

Psycholinguistics meets Continual Learning: Measuring Catastrophic Forgetting in Visual Question Answering

Claudio Greco¹

claudio.greco@unitn.it

Barbara Plank²

bplank@itu.dk

Raquel Fernández³

raquel.fernandez@uva.nl

Raffaella Bernardi^{1,4}

raffaella.bernardi@unitn.it

¹CIMEC and ⁴DISI
University of Trento

²Dept. of Computer Science
IT University of Copenhagen

³ILLC
University of Amsterdam

Abstract

We study the issue of catastrophic forgetting in the context of neural multimodal approaches to Visual Question Answering (VQA). Motivated by evidence from psycholinguistics, we devise a set of linguistically-informed VQA tasks, which differ by the types of questions involved (Wh-questions and polar questions). We test what impact task difficulty has on continual learning, and whether the order in which a child acquires question types facilitates computational models. Our results show that dramatic forgetting is at play and that task difficulty and order matter. Two well-known current continual learning methods mitigate the problem only to a limiting degree.

1 Introduction

Supervised machine learning models are incapable of continuously learning new tasks, as they forget how to perform the previously learned ones. This problem, called *catastrophic forgetting*, is prominent in artificial neural networks (McClelland et al., 1995). *Continual Learning* (CL) addresses this problem by trying to equip models with the capability to continuously learn new tasks over time (Ring, 1997). Catastrophic forgetting and CL have received considerable attention in computer vision (e.g., Zenke et al., 2017; Kirkpatrick et al., 2017), but far less attention within Natural Language Processing (NLP).

We investigate catastrophic forgetting in the context of multimodal models for Visual Question Answering (Antol et al., 2015) motivated by evidence from psycholinguistics. VQA is the task of answering natural language questions about an image. Evidence from child language acquisition indicates that children learn Wh-questions before polar (Yes/No) questions (Moradlou and Ginzburg, 2016; Moradlou et al., 2018). Motivated by this finding, we design a set of

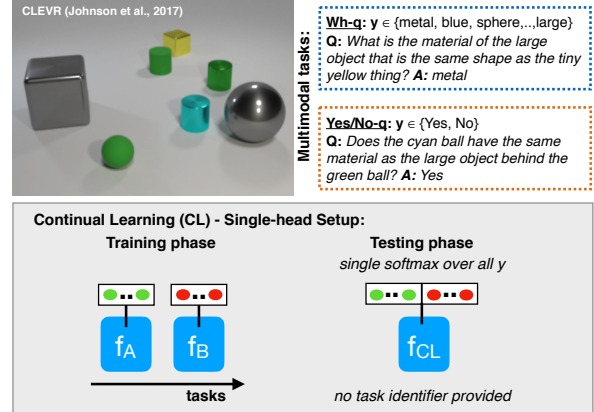


Figure 1: Overview of our linguistically-informed CL setup for VQA.

linguistically-informed experiments: i) to investigate whether the order in which children acquire question types facilitates continual learning for computational models and, accordingly, the impact of *task order* on catastrophic forgetting; ii) to measure how far two well-known CL approaches help to overcome the problem (Robins, 1995; Kirkpatrick et al., 2017)¹.

Contributions: Our study contributes to the literature on CL in NLP. In particular: i) we introduce a CL setup based on linguistically-informed task pairs which differ with respect to question types and level of difficulty; ii) we show the importance of task order, an often overlooked aspect, and observe asymmetric synergetic effects; iii) our results show that our VQA model suffers from extreme forgetting; rehearsal gives better results than a regularization-based method. Our error analysis shows that the latter approach encounters problems even in discerning Task A after having been trained on Task B. Our study opens the door to deeper investigations of CL on linguistic

¹Code and data are available at the link <http://continual-vista.github.io/>.

skills with different levels of difficulty based of psycholinguistics findings.

2 Task Setup

As a first step towards understanding the connection between linguistic skills and the impact on CL, we design a set of experiments within VQA where tasks differ with respect to the *type of question* and the *level of difficulty* according to the psycholinguistics literature. The overall setup is illustrated in Figure 1 and described next.

Dataset CLEVR (Johnson et al., 2017a) allows to study the ability of VQA agents. It requires compositional language and basic spatial reasoning skills. Every question in CLEVR is derived by a Functional Program (FP) from a scene graph of the associated image. The scene graph defines the objects and attributes in the image. The FP contains functions corresponding to skills, e.g., querying object attributes or comparing values (see Fig. 1, upper). Questions are categorized by their type. CLEVR consists of five question types whose answer labels range over 15 attributes, 10 numbers, and “yes”/“no” (in total 27 labels).

Multimodal Tasks We select the CLEVR sub-tasks ‘query_attribute’ and ‘equal_attribute’ with attributes *color*, *shape*, *material*, and *size*. The two types of questions differ by answer type $y \in \mathcal{Y}$:

- **Wh-questions** (Wh-q): Questions about the *attribute* of an object, e.g., “What is the material of the large object...?”, where $y \in \{blue, cube, small, \dots, metal\}$ spans over $|color| = 8$, $|shape| = 3$, $|size| = 2$ and $|material| = 2$ (in total $|\mathcal{Y}| = 15$).
- **Yes/No questions** (Y/N-q): Questions that *compare* objects with respect to an attribute, e.g., “Does the cyan ball have the same material as ...?”, with $y \in \{yes, no\}$ (in total $|\mathcal{Y}| = 2$).

Task Order We learn Task A followed by Task B (TASKA \rightarrow TASKB), but experiment with *both* directions, i.e., by first assigning Wh-q to Task A and Y/N-q to Task B, and vice versa. We expect that the inherent difficulty of a task and the order in which tasks are learned have an impact on CL.

Single-head Evaluation CL methods can be tested in two ways. We opt for a *single-head* evaluation setup (see Fig. 1, lower) with an output space over labels for all tasks (here: all CLEVR

labels). In contrast, in a *multi-head* setup predictions are restricted to task labels, as the task identifier is provided. Single-head is more difficult yet more realistic (Chaudhry et al., 2018).

3 Models and Experiments

VQA Model We take the model proposed by Yang et al. (2016) as a starting point, using the code released by Johnson et al. (2017b) (LSTM+CNN+SA). Questions are encoded with a recurrent neural network with Long Short-Term Memory (LSTM) units. Images are encoded with a ResNet-101 Convolutional Neural Network (CNN) pre-trained on ImageNet (He et al., 2016). The two representations are combined using Spatial Attention (SA) (Yang et al., 2016) to focus on the most salient objects and properties in the image and text. The final answer distribution is predicted with a Multilayer Perceptron (MLP).

Baselines In order to measure catastrophic forgetting, we first consider per-task baselines: A random baseline (i.e., random stratified sample of the label distribution per task) and the results of a model trained independently on each task (i.e., over task-specific \mathcal{Y}). For CL, we report again a random baseline (this time a random stratified sample drawing predictions according to the answer distribution of both tasks), and we consider the *Naive* and *Cumulative* baselines proposed by Maltoni and Lomonaco (2018). The *Naive* model is fine-tuned across tasks: It is first trained on Task A and then on Task B starting from the previously learned parameters. The *Cumulative* model is trained from scratch on the training sets of both Task A and Task B. This is a kind of upper bound, or performance that a CL model should achieve.

Continual Learning Models In CL there are two broad families of methods: Those that assume memory and access to explicit previous knowledge (instances), and those that have only access to compressed knowledge, such as previously learned parameters. These two families correspond to rehearsal and regularization, respectively. A widely-used regularization-based approach is *Elastic Weight Consolidation* (EWC, Kirkpatrick et al., 2017). A regularization term, parametrized by λ , is added to the loss function aiming the model to converge to parameters where it has a low error for both tasks. In the *Rehearsal* approach (Robins, 1995), the model is first trained on

Task A, then the parameters are fine-tuned through batches taken from a dataset containing a small number of examples of Task A and the training set of Task B. The selection of training examples of Task A is done through uniform sampling.

Data and Training Details Since CLEVR has no published ground-truth answers for the test set, we split the original validation set into a validation and a test set. To avoid performance impact due to different training data sizes, we downsample the training sets to the same size (Y/N-q data size), resulting in 125,654 training instances per task. The validation and test sets contain, respectively, 26,960 and 26,774 data points for Wh-q and 13,417 and 13,681 data points for Y/N-q.

For the baselines, we select the model which reaches maximum accuracy on the validation set of each task. For CL, we choose the model with the highest CL score computed according to the validation set of each task pair. Details on hyperparameters and evaluation metrics are provided in the supplementary material (SM).

4 Results and Analysis

The main results are provided in Table 1. There are several take-aways.

Task Difficulty The results of the per-task models (cf. first two rows in Table 1) show that there is a large performance gap between the two tasks. Wh-q is easier (.81) than Y/N-q (.52), regardless of the fact that a priori the latter should be easier (as shown by the respective task-specific random baselines). The Y/N-q task-specific model performs only slightly above chance (.52, in line with what Johnson et al. (2017a) report for ‘equal_attribute’ questions). This shows that despite the limited output space of the Y/N-q task, such type of questions in CLEVR are complex and require reasoning skills (Johnson et al., 2017a).

Catastrophic Forgetting We observe that extreme forgetting is at play. *Naive* forgets the previously learned skill completely: When tested on Task A after having been fine-tuned on Task B, it achieves 0.0 accuracy on the first task for *both directions* (I and II, cf. Table 1 lower). The *Cumulative* model by nature cannot forget, since it is trained on both tasks simultaneously, achieving .81 and .74 on Wh-q and Y/N-q, respectively. Interestingly, we observe an *asymmetric synergetic effect*. Being exposed to the Wh-q task helps the

Random (per-task)	WH: 0.09		Y/N: 0.50	
LSTM+CNN+SA	WH: 0.81		Y/N 0.52	
CL SETUPS:	I) WH→Y/N		II) Y/N→WH	
	Wh	Y/N	Y/N	Wh
Random (both tasks)	0.04	0.25	0.25	0.04
Naive	0.00	0.61	0.00	0.81
EWC	0.25	0.51	0.00	0.83
Rehearsal	0.75	0.51	0.51	0.80
Cumulative	0.81	0.74	0.74	0.81

Table 1: Mean accuracy over 3 runs: Trained on each task independently (first two rows; per-task label space \mathcal{Y}) vs. CL setups (single-head label space over all \mathcal{Y}).

Cumulative model improve on Y/N-q, reaching results beyond the task-specific model (from .52 to .74). The effect is not symmetric as the accuracy on Wh-q does not further increase.

Does CL Help? Current CL methods show only limiting (or no) effect. *EWC* performs bad overall: In the II) setup (Y/N→WH, harder task first), *EWC* does not yield any improvement over the *Naive* model; in the WH→Y/N setup, the model’s result on Task A is above chance level (.25 vs. .04) but far off per-task performance (.81). The *Rehearsal* model forgets less than *Naive* and *EWC* in both setups: In the Y/N→WH setup, it is above chance level (.51 vs. .25) reaching per-task random baseline results on Y/N questions (i.e., the model is able to identify Task A, despite the harder single-head setting, in contrast to the *Naive* and *EWC* models). There is no boost derived from being exposed to the Wh-q task in any of the two setups.

Task Order The results in Table 1 show that *the order of tasks plays an important role*: WH→Y/N facilitates CL more than the opposite order: less forgetting is at place when WH is learned first. This confirms psycholinguistic evidence. Overall, *Rehearsal* works better than *EWC*, but mitigates forgetting only to a limiting degree.

Analysis To get a deeper understanding of the models, we analyze the penultimate hidden layer on a sample of 512 questions from the test sets of both tasks (cf. Fig. 2) and relate the representations to confusion matrices of the whole test sets (provided in the SM) and test results (Table 1).

First of all, the model trained on Wh-q discriminates Wh-questions about different attributes very well, reflected in overall high accuracy (.81). It otherwise clusters all instances from the other task

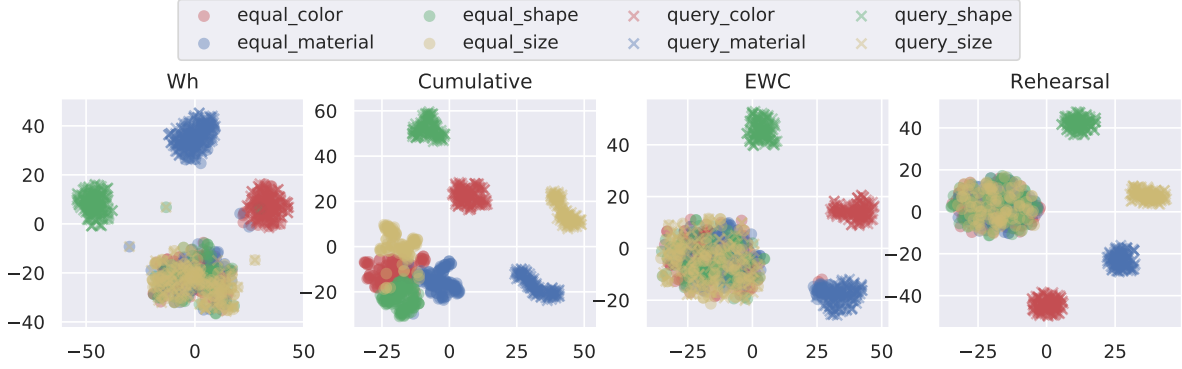


Figure 2: Analysis of the neuron activations on the penultimate hidden layer for the I) WH \rightarrow Y/N setup. “equal_{[shape,color,material,size]}” refers to Y/N-q, “query_{[..]}” refers to WH-questions.

(Y/N-q, which it has not been trained on) around Wh-questions related to size.

The *Cumulative* model, in contrast, is able to further tease the different kinds of Y/N questions apart. Questions about different attributes become distinguishable in the plot, although overall Y/N questions remain closer together than the clusters for Wh-q. This is in line with the lower performance of *Cumulative* on Y/N-q. Our examination of the confusion matrices confirms that the two question types are never confused by the *Cumulative* model. In contrast, the *Naïve* model is very prone to this type of mistake (see plot in SM).

As for the CL models, Fig. 2 (two right-most plots) shows that *EWC* learns representations which are rather similar to those learned by the model trained on Wh-q independently: Y/N questions result in a big hard-to-distinguish “blob”, and are confused with Wh-q about size, as visible in Fig. 2 and the confusion matrix analysis (in the SM). In contrast, *Rehearsal* remembers how to distinguish among all kinds of Wh-q and between Wh-q and Y/N-q. The error analysis confirms that the model hardly makes any mistakes related to task confusion. However, despite the higher performance than *EWC*, *Rehearsal* is still not able to discern well between different kinds of Y/N-q.

5 Related Work

Early work on life-long learning (Chen et al., 2015; Mitchell et al., 2015) is related to ours, but typically concerns a single task (e.g., relation extraction). Lee (2017) aims to transfer conversational skills from a synthetic domain to a customer-specific application in dialogue agents, while Yogatama et al. (2019) show that current models for different NLP tasks are not able to

properly reuse previously learned knowledge.

In general, continual learning has been mostly studied in computer vision. To the best of our knowledge, little has been done on catastrophic forgetting in VQA. A study on forgetting in the context of VQA and closest to ours is Perez et al. (2018). They show that their model forgets after being fine-tuned on data including images with objects of colors other than those previously seen. We took this work as starting point and extended it to consider different types of questions and to test different CL methods beyond fine-tuning.

6 Conclusion

We assessed to what extent a multimodal model suffers from catastrophic forgetting in a VQA task. We built two tasks involving different linguistic characteristics which are known to be learned sequentially by children and on which multimodal models reach different performance.

Our results show that dramatic forgetting is at play in VQA, and for the tested task pairs we empirically found *Rehearsal* to work better than a regularization-based method (*EWC*). More importantly, we show that the *order* in which models learn tasks is important, WH \rightarrow Y/N facilitates continual learning more than the opposite order, thereby confirming psycholinguistic evidence.

Our error analysis highlights the importance of taking the kind of mistakes made by the models into account: A model that does not detect Task A after having been exposed to Task B should be penalized more than a model that answers Task A with wrong task-related labels, but is still capable of identifying the task. Most importantly, our study revealed that differences in the inherent difficulty of the tasks at hand can have a strong im-

pact on continual learning. Regularization-based methods like *EWC* appear to work less well when applied to tasks with different levels of difficulty, as in our experiments. We reserve a deeper investigation of this aspect to future research.

Acknowledgements

We kindly acknowledge the support of NVIDIA Corporation with the donation of the GPUs used in our research to the University of Trento and IT University of Copenhagen. R. Fernández was funded by the Netherlands Organisation for Scientific Research (NWO) under VIDI grant nr. 276-89-008, *Asymmetry in Conversation*.

References

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual question answering. In *International Conference on Computer Vision (ICCV)*.
- Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip Torr. 2018. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *ECCV*.
- Zhiyuan Chen, Nianzu Ma, and Bing Liu. 2015. Lifelong learning for sentiment classification. In *ACL Short paper*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017a. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Judy Hoffman, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. 2017b. Inferring and executing programs for visual reasoning. In *ICCV*.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *PNAS*.
- Sungjin Lee. 2017. Toward continual learning for conversational agents. In *ACL*.
- Davide Maltoni and Vincenzo Lomonaco. 2018. Continuous learning in single-incremental-task scenarios. *arXiv preprint arXiv:1806.08568*.
- James L McClelland, Bruce L McNaughton, and Randall C O’reilly. 1995. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Review*, 102(3).
- T. Mitchell, W. Cohen, E. Hruscha, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohammad, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. 2015. Never-ending learning. In *AAAI*.
- Sara Moradlou and Jonathan Ginzburg. 2016. Young children’s answers to questions. In *Workshop on the Role of Pragmatic Factors on Child Language Processing*.
- Sara Moradlou, Xiaobei Zheng, Ye Tian, and Jonathan Ginzburg. 2018. Wh-questions are understood before polars. In *Proceedings of Architectures and Mechanisms for Language Processing (AMLaP)*.
- Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. 2018. Film: Visual reasoning with a general conditioning layer. In *AAAI*.
- Mark Ring. 1997. CHILD: A first step towards continual learning. *Machine Learning*, 28(1).
- Anthony Robins. 1995. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146.
- Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. 2016. Stacked attention networks for image question answering. In *CVPR*.
- Dani Yogatama, Cyprien de Masson d’Autume, Jerome Connor, Tomas Kocisky, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, et al. 2019. Learning and evaluating general linguistic intelligence. *arXiv preprint arXiv:1901.11373*.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual learning through synaptic intelligence. In *ICML*.