CARLS: Cross-platform Asynchronous Representation Learning System

Chun-Ta Lu*, Yun Zeng*, Da-Cheng Juan*, Yicheng Fan, Zhe Li, Jan Dlabal, Yi-Ting Chen, Arjun Gopalan, Allan Heydon, Chun-Sung Ferng, Reah Miyara, Ariel Fuxman, Futang Peng, Zhen Li, Tom Duerig, Andrew Tomkins

Google Research
Mountain View, CA
{chunta,xzeng,dacheng,yichengfan,lzhe,dlabal,yitingchen,arjung,aheydon,csferng,reahm,afuxman,futangpeng,zhenli, tduerig,tomkins}@google.com

ABSTRACT

Innovation on learning systems has been a key propellant in advancing and scaling up state-of-the-art neural models. In this work, we propose CARLS, a novel framework for augmenting the capacity of existing deep learning frameworks by enabling multiple components—model trainers, knowledge makers and knowledge banks—to concertedly work together in an asynchronous fashion across hardware platforms. The proposed CARLS is particularly suitable for learning paradigms where model training benefits from additional knowledge inferred or discovered during training, such as node embeddings for graph neural networks or reliable pseudo labels from model predictions. We also describe three learning paradigms—semi-supervised learning, curriculum learning and multimodal learning—as examples that can be scaled up efficiently by CARLS. One version of CARLS has been open-sourced and available for download at:

github.com/tensorflow/neural-structured-learning/tree/master/research/carls.

1 INTRODUCTION

During the past decade, several deep learning frameworks [2, 10, 36] have been made available to public and accelerating the advancement of many research domains in deep learning [16, 18, 29]. Despite their successes, cutting-edge technology and emerging applications are always pushing the limits of training capacity supported by the existing frameworks. For example, the largest T5 model [40] with 11 billion parameters has achieved state-of-the-art performance on multiple NLP benchmarks, and a near-human score on the SuperGLUE natural language understanding benchmark. GPT-3 [5], an auto-regressive language model with 175 billion parameters, achieves strong performance on many NLP benchmarks without any gradient updates or fine-tuning. Most recently, the switch transformer [17] advanced the size of the T5 model to 7× times, leading to a giant model speaking 102 different languages.

While most of the above efforts focus on training one monolithic model, human brains often solve a cognition problem using multiple pathways with multi-sensory inputs and multiple models [27]. Such a multimodal framework has already shown its potential in boosting the performance over the existing monolithic models (e.g., [23, 32, 37]) and in solving novel problems [19]. For
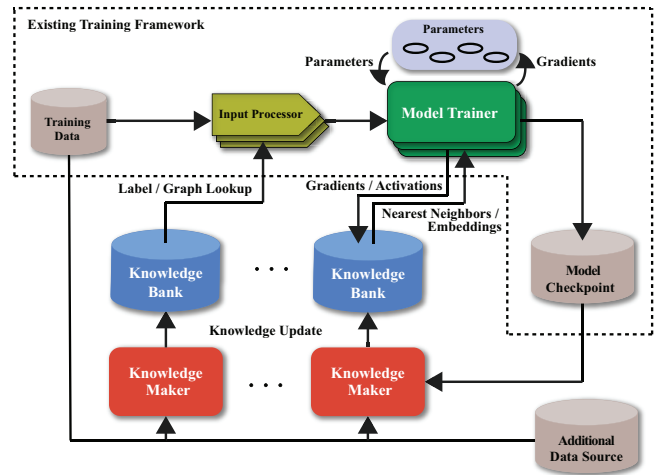

Figure 1: System Overview. CARLS employs three major components: Model Trainer, Knowledge Maker and Knowledge Bank. Models Trainers are the main jobs training and updating model parameters, and Knowledge Makers are running in parallel to compute auxiliary information required for gradient calculation and store the information in Knowledge Bank. Note that these components can be deployed on different platforms (GPUs, TPUs, etc) and communicate with each other in an asynchronous fashion.

example, for vision-language related problems such as visual question answering [3], caption-based image retrieval [50], and visual common sense reasoning [51], the vision and language modalities are encoded by its own designated encoder and the two encoders are co-trained jointly. However, implementing these multimodal learning tasks at scale imposes unique challenges, since the computational complexity usually grows with the number of encoders trained, and the advancement of hardware alone cannot keep up with the fast growth of model (or data) size and complexity. Such a scaling issue manifests itself in other machine learning paradigms such as graph-based neural networks [6, 28, 47], where each node in a graph is processed by a kernel that can be a simple multilayer perceptron or a complex attention model like BERT [15]. Furthermore, the computational cost grows linearly with the size of neighborhoods if no sampling technique is used.

---

* These authors contributed equally.

In this work, we propose *CARLS*: a **C**ross-platform, **A**synchronous **R**epresentation **L**earning **S**ystem that addresses the challenges in scaling up these "giant" models (*e.g.*, multimodal learning or graph neural networks). CARLS is able to (a) accommodate training of very large models with extremely high computational complexity—*e.g.*, graph neural networks with a graph of more than ten billion nodes, each node encoded by a BERT kernel—and (b) accelerate the training by shifting partial workload of inferring information required for gradient calculation from model trainers to "Knowledge Makers" deployed in CARLS.

CARLS contains three major components that operate in parallel: *Model Trainer*, *Knowledge Maker* and *Knowledge Bank*. Figure 1 illustrates an overview of the proposed CARLS system. Model trainers are the main training jobs with an extended functionality to retrieve auxiliary information generated by the knowledge makers that are usually implemented by a fleet of servers or pipelines. The model trainers and knowledge makers share the information via the knowledge bank, which stores the results generated by knowledge makers and enables efficient, asynchronous training. Model trainers and knowledge makers may be deployed using different hardware platforms (*e.g.*, CPUs, GPUs or TPUs). Furthermore, model trainers and knowledge makers are communicating in an asynchronous fashion, meaning they operate independently without affecting the training speed.

The key observation to design CARLS for large-scale, efficient learning is that the compute can be parallelized at model-level (or logic-level) for many neural architectures; in other words, the interactions among multiple models are usually limited to a few "connecting points" such as the joint of each encoder output, or the aggregation of neighborhood embeddings in a graph neural network. By employing the knowledge banks as the interface between individual modules, CARLS delegates the computation cost of the individual models from the model trainer to the knowledge makers. One potential issue of such an asynchronous mechanism is data freshness—some knowledge makers may generate results based on slightly outdated information. In practice, we find the impacts of such an issue are controllable and not significant. In summary, the main advantages of the proposed CARLS are two-fold:

- **Scalability**: CARLS enables efficient operations to store and lookup intermediate knowledge for training complex models. For each training step, the intermediate knowledge can be fetched from a remote knowledge bank, rather than being recomputed across multiple training batches.
- **Flexibility**: CARLS allows constructing and updating the knowledge dynamically based on the current model state at each training step, as opposed to information preprocessed from a static input dataset or by a pretrained model.

The proposed CARLS system is validated in a number of learning tasks which are considered to be very challenging or even infeasible before CARLS. For example, training a graph neural network (GNN) stacked on top of a dense encoder (*e.g.*, ResNet [21] or Transformer [46]), where each node features are encoded by the encoder and aggregated by a GNN, the training time becomes prohibitively expensive especially when the neighbor size grows large. With CARLS, we are able to train a graph-regularized model whose neighbor size is 10 times larger than the largest one currently proposed by the literature [25], without introducing any slowdown in the training speed. We also demonstrate three learning paradigms as examples that can be scaled up by CARLS, including semi-supervised learning, curriculum learning, and multimodal learning.

The remainder of this paper is organized as follows: after reviewing related work in Section 2, we elaborate the design of CARLS in Section 3. In Section 4, we demonstrate the learning paradigms that can be scaled up efficiently by CARLs. Finally, we conclude this paper in Section 5.

2 RELATED WORK

Large-scale learning systems (*e.g.*, models with more than 10G parameters) are often needed in language/image modeling, recommendation systems, or models with very large sparse feature set input (*e.g.*, click through rate prediction models in digital advertising [54, 55]). Directly relying on the design of existing deep learning libraries (*e.g.*, [2, 31]) is often limited in practice and challenges exist both in training and serving these large models.

Existing large language/image models are characterized by their sophisticated internal building blocks (*e.g.*, the self-attention mechanism [46] or the residual network [21]) that are often stacked up to hundreds of layers (*e.g.*, [5, 14, 38, 39]). Solely relying on the innovations of hardware (*e.g.*, GPU or TPU) can hardly keep up with the explosion of model size. To efficiently train these models, parallelisms/sparsity among the partitions of the model are often explored (*e.g.*, [30, 35]). The mixture-of-experts (MoE) model is an example model with sparsity that allows different sub-models to be trained in parallel [35, 41]. Most recently, the switch transformer [17] is able to handle models with 1T parameters by solving a big MoE problem. In order to serve these large models in small devices or more efficiently, the distillation technique [22] or teacher-student networks are usually applied to compress the model.

Among industrial scale recommendation systems, the number of unique items (*e.g.*, keywords, products) can easily reach up to billions or even trillions, which are usually represented as the softmax/logistic output layer of a model. Since these output layers can be regarded as nearest neighbor search in an embedding space (one embedding for each item), a popular approach to efficient model training and inference is to compress the output embeddings via various hashing techniques [34, 42, 44, 53]. However, these compression/approximation approaches often come with the price of sacrificing modeling accuracy.

There are also more recent interests in decoupling the memory part (*e.g.*, sparse feature embeddings for both input and output layers) of a model from its structural part (*e.g.*, [54]), therefore achieving efficient model training/inference, automatic model growth, as well as efficient multi-task learning. For example, the DynamicEmbedding system [52] employs a *memory server* for serving the embedding data in a model with almost unlimited sparse feature set; The Taskology system [33] employs an *inference server* to facilitate multi-task training with consistency constraints. In this work, we further extend this line of research by proposing a highly distributed system with both memory and inference components.

3 PROPOSED CARLS

In this section, we provide the details of the proposed CARLS framework, which consists of three major components that enable asynchronous representation learning across heterogeneous platforms:

- **Model Trainer**: responsible for both training various models and communicating with the knowledge bank to fetch augmented information to facilitate the training.
- **Knowledge Maker**: responsible for generating the knowledge requested by the model trainer. Specifically, a knowledge maker loads the latest checkpoint (saved by the model trainer) to make inference on auxiliary data, or discover new neighborhoods from examples with close representations learnt from the model.
- **Knowledge Bank**: responsible for storing and refreshing the knowledge generated by knowledge makers. A knowledge bank also serves the fetch requests from model trainers.

Two key characteristics of CARLS are that (a) Model Trainers, Knowledge Banks, and Knowledge Makers can be operated asynchronously, and (b) these components can be deployed on heterogeneous hardware platforms (CPUs, GPUs, TPUs, etc). In the following sections, we describe the features that have been implemented, and leave a wide range of scenarios that are not yet implemented but accommodated by CARLS as a straight-line future work.

3.1 Knowledge Maker

Knowledge makers play a critical role in extending the capacity of a learning system. Conventionally, the model trainers process the data and compute all the information required for calculating losses. In CARLS, knowledge makers are designed to dynamically calculate (or discover) the knowledge that can be exploited by model trainers for calculating the losses. Therefore, part of model trainers' workload can be shifted to knowledge makers, which helps the model trainers run faster or save capacity to accommodate larger models or training with more data. Below we provide three example types of knowledge that are suitable to be processed or discovered by knowledge makers:

- **Graph structure and node embedding**: knowledge makers can load parts of the model—such as the node encoder of graph neural network—to update the knowledge. The graph structure can also be dynamically updated with the similarity between the computed node embeddings, as opposed to a given static graph (Section 4.1).
- **Augmented labels**: knowledge makers can load the whole model to (a) make inference on the missing labels for augmenting the training data, or (b) clean up noisy labels for providing a cleaner training dataset (Section 4.2).
- **External Knowledge**: another advanced application of CARLS is to integrate the external knowledge/memeory for training [19, 48]. In this case, the responsibility of knowledge makers is to pre-compute and process the related knowledge for enabling fast serving by the knowledge bank (Section 4.3)

Knowledge makers keep the same machine states as model trainers by periodically loading the parameters from the latest checkpoints or parameter servers. This way, the inferences made by the knowledge makers will be consistent with model trainers. Knowledge makers can be implemented by a MapReduce [13] or Flume [7] pipeline, or by sending remote procedure calls (RPCs) to a remote inference server.

3.2 Knowledge Bank

The knowledge inferred by knowledge makers are stored in knowledge banks for fast serving model-training jobs. Depending on the application type, a knowledge bank can be an external storage system—such as Bigtable [8] or Spanner [11]—or a specialized embedding server [52] responsible for updating internal parameters during back-propagation stage. Some applications may require knowledge banks perform certain computations, such as retrieving top-N nearest neighbors given a query. Note that the number of instances stored in knowledge banks can be very large (multi-millions to multi-billions). To keep the computational latency constant—not growing as the data size grows, the knowledge banks are sharded and deployed in a distributed fashion [45].

Below is a list of example applications enabled by knowledge banks:

- **Feature Lookup**: an instance's features (*e.g.*, neighbor IDs from a graph, or labels) are stored as a protocol buffer [1] and keyed by the instance's unique ID. In this application, knowledge banks can be implemented directly by external storage systems without special treatments to store these feature information.
- **Embedding Lookup and Update**: in certain applications—*e.g.*, node embeddings in the graph neural network or word embeddings in language modeling, the embeddings can be stored in a knowledge bank and will be continuously updated throughout the training process. To support back-propagate the gradients with respect to the embeddings, we leverage the DynamicEmbedding [52] (as the implementation of knowledge banks) to update the embeddings with the gradient values.
- **Nearest Neighbors Lookup**: efficient nearest neighbor search in the embedding space is important for solving large scale classification and retrieval tasks [20]. Unlike existing learning paradigms that either take the pre-computed nearest neighbors as input features or search over a very limited scope (*e.g.*, within the same batch), CARLS enables searching over the embeddings kept in the knowledge bank, which is essentially the entire dataset—both labeled and unlabeled samples. Furthermore, using CARLS enables nearest-neighbor search over the dataset that grows or is updated during training. To facilitate efficient nearest neighbors search as the dataset grows, the computation is distributed into multiple shards and ScaNN [20] can be applied for search space pruning and quantization.

Lazy update for asynchronous gradient update. While data consistency, isolation, and durability are delegated to the design of existing storage systems [8, 11] and computation efficiency and parallelism can be handled by existing server designs [52], special

treatment is still required when multiple training jobs are attempting to update the gradients of the same embedding entry stored in knowledge bank. In this case, simply guaranteeing atomicity may not be sufficient since this mechanism favors the last model that updates the gradients and ignores the contribution from other models. CARLS resolves this issue by employing a lazy update scheme: caching the results of gradient update until the next lookup request arrives, or an expiration time is reached. The update is based on the average of all cached gradients with possible outlier detection. With this lazy update mechanism, the overall training process is more stable compared with simple stochastic gradient descent.

3.3   Model Trainer

A model trainer is typically a machine-learning model equipped with a communication module to fetch the knowledge (of inferences) from knowledge banks. The communication module is usually implemented as customized operations or layers, enabling the funtionalities as listed in Section 3.2. The communication module sends the request to the knowledge bank for retrieving additional information discovered during training. The inputs of communication modules can come from the training data (*e.*g., sample IDs for fetching neighbor IDs), or the embedding values at a hidden layer (for retrieving the nearest neighbors). Note that model trainers can be implemented with various platforms (*e.*g., TensorFlow or Pytorch) in a distributed fashion. Special cases need to be taken for the communication module when the trainers are running on TPU machines [24], which works best when the entire model fits into the on-chip memory.

4   LEARNING PARADIGMS

In the section, we provide three learning paradigms as examples that can be scaled up efficiently by CARLS, including but not limited to semi-supervised learning, curriculum learning, and multimodal learning.

4.1   Semi-Supervised Learning

Semi-supervised learning (SSL) is a machine learning paradigm that harnesses both labeled and unlabeled data to improve model prediction performance, and is particularly useful when labeled data is limited or expensive to obtain. However, despite the advantages and successes brought by SSL, using abundant unlabeled data also significantly increases the training computation. For example, graph-based methods are one type of popular SSL approaches, where each node's information (*e.*g., labels) can be iteratively updated by aggregating the information from neighboring nodes whether they are labeled or unlabeled [6, 25]. This means, in addition to computing or processing the labeled data, additional computation is required to process these neighboring nodes. In this case, CARLS deploy a fleet of knowledge makers to process these neighboring nodes, shifting the additional computation conventionally done by the model trainer to knowledge makers. Furthermore, CARLS eliminates redundant computations (*e.*g., one node being a common neighbor among several other nodes only needs to be processed once) by caching the results in the knowledge banks.
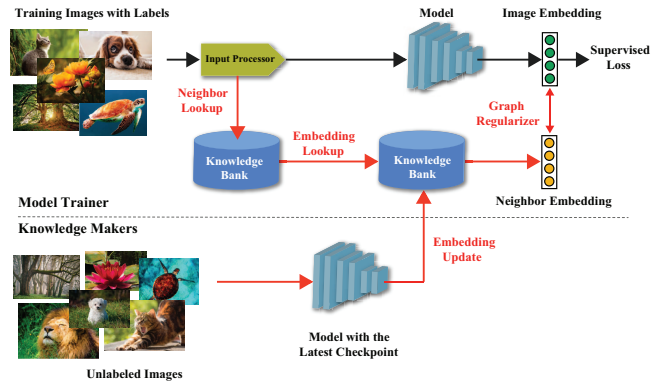

Figure 2: Training a graph-regularized model at scale with CARLS. Flow in red is added to enable graph regularization and required only during training. The model trainer and knowledge makers are running in parallel with a shared knowledge bank (KB). Knowledge makers keep loading the latest checkpoint generated by model trainers for making inferences and update the embeddings in KB. The input processor of the model trainer will look up training instances' neighborhood from the KB. After retrieving the neighbor info (e.g., neighbor IDs, edge weights), the neighbor IDs will be used to lookup neighbors' embeddings from the KB.

Figure 2 illustrates a graph-regularized model trained with CARLS. The training objective is to minimize the sum of the supervised loss and the graph regularizer, which is a pairwise distance between the embeddings of neighboring nodes. The neighbourhood can be constructed offline using existing signals (e.g., co-clicked image pairs [25] or images in the same board [49]), or created online during training (e.g., image augmentation [12]), and looked up from the knowledge bank in the trainer.

If the neighbor embeddings are directly computed in the trainer, the training time and the trainer memory consumption will increase linearly in proportional to the number of neighbors [25]. To make the graph-regularized model scalable, the embeddings of each node/image are computed by knowledge makers and updated to the knowledge bank in parallel with the model trainer.

Figure 3 illustrates another example of applying CARLS to enable efficient graph operations for large-scale graph learning. In this example, the model is a stack of an encoder (*e.*g.ResNet or Transformer) and a graph neural network (GNN). It can be very challenging to train such a model, especially when the size of the subgraph is large, without the support of CARLS.

4.2   Curriculum Learning

Curriculum learning [4] refers to a learning paradigm where the challenge of a learning task gradually increases over time for improving the model performance. The increased challenge usually forces models to focus on learning additional information that is not well captured from the previous iterations [43]. This paradigm maps favorably to the CARLS framework where the additional information can be stored in the knowledge bank and the knowledge makers are responsible for generating this information required to
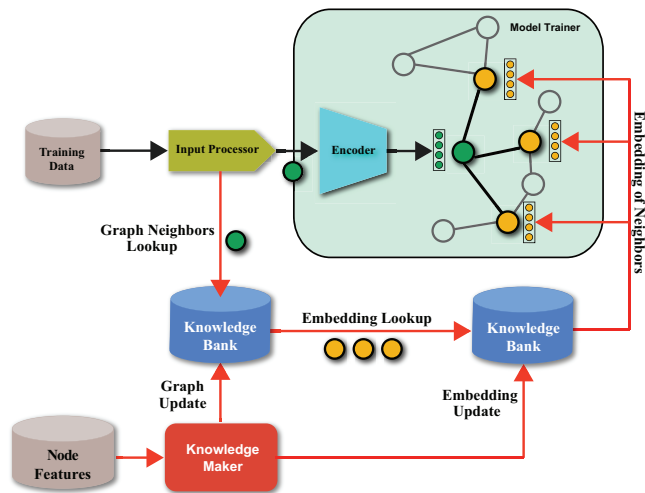
Figure 3: An illustration of the training pipeline of CARLS for a graph neural network stacked on top of a node encoder. The node encoder encodes the node features into embeddings and feed to the GNN. In the model trainer, each training input node can look up its sub-graph and the embeddings of each node in the subgraph from the KB. The knowledge makers can load the latest checkpoint to compute and update the node embeddings as well as the graph (based on the node embeddings) to the KB.
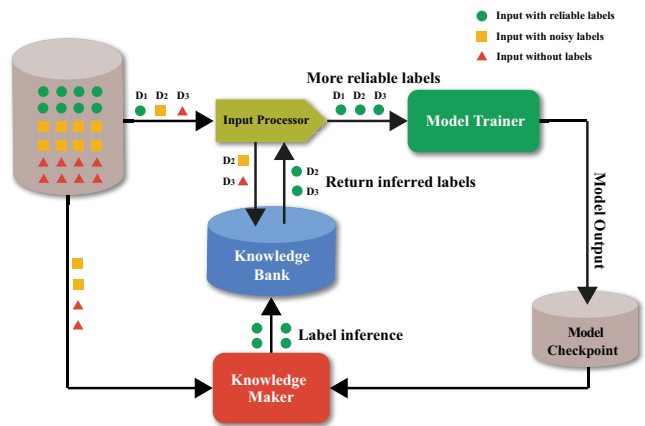

Figure 4: An illustration of the workflow of curriculum learning using CARLS. The labels of the input examples can be improved by the knowledge maker (*e.g.*, using online label mining for unreliable labels or graph agreement model for inferring missing labels).

increase the training complexity. Fig. 4 illustrates the workflow of using CARLS for curriculum learning. Specifically, we demonstrate two examples: *online label mining* deals with noisy labels, and the *graph agreement model* deals with missing labels for some inputs.

4.2.1 Online label mining. In large scale applications of image classification, the labels of each image are usually inferred from
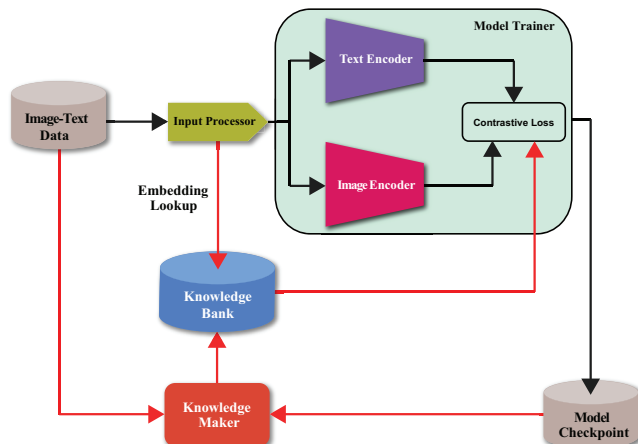

Figure 5: An illustration of the training pipeline of CARLS for a deep image-text two-tower model. Flow in red is added to enable lookup image/text embeddings and is required only during training. The distributed knowledge makers will compute the image/text embeddings and update the embeddings to the knowledge bank. The model trainer can look up the embeddings from the KB instead of computing the image/text embeddings through the image/text encoder.

a rather noisy input (*e.g.*, search queries or text captions associated to an image) that are often inconsistent with each other (*e.g.*, wrong or missing labels). In such cases, multiple iterations are often needed to alternate between i) training the model based on noisy data and ii) refining the model based on the latest trained model. Such a process can be significantly accelerated if we can do i) and ii) in parallel. In CARLS, i) can be achieved via a model trainer and ii) by knowledge makers, and the knowledge bank stores the newest inferred labeled for each instance.

4.2.2 Graph Agreement Model. Another example of curriculum learning is the graph agreement model for semi-supervised learning [43], where the labels of unknown training examples are inferred based on the nearest neighbors from labeled examples, where the nearest neighbors are calculated based on the embedding of the input (from the hidden layer of the trained model). In CARLS, this nearest neighbors search can be done by a knowledge maker and the inferred labels can be stored in the knowledge bank.

4.3 Multimodal Learning

In addition to training a single deep model, CARLS can also be used for multimodal learning such as training an image-text two-tower model [23] and the modality nets [26] that has four models for language (text data), images, audio, and categorical data.

Figure 5 illustrates an example of a image-text deep two-tower model trained with CARLS, where the image and text encoders can be learned via a contrastive loss [9] that pushes the embeddings of matched image-text pair together while pushing those of non-matched image-text pair apart. The image-text matching scores are computed as the cosine-similarity of the image embeddings and the text embeddings from the corresponding encoders. Although

when the image/text encoder goes deeper it can often achieve better performance [23], it is more challenging to train as the computational complexity grows significantly. Instead of computing the image/text embeddings directly in the model trainer, we can utilize CARLS to lookup historical embeddings from the knowledge bank[1].

In addition to reducing the computation in the trainer, CARLS can also be used to improve the model performance. As shown in [23] that the model performance increases along with the number of random negatives used in the contrastive learning. As the embeddings of the random negatives can be looked up from the knowledge bank, we can easily scale up the number of random negatives. Besides, because each encoder can be considered as an independent model in the knowledge makers (by loading parts of the network from the checkpoint of the model), the image/text embeddings can be computed and updated to the knowledge bank independently. Thus, image/text data feed to the knowledge makers can go beyond the training image-text pairs.

5 CONCLUSION

In this paper, we proposed CARLS, a cross-platform, asynchronous learning framework that significantly extends the capacity of current deep learning frameworks by enabling dynamic knowledge sharing between model trainers and knowledge makers. The proposed CARLS is particularly suitable for training large-scale neural models that benefit from additional knowledge inferred or discovered during training, such as graph neural networks and multi-modality models. CARLS also enables existing frameworks to dynamically learn with changes inferred or occurred during training, facilitating curriculum learning to be more efficient. As a straight-line future work, we will extend CARLS to train more variety of learning models, design new APIs to streamline system setup and provide better support for heterogeneous hardware.

REFERENCES

[1] 2016. Protocol Buffers Language Guide.
[2] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: a system for large-scale machine learning.. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, Vol. 16. 265–283.
[3] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*. 2425–2433.
[4] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*. 41–48.
[5] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165* (2020).
[6] Thang D Bui, Sujith Ravi, and Vivek Ramavajjala. 2018. Neural Graph Learning: Training Neural Networks Using Graphs. In *ACM International Conference on Web Search and Data Mining*.
[7] Craig Chambers, Ashish Raniwala, Frances Perry, Stephen Adams, Robert R Henry, Robert Bradshaw, and Nathan Weizenbaum. 2010. FlumeJava: easy, efficient data-parallel pipelines. *ACM Sigplan Notices* 45, 6 (2010), 363–375.
[8] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. 2006. Bigtable: A Distributed Storage System for Structured Data. In *7th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 205–218.
[9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
[10] Ronan Collobert, Samy Bengio, and Johnny Mariéthoz. 2002. *Torch: a modular machine learning software library*. Technical Report. Idiap.
[11] James C. Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, JJ Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alexander Lloyd, Sergey Melnik, David Mwaura, David Nagle, Sean Quinlan, Rajesh Rao, Lindsay Rolig, Dale Woodford, Yasushi Saito, Christopher Taylor, Michal Szymaniak, and Ruth Wang. 2012. Spanner: Google's Globally-Distributed Database. In *OSDI*.
[12] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. 2019. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 113–123.
[13] Jeffrey Dean and Sanjay Ghemawat. 2004. MapReduce: Simplified Data Processing on Large Clusters. In *OSDI'04: Sixth Symposium on Operating System Design and Implementation*. San Francisco, CA, 137–150.
[14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR* abs/1810.04805 (2018).
[15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
[16] Andre Esteva, Alexandre Robicquet, Bharath Ramsundar, Volodymyr Kuleshov, Mark DePristo, Katherine Chou, Claire Cui, Greg Corrado, Sebastian Thrun, and Jeff Dean. 2019. A guide to deep learning in healthcare. *Nature medicine* 25, 1 (2019), 24–29.
[17] William Fedus, Barret Zoph, and Noam Shazeer. 2021. Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity. *arXiv preprint arXiv:2101.03961* (2021).
[18] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. 2016. *Deep learning*. Vol. 1. MIT press Cambridge.
[19] Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al. 2016. Hybrid computing using a neural network with dynamic external memory. *Nature* 538, 7626 (2016), 471–476.
[20] Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. 2020. Accelerating Large-Scale Inference with Anisotropic Vector Quantization. In *International Conference on Machine Learning*. https://arxiv.org/abs/1908.10396
[21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep Residual Learning for Image Recognition. In *Computer Vision and Pattern Recognition (CVPR)*.
[22] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
[23] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. arXiv:cs.CV/2102.05918
[24] Norman P Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, et al. 2017. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th annual international symposium on computer architecture*. 1–12.
[25] Da-Cheng Juan, Chun-Ta Lu, Zhen Li, Futang Peng, Aleksei Timofeev, Yi-Ting Chen, Yaxi Gao, Tom Duerig, Andrew Tomkins, and Sujith Ravi. 2020. Ultra Fine-Grained Image Semantic Embedding. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 277–285.
[26] Lukasz Kaiser, Aidan N Gomez, Noam Shazeer, Ashish Vaswani, Niki Parmar, Llion Jones, and Jakob Uszkoreit. 2017. One model to learn them all. *arXiv preprint arXiv:1706.05137* (2017).
[27] Eric R Kandel, James H Schwartz, Thomas M Jessell, Department of Biochemistry, Molecular Biophysics Thomas Jessell, Steven Siegelbaum, and AJ Hudspeth. 2000. *Principles of neural science*. Vol. 4. McGraw-hill New York.
[28] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
[29] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep Learning. *Nature* 521 (2015), 436–444.
[30] Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2020. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668* (2020).
[31] Mu Li, David G Andersen, Jun Woo Park, Alexander J Smola, Amr Ahmed, Vanja Josifovski, James Long, Eugene J Shekita, and Bor-Yiing Su. 2014. Scaling distributed machine learning with the parameter server. In *11th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 14)*. 583–598.

---

[1]In order to back-propagate the gradients back to the encoders, a subset of the examples in a training mini-batch should still be computed through the encoders in the model trainer.

[32] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 2020. 12-in-1: Multi-task vision and language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10437–10446.

[33] Yao Lu, Sören Pirk, Jan Dlabal, Anthony Brohan, Ankita Pasad, Zhao Chen, Vincent Casser, Anelia Angelova, and Ariel Gordon. 2020. Taskology: Utilizing Task Relations at Scale. *arXiv preprint arXiv:2005.07289* (2020).

[34] Tharun Medini, Qixuan Huang, Yiqiu Wang, Vijai Mohan, and Anshumali Shrivastava. 2019. Extreme classification in log memory using count-min sketch: A case study of amazon search with 50m products. *arXiv preprint arXiv:1910.13830* (2019).

[35] Deepak Narayanan, Aaron Harlap, Amar Phanishayee, Vivek Seshadri, Nikhil R Devanur, Gregory R Ganger, Phillip B Gibbons, and Matei Zaharia. 2019. PipeDream: generalized pipeline parallelism for DNN training. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*. 1–15.

[36] PyTorch. 2018. http://pytorch.org http://pytorch.org.

[37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. [n. d.]. Learning Transferable Visual Models From Natural Language Supervision. *Image* 2 ([n. d.]), T2.

[38] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. (2018).

[39] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.

[40] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* 21, 140 (2020), 1–67. http://jmlr.org/papers/v21/20-074.html

[41] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538* (2017).

[42] Liuyihan Song, Pan Pan, Kang Zhao, Hao Yang, Yiming Chen, Yingya Zhang, Yinghui Xu, and Rong Jin. 2020. Large-Scale Training System for 100-Million Classification at Alibaba. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2909–2930.

[43] Otilia Stretcu, Krishnamurthy Viswanathan, Dana Movshovitz-Attias, Anthony Platanios, Sujith Ravi, and Andrew Tomkins. 2019. Graph agreement models for semi-supervised learning. (2019).

[44] Yukihiro Tagami. 2017. Annexml: Approximate nearest neighbor search for extreme multi-label classification. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. 455–464.

[45] Maarten van Steen and Andrew S. Tanenbaum. 2017. *Distributed Systems*. CreateSpace Independent Publishing Platform.

[46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems (NIPS)*. 5998–6008.

[47] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).

[48] Jason Weston, Sumit Chopra, and Antoine Bordes. 2014. Memory networks. *arXiv preprint arXiv:1410.3916* (2014).

[49] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 974–983.

[50] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics* 2 (2014), 67–78.

[51] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6720–6731.

[52] Yun Zeng, Siqi Zuo, and Dongcai Shen. 2020. DynamicEmbedding: Extending TensorFlow for Colossal-Scale Applications. *arXiv preprint arXiv:2004.08366* (2020).

[53] Xingcheng Zhang, Lei Yang, Junjie Yan, and Dahua Lin. 2018. Accelerated training for massive classification via dynamic class selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.

[54] Weijie Zhao, Jingyuan Zhang, Deping Xie, Yulei Qian, Ronglai Jia, and Ping Li. 2019. AIBox: CTR prediction model training on a single node. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 319–328.

[55] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1059–1068.