Catastrophic Interference in Predictive Neural Network Models of Distributional Semantics

Willa M. Mannering

Michael N. Jones

Indiana University, Bloomington

Correspondence:

Michael Jones
Department of Psychological and Brain Sciences
Indiana University
Bloomington, IN, 47405
Email: jonesmn@indiana.edu

Abstract

The semantic memory literature has recently seen the emergence of predictive neural network models that use principles of reinforcement learning to create a “neural embedding” of word meaning when trained on a language corpus. These models have taken the field by storm, partially due to the resurgence of connectionist architectures, but also due to their remarkable success at fitting human data. However, predictive embedding models also inherit the weaknesses of their ancestors. In this paper, we explore the effect of catastrophic interference (CI), long known to be a flaw with neural network models, on a modern neural embedding model of semantic representation (*word2vec*). We use homonyms as an index of bias as a function of the order in which a corpus is learned. If the corpus is learned in random order, the final representation will tend towards the dominant sense of the word (*bank* $\rightarrow$ *money*) as opposed to the subordinate sense (*bank* $\rightarrow$ *river*). However, if the subordinate sense is presented to the network after learning the dominant sense, CI almost completely erases the dominant sense and the final representation strongly tends towards the more recent subordinate sense. We demonstrate the impact of CI and sequence of learning on the final neural embeddings learned by *word2vec* in both an artificial language and in an English corpus and evaluate the effectiveness of a recently proposed solution to CI from neuroscience, elastic weight consolidation, on mitigating the effects of CI.

Catastrophic Interference in Predictive Neural Network

Models of Distributional Semantics

Distributional models of semantic memory (DSMs; e.g. Landauer & Dumais, 1997) have been hugely successful in cognitive science, explaining how humans transform first-order statistical experience with language into deep representations of word meaning. These models are all based on the distributional hypothesis from linguistics (Harris, 1970) and specify mechanisms to formalize the classic notion that “you shall know a word by the company it keeps” (Firth, 1957). There are dozens of DSMs in the cognitive literature now, with learning mechanisms inspired by different theoretical camps ranging from Hebbian learning to probabilistic inference (for reviews, see Günther, Rinaldi, & Marelli, 2019; Jones, Willits, & Dennis, 2015). The commonality to all these models is that they use co-occurrence counts of words across contexts in linguistic corpora and exploit these statistical redundancies to construct semantic representations.

There has been a recent resurgence of neural network models across cognitive science and machine learning. This resurgence has included very successful predictive DSMs within connectionist architectures based on principles of error-driven learning core to theories of reinforcement learning. Earlier work had explored neural networks to learn distributed semantic representations from artificial languages using recurrent networks (e.g., Elman, 1990) and feed-forward networks (Rogers & McClelland, 2004). But neural networks were essentially non-existent in the 1990s and early 2000s as models that scale up and learn from natural language corpora. Rather, the field became fixated on algebraic models based on dimensional reduction mechanisms such as singular value decomposition applied to a matrix of word counts across contexts in a corpus (e.g., classic Latent Semantic Analysis; Landauer & Dumais, 1997).

However, the recent literature has hugely reversed this trend, with the emergence of *predictive neural network models*, and the observation that they tend to outperform classic models on most tasks of semantic similarity that are commonly used in natural language processing. Rather than counting words in contexts, these models predict the word given the context, and backpropagate the error signal through hidden layers in a neural network. Although a rather simple architecture, and essentially the same one studied by Rogers and McClelland (2004) in their seminal work, these predictive models have rapidly risen to the top of the DSM battle in their ability to account for human data across a range of tasks.

The standard predictive network currently discussed in the literature is Mikolov et al.'s (2013) *word2vec* model. Word2vec is a feedforward neural network with localist input and output layers that contain one node per word in the vocabulary, and a hidden layer of ~ 300 nodes that is fully connected to both input and output layers. The word2vec architecture has two possible model directions: The context may be used to predict the word (referred to as a CBOW), or the word may be used to predict the context (a skipgram). Although the theoretical claims here apply broadly to neural network models, and hence, both directions in word2vec, we will use skipgram as our demonstration example because it maps conceptually onto most connectionist models of cognition. When a linguistic context is sampled (e.g., “save money bank”) the target node is activated at the input layer (+bank) and activation is forward propagated to the output layer, with the desired output being the observed context words (+save, +money). The error signal (observed output – desired output) is applied to the network with backpropagation (Rumelhart, Hinton, & Williams, 1985) to correct the weights and make it more likely the next time this target word is encountered that the correct output pattern will be generated. Although prediction is used to train the network, it is the final pattern of weights across the input-to-hidden layer that are exported and used as deep

semantic representations of word meaning. Two words that have similar vector pattern across these weights are predicted by similar contexts, even if they never co-occur with each other, akin to (but superior in data fit) the second-order inference vectors learned by traditional DSMs.

Word2vec has received considerable attention in the machine learning literature due to its ability to outperform all previous models (Baroni et al., 2014). To cognitive science, this success is of considerable interest as word2vec implements a potentially biologically plausible neural architecture and links to classic theories of reinforcement learning (e.g., Rescorla & Wagner, 1972), drawing theoretical connections to other areas of cognition with a unified mechanism. One feature of particular interest in these models is that they are incremental learners, in contrast to earlier algebraic DSMs which were largely batch learners. The incremental learning of predictive neural networks has been taken by some as additional evidence in favor of cognitive plausibility of the models. Mandera, Keuleers, and Brysbaert (2017) argue that a predictive model is a theoretical leap forward: “...because the model can be specified using prediction-based incremental learning principles in a neural network, it solves the computational problem posed by the count models in a way that is to a large extent psychologically plausible” (p. 75).

While the enthusiasm surrounding predictive neural networks is certainly warranted given their recent success at fitting human data, it is important to remember that the models also inherit the weaknesses of their predecessor neural networks. One weakness in particular that models such as word2vec are likely to exhibit is *catastrophic interference (CI)*: The tendency of neural networks to completely lose previously learned associations when encoding new ones. In this sense, the positive attribute of predictive neural networks being incremental learners is also what opens them up to a potentially serious flaw that does not exist in their batch learning algebraic counterparts. The goal of this paper is to document how CI is affecting the semantic representations

learned by predictive neural networks, using word2vec in particular, and to evaluate a proposed solution to CI which has had success mitigating the effects of CI on predictive neural network outputs within categorization tasks.

Stability-Plasticity Dilemma in Neural Networks

The *stability-plasticity dilemma* (Grossberg, 1982) refers to the problem of any learning system to learn new stimuli while preventing the new learning from distorting existing memory. We need to balance memory for individual exemplars with abstraction, recency with primacy, and it is optimal to preferentially strengthen memories that are more likely to be needed (see Jones, Dye, & Johns, 2017 for a review). While all cognitive systems gradually forget information, biological organisms exhibit gradual forgetting of old information as new information is acquired. In contrast, artificial neural networks have long been known to forget catastrophically. Catastrophic interference (CI) is thus defined as the sudden and complete loss of previously learned associations when learning new associations (see French, 1999 for a review). CI is a consequence of using backpropagation as a learning mechanism to reuse neural connections to tune learning and is a key flaw to all feedforward neural embedding architectures that are currently used as DSMs.

McClosky and Cohen's (1989) seminal work trained a standard neural network to learn single-digit "ones" arithmetic facts (e.g., $1 + 1$, $9 + 1$) using backpropagation until the network had perfectly learned the associations. They next trained the same network on a new set of single-digit "twos" facts (e.g., $2 + 1$), until the network had been trained to respond correctly to all of them. While the network was able to correctly answer the twos facts, it had completely lost the previously learned ones facts—the associations that had been trained to zero error were now lost completely. The learning of new associations with backpropagation overwrote the previous learning. The CI

pattern was duplicated in a second experiment by McClosky and Cohen by simulating standard tasks of paired associate word learning. Further, Ratcliff (1990) demonstrated that in standard sequential learning paradigms, backpropagation networks catastrophically forget previous items as new items are sequentially learned, unlike humans performing the same tasks.

Additionally, the backpropagation algorithm has been criticized for being biologically implausible in multiple ways. Some popular criticisms are, among others, (1) that the computation is only linear, whereas biological systems contain both linear and non-linear processes, (2) that biological neurons communicate via binary values not continuous values, (3) if backpropagation were used in the brain, the feedback paths would need to use precisely symmetric weights to the feedforward paths (known as the Weight Transport Problem: Bengio et al., 2015; Grossberg, 1987). While there are many concerns about the biological plausibility of backpropagation *as it is now* there is a huge ongoing effort within the neuroscience and machine learning communities to alter the algorithm to make it more plausible. One such endeavor, referred to as elastic weight consolidation, we explore in detail in this paper and evaluate the application to the semantic modeling field.

Using Homonyms to Measure Representational Bias in Semantic Space

The standard architecture used by neural embedding models such as word2vec is susceptible to CI, but it is unclear if, or to what extent, CI would affect the final semantic representation. Word2vec uses predictive association for learning (e.g., bank $\rightarrow$ save + money) and backpropagation for error correction. But the final representations of word meanings are contained in the vector of association weights across the input-to-hidden layer. It is reasonable to expect that if a word predicts very different contexts, such as homonyms (bank $\rightarrow$ save + money; bank $\rightarrow$ sand + river) that the final semantic representation for a word will be heavily dependent

on the most recently learned sense association, potentially producing a great loss of the previously learned sense association. Hence, homonyms act similarly to classic XOR stimuli in experimental studies: The same input pattern predicts two (or more) orthogonal output patterns.

Figure 1. A schematic illustration of the semantic space of a homograph (*bank*) as a function of sense learning order. The left panel shows the ideal situation if the contexts of a corpus were randomly sampled, with *bank* equidistant to its two most dominant senses and relatively far from its rare verb sense. The middle panel shows the case where river contexts have been sampled most recently—in this case, *bank* is biased by backpropagation to the more recent sense despite the fact that it is equally frequent in the corpus as the money sense. The right panel illustrates the potential case where the rare verb sense (“The pilot had to bank the airplane”) is sampled most recently—in this case, the recency bias overpowers frequency.

The first goal of this paper is to explore the impact of the training sequence of context/target pairs on the final representational space in the word2vec skipgram architecture. Hence, we use homonyms as an index of movement in semantic space. For a homonym with two equally frequent distinct meaning senses, the semantic representation of the target word in word2vec should be equidistant between the two opposing meanings in semantic space. If the homonym has a dominant and subordinate sense, then the final meaning will tend towards the more frequent dominant sense in semantic space. However, if contexts containing the subordinate meaning are the most recently learned, then the semantic representation may erroneously tend towards recency over the more frequent meaning. Consequently, we can use homonyms as an elegant measure of how a word’s meaning differs from a randomly sampled corpus when we make certain sense contexts more or

less recently presented to the backpropagation algorithm. The second goal of this paper is to evaluate a possible solution to CI, elastic weight consolidation (EWC: Kirkpatrick et al., 2017), which has been impressively successful on machine learning tasks and can be considered a “vaccination” for predictive neural networks that could potentially prevent the effects of CI.

This paper is split into two sections: the first section addresses our first goal, to determine the effects of CI on final semantic representations, and the second section addresses our second goal, to determine the effectiveness of implementing EWC. Both the first and second sections of this paper are split into two experiments. Experiment 1 in both sections uses a simple engineered language presented to the models. The corpus learned is the same in all conditions, but the ordering of the contexts is varied (cf. Ratcliff, 1990). Experiment 2 scales the principles up to a natural language corpus where contexts containing the dominant and subordinate senses of the target word are presented to the models, either in random order, or in sequentially manipulated orders.

1. Estimating the Impact of CI on Semantic Representations

Experiment 1a: CI in an Artificial Language

As an initial evaluation of CI in word2vec, we created a simple artificial language inspired by Elman (1990). In the language, there is a single homonym, *bass*, with two distinct senses—*bass* [fish], or *bass* [guitar]. There are two actors (man/woman) who may occur in either sense context. Hence, a corpus containing the simple language was created by randomly sampling from the Markov grammar displayed in Figure 2.

We generated a corpus of 8,000 sentences from this grammar (e.g., “man catch bass,” “woman play bass,” ...). The corpus was generated to ensure that *bass* occurs an equal number of times in the fish and guitar contexts, and a neural embedding model will form a semantic representation for *bass* that is an average of its two distinct sense contexts. To measure the

Figure 2. The artificial language used to train the model. Bass is the target homonym, and its position in semantic space relative to the two sense-pure synonyms (acoustic/trout) is measured as a function of corpus learning sequence.

semantic position of *bass* relative to its two distinct senses, we compute the cosine to its two sense-pure synonyms, *trout* and *acoustic*. When the corpus is sampled randomly, *bass* has an equal similarity to both *trout* and *acoustic*. However, when the order of senses in the corpus is not random but favors a more recent sense, the position of *bass* is expected to tend towards the sense-pure synonym of that sense.

The word2vec skipgram architecture (Mikolov, et al., 2013) was trained on the corpus of 8,000 sentences sampled from the artificial language in three distinct orderings: random, fish context first, and guitar context first. The *exact* same set of 8,000 sentences was presented to the model in each condition, but the sequencing was varied (i.e., LSA would produce the exact same representation for all three). In the sequenced versions, ordering was randomized but nested within

the constraint that one sense occurred in the first half of the corpus, and the other sense was restricted to the second half of the corpus. Because word2vec has stochastic weight initialization, we ran the model for 200 replications on each ordering of the corpus. Each training run, the model was presented with the full corpus (8,000 training epochs) and we recorded the vector cosine of *bass* to *trout* and *acoustic*.

The skipgram model was trained by minimizing objective function predicting a target word given its surrounding context, defined as the mean log likelihood:

$$\frac{1}{T} \sum_{t=1}^T \sum_{j \neq 0} \log p(w_{t+j} | w_t) \quad (1)$$

where $\log p(w_{t+j} | w_t)$ is implemented with a softmax:

$$p(w_o | w_i) = \frac{\exp(v'_{w_o} v_i)}{\sum_{w=1}^W \exp(v'_{w_o} v_i)} \quad (2)$$

Results

Figure 3 shows the vector cosine of *bass* to *trout* and *acoustic* in the final semantic space for word2vec across 200 replications, while varying ordering. When the model is trained on a random ordering of the corpus, *bass* is equidistant to *trout* and *acoustic*. However, when all the fish contexts are presented first, followed by the guitar contexts, the resulting semantic spaces showed a very strong bias towards the more recent sense.

For example, if the contexts representing the fish sense of *bass* were presented first, the model learned that *bass* was synonymous with *trout* until the second half of the corpus when the new sense was introduced. With the introduction of a guitar sense of *bass*, word2vec then had to reuse the connection weights to learn the new context association. As a result, *bass* became

synonymous with *acoustic*, and the previous prediction of fish contexts was almost completely erased.

Figure 3. Cosine similarity of the homonym *bass* to the fish or guitar sense as a function of sense presentation order. When the corpus is sampled randomly, *bass* is equally pulled between the two senses. However, when one sense is presented earlier in the corpus, the final representation shows a strong bias towards the more recently learned sense.

The final similarity of *bass* after learning the corpus in the order of fish-then-guitar was maximal to the sense-pure synonym (*acoustic*) and was reduced to nearly zero for the originally learned sense-pure synonym (*trout*). In contrast, the reverse preference is seen when the ordering of senses is reversed. Note that a batch model such as LSA would produce the behavior seen in the random case across all three orderings. Very similarly to the classic McClosky and Cohen (1989) experiments, word2vec is largely overwriting the learning of the first context with training on the second context, and this pattern is reflected in the final semantic representation of the homonym. Figure 4 shows the continuous, average learning curve over 5 runs of the model. In this figure, the guitar sense is learned first and the fish sense is learned second. When the fish sense

is introduced to the model, all previous knowledge of the guitar sense is lost. This figure shows a similar pattern to Figure 3.

Figure 4. Cosine similarity of the homonym *bass* to the fish or guitar sense as a function of sense presentation order. The effects of CI are averaged across five runs of the model here, and the qualitative trend is consistent across the parameter space.

The competing sense proximity is random in Phase 1, but when this sense is then reinforced in Phase 2, the model's forgetting of the first sense very rapidly parallels the rapid learning of the new sense. However, an interesting feature to note is that while the first sense is no longer presented to the model in Phase 2, the model still shows continued and active unlearning of the initial sense, until it is effectively wiped out more so than if it were simply never presented. This is somewhat in contradiction to the recent finding that repeated training on previously learned (but

forgotten) tasks shows a relearning savings with strategy-learning neural nets (Davidson & Mozer, 2019; Goodfellow et al., 2015).

We next varied the sense dominance of the homonym *bass*, to present a dominant and subordinate sense frequency in the artificial corpus. Rather than sampling randomly from the two senses, the subordinate sense of the homonym was sampled to be one-third as likely as the dominant sense. Hence, one sense is two-thirds more prevalent in the experienced sentences, and we test word2vec’s sense preference as a function of training sequence.

The results are presented in Figure 5. In the two randomly presented conditions (two left clusters of bars), word2vec prefers the dominant sense of the word over the subordinate sense—the representation of *bass* was more similar (closer) to the dominant sense from the corpus. But a dramatic reversal can be seen in the ordered conditions (two right clusters of bars). In both of these cases, the dominant sense is presented first, followed by one-third as many sentences from the subordinate sense second. Consequently, in both cases, CI shows its effect clearly in the final representations: there was a strong bias towards the more recently learned sense in the semantic space, despite the fact that it was less frequent than the dominant sense in the corpus. Even if the fish sense of *bass* was the dominant sense, recent presentations of the subordinate guitar sense almost completely overwrites that learning. Again, note that a batch model such as LSA would simply show a preference for the dominant sense across all of these presentation orders, but word2vec is influenced by CI to produce a recency bias that overrides the standard frequency bias.

Experiment 1b: CI in a Natural Corpus

We next test the influence of CI in the representational space learned by word2vec when trained on natural language from the classic TASA corpus (Landauer & Dumais, 1997). TASA is ideal because it contains linguistic contexts from textbooks with metadata that tags the topic of

the material (e.g., Science, Language Arts, Health, etc.), which allowed us to select and track a sample of homonyms with distinct senses.

Figure 5. Cosine similarity of the homonym bass to the fish or guitar sense as a function of order and varied sense dominance. When sampled randomly, bass tends towards the more dominant sense. However, if the subordinate sense is presented more recently, it outweighs the dominant sense, losing the association of frequency to recency.

We used the homonym norms from Armstrong, Tokowicz, and Plaut (2012) to identify a sample of 14 homonyms that also exist in TASA with very distinct meaning senses as rated by human subjects. Of the sample of homonyms, half have roughly equal frequencies in TASA between the two senses, and half have a clear and dominant sense (as per frequency in TASA). An example of a sense-balanced homonym is *pupil*—the “part of the iris of the eye” sense occurs in an equal number of science contexts across TASA as the occurrence of the “young learner in

school” sense in social studies contexts. An example of a sense-imbalanced homonym is *firm*—the “business facility” sense occurred in business contexts across TASA eight times more often than did the “having a solid structure” sense in science contexts.

Method

If the occurrence of the homonym in one sense exceeded the other sense by a factor of two, we classified it as sense biased, and classified it as sense balanced otherwise. Six of the homonyms were classified as balanced (*hamper, capital, net, slip, plane, pupil*), and the remaining eight were classified as biased (*firm, compact, hull, compound, pitch, cap, gum, bull*).

The distance of the homonym in semantic space must be measured with respect to some sense-pure synonym, analogous to the artificial grammar simulations that measured *bass* in respect to *trout* or *acoustic*. To do this, we first trained word2vec on the single sense of each homonym, e.g., training the model separately on the “student” and “eye” contexts of *pupil*. For each separate sense space, we determined the target word’s most similar semantic associate. Due to the stochastic noise added by the weight initialization process, we identified each target word’s closest 10 neighbors across 20 replications for sense 1 (e.g., *pupil*→*student, teacher, learn, classroom, books, etc.*) and sense 2 (e.g., *pupil*→*iris, cornea, eye, vision, etc.*), and selected the single associate that was most often the highest ranked across the resamples. This word will be considered the target word’s sense-pure semantic associate as were *trout* and *acoustic* in the artificial language, and will be used as an anchor point to measure the target word’s representation as a function of corpus training order.

With the sense-pure semantic associate identified for each sense of each homonym, we then trained on the entire corpus, which mixed the sense contexts together. In the random condition, the contexts for the target word were simply sampled in random order, as in usual

applications of word2vec. In the other two classes of simulations, we first trained on one sense of the target word first, and then the other sense.

Results

Figure 6 displays the results for the sense-balanced homonyms, analogous to Figure 3 from the artificial corpus. In the random training order, the target homonym is evenly pulled between its two senses. The homonym representation learned by the model has an equal similarity to each sense- pure associate. However, the next two sets of bars show the effect of CI that comes with sequential sense training. In both cases, the target homonym becomes more similar to the recently learned sense and less similar to the previously learned sense. Note that the forgetting is not *completely* catastrophic. However, this is a very profound effect on the final learned representation of the homonym considering that this is the exact same input corpus across all three training orders.

Figure 6. Similarity of the homonym representation from TASA for the sense balanced cases, trained in random order or with sequenced senses. The CI effect is illustrated in the second two clusters of bars where the target homonym is more similar to the recently encountered sense and less similar to the previously encountered sense.

Figure 7 plots the results for the sense-imbalanced homonyms from TASA. Each target homonym here has a bias in TASA towards its more frequent sense (e.g., *firm* is encountered in many more financial contexts in TASA than it is in science contexts). When trained in random order, the target homonym is more similar to the dominant sense in the corpus. The second two clusters of bars show the results for the sequenced training where one sense is learned before the other. If the dominant sense is also the more recently learned sense (middle cluster of bars), then the target homonym becomes even more similar to the dominant sense and less similar to the subordinate sense. However, the effect of CI is seen prominently as a reversal in the far right cluster of bars. When the subordinate sense is learned after the dominant sense, the similarity of the target homonym favors recency over dominance. In this case, the target homonym actually becomes more similar to the recently presented subordinate sense than it is to the dominant sense. Again, the forgetting is not completely catastrophic. But the pattern is very powerful: CI produces an effect that makes recency overpower frequency. The model believes that the target homonym's meaning is more similar to the subordinate sense over the dominant sense, simply because it was encountered more recently.

Discussion

The resurgence of predictive neural network models has already led to reconceptualization of the mechanisms underlying cognitive phenomena. In addition, the utility of these models to machine learning demonstrates the importance of basic cognitive science to solving applied problems. However, it is crucial to remember our history: CI is a problem that was never fully addressed in the original connectionist models of the 90s. We have made a clear first demonstration

Figure 7. Similarity of the homonym representation for the sense imbalanced cases. In random order, the target homonym is most similar to the more frequent sense (left bars). If the dominant sense is also more recent, the homonym moves closer to the dominant sense and further from the subordinate one (right bars). However, CI produces an effect that reverses this trend (middle bars): if the subordinate sense is more recently encountered, the homonym representation favors recency over frequency.

here that CI can have a strong impact on the semantic representations learned by neural embedding models such as word2vec. This is particularly important because the potential source of error goes beyond our basic science and into the many applications of word2vec. McCloskey and Cohen (1989) and Ratcliff (1990) used CI to demonstrate an inherent flaw with neural networks as theoretical models of cognition. But beyond that, neural embedding models such as word2vec are being used in a massive number of NLP and knowledge mining applications, potentially bringing the errors that come with CI along for the ride.

Recently, algorithms that alter the plasticity of weights and weight update functions of neural networks have become the de facto solutions for CI. One such algorithm is Elastic Weight

Consolidation (EWC) which has shown great promise at learning new associations while insulating previously learned associations against forgetting (Kirkpatrick et al., 2017). EWC constrains the weight space of a deep learning network within the optimal parameter space of a previously learned task, essentially having the effect of a spring (or elastic) on already learned weights that are important to a previously learned task. Kirkpatrick et al. demonstrated that EWC networks could learn new strategies and associative patterns with minimal loss to previously learned but orthogonal associative patterns. However, EWC and its relatives have not yet been implemented in predictive networks of semantic learning. The goal of the second section of this paper is to evaluate the effectiveness of equipping neural networks with EWC to prevent the effects of CI in final semantic representations.

2. Vaccinating Against CI with Elastic Weight Consolidation

The goal of elastic weight consolidation (EWC) is to allow a predictive neural network to learn two sequential tasks, Task A and then Task B, without incurring CI. To do this, Kirkpatrick et al. (2017) introduced a method to constrain the parameters of a neural network after learning Task A so that the network can subsequently learn Task B without forgetting Task A. The new loss function they introduce is a quadratic penalty that differentially constrains parameters in the neural network depending on how important each parameter is to completing Task A. To determine which weights in the network are important for Task A they calculate the Fisher Information for each parameter—a mathematical method to measure the amount of information a variable carries about a parameter. The resulting loss function to be minimized is:

$$\mathcal{L}(\theta) = \mathcal{L}_B(\theta) + \sum_i \frac{\lambda}{2} F_i (\theta_i - \theta_{A,i}^*)^2 \quad (3)$$

where $\mathcal{L}_B(\theta)$ is the loss for Task B, λ controls how important Task B is compared to Task A, F is the Fisher Information calculated for each parameter, and θ represents the parameters in the

network. Kirkpatrick et al. (2017) showed that EWC was able to insulate against CI when training a predictive neural network on the MNIST (LeCun et al., 1998) data set of handwritten images and on standard Atari games. While EWC has been tested several times on categorization tasks, this paper will present the first implementation for use with distributional semantic models.

EWC has potential as a “vaccine” for predictive DSMs, that is, networks may be insulated from CI without having to implement new architectures. There is reason to suspect that EWC may have limited effectiveness when translated to the field of semantic modeling. EWC calculates the relevance of each model parameter to Task A based on the actual class of the training data. However, in the case of semantic modeling, we are not necessarily interested in the final predicted class of the training data but in the internal representations created by models as they learn. It is the goal of this section to determine how EWC affects the internal representations of neural embedding models using the same experimental structure used in the first section.

Experiment 2a: Effects of CI on EWC model in an Artificial Language

The same artificial language was created by the Markov grammar shown in Figure 2. In the first condition, a corpus of 8,000 sentences was generated from this grammar (“man catch bass”, “woman play bass”). Each sense of the word *bass* was equally frequent; the fish-sense made up half of the total sentences and the guitar-sense made up the other half. In the second (sense biased) condition, the subordinate sense was 1/3 as frequent as the dominant sense.

Results

EWC adds one additional parameter to the word2vec model, λ , which controls the importance of Task A compared to Task B. For both conditions in the first experiment, the value of λ had little to no effect on the final semantic representations. The results shown below are representative of the results obtained by any value of λ .

Figure 8 shows the cosine similarity of the vectors produced by word2vec and word2vec vaccinated with EWC in the case where sense 1 and sense 2 of *bass* are equally frequent. The pattern produced by word2vec is consistent with the findings in the first section of this paper. When the model was trained in random order, the *bass-sense1* and *bass-sense2* similarities produced were approximately equal. When trained in sequential order, the sense sampled most recently had a higher similarity to *bass*. The same procedure was repeated using word2vec equipped with EWC. After exploring various parameter settings, we found that implementing EWC had virtually no effect on the results of the first experiment.

Figure 8. The y-axis represents the cosine similarity of vectors produced by word2vec and word2vec equipped with EWC. The x-axis represents one of three training orders: random order, sequential order with sense 1 first then sense2, and sequential order with sense 2 first then sense 1. Sense 1 and Sense2 are equally frequent in this case. CI is present in both word2vec and word2vec equipped with EWC.

Figure 9 shows the cosine similarity of the vectors produced by word2vec and word2vec equipped with EWC in the case where one sense is dominant and the other is subordinate. The

pattern produced by word2vec is once again consistent with our findings in the first section of this paper. When trained in random order, the dominant sense is more similar to *bass* than the subordinate sense. When trained in sequential order, the effects of CI reverse the frequency effects; when the subordinate sense of *bass* is trained recently it becomes more similar to *bass* than the dominant sense. When the same procedure was performed using EWC, we saw similar results to the first condition. The addition of EWC did not change the performance of word2vec.

Figure 9. The y-axis represents the cosine similarity of vectors produced by word2vec and word2vec equipped with EWC. The x-axis represents one of three training orders: random order, sequential order with the dominant first then the subordinate sense, and sequential order with the subordinate sense first then the dominant sense. CI is present in both word2vec and word2vec equipped with EWC.

Experiment 2b: Effects of CI on EWC model in a Natural Corpus

We used the same natural language corpus (TASA) in the first and second sections of this paper. We also used the same set of homonyms to test both word2vec and word2vec equipped with EWC as was used in the first section of the paper. Once again, the 14 homonyms were divided into

two groups: sense-balanced and sense-imbalanced. We classified the two senses of a homonym as sense-imbalanced if one sense was at least twice as frequent in the TASA corpus, otherwise the two senses of a homonym were classified as sense-balanced.

Results

Figure 10 shows the results of training word2vec and word2vec vaccinated with EWC on the sense-balanced homonyms from TASA. The pattern of cosine similarities produced by word2vec are consistent with the results from the artificial language. When trained in random order the target words have approximately equal similarities to both of its senses. When trained sequentially, we see the same issue that occurred in the first experiment—the sense that was trained last becomes more similar to the target word. The model equipped with EWC once again does not show improvement over word2vec.

Figure 10. The y-axis represents the cosine similarity of vectors produced by word2vec and word2vec equipped with EWC when trained on sense-balanced homonyms from the TASA corpus. The x-axis represents one of three training orders: random order, sequential order with sense 1 first then sense2, and sequential order with sense 2 first then sense 1. CI is present in both models.

Figure 11 shows the results of training word2vec and word2vec with EWC on the sense-imbalanced homonyms from TASA. The cosine similarities produced from word2vec are consistent again with the results from the artificial language. Similarly, the cosine similarities produced by word2vec with EWC do not improve over the results from word2vec.

Figure 11. The y-axis represents the cosine similarity of vectors produced by word2vec and word2vec with EWC when trained on sense-imbalanced homonyms from the TASA corpus. The x-axis represents one of three training orders: random order, sequential order with the dominant first then the subordinate sense, and sequential order with the subordinate sense first then the dominant sense. CI is present in both models.

Discussion

While Kirkpatrick et al. (2017) were able to show promising results from EWC on categorization tasks, we found no such benefit when the method was applied to semantic modeling domain. This may be because the goal of EWC is to prevent the weights of a predictive neural network from changing based on how much information each weight carries about the *true class*

of each training item. The connection between training items and their class is very straightforward in categorization tasks but is not as clear in semantic modeling tasks. Due to computational constraints, when predictive neural networks like word2vec learn a word representation from a natural corpus, they do not explicitly predict the class of a word but attempt to predict which words belong or don't belong in context with a target word. This method for learning word representations is called noise contrastive estimation and allows predictive neural networks to operate on the scale of natural language. In addition, the larger the output layer (which is equal to the number of vocabulary items), the more complex and difficult the problem of credit assignment is.

General Discussion

Humans are not fed randomized linguistic contexts and provided with explicit error feedback for their predictions. Even before school age, children learn in temporally, spatially, and semantically clustered environments, and they generate their own datasets by exploring the environment (Smith & Gasser, 2005). Throughout the lifespan, human learners are faced with the problem of acquiring conceptual knowledge from input that is clustered serially, and they must integrate new semantic knowledge without catastrophically damaging existing knowledge. In fact, many of the best practices for training machine learning systems (randomization of training instances, uniform sampling of competing categories, etc.) are virtually opposite from the environmental structure that humans naturally learn from. If the goal of our basic science is to understand the psychological mechanisms that humans have evolved to learn semantics from the environment, CI is a litmus test for the plausibility of an architecture. In addition, the field of machine learning is now seeing explicit calls to researchers to build machines that learn like people (Kriegeskorte, 2015; Lake, Salakhutdinov, & Tenenbaum, 2015; Tenenbaum, Kemp, Griffiths, &

Goodman, 2011), and CI is likely to be a problem with future applications of these models, from self-driving cars to augmented reality.

The results shown in the first section of this paper indicate the CI is indeed producing a strong bias in the final semantic representations of neural embedding models. But this raises the question: How big of a problem is CI really likely to be in the regular application of neural embedding models? Firstly, CI is still very much a theoretical flaw with these architectures as theories of cognition, in that they show forgetting and interference patterns that are extremely unlike biological systems. But the version of word2vec that we ran here was essentially stripped down to a core predictive neural network to test the basic theoretical mechanism of prediction and error feedback. The full word2vec algorithm includes negative sampling, frequency subsampling, and other machine learning tricks to speed learning and scale up to large amounts of text. Frequency subsampling adjusts the sampling of contexts of words inversely proportionate to their normative frequency. The idea is to sample word tokens roughly an even amount, and so the model needs to account for the Zipfian law of word frequencies or it will be spending far too much time on already well-learned frequent words, and not enough on rarer words further out on the Zipfian tail. The addition of frequency subsampling can potentially allow CI to do even more damage to the final semantic space than without it. The practice means that fewer samples are taken from high-frequency words, meaning that there is a greater likelihood that a single rare sense of the word lemma could be sampled most frequently, undoing previous learning for the more common sense of the word. Frequency subsampling assumes a single sense for a word but does not take into account the fact that there is also a skewed distribution of sense frequencies within homonyms and polysemes, which constitute half the lexicon (Lin & Ahrens, 2005). Clearly, more research is

needed to evaluate how much variance there is in word2vec’s estimate of a word’s position in semantic space across multiple replications, and in different sequential orders, of a training corpus.

The results of the second section of this paper suggest that efforts to mitigate the effects of CI need to be interdisciplinary. Within the machine learning and the neuroscience community, insulating, or “vaccinating”, predictive neural networks from CI is an emerging area that has seen some innovation in recent years. However, suggested solutions so far only consider the problem as it relates to strictly machine learning tasks such as categorization or image classification tasks. The second section of this paper has shown that solutions from the machine learning community are not guaranteed to work when applied to tasks from different fields.

Additionally, EWC as it is now is not theoretically plausible for any task which requires unsupervised learning because the new loss function must be “switched on” when the network is learning a second task—the network needs to be externally told when senses are changing, which skirts the problem of learning the signal. This is especially cumbersome in NLP where it is impossible to supervise learning to the extent which EWC requires. Furthermore, EWC is unable to scale up well with its current implementation. Because it was designed to prevent CI in categorization tasks, EWC requires each training item to have a true class. This requirement necessitates the use of an output layer which represents the probability distribution of the classes for each word. These types of output layers are used in classic softmax learning mechanisms which are feasible for classification tasks or artificial languages but run into severe computational issues once natural language is introduced. Thus, the connection between modern learning methods, such as noise contrastive estimation, to an output layer of this type—which are unused by modern methods because they are computationally expensive—is unclear. Similarly, calculating the Fisher

Information for each node in a network gets very computationally expensive when the vocabulary and network increase to human scales.

Currently, the most successful approach in the field of machine learning to mitigate the effect of CI on predictive neural networks is known as *dropout* (Hinton et al., 2012). Dropout is an addition to stochastic gradient descent training in which the input and hidden layers of a network are multiplied by a binary mask with each training instance that is learned. In effect, this allows many different networks to be trained on subsets of the training set and the resulting model predictions to be averaged together, effectively regularizing the model’s predictions over sequenced training sets. Although it has been claimed that dropout acts somewhat similar to how hippocampal-neocortical complimentary encoding systems function (Goodfellow et al., 2015), the link is rather tenuous, and is the same theoretical claim that is made with EWC and various other algorithms that are computationally quite different. Other practical approaches include Gradient Episodic Memory (Lopez-Paz & Razanto, 2017), which basically retains exemplars of previous tasks when learning new ones. But none of these approaches is likely to be adaptable to the learning of clustered multiple senses of words, and all lack cognitive and neural plausibility (see Parisi et al., 2019, for a review).

We have shown in the second section of this paper that “vaccine” methods used to fix neural networks are not theoretically or practically feasible in the field of semantic modeling. So then, what can be done to insulate distributional semantic models from CI? We believe the answer lies in alternative architectures. The most promising approaches take their organizations from the anatomical and functional organization of the brain. For example, complimentary systems theory (McClelland, McNaughton, & O’Reilly, 1995) posits that a slow neocortical processing system and faster hippocampal encoding system in the human brain allows deep processing of predictive

information while avoiding problems that come with CI. The neocortical system learns deep abstractions but is susceptible to CI, however, it is complimented by the hippocampal replay of recent episodic experience. More recent neurobiological models of memory, such as the cascade model (Benna & Fusi, 2016), manage to avoid problems of CI by implementing a bimodal system that initially learns very rapidly but gradually transfers information to a slower-learning mechanism. For reinforcement learning of procedural tasks in machine learning, the most successful convolution networks of visual perception use “experiential playback” to allow them to randomize over the training input that is otherwise experienced in a serial manner (Mnih et al., 2015), very much like complimentary learning systems theory.

Other architectures, such as holographic neural networks, random vector accumulators, and exemplar-based models, are also theoretically immune to CI and incorporate different theoretical frameworks of learning (see Jones, et al., 2015). Holographic neural networks use convolution as an association mechanism to learn words rather than backpropagation and are able to learn complex non-linear patterns with a single layer which makes them more space efficient than predictive neural networks (Jones & Mewhort, 2007). Random vector accumulation models should also be immune to CI because they utilize principles of associative learning and do not rely on an error signal—learning via a simple Hebbian co-occurrence learning mechanism. Exemplar-based models, unlike other DSM models which store an abstract semantic representation, store only episodic context (Jones, 2018). These models construct semantic meaning from the aggregation of episodic context when presented with a memory cue (Jamieson et al., 2018). Alternate architectures should be evaluated to determine the effect CI has on their internal semantic representations.

While alternate architectures are promising solutions, they have faced criticisms in the past. Some architectures are known to have problems with metric space compression—causing most word similarities to be compressed between 0 and 1—which limits the ability of the model to discriminate between related and unrelated words (Asr & Jones, 2017). It was initially believed that predictive neural networks were able to more accurately discriminate between words because of back-propagation or the connectionist architectures they commonly use (which is one of the reasons this architecture has become so popular). However, recently the role of negative sampling in DSMs has been explored in more depth by Johns, Mewhort, & Jones (2019) who find that the success predictive neural networks have at discriminating between words is due to the inclusion of negative information in the training data—not the use of connectionist architecture or predictive error correction. In fact, Johns et al. demonstrated when negative sampling information is included in the training data for other DSMs, including random vector accumulators, their ability to discriminate words is on par with predictive neural networks. This indicates that it is not error correction that is producing the benefit of predictive models, but the benefit can be seen in errorless Hebbian learning models if they also implement negative sampling.

Up until now, the fields of machine learning and cognitive science have both been facing similar problems with CI in predictive neural networks. Unfortunately, there has been little to no interdisciplinary communication to propose solutions. When we consider CI from a cognitive science perspective, we find that there are several possible solutions which haven't been considered yet. These solutions, which are arguably more elegant and theoretically immune to CI than continuously trying to “vaccinate” predictive neural networks, have the potential to introduce new mechanisms for artificial learning, assisting with new technological advances that require

sequential learning and providing a framework for learning that does not exhibit the downfalls brought on by predictive neural networks.

References

- Armstrong, B. C., Tokowicz, N., & Plaut, D. C. (2012). eDom: Norming software and relative meaning frequencies for 544 English homonyms. *Behavior research methods*, 44(4), 1015-1027.
- Asr, F. T., & Jones, M. N. (2017). An Artificial Language Evaluation of Distributional Semantic Models. Proceedings of the ACL Conference on Natural Language Learning (CoNLL).
- Baroni, M., Dinu, G., & Kruszewski, G. (2014). Don't count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings Association of Computational Linguistics* (Vol. 1, pp. 238-247).
- Bengio, Y., Lee, D.H., Bornschein, J., Mesnard, T., & Lin, Z. (2015). Towards Biologically Plausible Deep Learning. arXiv: 1502.04156 [cs]
- Benna, M. K., & Fusi, S. (2015) Computational principles of biological memory. arXiv: 1507.07580 [q-bio]
- Davidson, G., & Mozer, M. C. (2019). Sequential mastery of multiple tasks: Networks naturally learn to learn. *arXiv preprint arXiv:1905.10837*.
- Elman, J. L. (1990). Finding structure in time. *Cognitive science*, 14(2), 179-211.
- Firth, J. R. (1957). *A synopsis of linguistic theory* (pp. 1930–1955). Oxford.
- French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4), 128-135.
- Goodfellow, I. J., Mirza, M., Xiao, D., Courville, A., & Bengio, Y. (2013). An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*.
- Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review*, 87, 1-51.

- Grossberg, S. (1987). Competitive learning: from interactive activation to adaptive resonance. *Cognitive Science*, 11, 23-63.
- Günther, F., Rinaldi, L., & Marelli, M. (2019). Vector-space models of semantic representation from a cognitive perspective: A discussion of common misconceptions. *Perspectives on Psychological Science*, 1745691619861372.
- Hasselmo, M. E. (2017). Avoiding catastrophic forgetting. *Trends in Cog Sci*, 21.
- Harris, Z. (1970). Distributional structure. In *Papers in structural and transformational Linguistics* (pp. 775–794).
- Hu, X., Cai, Z., Graesser, A. C., & Ventura, M. (2005). Similarity between semantic spaces. In *Proceedings of the 27th annual conference of the cognitive science society* (pp. 995-1000). LEA Hillsdale, NJ.
- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. R. (2012). Improving neural networks by preventing co-adaptation of feature detectors. arXiv preprint arXiv:1207.0580.
- Jamieson, R. K., Johns, B. T., Avery, J. E., & Jones, M. N. (2018). An instance theory of semantic memory. *Computational Brain & Behavior*, 1(2), 119-136.
- Johns, B. T., Mewhort, D. J. K., & Jones, M. N. (2019). The role of negative information in distributional semantic learning. *Cognitive Science*.
- Jones, M. N. (2018). When does abstraction occur in semantic memory: insights from distributional models. *Language, Cognition and Neuroscience*, 1-9.
- Jones, M. N., Dye, M., & Johns, B. T. (2017). Context as an organizational principle of the lexicon. In B. Ross (Ed.), *The Psychology of Learning and Motivation*, 67:43.

- Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological review*, 114(1), 1-37.
- Jones, M. N., Willits, J. A., & Dennis, S. (2015). Models of semantic memory. In J. R. Busemeyer & J. T. Townsend (Eds.) *Oxford Handbook of Mathematical and Computational Psychology*, 232-254.
- Kirkpatrick, et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proc Natl Acad. Sci. U.S.A.*, 114, 3521-3526.
- Kriegeskorte, N. (2015). Deep neural networks: A new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science*, 1, 417-446.
- Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266), 1332-1338.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211-240.
- LeCun, Y., Cortes, C., & Burgess, C. J. C. (2012). The mnist database of handwritten images.
- Lin, C. C., & Ahrens, K. (2005). How many meanings does a word have? Meaning estimation in Chinese and English. *Language acquisition, change and emergence: Essays in evolutionary linguistics*, 437-464.
- Mandera, P., Keuleers, E., & Brysbaert, M. (2017). Explaining human performance in psycholinguistic tasks with models of semantic similarity based on prediction and counting: A review and empirical validation. *Journal of Memory and Language*, 92, 57-78.

- McCloskey, M. & Cohen, N. (1989) Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bower (ed.) *The Psychology of Learning and Motivation*, 24, 109-164.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (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. *Psychological Review*, 102(3), 419-457.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Petersen, S. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems* (pp. 3111-3119).
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks*.
- Ratcliff, R. (1990) Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological Review*, 97, 285-308.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research and theory*, 2, 64-99.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. MIT press.

- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1988). Learning representations by back-propagating errors. *Cognitive modeling*, 5(3), 1.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279-1285.

Appendix A: Implementation details for models used in all experiments

Experiment 1a: CI in an Artificial Language The Python module TensorFlow was used to implement a bare-bones version of a predictive neural network. This version does not implement subsampling like the full word2vec model—it is unneeded because the vocab size is so small. This model used noise contrastive estimation as a loss function and optimized the function using the Adam Optimizer. The embedding size was 10, the learning rate was 0.02, the window size was 2, and one negative sample was used to train the model. Negative information and noise contrastive estimation were used because they are standard in the language learning community (word2vec uses these as well).

Experiment 1b: CI in a Natural Corpus The word2vec model from Gensim was used to train word2vec on the natural language corpus. Gensim is an open-source python module that includes popular, pre-implemented natural language processing models. In this case, noise contrastive estimation was used as the loss function, one negative sample was used, the window size was 5, and the learning rate was 0.16. Gensim and Mikolov et al.’s word2vec both use subsampling to improve the accuracy of the final representations that the other models in this experiment did not include.

Experiment 2a: Effects of CI on EWC model in an Artificial Language TensorFlow was once again used to implement a bare-bones version of a predictive neural network with the addition of EWC. The original implementation of EWC uses a version of a neural network that computes the loss function using Stochastic Gradient Descent and a SoftMax output layer. While this version of a neural network is not common in natural language processing (because it is more computationally expensive), it was used for the artificial experiment in order to calculate the EWC loss. Therefore, the loss function used in this implementation was SoftMax cross entropy (however

it is possible to use noise contrastive estimation). This choice was made in order to be as close to the original EWC implementation as possible. The neural net Kirkpatrick et al. used did not include negative information or noise contrastive estimation.

The SoftMax loss function seems to be more affected by CI than noise contrastive estimation—possibly because of the inclusion of negative sampling information. Thus, the difference in average cosine similarity between the two sections can be explained by the difference in loss functions. The learning rate for this model was 0.02, the embedding size was 10, and the window size was 2.

Experiment 2b: Effects of CI on EWC model in a Natural Corpus Experiment 2 used the same bare-bones TensorFlow implementation used in Experiment 1 of the first section. Originally, we attempted to use the same implementation as Experiment 1 of the second section (with SoftMax and cross entropy) but this version was unable to scale up to natural language¹. To scale up to the natural language, we had to include negative information and noise contrastive estimation. In this case, because there is no subsampling method implemented—which prioritizes the learning of less frequent words—more negative samples and a higher learning rate had to be used to achieve cosine similarity levels approximately equal to the results produced by Gensim. We did not need to do this in the artificial language because the vocabulary was so small. We used 64 negative samples, a learning rate of 0.3, and Stochastic Gradient Descent as the optimization function. The number of negative samples and learning rate, while higher than the

---

¹ A "standard" neural network is usually trained with a *cross-entropy* cost function which requires the values of the output neurons to represent probabilities - which means that the output "scores" computed by the network for each class have to be normalized, converted into actual probabilities for each class. This normalization step is achieved by means of the **SoftMax** function. If a word learning model is trained via this method, then each word in the corpus must be represented by an output neuron. This is adequate when the vocab is only 7 words, but for natural language this set up was too costly—the **SoftMax** function is too computationally expensive.

Gensim implementation, are not uncommon parameter settings to use in the field. The embedding size for this model was 50 and the window size was 5.