Continual Learning from the Perspective of Compression

Anonymous Authors[1]

Abstract

Connectionist models such as neural networks suffer from catastrophic forgetting. In this work, we study this problem from the perspective of information theory and define forgetting as the increase of description lengths of previous data when they are compressed with a sequentially learned model. In addition, we show that continual learning approaches based on variational posterior approximation and generative replay can be considered as approximations to two prequential coding methods in compression, namely, the Bayesian mixture code and maximum likelihood (ML) plug-in code. We compare these approaches in terms of both compression and forgetting and empirically study the reasons that limit the performance of continual learning methods based on variational posterior approximation. To address these limitations, we propose a new continual learning method that combines ML plug-in and Bayesian mixture codes.

1. Introduction

When sequentially learning a non-stationary data stream, neural networks usually suffer from the problem of catastrophic forgetting when the data from the past are not available anymore (McCloskey & Cohen, 1989). Recently, this problem has received renewed attention within the field of deep learning, and a flurry of methods (see (Parisi et al., 2019) for a survey) have been proposed to address this issue in order to achieve the goal of continual learning (CL) (Ring, 1998).

Intuitively speaking, forgetting can be understood as loss of information about past data. Previous studies on this problem have indirectly measured forgetting by degradation of task-specific performance metrics such as classification accuracy (Lopez-Paz & Ranzato, 2017), regression error (Javed & White, 2019) or the Frechet Inception Distance (FID) (Lesort et al., 2019). In order to explicitly study the problem of information loss, here we place continual learning within the framework of information theory and Minimum Description Length (MDL) (Grünwald, 2007).

According to the MDL principle, any regularity within a dataset can be used for compression, and the best model for a given dataset is the one that results in the minimum *total* description length of the dataset together with the model itself. Therefore, MDL formalizes the Occam's Razor principle for machine learning and provides a criterion for model selection. Other inspirations and foundations for MDL include the "comprehension is compression" hypothesis by Chaitin (Chaitin, 2002) and Solomonoff's theory of inductive inference (Solomonoff, 1964), which proved that the most likely explanation of some observations of the world is the smallest computer program (in terms of its Kolmogrov Complexity (Kolmogorov, 1968)) that generates these observations, assuming the world is generated by an unknown computer.

The compression perspective has provided insights in various problems in machine learning, including generalization (Arora et al., 2018), unsupervised modeling (Ollivier, 2014) and information bottleneck (Tishby & Zaslavsky, 2015). Recently, (Blier & Ollivier, 2018) showed that one particular coding method from the MDL toolbox called *prequential coding* outperforms the other coding methods at explaining why deep learning models generalize well despite their large number of parameters, displaying strong correlation between generalization and compression.

In this work, we show that the compression perspective can also shed light on continual learning. Our contributions are:

- We provide a formal definition of forgetting based on information theory, which allows us to evaluate continual learning methods for generative models in a principled way.
- We show that two major paradigms of continual learning (variational continual learning and generative replay) can be interpreted as approximations of two prequential coding methods (Bayesian mixture coding and ML plug-in coding) in the MDL framework, thus establishing a connection between continual learning and compression.
- We compare different CL methods for continual generative modeling using the proposed definition of forgetting and empirically study the limitations of the

variational continual learning approaches.

- We introduce a new CL method that combines the prediction strategies of the two prequential coding methods and show that it improves over the variational continual learning.

2. Minimum Description Length

Formally, given a dataset $x^n := \{x_1, \dots, x_n\}$ where each x_i is from a space of observation $\mathcal{X}$, the overall description length of x^n under the model with parameters θ can be separated into two parts:

$$L(x^n) = L(x^n|\theta) + L(\theta) \tag{1}$$

The first part $L(x^n|\theta)$ describes the codelength (number of bits) of x^n when compressed with the help of a model parametrized by θ, and the second part $L(\theta)$ corresponds to the codelength of the model itself. A model with more capacity can fit the data better, thus resulting in shorter codelength of the first part. However, higher capacity might also lead to longer description length of the model itself. Therefore, the best explanation for x^n, according to MDL, is given by the model that minimizes the total length $L(x^n)$. For probabilistic models, the code length $L(x^n|\theta)$ can be computed using the Shannon-Fano code.

Shannon-Fano Code Let $p(x|\theta)$ be a probability density (or mass) function over $\mathcal{X}$ determined by parameters θ, then for any $x \in \mathcal{X}$, there exists a prefix code that losslessly compresses x with code length

$$L(x|\theta) = -\log_2 p(x|\theta) \tag{2}$$

In this work, we use the codelength function L only as a metric to measure the amount of information and do not care about the actual encoding of data, hence we adopt the idealized view of code lengths which does not require them to be integers.

Bayesian Code One way to encode the parameters θ is to consider these parameters also as random variables and use the following Bayesian coding scheme:

$$\log p(x^n) = \log \int_\theta p(x^n|\theta)p(\theta) \tag{3}$$

$$\geq \int_\theta q(\theta|x^n) \log p(x^n|\theta) d\theta - D_{\mathrm{KL}}\left[q(\theta|x^n) \| p(\theta)\right] \tag{4}$$

where (4) is the evidence lower bound (ELBO) used in variational Bayesian methods. The equality holds when $q(\theta|x^n) = p(\theta|x^n)$. The first term corresponds to $L(x^n|\theta)$, describing the expected code length of x^n if encoded with θ sampled from the approximate posterior $q(\theta|x^n)$. The second term corresponds to $L(\theta)$, describing the code length of θ using the bits-back coding (Hinton & Van Camp, 1993) with a fixed prior distribution $p(\theta)$.

Prequential (Online) Code Another way to compute the total description length $L(x^n)$ without explicitly separating $L(x^n)$ into two parts is to use the prequential or online code, which considers a dataset x^n as a sequence of observations and encode each observation x_t at time t by predicting its value based on past observations x^{t-1}. Note that for any x^t, we can write its Shannon-Fano code as a sum of negative log conditional probabilities:

$$-\log p(x^n) = \sum_{t=1}^{n} -\log p(x_t|x^{t-1}) \tag{5}$$

Hence the total description length can be computed by accumulating the code length $-\log p(x_t|x^{t-1})$ at every step t using a prediction strategy. Formally, a prediction strategy is a function $S : \bigcup_{0 \leq t \leq n} \mathcal{X}^t \to \mathcal{P}_{\mathcal{X}}$ that maps any initial observations x^{t-1} to a distribution on $\mathcal{X}$ corresponding to the conditional model $p(X_t|x^{t-1})$. For example, *Bayesian mixture code* uses the following prediction strategy:

$$p(X_t|x^{t-1}) = \int_\theta p(X_t|\theta)p(\theta|x^{t-1}) \tag{6}$$

where $p(\theta|x^{t-1})$ is the posterior of model parameters given all previous data x^{t-1}.
Another example of prequential coding is the *maximum likelihood (ML) plug-in code*, which encodes each observation by

$$p(X_t|x^{t-1}) = p(X_t|\theta(x^{t-1})) \tag{7}$$

where $\theta(x^{t-1})$ is given by a maximum likelihood estimator:

$$\theta(x^{t-1}) = \arg\max_\theta \log p(x^{t-1}|\theta) \tag{8}$$

As we will show in the next section, this online incremental view of prequential coding is well-aligned with the setting of continual learning.

3. Continual Learning

Continual learning studies the scenario where data are provided at different time in a sequential manner, and the learning algorithm incrementally updates the model parameters as more data are presented. However, due to reasons such as privacy or limited memory budget, the learning algorithm may not always store all the data x^{t-1} it has seen so far. Alternatively, even if all data are stored, naively retraining on all previous data will increase the run time per step over time, which may not be affordable for long data streams. On the other hand, if the model is only trained on the data

presented recently, it usually suffers from "catastrophic forgetting". This problem is especially severe when the data stream is non-stationary. Here we provide a definition of forgetting that allows us to precisely measure the amount of information forgotten.

3.1. An Information-Theoretic Definition of Forgetting

We can see from (1) that when the data x^n is compressed with the help of a model, the information of x^n is divided into two parts, with one part contained in the code words of x^n and the other part contained in the model. Hence, the model alone is not enough to recover the data losslessly, it requires also the code words of the compressed data, which can be understood as the information of x^n not captured by the model. This intuition leads to our definition of forgetting: for any data x, the amount of forgetting about x caused by changing the model parameters from θ_i to θ_j is given as the increase of the codelength of x:

$$\mathcal{F}[\theta_j/\theta_i](x) := \max(L(x|\theta_j) - L(x|\theta_i), 0) \tag{9}$$

In other words, the number of extra bits required for the model to losslessly recover x is the amount of information the model has forgotten about x. The $\max(\cdot, 0)$ function in (9) is used to prevent "negative forgetting", if the codelength does not increase, there is no forgetting.

In the prequential setting, at any time t, we can measure the *cumulative average forgetting* of a learning algorithm over the entire sequence x^t seen so far by the following metric:

$$\mathcal{C}(x^t) := \frac{1}{t}\sum_{i=1}^{t} \mathcal{F}[\theta_t/\theta_i](x_i) \tag{10}$$

$$= \frac{1}{t}\sum_{i=1}^{t} \max(L(x_i|\theta_t) - L(x_i|\theta_i), 0) \tag{11}$$

where θ_i are the parameters given by the learning algorithm at time i, after seeing x^i.

3.2. Prequential Interpretation of CL Approaches

A primary goal of continual learning is to overcome or to alleviate catastrophic forgetting. Many CL methods have been proposed for this purpose (Parisi et al., 2019). Some common paradigms emerged from these methods. One paradigm is based on the sequential nature of Bayesian inference and it applies variational methods to approximating the parameter posterior. The approximated posterior is then used as the prior for future learning. Examples of this paradigm include Bayesian Online Learning (Opper, 1998), Variational Continual Learning (VCL) (Nguyen et al., 2018) and Online Structured Laplace Approximations (Ritter et al., 2018). Some regularization-based CL methods such as Elastic Weight Consolidation (EWC) (Kirkpatrick et al., 2017) and Uncertainty-guided Continual Bayesian Neural Networks (UCB) (Ebrahimi et al., 2020) are also motivated by this idea. Another paradigm combats forgetting by training a generative model of past data. When new data are provided later, the samples from the generative model are replayed instead of the original data. Representative works of this class are Deep Generative Replay (DGR) (Shin et al., 2017), Memory Replay GAN (Wu et al., 2018) and Continual Unsupervised Representation Learning (CURL) (Rao et al., 2019). Farquhar & Gal (2019) proposed a unifying Bayesian view that encompasses both paradigms. From the perspective of MDL, these two types of approaches can be seen as approximations of the two prequential coding methods introduced before. We discuss their relationships below using VCL and Generative Replay as examples.

VCL The first paradigm corresponds to the Bayesian mixture code, which uses (6) as its prediction strategy. Since the exact posterior $p(\theta|x^{t-1})$ in (6) depends on all previous data x^{t-1}, which might not always be available in the continual learning scenario, VCL therefore uses a variational model $q(\theta|x^{t-1})$ to approximate the real posterior $p(\theta|x^{t-1})$. By the Bayes' rule, we have

$$p(\theta|x^t) = \frac{p(x_t|\theta)p(\theta|x^{t-1})}{Z_t} \tag{12}$$

where $Z_t := \int_\theta p(x_t|\theta)p(\theta|x^{t-1})$ is a normalizing factor that does not depend on θ. Therefore, $q(\theta|x^t)$ can be updated recursively by minimizing the following KL divergence at every step t:

$$q(\theta|x^t) = \arg\min_{q(\theta)} D_{\mathrm{KL}}(q(\theta)||\frac{1}{Z_t}p(x_t|\theta)q(\theta|x^{t-1})) \tag{13}$$

The first approximate posterior $q(\theta|x^0)$ is usually a fixed prior chosen before learning starts: $q(\theta|x^0) = p(\theta)$. Since Z_t does not depend on θ, it is not required for the optimization above.

Replay The second paradigm, on the other hand, corresponds to the ML plug-in code, which uses the prediction strategy defined in (7) and (8). The objective function in (8) requires all previous data x^{t-1}. In order to relax this requirement for continual learning, replay-based methods approximate the real objective function using a generative model $p(x|\theta)$. Assuming $x_i \in x^t$ are independent given θ, we have:

$$\begin{aligned}\log p(x^t|\theta) &= \sum_{i=1}^{t} \log p(x_i|\theta) = \log p(x_t|\theta) + \sum_{i=1}^{t-1} \log p(x_i|\theta) \\ &\approx \log p(x_t|\theta) + (t-1)\mathbb{E}_{p(x|\theta_{t-1}))}[\log p(x|\theta)] \\ &= \log p(x_{t-1}|\theta) + (t-1)\int_x p(x|\theta_{t-1}) \log p(x|\theta)\end{aligned} \tag{14}$$

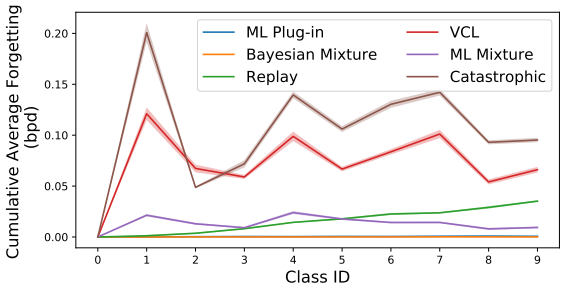

Figure 1: Cumulative average forgetting (defined in (10)) measured after learning each digit in MNIST. The resulting codelength was divided by the number of pixels in an image to get bits per dimension (bpd). The means and standard errors of the mean (SEM) were computed over ten random mini-batches of 32 images from each class. The shaded region correspond to 3 times of SEM. Here "Catastrophic" corresponds to the method that trains the model only on images from the current class. The "ML Mixture" method is introduced in Sec. 5. Both ML plug-in and Bayesian Mixture have near zero forgetting, so their curves are overlapping with each other.

where θ_{t-1} are the parameters of the generative model at time $t-1$ and the expected value can be computed by Monte Carlo sampling. In this way the generative model can be recursively updated by the following rule:

$$\theta_t = \arg\max_{\theta} \log p(x_t|\theta) + (t-1) \int_x p(x|\theta_{t-1}) \log p(x|\theta) \tag{15}$$

Measuring Forgetting and Compression We empirically compare the two prequential coding methods and their continual learning approximations in terms of both forgetting and compression on the MNIST dataset. Data are presented in a class-incremental fashion: first only images of digit zero, then only images of digit one, and so on. Throughout the paper, we use a variational autoencoder (VAE)(Kingma & Welling, 2013) for encoding and generative modelling. In this experiment, one observation corresponds to an entire class of images.

Fig.1 shows the cumulative average forgetting after learning each digit. Since ML plug-in and Bayesian Mixture methods are trained on all data seen so far, they have almost zero forgetting. VCL and Replay, on the other hand, do not have access to previous data, so information are lost over time. However, thanks to the variational posterior and generated samples, they achieve less forgetting than the Catastrophic baseline, which only trains the model on images of the current digit.

The results on prequential coding as defined in (5) are shown in Fig. 2. After all the images of one class are presented, the model is updated and used to encode all the images of the next class. The left plot shows the codelength of each class, since models are only trained on the images from previous classes, the results shown here correspond to the concept of forward transfer in continual learning. We can see that Bayesian Mixture and VCL have high codelength at the beginning, when there is no forgetting yet. The reason for this phenomenon is discussed in Sec. 4.1. The plot on the right shows the prequential codelength of the entire MNIST dataset measured in bits per dimension.

4. Limitations of VCL

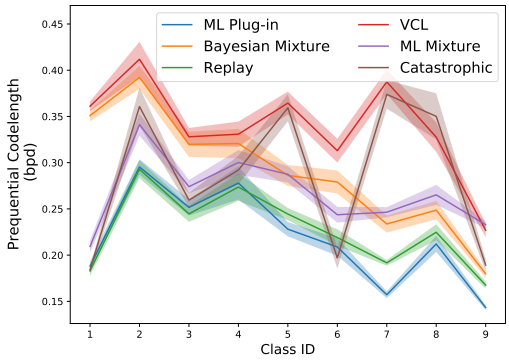

(a) Prequential codelength of images in class i, measured after the model has seen all images in class $i-1$. Since the images of the first class are encoded with a predefined model for all methods, their codelength does not affect the comparison results.

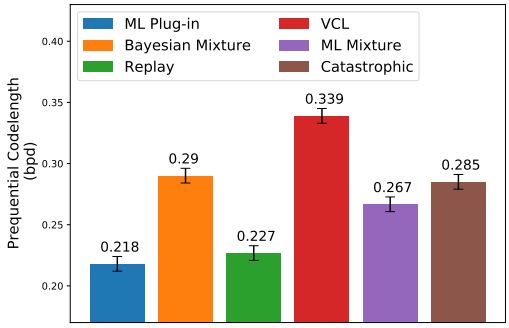

(b) The total prequential codelength of the MNIST dataset, averaged over number of samples.

Figure 2: Compression performance of different coding methods on the MNIST dataset. The shaded region correspond to 3 times of SEM.

It can be seen from both Fig.1 and 2, VCL performs poorly compared to generative replay. A systematic comparison

between VCL and replay on discriminative tasks was carried out by Farquhar & Gal (2018). VCL was reported to work similarly to generative replay on simple tasks like Permuted MNIST. On the more challenging Split MNIST task, VCL requires both a multi-head output layer and a coreset of past examples in order to be comparable with generative replay. For generative modeling, the results in Fig. 1 confirm the limitation of VCL at retaining information of past data. In this section, we analyze the reasons that limit the performance of VCL.

4.1. Overhead of Bayesian Mixture due to Prior

With deep generative models, the integration in Bayesian mixture code defined in (6) is intractable because $p(x|\theta)$ is a neural network generator. Therefore, the ELBO is used in practice as an approximation, which adapts $q(\theta)$ in (4) to minimize the total description length. We make two observations here. First, in VCL, the best achievable description length is predetermined by the selected prior $p(\theta)$. Second, compared to ML plug-in code, which directly minimizes the encoding length $L(x|\theta)$, the approximate Bayesian code has an extra KL regularization, indicating a longer description length $L(x|\theta)$ for the code words. Furthermore, since we do not explicitly encode data and keep their code words in continual learning, a larger encoding length means more information is lost about the data.

This overhead is negligible asymptotically if the data sequence is very long. This can be seen from the relative weights of the two terms in (4). The relative weight of the KL regularization diminishes when the data sequence grows longer. This is why we see a high initial codelength for Bayesian Mixture in Fig. 2 (a), but gradually it closes the gap to ML plug-in.

However, for a sequence with fixed length, the overhead due to prior can be significant and is very sensitive to the choice of the prior. Figure 3 compares the prequential codelength of Bayesian Mixture code using Gaussion priors with different standard deviation σ. One can see that good performance can be achieved only when σ is within a small region around 0.1. And the best description length is obtained when the prior is not used. Similar effects were observed for VCL in (Swaroop et al., 2019), where the authors showed the choice of prior variance has great impact on the performance of VCL.

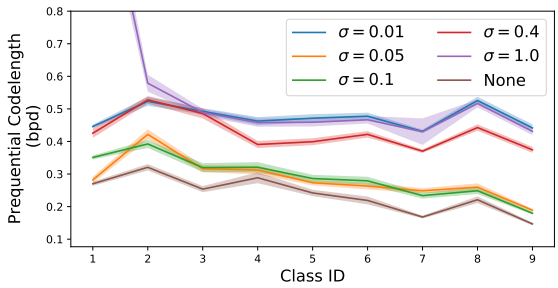

Figure 3: Prequential description length of Bayesian Mixture code using Gaussian priors with different standard deviations σ. Here, "None" means no prior is used, and the KL term in (4) is not optimized.

4.2. The Variational Gap

As we describe before, in continual learning, the true posterior $p(\theta|x^{t-1})$ is not available, thus VCL approximates it with a variational posterior $q(\theta|x^{t-1})$. The underlying assumption is that when their KL divergence is small enough: $D_{\mathrm{KL}}[q(\theta|x^{t-1})\|p(\theta|x^{t-1})] < \epsilon$, the following approximations are good enough for continual learning:

$$q(\theta|x^{t-1}) \approx p(\theta|x^{t-1}) \tag{16}$$

$$\mathcal{L}_t^{\mathrm{VCL}}(q) = \int_\theta q(\theta|x^t) \log p(X_t|\theta) - D_{\mathrm{KL}}[q(\theta|x^t)\|p(\theta|x^{t-1})] \tag{17}$$

$$\approx \int_\theta q(\theta|x^t) \log p(X_t|\theta) - D_{\mathrm{KL}}[q(\theta|x^t)\|q(\theta|x^{t-1})] \tag{18}$$

However, for VCL with a simple distribution family such as diagonal Gaussian, this assumption can hardly hold, as the real posterior is proportional to $p(x^{t-1}|\theta)p(\theta)$, which is usually complex and multimodal. Since the only difference between VCL and Bayesian Mixture code is the substitution of (17) by (18), the error between these two terms is entirely responsible for their large performance gap observed in Fig.1 and 2. One potential remedy to this problem is to use more complex distributions on the parameter space, for example, normalizing flow as a Bayesian hypernetwork (Krueger et al., 2017). However, unlike Gaussian distributions, the KL divergence of complex distributions usually cannot be calculated analytically. It is also unclear whether sampling in the parameter space to minimize $D_{\mathrm{KL}}[q(\theta|x^{t-1})\|p(\theta|x^{t-1})]$ would be any cheaper than sampling in the data space for replay.

Furthermore, the approximation error between (17) and (18) could be arbitrarily large even if $D_{\mathrm{KL}}[q(\theta|x^{t-1})\|p(\theta|x^{t-1})] < \epsilon$, which only states that $\log \frac{q(\theta|x^{t-1})}{p(\theta|x^{t-1})}$ is small in its expectation over the distribution $q(\theta|x^{t-1})$. As shown in (19), the approximation error is the expectation of the same log ratio over $q(\theta|x^{t-1})p(x_t|\theta)/Z_t$. Unfortunately, this distribution could be very different from $q(\theta|x^{t-1})$. It depends on $p(x_t|\theta)$, which is not available at the time $q(\theta|x^{t-1})$ is solved. One sufficient condition to bound the substitution error is $\max_\theta |\log \frac{q(\theta|x^{t-1})}{p(\theta|x^{t-1})}| < \epsilon$. However, this involves an

minimax optimization of the log ratio, which is a hard problem itself.

$$\Delta(\mathcal{L}_t^{\mathrm{VCL}}) = -\int_\theta q(\theta|x^t)\log\frac{q(\theta|x^{t-1})}{p(\theta|x^{t-1})} = -\int_\theta \frac{1}{Z_t}q(\theta|x^{t-1})p(x_t|\theta)\log\frac{q(\theta|x^{t-1})}{p(\theta|x^{t-1})} \tag{19}$$

5. ML Mixture Code

In order to overcome the two limitations discussed above, we can combine the prediction strategies of Bayesian mixture and ML plug-in using a hierarchical model with hyperparameters ϕ. For diagonal Gaussian distribution, ϕ stands for the mean and variance of θ. The resulting prediction strategy, which we call *maximum likelihood (ML) mixture*, is as follows:

$$p(X_t|x^{t-1}) = p(X_t|\phi(x^{t-1})) = \int_\theta p(X_t|\theta)p(\theta|\phi(x^{t-1})) \tag{20}$$

where $\phi(x^t)$ is learned to maximize the following objective:

$$\begin{aligned}\mathcal{L}_t^{\mathrm{MLM}}(\phi) &= \log p(x^t|\phi) = \sum_{i=1}^{t}\log\int_\theta p(x_i|\theta)p(\theta|\phi)\\ &= \sum_{i=1}^{t}\int_\theta q_t^*(\theta|x_i)\log\frac{p(x_i|\theta)p(\theta|\phi)}{q_t^*(\theta|x_i)}\\ &\geq \mathcal{L}_{t-1}^{\mathrm{MLM}} + \int_\theta q_t(\theta|x_t)\log\frac{p(x_t|\theta)p(\theta|\phi)}{q_t(\theta|x_t)}\\ &\quad - \sum_{i=1}^{t-1}\int_\theta q_{t-1}(\theta|x_i)\log\frac{p(\theta|\phi_{t-1})}{p(\theta|\phi)}\\ &\approx \mathcal{L}_{t-1}^{\mathrm{MLM}} + \int_\theta q_t(\theta|x_t)\log\frac{p(x_t|\theta)p(\theta|\phi)}{q_t(\theta|x_t)}\\ &\quad - (t-1)D_{\mathrm{KL}}[p(\theta|\phi_{t-1})\|p(\theta|\phi)]\end{aligned} \tag{21}$$

ϕ_{t-1} minimizes $\mathcal{L}_{t-1}^{\mathrm{MLM}}(\phi)$ and q_t^* denotes the theoretical optimal posterior that maximizes the ELBO, which means $q_t^*(\theta|x_i) = p(\theta|x_i, \phi_t)$. The approximation in (21) is based on the assumption that the previous objective $\mathcal{L}_{t-1}^{\mathrm{MLM}}$ is optimized good enough such that

$$p(\theta|\phi_{t-1}) \approx \frac{1}{(t-1)}\sum_{i}^{t-1} q_{t-1}(\theta|x_i) \tag{22}$$

For continual learning, in order to remove any dependency on previous data or previous approximate posterior, we optimize (21) sequentially in a greedy manner, namely, only the last two terms of (21) are optimized at time t:

$$\begin{aligned}\phi_t = &\arg\max_\phi \int_\theta q_t(\theta|x_t)\log\frac{p(x_t|\theta)p(\theta|\phi)}{q_t(\theta|x_t)}\\ &- (t-1)D_{\mathrm{KL}}[p(\theta|\phi_{t-1})\|p(\theta|\phi)]\end{aligned} \tag{23}$$

Like VCL, the optimization involved in ML Mixture code takes the form of a KL regularization in the parameter space. It addresses the two above-mentioned limitations of VCL. First, the prior overhead described in Sec. 4.1 is resolved because the mixture distribution is now optimized via maximum likelihood. Second, the approximation in (22) is more plausible than (16), because (16) tries to approximate a fixed multimodal distribution with a simple distribution, while (22) has optimizable components on both sides of the approximation.

The performance of ML Mixture code in terms of forgetting and compression are also presented in Fig.1 and 2. One can see that although both VCL and ML Mixture use diagonal Gaussian for the distribution of model parameters (corresponding to $q(\theta|x^{t-1})$ in VCL and $p(\theta|\phi)$ in ML Mixture), ML Mixture achieves much less forgetting and smaller prequential codelength.

5.1. The Intuitive Interpretation

The ML Mixture code is effectively performing ML plug-in coding with a VAE. Unlike normal VAEs, the latent code includes the entire parameter vector of the decoder. This means that the mapping from the latent code to the input space is predetermined by the decoder architecture. Intuitively, ML Mixture projects each data x_i from the input space to a local posterior distribution over the parameter space through the encoder $q(\theta|x_i)$. It then uses the prior distribution $p(\theta|\phi)$ to model the projected data in the parameter space. For each x_i, there exists a large number of simple distributions $q(\theta|x_i)$ that can be used as the encoder. Potentially, one can select the encoder so that the mirrored data points in the parameter space obeys a simple distribution, i.e. $\frac{1}{n}\sum_i^n q(\theta|x_i)$ is a simple distribution. In this case, $p(\theta|\phi)$ which is used to model the mixture distribution can also be kept simple. One supporting evidence for this possibility is the success of normal VAE/WAE(Tolstikhin et al., 2017). The prior distribution over the joint vector of latent code and the decoder parameters of VAE/WAE is a diagonal Gaussian, if we see the deterministic parameters as Gaussian random variables with zero variance.

5.2. Limitation and Future Improvements

There are a few limitations of the ML Mixture code that need to be solved before it is practical for continual learning. First, as we greedily optimize only the last two terms in (21) in a sequential manner, the previous approximate posterior

q_{t-1} is fixed once learned. This locks the prior distribution to a local region biased to the early tasks. As a result, ML Mixture code learns well and retains the knowledge from early experiences. However, the description length is only marginally improved at the later tasks (see Appendix A). A potential fix is to introduce certain level of generative replay to optimize the first term in (21). Second, both VCL and ML Mixture use approximate objectives to remove dependency on previous data. Although (22) could be more plausible than (16), the error introduced in the objective could still be arbitrarily large as pointed out earlier in Sec. 4.2. In the future, it would be preferable to develop an alternative objective for which the approximation error is bounded.

References

Arora, S., Ge, R., Neyshabur, B., and Zhang, Y. Stronger generalization bounds for deep nets via a compression approach. *arXiv preprint arXiv:1802.05296*, 2018.

Blier, L. and Ollivier, Y. The description length of deep learning models. In *Advances in Neural Information Processing Systems*, pp. 2216–2226, 2018.

Chaitin, G. J. On the intelligibility of the universe and the notions of simplicity, complexity and irreducibility. *arXiv preprint math/0210035*, 2002.

Ebrahimi, S., Elhoseiny, M., Darrell, T., and Rohrbach, M. Uncertainty-guided continual learning with bayesian neural networks. In *International Conference on Learning Representations*, 2020.

Farquhar, S. and Gal, Y. Towards robust evaluations of continual learning. *arXiv preprint arXiv:1805.09733*, 2018.

Farquhar, S. and Gal, Y. A unifying bayesian view of continual learning. *arXiv preprint arXiv:1902.06494*, 2019.

Grünwald, P. D. *The minimum description length principle*. The MIT Press, 2007.

Hinton, G. E. and Van Camp, D. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pp. 5–13, 1993.

Javed, K. and White, M. Meta-learning representations for continual learning. In *Advances in Neural Information Processing Systems*, pp. 1818–1828, 2019.

Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

Kolmogorov, A. N. Three approaches to the quantitative definition of information. *International journal of computer mathematics*, 2(1-4):157–168, 1968.

Krueger, D., Huang, C.-W., Islam, R., Turner, R., Lacoste, A., and Courville, A. Bayesian hypernetworks. *arXiv preprint arXiv:1710.04759*, 2017.

Lesort, T., Caselles-Dupré, H., Garcia-Ortiz, M., Stoian, A., and Filliat, D. Generative models from the perspective of continual learning. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2019.

Lopez-Paz, D. and Ranzato, M. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, pp. 6467–6476, 2017.

McCloskey, M. and Cohen, N. J. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pp. 109–165. Elsevier, 1989.

Nguyen, C. V., Li, Y., Bui, T. D., and Turner, R. E. Variational continual learning. In *International Conference on Learning Representations*, 2018.

Ollivier, Y. Auto-encoders: reconstruction versus compression. *arXiv preprint arXiv:1403.7752*, 2014.

Opper, M. A bayesian approach to on-line learning. *On-line learning in neural networks*, pp. 363–378, 1998.

Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., and Wermter, S. Continual lifelong learning with neural networks: A review. *Neural Networks*, 2019.

Rao, D., Visin, F., Rusu, A., Pascanu, R., Teh, Y. W., and Hadsell, R. Continual unsupervised representation learning. In *Advances in Neural Information Processing Systems*, pp. 7645–7655, 2019.

Ring, M. B. Child: A first step towards continual learning. In *Learning to learn*, pp. 261–292. Springer, 1998.

Ritter, H., Botev, A., and Barber, D. Online structured laplace approximations for overcoming catastrophic forgetting. In *Advances in Neural Information Processing Systems*, pp. 3738–3748, 2018.

Shin, H., Lee, J. K., Kim, J., and Kim, J. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, pp. 2990–2999, 2017.

Solomonoff, R. J. A formal theory of inductive inference. part i. *Information and control*, 7(1):1–22, 1964.

Swaroop, S., Nguyen, C. V., Bui, T. D., and Turner, R. E. Improving and understanding variational continual learning. *arXiv preprint arXiv:1905.02099*, 2019.

Tishby, N. and Zaslavsky, N. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pp. 1–5. IEEE, 2015.

Tolstikhin, I., Bousquet, O., Gelly, S., and Schoelkopf, B. Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558*, 2017.

Wu, C., Herranz, L., Liu, X., wang, y., van de Weijer, J., and Raducanu, B. Memory replay gans: Learning to generate new categories without forgetting. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 5962–5972. Curran Associates, Inc., 2018.

A. Description length on each digit through out training

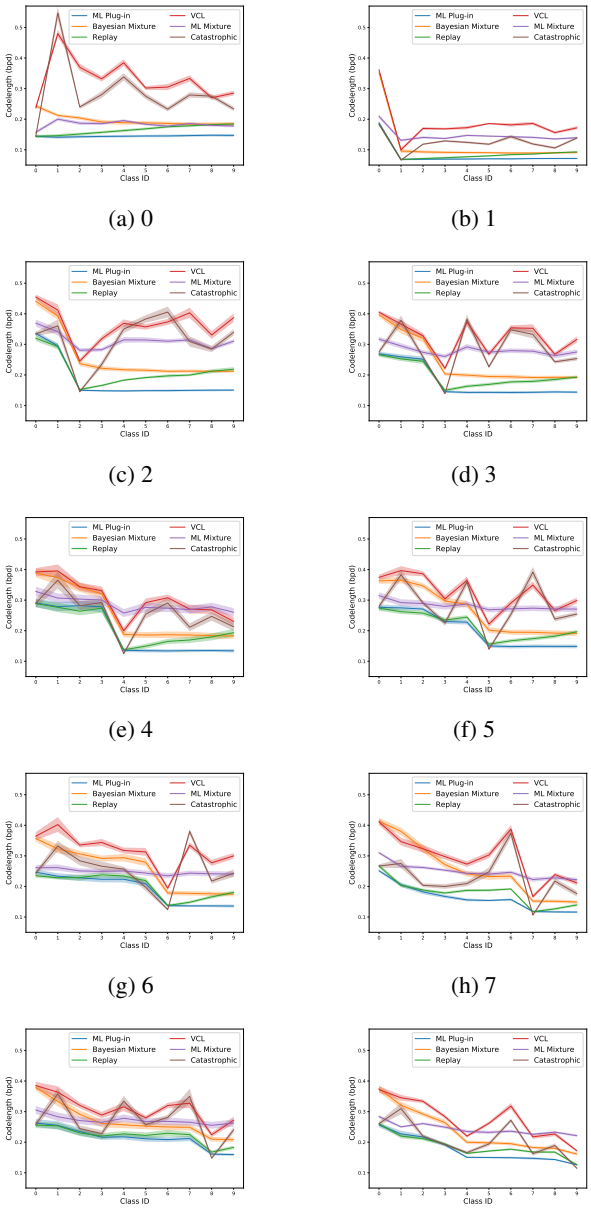

(a) 0 (b) 1 (c) 2 (d) 3 (e) 4 (f) 5 (g) 6 (h) 7 (i) 8 (j) 9

Figure 4: The change of description length on different class through out training.