

# Causes of Catastrophic Forgetting in Class-Incremental Semantic Segmentation

Tobias Kalb<sup>1</sup>[0000–0002–1993–0086] and Jürgen Beyerer<sup>2,3</sup>[0000–0003–3556–7181]

<sup>1</sup> Porsche Engineering Group GmbH, Porschestraße 911, 71287 Weissach, Germany  
[tobias.kalb@porsche-engineering.de](mailto:tobias.kalb@porsche-engineering.de)

<sup>2</sup> Fraunhofer Institute of Optonics, Systems Technologies and Image Exploitation  
IOSB

<sup>3</sup> Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany  
[juergen.beyerer@iosb.fraunhofer.de](mailto:juergen.beyerer@iosb.fraunhofer.de)

**Abstract.** Class-incremental learning for semantic segmentation (CiSS) is presently a highly researched field which aims at updating a semantic segmentation model by sequentially learning new semantic classes. A major challenge in CiSS is overcoming the effects of catastrophic forgetting, which describes the sudden drop of accuracy on previously learned classes after the model is trained on a new set of classes. Despite latest advances in mitigating catastrophic forgetting, the underlying causes of forgetting specifically in CiSS are not well understood. Therefore, in a set of experiments and representational analyses, we demonstrate that the semantic shift of the background class and a bias towards new classes are the major causes of forgetting in CiSS. Furthermore, we show that both causes mostly manifest themselves in deeper classification layers of the network, while the early layers of the model are not affected. Finally, we demonstrate how both causes are effectively mitigated utilizing the information contained in the background, with the help of knowledge distillation and an unbiased cross-entropy loss.

**Keywords:** Continual Learning · Semantic Segmentation · Class-Incremental.

## 1 Introduction

Semantic segmentation is a long-standing problem in computer vision, which aims to assign a semantic label to each pixel of an image. However, a fundamental constraint of the traditional semantic segmentation benchmarks is that they assume that all classes are known beforehand and are all learned at once. This assumption limits the use of semantic segmentation for practical applications, as within a realistic scenario the model should be able to learn new classes without requiring a complete retraining of all previous ones. For this reason, the recently emerged field of class-incremental semantic segmentation (CiSS) focuses on achieving this goal of incrementally learning new classes. The two main challenges that CiSS has to overcome are catastrophic forgetting [11,26] of old classes and the semantic shift of the background class. Recently, several methods were proposed to mitigate the limitations brought on by these challenges

[40,28,19,3,8,29]. Contrary to methods in class-incremental image classification, CiSS methods are mostly based on the concept of knowledge distillation [14]. While significant progress has been made in mitigating the effects of catastrophic forgetting and the semantic shift of the background class, there is limited understanding of how the drop in accuracy manifests itself within the CNN-based semantic segmentation model. The focus of our work is to identify the causes of forgetting in CiSS. Specifically, we aim at revealing how activation drift, inter-task-confusion and task-recency bias affect the performance in CiSS and how existing approaches overcome these effects. Our main contributions are:

1. We study the impact of the semantic shift of the background class on PascalVoc2012 [10] in 3 different task protocols: *overlapped*, *disjoint* and a novel *full disjoint* setup.
2. We analyse the degree of activation drift in various layers by stitching them with the previous task’s network and reveal that semantic shift of the background class is the main cause of the catastrophic drop in performance in CiSS. Forgetting mainly happens in the decoder layers of the model, where discriminating features for old classes of the encoder are assigned to new visually similar classes or to the background class. However, the re-appearance of previous classes in the background of subsequent training tasks also reduces the internal activation drift in the encoder.
3. We introduce Decoder Retraining Accuracy to analyse the degree to which the decoder of the network contributes to inter-task confusion. We observe that methods that do not use any form of replay fail to learn discriminating features for all classes. Specifically, the model is not able to distinguish old classes from new classes that are visually closely related, e.g. *train* and *bus*.
4. In our final analysis, we apply the UNCE loss from [3] to the full disjoint setting and provide observations that support our inferences.

## 2 Related Work

### 2.1 Continual Learning

The majority of research in continual learning is focused on developing methods to overcome the effects of catastrophic forgetting. This is achieved by dedicating a subset of parameters to each task either dynamically or explicitly, by penalizing updates on important parameters for the previous task during training on a new task [1,18,7,43], directly storing data from previous task for later replay [39,35,34,13], as well as knowledge-distillation-based methods, which utilize the activations of the old network as regularization term during training on new data [23]. In a recent empirical survey by Masana *et al.* [25] the best performing methods on class-incremental image classification all utilize replay to cope with forgetting [25]. Other surveys and investigations confirm the notion that class-incremental learning requires some form of replay [41] and that prior regularization methods might fail at learning discriminating features in the class-incremental setting [22,16]. For a state-of-the-art overview on class-incremental image classification we refer to [25].

## 2.2 Class-Incremental Semantic Segmentation

When comparing the most successful approaches in class-incremental image classification [25] and CiSS [40,28,19,3,8,29,17], it is noticeable that contrary to classification in CiSS most of the recent methods build on the idea of Learning without Forgetting (LwF) [23] and utilize a knowledge distillation-based loss. In most CiSS benchmarks knowledge distillation-based approaches even outperform replay-based methods [17]. As a result, state-of-the-art CiSS approaches mostly rely on some form of knowledge distillation [14] and do not require any form of replay. These approaches have been proven to be effective on their respective benchmarks. However, they require that previously learned classes will reappear in the future training images, because otherwise they would still suffer from forgetting, as we will confirm in our experiments. Therefore, recent approaches also investigated how to integrate exemplar replay into CiSS [24,9,17] or utilizing unlabelled auxiliary data [42,2].

## 2.3 Studying effects of Catastrophic Forgetting

Prior work on understanding catastrophic forgetting in deep learning for image classification used representational analysis techniques such as centered kernel alignment [20] and linear probing to conclude that deeper layers are disproportionately the cause of forgetting in CNNs [6,33]. We will confirm this is also true for CiSS. Other work analyzed the loss-landscape of continually trained models [30] or investigated how the task sequence [32] and task similarity [33] impact forgetting. Contrary to prior work, we concentrate on identifying the causes of forgetting that arise specifically in CiSS, as the semantic shift of background class and the multi-class nature of semantic segmentation introduce new challenges into the continual learning setting that are not yet well understood.

## 3 Problem Formulation

Before going into details about the causes of forgetting, we define the general task of CiSS. The goal of semantic segmentation is to assign a class out of a set of pre-defined classes  $\mathcal{C}$  to each pixel in a given image. A training task  $T = \{(x_m, y_m)\}_{m=1}^M$  consists of a set of  $M$  images  $x \in \mathcal{X}$  with  $\mathcal{X} = \mathbb{R}^{H \times W \times 3}$  and corresponding labels  $y \in \mathcal{Y}$  with  $\mathcal{Y} = \mathcal{C}^{H \times W}$ . Given the task  $T$  the goal in semantic segmentation is to learn a mapping  $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^{H \times W \times |\mathcal{C}|}$  from the image space  $\mathcal{X}$  to a posterior probability vector  $q$ . The output segmentation mask for a single pixel  $i$  of the image is obtained as  $\bar{y}_i = \arg \max_{c \in \mathcal{C}} q_{i,c}$ . In the class-incremental learning setting the model  $f_\theta$  is not trained on a single task  $T$  but on a sequence of tasks  $T_t$ . Each task  $T_t$  extends the previous set of classes  $\mathcal{C}_{t-1}$  by a set of novel classes  $\mathcal{S}_t$  resulting in the new label set  $\mathcal{C}_t = \mathcal{C}_{t-1} \cup \mathcal{S}_t$ . In this setting, the labels of classes  $\mathcal{C}_{t-1}$  are not included in the training set of  $T_t$ , with the exception of the background class  $b$ . After learning a task  $T_t$ , the model is required to correctly discriminate between all the observed classes  $\mathcal{C}_t$ .

### 3.1 Causes of Forgetting in Class-Incremental Learning

The fundamental challenge in class-incremental learning is that during optimization of the model on a new  $T_t$ , the model is optimized without regard to the previous classes  $\mathcal{C}_{t-1}$ , which leads to catastrophic forgetting of previous classes. Masana *et al.* [25] stated four causes of forgetting in class-incremental classification.

- **Weight Drift:** During optimization on  $T_t$ , the weights of the model that were relevant to the previous task  $T_{t-1}$  are updated without regard to the previous task, resulting in drop of performance on task  $T_{t-1}$ .
- **Activation Drift:** A change of the weights of the model directly results in a change of internal activations and to the output of the model.
- **Inter-task confusion:** The objective in class-incremental learning is to correctly discriminate between all the observed classes  $\mathcal{C}_t$ . However, as the classes are never jointly trained, the learned features are not optimized to discriminate classes from different tasks, as shown in Fig. 1.
- **Task-recency bias:** In the class-incremental setting, the model is optimized to predict new classes without regarding the old classes. This leads to a strong bias for the most recently learned classes. This bias can easily be seen in the confusion matrix, as shown in Fig. 3a. Furthermore, especially in the final classification layer, the weights and biases of the classifier layer have higher magnitudes of the weights vectors for new classes [15]. In CiSS with the addition of the background class, an additional bias to the background class can be observed.

In the following we want to investigate how these causes manifest themselves in CiSS and how existing approaches mitigate these causes.




Fig. 1: Visualization of task confusion in class-incremental learning. As classes of Task 1 (circles) and classes of Task 2 (squares) are never trained at the same time, the classifier never learns to discriminate between circles and squares, which causes inter-task confusion.

## 4 Methods to measure forgetting

Purely accuracy based evaluation only allows restrictive insight into the causes of forgetting of a model. Therefore, in this section, we present the key methods

that we use in our analysis, including layer matching [5] and decoder retraining accuracy. We utilize these methods to measure the activation shift between a model  $f_0$  and  $f_1$ . The model  $f_0$  is trained on  $T_0$ , whereas  $f_1$  is initialized with the parameters of  $f_0$  and incrementally trained on  $T_1$ .

#### 4.1 Layer Matching with Dr. Frankenstein

The Dr. Frankenstein toolset proposed by Csizs  rik *et al.* [5] aims to match the activations of two neural networks on a given layer by joining them with a stitching layer, compare Fig. 2. The goal of the stitching layer is to transform the activations of a specific layer of  $f_0$  to the corresponding activations of a model  $f_1$ . The stitching layer is initialized using least-squares-matching and is optimized using the loss function the network was trained with. In order to measure the similarity of the learned representations, the accuracy of the resulting Frankenstein Network is evaluated on the test set and compared to the initial accuracy of the model  $f_0$ . The higher the resulting relative accuracy is, the closer the learned representations of the models are to each other. In our analysis we omit the stitching layer and directly use the activations of  $f_0$  in  $f_1$ , as we noticed in our experiments that the initial activations are already very similar. We attribute this to the fact that the models are closely related because  $f_1$  is initialized with the parameters of  $f_0$ . Our setup can be seen on the right side in Sect. 4.1. We denote the resulting Frankenstein Network when layer  $n$  of model  $f_1$  is stitched to layer  $n + 1$  of model  $f_0$  as  $f_{1,0}^n$ , i.e. the network depicted in Sect. 4.1 would be  $f_{1,0}^2$ . If the accuracy of the resulting Frankenstein network is not affected, this is clear evidence that the internal representations of  $f_1$  were not altered drastically during training on  $T_1$ . This analysis will give insight into how much the activation at a specific layer has changed after training continually, but will give no insight into a possible positive backward transfer.

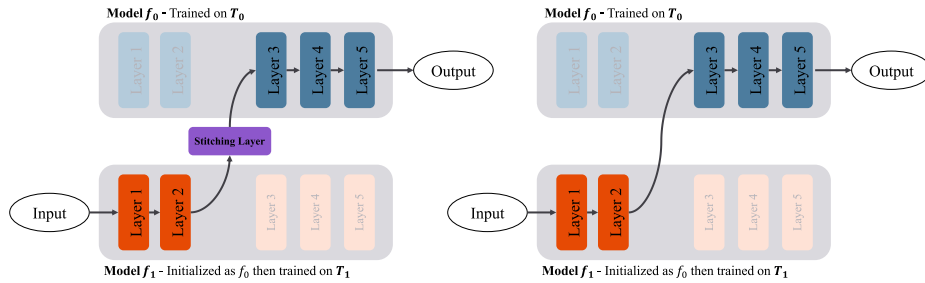


Fig. 2: Comparing a) the original Dr. Frankenstein layer matching approach of using an additional stitching layer as proposed by [5] and b) our approach for measuring activation drift in continual learning by directly propagating the activations of  $f_1$  to  $f_0$  without a stitching layer.

## 4.2 Decoder Retrain Accuracy

Inspired by use of linear probing for continual learning [6], which measures representational forgetting by calculating the difference in accuracy an optimal linear classifier achieves on an old task before and after introducing a new task, we propose Decoder Retraining Accuracy. To measure the Decoder Retraining Accuracy, we freeze the encoder of the model and retrain the decoder on all classes with the same training configuration and measure the mIoU on a test set. Instead of measuring the representational drift as in [6] this will give a measure of how useful the learned representations of the encoder are to discriminate between classes of different tasks, effectively measuring the contribution of inter-task-confusion.

## 5 Experiments

*Datasets:* We conduct our experiments using the PascalVoc-2012 [10] dataset, which contains 20 object classes and a background class. We follow the established CiSS PascalVoc-15-5 split that is widely used [8,28,29,3]. The PascalVOC-15-5 split is a two step incremental learning task, which consists of learning 15 classes (1–15) in the first step  $T_0$  and the remaining 5 classes (16–20) in the second step  $T_1$ .<sup>4</sup> We follow the two distinct class-incremental settings for *Disjoint* and *Overlapped* proposed by Cermelli *et al.* [3]. In both settings, only the set of current classes  $\mathcal{S}_t$  is labelled, while the rest is labelled as background  $b$ . However, in the *Disjoint* setting, the images of the current task  $T_t$  only contain pixels of classes  $\mathcal{C}_t$ , meaning that images that contain pixels belonging to classes of future tasks will be discarded in the training set of  $T_t$ . In the *Overlapped* setting, pixels can belong to any of the classes, but classes that do not belong to current training set will be labeled as background. Finally, in order to study the impact of the semantic shift that the background class is subjected to in the *Disjoint* and *Overlapped* setting, we introduce the *Fully Disjoint* setting. In this setting each task only contains pixels belonging to the current set of classes, which therefore avoids the interference that originates from the semantic shift of the background class. We utilize this setting to study the impact of the semantic background shift on catastrophic forgetting, but we note that this is an unrealistic scenario, as in semantic segmentation classes naturally re-appear.

*Model:* Similar to [17] we use ERFNet [36] in our evaluation, as the underlying effects of forgetting are similar to more established models like DeepLabV3+ [4], while at the same time ERFNet is more susceptible to forgetting due to its smaller size, which exaggerates the effect of the causes of forgetting. However, we confirm our findings in Appendix B with DeepLabv3+ and U-Net [37]. We do not use any pre-trained models in our experiments, as pre-training is known to

<sup>4</sup> As the focus of this paper is to understand the general causes of forgetting in CiSS, we leave the study of the impact of different splits, more classes and longer task sequences to future work.

increase robustness to catastrophic forgetting [12,27]. Instead, we use the same randomly initialized weights for every method in our experiments.

*Methods Compared:* In our evaluation, we focus on evaluating naive fine-tuning approaches, representative regularization and replay methods as these cover the basic concepts that are considered as potential solutions for CiSS. For prior regularization methods we consider EWC [18] and MAS [1]. For data regularization methods we use LwF [23], with the modification of [19], in which the distillation loss is only applied to the parts of the image that are labelled as background. We explicitly do not consider incremental improvements of knowledge distillation based approaches of [8,3,29,24], as our focus is on evaluating the underlying causes of forgetting. For replay we save 20 samples for each class in the buffer. For every experiment we also list results for the *offline* model, which is jointly trained on all classes in one step. Further information regarding the implementation details and selected hyperparameters can be found in the supplementary material.

### 5.1 Semantic Background Shift and Class Confusion

First, we study the impact of the semantic background shift on forgetting by comparing the results of the selected CiSS methods on the *Overlapped*, *Disjoint*, and *Full Disjoint* tasks. These tasks have a varying degree of semantic shift of the background class. The results are displayed in Table 1. For the *Overlapped* and *Disjoint* tasks, it can be noted that only LwF and Replay effectively learn to discriminate between all classes. EWC and MAS effectively mitigate forgetting of old classes (0–15) compared to Fine-Tuning, but they also inhibit the learning of new classes (16–20). The reason for the low mIoU for EWC and MAS on all classes can be inferred from the confusion matrix, in which EWC and MAS exhibit a strong bias to the background class and a minor bias towards a few new classes, shown in in Figs. 3b and 3c. LwF and Replay reduce both biases. In the *Overlapped* and *Disjoint* setting, LwF and Replay achieve similar performance, as in this setting LwF can effectively replay old classes by discovering them in the background of new images. However, once these classes do not re-appear in the background, as is the case in the *Full Disjoint* task, LwF develops a strong bias towards selected new classes. In contrast to this, Replay benefits from the *Full Disjoint* setting as the training is no longer affected by the semantic background shift. Similarly, MAS and EWC also show significant improvement in this setting, as they benefit from the fact that old classes do not appear as background in the new task, thus not interfering with previously learned knowledge. This is especially noticeable in the confusion matrices of the *Full Disjoint* setting, in which the bias towards the background class is greatly reduced<sup>5</sup>. This demonstrates that the semantic shift of the background class is a significant cause of forgetting for prior regularization methods like EWC and MAS and that it has a noticeable effect on replay as well. The results also indicate that

<sup>5</sup> The confusion matrices are shown in the supplementary material.

knowledge distillation is the most effective method to combat this semantic shift.

Finally, upon a closer look at the semantics of the false positives, we see that old classes that are falsely assigned to a new class share semantic and visual properties. In this case *bus* (6), *car* (7), *boat* (5) are assigned to *train* (19), whereas *cow* (10) and *horse* (13) are classified as *sheep* (17). The remaining classes that do not share such a relationship with new classes are falsely classified as background. This confusion can only be alleviated by either Replay or LwF when old classes re-appear in the background in subsequent tasks. To clarify which layers of the model contribute the most to the bias for the background class and newly learned classes, next up we investigate the internal activation drift.

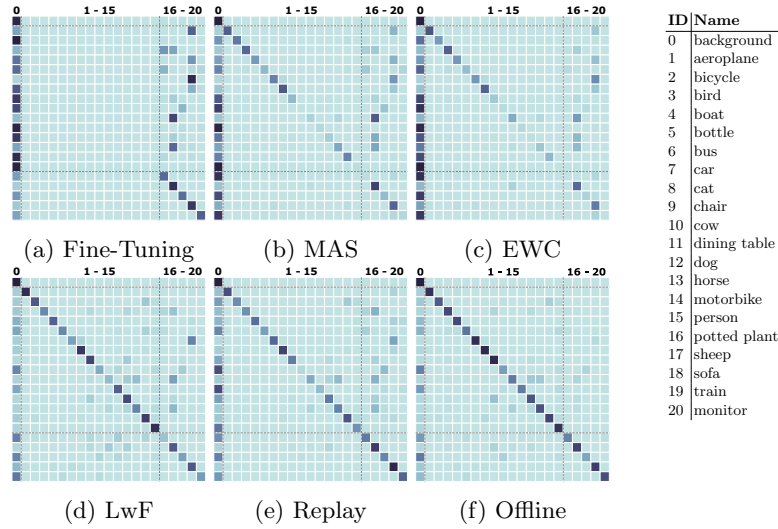


Fig. 3: Confusion matrices after training on PascalVoc-15-5 (disjoint). The confusion matrix for Fine-Tuning a) shows a severe bias to the background class and the classes of the most recent task (16-20). EWC [18] and MAS [1] decrease the bias in exchange for worse accuracy on the most recent classes. Replay and LwF [23] reduce the bias towards new classes and the background.

## 5.2 Measuring the effect of Activation Drift

We use the Dr. Frankenstein tool set to measure the activation drift for each layer between the model before and after learning  $T_1$ , to investigate which layers are affected the most by the internal activation drift. We follow the setup from Fig. 2 without an additional stitching layer, meaning that the activations of the layer  $n$  under examination  $f_1$  are directly propagated to the layer  $n + 1$  in  $f_0$ .



Table 1: Results of Semantic Segmentation on Pascal-VOC 2012 in Mean IoU (%) on the overlapped, disjoint and full disjoint settings.

PascalVoc 15-5 Semantic Segmentation									
Method	Overlapped			Disjoint			Full Disjoint		
	0-15	16-20	all	0-15	16-20	all	0-15	16-20	all
Fine-Tuning	4.5	22.2	8.8	4.6	23.0	9.0	5.1	16.3	7.8
MAS [1]	24.1	10.8	21.0	30.6	12.9	26.4	35.6	12.9	30.2
EWC [18]	23.8	11.8	21.0	28.1	10.1	23.8	35.2	10.9	29.4
Replay	41.3	31.5	39.0	42.2	29.1	39.1	48.2	28.8	43.6
LwF [23]	45.8	28.2	41.6	44.4	25.4	39.9	35.9	12.6	30.4
Offline	55.7	47.6	53.8	55.7	47.6	53.8	55.7	47.6	53.8

The resulting Frankenstein Network  $f_{1,0}^n$  is then evaluated on the data of the first task (classes 0–15). The mIoU relative to the initial performance on the first task is shown in Fig. 4. Overall, we observe that the activations of the early layers of the network (layers 0–4) stay very stable for every approach and that later layers are disproportionately affected by activation drift, especially the layers of decoder. A similar observation was made in prior work on image classification [6,31]. This confirms that forgetting in CiSS is also mostly affecting deeper layers of the network. Furthermore, the results show that EWC and MAS effectively prevent severe activation drift in the deeper layers of the encoder, dropping only to about 90% of the initial mIoU on the disjoint task, compared to the 30% of the stitched fine-tuning model. This suggests that forgetting for EWC and MAS is less severe as accuracy in Table 1 would reveal. The reason for this could be two-fold: Firstly, the bad accuracy could be attributed to the classifier being biased towards new classes (task-recency bias) or secondly that the regularization methods fail to learn meaningful features that help to discriminate between old and new classes as they are never trained jointly. While a biased classifier is fixed more easily, inter-task confusion is a fundamental shortcoming of prior regularization methods [22].

Another striking phenomenon is the severe change of activations at the third decoder layer (layer 18) that Fine-Tuning, MAS and EWC show on the *Disjoint* task. The predictions of the specific Frankenstein Networks  $f_{1,0}^{17}$  and  $f_{1,0}^{18}$  in Fig. 5, show that  $f_{1,0}^{17}$  is able to correctly classify old classes (*bike*, *person*) as such, but  $f_{1,0}^{18}$  assigns the background class to these regions. Therefore, it can be concluded that the sudden activation change of MAS originates from the fact that features that were evidence for old classes in  $f_0$  are now attributed evidence for the background class. This validates that the semantic shift of the background class is mostly affecting the later layers of the decoder and that the features for old classes are in fact not forgotten, but assigned to the background class. Similar observations can be made for EWC and Fine-Tuning. In Appendix D we measure the similarity with Centered Kernel Alignment to support our observations. When completely avoiding the semantic background shift in the *Full Disjoint* task, we observe that activation drift for the fine-tuned model is

much more pronounced in the middle layers of the encoder (layer 8-15), which implies that the re-appearance of old classes, even though they are labelled as background, is mitigating the activation drift in the earlier layers of the model.




Fig. 4: Activation drift between  $f_1$  to  $f_0$  measured by relative mIoU on the first task of the Frankenstein Networks stitched together at specific layers (horizontal axis). The layers of the encoder are layer 0–15 (grey area), the decoder layers are 17-20 (white area). The activations in the early layers of the encoder stay very stable for all methods, whereas EWC, MAS and Fine-Tuning have a severe drift in activations in the decoder layers of the network, which is clear evidence that forgetting is mostly affecting later layers in the *Disjoint* setting.

### 5.3 The Impact of Inter-task Confusion on the Encoder

As we observed in Sect. 5.2 that the early layers of a model trained with a continual learning method do not suffer from severe activation drift, in this experiment we investigate how useful the learned features of the encoder of the different methods are to discriminate between all classes. Therefore, we measure Decoder Retrain Accuracy, introduced in Sect. 4.2, for which the decoder of the model is retrained on all classes and subsequently evaluated on the test set. The first observation to be made when looking at the retraining accuracy in Table 3, is that all methods improve after decoder retraining, though EWC, MAS and Fine-Tuning show bigger improvements than LwF and Replay. This again confirms that forgetting in the encoder is not as severe for Fine-Tuning, EWC and MAS as the accuracy indicates. Furthermore, it also verifies that MAS and EWC are effectively preserving important features for old classes in the encoder, but that the biased decoder layer might wrongly attribute important features for old classes to the background class or new classes, which leads to a severe amount of misclassifications.

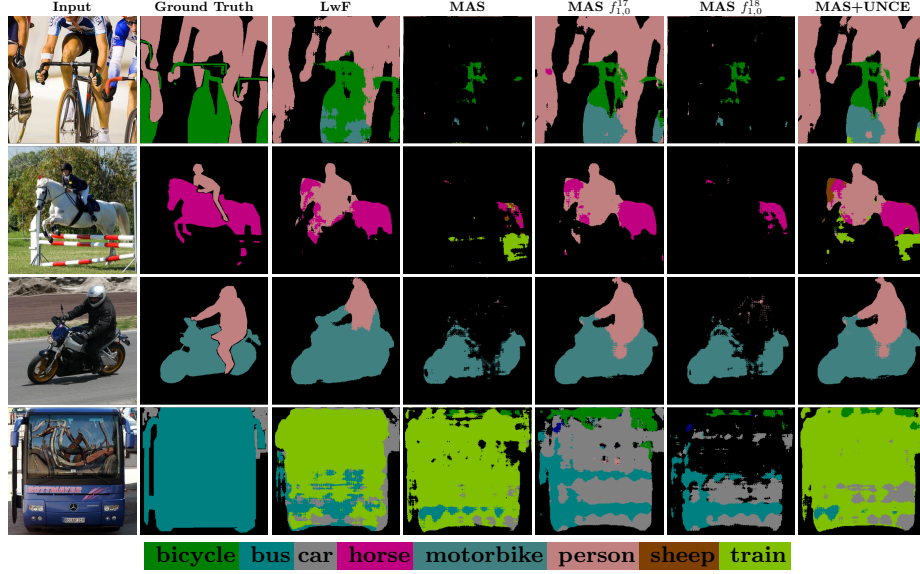


Fig. 5: Visualizations of the segmentation maps for LwF, MAS and the resulting Frankenstein Networks of MAS  $f_{1,0}^{17}$  and  $f_{1,0}^{18}$ . The predictions of  $f_{1,0}^{17}$  and  $f_{1,0}^{18}$  show that up until layer 17 the information for previously learned classes *person* and *horse* is still available, but is assigned to the *background* in layer 18.

Still, as EWC, MAS and Fine-Tuning do not achieve a comparable mIoU as LwF or Replay after decoder retraining, it can be concluded that the learned features of the encoder are less useful for discriminating between all classes. Specifically the aforementioned related classes *bus* (6), *car* (7), *boat* (5), *train* (19), as well as *cow* (10), *horse* (13), *sheep* (17) cannot be effectively classified after retraining, compare Fig. 6. We hypothesize that Replay does not suffer from inter-task confusion since old classes are taken into account when optimizing for new classes, leading to more discriminative features. The same holds for LwF in the *Overlapped* and *Disjoint* setting, in which old classes are effectively replayed, by using soft-labels for old classes that are discovered in the background.

#### 5.4 Mitigating background bias and the task recency bias

A simple method to reduce the recency bias in the classification layer that is used in class-incremental classification is to calculate the cross-entropy loss (CE) only for classes of the current training set [25]. This enforces that errors are only back-propagated for probabilities that are related to the current set of classes:

$$\ell_{ce}(y, q) = -\frac{1}{I} \sum_{i \in \mathcal{I}} \sum_{c \in \mathcal{C}_t} y_{i,c} \log(q_{i,c}) \quad (1)$$




Fig. 6: Confusion matrices before (a) and after b), c) retraining the decoder on all classes of PascalVoc2012.

However, in the case of CiSS, this addition has proven to be less effective than the standard cross-entropy loss [8]. Therefore, an unbiased cross-entropy loss (UNCE) is proposed in [3], which accounts for the uncertainty of the content of the background class. This is achieved by comparing the pixels that are labelled as background with the probability of having either an old class or the background predicted by the model:

$$\ell_{ce}(y, q) = -\frac{1}{I} \sum_{i \in \mathcal{I}} \sum_{c \in \mathcal{C}} y_{i,c} \log(\hat{q}_{i,c}) \quad (2)$$

$$\hat{q}_{i,c} = \begin{cases} \sum_{k \in \mathcal{C}_{t-1}} q_{i,k} & \text{if } c = b \\ q_{i,c} & \text{otherwise} \end{cases} \quad (3)$$

In addition, Weight Normalization Layers [38] were also successfully used in classification tasks to address the recency bias [21]. In the next experiment we study the impact of UNCE and UNCE combined with Weight Normalization to combat the recency and background bias in CiSS.

The results in Table 2 show that UNCE improves the accuracy for all approaches on the *Disjoint* setting. Specifically, the prior regularization methods MAS and EWC show a significantly higher accuracy compared to the basic cross-entropy loss. This can be attributed to the fact that UNCE effectively mitigates the background bias, as we can see in the confusion matrix in Fig. 3f and the segmentation maps in Fig. 5. In addition, the severe activation drift that we observed in Sect. 5.2 between layer 17 and 18 for MAS completely vanishes with the use of UNCE. Therefore, UNCE effectively resolves the confusion of the old classes with the background class. This confirms the assumption that a major cause of forgetting in Sect. 5.1 was in fact a bias of the classifier towards the background and the new classes. However, the confusion matrices shows that while the background bias is severely reduced by using UNCE, the semantic confusion of old and new classes is amplified.

In the *Full Disjoint* setting the use of UNCE does not improve the performance as much as it does in the *Disjoint* setting. The reason is that in the *Full Disjoint*

setting the pixels of old classes do not re-occur and thus the de-biasing effect of UNCE is decreased. Therefore, the content of the background class plays an important role to mitigate forgetting. Of the selected approaches, only Replay benefits from the addition of the Weight Normalization layer. Finally, we note that EWC and MAS, with the addition of UNCE, show competitive performance to the remaining approaches and more recent approaches like MiB [3], even without the use of knowledge-distillation or replay. However, we hypothesize that for longer task sequence and more classes MiB will outperform prior regularization methods, as they will not be able to learn discriminative features.

Table 2: Results on Pascal-15-5 in mIoU (%) on the disjoint and full-disjoint settings with: Cross-Entropy Loss (CE), Unbiased Cross-Entropy (UNCE) and UNCE combined with a Weight Normalization (UNCE+WN). UNCE effectively reduces forgetting for all approaches, especially for EWC and MAS.

Method	PascalVoc 15-5 (disjoint)									PascalVoc 15-5 (full disjoint)								
	CE			UNCE			UNCE+WN			CE			UNCE			UNCE+WN		
	0-15	16-20	all	0-15	16-20	all	0-15	16-20	all	0-15	16-20	all	0-15	16-20	all	0-15	16-20	all
Fine-Tuning	4.6	23.0	9.0	10.4	21.8	13.1	16.5	21.6	17.7	5.1	16.3	7.8	6.0	15.5	8.3	7.7	15.2	9.5
EWC [18]	28.1	10.1	23.8	48.2	11.6	39.4	17.0	9.5	15.2	35.2	10.9	29.4	41.1	9.8	33.6	34.8	9.7	28.8
MAS [1]	30.6	12.9	26.4	45.8	14.4	38.3	41.0	13.9	34.6	35.6	12.9	30.2	39.1	12.3	32.7	32.5	11.8	27.6
LwF [23]	44.4	25.4	39.9	45.3	22.9	40.0	46.6	19.7	40.2	35.9	12.6	30.4	38.0	13.8	32.2	38.8	13.3	32.8
Replay	42.2	29.1	39.1	47.2	31.4	43.5	48.1	31.9	44.3	48.2	28.8	43.6	47.7	28.0	43.0	48.8	28.5	44.0
MiB [3]	-	-	-	48.6	21.7	42.2	49.4	24.1	43.3	-	-	-	47.6	19.7	41.0	48.6	20.7	42.0

Table 3: Decoder Retraining Results on Pascal-VOC.  $mIoU_I$  and  $mIoU_R$  denote the mIoU (%) before and after retraining.

PascalVoc 15-5 - Decoder Retraining Accuracy						
Method	Overlapped		Disjoint		Full Disjoint	
	$mIoU_I$	$mIoU_R$	$mIoU_I$	$mIoU_R$	$mIoU_I$	$mIoU_R$
Fine-Tuning	8.8	28.0	9.0	27.9	7.8	22.0
MAS [1]	21.0	34.3	26.4	36.2	30.2	37.3
EWC [18]	21.0	34.3	23.8	35.1	29.4	36.9
LwF [23]	41.6	45.3	39.9	43.3	30.4	38.1
Replay	39.0	42.6	39.1	42.9	43.6	45.6
Offline	53.8	54.6	53.8	54.6	53.8	54.6

Table 4: Classification Results on PascalVoc-15-5.

PascalVoc 15-5 Classification			
Method	Full Disjoint		
	0-15	16-20	all
Fine-Tuning	13.6	27.6	17.1
MAS [1]	32.0	27.8	31.0
EWC [18]	27.2	25.4	26.8
LwF [23]	39.6	32.7	37.9
Replay	42.1	34.2	40.1
Offline	51.3	54.8	52.2

## 5.5 The Role of the Background Class to overcome forgetting

The prior observations show that in CiSS the semantic shift of the background class is a major cause of a rapid drop in performance if not addressed correctly. However, if the uncertainty of the content of the background class is taken into account by either UNCE, Knowledge Distillation or both, the appearance of old classes in the background can to some extent be used for replay. In the experiments of the *Full Disjoint* setting we see that once classes do not reoccur, these methods are less effective, whereas explicit replay of classes benefits from avoiding the semantic shift. The ranking of the methods of *Full Disjoint* setting

in CiSS is also similar to the ranking of the same methods for class-incremental image classification, compare Table 4. This indicates that the discrepancy in performance of LwF and replay in image classification is due to the missing background class. Looking at it the other way around, this could also mean that introducing an out-of-set class for image classification could help to reduce forgetting in the class-incremental setting without requiring explicit replay via stored samples, as the re-appearing classes in the out-of-set-class play a similar role as explicit replay.

## 6 Conclusion

We studied the major causes of catastrophic forgetting in CiSS, answering how it manifests itself in the hidden representations of the network and how the background class both causes severe forgetting and decreases activation drift. Using representational similarity techniques, we demonstrated that forgetting is concentrated at deeper layers and that re-appearing classes mitigate activation drift in the encoder even when they are labelled as background. Moreover, we show that EWC and MAS are effectively reducing representational drift in the later layers of the encoder, but suffer from severe background and recency bias, which leads to the sudden drop in accuracy. These biases manifest themselves in deeper layers of the networks by assigning previous discriminating features for the previous classes to the background class or visually related classes.

The background bias can be effectively alleviated using an unbiased cross-entropy loss, which leads to a significant improvement for all methods, when classes re-appear in the background of new training data. Finally, we find that only methods that in some form replay old classes during training of new classes can learn to correctly discriminate between all classes after incremental training, as otherwise the model fails to learn to discriminate between new and old classes that share similar visual features. Overall, the results of our work provide a foundation for deeper understanding of the principles of forgetting in CiSS and open the door to future directions to explore methods for its mitigation.

**Acknowledgments** The research leading to these results is funded by the German Federal Ministry for Economic Affairs and Climate Action within the project “KI Delta Learning“ (Förderkennzeichen 19A19013T). The authors would like to thank the consortium for the successful cooperation.

## References

1. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., Tuytelaars, T.: Memory aware synapses: Learning what (not) to forget. In: Proceedings of the European Conference on Computer Vision (ECCV) (9 2018)
2. Cermelli, F., Fontanel, D., Tavera, A., Ciccone, M., Caputo, B.: Incremental learning in semantic segmentation from image labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4371–4381 (June 2022)
3. Cermelli, F., Mancini, M., Bulò, S.R., Ricci, E., Caputo, B.: Modeling the background for incremental learning in semantic segmentation. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition pp. 9230–9239 (2 2020). <https://doi.org/10.1109/CVPR42600.2020.00925>, <https://arxiv.org/abs/2002.00718>
4. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European Conference on Computer Vision (ECCV) (September 2018)
5. Csiszárík, A., Kőrösi-Szabó, P., Matszangosz, A., Papp, G., Varga, D.: Similarity and matching of neural network representations. In: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) Advances in Neural Information Processing Systems. vol. 34, pp. 5656–5668. Curran Associates, Inc. (2021), <https://proceedings.neurips.cc/paper/2021/file/2cb274e6ce940f47beb8011d8ecb1462-Paper.pdf>
6. Davari, M., Asadi, N., Mudur, S., Aljundi, R., Belilovsky, E.: Probing representation forgetting in supervised and unsupervised continual learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16712–16721 (June 2022)
7. Delange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G., Tuytelaars, T.: A continual learning survey: Defying forgetting in classification tasks. IEEE Transactions on Pattern Analysis and Machine Intelligence pp. 1–1 (2021). <https://doi.org/10.1109/TPAMI.2021.3057446>
8. Douillard, A., Chen, Y., Dapogny, A., Cord, M.: Plop: Learning without forgetting for continual semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
9. Douillard, A., Chen, Y., Dapogny, A., Cord, M.: Tackling catastrophic forgetting and background shift in continual semantic segmentation. CoRR **abs/2106.15287** (2021), <https://arxiv.org/abs/2106.15287>
10. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html> (2012)
11. French, R.: Catastrophic forgetting in connectionist networks. Trends in Cognitive Sciences **3**(4), 128–135 (1999). [https://doi.org/10.1016/s1364-6613\(99\)01294-2](https://doi.org/10.1016/s1364-6613(99)01294-2)
12. Gallardo, J., Hayes, T.L., Kanan, C.: Self-Supervised Training Enhances Online Continual Learning. In: British Machine Vision Conference (BMVC) (2021)
13. Hayes, T.L., Kafle, K., Shrestha, R., Acharya, M., Kanan, C.: Remind your neural network to prevent catastrophic forgetting. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
14. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. In: NIPS Deep Learning and Representation Learning Workshop (2015), <http://arxiv.org/abs/1503.02531>

15. Hou, S., Pan, X., Loy, C.C., Wang, Z., Lin, D.: Learning a unified classifier incrementally via rebalancing. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
16. Hsu, Y.C., Liu, Y.C., Ramasamy, A., Kira, Z.: Re-evaluating continual learning scenarios: A categorization and case for strong baselines. In: NeurIPS Continual learning Workshop (2018), <https://arxiv.org/abs/1810.12488>
17. Kalb, T., Roschani, M., Ruf, M., Beyerer, J.: Continual learning for class- and domain-incremental semantic segmentation. In: 2021 IEEE Intelligent Vehicles Symposium (IV). pp. 1345–1351 (2021). <https://doi.org/10.1109/IV48863.2021.9575493>
18. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., Hadsell, R.: Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* **114**(13), 3521–3526 (2017). <https://doi.org/10.1073/pnas.1611835114>, <https://www.pnas.org/doi/abs/10.1073/pnas.1611835114>
19. Klingner, M., Bär, A., Donn, P., Fingscheidt, T.: Class-incremental learning for semantic segmentation re-using neither old data nor old labels. In: 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC). pp. 1–8. IEEE (2020)
20. Kornblith, S., Norouzi, M., Lee, H., Hinton, G.: Similarity of neural network representations revisited. In: 36th International Conference on Machine Learning, ICML 2019. vol. 2019-June, pp. 6156–6175 (5 2019), <https://arxiv.org/abs/1905.00414>
21. Lesort, T., George, T., Rish, I.: Continual learning in deep networks: an analysis of the last layer (2021). <https://doi.org/10.48550/ARXIV.2106.01834>, <https://arxiv.org/abs/2106.01834>
22. Lesort, T., Stoian, A., Filliat, D.: Regularization shortcomings for continual learning (2020)
23. Li, Z., Hoiem, D.: Learning without Forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **40**(12), 2935–2947 (2018). <https://doi.org/10.1109/TPAMI.2017.2773081>
24. Maracani, A., Michieli, U., Toldo, M., Zanuttigh, P.: Recall: Replay-based continual learning in semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 7026–7035 (10 2021)
25. Masana, M., Liu, X., Twardowski, B., Menta, M., Bagdanov, A.D., van de Weijer, J.: Class-incremental learning: survey and performance evaluation. *arXiv preprint arXiv:2010.15277* (2020)
26. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation - Advances in Research and Theory* **24**(C), 109–165 (Jan 1989). [https://doi.org/10.1016/S0079-7421\(08\)60536-8](https://doi.org/10.1016/S0079-7421(08)60536-8)
27. Mehta, S.V., Patil, D., Chandar, S., Strubell, E.: An empirical investigation of the role of pre-training in lifelong learning (2021). <https://doi.org/10.48550/ARXIV.2112.09153>, <https://arxiv.org/abs/2112.09153>
28. Michieli, U., Zanuttigh, P.: Incremental learning techniques for semantic segmentation. *Proceedings - 2019 International Conference on Computer Vision Workshop, ICCVW 2019* pp. 3205–3212 (7 2019). <https://doi.org/10.1109/ICCVW.2019.00400>, <http://arxiv.org/abs/1907.13372>
29. Michieli, U., Zanuttigh, P.: Continual semantic segmentation via repulsion-attraction of sparse and disentangled latent representations. In: *Proceedings of the*



- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1114–1124 (6 2021)
30. Mirzadeh, S.I., Farajtabar, M., Gorur, D., Pascanu, R., Ghasemzadeh, H.: Linear mode connectivity in multitask and continual learning. In: International Conference on Learning Representations (2021), [https://openreview.net/forum?id=Fmg\\_fQYUejf](https://openreview.net/forum?id=Fmg_fQYUejf)
  31. Neyshabur, B., Sedghi, H., Zhang, C.: What is being transferred in transfer learning? In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 512–523. Curran Associates, Inc. (2020), <https://proceedings.neurips.cc/paper/2020/file/0607f4c705595b911a4f3e7a127b44e0-Paper.pdf>
  32. Nguyen, C.V., Achille, A., Lam, M., Hassner, T., Mahadevan, V., Soatto, S.: Toward Understanding Catastrophic Forgetting in Continual Learning (8 2019), <http://arxiv.org/abs/1908.01091>
  33. Ramasesh, V.V., Dyer, E., Raghu, M.: Anatomy of catastrophic forgetting: Hidden representations and task semantics. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=LhY8QdUGSuw>
  34. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: iCaRL: incremental classifier and representation learning. In: CVPR (2017)
  35. Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T., Wayne, G.: Experience replay for continual learning. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 32. Curran Associates, Inc. (2019), <https://proceedings.neurips.cc/paper/2019/file/fa7cdfad1a5aaf8370ebeda47a1ff1c3-Paper.pdf>
  36. Romera, E., Álvarez, J.M., Bergasa, L.M., Arroyo, R.: Erfnet: Efficient residual factorized convnet for real-time semantic segmentation. *IEEE Transactions on Intelligent Transportation Systems* **19**(1), 263–272 (2018). <https://doi.org/10.1109/TITS.2017.2750080>
  37. Ronneberger, O., et al.: U-net: Convolutional networks for biomedical image segmentation (2015). <https://doi.org/10.48550/ARXIV.1505.04597>, <https://arxiv.org/abs/1505.04597>
  38. Salimans, T., Kingma, D.P.: Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 29. Curran Associates, Inc. (2016), <https://proceedings.neurips.cc/paper/2016/file/ed265bc903a5a097f61d3ec064d96d2e-Paper.pdf>
  39. Shin, H., Lee, J.K., Kim, J., Kim, J.: Continual learning with deep generative replay. *Advances in neural information processing systems* **30** (2017)
  40. Tasar, O., Tarabalka, Y., Alliez, P.: Incremental Learning for Semantic Segmentation of Large-Scale Remote Sensing Data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **12**(9), 3524–3537 (2019). <https://doi.org/10.1109/JSTARS.2019.2925416>
  41. van de Ven, G.M., Siegelmann, H.T., Tolia, A.S.: Brain-inspired replay for continual learning with artificial neural networks. *Nature Communications* **11**(1) (2020). <https://doi.org/10.1038/s41467-020-17866-2>
  42. Yu, L., Liu, X., van de Weijer, J.: Self-training for class-incremental semantic segmentation. *IEEE Transactions on Neural Networks and Learning Systems* pp. 1–12 (2022). <https://doi.org/10.1109/TNNLS.2022.3155746>
  43. Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. In: Precup, D., Teh, Y.W. (eds.) *Proceedings of the 34th International Conference*

on Machine Learning. Proceedings of Machine Learning Research, vol. 70, pp. 3987–3995. PMLR, International Convention Centre, Sydney, Australia (8 2017), <http://proceedings.mlr.press/v70/zenke17a.html>