potentially induced by the former) that naturally results from minimizing our loss function $\text{KL}[q_\mu(\theta | D_r) \parallel p(\theta | D_r)]$, which differs from the works of [9, 13] needing to “patch up” their loss functions (Sec. 1). Generative models can be used to model the approximate posterior belief $q_\mu(\theta | D_r)$ in the EUBO (3) in the same way as that in the ELBO (1).

²Similarly, $\text{KL}[p(y | D_r) \parallel q_\mu(y | D_r)] \leq \text{KL}[p(\theta | D_r) \parallel q_\mu(\theta | D_r)]$ holds.

(a) Unlearning from D_e by minimizing EUBO | (b) Retraining with D_r by maximizing ELBO
 Figure 1: Plots of EUBO and ELBO when (a) unlearning from D_e and (b) retraining with D_r .

Figure 2: Plot of $q(\theta|D)$ learned using VI. Gray shaded region corresponds to values where $q(\theta|D) = \max_{\theta} p(\theta|D)$. Vertical blue strips on horizontal axis show 100 samples of $q(\theta|D)$.

3.3 Approximate Bayesian Unlearning with Approximate Posterior Belief $q(\theta|D)$

Often, in model training, only an approximate posterior belief $q(\theta|D)$ (instead of the exact $p(\theta|D)$ in Sec. 3.2) of model parameters given full data D can be learned, say, using VI by maximizing the ELBO (Sec. 2). Our proposed unlearning methods are parsimonious in requiring $q(\theta|D)$ and erased data D_e to be available, which makes unlearning even more challenging since there is no further information about $p(\theta|D)$ nor the difference between $p(\theta|D)$ vs. $q(\theta|D)$. So, we estimate the unknown $p(\theta|D_r)$ (2) with

$$p(\theta|D_r) / q(\theta|D) = p(\theta|D_e) \quad (4)$$

and minimize the KL divergence between the approximate posterior belief recovered by directly unlearning from erased data D_e vs. $p(\theta|D_r)$ (4) instead. We will discuss two novel tricks below to alleviate the undesirable consequence of using $p(\theta|D_r)$ instead of the unknown $p(\theta|D)$ (2).

3.3.1 EUBO with Adjusted Likelihood

Let the loss function $\text{KL}[q_u(\theta|D_r) \parallel p(\theta|D_r)]$ be minimized w.r.t. the approximate posterior belief $q_u(\theta|D_r)$ that is recovered by directly unlearning from erased data. Similar to Proposition 2, $q_u(\theta|D_r)$ can be optimized by minimizing the following EUBO:

$$\mathcal{L}, \quad q_u(\theta|D_r) \log p(D_e|\theta) + \text{KL}[q_u(\theta|D_r) \parallel q(\theta|D)] \quad (5)$$

which follows from simply replacing the unknown $p(\theta|D)$ in U (3) with $q(\theta|D)$. We discuss the difference between $p(\theta|D)$ vs. $q(\theta|D)$ in the remark below:

Remark 1. We analyze two possible sources of inaccuracy in $q(\theta|D)$ that is learned using VI by minimizing the loss function $\text{KL}[q(\theta|D) \parallel p(\theta|D)]$ (Sec. 2). Firstly, $q(\theta|D)$ often underestimates the variance of $p(\theta|D)$: Though $q(\theta|D)$ tends to be close to $p(\theta|D)$ at values of θ where $p(\theta|D)$ is close to 0, the reverse is not enforced (see, for example, Fig. 2). $q(\theta|D)$ can differ from $p(\theta|D)$ at values of θ where $q(\theta|D)$ is close to 0. Secondly, if $q(\theta|D)$ is learned through stochastic optimization of the ELBO (i.e., with stochastic samples of $q(\theta|D)$ in each iteration of SGA), then it is unlikely that the ELBO is maximized using samples of $q(\theta|D)$ with small $q(\theta|D)$ (Fig. 2). Thus, both sources of inaccuracy primarily occur at values of θ with small $q(\theta|D)$. Though it can also be inaccurate at values of θ with large $q(\theta|D)$, such an inaccuracy can be reduced by representing $q(\theta|D)$ with a complex distribution (Sec. 2).

Remark 1 motivates us to curb unlearning at values with small $q(\theta|D)$ by proposing our first novel trick of an adjusted likelihood of the erased data:

$$p_{\text{adj}}(D_e|\theta; q), \quad \begin{cases} p(D_e|\theta) & \text{if } q(\theta|D) > \max_{\theta} q(\theta|D); \\ 1 & \text{otherwise (i.e., shaded area in Fig.;)} \end{cases} \quad (6)$$

³With a slight abuse of notation, we let $q(\theta|D)$ be the approximate posterior belief that maximizes the ELBO L (1) (Sec. 2) from Sec. 3.3 onwards.

Figure 3: Plot of approximate posterior beliefs with varying α obtained by minimizing (a) EUBO (i.e., $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha)$) and (b) reverse KL (i.e., $\mathbf{q}_\nu(\mathbf{j}D_r; \alpha)$); horizontal axis denotes α . In (a), a huge probability mass $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha = 0)$ is at large values of $\mathbf{j}D$ beyond the plotting area and the top of the plot of $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha = 10^{-20})$ is cut off due to lack of space.

for any α where $\alpha \in [0; 1]$ controls the magnitude of a threshold under which $\mathbf{j}D$ is considered small. To understand the effect of let $\mathbf{p}_{\text{adj}}(\mathbf{j}D_r; \alpha) / \mathbf{q}(\mathbf{j}D) = \mathbf{p}_{\text{adj}}(D_{e_j}; \alpha)$, i.e., by replacing $\mathbf{p}(D_{e_j})$ in (4) with $\mathbf{p}_{\text{adj}}(D_{e_j}; \alpha)$. Then, using (6),

$$\mathbf{p}_{\text{adj}}(\mathbf{j}D_r; \alpha) / \mathbf{q}(\mathbf{j}D) = \begin{cases} \mathbf{p}(D_{e_j}) & \text{if } \mathbf{q}(\mathbf{j}D) > \max_{\mathbf{j}D} \mathbf{q}(\mathbf{j}D); \\ \mathbf{q}(\mathbf{j}D) & \text{otherwise (i.e., shaded area in Fig. 2)}. \end{cases} \quad (7)$$

According to (7), unlearning is therefore focused at values $\mathbf{j}D$ with sufficiently large $\mathbf{q}(\mathbf{j}D)$ (i.e., $\mathbf{q}(\mathbf{j}D) > \max_{\mathbf{j}D} \mathbf{q}(\mathbf{j}D)$). Let the loss function $\text{KL}[\mathbf{q}_\mu(\mathbf{j}D_r; \alpha) \parallel \mathbf{p}_{\text{adj}}(\mathbf{j}D_r; \alpha)]$ be minimized w.r.t. the approximate posterior belief $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha)$ that is recovered by directly unlearning from erased data D_e . Similar to (5), $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha)$ can be optimized by minimizing the following EUBO:

$$\mathcal{L}_{\text{adj}}(\alpha), \quad \mathbf{q}_\mu(\mathbf{j}D_r; \alpha) \log \mathbf{p}_{\text{adj}}(D_{e_j}; \alpha) + \text{KL}[\mathbf{q}_\mu(\mathbf{j}D_r; \alpha) \parallel \mathbf{q}(\mathbf{j}D)] \quad (8)$$

which follows from replacing $\mathbf{p}(D_{e_j})$ in (5) with $\mathbf{p}_{\text{adj}}(D_{e_j}; \alpha)$. Note that $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha)$ can be represented by a simple distribution (e.g., Gaussian) or a complex one (e.g., generative neural network, IAF). We initialize $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha)$ at $\mathbf{q}(\mathbf{j}D)$ for achieving empirically faster convergence. When $\alpha = 0$, $\mathcal{L}_{\text{adj}}(\alpha = 0)$ reduces to \mathcal{L} (5), i.e., EUBO is minimized without the adjusted likelihood. As a result, $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha = 0) = \mathbf{q}_\mu(\mathbf{j}D_r)$. As α increases, unlearning is focused on a smaller and smaller region of $\mathbf{j}D$ with sufficiently large $\mathbf{q}(\mathbf{j}D)$. When α reaches 1, no unlearning is performed since $\mathbf{p}_{\text{adj}}(\mathbf{j}D_r; \alpha = 1) = \mathbf{q}(\mathbf{j}D)$, which results in $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha = 1) = \mathbf{q}(\mathbf{j}D)$ minimizing the loss function $\text{KL}[\mathbf{q}_\mu(\mathbf{j}D_r; \alpha = 1) \parallel \mathbf{p}_{\text{adj}}(\mathbf{j}D_r; \alpha = 1)]$.

Example 1. To visualize the effect of varying α on $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha)$, we consider learning the shape of a Gamma distribution with a known rate (i.e., β): D are 20 samples of the Gamma distribution, D_e are the smallest 5 samples in D , and the (non-conjugate) prior belief and approximate posterior beliefs of β are all Gaussians. Fig. 3a shows the approximate posterior belief $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha)$ with varying α as well as $\mathbf{q}(\mathbf{j}D)$ and $\mathbf{p}_{\text{adj}}(\mathbf{j}D_r; \alpha)$ learned using VI. As explained above, $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha = 1) = \mathbf{q}(\mathbf{j}D)$. When $\alpha = 0.001$, $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha = 0.001)$ is close to $\mathbf{q}(\mathbf{j}D_r)$. However, as α decreases to 0, $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha)$ moves away from $\mathbf{q}(\mathbf{j}D_r)$.

The optimized $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha)$ suffers from the same issue of underestimating the variance σ^2 learned using VI (see Remark 1), especially when α tends to 0 (e.g., see $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha = 10^{-20})$ in Fig. 3a). Though this issue can be mitigated by tuning the adjusted likelihood (6), we may not want to risk facing the consequence of picking a value α that is too small. So, in Sec. 3.3.2, we will propose another novel trick that is not inconvenienced by this issue.

3.3.2 Reverse KL

Let the loss function be the reverse KL divergence $\text{KL}[\mathbf{p}(\mathbf{j}D_r) \parallel \mathbf{q}_\nu(\mathbf{j}D_r)]$ that is minimized w.r.t. the approximate posterior belief $\mathbf{q}_\nu(\mathbf{j}D_r)$ recovered by directly unlearning from erased data D_e . In contrast to the optimized $\mathbf{q}_\mu(\mathbf{j}D_r; \alpha)$ from minimizing EUBO(8), the optimized $\mathbf{q}_\nu(\mathbf{j}D_r)$ from minimizing the reverse KL divergence overestimates (instead of underestimates) the variance of $\mathbf{p}(\mathbf{j}D_r)$ [1]: If $\mathbf{p}(\mathbf{j}D_r)$ is close to 0, then $\mathbf{q}_\nu(\mathbf{j}D_r)$ is not necessarily close to 0. From (4), the reverse KL divergence can be expressed as follows:

$$\text{KL}[\mathbf{p}(\mathbf{j}D_r) \parallel \mathbf{q}_\nu(\mathbf{j}D_r)] = C_0 - C_1 \mathbb{E}_{\mathbf{q}(\mathbf{j}D)} [(\log \mathbf{q}_\nu(\mathbf{j}D_r)) = \mathbf{p}(D_{e_j})] \quad (9)$$

where C_0 and C_1 are constants independent of $(\mathbf{j}D_r)$. So, the reverse KL divergence (8) can be minimized by maximizing $E_{q(\mathbf{j}D)}[(\log \mathbf{q}_v(\mathbf{j}D_r)) = p(D_{ej})]$ with stochastic gradient ascent (SGA). We also initialize $\mathbf{q}_v(\mathbf{j}D_r)$ at $q(\mathbf{j}D)$ for achieving empirically faster convergence. Since stochastic optimization is performed with samples of $q(\mathbf{j}D)$ in each iteration of SGA, it is unlikely that the reverse KL divergence (8) is minimized using samples of with small $q(\mathbf{j}D)$. This naturally curbs unlearning at values of with small $q(\mathbf{j}D)$, as motivated by Remark 1. On the other hand, it is still possible to employ the same trick of adjusted likelihood (Sec. 3.3.1) by minimizing the reverse KL divergence $KL[p_{\text{adj}}(\mathbf{j}D_r; \cdot) \parallel \mathbf{q}_v(\mathbf{j}D_r; \cdot)]$ as the loss function or, equivalently, maximizing $E_{q(\mathbf{j}D)}[(\log \mathbf{q}_v(\mathbf{j}D_r; \cdot)) = p_{\text{adj}}(D_{ej}; \cdot)]$ where $p_{\text{adj}}(D_{ej}; \cdot)$ and $p_{\text{adj}}(\mathbf{j}D_r; \cdot)$ are previously defined in (6) and (7), respectively. The role of here is the same as that in (8).

To illustrate the difference between minimizing the reverse KL divergence (9) and EUBO (8), Fig. 3b shows the approximate posterior beliefs $q(\mathbf{j}D_r; \cdot)$ with varying β . It can be observed that $\mathbf{q}_v(\mathbf{j}D_r; \beta = 1) = q(\mathbf{j}D)$ (i.e., no unlearning). In contrast to minimizing EUBO (Fig. 3a), as β decreases to 0, $\mathbf{q}_v(\mathbf{j}D_r; \cdot)$ does not deviate that much from $q(\mathbf{j}D_r)$, even when $\beta = 0$ (i.e., the reverse KL divergence is minimized without the adjusted likelihood). This is because the optimized $\mathbf{q}_v(\mathbf{j}D_r; \cdot)$ is naturally more protected from both sources of inaccuracy (Remark 1), as explained above. Hence, we do not have to be as concerned about picking a small value β which is also consistently observed in our experiments (Sec. 4).

4 Experiments and Discussion

This section empirically demonstrates our unlearning methods on Bayesian models such as sparse Gaussian process and logistic regression using synthetic and real-world datasets. Further experimental results on Bayesian linear regression and with a bimodal posterior belief are reported in Appendices C and D, respectively. In our experiments, each dataset comprises pairs of inputs corresponding output/observation y_x . We use RMSProp as the SGA algorithm with a learning rate of 0.01. To assess the performance of our unlearning methods (i.e., by directly unlearning from erased data, we consider the difference between their induced predictive distributions vs. that obtained using VI from retraining with remaining data D_r , as motivated from Sec. 3.2. To do this, we use a performance metric that measures the KL divergence between the approximate predictive distributions

$$q_u(y_x | \mathbf{j}D_r), \quad p(y_x | \cdot) \parallel q_u(\mathbf{j}D_r; \cdot) \quad \text{or} \quad q_v(y_x | \mathbf{j}D_r), \quad p(y_x | \cdot) \parallel q_v(\mathbf{j}D_r; \cdot) \quad \text{d}$$

vs. $q(y_x | \mathbf{j}D_r), \quad p(y_x | \cdot) \parallel q(\mathbf{j}D_r; \cdot)$ where $q_u(\mathbf{j}D_r; \cdot)$ and $q_v(\mathbf{j}D_r; \cdot)$ are optimized by, respectively, minimizing EUBO (8) and reverse KL (rKL) divergence (9) while requiring only $q(\mathbf{j}D)$ and erased data D_e (Sec. 3.3), and $q(\mathbf{j}D_r)$ is learned using VI (Sec. 2). The above predictive distributions are computed via sampling with 100 samples of \cdot . For tractability reasons, we evaluate the above performance metric over finite input domains, specifically, over the D_r and D_e , which allows us to assess the performance of our unlearning methods on both the erased and remaining data, i.e., whether they can fully unlearn from D_e yet not forget nor catastrophically unlearn from D_r , respectively. For example, the performance of our EUBO-based unlearning method is shown as an average (with standard deviation) of the KL divergences between $q_u(y_x | \mathbf{j}D_r)$ vs. $q(y_x | \mathbf{j}D_r)$ over all $(x; y_x) \in D_e$. We also plot an average (with standard deviation) of the KL divergences between $q(y_x | \mathbf{j}D)$ vs. $q(y_x | \mathbf{j}D_r)$ over D_r and D_e as baselines (i.e., representing no unlearning), which is expected to be larger than that of our unlearning methods (i.e., if performing well) and labeled as full in the plots below.

4.1 Sparse Gaussian Process (GP) Classification with Synthetic Moon Dataset

This experiment is about unlearning a binary classifier that is previously trained with the synthetic moon dataset (Fig. 4a). The probability of input $\mathbf{x} \in \mathbb{R}^2$ being in the 'blue' class (i.e., $y_x = 1$ and denoted by blue dots in Fig. 4a) is defined as $\frac{1}{1 + \exp(-f_x)}$ where f_x is a latent function modeled by a sparse GP [7], which is elaborated in Appendix E. The parameters of the sparse GP consist of 20 inducing variables; the approximate posterior beliefs are thus multivariate Gaussians (with full covariance matrices), as shown in Appendix E. By comparing Figs. 4b and 4c, it can be observed that after erasing D_e (i.e., mainly in 'yellow' class) $q(y_x = 1 | \mathbf{j}D_r)$ increases at D_e . Figs. 4d and 4e show results of the performance of our EUBO- and rKL-based unlearning methods




Figure 4: Plots of (a) synthetic moon dataset with erased data (crosses) and remaining data D_r (dots), and of predictive distributions obtained using VI from (b) training with full data and (c) retraining with D_r . Graphs of averaged KL divergence vs. achieved by EUBO, rKL, and $q(y_x | D)$ (i.e., baseline labeled as full) over (d) D_r and (e) D_e . Plots of predictive distributions (f-h) $q_u(y_x = 1 | D_r)$ and (i-k) $q_v(y_x = 1 | D_r)$ induced, respectively, by EUBO and rKL for varying

over D_r and D_e with varying α , respectively. When $\alpha = 10^{-9}$, EUBO performs reasonably well (compare Figs. 4g vs. 4c) as its averaged KL divergence is smaller than that of $q(y_x | D)$ (i.e., baseline labeled as full). When $\alpha = 0$, EUBO performs poorly (compare Figs. 4h vs. 4c) as its averaged KL divergence is much larger than that of $q(y_x | D)$, as shown in Figs. 4d and 4e. This agrees with our discussion of the issue with picking too small a value for EUBO at the end of Sec. 3.3.1. In particular, catastrophic unlearning is observed as the input region containing, mainly in 'yellow' class has a high probability in 'blue' class after unlearning in Fig. 4h. On the other hand, when $\alpha = 0$, rKL performs well (compare Figs. 4k vs. 4c) as its KL divergence is much smaller than that of $q(y_x | D)$, as seen in Figs. 4d and 4e. This agrees with our discussion at the end of Sec. 3.3.2 that rKL can work well without needing the adjusted likelihood.

One may question how the performance of our unlearning methods would vary when erasing a large quantity of data or with different distributions of erased data (e.g., erasing the data randomly vs. deliberately erasing all data in a given class). To address this question, we have discovered that a key factor in unlearning their unlearning performance in these scenarios is the difference between the posterior beliefs of model parameters given remaining data D_r vs. that given full data D , especially at values of α with small $q(y_x | D)$ since unlearning in such a region is curbed by the adjusted likelihood and reverse KL. In practice, we expect such a difference not to be large due to the small quantity of erased data and redundancy in real-world datasets. We will present the details of this study in Appendix F due to lack of space by considering how m_{D_r} reduces the entropy of given D_r .

4.2 Logistic Regression with Banknote Authentication Dataset

The banknote authentication dataset of size $|D| = 1372$ is partitioned into erased data of size $|D_e| = 412$ and remaining data of size $|D_r| = 960$. Each input x comprises 4 features extracted from an image of a banknote and its corresponding binary label indicates whether the banknote is genuine or forged. We use a logistic regression model with parameters that is trained with this dataset. The prior beliefs of the model parameters are independent Gaussians.

Unlike the previous experiment, the erased data here is randomly selected and hence does not reduce the entropy of model parameters given D_r much, as explained in Appendix F; a discussion on erasing informative data (such as that in Sec. 4.1) is in Appendix F. As a result, Figs. 5a and 5b show a very small averaged KL divergence of about 10^{-3} between $q(y_x | D)$ vs. $q(y_x | D_r)$ (i.e., baselines) over D_r and D_e .⁴ Figs. 5a and 5b also show that our unlearning methods do not perform well when using multivariate Gaussians to model the approximate posterior beliefs. While rKL still gives a useful $q_v(y_x | D_r; \cdot)$ achieving an averaged KL divergence close to that of $q(y_x | D)$, EUBO gives a useless $q_u(y_x | D_r; \cdot)$ incurring a large averaged KL divergence when α is small. On the other hand,

⁴Note that the log plots can only properly display the upper confidence interval as standard deviation (shaded area) and hence do not show the lower confidence interval.

(a) D_r (b) D_e (c) D_r (d) D_e

Figure 5: Graphs of averaged KL divergence achieved by EUBO, rKL, and (jD) (i.e., baseline labeled as ll) over D_r and D_e for the banknote authentication dataset with the approximate posterior beliefs of model parameters represented by (a-b) multivariate Gaussians and (c-d) normalizing flows.

when more complex models like normalizing flows with the MADE architecture [26] are used to represent the approximate posterior beliefs, EUBO and rKL can unlearn well (Figs. 5c and 5d).

4.3 Logistic Regression with Fashion MNIST Dataset

The fashion MNIST dataset of size $|D| = 60000$ (28 × 28 images of fashion items in 10 classes) is partitioned into erased data of size $|D_e| = 10000$ and remaining data of size $|D_r| = 50000$. The classification model is a neural network with fully-connected hidden layers of 128, 128, 64 hidden neurons and a softmax layer to output 10-class probabilities. The model can be interpreted as one of logistic regression of 64 features generated from the hidden layer neurons. Since modeling all weights of the neural network as random variables can be costly, we model only weights in the transformation of the 64 features to the inputs of the softmax layer. The other weights remain constant during unlearning and retraining. The prior beliefs of the network weights are 0. The approximate posterior beliefs are modeled with independent Gaussians. Though a large part of the network is fixed and we use simple models to represent the approximate posterior beliefs, we show that unlearning is still fairly effective.

As discussed in Sec. 4.1, 4.2, and Appendix F, the random selection of erased data redundancy in D lead to a small averaged KL divergence of about between $q(y_x | D)$ vs. $q(y_x | D_r)$ (i.e., baselines) over D_r and D_e (Figs. 6a and 6b) despite choosing a relatively large β . Figs. 6a and 6b show that when $\beta = 10^{-9}$, EUBO and rKL achieve averaged KL divergences comparable to that of $q(\cdot | D)$ (i.e., baseline labeled as ll), hence making their unlearning insignificant. However, at $\beta = 0$, the unlearning performance of rKL improves by achieving a smaller averaged KL divergence than that of $q(\cdot | D)$, while EUBO's performance deteriorates. Their performance can be further improved by using more complex models to represent their approximate posterior beliefs like that in Sec. 4.2, albeit high-dimensional. Figs. 6c and 6d show the class probabilities for two images evaluated at the mean of the approximate posterior beliefs μ . We observe that rKL induces the highest class probability for the same class as that of $q(\cdot | D_r)$. The class probabilities for other images are shown in Appendix G. The two images are taken from a separate 6000 test images (i.e., different from D) where rKL with $\beta = 0$ yields the same predictions as $q(\cdot | D_r)$ and $q(\cdot | D)$ in, respectively 99.34% and 99.22% of the test images, the latter of which are contained in the former.

4.4 Sparse Gaussian Process (GP) Regression with Airline Dataset

This section illustrates the scalability of unlearning to the massive airline dataset of 20 million flights [15]. Training a sparse GP model with this massive dataset is made possible through stochastic VI [15]. Let X_u denote the set of inducing inputs in the sparse GP model and \mathbf{z} be a vector of corresponding latent function values (i.e., inducing variables). The posterior belief $q(\mathbf{z} | X_u, D)$ is approximated as $q(\mathbf{z} | X_u, D) \approx q(\mathbf{z} | X_u) p(\mathbf{z} | X_u)$ where $\mathbf{z} \sim \mathcal{N}(\mathbf{z} | \mathbf{0}, \mathbf{I})$. Let the sets X_D and X_{D_e} denote inputs in the full and erased data, respectively. Then, the ELBO can be decomposed to

$$L = \int_{X_D} \int_{X_{D_e}} \int_{\mathbf{z}} q(\mathbf{z} | X_u, D) p(\mathbf{z} | X_u) \log p(y_x | f_x) df_x df_{X_u} - \text{KL}[q(\mathbf{z} | X_u, D) \| p(\mathbf{z} | X_u)] \quad (10)$$

where $\int_{X_D} \int_{X_{D_e}} p(f_x | f_{X_u}) \log p(y_x | f_x) df_x$ can be evaluated in closed form [11]. To unlearn such a trained model from D_e ($|D_e| = 100K$ here), the EUBO (8) can be expressed in a similar way as the ELBO:

$$\mathcal{L}_{\text{adj}}(\theta) = \int_{X_D} \int_{X_{D_e}} \int_{\mathbf{z}} q_{\theta}(\mathbf{z} | X_u, D_r; \theta) p(f_x | f_{X_u}) \log p_{\text{adj}}(y_x | f_x; \theta) df_x df_{X_u} + \text{KL}[q_{\theta}(\mathbf{z} | X_u, D_r; \theta) \| q(\mathbf{z} | X_u, D)]$$

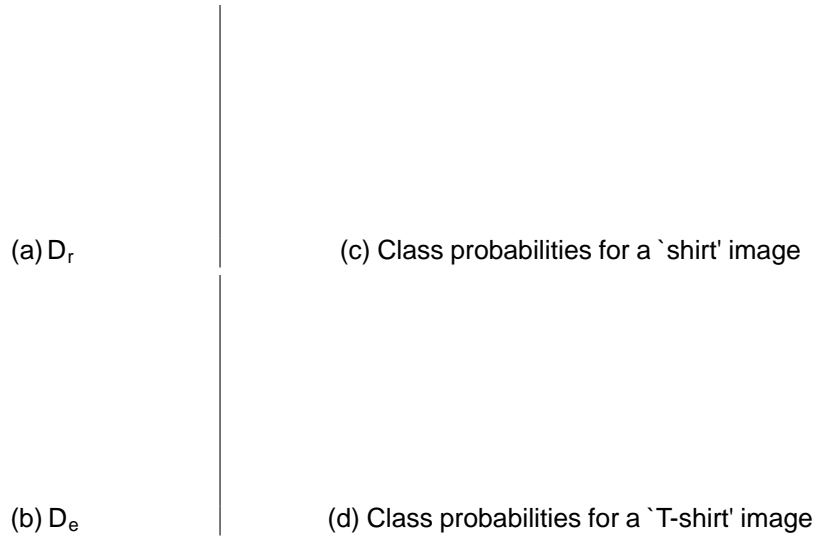


Figure 6: Graphs of averaged KL divergence achieved by EUBO, rKL, and (jD) (i.e., baseline labeled as full) over (a) D_r and (b) D_e . (c-d) Plots of class probabilities for two images in the fashion MNIST dataset obtained using (jD) , $q(jD_r)$, optimized $q(jD_r; \epsilon = 0)$ and $q_u(jD_r; \epsilon = 0)$.

Table 1: KL divergence achieved by EUBO (top row) and rKL (bottom row) with varying ϵ for airline dataset.

	10^{-11}	10^{-13}	10^{-20}	0
$KL[q_u(f_{X_u} D_r; \epsilon) \parallel q(f_{X_u} D_r)]$	219449	194300	138496	262971
$KL[q_v(f_{X_u} D_r; \epsilon) \parallel q(f_{X_u} D_r)]$	418:42	367:12	54345	45511

where $p_{\text{adj}}(y_x|f_x; \epsilon) = p(y_x|f_x)$ if $q(f_x; f_{X_u}|D) > \max_{f_{X_u}} q(f_x; f_{X_u}|D)$, and $p_{\text{adj}}(y_x|f_x; \epsilon) = 1$ otherwise. EUBO can be minimized using stochastic gradient descent with random subsets (i.e., mini-batches of size 10K) of D_e in each iteration. For rKL, we use the entire D_e in each iteration. Since $q_u(f_{X_u}|D_r; \epsilon)$, $q_v(f_{X_u}|D_r; \epsilon)$, and $q(f_{X_u}|D_r)$ in (10) [11] are all multivariate Gaussians, we can directly evaluate the performance of EUBO and rKL with varying ϵ through their respective $KL[q_u(f_{X_u}|D_r; \epsilon) \parallel q(f_{X_u}|D_r)]$ and $KL[q_v(f_{X_u}|D_r; \epsilon) \parallel q(f_{X_u}|D_r)]$ which, according to Table 1, are smaller than $KL[q(f_{X_u}|D) \parallel q(f_{X_u}|D_r)]$ of value 434409 (i.e., baseline representing no unlearning), hence demonstrating reasonable unlearning performance.

5 Conclusion

This paper describes novel unlearning methods for approximately unlearning a Bayesian model from a small subset of training data to be erased. Our unlearning methods are parsimonious in requiring only the approximate posterior belief of model parameters given the full data (i.e., obtained in model training with VI) and erased data to be available. This makes unlearning even more challenging due to two sources of inaccuracy in the approximate posterior belief. We introduce novel tricks of adjusted likelihood and reverse KL to curb unlearning in the region of model parameters with low approximate posterior belief where both sources of inaccuracy primarily occur. Empirical evaluations on synthetic and real-world datasets show that our proposed methods (especially reverse KL without adjusted likelihood) can effectively unlearn Bayesian models such as sparse GP and logistic regression from erased data. In practice, for the approximate posterior beliefs recovered by unlearning from erased data using our proposed methods, they can be immediately used in ML applications and continue to be improved at the same time by retraining with the remaining data at the expense of parsimony. In our future work, we will apply our proposed methods to unlearning more sophisticated Bayesian models like the entire family of sparse GP models [7, 8, 16, 17, 18, 19, 20, 22, 23, 25, 31, 32, 34] and deep GP models [33].

Broader Impact

As discussed in our introduction (Sec. 1), a direct contribution of our work to the society in this information age is to the implementation of personal data ownership (i.e., enforced by the General Data Protection Regulation in the European Union [24]) by studying the problem of machine unlearning for Bayesian models. Such an implementation can boost the confidence of users about sharing their data with an application/organization when they know that the trace of their data can be reduced/erased, as requested. As a result, organizations/applications can gather more useful data from users to enhance their service back to the users and hence to the society.

Our unlearning work can also contribute to the defense against data poisoning attacks (i.e., injecting malicious training data). Instead of retraining the tampered machine learning model from scratch to recover the quality of a service, unlearning the model from the detected malicious data may incur much less time, which improves the user experience and reduces the cost due to the service disruption.

In contrast, the ability to unlearn machine learning models may also open the door to new adversarial activities. For example, in the context of data sharing, multiple parties share their data to train a common machine learning model. An unethical party can deliberately share a low-quality dataset instead of its high-quality one. After obtaining the model trained on datasets from all parties (including the low-quality dataset), the unethical party can unlearn the low-quality dataset and continue to train the model with its high-quality dataset. By doing this, the unethical party achieves a better model than other parties in the collaboration. Therefore, the possibility of machine unlearning should be considered in the design of different data sharing frameworks.

Acknowledgments and Disclosure of Funding

This research/project is supported by the National Research Foundation, Singapore under its Strategic Capability Research Centres Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

- [1] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [2] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *J. American Statistical Association* 112(518):859–877, 2017.
- [3] L. Bourtole, V. Chandrasekaran, C. Choquette-Choo, H. Jia, A. Travers, B. Zhang, D. Lie, and N. Papernot. Machine unlearning. arXiv:1912.03817, 2019.
- [4] Y. Cao and J. Yang. Towards making systems forget with machine unlearning. *Proc. IEEE S&P*, pages 463–480, 2015.
- [5] J. Chen, N. Cao, B. K. H. Low, R. Ouyang, C. K.-Y. Tan, and P. Jaillet. Parallel Gaussian process regression with low-rank covariance matrix approximations. *Proc. UAI*, pages 152–161, 2013.
- [6] J. Chen, B. K. H. Low, P. Jaillet, and Y. Yao. Gaussian process decentralized data fusion and active sensing for spatiotemporal traffic modeling and prediction in mobility-on-demand systems. *IEEE Trans. Autom. Sci. Eng.* 12:901–921, 2015.
- [7] J. Chen, B. K. H. Low, and C. K.-Y. Tan. Gaussian process-based decentralized data fusion and active sensing for mobility-on-demand system. *Proc. RSS* 2013.
- [8] J. Chen, B. K. H. Low, C. K.-Y. Tan, A. Oran, P. Jaillet, J. M. Dolan, and G. S. Sukhatme. Decentralized data fusion and active sensing with mobile sensors for modeling and predicting spatiotemporal traffic phenomena. *Proc. UAI*, pages 163–173, 2012.
- [9] M. Du, Z. Chen, C. Liu, R. Oak, and D. Song. Lifelong anomaly detection through unlearning. In *Proc. CCS* pages 1283–1297, 2019.

- [10] D. Dua and C. Graff. UCI machine learning repository, 2017.
- [11] Y. Gal and M. van der Wilk. Variational inference in sparse Gaussian process regression and latent variable models—a gentle tutorial. *Xiv preprint arXiv, 1402*, 2014.
- [12] A. Ginart, M. Guan, G. Valiant, and J. Y. Zou. Making AI forget you: Data deletion in machine learning. In *Proc. NeurIPS* pages 3513–3526, 2019.
- [13] A. Golatkar, A. Achille, and S. Soatto. Eternal sunshine of the spotless net: Selective forgetting in deep neural networks. *Proc. CVPR*, 2020.
- [14] C. Guo, T. Goldstein, A. Hannun, and L. van der Maaten. Certified data removal from machine learning models. *arXiv:1911.03030*, 2019.
- [15] J. Hensman, N. Fusi, and N. D. Lawrence. Gaussian processes for big data. In *ICML*, pages 282–290, 2013.
- [16] Q. M. Hoang, T. N. Hoang, and B. K. H. Low. A generalized stochastic variational Bayesian hyperparameter learning framework for sparse spectrum Gaussian process regression. In *AAAI*, pages 2007–2014, 2017.
- [17] Q. M. Hoang, T. N. Hoang, B. K. H. Low, and C. Kingsford. Collective model fusion for multiple black-box experts. In *Proc. ICML*, pages 2742–2750, 2019.
- [18] T. N. Hoang, Q. M. Hoang, and B. K. H. Low. A unifying framework of anytime sparse Gaussian process regression models with stochastic variational inference for big data. In *ICML*, pages 569–578, 2015.
- [19] T. N. Hoang, Q. M. Hoang, and B. K. H. Low. A distributed variational inference framework for unifying parallel sparse Gaussian process regression models. In *Proc. ICML*, pages 382–391, 2016.
- [20] T. N. Hoang, Q. M. Hoang, B. K. H. Low, and J. P. How. Collective online learning of Gaussian processes in massive multi-agent systems. In *Proc. AAAI*, pages 7850–7857, 2019.
- [21] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. Improved variational inference with inverse autoregressive flow. *Proc. NeurIPS* pages 4743–4751, 2016.
- [22] B. K. H. Low, N. Xu, J. Chen, K. K. Lim, and E. B. Özgül. Generalized online sparse Gaussian processes with application to persistent mobile robot localization. In *Proc. ECML/PKDD Nectar Track*, pages 499–503, 2014.
- [23] B. K. H. Low, J. Yu, J. Chen, and P. Jaillet. Parallel Gaussian process regression for big data: Low-rank representation meets Markov approximation. In *Proc. AAAI*, pages 2821–2827, 2015.
- [24] A. Mantelero. The EU proposal for a general data protection regulation and the roots of the 'right to be forgotten'. *Computer Law & Security Review* 29(3):229–235, 2013.
- [25] R. Ouyang and B. K. H. Low. Gaussian process decentralized data fusion meets transfer learning in large-scale distributed cooperative perception. In *Proc. AAAI*, pages 3876–3883, 2018.
- [26] G. Papamakarios, T. Pavlakou, and I. Murray. Masked autoregressive flow for density estimation. In *Proc. NeurIPS* pages 2338–2347, 2017.
- [27] J. Quiñero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *JMLR*, 6:1939–1959, 2005.
- [28] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [29] D. J. Rezende and S. Mohamed. Variational inference with normalizing flows. In *Proc. ICML*, pages 1530–1538, 2015.

- [30] S. Schelter. “Amnesia” – Towards machine learning models that can forget user data very fast. In Proc. International Workshop on Applied AI for Database Systems and Applications, 2019.
- [31] T. Teng, J. Chen, Y. Zhang, and B. K. H. Low. Scalable variational Bayesian kernel selection for sparse Gaussian process regression. Proc. AAAI pages 5997–6004, 2020.
- [32] N. Xu, B. K. H. Low, J. Chen, K. K. Lim, and E. B. Özgül. GP-Localize: Persistent mobile robot localization using online sparse Gaussian process observation models. Proc. AAAI pages 2585–2592, 2014.
- [33] H. Yu, Y. Chen, Z. Dai, K. H. Low, and P. Jaillet. Implicit posterior variational inference for deep Gaussian processes. Proc. NeurIPS pages 14475–14486, 2019.
- [34] H. Yu, T. N. Hoang, B. K. H. Low, and P. Jaillet. Stochastic variational inference for Bayesian sparse Gaussian process regression. Proc. IJCNN 2019.

A Proof of Proposition 1

We first follow the proof of the log-sum inequality to prove the following inequality:

$$q_u(y|D_r) \log \frac{q_u(y|D_r)}{p(y|D_r)} \geq \int q_u(j|D_r) p(y_j) \log \frac{q_u(j|D_r)}{p(j|D_r)} dy \quad (11)$$

where $q_u(j|D_r)$, $E_{q_u(j|D_r)}[p(y_j)] = \int q_u(j|D_r) p(y_j) dy$ and $p(y|D_r)$, $E_{p(j|D_r)}[p(y_j)] = \int p(j|D_r) p(y_j) dy$.

Proof. Define the function $f(t) = t \log t$ which is convex. Then,

$$\begin{aligned} & \int q_u(j|D_r) p(y_j) \log \frac{q_u(j|D_r)}{p(j|D_r)} dy \\ &= \int p(j|D_r) p(y_j) f \left(\frac{q_u(j|D_r)}{p(j|D_r)} \right) dy \\ &= E_{p(j|D_r)}[p(y_j)] \int \frac{p(j|D_r) p(y_j)}{E_{p(j|D_r)}[p(y_j)]} f \left(\frac{q_u(j|D_r)}{p(j|D_r)} \right) dy \\ &= E_{p(j|D_r)}[p(y_j)] \int \frac{p(j|D_r) p(y_j) q_u(j|D_r)}{E_{p(j|D_r)}[p(y_j)] p(j|D_r)} dy \\ &= E_{p(j|D_r)}[p(y_j)] \int \frac{p(y_j) q_u(j|D_r)}{E_{p(j|D_r)}[p(y_j)]} dy \\ &= E_{p(j|D_r)}[p(y_j)] \int \frac{E_{q_u(j|D_r)}[p(y_j)]}{E_{p(j|D_r)}[p(y_j)]} dy \\ &= E_{q_u(j|D_r)}[p(y_j)] \log \frac{E_{q_u(j|D_r)}[p(y_j)]}{E_{p(j|D_r)}[p(y_j)]} \\ &= q_u(y|D_r) \log \frac{q_u(y|D_r)}{p(y|D_r)} \end{aligned}$$

where the inequality is due to Jensen's inequality. □

Then, integrating both sides of (11) w.r.t.

$$\begin{aligned} & \int q_u(y|D_r) \log \frac{q_u(y|D_r)}{p(y|D_r)} dy \geq \int \int q_u(j|D_r) p(y_j) \log \frac{q_u(j|D_r)}{p(j|D_r)} dy dy \\ & \int q_u(y|D_r) \log \frac{q_u(y|D_r)}{p(y|D_r)} dy \geq \int q_u(j|D_r) \int p(y_j) dy \log \frac{q_u(j|D_r)}{p(j|D_r)} dy \\ & \int q_u(y|D_r) \log \frac{q_u(y|D_r)}{p(y|D_r)} dy \geq \int q_u(j|D_r) \log \frac{q_u(j|D_r)}{p(j|D_r)} dy \\ & KL[q_u(y|D_r) \| p(y|D_r)] \geq KL[q_u(j|D_r) \| p(j|D_r)] : \end{aligned}$$

B Proof of Proposition 2

From (2),

$$\begin{aligned} \log p(D_{ej}|D_r) &= \log \frac{p(D_{ej}) p(j|D_r)}{p(j|D)} \\ &= \log \frac{q_u(j|D_r) p(D_{ej}) p(j|D_r)}{q_u(j|D_r) p(j|D)} \\ &= \log p(D_{ej}) + \log \frac{q_u(j|D_r)}{p(j|D)} - \log \frac{q_u(j|D_r)}{p(j|D_r)} : \end{aligned}$$

Then, taking an expectation of both sides w.r.t. $q_u(\cdot | D_r)$,

$$\begin{aligned} \log p(D_{ej} | D_r) &= \int q_u(\cdot | D_r) \log p(D_{ej} | \cdot) d\mu + \int q_u(\cdot | D_r) \log \frac{q_u(\cdot | D_r)}{p(\cdot | D_r)} d\mu - \int q_u(\cdot | D_r) \log \frac{q_u(\cdot | D_r)}{p(\cdot | D_r)} d\mu \\ &= \int q_u(\cdot | D_r) \log p(D_{ej} | \cdot) d\mu + \text{KL}[q_u(\cdot | D_r) \parallel p(\cdot | D_r)] - \text{KL}[q_u(\cdot | D_r) \parallel p(\cdot | D_r)] \\ &= U - \text{KL}[q_u(\cdot | D_r) \parallel p(\cdot | D_r)] : \end{aligned}$$

Therefore,

$$U = \log p(D_{ej} | D_r) + \text{KL}[q_u(\cdot | D_r) \parallel p(\cdot | D_r)] - \log p(D_{ej} | D_r)$$

since $\text{KL}[q_u(\cdot | D_r) \parallel p(\cdot | D_r)] \geq 0$. So, U is an upper bound of $\log p(D_{ej} | D_r)$.

C Bayesian Linear Regression

We perform unlearning of a simple Bayesian linear regression model: $y = ax^3 + bx^2 + cx + d + \epsilon$, where $a = 2, b = -3, c = 1, d = 0$ are the model parameters and the noise is $\epsilon \sim N(0, 0.05^2)$. Though the exact posterior belief of θ is known to be a multivariate Gaussian, we choose to use a low-rank approximation (i.e., multivariate Gaussian with a diagonal covariance matrix) and represent the approximate posterior beliefs of the model parameters with independent Gaussians so that the approximation is not exact.

Fig. 7a shows the remaining data D_r and erased data D_e . Note that the erased data D_e is informative to the approximate posterior beliefs of the model parameters and D_e are clustered. So, the difference between the samples drawn from predictive distribution $q(y | D_r)$ (Fig. 7b) vs. $q(y | D_r)$ (Fig. 7c) is large.

Table 2: KL divergences achieved by EUBO (left column) and rKL (right column) with varying synthetic linear regression dataset.

	$\text{KL}[q_u(\cdot D_r; \cdot) \parallel q(\cdot D_r)]$	$\text{KL}[q_r(\cdot D_r; \cdot) \parallel q(\cdot D_r)]$
0:5	0.1143	0.1012
0:1	0.0899	0.0600
0:0	2.6668	0.0158

From Table 2, the KL divergences achieved by EUBO and rKL with 0:1; 0:5 are smaller than $\text{KL}[q(\cdot | D) \parallel q(\cdot | D_r)]$ of value 0.1170 (i.e., baseline representing no unlearning), hence demonstrating reasonable unlearning performance. When 0, EUBO suffers from catastrophic unlearning, but rKL does not. The KL divergences in Table 2 also agree with the plots of samples drawn from the predictive distributions induced by EUBO and rKL in Fig. 7 by comparing with the samples drawn from the predictive distribution obtained using VI from retraining with D_r in Fig. 7c.

D Bimodal Posterior Belief

Let the posterior belief of model parameter θ given full data D be a Gaussian mixture (i.e., a bimodal distribution):

$$p(\cdot | D) = 0.5 \mathcal{N}(\cdot; 0; 1) + 0.5 \mathcal{N}(\cdot; 2; 1) \quad (12)$$

where $\mathcal{N}(\cdot; \mu; \sigma^2)$ is a Gaussian p.d.f. with mean μ and variance σ^2 . We deliberately choose the likelihood of the erased data D_e to be

$$p(D_{ej} | \cdot) = 1 + \frac{\mathcal{N}(\cdot; 2; 1)}{\mathcal{N}(\cdot; 0; 1)} \quad (13)$$

so that the posterior belief of θ given the remaining data D_r is a Gaussian:

$$p(\cdot | D_r) / \frac{p(\cdot | D)}{p(D_{ej} | \cdot)} = \mathcal{N}(\cdot; 0; 1) \quad (14)$$

(a) Dataset (b) Samples from $q(y_x | D)$ (c) Samples from $q(y_x | D_r)$

(d) EUBO with $\beta = 0.5$ (e) EUBO with $\beta = 0.1$ (f) EUBO with $\beta = 0$

(g) rKL with $\beta = 0.5$ (h) rKL with $\beta = 0.1$ (i) rKL with $\beta = 0$

Figure 7: Plots of (a) synthetic linear regression dataset with erased data (crosses) and remaining data D_r (dots), and samples from predictive distributions obtained using VI from (b) training with full data D and (c) retraining with D_r . Plots of samples from predictive distributions (d-f) $q(y_x | D_r)$ and (g-i) $q_v(y_x | D_r)$ induced, respectively, by EUBO and rKL with varying

where the proportionality is due to (2).

We assume to only have access to the likelihood of the erased data (11) in the exact posterior beliefs of y given the full data (12) and that given the remaining data (14) are not available. Instead, we have access to an approximate posterior belief $q(y | D)$ given the full data obtained using VI by minimizing $KL[q(y | D) \parallel p(y | D)]$ or, equivalently, maximizing the ELBO (Section 2):

$$q(y | D) = \mathcal{N}(y; 1.004, 1.390^2) \quad (15)$$

Given the likelihood $p(D_e | y)$ of the erased data (13) and the approximate posterior belief $q(y | D)$ given the full data (15), unlearning from D_e is performed using EUBO and rKL to obtain

$$q_u(y | D_r; \beta = 0) = \mathcal{N}(y; 0.060, 1.000^2) \quad \text{and} \quad q_v(y | D_r; \beta = 0) = \mathcal{N}(y; 0.062, 1.018^2);$$

respectively. Hence, both EUBO and rKL perform reasonably well since their respective $q_u(y | D_r; \beta = 0)$ and $q_v(y | D_r; \beta = 0)$ are close to $p(y | D_r) = \mathcal{N}(y; 0, 1)$ (14) when $p(y | D)$ is a bimodal distribution.

E Gaussian Process (GP) Classification with Synthetic Moon Dataset: Additional Details and Experimental Results

This section discusses the sparse GP model that is used in the classification of the synthetic moon dataset in Sec. 4.1. Let $f: \mathbb{R}^2 \rightarrow \{0, 1\}$ be the class label of $2 \times X$ \mathbb{R}^2 ; $y_x = 1$ denotes the 'blue' class plotted as blue dots in Fig. 4a. The probability $p(y_x)$ is defined as follows:

$$\begin{aligned} p(y_x = 1 | f_x) &= \frac{1}{1 + \exp(f_x)} \\ p(y_x = 0 | f_x) &= \frac{\exp(f_x)}{1 + \exp(f_x)} \end{aligned} \quad (16)$$

(a) $p(x_j|D)$ (b) $\frac{2}{x_j|D}$ (c) $p(x_j|D_r)$ (d) $\frac{2}{x_j|D_r}$

Figure 8: Plots of approximate posterior beliefs (a-b) $q(f_x | jD)$ and (c-d) $q(f_x | jD_r)$.

where f_x is modeled using a GP [28], that is, every finite subset $\mathcal{D} \subseteq \mathcal{X} \times \mathcal{Y}$ follows a multivariate Gaussian distribution. A GP is fully specified by its prior mean (i.e., assumed to be 0) and covariance k_{xx^0} , $\text{cov}(x; x^0)$, the latter of which can be defined by the widely-used squared exponential covariance function k_{xx^0} , $\frac{1}{2} \exp(-0.5k(x - x^0)^2)$ where $k = \text{diag}[l_1^2; l_2^2]$ and $\frac{1}{2}$ are the length-scale and signal variance hyperparameters, respectively. In this experiment, we set $l_1 = 1:56$, $l_2 = 1:35$, and $\frac{1}{2} = 4:74$.

We employ a sparse GP model, namely, the deterministic training condition (DTC) [27] approximation of the GP model with a set of 20 inducing inputs. These inducing inputs are randomly selected from \mathcal{X} and remain the same (and fixed) for both model training and unlearning. Given the latent function values (i.e., also known as inducing variables) f_{X_u} , $(f_x)_{\mathcal{X} \times \mathcal{Y}}$ at these inducing inputs, the posterior belief of the latent function value at a new input x is a Gaussian $p(f_x | f_{X_u}) = N(k_{xX_u} K_{X_u X_u}^{-1} f_{X_u}; k_{xx} - k_{xX_u} K_{X_u X_u}^{-1} k_{X_u x})$ where k_{xX_u} , $(k_{xx^0})_{\mathcal{X} \times \mathcal{Y}}$, $k_{X_u x} = k_{xX_u}^T$, and $K_{X_u X_u} = (k_{xx^0})_{\mathcal{X} \times \mathcal{Y}}$.

Using $p(f_x | f_{X_u})$ and $q(f_{X_u} | D)$, $N(x_u; x_u)$, it can be derived that the approximate posterior belief $q(f_x | jD)$ of f_x given full data D is also a Gaussian with the following respective posterior mean and variance:

$$x_j|D, k_{xX_u} K_{X_u X_u}^{-1} x_u; \tag{17}$$

$$\frac{2}{x_j|D}, k_{xx} - k_{xX_u} K_{X_u X_u}^{-1} k_{X_u x} + k_{xX_u} K_{X_u X_u}^{-1} x_u K_{X_u X_u}^{-1} k_{X_u x}; \tag{18}$$

The approximate posterior belief $q(f_x | jD_r)$ of f_x from retraining with remaining data D_r using VI (specifically, using $q(f_{X_u} | jD_r)$) can be derived in the same way as that of $q(f_x | jD)$.

The parameters x_u , x_u of the approximate posterior belief $q(f_{X_u} | jD)$ is optimized by maximizing the ELBO with stochastic gradient ascent (let f_{X_u} in (1) in Sec. 2):

$$E_{f_{X_u} \sim q(f_{X_u} | jD)} [\log p(D | f_{X_u}) - \log q(f_{X_u} | jD) + \log p(f_{X_u})]$$

where $p(D | f_{X_u})$ is computed using (16), (17) and (18).

Fig. 8 visualizes $q(f_x | jD)$ (Figs. 8a and 8b) and $q(f_x | jD_r)$ (Figs. 8c and 8d) whose corresponding predictive distributions $q(y_x = 1 | jD)$ and $q(y_x = 1 | jD_r)$ are shown in Figs. 4b and 4c, respectively. On the other hand, Figs. 9 and 10 visualize the approximate posterior beliefs $q_u(f_x | jD_r; \epsilon)$ and $q_v(f_x | jD_r; \epsilon)$ induced, respectively, by EUBO and rKL whose corresponding predictive distributions $q_u(y_x = 1 | jD_r)$ and $q_v(y_x = 1 | jD_r)$ are shown in Figs. 4f-k. Similar to the comparison between predictive distributions $q_u(y_x = 1 | jD_r)$ vs. $q(y_x = 1 | jD_r)$ in Sec. 4.1, it can be observed that the approximate posterior belief $q_u(f_x | jD_r; \epsilon = 10^{-9})$ induced by EUBO is similar to $q(f_x | jD_r)$ obtained using VI from retraining with D_r (compare Figs. 9c vs. 8c and Figs. 9d vs. 8d). However, $q_u(f_x | jD_r; \epsilon = 0)$ induced by EUBO differs from $q(f_x | jD_r)$ obtained using VI from retraining with D_r (compare Figs. 9e vs. 8c and Figs. 9f vs. 8d). On the other hand, both the approximate posterior beliefs $q_v(f_x | jD_r; \epsilon = 10^{-9})$ and $q_v(f_x | jD_r; \epsilon = 0)$ induced by rKL are similar to $q(f_x | jD_r)$ obtained using VI from retraining with D_r (compare Fig. 10 vs. Figs. 8c-d).

F A Note on Erasing Informative Data

In this section, we study the performance of our unlearning methods when erasing a large quantity of data or with different distributions of erased data (i.e., erasing the data randomly vs. deliberately

(a) Mean of $q_u(f_x | D_r; \epsilon = 10^{-5})$ (b) Variance of $q_u(f_x | D_r; \epsilon = 10^{-5})$

(c) Mean of $q_u(f_x | D_r; \epsilon = 10^{-9})$ (d) Variance of $q_u(f_x | D_r; \epsilon = 10^{-9})$

(e) Mean of $q_u(f_x | D_r; \epsilon = 0)$ (f) Variance of $q_u(f_x | D_r; \epsilon = 0)$

Figure 9: Plots of approximate posterior beliefs $q_u(f_x | D_r; \epsilon)$ induced by EUBO for varying ϵ .

(a) Mean of $q_v(f_x | D_r; \epsilon = 10^{-5})$ (b) Variance of $q_v(f_x | D_r; \epsilon = 10^{-5})$

(c) Mean of $q_v(f_x | D_r; \epsilon = 10^{-9})$ (d) Variance of $q_v(f_x | D_r; \epsilon = 10^{-9})$

(e) Mean of $q_v(f_x | D_r; \epsilon = 0)$ (f) Variance of $q_v(f_x | D_r; \epsilon = 0)$

Figure 10: Plots of approximate posterior beliefs $q_v(f_x | D_r; \epsilon)$ induced by rKL for varying ϵ .

erasing all data in a given class). Let us consider the experiment in Sec. 4.1 on the sparse GP model (i.e., the model parameters in (1) in Sec. 2 are inducing variables ϕ_u) in the classification of the synthetic moon dataset as it allows us to easily visualize both the approximate posterior beliefs of the latent function f_x and the predictive distributions of the output/observation. A key factor in affecting the performance of our unlearning methods in the above-mentioned scenarios is the difference between the approximate posterior belief of model parameters/inducing variables given remaining data D_r vs. that given full data D . We quantify such a difference by how much the erased data reduces the entropy of model parameters/inducing variables given remaining data D_r :

$$I, H(f_{X_u}|D_r) - H(f_{X_u}|D) = \int q(f_{X_u}|D_r) \log q(f_{X_u}|D_r) df_{X_u} - \int q(f_{X_u}|D) \log q(f_{X_u}|D) df_{X_u} \quad (19)$$

Note that (19) is not the same as the mutual information (i.e., information gain) between y_{D_e} and $(y_x)_{(x;y_x) \in D_e}$ given D_r , which is equal to $H(f_{X_u}|D_r) - E_{p(y_{D_e}|D_r)} [H(f_{X_u}|D_r; y_{D_e})]$ with an expensive-to-evaluate expectation term. Furthermore, the outputs/observations are known from D_e . These therefore prompt us to choose (19) as the measure of how much the erased data D_e reduces the entropy of model parameters/inducing variables given remaining data D_r .

We investigate 4 different scenarios in the order of increasing

1. Randomly selected D_e ($I = 0:27$): The erased data of size $|D_e| = 20$ are randomly selected from D . Hence, they are not necessarily near the decision boundary. D_e does not reduce the entropy of model parameters/inducing variables given D_r much;
2. Partially 'yellow' D_e ($I = 1:59$): The erased data of size $|D_e| = 30$ are labeled with the 'yellow' class and comprise inputs with the largest possible first component. Such a choice ensures that the erased data group together to cover a part of the decision boundary, as shown in Fig. 11d;
3. Largely 'yellow' D_e ($I = 2:06$): The erased data of size $|D_e| = 40$ are labeled with the yellow class and comprise inputs with the largest possible first component. As the quantity of the erased data increases from 30 (i.e., partially 'yellow' D_e) to 40, D_e covers a larger part of the decision boundary (compare Figs. 11g vs. 11d); and
4. Fully 'yellow' D_e ($I = 3:86$): The erased data of size $|D_e| = 50$ comprise all data in the yellow class. In this case, D_e reduces the entropy of the model parameters/inducing variables f_{X_u} given D_r the most when compared to the above scenarios.

As I increases, the difference between the approximate posterior belief given remaining data D_r vs. that given full data D increases. Though it is difficult to visualize such a difference directly, Proposition 1 tells us that this difference can be alternatively understood by comparing the predictive distributions $q(y_x = 1|D_r)$ in Table 3 vs. $q(y_x = 1|D)$ in Fig. 4b.

Fig. 11 shows results of averaged KL divergences (i.e., performance metric described in Sec. 4) achieved by EUBO, rKL, and $q(f_{X_u}|D)$ over D_r and D_e for the 4 scenarios above. Table 3 also analyzes the performance of our unlearning methods qualitatively by plotting the means of the approximate posterior beliefs $q(f_x|D_r; \cdot)$ and $q(f_x|D; \cdot)$ induced, respectively, by EUBO and rKL with the corresponding predictive distributions $q(y_x = 1|D_r)$ and $q(y_x = 1|D)$, together with the mean of the approximate posterior belief $q(f_x|D_r)$ with the corresponding predictive distribution $q(y_x = 1|D_r)$ obtained using VI from retraining with remaining data D_r . The following observations result:

Fig. 11 shows that as I increases across the 4 scenarios, the averaged KL divergence between $q(y_x|D)$ vs. $q(y_x|D_r)$ over D_r and D_e (i.e., baseline labeled as 1) generally increases.

In the scenario of randomly selected D_e (i.e., I is small), we expect the difference between the predictive distributions $q(y_x|D)$ vs. $q(y_x|D_r)$ over D_r and D_e to be small, which is reflected in the very small averaged KL divergences of about 0.002 and 0.004 achieved by $q(f_{X_u}|D)$ (i.e., baseline labeled as 1) in Figs. 11b and 11c, respectively. It can also be observed that though EUBO and rKL with 2×10^{-5} ; 10^{-9} achieve smaller averaged KL divergences than that of $q(f_{X_u}|D)$ (i.e., baseline), EUBO's averaged KL divergence increases beyond that of the baseline when $\epsilon = 0$, but remains very small. As a result, the first row in Table 3 shows that when $\epsilon = 10^{-9}$ or $\epsilon = 0$, the predictive distributions

$q_u(y_x = 1 | D_r)$ and $q_v(y_x = 1 | D_r)$ induced, respectively, by EUBO and rKL are similar to $q(y_x = 1 | D_r)$ obtained using VI from retraining with D_r . Hence, we can conclude that both EUBO and rKL perform reasonably well in this scenario, even when $\epsilon = 0$.

In the scenarios of partially and largely 'yellow' D_e , I is much larger than that in the scenario of randomly selected D_e . So, we expect an increase in the difference between the predictive distributions $q(y_x | D)$ vs. $q(y_x | D_r)$ over D_r and D_e . It can be observed from Figs. 11e-f and 11h-i that when $\epsilon = 0$, EUBO performs poorly as its averaged KL divergence is larger than that of $q(f_{x_u} | D)$ (i.e., baseline labeled 'all'), while rKL performs well as its averaged KL divergence is much smaller than that of the baseline. On the other hand, when $\epsilon = 10^{-9}$, both EUBO and rKL perform well, which can also be observed from the second and third rows of Table 3. These plots also show that while the predictive distributions $q(y_x = 1 | D_r)$ induced by rKL with $\epsilon = 10^{-9}$ are not as similar to $q(y_x = 1 | D_r)$ as $q_u(y_x = 1 | D_r)$ induced by EUBO with $\epsilon = 10^{-9}$, the performance of rKL with $\epsilon = 0$ is more robust.

In the scenario of fully 'yellow' D_e (i.e., I is largest), the difference between the predictive distributions $q(y_x | D)$ vs. $q(y_x | D_r)$ over D_r and D_e is larger than that in the above scenarios. Except for EUBO with $\epsilon = 0$, the predictive distributions $q_u(y_x | D_r)$ and $q_v(y_x | D_r)$ induced, respectively, by EUBO and rKL are closer to $q(y_x | D)$ than $q(f_{x_u} | D)$ as they achieve smaller averaged KL divergences than that of $q(f_{x_u} | D)$, as shown in Figs. 11k-l. However, the fourth row of Table 3 shows that both EUBO and rKL do not perform that well. Nevertheless, it can be observed that when $\epsilon = 0$, the predictive distribution $q_v(y_x = 1 | D_r)$ induced by rKL is still usable while $q_u(y_x = 1 | D_r)$ induced by EUBO is useless.

To summarize, when only an approximate posterior belief $q(f_{x_u} | D)$ of model parameters $\theta = f_{x_u}$ given full data D (i.e., obtained in model training with VI) is available, both EUBO and rKL can perform well if the difference between the approximate posterior belief of model parameters given remaining data D_r vs. that given full data D is sufficiently small. In practice, this is expected due to the small quantity of erased data and redundancy in real-world datasets. In the case where the erased data is highly informative, the approximate posterior belief $q(f_{x_u} | D_r; \epsilon = 0)$ induced by rKL remains usable by being close to $q(f_{x_u} | D)$ and hence sacrificing its unlearning performance. On the other hand, EUBO may suffer from poor unlearning performance when ϵ is too small.

The above remark highlights the limitation of our unlearning methods when the erased data is highly informative and only the approximate posterior belief $q(f_{x_u} | D)$ is available. Such a limitation is due to the lack of information about the difference between the exact posterior belief $p(\theta | D)$ vs. the approximate one $q(\theta | D)$ (Sec. 3.3), which motivates future investigation into maintaining additional information about this difference during the model training with VI to improve the unlearning performance. In practice, an ML application may require an unlearning method to be time-efficient in order to satisfy the constraint on the response time to a user's request for her data to be erased while not rendering the model useless (e.g., due to catastrophic unlearning). After processing the user's request, the ML application can continue to improve the approximate posterior belief recovered by unlearning from erased data (i.e., using our proposed EUBO or rKL) by retraining with the remaining data at the expense of parsimony (i.e., in terms of time and space costs).

One may wonder how our unlearning methods can handle multiple users' request arriving sequentially over time. To avoid approximation errors from accumulating, we can adopt the approach of unlearning by aggregating all the (past and new) users' erased data D_e into performing unlearning (i.e., using only $q(\theta | D)$ and D_e) as and when necessary. As expected, our unlearning methods can perform well, provided that the aggregated erased data remains sufficiently small or contains enough redundancy.

G Logistic Regression with Fashion MNIST Dataset: Additional Experimental Results

In this section, we will present the following:

Additional visualizations of the class probabilities for image D_i evaluated at the mean of the approximate posterior beliefs obtained using EUBO and rKL with $\epsilon = 0$ in Fig. 13, and

$I = 0:27$
 (a) Randomly selected D_e (b) D_r (c) D_e

$I = 1:59$
 (d) Partially `yellow' D_e (e) D_r (f) D_e

$I = 2:06$
 (g) Largely `yellow' D_e (h) D_r (i) D_e

$I = 3:86$
 (j) Fully `yellow' D_e (k) D_r (l) D_e

Figure 11: Plots of (a,d,g,j) synthetic moon dataset with erased D_e (crosses) and remaining data D_r (dots) in 4 different scenarios. Graphs of averaged KL divergence as achieved by EUBO, reverse KL (rKL), and $q(jD)$ (i.e., baseline labeled as $q(jD)$) over D_r and D_e in the following 4 scenarios: (b-c) randomly selected D_e , (e-f) partially `yellow' D_e , (h-i) largely `yellow' D_e , and (k-l) fully `yellow' D_e .

Table 3: Plots of the mean of approximate posterior beliefs $q(f_x | D_r)$ with the corresponding predictive distribution $q(y_x = 1 | D_r)$ obtained using VI from retraining with remaining data, and also the means of approximate posterior beliefs $q(f_x | D_r; \cdot)$ and $q_v(f_x | D_r; \cdot)$ induced, respectively, by EUBO and rKL with the corresponding predictive distributions $q(y_x = 1 | D_r)$ and $q_v(y_x = 1 | D_r)$ for $\sigma^2 \in [10^{-9}; 0]$. The 1-st, 2-nd, 3-rd, and 4-th rows correspond to the following respective scenarios: randomly selected D_e , partially 'yellow' D_e , largely 'yellow' D_e , and fully 'yellow' D_e .

Dataset	Retrained		EUBO		rKL	
	Mean $q(f_x D_r)$	$q(y_x = 1 D_r)$	Mean $q(f_x D_r; \cdot)$	$q_v(y_x = 1 D_r)$	Mean $q_v(y_x = 1 D_r)$	$q_v(y_x = 1 D_r)$
						$= 10^{-9}$
						$= 0$
						$= 10^{-9}$
						$= 0$
						$= 10^{-9}$
						$= 0$
						$= 10^{-9}$
						$= 0$

(a) D_r (b) D_e (c) D_r (d) D_e

Figure 12: Graphs of averaged KL divergence vs. achieved by EUBO, rKL, and $q(\mathbf{j}|D)$ (i.e., baseline labeled as \mathbf{all}) over D_r and D_e for the fashion MNIST dataset. The approximate posterior beliefs of the model parameters/weights are represented by (a-b) independent Gaussians (i.e., diagonal covariance matrices) and (c-d) multivariate Gaussians (i.e., full covariance matrices).

Comparison of the unlearning performance obtained using approximate posterior beliefs modeled with independent Gaussians (i.e., diagonal covariance matrices) vs. that modeled with multivariate Gaussians (i.e., full covariance matrices).

Fig. 13 shows the class probabilities for the images \mathbf{S}_i reevaluated at the mean of the approximate posterior beliefs with $\beta = 0$. Figs. 13a-d and 13g show that rKL induces the highest class probability for the same class as that of $q(\mathbf{j}|D_r)$. In Figs. 13e-f and 13h, the class probabilities obtained using optimized $q(\mathbf{j}|D_r; \beta = 0)$ resemble that obtained using $q(\mathbf{j}|D)$, though the probability of the correct class is reduced due to unlearning.

Fig. 12 shows the averaged KL divergences of EUBO, rKL, and $q(\mathbf{j}|D)$ where the approximate posterior beliefs are modeled with independent Gaussians (i.e., diagonal covariance matrices) in Figs. 12a-b and multivariate Gaussians (i.e., full covariance matrices) in Figs. 12c-d. It can be observed that the averaged KL divergences between $q(\mathbf{j}|D)$ vs. $q(\mathbf{y}_x|\mathbf{j}|D_r)$ over D_r and D_e (i.e., baselines labeled as \mathbf{all}) decrease when multivariate Gaussians with full covariance matrices are used to model the approximate posterior beliefs instead (compare the baselines labeled as Figs. 12c-d vs. that in Figs. 12a-b). Furthermore, in such a case, the unlearning performance of both EUBO and rKL improve as their averaged KL divergences are not as large (relative to the baselines) as that using independent Gaussians.

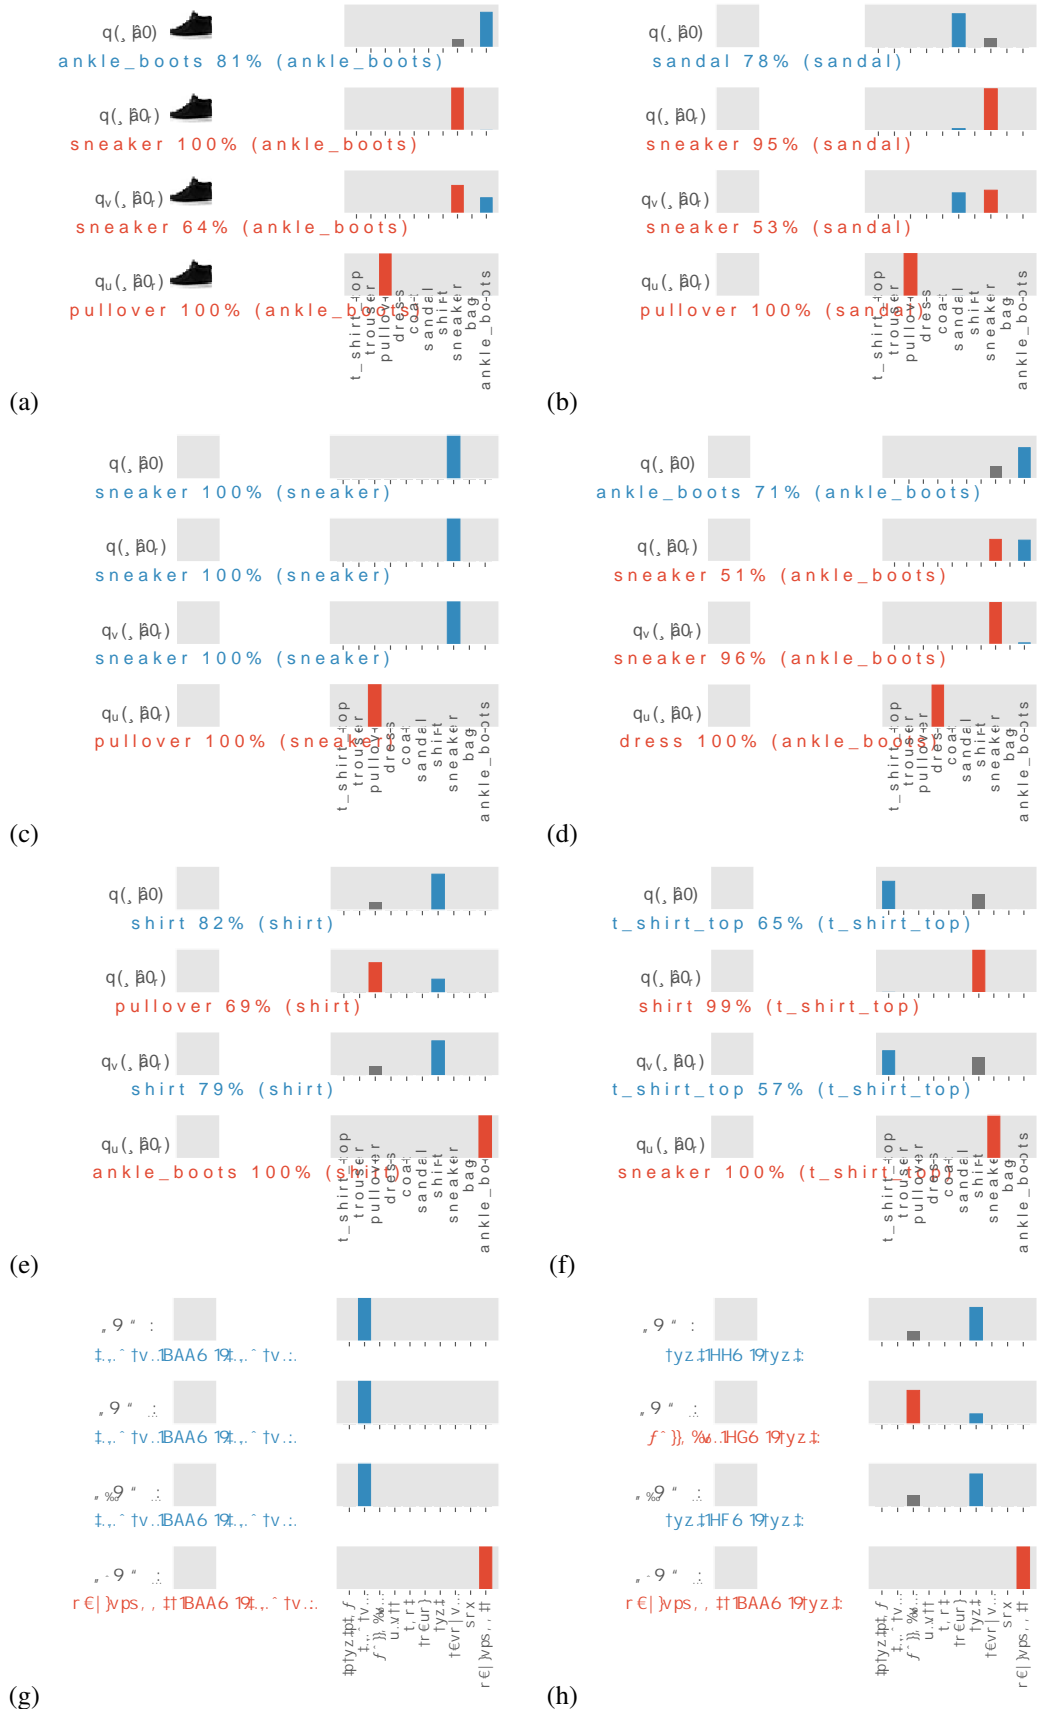


Figure 13: Plots of class probabilities for images in D_r obtained using $q(\cdot | D)$, $q(\cdot | D_r)$, optimized $q_v(\cdot | D_r; \lambda = 0)$ and $q_u(\cdot | D_r; \lambda = 0)$.