



Studying Catastrophic Forgetting in Neural Ranking Models

Jesus Lovon, Laure Soulier, Karen Pinel-Sauvagnat, Lynda Tamine

► To cite this version:

Jesus Lovon, Laure Soulier, Karen Pinel-Sauvagnat, Lynda Tamine. Studying Catastrophic Forgetting in Neural Ranking Models. 43rd European Conference on Information Retrieval (ECIR 2021), Apr 2021, Lucca (on line), Italy. pp.375-390, 10.1007/978-3-030-72113-8_25 . hal-03156630

HAL Id: hal-03156630

<https://hal.science/hal-03156630>

Submitted on 2 Mar 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Studying Catastrophic Forgetting in Neural Ranking Models

Jesús Lovón-Melgarejo¹ Laure Soulier² Karen Pinel-Sauvagnat¹ Lynda Tamine¹

¹Université Paul Sabatier, IRIT, Toulouse, France

²Sorbonne Université, CNRS, LIP6, F-75005 Paris, France,

{jesus.lovon, sauvagnat, tamine}@irit.fr, laure.soulier@lip6.fr

Abstract. Several deep neural ranking models have been proposed in the recent IR literature. While their transferability to one target domain held by a dataset has been widely addressed using traditional domain adaptation strategies, the question of their cross-domain transferability is still under-studied. We study here in what extent neural ranking models catastrophically forget old knowledge acquired from previously observed domains after acquiring new knowledge, leading to performance decrease on those domains. Our experiments show that the effectiveness of neural IR ranking models is achieved at the cost of catastrophic forgetting and that a lifelong learning strategy using a cross-domain regularizer successfully mitigates the problem. Using an explanatory approach built on a regression model, we also show the effect of domain characteristics on the rise of catastrophic forgetting. We believe that the obtained results can be useful for both theoretical and practical future work in neural IR.

Keywords: Neural ranking · Catastrophic forgetting · Lifelong learning

1 Introduction

Neural ranking models have been increasingly adopted in the information retrieval (IR) and natural language processing (NLP) communities for a wide range of data and tasks [35,40]. One common underlying issue is that they learn relationships that may hold only in the domain from which the training data is sampled, and generalize poorly in unobserved domains¹ [6,40]. To enhance the transferability of neural ranking models from a source domain to a target domain, transfer learning strategies such as fine-tuning [53], multi-tasking [29], domain adaptation [41], and more recently adversarial learning [7], have been widely used². However, these strategies have by essence two critical limitations reported in the machine learning literature [6,22]. The first one, as can

¹ According to Jialin and Qiang [41], a domain consists of at most two components: a feature space over a dataset and a marginal probability distribution within a task.

² We consider the definition of transfer learning in [41]. Please note that several other definitions exist [13].

be acknowledged in the NLP and IR communities [7,29], is that they require all the domains to be available simultaneously at the learning stage (except the fine-tuning). The second limitation, under-studied in both communities, is that the model leans to *catastrophically forget* existing knowledge (source domain) when the learning is transferred to new knowledge (target domain) leading to a significant drop of performance on the source domain. These limitations are particularly thorny when considering open-domain IR tasks including, but not limited to, conversational search. In the underlying settings (e.g., QA systems and chatbots [15,25,33,43]), neural ranking models are expected to continually learn features from online information streams, sampled from either observed or unobserved domains, and to scale across different domains but without forgetting previously learned knowledge.

Catastrophic forgetting is a long-standing problem addressed in machine learning using *lifelong learning* approaches [6,42]. It has been particularly studied in neural-network based classification tasks in computer vision [22,26] and more recently in NLP [32,37,46,49]. However, while previous work showed that the level of catastrophic forgetting is significantly impacted by dataset features and network architectures, we are not aware of any existing research in IR providing clear lessons about the transferability of neural ranking models across domains, nor basically showing if state-of-the-art neural ranking models are actually faced with the catastrophic forgetting problem and how to overcome it if any. Understanding the conditions under which these models forget accumulated knowledge and whether a lifelong learning strategy is a feasible way for improving their effectiveness, would bring important lessons for both practical and theoretical work in IR. This work contributes to fill this gap identified in the neural IR literature, by studying the transferability of ranking models. We put the focus on catastrophic forgetting which is the bottleneck of lifelong learning.

The main contributions of this paper are as follows. 1) We show the occurrence of catastrophic forgetting in neural ranking models. We investigate the transfer learning of five representative neural ranking models (DRMM[14], PACRR[17], KNRM[50], V-BERT[31] and CEDR [31]) over streams of datasets from different domains³ (MS MARCO [3], TREC Microblog [45] and TREC COVID19 [47]); 2) We identify domain characteristics such as relevance density as signals of catastrophic forgetting ; 3) We show the effectiveness of constraining the objective function of the neural IR models with a forget cost term, to mitigate the catastrophic forgetting.

2 Background and Related Work

From Domain Adaptation to Lifelong Learning of Neural Networks.

Neural networks are learning systems that must commonly, on the one hand, exhibit the ability to acquire new knowledge and, on the other hand, exhibit robustness by refining knowledge while maintaining stable performance on existing knowledge. While the acquisition of new knowledge gives rise to the well-known

³ In our work, different domains refer to different datasets characterized by different data distributions w.r.t. to their source and content as defined in [41].

domain shift problem [18], maintaining model performance on existing knowledge is faced with the *catastrophic forgetting* problem. Those problems have been respectively tackled using *domain adaptation* [41] and *lifelong learning* strategies [6,42]. Domain adaptation, commonly known as a specific setting of *transfer learning* [41], includes machine learning methods (e.g., fine-tuning [49] and multi-tasking [29]) that assume that the source and the target domains from which are sampled respectively the training and testing data might have different distributions. By applying a transfer learning method, a neural model should acquire new specialized knowledge from the target domain leading to optimal performance on it.

One of the main issues behind common transfer learning approaches is catastrophic forgetting [11,12]: the newly acquired knowledge interferes with, at the worst case, overwrites, the existing knowledge leading to performance decrease on information sampled from the latter. Lifelong learning [6,42] tackles this issue by enhancing the models with the ability to continuously learn over time and accumulate knowledge from streams of information sampled across domains, either previously observed or not. The three common lifelong learning approaches are [42]: 1) regularization that constrains the objective function with a forget cost term [22,26,49]; 2) network expansion that adapts the network architecture to new tasks by adding neurons and layers [5,44]; and 3) memory models that retrain the network using instances selected from a memory drawn from different data distributions [2,32].

On the Transferability of Neural Networks in NLP and IR. Transferability of neural networks has been particularly studied in classification tasks, first in computer vision [4,54] and then only recently in NLP [19,38,39]. For instance, Mou et al. [39] investigated the transferability of neural networks in sentence classification and sentence-pair classification tasks. One of their main findings is that transferability across domains depends on the level of similarity between the considered tasks. Unlikely, previous work in IR which mainly involves ranking tasks, have only casually applied transfer learning methods (e.g., fine-tuning [53], multi-tasking [29] and adversarial learning [7]) without bringing generalizable lessons about the transferability of neural ranking models. One consensual result reported across previous research in the area, is that traditional retrieval models (e.g., learning-to-rank models [28]) that make fewer distributional assumptions, exhibit more robust cross-domain performances [7,40]. Besides, it has been shown that the ability of neural ranking models to learn new features may be achieved at the cost of poor performances on domains not observed during training [35]. Another consensual result is that although embeddings are trained using large scale corpora, they are generally sub-optimal for domain-specific ranking tasks [40].

Beyond domain adaptation, there is a recent research trend in NLP toward lifelong learning of neural networks, particularly in machine translation [46], and language understanding tasks [37,49,51]. For instance, Xu et al. [51] recently revisited the domain transferability of traditional word embeddings [34] and proposed *lifelong domain embeddings* using a meta-learning approach. The pro-

posed meta-learner is fine-tuned to identify similar contexts of the same word in both past domains and the new observed domain. Thus, its inference model is able to compute the similarity scores on pairs of feature vectors representing the same word across domains. These embeddings have been successfully applied to a topic-classification task. Unlikely, catastrophic forgetting and lifelong learning are still under-studied in IR. We believe that a thorough analysis of the transferability of neural ranking models from a lifelong learning perspective would be desirable for a wide range of emerging open-domain IR applications including but not limited to conversational search [15,33,25,43].

3 Study Design

Our study mainly addresses the following research questions:

RQ1: Does catastrophic forgetting occur in neural ranking models?

RQ2: What are the dataset characteristics that predict catastrophic forgetting?

RQ3: Is a regularization-based lifelong learning method effective to mitigate catastrophic forgetting in neural ranking models?

3.1 Experimental Set Up

Given a neural model M designed for an ad-hoc ranking task, the primary objectives of our experiments are twofold: O1) measuring the catastrophic forgetting of model M while applying a domain adaptation method \mathcal{D} , in line of RQ1 and RQ2; and O2) evaluating the effect of a lifelong learning method \mathcal{L} to mitigate catastrophic forgetting in model M , in line of RQ3. We assume that model M learns a ranking task across a stream of n domain datasets $\{D_1, \dots, D_n\}$ coming in a sequential manner one by one. At a high level, our experimental set up is:

1. Set up an ordered dataset stream **setting** $D_1 \rightarrow \dots D_{n-1} \rightarrow D_n$
2. Learn oracle models $M_i^*, i = 1 \dots n$, with parameters $\hat{\theta}^{i*}$ by training **the neural ranking model** M on training instances of **dataset** $D_i, i = 1 \dots n$.
3. Measure the retrieval performance $R_{i,i}^*$ of each oracle model M_i^* on testing instances of the same dataset D_i .
4. Launch a **domain adaptation method** \mathcal{D} w.r.t. to objective O1 (resp. a **lifelong learning method** \mathcal{L} w.r.t. to objective O2) along the considered setting as follows:
 - Initialize ($k = 1$) model M_k , with $\hat{\theta}^{1*}$, parameters of model M_1^* (trained on the dataset base D_1).
 - Repeat
 - Apply to model M_k a method \mathcal{D} w.r.t to objective O1 (resp. method \mathcal{L} w.r.t. to objective O2) to transfer knowledge to the right dataset D_{k+1} (forward transfer). The resulting model is noted M_{k+1} with parameters $\hat{\theta}^{k+1}$. Its performance on dataset D_{k+1} is noted $R_{k+1,k+1}$.
 - Measure the retrieval performance $R_{k+1,k}$ of model M_{k+1} obtained on the testing instances of left dataset D_k (backward transfer)
 - Move to the next right dataset : $k = k + 1$
 - Until the end of the dataset stream setting ($k = n$).
5. **Measure catastrophic forgetting** in model M .

This experimental pipeline, illustrated in Figure 1, follows general guidelines adopted in previous work [2,20,26]. We detail below the main underlying components highlighted in bold.

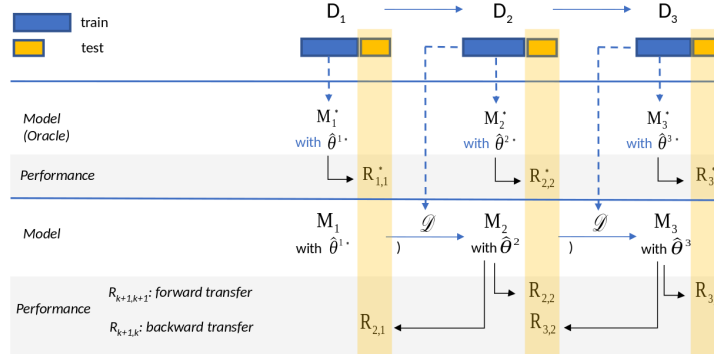


Fig. 1: Experimental pipeline using a 3-dataset stream setting for a given model M

Neural ranking models We evaluate catastrophic forgetting in five (5) state-of-the-art models selected from a list of models critically evaluated in Yang et al. [52]: 1) interaction-based models: DRMM [14] and PACRR [17] and KNRM [50]; 2) BERT-based models: Vanilla BERT [31] and CEDR-KNRM [31]. We use the OpenNIR framework [30] that provides a complete neural ad-hoc document ranking pipeline. Note that in this framework, the neural models are trained by linearly combining their own score (S_{NN}) with a BM25 score (S_{BM25}).

Datasets and settings. We use the three following datasets: 1) MS MARCO (*ms*) [3] a passage ranking dataset which includes more than 864 K question-like queries sampled from the Bing search log and a large-scale web document set including 8841823 documents; 2) TREC Microblog (*mb*) [27], a real-time ad-hoc search dataset from TREC Microblog 2013 and 2014, which contains a public Twitter sample stream between February 1 and March 31, 2013 including 124969835 tweets and 115 queries submitted at a specific point in time; 3) TREC CORD19 (*c19*) [47] an ad-hoc document search dataset which contains 50 question-like queries and a corpora of 191175 published research articles dealing with SARS-CoV-2 or COVID-19 topics. It is worth mentioning that these datasets fit with the requirement of cross-domain adaptation [41] since they have significant differences in both their content and sources. Besides, we consider four settings (See Table 1, column **Setting**) among which three 2-dataset ($n = 2$) and one 3-dataset ($n = 3$) settings. As done in previous work [2,26], these settings follow the patterns $(D_1 \rightarrow D_2)$ or $(D_1 \rightarrow D_2 \rightarrow D_3)$ where dataset orders are based on the decreasing sizes of the training sets assuming that larger datasets allow starting with well-trained networks.

Domain adaptation and lifelong learning methods. We adopt fine-tuning (training on one domain and fine-tuning on the other) as the representative domain adaptation task \mathcal{D} since it suffers from the catastrophic forgetting problem [2,22]. Additionally, we adopt the Elastic Weight Consolidation (EWC) [22] as the lifelong learning method \mathcal{L} . The EWC constrains the loss function with an additional forget cost term that we add to the objective function of each of the five neural models studied in this work. Basically speaking, EWC

constrains the neural network-based model to remember knowledge acquired on left datasets by reducing the overwriting of its most important parameters as:

$$\mathcal{L}(\hat{\theta}^k) = \mathcal{L}(\hat{\theta}^k) + \sum_{1 \leq i < k} \frac{\lambda}{2} \mathcal{F}_i(\hat{\theta}^k - \hat{\theta}^i)^2 \quad (1)$$

where $\mathcal{L}(\hat{\theta}^k)$ is the loss of the neural ranking model with parameters θ^k obtained right after learning on (D_k) , λ is the importance weight of the models parameters trained on left datasets $(D_i, i < k)$ with the current one (D_k) , \mathcal{F} is the Fisher information matrix.

Measures. Given the setting $(D_1 \rightarrow \dots \rightarrow D_n)$, we use the *remembering* measure (REM) derived from the *backward transfer measure* (BWT) proposed by Rodriguez et al. [10] as follows:

- **BWT:** measures the intrinsic effect (either positive or negative) that learning a model M on a new dataset (right in the setting) has on the model performance obtained on an old dataset (left in the setting), referred as *backward transfer*. Practically, in line with a lifelong learning perspective, this measure averages along the setting the differences between the performances of the model obtained right after learning on the left dataset and the performances of the oracle model trained and tested on the same left dataset. Thus, while positive values handle positive backward transfer, negative values handle catastrophic forgetting. Formally, the BWT measure is computed as:

$$BWT = \frac{\sum_{i=2}^n \sum_{j=1}^{i-1} (R_{i,j} - R_{j,j}^*)}{\frac{n(n-1)}{2}} \quad (2)$$

$R_{i,j}$ is the performance measure of model M_i obtained right after learning on on dataset D_j . $R_{j,j}^*$ is the performance of the oracle model M_j^* trained on dataset D_j and tested on the same dataset. To make fair comparisons between the different studied neural models, we normalize the differences in performance $(R_{i,j} - R_{j,j}^*)$ on model agnostic performances obtained using *BM25* model on each left dataset D_j . In our work, we use the standard IR performance measures MAP, NDCG@20 and P@20 to measure $R_{i,j}$ but we only report the REM values computed using the MAP measure, as they all follow the same general trends.

- **REM:** because the BWT measure assumes either positive values for positive backward transfer and negative values for catastrophic forgetting, it allows to map with a positive remembering value in the range $[0, 1]$ as follows:

$$REM = 1 - |\min(BWT, 0)| \quad (3)$$

A REM value equals to 1 means that the model does not catastrophically forget.

To better measure the intrinsic ability of the neural ranking models to remember previously acquired knowledge, we deploy in the OpenNIR framework two runs for each neural model based on the score combination ($score_G = \alpha \times S_{NN} + (1 - \alpha) \times S_{BM25}$). The first one by considering the neural model after a re-ranking setup ($0 < \alpha < 1$) leading to compute an overall *REM* measure on the ranking model. The second one by only considering the neural ranking based on the S_{NN} score by totally disregarding the BM25 scores ($\alpha = 1$). *REM_N* denotes the remembering measure computed in this second run.

3.2 Implementation details

We use the OpenNIR framework with default parameters and the pairwise hinge loss function [8]. To feed the neural ranking models, we use the GloVe pre-trained embeddings (42b tokens and 300d vectors). The datasets are split into training and testing instance sets. For MS MARCO, we use the default splits provided in the dataset. For TREC CORD19 and TREC Microblog, where no training instances are provided, we adopt the splits by proportions leading to 27/18 and 92/23 training/testing queries respectively. In practice, we pre-rank documents using the BM25 model. For each relevant document-query pair (positive pair), we randomly sample a document for the same query with a lower relevance score to build the negative pair. We re-rank the top-100 BM25 results and use $P@20$ to select the best-performing model. For each dataset, we use the optimal BM25 hyperparameters selected using grid-search. In the training phase, we consider a maximum of 100 epochs or early-stopping if no further improvement is found. Each epoch consists of 32 batches of 16 training pairs. All the models are optimized using Adam [21] with a learning rate of 0.001. BERT layers are trained at a rate of $2e-5$ following previous work [31]. For the EWC, we fixed $\lambda = 0.5$. The code is available at <https://github.com/jeslev/OpenNIR-Lifelong>.

4 Results

4.1 Empirical Analysis Of Catastrophic Forgetting in Neural Ranking Models

Within- and Across-Model Analysis (RQ1). Our objective here is to investigate whether each of the studied neural models suffer from catastrophic forgetting while it is fine-tuned over a setting ($D_1 \rightarrow D_2$ or $D_1 \rightarrow D_2 \rightarrow D_3$). To carry out a thorough analysis of each model-setting pair, we compute the following measures in addition to the REM/REMN measures: 1) the MAP@100 performance ratio ($PR = \frac{1}{(n-1)} \sum_{i=2}^n \frac{R_{i,i}}{R_{i,i}^*}$) of the model learned on the right dataset and normalized on the oracle model performance; 2) the relative improvement in MAP@100 Δ_{MAP} (resp. Δ_{MAPN}) achieved with the ranking based on the global relevance score $Score_G$ (resp. $Score_{NN}$) trained and tested on the left dataset over the performance of the BM25 ranking obtained on the same testing dataset. Table 1 reports all the metric values for each model/setting pairwise. In line with this experiment’s objective, we focus on the **”Fine-tuning”** columns.

Looking first at the PR measure reported in Table 1, we notice that it is greater than 0.96 in 100% of the settings, showing that the fine-tuned models are successful on the right dataset, and thus allow a reliable investigation of catastrophic forgetting as outlined in previous work [38]. It is worth recalling that the general evaluation framework is based on a pre-ranking (using the BM25 model) which is expected to provide positive training instances from the left dataset to the neural ranking model being fine-tuned on the right dataset. The joint comparison of the REM (resp. $REMN$) and Δ_{MAP} (resp. Δ_{MAPN}) measures lead us to highlight the following statements:

- We observe that only CEDR and VBERT models achieve positive improvements w.r.t to both the global ranking ($\Delta_{MAP} : +19.6\%, +17.4\%$ resp.)

Table 1: Per model-setting results in our fine-tuning and EWC-based lifelong learning experiments. All the measures are based on the MAP@100 metric. The improvements $\Delta_{MAP(MAPN)}$ and $\Delta_{REM(REMN)}$ are reported in percent (%).

Model Setting		Fine-tuning			EWC-based lifelong learning		
		$REM(REMN)$	$\Delta_{MAP(MAPN)}$	PR	$REM(REMN)$	$\Delta_{REM(REMN)}$	PR
DRMM	$ms \rightarrow c19$	1.000(1.000)	+2.2(-73.6)	1.008	1.000(1.000)	0(0)	1.005
	$ms \rightarrow mb$	0.962(0.943)	+2.2(-73.6)	1.021	0.971(0.974)	+0.9(+ 3.3)	1.011
	$mb \rightarrow c19$	1.000(0.965)	-1.7(-7.7)	0.993	1.000(0.662)	0(-31.4)	0.995
	$ms \rightarrow mb \rightarrow c19$	0.976(0.938)	+2(-73.6)	1.011	0.979(1.000)	+0.3(+6.6)	1.004
PACRR	$ms \rightarrow c19$	1.000(0.760)	+2.5(-30.1)	1.000	1.000(0.756)	0(-0.5)	1.000
	$ms \rightarrow mb$	1.000(1.000)	+2.5(-30.1)	0.999	1.000(1.000)	0(0)	1.014
	$mb \rightarrow c19$	1.000(0.523)	0(+10)	1.000	1.000(0.940)	0(+ 79.7)	1.002
	$ms \rightarrow mb \rightarrow c19$	1.000(0.759)	+2.5(-30)	1.000	1.000(0.874)	0(+ 15.2)	1.015
KNRM	$ms \rightarrow c19$	1.000(1.000)	-12.1(-89)	1.069	1.000(1.000)	0(0)	1.058
	$ms \rightarrow mb$	1.000(1.000)	-12.1(-89)	0.991	1.000(1.000)	0(0)	0.991
	$mb \rightarrow c19$	1.000(0.810)	-2(-13.8)	1.135	1.000(0.902)	0(+ 11.4)	1.141
	$ms \rightarrow mb \rightarrow c19$	1.000(1.000)	-12.1(-89)	1.086	1.000(0.963)	0(-3.7)	1.087
VBERT	$ms \rightarrow c19$	0.930(1.000)	-10.6(0)	1.028	1.000(1.000)	+7.5(0)	0.990
	$ms \rightarrow mb$	1.000(0.883)	-10.6(0)	1.030	1.000(1.000)	0(+ 13.3)	0.992
	$mb \rightarrow c19$	0.913(1.000)	+17.4(+25.8)	0.963	1.000(1.000)	+9.5(0)	1.010
	$ms \rightarrow mb \rightarrow c19$	0.989(0.922)	-10.6(0)	1.011	1.000(1.000)	+1.1(+8.5)	0.987
CEDR	$ms \rightarrow c19$	0.826(1.000)	+2.6(+14.2)	1.013	1.000(1.000)	+21.1(0)	1.008
	$ms \rightarrow mb$	0.510(0.920)	+2.6(+14.2)	1.003	1.000(1.000)	+96.1(+8.7)	0.976
	$mb \rightarrow c19$	0.940(1.000)	+19.6(+29.2)	1.011	1.000(1.000)	+6.4(0)	0.984
	$ms \rightarrow mb \rightarrow c19$	0.771(0.946)	+2.6(+14.2)	0.996	0.891(1.000)	+15.6(+5.7)	0.961

and the neural ranking (Δ_{MAP} : +29.2%, +25.8% resp.), particularly under the setting where mb is the left dataset ($mb \rightarrow c19$). Both models are able to bring effectiveness gains additively to those brought by the exact-based matching signals in BM25. These effectiveness gains can be viewed as new knowledge in the form of semantic matching signals which are successfully transferred to the left dataset ($c19$) while maintaining stable performances on the left dataset (mb) ($REM=0.940$ and 0.913 for resp. CEDR and VBERT). This result is consistent with previous work suggesting that the regularization used in transformer-based models has an effect of alleviating catastrophic forgetting [23].

- We notice that the CEDR model achieves positive improvements w.r.t to the neural ranking score (Δ_{MAPN} : +14.2%) in all the settings (3/4) where ms is the left dataset while very low improvements are achieved w.r.t. to the global score (Δ_{MAP} : +2.6%). We make the same observation for the PACRR model but only for 1/4 model-setting pair (Δ_{MAPN} : +10% vs. Δ_{MAPN} : 0%) with mb as the left dataset. Under these settings, we can see that even the exact-matching signals brought by the BM25 model are very moderate (leading to a very few positive training instances), the CEDR and, to a lower extent, the PACRR models, are able to inherently bring significant new knowledge in terms of semantic matching signals at however the cost of significant forget on the global ranking for CEDR (REM is the range [0.510; 0.826]) and on the neural ranking for PACRR ($REM=0.523$).

- All the models (DRMM, PACRR, KNRM and VBERT (for 3/4 settings) that do not significantly beat the BM25 baseline either by using the global score (Δ_{MAP} in the range $[-12.1\%; +2.2\%]$) nor by using the neural score (Δ_{MAPN} in the range $[-89\%; +0\%]$), achieve near upper bound of remembering (both REM and REMN are in the range $[0.94; 1]$). Paradoxically, this result does not allow us to argue about the ability of these models to retain old knowledge. Indeed, the lack or even the low improvements over both the exact matching (using the BM25 model) and the semantic-matching (using the neural model) indicate that a moderate amount of new knowledge or even no knowledge about effective relevance ranking has been acquired from the left dataset. Thus, the ranking performance of the fine-tuned model on the left dataset only depends on the level of mismatch between the data available in the right dataset for training and the test data in the left dataset. We can interestingly see that upper bound remembering performance ($REM = 1$) is particularly achieved when *ms* is the left dataset (settings $ms \rightarrow c19$, $ms \rightarrow mb$, $ms \rightarrow mb \rightarrow c19$). This could be explained by the fact that the relevance matching signals learned by the neural model in in-domain knowledge do not degrade its performances on general-domain knowledge.

Assuming a well-established practice in neural IR which consists in linearly interpolating the neural scores with the exact-based matching scores (e.g., BM25 scores), these observations give rise to three main findings: 1) the more a neural ranking model is inherently effective in learning additional semantic matching signals, the more likely it catastrophically forgets. In other terms, intrinsic effectiveness of neural ranking models comes at the cost of forget; 2) transformer-based language models such as CEDR and VBERT exhibit a good balance between effectiveness and forget as reported in previous work in NLP [38]; 3) given the variation observed in REM and REMN, there is no clear trend about which ranking (BM25-based ranking vs. neural ranking) impacts more importantly the level of overall catastrophic forgetting of the neural models

Across Dataset Analysis (RQ2). Our objective here is to identify catastrophic forgetting signals from the perspective of the left dataset. As argued in [1], understanding the relationships between data characteristics and catastrophic forgetting allows to anticipating the choice of datasets in lifelong learning settings regardless of the neural ranking models. We perform a regression model to explain the REM metric (dependent variable) using nine datasets characteristics (independent variables). The latter are presented in Table 2 and include dataset-based measures inspired from [1,9] and effectiveness-based measures using the BM25 model. To artificially-generate datasets with varying data characteristics, we follow the procedure detailed in [1]: we sample queries within each left dataset in the settings presented in Table 1 (15 for *mb* and 50 for *ms*) to create sub-datasets composed of those selected queries and the 100 top corresponding documents retrieved by the BM25 model.

Then, we replace in each setting the left dataset by the corresponding sub-dataset. We estimate for each model-setting pair the REM value as well as the characteristic values of the left sub-dataset. We repeat this procedure 300

Table 2: Linear regression explaining catastrophic forgetting (REM metric) at the left dataset level. Significance: *** : $p \leq 0.001$; ** : $0.001 < p \leq 0.01$; * : $0.01 < p \leq 0.5$

Characteristic Description			Coeff
R^2			0.544
Independent variables	Dataset	Constant	0.7014***
		RS Retrieval space size: $\log_{10}(D \times Q)$	-0.1883
		RD Relevance density: $\log_{10} \frac{Q^{rels}}{D \times Q}$	-0.3997*
		SD Score relevance divergence: $KL(RSV_{D+}, RSV_{D-})$	0.0009
		Vocab Size of the vocabulary	-0.0932*
		DL Average length of documents	-0.0349
		QL Average length of queries	0.1803*
		QD Average query difficulty: $avg_q(\frac{1}{ q } \sum_{w \in q} idf_w)$	0.0044
		MAP Effectiveness of the BM25: MAP	-0.0220*
		std-AP Variation of BM25 effectiveness (AP metric): $\sigma_q(AP_q)$	0.0543*
Residual Variables	$Dataset_i$	MSmarco	0.18038
		Microblog	0.5211**
	M_j	DRMM	0.1798***
		PACRR	0.1965***
		KNRM	0.1924***
		VBERT	0.1313***
		CEDR	0.0014

times to obtain 300 new settings per model, based on the 300 sampled sub-datasets. This leads to 300 (sub-setting-model) pairs with a variation for both the dependent and the independent variables. Finally, we build the following explanatory regression model, referring to the "across dataset analysis" in [1]:

$$REM_{ij} = \sum_k C_k f_{ik} + Dataset_i + M_j + \epsilon_i \quad (4)$$

where i denotes the i^{th} sub-setting and j refers to the neural ranking model M_j . C_k and f_{ik} denote respectively the weight and the value of the k^{th} characteristic of the left dataset in the i^{th} sub-setting. Please note, that dataset feature values are independent of the model M_j . $Dataset_i$ and M_j are the residual variables of resp. the left dataset and the model. The characteristic values f_{ik} are centered before the regression as suggested in Adomavicius and Zhang [1].

Table 2 presents the result of the regression model. From R^2 and *Constant*, we can see that our regression model can explain 54.4% of the variation of the REM metric, highlighting an overall good performance in explaining the remembering metric with a good level of prediction (0.7014). From the independent variables, we can infer that the difficulty of the dataset positively impacts the remembering (namely, decreasing the catastrophic forgetting). More precisely, lower the relevance density (RD), the BM25 effectiveness (MAP) and higher the variation in terms of BM25 performances over queries (std-AP) are, the higher the REM metric is. This suggests that relevance-matching difficulty provides positive feedback signals to the neural model to face diverse learning instances, and therefore to better generalize over different application domains. This is however true to the constraint that the vocabulary of the dataset (*Vocab*) is

not too large so as to boost neural ranking performance as outlined in [16,36]. Looking at the residual variables ($Dataset_j$ and M_j), we can corroborate the observations made at a first glance in RQ1 regarding the model families clearly opposing (DRMM-PACRR-KNRM-VBERT) and CEDR since the former statistically exhibit higher REM metrics values than CEDR.

4.2 Mitigating Catastrophic Forgetting (RQ3)

From RQ1, we observed that some models are more prone to the catastrophic forgetting problem than others. Our objective here is to examine whether an EWC-based lifelong strategy can mitigate the problem. It is worth mentioning that this objective has been targeted in previous research in computer vision but without establishing a consensus [24,46,48]. While some studies reveal that EWC outperforms domain adaptation strategies in their settings [24,46], others found that it is less effective [48]. To achieve the experiment’s objective, we particularly report the following measures in addition to the REM/REM_N measures: 1) $\Delta_{REM(REM_N)}$ that reveals the improvement (positive or negative) of the REM/REM_N measures achieved using an EWC-based lifelong learning strategy over the REM/REM_N measure achieved using a fine-tuning strategy; 2) the PR measure introduced in Section 4.1. Unlikely, our aim through this measure here, is to highlight the performance stability of the learned model on the right dataset while avoiding catastrophic forgetting on the left dataset.

We turn now our attention to the **”EWC-based lifelong learning”** columns in Table 1. Our experiment results show that among the 9 (resp. 11) settings that exhibit catastrophic forgetting in the combined model (resp. neural model), EWC strategy allows to improve 9/9 i.e., 100% (resp. 9/11 i.e., 88%) of them in the range $[+0.3\%, +96.1\%]$ (resp. $[+3.3\%, +79.7\%]$). Interestingly, this improvement in performance on the left dataset does not come at the cost of a significant decrease in performance on the right dataset since 100% of the models achieve a PR ratio greater than 0.96. Given, in the one hand, the high variability of the settings derived from the samples, and in the other hand, the very low number of settings (10% i.e., 2/20) where a performance decrease is observed in the left dataset, we could argue that the EWC-based lifelong learning is not inherently impacted by dataset order leading to a general effectiveness trend over the models. We emphasize this general trend by particularly looking at the CEDR model which we recall (See Section 4.1, RQ1), clearly exhibits the catastrophic forgetting problem. As can be seen from Table 1, model performances on the left datasets are significantly improved ($+6.4\% \leq \Delta_{REM} \leq +96.1\%$; $0\% \leq \Delta_{REM_N} \leq +8.7\%$) while keeping model performances on the right dataset stable ($0.961 \leq PR \leq 1.008$). This property is referred to as the stability-plasticity dilemma [42].

To get a better overview of the effect of the EWC strategy, we compare in Figure 2 the behavior of the CEDR and KNRM models which exhibit respectively low level ($REM = 0.510$) and high level of remembering ($REM = 1$) particularly in the setting $ms \rightarrow mb$. The loss curves in Figure 2(a) highlight a peak after the 20th epoch for both CEDR and KNRM. This peak denotes the beginning of the fine-tuning on the mb dataset. After this peak, we can observe

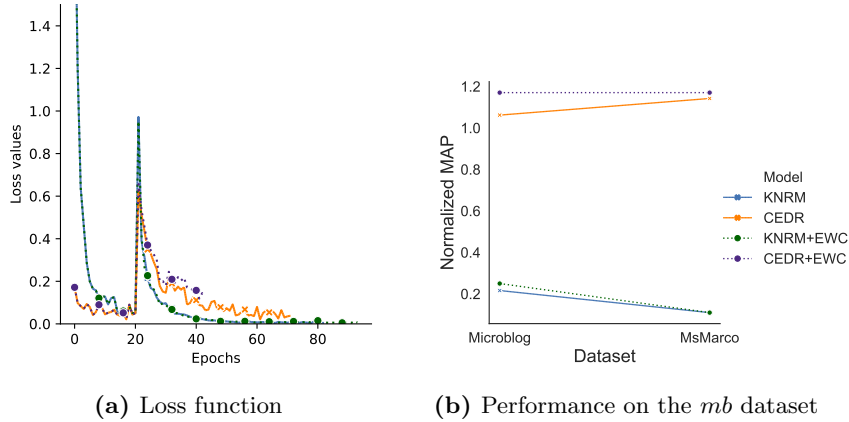


Fig. 2: Impact of the EWC strategy on loss and performance for the $ms \rightarrow mb$ setting.

that the curve representing the EWC-based CEDR loss (in purple) is slightly above the CEDR loss (in orange), while both curves for the KNRM model (green and blue resp. for with and without EWC) are overlaid. Combined with the statements outlined in RQ1 concerning the ability of the CEDR model to accumulate knowledge, this suggests that EWC is able to discriminate models prone to catastrophic forgetting and, when necessary, to relax the constraint of good ranking prediction on the dataset used for the fine-tuning to avoid over-fitting. This small degradation of knowledge acquisition during the fine-tuning on the ms dataset is carried out at the benefit of the previous knowledge retention to maintain retrieval performance on the mb dataset (Figure 2(b)). Thus, we can infer that the EWC strategy applied on neural ranking models plays fully its role to mitigate the trade-off between stability and plasticity.

5 Conclusion

We investigated the problem of catastrophic forgetting in neural-network based ranking models. We carried out experiments using 5 SOTA models and 3 datasets showing that neural ranking effectiveness comes at the cost of forget and that transformer-based models allow a good balance between effectiveness and remembering. We also show that the EWC-based strategy mitigates the catastrophic forgetting problem while ensuring a good trade-off between transferability and plasticity. Besides, datasets providing weak and varying relevance signals are likely to be remembered. While previous work in the IR community mainly criticized neural models regarding effectiveness [35,40,52], we provide complementary insights on the relationship between effectiveness and transferability in a lifelong learning setting that involves cross-domain adaptation. We believe that our study, even under limited setups, provides fair and generalizable results that could serve future research and system-design in neural IR.

6 Acknowledgement

We would like to thank projects ANR COST (ANR-18-CE23-0016) and ANR JCJC SESAMS (ANR-18-CE23-0001) for supporting this work.

References

1. Adomavicius, G., Zhang, J.: Impact of data characteristics on recommender systems performance. *ACM Trans. Manage. Inf. Syst.* **3**(1) (2012)
2. Asghar, N., Mou, L., Selby, K.A., Pantasdo, K.D., Poupart, P., Jiang, X.: Progressive memory banks for incremental domain adaptation. In: *ICLR*. vol. abs/1811.00239 (2020)
3. Bajaj, P., Campos, D., Craswell, N., Deng, L., Gao, J., Liu, X., Majumder, R., McNamara, A., Mitra, B., Nguyen, T., et al.: Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268* (2016)
4. Bengio, Y.: Deep learning of representations for unsupervised and transfer learning. In: *UTLW'11*. p. 17–37 (2011)
5. Cai, H., Chen, H., Zhang, C., Song, Y., Zhao, X., Yin, D.: Adaptive parameterization for neural dialogue generation. In: *EMNLP-IJCNLP*. pp. 1793–1802 (Nov 2019)
6. Chen, Z., Liu, B.: *Lifelong Machine Learning, Second Edition*. Synthesis Lectures on Artificial Intelligence and Machine Learning (2018)
7. Cohen, D., Mitra, B., Hofmann, K., Croft, W.B.: Cross domain regularization for neural ranking models using adversarial learning. In: *ACM SIGIR* (May 2018)
8. Dehghani, M., Zamani, H., Severyn, A., Kamps, J., Croft, W.B.: Neural ranking models with weak supervision. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. pp. 65–74 (2017)
9. Deldjoo, Y., Di Noia, T., Di Sciascio, E., Merra, F.A.: How dataset characteristics affect the robustness of collaborative recommendation models. In: *ACM SIGIR*. p. 951–960 (2020)
10. Díaz-Rodríguez, N., Lomonaco, V., Filliat, D., Maltoni, D.: Don't forget, there is more than forgetting: new metrics for Continual Learning. *ArXiv e-prints* (Oct 2018)
11. French, R.M.: Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences* **3**(4), 128 – 135 (1999)
12. Goodfellow, I.J., Mirza, M., Da, X., Courville, A.C., Bengio, Y.: An empirical investigation of catastrophic forgetting in gradient-based neural networks. *CoRR* **abs/1312.6211** (2014)
13. Gulrajani, I., Lopez-Paz, D.: In search of lost domain generalization. In: *International Conference on Learning Representations* (2021), <https://openreview.net/forum?id=1QdXeXDwItI>
14. Guo, J., Fan, Y., Ai, Q., Croft, W.B.: A deep relevance matching model for ad-hoc retrieval. In: *CIKM '16*. p. 55–64. Association for Computing Machinery (2016)
15. Hancock, B., Bordes, A., Mazare, P.E., Weston, J.: Learning from dialogue after deployment: Feed yourself, chatbot! In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. pp. 3667–3684 (Jul 2019)
16. Hofstätter, S., Rekabsaz, N., Eickhoff, C., Hanbury, A.: On the effect of low-frequency terms on neural-ir models. In: *SIGIR*. p. 1137–1140 (2019)
17. Hui, K., Yates, A., Berberich, K., de Melo, G.: Pacrr: A position-aware neural ir model for relevance matching. *arXiv preprint arXiv:1704.03940* (2017)
18. J, Q.C., M, S., A, S., D., L.N.: *Dataset Shift in Machine Learning* (2009)
19. Jha, R., Lovering, C., Pavlick, E.: When does data augmentation help generalization in nlp? (2020)

20. Kemker, R., McClure, M., Abitino, A., Hayes, T.L., Kanan, C.: Measuring catastrophic forgetting in neural networks. In: AAAI-18. pp. 3390–3398 (2018)
21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
22. Kirkpatrick, J., Pascanu, R., Rabinowitz, N.C., 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. CoRR **abs/1612.00796** (2016)
23. Lee, C., Cho, K., Kang, W.: Mixout: Effective regularization to finetune large-scale pretrained language models. In: 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020. OpenReview.net (2020), <https://openreview.net/forum?id=HkgaETNtDB>
24. Lee, S.W., Kim, J.H., Jun, J., Ha, J.W., Zhang, B.T.: Overcoming catastrophic forgetting by incremental moment matching. p. 4655–4665. NIPS’17, Curran Associates Inc., Red Hook, NY, USA (2017)
25. Li, J., Miller, A.H., Chopra, S., Ranzato, M., Weston, J.: Learning through dialogue interactions by asking questions. In: ICLR 2017 (2017)
26. Li, Z., Hoiem, D.: Learning without forgetting. IEEE Transactions on Pattern Analysis and Machine Intelligence (12), 2935–2947 (2018)
27. Lin, J., Efron, M.: Overview of the trec-2013 microblog track. In: Text REtrieval Conference (TREC), Gaithersburg, Maryland, USA (2013)
28. Liu, T.Y.: Learning to rank for information retrieval. Found. Trends Inf. Retr. **3**(3), 225–331 (Mar 2009)
29. Liu, X., Gao, J., He, X., Deng, L., Duh, K., Wang, Y.: Representation learning using multi-task deep neural networks for semantic classification and information retrieval. In: NAACL HLT 2015. pp. 912–921 (2015)
30. MacAvaney, S.: OpenNIR: A complete neural ad-hoc ranking pipeline. In: WSDM 2020 (2020)
31. MacAvaney, S., Yates, A., Cohan, A., Goharian, N.: Cedr: Contextualized embeddings for document ranking. In: ACM SIGIR. pp. 1101–1104 (2019)
32. de Masson d’Autume, C., Ruder, S., Kong, L., Yogatama, D.: Episodic memory in lifelong language learning. CoRR **abs/1906.01076** (2019)
33. Mazumder, S., Ma, N., Liu, B.: Towards a continuous knowledge learning engine for chatbots. CoRR **abs/1802.06024** (2018)
34. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. In: ICLR 2013 (2013)
35. Mitra, B., Craswell, N.: An introduction to neural information retrieval. Foundations and Trends in Information Retrieval **13**(1), 1–126 (December 2018)
36. Mitra, B., Craswell, N.: An updated duet model for passage re-ranking. CoRR **abs/1903.07666** (2019), <http://arxiv.org/abs/1903.07666>
37. Mosbach, M., Andriushchenko, M., Klakow, D.: On the Stability of Fine-tuning BERT: Misconceptions, Explanations, and Strong Baselines. arXiv e-prints p. arXiv:2006.04884 (Jun 2020)
38. Mosbach, M., Andriushchenko, M., Klakow, D.: On the stability of fine-tuning bert: Misconceptions, explanations, and strong baselines (2020)
39. Mou, L., Meng, Z., Yan, R., Li, G., Xu, Y., Zhang, L., Jin, Z.: How transferable are neural networks in NLP applications? In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing. pp. 479–489 (Nov 2016)
40. Onal, K.D., Zhang, Y., Altingövd, I.S., Rahman, M.M., Senkul, P., Braylan, A., Dang, B., Chang, H.L., Kim, H., McNamara, Q., Angert, A., Banner, E., Khetan,

- V., McDonnell, T., Nguyen, A.H., Xu, D., Wallace, B.C., de Rijke, M., Lease, M.: Neural information retrieval: at the end of the early years. *Information Retrieval Journal* **21**, 111–182 (2017)
41. Pan, S.J., Yang, Q.: A survey on transfer learning. *IEEE Trans. on Knowl. and Data Eng.* **22**(10), 1345–1359 (Oct 2010)
 42. Parisi, G.I., Kemker, R., Part, J.L., Kanan, C., Wermter, S.: Continual lifelong learning with neural networks: A review. *Neural Networks* **113**, 54–71 (2019)
 43. Roller, S., Boureau, Y.L., Weston, J., Bordes, A., Dinan, E., Fan, A., Gunning, D., Ju, D., Li, M., Poff, S., et al.: Open-domain conversational agents: Current progress, open problems, and future directions. *arXiv preprint arXiv:2006.12442* (2020)
 44. Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R.: Progressive neural networks. *CoRR* (2016)
 45. Soboroff, I., Ounis, I., Macdonald, C., Lin, J.J.: Overview of the TREC-2012 microblog track. In: *Proceedings of The Twenty-First Text REtrieval Conference, TREC 2012*. NIST Special Publication, vol. 500-298 (2012)
 46. Thompson, B., Gwinnup, J., Khayrallah, H., Duh, K., Koehn, P.: Overcoming catastrophic forgetting during domain adaptation of neural machine translation. In: *NAACL*. pp. 2062–2068 (2019)
 47. Wang, L.L., Lo, K., Chandrasekhar, Y., Reas, R., Yang, J., Eide, D., Funk, K., Kinney, R., Liu, Z., Merrill, W., et al.: Cord-19: The covid-19 open research dataset. *ArXiv* (2020)
 48. Wen, S., Itti, L.: Overcoming catastrophic forgetting problem by weight consolidation and long-term memory. *CoRR* **abs/1805.07441** (2018), <http://arxiv.org/abs/1805.07441>
 49. Wiese, G., Weissenborn, D., Neves, M.: Neural domain adaptation for biomedical question answering. In: *CoNLL 2017*. pp. 281–289 (Aug 2017)
 50. Xiong, C., Dai, Z., Callan, J., Liu, Z., Power, R.: End-to-end neural ad-hoc ranking with kernel pooling. In: *ACM SIGIR*. pp. 55–64 (2017)
 51. Xu, H., Liu, B., Shu, L., Yu, P.S.: Lifelong domain word embedding via meta-learning. In: *IJCAI-18*. pp. 4510–4516 (7 2018)
 52. Yang, W., Lu, K., Yang, P., Lin, J.: Critically examining the” neural hype” weak baselines and the additivity of effectiveness gains from neural ranking models. In: *ACM SIGIR*. pp. 1129–1132 (2019)
 53. Yang, W., Xie, Y., Tan, L., Xiong, K., Li, M., Lin, J.: Data augmentation for BERT fine-tuning in open-domain question answering. *CoRR* **abs/1904.06652** (2019)
 54. Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How transferable are features in deep neural networks? In: *NIPS’14*. p. 3320–3328 (2014)