

Improving Prediction Backward-Compatibility in NLP Model Upgrade with Gated Fusion

Yi-An Lai[♣] Elman Mansimov[♣] Yuqing Xie^{♡,*} Yi Zhang[♣]

[♣]AWS AI Labs [♡]University of Waterloo

{yianl,mansimov,yizhngn}@amazon.com

yuqing.xie@uwaterloo.ca

Abstract

When upgrading neural models to a newer version, new errors that were not encountered in the legacy version can be introduced, known as *regression*¹ errors. This inconsistent behavior during model upgrade often outweighs the benefits of accuracy gain and hinders the adoption of new models. To mitigate regression errors from model upgrade, distillation and ensemble have proven to be viable solutions without significant compromise in performance. Despite the progress, these approaches attained an incremental reduction in regression which is still far from achieving backward-compatible model upgrade. In this work, we propose a novel method, *Gated Fusion*, that promotes backward compatibility via learning to mix predictions between old and new models. Empirical results on two distinct model upgrade scenarios show that our method reduces the number of regression errors by 62% on average, outperforming the strongest baseline by an average of 25%.

1 Introduction

In order to achieve a smooth continuous improvement of NLP applications, it is critical to guarantee consistent operation of the system after an upgrade. New errors introduced during the model upgrade interfere with the existing user experience and are considered to be a *regression* in the quality. Due to the difficulty of modularizing or explaining the behavior of deep neural networks, traditional software regression tests are inapplicable to neural based systems. The cost of arduous error analysis and model patching often exceeds the benefits of model upgrades. Developing methods that ensure backward compatibility during model upgrades without compromise in performance becomes a valuable research direction (Yan et al., 2021; Xie et al., 2021; Träuble et al., 2021; Cai et al., 2022).

*Work done during author’s internship at AWS AI Labs.

¹Within this work, *regression* denotes performance degradation in software systems, instead of the statistical technique for estimating relationships among variables.

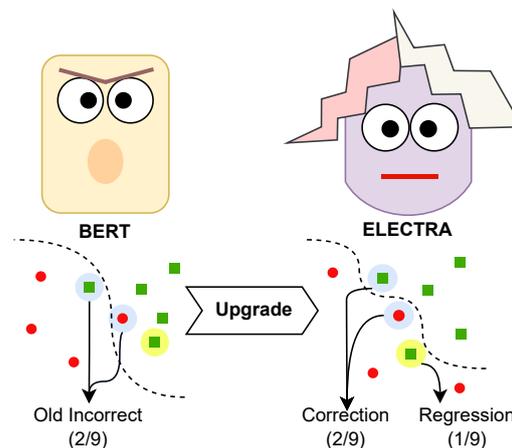


Figure 1: Illustration of regression errors when upgrading from BERT (Devlin et al., 2019) to ELECTRA (Clark et al., 2020) for classification. Red circles and green squares denote examples of different classes. Dashed lines represent decision boundaries.

The *prediction backward-compatible model upgrade* problem aims to improve consistency of correct classification predictions between legacy and upgraded models without accuracy loss. Yan et al. (2021) first studied backward compatibility during model upgrade on image classification tasks. They proposed to enforce the positive congruence of the new model with the old one by applying a knowledge distillation objective (Hinton et al., 2015) objective with re-weighting of training samples. Later, Xie et al. (2021) extended the work of Yan et al. (2021) by investigating the backward compatibility in NLP classification tasks. They found that their proposed distillation-based approach can help decrease the regression errors of specific linguistic phenomena in NLP classification tasks.

Despite progress with both distillation- and ensemble-based regression-mitigation approaches, there are limitations that prevent them from broad practical adoption in ML operations. Distillation-based methods attempt to transfer the prediction power of the old model to the new one on potential

regression instances (Hinton et al., 2015). However, given the huge complexity of current neural architectures and relatively scarce training data in downstream tasks, models could have insufficient data to reliably estimate the probable regression cases and carry out the transfer on them (Xie et al., 2021; Cai et al., 2022). On the other hand, model ensemble aggregates predictions from differently-trained new models but bears no connection with the legacy version (Yan et al., 2021). These limitations reveal the two major challenges when striving to ensure backward compatibility. First, the new model could have distinct inductive bias and prediction behavior than the old system, rooted from inherent differences such as architecture, model size, and pretraining procedure (Liu et al., 2021). Second, during new model training, a reliable mechanism is needed in place to bridge the gap between two models and mitigate potential inconsistencies.

Inspired by the strength and weakness of prior approaches, we propose *Gated Fusion* to integrate old and new models via gating mechanism (Hochreiter and Schmidhuber, 1997; Chung et al., 2014; Gu et al., 2016), essentially a light-weight ensemble of legacy and upgrade models connected via a learned fusion gate. Specifically, we add a *learned* gate on top of the new model. We combine the logits from old and new models according to the weight from the gate. We train our Gated Fusion model by minimizing the standard cross-entropy error. The intuition is that the gate could learn to put more weights on the old model when the new model cannot produce correct predictions, effectively doing fall-backs that optimizes backward compatibility.

Empirical results demonstrate that our proposed approach outperforms other competing methods significantly, where we can obtain on average 62% reduction of total negative flips, i.e. new errors caused by the model upgrade, without any degradation in accuracy performance. The effectiveness of Gated Fusion is validated across three diverse classification tasks and two distinct model upgrade scenarios (a) upgrade to larger model size (b) upgrade to advanced pretrained model, where consistent results are attained across the board.

Our main contributions are as follows:

- We propose Gated Fusion that integrates old and new models via gating mechanism for backward-compatible model upgrade;
- We evaluate competing methods on two distinct and challenging model upgrade scenarios

across three diverse classification tasks;

- Empirical results show that our proposed approach significantly outperforms competing methods and achieves regression reductions by a large margin across the board.

2 The Backward-Compatible Model Upgrade Problem

The goal of backward-compatible model upgrade is to minimize regression errors without compromising the accuracy performance during model upgrade (Yan et al., 2021; Xie et al., 2021). In this work, we aim to improve the backward compatibility of model predictions in the NLP classification tasks. Following Xie et al. (2021), we study the scenario where the underlying pretrained language model (LM) is being upgraded.

Let x be a natural language input with a class label $y \in \{1, 2, \dots, C\}$. $\mathcal{D} = \{x_i, y_i\}_{i=1}^N$ denotes a set of N examples with corresponding labels. A classifier f estimates the class probabilities given the input $\vec{f}(x) = (p(y = 1|x), \dots, p(y = C|x))^\top$. When upgrading from an *old* model f_{old} to a *new* model f_{new} , normally with distinct architectures and trained on the same data, an improved model f^* is produced based on f_{old} and f_{new} . Our goal is for f^* to minimize *regression errors* as an additional objective, while still achieving comparable performance to f_{new}^o , the new model trained in the vanilla setting. Note that f^* could be multiple times larger than f_{new}^o , with model ensemble of f_{new}^o as one example (Yan et al., 2021).

Measuring Backward Compatibility. The backward compatibility is measured via quantifying regression errors on a given regression measurement set $\mathcal{D}_{reg} = \{x_i, y_i\}_{i=1}^M$. \mathcal{D}_{reg} could be a hidden customer test set comprising critical use cases, a set of behavioral testing examples for targeted evaluation (Ribeiro et al., 2020), or the development split from the dataset of interest. In this work, we take the development set as our \mathcal{D}_{reg} for evaluation.

For classification, regression errors are characterized by *negative flips*, denoted as \mathcal{R}_{NF} – the portion of samples in \mathcal{D}_{reg} that flip from correct prediction $f_{old}(x_i) = y_i$ to incorrect output $f_{new}(x_i) \neq y_i$ during model upgrade:

$$\mathcal{R}_{NF}(\mathcal{D}_{reg}, \vec{f}_{old}, \vec{f}_{new}) = \frac{|\{x | f_{old} = y, f_{new} \neq y\}|}{|\mathcal{D}_{reg}|}. \quad (1)$$

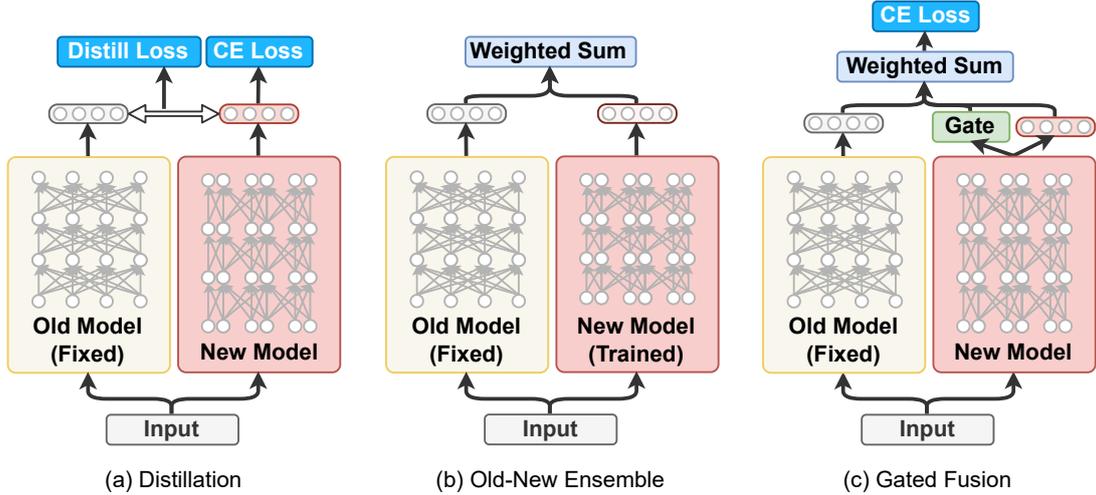


Figure 2: Methods to improve prediction backward compatibility during model upgrade. (a) Distillation-based approach to align predicted logits on potential regression instances (Xie et al., 2021). (b) Ensemble of old and new models via weighted sum of either predicted logits or probabilities. (c) Our proposed Gated Fusion that learns a gate as a soft switch to dynamically determine whether to fall back to previous predictions.

One thing to emphasize is that maximizing classifier performance does not necessarily help in minimizing \mathcal{R}_{NF} (Yan et al., 2021; Xie et al., 2021).

3 Gated Fusion: Methodology

3.1 Method Overview

To improve backward compatibility in model upgrade, it’s crucial to have a mechanism that detects potential regression errors and mitigates them when making predictions. We propose Gated Fusion (GF) to achieve this by learning a gate as a soft switch to choose between generating predictions by the new model or resorting to outputs of the old model. Gated Fusion is inspired by the gating mechanism widely used in other applications. For example, mixing word copying mode with word generation mode for language modeling (Merity et al., 2016) and summarization (See et al., 2017).

Our proposed Gated Fusion f_{GF}^* consists of the old model f_{old} , the new model f_{new} , and a gating network g_θ . The old model f_{old} is the legacy version before upgrade where the parameters are fixed. The new model f_{new} has the same architecture as f_{new}^o and is randomly initialized. The gating network g_θ is a multi-layer feed-forward network with sigmoid function. It produces a scalar weight α_{gate} in the range $[0, 1]$ from the output layer of f_{new} , denoted as \mathcal{E}_{new} :

$$\alpha_{gate}(x) = g_\theta(\mathcal{E}_{new}(x)). \quad (2)$$

We use α_{gate} to combine the logits of old and new

models as our final outputs:

$$l_{GF}^*(y|x) = (1 - \alpha_{gate}) \cdot \frac{l_{old}(y|x)}{T} + \alpha_{gate} \cdot l_{new}(y|x), \quad (3)$$

where $l(y|x)$ denotes predicted logits from models and T is the temperature scaling to regularize the magnitude of old model’s logits. f_{new} and g_θ are then jointly trained end-to-end with cross-entropy loss between our output logits $l_{GF}^*(y|x)$ and label distributions on downstream tasks.

The intuition behind Gated Fusion is that when f_{new} makes a mistake while f_{old} produces the correct output, the gate g_θ will learn to put more weight on f_{old} in order to minimize the final classification loss. This process effectively mitigates potential negative flips introduced by the model upgrade and thus improves the backward compatibility of final predictions.

3.2 Training and Inference

In practice, training Gated Fusion with randomly initialized f_{new} would make the shallow gating network quickly converge to favor the fully-trained f_{old} . To prevent this, we only train f_{new} for the first few epochs to ensure its competence before jointly training g_θ and f_{new} using $l_{GF}^*(x)$. In addition, we found that stopping gradient flow from g_θ to f_{new} can further prevent the performance decrease of the new model within Gated Fusion:

$$\alpha_{gate}(x) = g_\theta(\text{stop_grad}(\mathcal{E}_{new}(x))). \quad (4)$$

At inference time, Gated Fusion produces logits from f_{old} and f_{new} as well as the gate value α_{gate}

to make output predictions:

$$f_{GF}^*(x) = \text{Softmax}\left((1 - \alpha_{gate}) \cdot \frac{l_{old}}{T} + \alpha_{gate} \cdot l_{new}\right). \quad (5)$$

3.3 Inference with Cache

Our proposed Gated Fusion requires f_{old} to be hosted together with the new model. In reality, one could have a resource-constrained setting and request the old model to be discarded at inference. We note that in real applications, repetitive inputs are commonly seen in live traffic (Batrinca and Treleaven, 2015) and the backward compatibility of model upgrade entails that correct predictions can be preserved on the legacy instances already seen and predicted by the old model.

To simulate real scenarios, we randomly cache old model’s logits on a portion of test inputs. When getting out-of-cache instances, we use new model’s output embedding $\mathcal{E}_{new}(x)$ as key and euclidean distance as metric to search for the nearest cached instance. The cached old-model logits can then be used for Gated Fusion to make predictions without hosting f_{old} at inference.

4 Experiments Setup

4.1 Model Upgrade Scenarios

We conduct experiments on two representative model upgrade scenarios: (a) upgrade to a larger pretrained model of the same type, where we use BERT_{base} to BERT_{large}. (b) upgrade to a distinct pretrained model with the same size. We use BERT_{base} to ELECTRA_{base} (Clark et al., 2020) as this challenging model upgrade scenario for backward-compatibility, as they are pretrained under different self-supervised learning paradigms. The former uses masked language modeling (MLM) with reconstruction loss, while the latter is pretrained in generative-contrastive (adversarial) fashion with distributional divergence as the loss (Liu et al., 2021).

4.2 Datasets and Implementation

We evaluate our approach across three datasets. They represent different sentence-level classification tasks, from single-sentence to sentence-pair classification, with varying dataset sizes. We use: (a) Stanford Sentiment Treebank (SST-2), a single-sentence task to classify movie review sentiment, with 67k train and 0.9k dev set (Socher et al., 2013). (b) Microsoft Research Paraphrase Corpus (MRPC)

(Dolan and Brockett, 2005), a sentence-pair classification task for identifying paraphrases, with 3.7k train and 0.4k dev set. (c) Question Natural Language Inference (QNLI), a question-paragraph pair task to determine whether the paragraph contains the answer to the question, with 100k train and 5.5k dev set. Datasets are taken from GLUE Benchmark (Wang et al., 2018) and processed with scripts from Hugging Face².

For implementation, we use the sequence classification and pre-trained model parameters from Hugging Face Transformers³. Experiments are done in PyTorch (Paszke et al., 2019) with Tesla V100 GPUs and results are averaged over 5 random seeds. Learning rate, batch size, and train epoch are tuned during training new model alone on given tasks and then fixed for all backward-compatible solutions. In Gated Fusion, we first train f_{new} alone for first $(N - 1)$ epochs and then jointly train g_{θ} and f_{new} with Gated Fusion logits l_{GF}^* in the last epoch. Further implementation details can be found in the Appendix.

4.3 Baselines

We compare our approach with several strong baselines. (a) Train the new model directly on the target task without any adjustment, i.e. f_{new}^o . (b) The specialized distillation method proposed in Xie et al. (2021), where the KL-divergence of prediction probabilities between old and new models is applied when $p_{old}(y = y_i|x_i) > p_{new}(y = y_i|x_i)$. (c) Model ensemble via majority-voting that was shown to be very effective (Yan et al., 2021; Xie et al., 2021). Similarly, we use 5-seed new model ensemble as a strong baseline. (d) The ensemble of the old and new models probabilities, $p^*(y|x) = (1 - \alpha) \cdot p_{old}(y|x) + \alpha \cdot p_{new}(y|x)$, as well as ensemble of the old and new models logits, $l^*(y|x) = (1 - \alpha) \cdot l_{old}(y|x) + \alpha \cdot l_{new}(y|x)$, where α is searched among [0.5, 0.6, 0.7, 0.8, 0.9] to maximize backward-compatibility while achieving accuracy on par with the vanilla f_{new}^o .

5 Results and Analysis

5.1 Upgrade to a Larger Pretrained Model

Our first model upgrade scenario scales up the size of underlying pretrained language models. We experiment with BERT_{base} to BERT_{large}, where

²<https://huggingface.co/datasets/glue>

³<https://huggingface.co/docs/transformers/index>

| $\text{BERT}_{base} \rightarrow \text{BERT}_{large}$ | SST-2 | | MRPC | | QNLI | |
|--|-----------------------------|-----------------------|-----------------------------|-----------------------|-----------------------------|-----------------------|
| | \mathcal{R}_{NF} | Accuracy | \mathcal{R}_{NF} | Accuracy | \mathcal{R}_{NF} | Accuracy |
| Old Model | - | 92.00 _{0.27} | - | 85.69 _{0.90} | - | 90.74 _{0.09} |
| New Model | 2.18 _{0.21} | 93.12 _{0.29} | 4.12 _{1.04} | 87.40 _{1.02} | 2.72 _{0.13} | 92.22 _{0.16} |
| Distillation (Xie et al., 2021) | 1.97 _{0.22} | 93.33 _{0.20} | 3.53 _{0.77} | 87.70 _{1.34} | 2.31 _{0.14} | 92.60 _{0.19} |
| New Model Ensemble | 2.00 _{0.31} | 93.30 _{0.24} | 2.25 _{0.61} | 88.87 _{0.77} | 1.98 _{0.21} | 92.97 _{0.22} |
| Old-New Probs Ensemble | 1.06 _{0.27} | 93.12 _{0.38} | 1.67 _{0.78} | 87.16 _{1.12} | 1.04 _{0.26} | 92.44 _{0.23} |
| Old-New Logits Ensemble | 1.06 _{0.27} | 93.12 _{0.38} | 1.67 _{0.78} | 87.16 _{1.12} | 1.04 _{0.26} | 92.44 _{0.23} |
| Gated Fusion | 0.78 _{0.20} | 93.05 _{0.09} | 1.18 _{0.52} | 87.45 _{0.52} | 0.73 _{0.13} | 92.24 _{0.24} |

Table 1: Negative flip rate \mathcal{R}_{NF} and model accuracy (%) of competing methods to optimize backward-compatibility without performance degradation during $\text{BERT}_{base} \rightarrow \text{BERT}_{large}$ model upgrade.

| $\text{BERT}_{base} \rightarrow \text{ELECTRA}_{base}$ | SST-2 | | MRPC | | QNLI | |
|--|-----------------------------|-----------------------|-----------------------------|-----------------------|-----------------------------|-----------------------|
| | \mathcal{R}_{NF} | Accuracy | \mathcal{R}_{NF} | Accuracy | \mathcal{R}_{NF} | Accuracy |
| Old Model | - | 92.00 _{0.27} | - | 85.69 _{0.90} | - | 90.74 _{0.09} |
| New Model | 1.63 _{0.20} | 95.00 _{0.06} | 3.73 _{0.36} | 88.58 _{0.57} | 2.82 _{0.32} | 92.90 _{0.26} |
| Distillation (Xie et al., 2021) | 1.49 _{0.24} | 95.02 _{0.21} | 3.68 _{0.79} | 88.82 _{0.94} | 2.58 _{0.17} | 93.03 _{0.16} |
| New Model Ensemble | 1.12 _{0.09} | 95.39 _{0.09} | 3.24 _{0.24} | 89.02 _{0.48} | 2.26 _{0.08} | 93.49 _{0.07} |
| Old-New Probs Ensemble | 1.40 _{0.17} | 95.07 _{0.15} | 3.14 _{0.42} | 88.53 _{0.48} | 0.98 _{0.20} | 93.04 _{0.21} |
| Old-New Logits Ensemble | 0.89 _{0.17} | 94.95 _{0.13} | 3.28 _{0.43} | 88.48 _{0.51} | 0.98 _{0.20} | 93.04 _{0.21} |
| Gated Fusion | 0.71 _{0.18} | 95.02 _{0.16} | 2.40 _{0.50} | 88.68 _{0.68} | 0.81 _{0.16} | 92.98 _{0.17} |

Table 2: Negative flip rate \mathcal{R}_{NF} and model accuracy (%) of competing methods to optimize backward-compatibility without performance degradation during $\text{BERT}_{base} \rightarrow \text{ELECTRA}_{base}$ model upgrade.

the model size is tripled (110M vs 340M) and the model depth is doubled (12 vs 24 layers).

Table 1 shows the results. For f_{new}^o , we can observe that the negative flip rates \mathcal{R}_{NF} are usually much larger than the accuracy gains across tasks, which could be the reason to hinder new model adoptions in real-world applications. Besides, when dividing \mathcal{R}_{NF} over the error rate ($1 - accuracy$), we can observe that around 30% to 40% of all f_{new}^o prediction errors are in fact the *new* errors introduced during model upgrade. For improving prediction backward-compatibility, our proposed Gated Fusion outperforms other competing methods to considerably reduce \mathcal{R}_{NF} without degradation on accuracy. Note that best α values found for the two variants of old-new ensemble are both 0.5, hence producing identical results.

Compared to the vanilla new model, gated fusion obtains absolute \mathcal{R}_{NF} reductions of -1.40% on SST-2, -2.94% on MRPC, and -1.99% on QNLI. These translate to reducing the total negative flip cases by 64.2%, 71.4%, 73.2%, respectively. Compared to the strongest baseline (old-new ensemble), we obtain further absolute \mathcal{R}_{NF} reductions of -0.28% on SST-2, -0.49% on MRPC, and

-0.31% on QNLI, which translate to further reducing 12.8%, 11.9%, and 11.4% of negative flip cases. These results show the effectiveness of our method to mitigate a significant amount of regression errors during model upgrade.

5.2 Upgrade to a Different Pretrained Model

A more challenging upgrade scenario is when old and new models are pretrained under distinctive paradigms, producing two representation spaces of fairly different characteristics (Meng et al., 2021b). We experiment with BERT_{base} to ELECTRA_{base} in this scenario, where two models have the same size but are pretrained under utterly different schemes, i.e. generative versus adversarial.

Table 2 shows the results. For f_{new}^o , compared with upgrading to BERT_{large} , we observe larger accuracy gains and lower \mathcal{R}_{NF} on SST-2 and MRPC. However, on QNLI, upgrading to ELECTRA_{base} achieves a higher accuracy gain but an even a higher \mathcal{R}_{NF} . This implies that boosting accuracy and improving backward compatibility could be two related but different objectives.

For mitigation strategies, Gated Fusion achieves the lowest negative flip rates across datasets without any accuracy loss. We obtain absolute \mathcal{R}_{NF} re-

| | SST-2 | | MRPC | | QNLI | |
|------------------------------------|----------------------|-----------------------|----------------------|-----------------------|----------------------|-----------------------|
| | \mathcal{R}_{NF} | Accuracy | \mathcal{R}_{NF} | Accuracy | \mathcal{R}_{NF} | Accuracy |
| Old Model: BERT _{base} | - | 92.00 _{0.27} | - | 85.69 _{0.90} | - | 90.74 _{0.09} |
| New Model: BERT _{large} | 2.18 _{0.21} | 93.12 _{0.29} | 4.12 _{1.04} | 87.40 _{1.02} | 2.72 _{0.13} | 92.22 _{0.16} |
| Model Ensemble: 5 seeds | 2.00 _{0.31} | 93.30 _{0.24} | 2.25 _{0.61} | 88.87 _{0.77} | 1.98 _{0.21} | 92.97 _{0.22} |
| Model Ensemble: 10 seeds | 1.79 _{0.17} | 93.69 _{0.15} | 2.01 _{0.29} | 89.46 _{0.51} | 2.01 _{0.14} | 92.97 _{0.19} |
| Model Ensemble: 20 seeds | 1.79 _{0.25} | 93.62 _{0.16} | 1.76 _{0.50} | 89.56 _{0.48} | 1.82 _{0.15} | 93.13 _{0.08} |
| Gated Fusion | 0.78 _{0.20} | 93.05 _{0.09} | 1.18 _{0.52} | 87.45 _{0.52} | 0.73 _{0.13} | 92.24 _{0.24} |
| New Model: ELECTRA _{base} | 1.63 _{0.20} | 95.00 _{0.06} | 3.73 _{0.36} | 88.58 _{0.57} | 2.82 _{0.32} | 92.90 _{0.26} |
| Model Ensemble: 5 seeds | 1.12 _{0.09} | 95.39 _{0.09} | 3.24 _{0.24} | 89.02 _{0.48} | 2.26 _{0.08} | 93.49 _{0.07} |
| Model Ensemble: 10 seeds | 1.24 _{0.18} | 95.30 _{0.16} | 3.63 _{0.50} | 88.58 _{0.20} | 2.21 _{0.12} | 93.57 _{0.15} |
| Model Ensemble: 20 seeds | 1.19 _{0.16} | 95.32 _{0.17} | 3.43 _{0.51} | 88.92 _{0.48} | 2.15 _{0.17} | 93.63 _{0.11} |
| Gated Fusion | 0.71 _{0.18} | 95.02 _{0.16} | 2.40 _{0.50} | 88.68 _{0.68} | 0.81 _{0.16} | 92.98 _{0.17} |

Table 3: Negative flip rate \mathcal{R}_{NF} and model accuracy (%) when increasing number of seeds used in new model ensemble, comparing with our proposed method (Gated Fusion).

| | SST-2 | MRPC | QNLI |
|----------------------------|-----------------------|-----------------------|-----------------------|
| Old: BERT _{base} | 92.00 _{0.27} | 85.69 _{0.90} | 90.74 _{0.09} |
| to BERT _{large} | 93.12 _{0.29} | 87.40 _{1.02} | 92.22 _{0.16} |
| Gated Fusion | 93.05 _{0.09} | 87.45 _{0.52} | 92.24 _{0.24} |
| - drop old model | 93.17 _{0.61} | 87.75 _{1.14} | 92.22 _{0.44} |
| to ELECTRA _{base} | 95.00 _{0.06} | 88.58 _{0.57} | 92.90 _{0.26} |
| Gated Fusion | 95.02 _{0.16} | 88.68 _{0.68} | 92.98 _{0.17} |
| - drop old model | 95.16 _{0.09} | 88.63 _{0.94} | 93.06 _{0.13} |

Table 4: Accuracy (%) when dropping the old model within Gated Fusion at inference time.

ductions of -0.92% on SST-2, -1.33% on MRPC, and -2.01% on QNLI over the vanilla setup, reducing 56.4%, 35.7%, and 71.3% of overall negative flips, respectively. Compared with upgrading to BERT_{large}, we observe that upgrading to ELECTRA_{base} has much smaller relative negative flip reductions on SST-2 and MRPC, showing that it could be indeed harder to improve backward-compatibility when upgrading to a distinct pre-trained model. In contrast, similar relative negative flip reductions are observed on QNLI across two upgrade scenarios. This could be attributed to the abundant training data of the downstream task.

5.3 Drop Old Model at Inference Time

Our proposed method requires the old model to be hosted together with the new model. A natural question is whether we could train Gated Fusion with the old model and then discard it at inference time to host the new model only.

We first experiment with directly dropping the old model within Gated Fusion at inference time.

| | SST-2 | |
|------------------------------------|----------------------|-----------------------|
| | \mathcal{R}_{NF} | Accuracy |
| Old Model: BERT _{base} | - | 92.00 _{0.27} |
| New Model: ELECTRA _{base} | 1.63 _{0.20} | 95.00 _{0.06} |
| Gated Fusion - 50% cache | 1.26 _{0.10} | 94.86 _{0.27} |
| Gated Fusion - 75% cache | 0.99 _{0.25} | 94.91 _{0.12} |
| Gated Fusion | 0.71 _{0.18} | 95.02 _{0.16} |

Table 5: Negative flip rate \mathcal{R}_{NF} and model accuracy (%) of Gated Fusion with $X\%$ cache of old model logits at inference time.

Results in Table 4 show that dropping old model in Gated Fusion can still achieve comparable accuracy across the board, suggesting no performance degradation. Nonetheless, we observe that the negative flip rates also fall back to similar positions as training the new model in the vanilla setting.

However, in real application scenario, live inputs are often repetitively seen across time and ensuring backward-compatibility means that correct predictions on same instances can be preserved after model upgrade. We experiment with the caching method introduced in section 3.3 to store old model’s logits on random $X\%$ of test instances where Gated Fusion can later access them for inference. Results in Table 5 show that with higher percentage of cache, \mathcal{R}_{NF} is gradually reduced towards \mathcal{R}_{NF} of the original Gated Fusion, which is equivalent to 100% cache. Still, we observe a notable gap in \mathcal{R}_{NF} between the partial caching and full settings. We leave the examination of ways to achieve the upper bound in reduction in \mathcal{R}_{NF} with smaller cache to the future work.

| | (Task, Label) | Examples |
|--|------------------------|--|
| BERT _{base} → BERT _{large} | (SST-2, Positive) | [Sentence] A study in shades of gray, offering itself up in subtle plot maneuvers ... |
| | (SST-2, Negative) | [Sentence] Manages to be both repulsively sadistic and mundane. |
| | (MRPC, Not Equivalent) | [Sentence 1] Vivace was founded in 1999 and has raised over \$118 million in three rounds of venture financing. [Sentence 2] During difficult times for technology venture capital, Vivace raised over \$118 million in three rounds of venture financing. |
| | (QNLI, Entailment) | [Question] Why was there a depreciation of the industrialized nations dollars? [Sentence] Anticipating that currency values would fluctuate unpredictably for a time, the industrialized nations increased their reserves (by expanding their money supplies) in amounts far greater than before. |
| BERT _{base} → ELECTRA _{base} | (SST-2, Positive) | [Sentence] Aside from minor tinkering , this is the same movie you probably loved in 1994, except that it looks even better. |
| | (SST-2, Negative) | [Sentence] It showcases carvey’s talent for voices, but not nearly enough and not without taxing every drop of one’s patience to get to the good stuff . |
| | (MRPC, Equivalent) | [Sentence 1] Blair’s Foreign Secretary Jack Straw was to take his place on Monday to give a statement to parliament on the European Union. [Sentence 2] Blair’s office said his Foreign Secretary Jack Straw would take his place on Monday to give a statement to parliament on the EU meeting the prime minister attended last week. |
| | (QNLI, Not Entailment) | [Question] What is the main executive body of the EU? [Sentence] This means that the Commission has a monopoly on initiating the legislative procedure, although the Council is the "de facto catalyst of many legislative initiatives". |

Table 6: Examples of regression errors present when upgrading to the vanilla new model f_{new}^o but fixed by our Gated Fusion approach, i.e. predictions of $(f_{old}, f_{new}^o, f_{GF}^*)$ are *(correct, incorrect, correct)*, respectively.

5.4 Limitations of New Model Ensemble

In previous works (Yan et al., 2021; Xie et al., 2021), new model ensemble via majority voting is shown to effectively reduce negative flips and posed as a difficult-to-beat baseline. Here, we increase the number of models in ensemble to examine its limitations. Results in Table 3 show that ensemble with more models generally help to obtain lower \mathcal{R}_{NF} . However, \mathcal{R}_{NF} converges quickly as number of models increased, where a notable gap remains between new model ensemble and Gated Fusion. Moreover, the results show once more that boosting accuracy does not necessarily improve the backward compatibility in model upgrade.

In principle, two sources could cause negative flips during model upgrade (a) the stochasticity during model training, including initializations, data loading order, and optimization process (Somepalli et al., 2022). (b) the distinctions between old and new model hypotheses, including architecture and pretraining data and procedure, leading to different representation space structures and prediction behaviors in terms of decision boundaries. Without an explicit connection to f_{old} , new model ensemble can only reduce negative flips primarily caused by the first factor, while our proposed Gated Fusion directly learns to mitigate regression errors regardless of their causes.

Besides, as large-scale generative models become more and more powerful and popular (Raffel

et al., 2020; Brown et al., 2020; Su et al., 2021), it would be difficult to fine-tune them multiple times on a target task for ensemble.

5.5 Analysis of Gated Fusion

Comparing f_{new}^o with f_{GF}^* , we can calculate the *fix rate* and *new fault rate* of our Gated Fusion method. During an upgrade, if there are 20 negative flips with f_{new}^o and 16 out of them can be mitigated by f_{GF}^* , we obtain the fix rate to be $16/20 = 80\%$. Similarly, if f_{GF}^* introduces another 4 new negative flips which are not present with f_{new}^o , the new fault rate is computed to be $4/20 = 20\%$. We calculate the 5-seed average of these two rates across different classification tasks and upgrade scenarios. In BERT_{base} to BERT_{large}, the averaged fix rates by Gated Fusion are 68.4% on SST-2, 83.8% on MRPC, and 82.9% on QNLI, with new fault rates being 4.1% on SST-2, 11.3% on MRPC, and 9.7% on QNLI. In BERT_{base} to ELECTRA_{base}, Gated Fusion achieves the averaged fix rates 58.0% on SST-2, 50.8% on MRPC, and 75.6% on QNLI, with new fault rates being 2.8% on SST-2, 15.2% on MRPC, and 4.0% on QNLI. These results show that, on average, Gated Fusion is able to eliminate 69.9% of total regression errors while adding only 7.9% new ones, comparing with doing model upgrade without any treatment, i.e. f_{new}^o .

Table 6 shows a few regression error cases fixed by our proposed approach. In general, Gated Fu-

sion can mitigate negative flips happened on different classes across diverse tasks as well as on inputs with variable lengths. With closer inspections of f_{GF}^* , we found that when f_{new} produces incorrect predictions and f_{old} gives correct outputs, g_θ is capable of putting larger weights on f_{old} to ensure the backward compatibility. We also observed that the gate g_θ is more prone to over-fitting when the downstream tasks have smaller training set, e.g. MRPC, or are more difficult in nature, e.g. single-sentence task SST-2 versus sentence-pair tasks, which causes Gated Fusion to introduce more new errors, i.e. higher new fault rates.

6 Discussion

Gated Fusion requires to host both old and new models at inference time, which could raise a concern regarding the increased computational burden. However, in practice, old model’s logits of previous inference instances can be cached in storage and later leveraged in our Gated Fusion. That is, we only need to host the new model with the gate at inference time and leverage old predictions from cache. And for the out-of-cache inputs, backward-compatibility would be less of an issue since users have not observed such examples to make conclusions on the underlying regression.

For real-world applications, there could be multiple model updates and thus multiple legacy versions. We note that in this scenario, user experience would be primarily grounded on predictions of the latest legacy version, which are also saved in cache. Our Gated Fusion can hence leverage them and make new model’s predictions compatible to those from the latest legacy version.

In addition, we emphasize that the main challenge in the regression reduction research problem is to find the best trade-off between model effectiveness and backward compatibility. In this work, we show that the weighted ensemble of old-new models with a learned gate, which we call Gated Fusion, achieves a better negative flip rate than previously explored methods for regression reduction, while straight-forward ensemble approaches cannot naturally weigh on this trade-off. We don’t claim to invent the gated ensemble of old and new models but rather that our main contribution is to show that by repurposing the classic gating mechanism, the gated ensemble can become the most competitive approach to the challenging model-upgrade regression reduction problem, with no overall per-

formance degradation on two realistic model update scenarios across three different datasets.

Recently, more and more NLP products have been deployed in the industry as this field matures. We would like to stress that as better NLP models are being developed, the backward-compatible model upgrade problem naturally emerges as the new research topic strongly motivated by the real-world challenges. While backward-compatibility is currently a niche research topic, we believe that there are many thrilling future directions worth to be investigated.

7 Related Work

Yan et al. (2021) first studied the backward compatibility of predictions during model upgrade on image classification tasks. Later, Xie et al. (2021) investigated the similar topic in natural language understanding and formulated it as a constrained optimization problem. They both show that customized variants of knowledge distillation (Hinton et al., 2015), which align the predictions of old and new models on potential regression errors, are effective approaches. A model ensemble has also shown to be surprisingly effective (Yan et al., 2021; Xie et al., 2021), despite no explicit connection between old and new models. This was credited to variance reduction in model predictions, making it less prone to over-fitting and reducing regression errors indirectly. In this work, we leverage the gating mechanism to combine old and new models to further reduce model upgrade regression errors by a large margin across classification tasks.

Cai et al. (2022) analyzed and proposed backward congruent re-ranking to reduce regression in model upgrades for structured predictions tasks such as dependency parsing and conversational semantic parsing. Träuble et al. (2021) proposed an efficient probabilistic approach to locate data instances whose old predictions could be incorrect and update them with ones from the new model. Zhou et al. (2022) looked into forward compatibility, where new classes can be easily incorporated without negatively impacting existing prediction behavior. More recently, Schumann et al. (2023) inspected classification model regression during training data updates and mitigated the problem by interpolating between weights of the old and new models. On top of that, learning cross-model compatible embeddings has been extensively explored in visual search (Chen et al., 2019; Hu et al., 2019;

Wang et al., 2020). Several techniques have been proposed to optimize cross-model interoperability of embeddings, including metric space alignment (Shen et al., 2020), architecture search (Duggal et al., 2021), and aligning class centers between models Meng et al. (2021a). In this work, we focus on improving backward compatibility during model upgrade in terms of prediction behavior on classification tasks, i.e. old and new models should produce consistently correct predictions.

Reducing regression during model upgrade is also related to continual learning (Parisi et al., 2019; De Lange et al., 2019; Sun et al., 2019; Chuang et al., 2020; Sachidananda et al., 2021), incremental learning (Chaudhry et al., 2018; Shan et al., 2020) and concept drifting (Gama et al., 2014; Žliobaitė et al., 2016; Ganin et al., 2016; Zhuang et al., 2020; Lazaridou et al., 2021). In these problems, models are required to learn from and deal with continuously changing data (in terms of examples, classes or tasks), and also need to prevent the forgetting of previously learnt knowledge. This could be one potential cause of regression observed at inference. However, in backward-compatible model upgrade, a new model, usually with distinct network architecture, is trained from scratch to perform the same task and is expected to behave similarly wherever the previous model predicts correctly.

The gating mechanism is widely adopted by recurrent neural networks to effectively control information flows across networks (Hochreiter and Schmidhuber, 1997; Cho et al., 2014; Van Oord et al., 2016; Dauphin et al., 2017; Lai et al., 2019) and contextualize embeddings (Peters et al., 2018; Lai et al., 2020). It is then repurposed to act as a switch for the mixture of different prediction modes, notably to combine input word copying based on the pointer network (Vinyals et al., 2015) with the word generation from output vocabulary (Gu et al., 2016; Merity et al., 2016; See et al., 2017). Our proposed approach is inspired by these works and leverages the gating mechanism to effectively combine old and new models to improve backward compatibility during model upgrade.

8 Conclusion

Ensuring backward compatibility during model upgrade has become a critical topic in real-world NLP applications. In this work, we proposed a new approach, *Gated Fusion*, that achieves significantly better backward compatibility without

compromising accuracy performance on two challenging upgrade scenarios for NLP classification. Experiments demonstrated that our approach outperforms competing methods and achieves negative flip rate reductions by up to 73.2%. Our future research includes improving backward compatibility beyond classification to span detection, model upgrades with very large language models, and upgrades on training data or label schema. We hope that this work can inspire further research and make progress towards smoother transitions of prediction powers as NLP systems evolve.

Limitations

Our proposed method mostly works on the upgrades of underlying pretrained language models for NLP classification tasks. Potential limitations include applying our approach on distant tasks such as question answering or information retrieval, upgrade to models from different architecture families such as recurrent neural nets, and the inapplicability of our method to more recent learning formulation such as in-context learning via prompting.

Ethics Statement

Prediction backward compatibility during model upgrade is an emerging research topic to ensure positive congruency and smoother transitions from existing models towards more performant systems. With primary evaluation on accuracy and negative flips, we acknowledge that our method may also inherit social biases and other toxicity persisted in the legacy models. On the other hand, we have noted that fairness and safety have been one of principal criteria when developing system upgrades. Investigations of the inheritance of persistent toxicity and mitigation of it during backward-compatible upgrades merit interests of future research.

Acknowledgements

We would like to acknowledge AWS AI Labs for inspiring discussions, honest feedback, and full support. We are also very grateful to reviewers for judicious comments and valuable suggestions.

References

Bogdan Batrinca and Philip C Treleaven. 2015. Social media analytics: a survey of techniques, tools and platforms. *Ai & Society*, 30(1):89–116.

- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Deng Cai, Elman Mansimov, Yi-An Lai, Yixuan Su, Lei Shu, and Yi Zhang. 2022. Measuring and reducing model update regression in structured prediction for nlp. *arXiv preprint arXiv:2202.02976*.
- Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. 2018. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 532–547.
- Ken Chen, Yichao Wu, Haoyu Qin, Ding Liang, Xuebo Liu, and Junjie Yan. 2019. R3 adversarial network for cross model face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9868–9876.
- Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Yung-Sung Chuang, Shang-Yu Su, and Yun-Nung Chen. 2020. Lifelong language knowledge distillation. *arXiv preprint arXiv:2010.02123*.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*.
- Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier. 2017. Language modeling with gated convolutional networks. In *International conference on machine learning*, pages 933–941. PMLR.
- Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. 2019. Continual learning: A comparative study on how to defy forgetting in classification tasks. *arXiv preprint arXiv:1909.08383*, 2(6).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- William B. Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Rahul Duggal, Hao Zhou, Shuo Yang, Yuanjun Xiong, Wei Xia, Zhuowen Tu, and Stefano Soatto. 2021. Compatibility-aware heterogeneous visual search.
- João Gama, Indrè Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. 2014. A survey on concept drift adaptation. *ACM computing surveys (CSUR)*, 46(4):1–37.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor OK Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. *arXiv preprint arXiv:1603.06393*.
- Geoffrey E. Hinton, Oriol Vinyals, and J. Dean. 2015. Distilling the knowledge in a neural network. *ArXiv*, abs/1503.02531.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Jie Hu, Rongrong Ji, Hong Liu, Shengchuan Zhang, Cheng Deng, and Qi Tian. 2019. Towards visual feature translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3004–3013.
- Yi-An Lai, Arshit Gupta, and Yi Zhang. 2019. Goal-embedded dual hierarchical model for task-oriented dialogue generation. *arXiv preprint arXiv:1909.09220*.
- Yi-An Lai, Garima Lalwani, and Yi Zhang. 2020. Context analysis for pre-trained masked language models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3789–3804.
- Angeliki Lazaridou, Adhiguna Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Sebastian Ruder, Dani Yogatama, et al. 2021. Pitfalls of static language modelling. *arXiv preprint arXiv:2102.01951*.
- Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. 2021. Self-supervised learning: Generative or contrastive. *IEEE Transactions on Knowledge and Data Engineering*.
- Qiang Meng, Chixiang Zhang, Xiaoqiang Xu, and Feng Zhou. 2021a. Learning compatible embeddings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9939–9948.

- Yu Meng, Chenyan Xiong, Payal Bajaj, Paul Bennett, Jiawei Han, Xia Song, et al. 2021b. Coco-lm: Correcting and contrasting text sequences for language model pretraining. *Advances in Neural Information Processing Systems*, 34:23102–23114.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. **Beyond accuracy: Behavioral testing of NLP models with CheckList**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Vin Sachidananda, Jason S Kessler, and Yi-An Lai. 2021. Efficient domain adaptation of language models via adaptive tokenization. *arXiv preprint arXiv:2109.07460*.
- Raphael Schumann, Elman Mansimov, Yi-An Lai, Nikolaos Pappas, Xibin Gao, and Yi Zhang. 2023. Backward compatibility during data updates by weight interpolation. *arXiv preprint arXiv:2301.10546*.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Guangxu Shan, Shiyao Xu, Li Yang, Shengbin Jia, and Yang Xiang. 2020. Learn#: A novel incremental learning method for text classification. *Expert Systems with Applications*, 147:113198.
- Yantao Shen, Yuanjun Xiong, Wei Xia, and Stefano Soatto. 2020. Towards backward-compatible representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. **Recursive deep models for semantic compositionality over a sentiment treebank**. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Gowthami Somepalli, Liam Fowl, Arpit Bansal, Ping Yeh-Chiang, Yehuda Dar, Richard Baraniuk, Micah Goldblum, and Tom Goldstein. 2022. Can neural nets learn the same model twice? investigating reproducibility and double descent from the decision boundary perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13699–13708.
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. 2021. Multi-task pre-training for plug-and-play task-oriented dialogue system. *arXiv preprint arXiv:2109.14739*.
- Fan-Keng Sun, Cheng-Hao Ho, and Hung-Yi Lee. 2019. Lamol: Language modeling for lifelong language learning. In *International Conference on Learning Representations*.
- Frederik Träuble, Julius Von Kügelgen, Matthäus Kleindessner, Francesco Locatello, Bernhard Schölkopf, and Peter Gehler. 2021. Backward-compatible prediction updates: A probabilistic approach. *Advances in Neural Information Processing Systems*, 34:116–128.
- Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. 2016. Pixel recurrent neural networks. In *International conference on machine learning*, pages 1747–1756. PMLR.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. *Advances in neural information processing systems*, 28.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. **GLUE: A multi-task benchmark and analysis platform for natural language understanding**. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355.
- Chien-Yi Wang, Ya-Liang Chang, Shang-Ta Yang, Dong Chen, and Shang-Hong Lai. 2020. Unified representation learning for cross model compatibility. *arXiv preprint arXiv:2008.04821*.
- Yuqing Xie, Yi an Lai, Yuanjun Xiong, Yi Zhang, and Stefano Soatto. 2021. Regression bugs are in your model! measuring, reducing and analyzing regressions in nlp model updates. *arXiv preprint arXiv:2105.03048*.

Sijie Yan, Yuanjun Xiong, Kaustav Kundu, Shuo Yang, Siqu Deng, Meng Wang, Wei Xia, and Stefano Soatto. 2021. Positive-congruent training: Towards regression-free model updates. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14299–14308.

Da-Wei Zhou, Fu-Yun Wang, Han-Jia Ye, Liang Ma, Shiliang Pu, and De-Chuan Zhan. 2022. Forward compatible few-shot class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9046–9056.

Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. 2020. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76.

Indrė Žliobaitė, Mykola Pechenizkiy, and Joao Gama. 2016. An overview of concept drift applications. *Big data analysis: new algorithms for a new society*, pages 91–114.

A Details on Experiment Settings

A.1 Model Training Hyper-parameters

We search among following hyper-parameter space for the training of the old model f_{old} and the new model in the vanilla setting f_{new}^o across all datasets:

- Learning Rate: $5e^{-6}$, $1e^{-5}$, $3e^{-5}$, $5e^{-5}$
- Batch Size: 16, 32
- Training Epochs: 3, 5, 8.

The selected hyper-parameters for each model with (*learning rate, batch size, training epoch*):

- BERT_{base}:
 - On SST-2: (lr $1e^{-5}$, batch 16, epoch 5)
 - On MRPC: (lr $3e^{-5}$, batch 16, epoch 5)
 - On QNLI: (lr $3e^{-5}$, batch 32, epoch 3)
- BERT_{large}:
 - On SST-2: (lr $1e^{-5}$, batch 16, epoch 5)
 - On MRPC: (lr $3e^{-5}$, batch 16, epoch 5)
 - On QNLI: (lr $3e^{-5}$, batch 32, epoch 3)
- ELECTRA_{base}:
 - On SST-2: (lr $1e^{-5}$, batch 16, epoch 5)
 - On MRPC: (lr $5e^{-5}$, batch 32, epoch 5)
 - On QNLI: (lr $3e^{-5}$, batch 32, epoch 3)

These model training hyper-parameters for a specific model on one specific dataset is then fixed and reused for all the competing methods to improve backward compatibility during model upgrade.

A.2 Distillation Hyper-parameters

The knowledge distillation method from Xie et al. (2021) imposed an additional loss $\lambda \cdot KL(l_{old}/T, l_{new}/T)$ on potential regression instances. We experimented the best possible hyper-parameters from the following:

- λ : 0.1, 1.0, 10.0
- Temperature T : 0.5, 1.0, 2.0

A.3 Details on Gated Fusion

We initialize the gate g_θ to be a two-layer feed-forward network with the architecture (*Dropout, Linear, LayerNorm, ReLU, Dropout, Linear, Sigmoid*) and fix the hidden size to be 64 across all our experiments.

During the training of Gated Fusion, we only train the f_{new} within f_{GF}^* for the first $(N - 1)$ epochs to ensure its competence, where N is the total training epochs. In the last training epoch, we jointly train g_θ and f_{new} using the Gated Fusion logits l_{GF}^* with the secondary learning rate $lr2$. To prevent the overfitting of the gate, we also apply

drop_gate where at each training step during the last epoch, there is $D\%$ to only train f_{new} within f_{GF}^* and $(1 - D)\%$ to train with l_{GF}^* .

The hyper-parameter space of Gated Fusion is listed as follows:

- Drop Gate (%): 40, 50, 60, 80
- Temperature T on old logits: 1.0, 1.2, 1.4, 1.6
- lr2: $5e^{-7}$, $1e^{-6}$, $3e^{-6}$, $1e^{-5}$, $3e^{-5}$

We found that to achieve good results, the gap in logit magnitude of f_{old} and f_{new} needs to be bridged by the temperature when upgrading from BERT_{base} to ELECTRA_{base}, with T being 1.6 on SST-2, 1.6 on MRPC, and 1.2 on QNLI. On the other hand, $T = 1$ gives good results across three datasets when upgrading from BERT_{base} to BERT_{large}. This could result from the distinct pretraining schemes between models where MLM seem to produce larger magnitude of output logits.