

Class-Incremental Mixture of Gaussians for Deep Continual Learning

Lukasz Korycki

Virginia Commonwealth University

koryckil@vcu.edu

Bartosz Krawczyk

Rochester Institute of Technology

bartosz.krawczyk@rit.edu

Abstract

Continual learning models for stationary data focus on learning and retaining concepts coming to them in a sequential manner. In the most generic class-incremental environment, we have to be ready to deal with classes coming one by one, without any higher-level grouping. This requirement invalidates many previously proposed methods and forces researchers to look for more flexible alternative approaches. In this work, we follow the idea of centroid-driven methods and propose end-to-end incorporation of the mixture of Gaussians model into the continual learning framework. By employing the gradient-based approach and designing losses capable of learning discriminative features while avoiding degenerate solutions, we successfully combine the mixture model with a deep feature extractor allowing for joint optimization and adjustments in the latent space. Additionally, we show that our model can effectively learn in memory-free scenarios with fixed extractors. In the conducted experiments, we empirically demonstrate the effectiveness of the proposed solutions and exhibit the competitiveness of our model when compared with state-of-the-art continual learning baselines evaluated in the context of image classification problems.

1. Introduction

While the initial research done in the domain of continual learning from stationary data was, in large part, oriented towards task-incremental solutions, more recent works attempt to address generalized cases consisting of purely class-incremental and data-incremental (also known as domain-incremental) settings [31, 39]. These scenarios are usually more universal but also more challenging and restrictive mainly due to the lack of task or even class labels. Such settings make many of the previously proposed solutions practically useless, for example, the methods based on memory-free regularization [24], which are not capable of discriminating between older and new classes, even if they address the catastrophic forgetting problem [30, 39]. Although the most standard experience replay methods can

be effectively applied in the class-incremental scenarios [5, 42], there has been also a search for alternative approaches that could provide natural capabilities required for such cases. A significant group of methods can be identified based on their reliance on centroids (or prototypes) combined with the nearest-centroid classification methods [27]. Since centroids can be independently added to the classifier, they are examples of methods that can be very smoothly incorporated into class-incremental scenarios, offering almost no interference in the latent space.

In this work, we explore an advanced version of these alternatives by proposing integration of the gradient-based Gaussian mixture model with a class-incremental deep continual learning framework, called **MIX**. In fact, it requires us to tackle three major problems at the same time: (i) gradient-based mixture training, (ii) combining it with a trainable deep feature extractor and, finally, (iii) making it suitable for class-incremental scenarios. To achieve these goals, we introduce a set of dedicated losses, configurations and methods, providing a probabilistic classifier on top of a feature extractor and within a model capable of learning end-to-end. This opens many potential research directions that could exploit the well-modeled statistical properties of Gaussians. In addition to that, we show that our class-incremental mixture model, analogously to the centroid-driven algorithms, is characterized by some inherent properties useful in continual learning scenarios. They allow it for much better separation of concepts at the level of the classification module, leading to significant improvements in memory-free scenarios when pre-trained extractors are used. Through an extensive empirical study, we analyze different configurations of our method, provide the reader with some intuition about its parameters and show its competitiveness in the context of other continual learning algorithms.

2. Related works

Continual learning: In continual learning, our focus should be on effective incorporation of the arriving data and retention of the acquired knowledge [20]. The main problem that learning algorithms will encounter here is

catastrophic forgetting [10, 33]. The most straightforward approaches involve replaying instances of previously seen tasks or classes while learning new ones [4, 12, 37, 42]. Instead of putting instance-level constraints on the learning directions, we can apply direct adjustments to the loss using dedicated regularization terms. The most commonly used approach involves utilizing the knowledge-distillation loss [14] combined with standard cross-entropy [7, 16, 25] or maintaining importance weights to distinguish parameters that are crucial for the retention [2, 21, 48]. These methods generally cannot be used in more realistic class-incremental or data-incremental scenarios (if they do not use memory buffers), since they cannot learn how to discriminate new instances from the older ones [17, 30, 39]. Other approaches may employ masking to isolate parameters per task to keep them static when learning new ones [18, 28, 29], use dynamic structures to expand the network for new concepts [41, 45, 46], utilize ensemble techniques [1, 36, 44] or meta-learning and hypernetworks [9, 19, 43]. Finally, interesting alternative approaches focus on hybridizing the neural networks with different machine learning methods, e.g. decision trees [23] or centroid-driven algorithms [27, 35, 47]. The latter group of methods has been found especially useful in one-class-incremental scenarios, since, as mentioned in the introduction, centroids can be stored independently per class, allowing for natural class-incremental learning without additional interference at the level of a classifier. In this work, we follow these approaches and replace basic centroids learned separately from the feature extractor with more complex end-to-end mixture models.

Mixture optimization: Various techniques can be applied for the task of fitting the mixture model to given data. The most standard approach utilizes the EM algorithm, which can be realized in both offline and online settings [8, 32]. While EM provides a stable framework for learning the mixtures – in terms of mathematical constraints and convergence – it is critically limited when it comes to working with high-dimensional data and feasible memory consumption [11]. On top of that, this algorithm is intrinsically incapable of being fully integrated with neural networks, preventing it from achieving joint end-to-end deep learning and benefiting from dedicated features. An alternative approach involves gradient-based optimization [11]. This method has been proved to be able to provide more scalable and flexible algorithms capable of working in challenging scenarios with high-dimensional data and in online settings. Most importantly, the gradient-based approach naturally enables combining the model as a classifier with a trainable deep feature extractor [40], allowing for extending the optimization process with the input space adjustments. Methods utilizing such a compound learning process showed much evidence of its usability in offline and

unsupervised scenarios, while at the same time encouraging researchers to develop further extensions and improvements [11, 34]. Given all of the characteristics, we decided to use this approach in our scenario of continual learning.

3. Mixture of Gaussians for Class-Incremental Learning

Formally, the general goal of our work is to incrementally learn a classification model defined as $\phi^{(t)} : \mathcal{X} \rightarrow \mathcal{C}$ that can effectively incorporate subsequent class batches $\langle (\mathbf{X}^{(1)}, c = 1), (\mathbf{X}^{(2)}, c = 2), \dots, (\mathbf{X}^{(t)}, c = t) \rangle$, where $\mathbf{X}^{(t)}$ contains instances x only for a given class c . After t classes the model $\phi^{(t)}$ should aim at minimizing the loss for the current class $c = t$ and all previously observed ones:

$$\mathcal{L}^{(t)} = \sum_{c=1}^t \sum_{n=1}^{N_c} \mathcal{L}^{(c)}(\phi^{(t)}(x_n^{(c)})), \quad (1)$$

where $x_n^{(c)} \in \mathbf{X}^{(c)}$ and $\mathcal{L}^{(c)}$ can be any supervised loss.

Additionally, since we are interested in deep learning, we define the whole trainable model as a tuple $\phi^{(t)} = \langle \mathcal{F}^{(t)}, \mathcal{G}^{(t)} \rangle$ consisting of a feature extractor $\mathcal{F}^{(t)}$ and a classifier $\mathcal{G}^{(t)}$ jointly aggregating knowledge from t classes. The model makes prediction by classifying the features provided from the extractor $\phi^{(t)}(x) = \mathcal{G}^{(t)}(\mathcal{F}^{(t)}(x)) = \mathcal{G}^{(t)}(\hat{x})$. In this work, we aim at employing the mixture of Gaussians as a jointly trained incremental classifier. Although the model learns from dedicated features \hat{x} , in the next section, we use x for the sake of simplicity of notation.

3.1. Generic supervised mixture model

Formally, in a standard unsupervised setting the density for a given point \mathbf{x} can be expressed using a multivariate normal distribution defined as:

$$\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) = \frac{1}{\sqrt{(2\pi)^D |\boldsymbol{\Sigma}_k|}} \times \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_k)^T \boldsymbol{\Sigma}_k^{-1} (\mathbf{x} - \boldsymbol{\mu}_k)\right), \quad (2)$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are its mean and covariance, and D is the size of the input (number of dimensions). The Gaussian mixture models (GMM) have been designed to approximate more complex multivariate densities by decomposing them into K components:

$$p(\mathbf{x}) = \sum_{k=1}^K \omega_k \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (3)$$

where each of them is defined using a single Gaussian defined above and ω_k are their weights. The combined model, equipped with more degrees of freedom, should be capable

of providing more accurate expressions of the overall observed distributions than a simpler approach utilizing only a single component. In such a framework, the fitting of the mixture to given data \mathbf{X} is based on minimizing the loss defined using the log-likelihood function:

$$\begin{aligned}\bar{\mathcal{L}} &= -\log p(\mathbf{X}|\omega, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \\ &= -\frac{1}{N} \sum_{n=1}^N \log \sum_{k=1}^K \omega_k \mathcal{N}(\mathbf{x}_n | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k),\end{aligned}\quad (4)$$

where we adjust the free parameters of the model – means $\boldsymbol{\mu}$, covariance matrices $\boldsymbol{\Sigma}$ and weights ω . To adapt the given framework to supervised scenarios we can simply specify a separate mixture model for each class c :

$$p(\mathbf{x}|c) = \sum_{k=1}^K \omega_k^{(c)} \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k^{(c)}, \boldsymbol{\Sigma}_k^{(c)}),\quad (5)$$

and focus on minimizing the aforementioned loss also per class $\bar{\mathcal{L}}^{(c)}$:

$$\hat{\mathcal{L}} = \sum_{c=1}^C \bar{\mathcal{L}}^{(c)} = -\log \sum_{c=1}^C p(\mathbf{X}^{(c)} | \omega^{(c)}, \boldsymbol{\mu}^{(c)}, \boldsymbol{\Sigma}^{(c)}),\quad (6)$$

where $\mathbf{X}^{(c)}$ are N_c class-specific observations.

In continual learning we should aim at minimizing the interference of current updates with previously created models to alleviate the detrimental effect of catastrophic forgetting. Therefore, it is worth mentioning here that GMMs create such an opportunity by allowing for maximizing the log-likelihood only for a currently learned class through $\bar{\mathcal{L}}^{(c)}$. It provides a perfect separation at the level of the classification model.

3.2. Mixture optimization for class-incremental deep learning

In order to apply gradient-based learning to GMM in class-incremental deep learning scenarios, we have to address several different issues. Some of them are common for all GMM models using gradient-based learning, while others are specific for the class-incremental deep learning settings.

In general, we say that our goal is to optimize the class-incremental joint model $\phi^{(l)} = \langle \mathcal{F}^{(l)}, \mathcal{G}^{(l)} \rangle$, defined at the beginning of Sec. 3, using some supervised loss \mathcal{L} . Since we set $\mathcal{G}^{(l)} = \mathcal{N}^{(l)}$, where $\mathcal{N}^{(l)}$ is a whole GMM model, we have $\phi^{(l)}(\mathbf{x}) = \mathcal{N}^{(l)}(\mathcal{F}^{(l)}(\mathbf{x}))$. The trainable parameters are weights $\partial\mathcal{L}/\partial\mathbf{W}$ and biases $\partial\mathcal{L}/\partial\mathbf{b}$ for the extractor, and means $\partial\mathcal{L}/\partial\boldsymbol{\mu}$, covariance matrices $\partial\mathcal{L}/\partial\boldsymbol{\Sigma}$ and component weights $\partial\mathcal{L}/\partial\omega$ for the classifier. All of the subsequent paragraphs focus on designing optimization in the classifier (mixture) space, as it was introduced in Sec. 3.1.

3.2.1 Loss design

Max-component: It has been shown that optimizing the full loss $\bar{\mathcal{L}}^{(c)}$ given in Eq. 4 may lead to some numerical instabilities, especially for high-dimensional data [11]. To address this issue a *max-component approximation* can be used. This approach is very straightforward. Since all $p(x|c, k)$ in Eq. 5 are positive, any component provides a lower bound for the whole sum used in $\bar{\mathcal{L}}^{(c)}$. If for every point x_n we find a component providing the highest log-likelihood and sum all of them, we will get the largest (max-component) lower bound [11]:

$$\mathcal{L}_{max}^{(c)} = -\frac{1}{N_c} \sum_{n=1}^{N_c} \max_k \log(\omega_k^{(c)} \mathcal{N}(\mathbf{x}_n^{(c)} | \boldsymbol{\mu}_k^{(c)}, \boldsymbol{\Sigma}_k^{(c)})).\quad (7)$$

Since we can state that $\mathcal{L}_{max}^{(c)} \geq \bar{\mathcal{L}}^{(c)}$, we are able to minimize $\bar{\mathcal{L}}^{(c)}$ by focusing only on $\mathcal{L}_{max}^{(c)}$. It is also worth mentioning that just like the general formula given in Eq. 6 may eliminate the interference with previously learned classes, the max-component approximation can limit the same issue at the level of class components, for example, in data-incremental scenarios [39], making this approach a natural candidate for continual learning settings.

Inter-contrastive loss: All of the introduced losses are limited to scenarios either without a feature extractor or with a fixed pre-trained one. Unfortunately, if we operate in a setting where we can modify the input space of the mixture model and we utilize any of the aforementioned metrics relying entirely on maximizing log-likelihood, we will inevitably end up with a local minimum that for a joint model $\phi^{(l)}$ exists for $\forall \mathbf{x}(\mathcal{G}^{(l)}(\mathbf{x}) = \mathbf{0})$. This issue can be solved by incorporating an inter-contrastive loss that will distance representations for different classes. We define the loss as:

$$\mathcal{L}_{ie}^{(c)} = \frac{1}{N_c} \max_{j \neq c} \sum_{n=1}^{N_c} \max_k \log(\omega_k^{(j)} \mathcal{N}(\mathbf{x}_n^{(c)} | \boldsymbol{\mu}_k^{(j)}, \boldsymbol{\Sigma}_k^{(j)})),\quad (8)$$

which boils down to finding the closest component in other classes, and then optimizing against the class that on average is the closest to the one currently being considered. We keep the log-likelihood to ensure a similar numerical space of loss values as the one for the positive part given in Eq. 7. However, now one should notice that minimizing such a loss may very easily destabilize learning since optimization will gravitate towards $\bar{\mathcal{L}}_{ie}^{(c)} \rightarrow -\infty$ preventing the model from actually fitting to the class examples. To avoid it we introduce a *tightness bound* τ that clips the contrastive loss value at some pre-defined point $\mathcal{L}_{ie}^{(c)}(\tau) = \max(\tau, \mathcal{L}_{ie}^{(c)})$. This basically means that we stop the decrease of the contrastive loss below the given bound, allowing for a more significant contribution of the actual fitting part $\mathcal{L}_{max}^{(c)}$. We

parametrize the τ value with a simple linear transformation $\tau = \bar{p}_{max}^{(c)} - 1/\tau_p$, where $\bar{p}_{max}^{(c)}$ is the average maximum density value observed across all class components (can be obtained on-the-fly) and τ_p is a tunable hyperparameter that takes values between $(0, 1)$. Such a loss can provide effective discrimination between components of different classes, as shown for an example in Appendix A.

Diverse components: While all of the introduced techniques and modifications ensure reliable discrimination between components of different classes, they do not consider differentiation between components of the same class or their quality. In fact, even in offline gradient-driven settings without dynamic feature extraction it is common to obtain mixtures reduced to a single component per class with all the others practically meaningless, e.g., due to zeroed weights [11]. In scenarios with a trainable extractor, this problem becomes even more significant as it is very easy for the optimizer to focus on maximizing log-likelihood from a single component, as both mixture model and flexible extractor lack constraints to prevent this. While in standard scenarios this problem can be successfully addressed with a good initialization method, e.g., using k-means [38], we observed that it was not enough in our case. As a consequence, we introduced two elements to the learning process.

Regionalization – before learning each class, we first divide it into K clusters using the k-means clustering. Then we force each component to fit only to the data from its cluster called a *region* $\mathcal{R}_k^{(c)}$. This replaces the max-component loss $\mathcal{L}_{max}^{(c)}$ defined in Eq. 7 with:

$$\mathcal{L}_{reg}^{(c)} = - \sum_{k=1}^K \frac{1}{N_k} \sum_{\mathbf{x} \in \mathcal{R}_k^{(c)}} \log(\omega_k^{(c)} \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k^{(c)}, \boldsymbol{\Sigma}_k^{(c)})). \quad (9)$$

Intra-contrastive loss – the regionalization approach is necessary yet not sufficient to provide sufficient diversification between same-class components. The reason for it is the same as for discrimination between different classes, as described in the previous paragraph. Analogously to the inter-contrastive loss, we add the intra-contrastive loss with the tightness bound τ :

$$\mathcal{L}_{ia}^{(c)}(\tau) = \sum_{k=1}^K \max(\tau, \max_{m \neq k} \frac{1}{N_k}) \times \sum_{\mathbf{x} \in \mathcal{R}_k} \log(\omega_m^{(c)} \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_m^{(c)}, \boldsymbol{\Sigma}_m^{(c)})), \quad (10)$$

which for each class region pushes away other same-class components that on average are closest to the currently considered one, based on the regionalization conducted in the previous step. Obviously, one can define separate τ for the inter- and intra-contrastive loss.

Such an approach can effectively increase the diversity of the same-class components, as given for an example in Appendix A. However, this approach imposes a hard constraint on how the representation and mixture may look, which limits the flexibility of the whole model. Regardless of these concerns, this method can still effectively improve the overall performance of a multi-component model over a method without the proposed improvement, as we will show in our extensive experiments.

Final component-based losses: To summarize, we distinguish two component-based losses. One uses the max-component approach (**MC**):

$$\mathcal{L}_{mc} = \sum_{c=1}^l \mathcal{L}_{max}^{(c)} + \mathcal{L}_{ie}^{(c)}(\tau_{ie}), \quad (11)$$

while the second loss adds the regionalization technique with the intra-contrastive part (**MCR**):

$$\mathcal{L}_{mcr} = \sum_{c=1}^l \mathcal{L}_{reg}^{(c)} + \beta(\mathcal{L}_{ie}^{(c)}(\tau_{ie}) + \mathcal{L}_{ia}^{(c)}(\tau_{ia})). \quad (12)$$

Cross-entropy loss: Last but not least, we can also attempt to directly optimize the whole standard loss $\hat{\mathcal{L}}$ given in Eq. 4, using a high-level supervised wrapper loss, e.g., based on cross-entropy (**CE**). In such a case, our loss is defined as:

$$\mathcal{L}_{ce} = - \sum_{c=1}^l \sum_{n=1}^{N_c} \mathbf{y}_n^{(c)} \log \hat{y}_n^{(c)}, \quad (13)$$

where \mathbf{y} is a one-hot target vector and $\hat{y}_n^{(c)}$ comes from the softmax function $\hat{y}_n^{(c)} = e^{p_n^{(c)}} / \sum_{c=1}^l e^{p_n^{(c)}}$ and $p_n^{(c)} = p(x_n | c)$ is a density value for a given class produced by the mixture model accordingly to Eq. 5.

3.2.2 Constraints

Other issues that have to be addressed when using gradient-based mixture training are the mathematical constraints that have to be enforced to preserve a valid mixture model. This is required since gradient-based learning does not constrain the possible values for means, covariance matrices and weights, and the last two have to remain in a specific range of values.

Component weights: For the GMM model its component weights ω_k have to sum up to one: $\sum_{k=1}^K \omega_k = 1$. To ensure that the effective weights satisfy this requirement we simply train auxiliary free parameters $\hat{\omega}_k$ and use the softmax-based normalization $\omega_k = e^{\hat{\omega}_k} / \sum_{j=1}^K e^{\hat{\omega}_j}$ to obtain required values [11, 15].

Covariance matrices: For a general case, the covariance matrices of the GMM model should be symmetric positive definite $v^T \Sigma v > 0$ for all nonzero vectors v . This can be enforced using the Cholesky decomposition [13] $\Sigma = AA^T$, where A is a triangular matrix with positive diagonal values $a_{ii} > 0$ and, at the same time, our trainable proxy parameter. To enforce positive diagonal values, after each gradient-based update we clamp them with $a_{ii} = \min(a_{ii}, d_{min})$ using some predefined d_{min} value. Finally, we also consider a case of a mixture using only the diagonal of the covariance – variance σ , which we control using the same clamp-based approach $\sigma_i = \min(\sigma_i, d_{min})$.

3.3. Memory buffer

In our work, we consider the class-incremental scenario with strictly limited access to previously seen observations (classes). Therefore, in all of the introduced losses we use all available data for the currently learned class t , while for the others we sample from the memory buffers \mathcal{M}_c that store an equal number of examples per each previously seen class. On the other hand, if the feature extractor is pre-trained and static we could remove the inter-contrastive loss and even get rid of the memory buffer, allowing for memory-free training, as we will show in our experimental study. The memory buffer is needed in a general case when we assume the joint training of the whole model.

3.4. Classification

Finally, in the presented model, the classification of an instance x_n can be performed using two approaches, either utilizing the softmax function $\hat{y}_n^{(c)} = e^{p_n^{(c)}} / \sum_{c=1}^L e^{p_n^{(c)}}$, where $p_n^{(c)} = p(x_n|c)$, or by taking the weighted support of the closest component $\hat{y}_n^{(c)} = \max_k \omega_k^{(c)} \mathcal{N}(x_n | \mu_k^{(c)}, \Sigma_k^{(c)})$. We will empirically show that these methods work best with specific losses designed in the previous sections.

4. Experimental study

In our experiments, we empirically explore all of the introduced methods and parameters and put our method in the performance context of different state-of-the-art baselines. We show how our model performs in end-to-end scenarios and with a pre-trained extractor, compared with other solutions. For more specific details regarding data, configurations and results, please refer to Appendix A and B, as well as to our repository containing source code for our methods and all experiments: github.com/lkorycki/mix. All of the experiments were conducted using 4 GPUs (Tesla V100) that were part of an internal cluster.

4.1. Setup

For the purpose of the evaluation we selected commonly utilized image classification datasets that were turned into

class-incremental sequences by presenting their classes subsequently to the models [30, 31]. We used: MNIST, FASHION, SVHN, CIFAR and IMAGENET datasets using various variants (number of classes, pre-trained features). For the analysis of different configurations of our model we used shorter sequences. We extended them with the longer benchmarks for the comparison with baselines.

In the final section of this work, we compared our class-incremental Gaussian mixture model (MIX-MCR, MIX-CE) with other classifiers dedicated for continual learning scenarios. We considered: standard experience replay (ER) [22], experience replay with subspaces (ERSB) [22], centroid-based iCaRL [35], two gradient-based sample selection methods (GSS and A-GEM) [3, 6], experience replay combined with knowledge distillation and regularization (DER) [5], and two purely regularization-based approaches – LWF [25] and SI [48]. Most of the algorithms were implemented as wrappers of the source code provided in [5, 26] under MIT License. For the last two we used their modifications adjusted for single-task learning [30]. As our lower bound we used a naively learning net (NAIVE), and for the upper bound we present results for the offline model (OFFLINE).

We evaluated the presented methods in a class-incremental setting, where all of the classes were presented to the models subsequently and were not shown again after their initial appearance. We measured the accuracy of a given algorithm after each class batch, utilizing holdout testing sets, and then, based on [20], used it to calculate the average incremental accuracy over the whole sequence:

$$\Omega_{all} = \frac{1}{T} \sum_{t=1}^T \alpha_t, \quad (14)$$

where α_t is the model performance after t classes and $T = C$ is the total number of classes. In addition to the whole aggregation, for the final comparison, we provided these values after each batch to present a more complete perspective of the obtained results.

4.2. Results

In this section, we present and describe all of the results that were obtained for the experiments introduced in the previous paragraphs. The first part consists of the analysis of different configurations of MIX, while the second one focuses on a comparison with other class-incremental algorithms.

Loss and classification: We analyzed different combinations of the proposed losses and classification methods. Based on Fig. 1, we can make three major observations. Firstly, the softmax classification works significantly better with the CE loss, and max-component can be more efficiently paired with MC and MCR than softmax. It was




Figure 1. Average incremental accuracy for different losses combined with different classification methods.




Figure 2. Average incremental accuracy for different combinations of the tightness parameter values (inter and intra).

evident for almost all cases (except for MC on CIFAR10) and resulted in almost 0.15 difference on average between softmax and max-component for CE, and about 0.05 for MC and MCR.

Secondly, the MCR loss performed better than MC, showing consistent improvements, especially for more complex datasets like SVHN, CIFAR10 or IMAGENET10, which resulted in more than 0.1 for a difference on average. This demonstrate that the regionalization and intra-contrastive loss are capable of providing meaningful improvements over simpler MC loss utilizing only max-component and inter-contrastive elements, and that ensuring higher diversity among class components can be beneficial to the model.

Finally, we can see that CE with softmax could provide very similar results as MCR with max-component, which means that the general GMM learning formula, wrapped with a high-level supervised loss, can be sometimes as useful as more complex MCR without the need for tuning additional parameters. One drawback of using CE, however, is the fact that it does not model the Gaussian mixtures well (see Appendix B for additional visualizations). The CE loss does not really have to fit the mixtures to the data since it is enough for it to ensure high classification quality. We can also observe a similar behavior for the MC loss. It may be prohibitive if one wants to obtain a reliable description of the latent space. The MCR loss achieves both objectives at the same time: high classification accuracy and high quality of the mixture models for features. This may be important if someone requires interpretable models or would like to extend the proposed algorithm with some Gaussian-oriented techniques that MCR may enable. Furthermore, we believe that analyzing its probabilistic properties in detail could be a

part of incremental works built on top of the mixture model. They could utilize its well-defined characteristics, e.g. by proposing new mixture-based losses.

Tightness: In Fig. 2, we presented a grid of values for the average incremental accuracy per each pair of inter- and intra-tightness for every dataset. One can clearly see that imposing the constraint (tightness) on the inter- and intra-contrastive loss values is beneficial to the learning process. Most of the benchmarks required $\tau_{p,ie}$ at the level of 0.0001 or 0.001 and slightly higher intra-tightness $\tau_{p,ia}$ around 0.001 or 0.01 to achieve the best results. At the same time, one should notice that imposing too high inter-tightness (0.01) leads to abrupt deterioration of quality, which is a result of blocking the contrastive part of the loss from pushing components of different classes from each other. The influence of setting too high intra-tightness is less important since we may simply end up with a single component that can still be effectively used for classification.

The examples for FASHION, given in Fig. 3 and 4, show how increasing the inter-tightness (the first one) and intra-tightness (the second one) affects learned representations and mixture models. We can observe the positive impact of the constraint and the potential for sweet spots providing a good balance between differentiating components between each other and fitting them to the actual data. It is evident that too low values introduce critical instabilities to the learning process (very high contrastive loss values overwhelming the fitting part), while too high thresholds lead either to the decline of discriminative properties of the model or degenerate solutions.




Figure 3. Visualization of the inter-tightness effect on learned representations ($K=1$) after 10 classes of FASHION.




Figure 4. Visualization of the intra-tightness effect on learned representations ($K=3$) after 6 classes of FASHION.

Baseline comparison: In the second section of our experimental study, we placed our algorithm in the class-incremental performance context by comparing it with the introduced baselines (Fig. 5). First of all, we can see that the MIX-MCR variant performed better than the MIX-CE for most of the datasets, while being very close to it for the longer sequences (difference between less than 0.01 and 0.03). This proves that MIX-MCR is capable of providing not only a better representation (mixture) model but also that it is more reliable from the accuracy perspective. This also means that it is worth trying to maximize the quality of the produced Gaussian models as an alternative to high-level cross-entropy for classification. Secondly, although our model cannot be distinguished as the best classifier (being worse than iCaRL on average, with a difference equal to about 0.04), it is, at the same time, reliably competitive when compared with the remaining baselines (ER, GSS, DER) with a difference about 0.01 and less than 0.03. Also, it does not fall into the same pitfalls as either the weakest replay method (A-GEM) or the regularization-based ones (LWF, SI), outperforming them by almost 0.4 for accuracy on average. We can see that MIX could be found among the best models for MNIST, FASHION, IMAGENET10, IMAGENET20A and IMAGENET20B, especially at the end of the datasets, providing relatively reliable performance throughout the whole sequences. On the other hand, it struggled with catching up with the best replay methods for SVHN and CIFAR-based datasets showing that there is still a potential for improvements when it comes to predictive accuracy.

The overall very poor performance of LWF and SI (but also A-GEM), which were not much better than the NAIVE approach, confirms the observations made in other publica-

tions that the regularization-based methods cannot handle the most challenging 1-class-incremental scenarios without memory buffers [39] even after improvements proposed in [30].

We can also see that for the scenarios with end-to-end training the models were much closer (0.01-0.3) to the OFFLINE upper bound for the shorter sequences (MNIST, FASHION, SVHN and IMAGENET10, except for CIFAR10) than for the longer ones (IMAGENET20A, IMAGENET20B, CIFAR20) with differences between 0.4-0.5, which shows that all of the state-of-the-art methods still struggle with bridging the gap between incremental learning and offline optimum.

Finally, the results for the memory-free scenarios with pre-trained models, given in the last row of Fig. 5, exhibit the main strength of the MIX algorithm. Since in these scenarios, it does not use the inter-contrastive loss, it can perfectly separate the incremental learning process for each class, preventing catastrophic forgetting at the level of the classifier. As a result, it does not have to rehearse the previous concepts at all ($\mathcal{M}_c=0$) while still being able to conduct very effective learning producing results very close to the OFFLINE upper bound (difference between about 0 and 0.1), regardless of the quality of the extractor (pre-trained on 10 and 20 or 100 and 200 classes). The MIX-MCR method outperforms all of the baselines for all cases except for IMAGENET200-PRE20, for which only iCaRL was able to provide slightly higher accuracy, even though they had a small advantage of having approximately one example per class in the buffer. It is not a coincidence that practically only iCaRL is close to our method on average (worse by about 0.1), since it uses a similar paradigm in the classification layer by storing prototypes/centroids that are

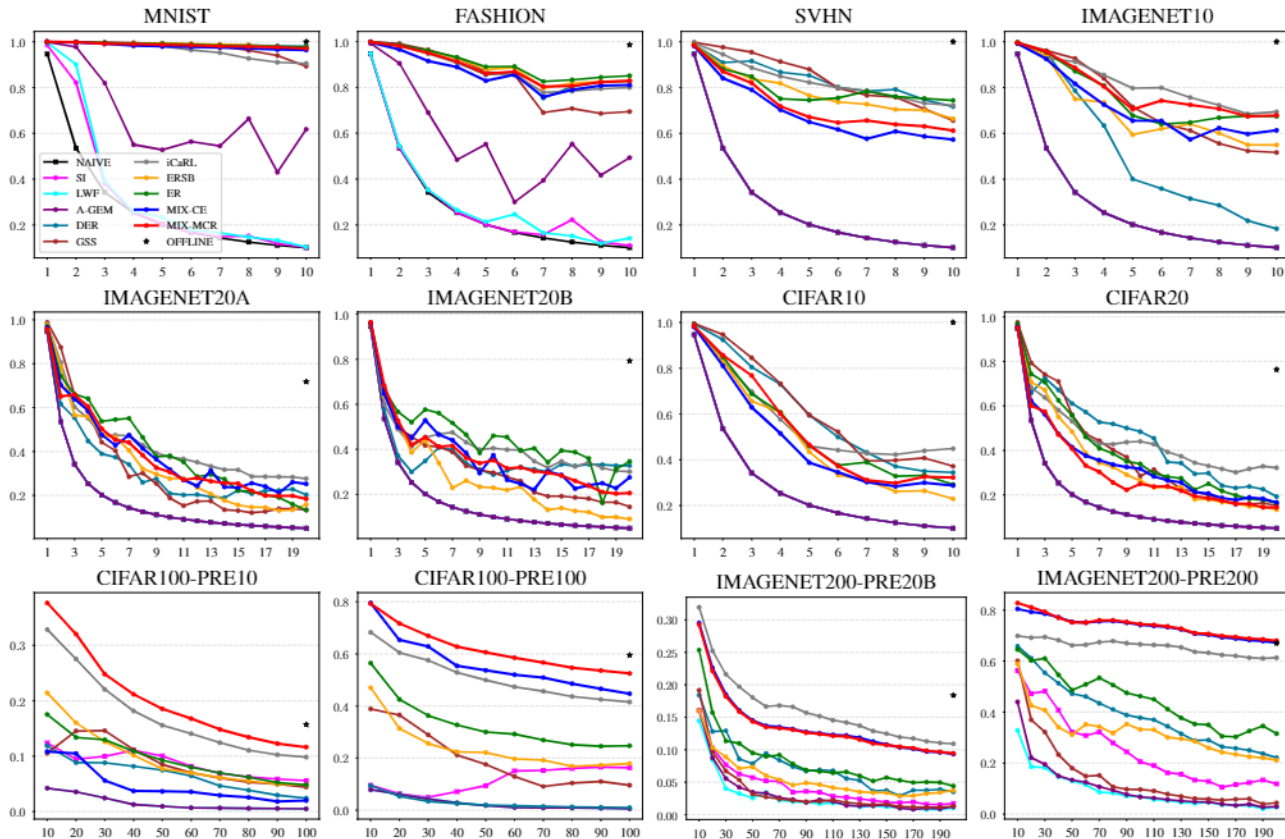


Figure 5. Incremental accuracy after each class batch for our methods and different baselines.

used for classification. All of the remaining algorithms cannot handle the memory-free scenario effectively, producing solutions worse by at least 0.2 on average. This can be a crucial property when one has to consider, for example, data privacy issues or mobile and edge computing.

All of the presented observations, conclusions and recommendations can be also found in a condensed form at the end of Appendix B.

5. Summary

In this work, we introduced a class-incremental mixture of Gaussians model (MIX) for deep continual learning. We proposed different variants of the algorithm to make it suitable for gradient-based optimization and, through an extensive experimental study, we exhibited its practical configurations and capabilities in the context of other state-of-the-art continual learning models.

In our future research, we will focus on replacing the regionalization approach with a more flexible method that do not assume any pre-training structure and allows the gradient-based procedure to fully explore potential solutions, e.g. annealing [11], and on removing the static tightness hyperparameter to increase flexibility even more – it

could be more beneficial to either find a better (parameter-free) distance function or propose an adaptive threshold. It is also an open question whether we can effectively train a gradient-based mixture using a full covariance matrix. Finally, we could consider some kind of hybridization of the mixture models with the feature extractor to benefit from the capabilities of the former to limit interference with previously learned concepts by utilizing max-component losses. All of these potential improvements combined could provide significant performance gains in the class-incremental continual learning scenarios.

References

- [1] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert Gate: Lifelong Learning with a Network of Experts. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7120–7129, 2017. 2
- [2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory Aware Synapses: Learning what (not) to forget. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [3] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual

- learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 5
- [4] Jihwan Bang, Heesu Kim, Young Joon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow Memory: Continual Learning with a Memory of Diverse Samples. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8214–8223, 2021. 2
- [5] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark Experience for General Continual Learning: a Strong, Simple Baseline. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 1, 5
- [6] Arslan Chaudhry, Marc Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient Lifelong Learning with A-GEM. In *International Conference on Learning Representations (ICLR)*, 2019. 5
- [7] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyang Wu, and Rama Chellappa. Learning Without Memorizing. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5133–5141, 2019. 2
- [8] Paulo Martins Engel and Milton Roberto Heinen. Incremental Learning of Multivariate Gaussian Mixture Models. In *Advances in Artificial Intelligence – SBIA 2010*, pages 82–91, 2010. 2
- [9] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017. 2
- [10] Robert M. French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135, 1999. 2
- [11] Alexander Rainer Tassilo Gepperth and Benedikt Pfülb. Gradient-Based Training of Gaussian Mixture Models for High-Dimensional Streaming Data. *Neural Process. Lett.*, 53:4331–4348, 2021. 2, 3, 4, 8
- [12] Tyler L. Hayes, Giri P. Krishnan, Maxim Bazhenov, Hava T. Siegelmann, Terrence J. Sejnowski, and Christopher Kanan. Replay in Deep Learning: Current Approaches and Missing Biological Elements. *Neural Computation*, 33(11):2908–2950, 2021. 2
- [13] Nicholas J. Higham. Cholesky factorization. *Wiley Interdisciplinary Reviews: Computational Statistics*, 1(2):251–254, 2009. 5
- [14] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the Knowledge in a Neural Network. *arXiv*, abs/1503.02531, 2015. 2
- [15] Reshad Hosseini and Suvrit Sra. Matrix Manifold Optimization for Gaussian Mixtures. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015. 4
- [16] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Lifelong Learning via Progressive Distillation and Retrospection. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [17] Yen-Chang Hsu, Yen-Cheng Liu, and Zsolt Kira. Re-evaluating Continual Learning Scenarios: A Categorization and Case for Strong Baselines. *arXiv*, abs/1810.12488, 2018. 2
- [18] Steven C. Y. Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, Picking and Growing for Unforgetting Continual Learning. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019. 2
- [19] Khurram Javed and Martha White. Meta-Learning Representations for Continual Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2
- [20] Ronald Kemker, Angelina Abitino, Marc McClure, and Christopher Kanan. Measuring Catastrophic Forgetting in Neural Networks. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2018. 1, 5
- [21] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017. 2
- [22] Łukasz Korycki and Bartosz Krawczyk. Class-Incremental Experience Replay for Continual Learning under Concept Drift. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 3644–3653, 2021. 5
- [23] Łukasz Korycki and Bartosz Krawczyk. Streaming Decision Trees for Lifelong Learning. In *Machine Learning and Knowledge Discovery in Databases. Research Track - European Conference (ECML/PKDD)*, pages 502–518, 2021. 2
- [24] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory G. Slabaugh, and Tinne Tuytelaars. Continual learning: A comparative study on how to defy forgetting in classification tasks. *CoRR*, abs/1909.08383, 2019. 1
- [25] Zhizhong Li and Derek Hoiem. Learning without Forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2018. 2, 5
- [26] Vincenzo Lomonaco, Lorenzo Pellegrini, Andrea Cossu, Antonio Carta, Gabriele Graffieti, Tyler L. Hayes, Matthias De Lange, Marc Masana, Jary Pomponi, Gido van de Ven, Martin Mundt, Qi She, Keiland Cooper, Jeremy Forest, Eden Belouadah, Simone Calderara, German I. Parisi, Fabio Cuzzolin, Andreas Tolia, Simone Scardapane, Luca Antiga, Subutai Amhad, Adrian Popescu, Christopher Kanan, Joost van de Weijer, Tinne Tuytelaars, Davide Bacciu, and Davide Maltoni. Avalanche: an end-to-end library for continual learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2021. 5
- [27] Zheda Mai, Ruiwen Li, Hyunwoo Kim, and Scott Sanner. Supervised Contrastive Replay: Revisiting the Nearest Class Mean Classifier in Online Class-Incremental Continual Learning. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 3584–3594, 2021. 1, 2
- [28] Arun Mallya and Svetlana Lazebnik. PackNet: Adding Multiple Tasks to a Single Network by Iterative Pruning. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7765–7773, 2018. 2

- [29] Arun Mallya and Svetlana Lazebnik. Piggyback: Adding Multiple Tasks to a Single, Fixed Network by Learning to Mask. *arXiv*, abs/1801.06519, 2018. [2](#)
- [30] Davide Maltoni and Vincenzo Lomonaco. Continuous learning in single-incremental-task scenarios. *Neural Networks*, 116:56–73, 2019. [1](#), [2](#), [5](#), [7](#)
- [31] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D. Bagdanov, and Joost van de Weijer. Class-incremental learning: survey and performance evaluation. *CoRR*, abs/2010.15277, 2020. [1](#), [5](#)
- [32] Volodymyr Melnykov and Ranjan Maitra. Finite mixture models and model-based clustering. *Statistics Surveys*, 4: 80 – 116, 2010. [2](#)
- [33] German I. Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019. [2](#)
- [34] Benedikt Pfülb and Alexander Rainer Tassilo Geppert. Overcoming Catastrophic Forgetting with Gaussian Mixture Replay. *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–9, 2021. [2](#)
- [35] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, G. Sperl, and Christoph H. Lampert. iCaRL: Incremental Classifier and Representation Learning. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5533–5542, 2017. [2](#), [5](#)
- [36] Amanda Rios and Laurent Itti. Lifelong Learning Without a Task Oracle. *2020 IEEE 32nd International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 255–263, 2020. [2](#)
- [37] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy P. Lillicrap, and Gregory Wayne. Experience Replay for Continual Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. [2](#)
- [38] Ting Su and Jennifer G. Dy. In search of deterministic methods for initializing K-means and Gaussian mixture clustering. *Intell. Data Anal.*, 11:319–338, 2007. [4](#)
- [39] Gido M. van de Ven and Andreas Savas Tolias. Three scenarios for continual learning. *arXiv*, abs/1904.07734, 2019. [1](#), [2](#), [3](#), [7](#)
- [40] Ehsan Variani, Erik McDermott, and Georg Heigold. A Gaussian Mixture Model layer jointly optimized with discriminative features within a Deep Neural Network architecture. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4270–4274, 2015. [2](#)
- [41] Tom Veniat, Ludovic Denoyer, and MarcAurelio Ranzato. Efficient continual learning with modular networks and task-driven priors. In *International Conference on Learning Representations (ICLR)*, 2021. [2](#)
- [42] Eli Verwimp, Matthias De Lange, and Tinne Tuytelaars. Rehearsal revealed: The limits and merits of revisiting samples in continual learning. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [1](#), [2](#)
- [43] Johannes von Oswald, Christian Henning, João Sacramento, and Benjamin F. Grewe. Continual learning with hypernetworks. In *International Conference on Learning Representations, ICLR*, 2020. [2](#)
- [44] Yeming Wen, Dustin Tran, and Jimmy Ba. BatchEnsemble: an Alternative Approach to Efficient Ensemble and Lifelong Learning. In *International Conference on Learning Representations (ICLR)*, 2020. [2](#)
- [45] Shipeng Yan, Jiangwei Xie, and Xuming He. DER: Dynamically Expandable Representation for Class Incremental Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [2](#)
- [46] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong Learning with Dynamically Expandable Networks. In *International Conference on Learning Representations*, 2018. [2](#)
- [47] Lu Yu, Bartłomiej Twardowski, Xialei Liu, Luis Herranz, Kai Wang, Yong mei Cheng, Shangling Jui, and Joost van de Weijer. Semantic Drift Compensation for Class-Incremental Learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6980–6989, 2020. [2](#)
- [48] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual Learning through Synaptic Intelligence. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017. [2](#), [5](#)