# PyTorchPipe: a framework for rapid prototyping of pipelines combining language and vision

Tomasz Kornuta\* IBM Research AI Almaden Research Center San Jose, CA 15120 tkornuta@gmail.com

## Abstract

Access to vast amounts of data along with affordable computational power stimulated the reincarnation of neural networks. The progress could not be achieved without adequate software tools, lowering the entry bar for the next generations of researchers and developers. The paper introduces PyTorchPipe (PTP), a framework built on top of PyTorch. Answering the recent needs and trends in machine learning, PTP facilitates building and training of complex, multi-modal models combining language and vision (but is not limited to those two modalities). At its core, PTP employs a component-oriented approach and relies on the concept of a pipeline, defined as a directed acyclic graph of loosely coupled components. A user defines a pipeline using yaml-based (thus human-readable) configuration files, whereas PTP provides generic workers for their loading, training, and testing using all the computational power (CPUs and GPUs) that is available to him. The paper covers the main concepts of PyTorchPipe, discusses its key features and briefly presents the currently implemented tasks, models and components.

# **1** Introduction

Deep learning [LBH15] enabled impressive improvement in performance across many problem domains, and neural network based models became dominating solutions, achieving state of the art performances, e.g., in speech recognition [GMH13], image classification [KSH12], object detection [RDGF16], instance segmentation [HGDG17], question answering [WCB14] or machine translation [BCB14]. This would not be possible without (1) access to large-scale datasets, (2) harnessing the computational power and parallel processing of GPUs and (3) the appropriate software. Starting from low-level libraries such as CUDA [SK10], the AI community has been developing tools and libraries offering more and more levels of abstractions above the multiply-accumulate operations and elementary tensor algebra, e.g. Torch [CBM02], Theano [BLP+12], Chainer [TOHC15] or TensorFlow [ABC<sup>+</sup>16]. Those API based tools facilitate building and training of deeper and more complex neural models and significantly lower the entry bar for new users and researchers. As the community started to shift towards multi-problem suites such as MS COCO [LMB<sup>+</sup>14], being a large-scale object detection, segmentation, and captioning dataset, or ParlAI [MFF<sup>+</sup>17] and AllenNLP [GGN<sup>+</sup>17], offering wide varieties of linguistic tasks, researchers started to feel the need for more specialized APIs that will enable comparison of different models under exactly the same settings and target the so-called *reproducibility crisis* [Hut18], with examples being object detection APIs such as Detectron [ $GRG^{+18}$ ] for PyTorch [ $PGC^{+17}$ ] or object detection API in TensorFlow [HRS<sup>+</sup>17]. The need for more abstraction from both problem domain (dataset) and model specialized for a given domain led to development of the next generation of tools, such as

<sup>\*</sup>Work done at IBM Research. Currently with NVIDIA.

Tensor2Tensor [VBB<sup>+</sup>18], MI-Prometheus [KMM<sup>+</sup>18], Pytia [SNJ<sup>+</sup>18], OpenSeq2Seq [KGG<sup>+</sup>18], NeMo [KLN<sup>+</sup>19] or Ludwig [Mol19]. Aside of the mentioned modularity and abstraction, all those tools offer powerful scripts, agnostic to dataset and model, that standardize training (and testing) procedures, once again lowering the entry bar and enabling the user to focus more on the model rather than on the training loop, validation procedures, utilization of GPUs etc.



Figure 1: Exemplary multi-modal flow diagram. In PTP each block is implemented as a separate component, whereas framework handles passing data between them and controls the execution order.

As the individual problem domains matured and saturated with diverse solutions, a new strong trend called Multi-Modal Machine Learning [BAM18] emerged, with tasks combining vision and language [MKK19] such as Image Captioning [KFF15] and Image/Visual Question Answering (VQA) [MF14, AAL+15] becoming slightly more popular than the others. A clearly different nature of those two modalities (spatial vs sequential) implies diverse transformation that the inputs need to undergo, e.g. convolutional layers in the case of images and tokenization, indexing, word embeddings etc. in the case of questions. As a result, the most general VQA architecture consists of four major modules: two encoders responsible for encoding of raw image and question into more useful representations, a reasoning module that combines them and decoder that produces the answer. In early VQA systems, reasoning modules implemented diverse multi-modal fusion mechanisms [MD18], from concatenation to pooling, e.g. MCB [FPY<sup>+</sup>16] and MLB [KOL<sup>+</sup>17], to diverse (co-)attention mechanisms, e.g. question-driven attention over image features [KE17]. More recently, researchers have focused on reasoning mechanisms such as Relational Networks [SRB+17, DCK18] and Memory, Attention and Composition (MAC) networks [HM18, MJA<sup>+</sup>18]. Still, they all can fit into the *architectural pattern* defined above. Similarly, one can assume adequate architectural patterns for Image Captioning or any other task, e.g. Visual Dialog [DKG<sup>+</sup>17] or Visual Image Generation [MKK19]. Such assumptions enable development of software facilitating their creation, for example, OpenSeq2Seq assumes model to follow the *Encoder(s)-Decoder* architecture, Ludwig imposes it to be composed of *Encoder(s)-(Combiner)-Decoder*, whereas Pytia assumes that user will create a *monolithic* model, but instead standardizes interfaces by enforcing that batches fetched by all datasets (Tasks) would have exactly the same signature consisting of (*image, context, text, output*).

However, in practice, some additional computations that need to be done. When looking at an exemplary VQA system presented in Fig. 1, one might notice that answers also undergo some transformations (from words to indices/one-hot encoding) and aside of loss there are some additional modules calculating statistics. Besides, for the *test split* the answers are typically not provided. Hence some of the modules should not be executed at all when working with that split. Finally, it is a common practice [JNC<sup>+</sup>18, TAHvdH18] to use transfer learning [PY09] and incorporate into the model image encoders and/or word embeddings pre-trained a priori on external datasets. Moreover,

typical approach during training is to *freeze* the encoder(s) weights at first (which enables faster training), but at the end to *unfreeze* them and fine-tune the whole model jointly (which gives a slight accuracy boost). Those observations lead to relaxation of assumptions regarding utilization of fixed architectural patterns and formulated the requirement for flexible decomposition of models into smaller modules that laid the core foundation for PyTorchPipe.

# 2 PyTorchPipe

PyTorchPipe (PTP)<sup>2</sup>, being the main contribution of this paper, is a software framework build on top of PyTorch [PGC<sup>+</sup>17]. PTP is a component-oriented framework that facilitates development and rapid prototyping of computational multi-modal pipelines and comparison of diverse neural network-based models. The most important feature of PTP is the **decomposition of complex "monolithic" models into graphs with many inter-connected computational nodes**, that increases their reusability while still enabling their joint training. As a consequence, PyTorchPipe enables:

- Plugging in/out "modules" (components) at run-time<sup>3</sup>,
- Importing the pre-trained models (or their "parts") and saving them during/after training,
- Freezing/unfreezing of the models (or their "parts") on demand (at run-time),
- Run-time parametrization of all "modules",
- Pipeline-agnostic, generic scripts for training and testing, enabling run-time parametrization (hyper-parameters) and utilization of many CPUs/GPUs on-demand.

Additionally, PTP offers several other functionalities useful for rapid prototyping and training of new models, such as: different optimizers, sampling methods, automated logging and statistics collection facilities, the export of statistics to files (and TensorBoard) and automated saving of the best performing pipelines/models during training. Moreover, PTP provides out-of-a-box library of diverse, parameterizable components that can be used for rapid creation of various prototypes.

Experiments	Configuration file(s) Experiment Logs and Statistics Dataset(s) Checkpoint(s)
Components	Tasks
Workers	Off-line Trainer Processor
Framework Core	Task Manager       Sampler Factory       Logging Facility       Naming Facility       Statistics Facilities         Configuration Registry       Global Parameters       Pipeline Manager       Component Factory
Ext. Component Libraries	TorchVision TorchText PIL NLTK Matplotlib
Ext. Core Libraries	PyTorch NumPy TensorBoardX CUDA

Figure 2: Architecture of the PyTorchPipe framework.

From the architectural point of view, PTP can be seen as a stack of layers, as presented in Fig. 2. The two bottom layers are external libraries that PTP and its components depend on. The **Framework Core** layer contains all the tools and utilities realizing core functionalities, i.e. loading configuration files into registry, managing registry and global parameters, building the pipeline based on loaded configuration, component factories, loggers, statistics collectors and aggregators. **Workers** later use those for building and deploying the actual pipelines and execution of experiments. The **Components** layer contains the implementations of components along with their default configurations that are

<sup>&</sup>lt;sup>2</sup>https://github.com/ibm/pytorchpipe

<sup>&</sup>lt;sup>3</sup>Run-time is defined here as the moment of execution of the experiment.

loaded and managed by the workers. Finally, the **Experiments** layer contains files associated with a given experiment that the user wants to perform, staring from the configuration file(s), all the logs and statistic collected and saved to files during experiment execution, files associated with a given dataset, loaded and saved checkpoint(s) with weights of the models incorporated in the pipeline etc. Please note that the layer architecture reflects three types of user roles classified by their expertise, from top:

- Experiment runners might have little to no knowledge of the underlying models, training procedures, methods (backpropagation) etc. but are able to define experiment by writing configuration file(s) incorporating the existing components and relying on existing workers<sup>4</sup>.
- Component developers must understand how a given model works, however, do not need to
  understand every aspect of the training/testing procedures nor how PTP operates internally.
- Worker developers must understand both how models and training/testing work, as well as be familiar with internal operation of PTP core, how pipelines are assembled etc.

### 2.1 Pipelines

What the previously mentioned tools typically define as a *Model*, in PTP is framed as a **Pipeline**: a directed acyclic graph (DAG) consisting of many inter-connected components. The components are loosely coupled and care only about the Input Streams they read from and Output Streams they publish to. There are three special classes of components: Task, Model and Loss. Tasks are components feeding batches of samples into pipelines, and due to the training methodologies that interlace training and validation batches there are treated is a slightly different way, thus are in fact *outside of pipeline*. Models are components containing trainable weights, whereas Losses are components calculating values of loss function for a given batch (and from which gradients are back-propagated). Once the components are implemented, a user defines a pipeline by writing a configuration file (Fig. 3). Components are executed following the priorities defined by the user. A pipeline can consist of any number of components of a given type, including many Models and many Losses (what enables e.g. Multi-Task Learning [Car97]) and other components providing required transformations and computations. The user connects the components by explicitly indicating the data streams connections; Naming Facilities play an important role here, enabling one to rename any stream at the run time, without the need to change the component code. The framework offers full flexibility and it is up to the programmer to choose the granularity of his/her components/models/pipelines.



Figure 3: Exemplary pipeline (left) with its YAML definition (right).

#### 2.2 Components

All components have two basic operation modes: **Initialization** and **Execution** (Fig. 4), that also reflect two main modes of PTP workers. During **Initialization** workers load configuration file(s) and create all instances of the components in the pipeline. Each component receives its subsection

<sup>&</sup>lt;sup>4</sup>Also note that all PTP components and workers come with default configuration files, so one does not have to look at the actual implementation when defining its own experiment.

of configuration file and uses it to set its internal variables, rename stream/global names, get/set global variables, and perform other component-specific operations (load dataset/create nn module etc.). As soon as all components are initialized the worker performs *handshaking*, i.e. makes sure that all streams are correctly connected by checking compatibility of their output-input definitions. Initialization happens only once and when finished, workers switch to **Execution** mode. During **Execution** a worker activates components one by one, according to their priorities. Each component processes data from input streams and publishes results to output streams, logs its operation, collects statistics (automatically exported to CSV files and, optionally, to Tensorboard) etc.

## 2.3 Workers

PTP workers are python scripts that are agnostic to tasks/models/components/pipelines that they are supposed to work with. Currently our framework offers three types of workers:

- **ptp-offline-trainer**: a trainer relying on classical methodology interlacing training and validation at the end of every epoch; creates separate instances of training and validation tasks and trains the models by feeding the created pipeline with batches of data.
- **ptp-online-trainer**: a flexible trainer creating separate instances of training and validation tasks; validation is performed using a single batch every userdefined interval; relying on the notion of an episode rather than epoch.
- **ptp-processor**: performing a single pass over all the samples (batches) returned by a given task; useful for collecting scores on test sets, answers for submissions for competitions etc.



Figure 4: Component operation modes.

In its core, to accelerate the computations on their own, PTP relies on PyTorch and all workers extensively use its mechanisms for distribution of computations on CPUs/GPUs, including multi-process data loaders and multi-GPU data parallelism. The Tasks, Models and other components are agnostic to those operations and the user indicates whether to use the former (data loaders) in configuration files or the latter (GPUs) by passing an adequate argument (-gpu) at run-time.

#### 2.4 Task/Model/Component Zoo

PTP is generally agnostic to operations performed by the components and data passed in data streams. However, the currently provided tasks, models and components focus mostly on applications that can be roughly classified as belonging to: (a) computer vision domain, (b) Natural Language Processing domain and (c) domain combining vision and language. Fig. 5 reflect this. Aside of selection of tasks covering various aspects from those three domains, PTP offers several models that are specialized for language or vision, but also includes several general usage components, such as Feed Forward Network (with variable number of Fully Connected layers with activation functions and dropout between them) or Recurrent Neural Network (with several cell types available; with activation functions and dropout; the component can also work as encoder or decoder). PTP also provides several ready to use non-trainable (but still parametrizable) components. That list includes components useful when working with language, vision or other types of streams (e.g., tensor transformations). There are also several general-purpose components, such as components calculating losses and statistics, and publishers and viewers that enable the user to digest and analyze contents of data streams. The diversity of those components illustrates the flexibility of the framework.

	Language		Vision			Multi-modal	
Tasks	WiLY (Language Identification)     WikiText-2 / WikiText-103 (Language Modelling)		MNIST (Image (	MNIST (Image Classification)     CIEAR-100 (Image Classification)		CLEVR (Visual Reasoning)     GOA (Visual Reasoning)	
	ANKI (Machine Translation)					ImageCLEF VQA-Med 2019	(Visual Question Answering)
	Language		Vision		Multi-modal Reasoning		General Usage
Models	Sencence Embeddings		<ul> <li>Convnet Encoder</li> </ul>		Question-driven attention over image		<ul> <li>Feed Forward Network</li> </ul>
	ୁ GloVe.6B.50d ୁ GloVe.4	2B.300d	<ul> <li>Generic Image Encoder</li> </ul>		• Multi-modal Low-rank Bilinear pooling		<ul> <li>Recurrent Neural Network</li> </ul>
	GloVe.6B.100d GloVe.840B.300d		ଂVGG-16 ResNet-152		• Multi-modal Compact Bilinear pooling		ି RNN TanH ଓ GRU
	GloVe.6B.200d GloVe.twitter.27B		ାResNet-50 ାDenseNet-121		<ul> <li>Milti-modal Factorized Bilinear pooling</li> </ul>		ାRNN Sigmoid ା LSTM
	GloVe.6B.300d OMIMIC.fastText.300d		• LeNet-5		Relational Network		• Seq2Seq
							• Attention Decoder
omponents	Language		Losses & Statistics			Publishers & Viewers	Transformations
	<ul> <li>Sentence Tokenizer</li> </ul>	• Label Indexer	<ul> <li>Accuracy Statistics</li> </ul>	• Precisi	on/Recall	• Global Variable Publisher	<ul> <li>List to Tensor</li> </ul>
	<ul> <li>Sentence Indexer</li> </ul>	BoW Encoder	<ul> <li>BLEU Statistics</li> </ul>	Statistics		• Image Viewer	<ul> <li>Reshape Tensor</li> </ul>
	Sentence One Hot Encoder	• Word Decoder	• Nl ecoder • Batch Size Statistics	• NLL Lo	NLL Loss	• Stream Viewer	<ul> <li>Reduce Tensor</li> </ul>
C						<ul> <li>Stream File Exporter</li> </ul>	<ul> <li>Concatenate Tensors</li> </ul>

Figure 5: Tasks, models and other components currently available in PTP.

# 3 Use case: the Image-CLEF Med-VQA 2019 challenge

The VQA-Med 2019 [BHD<sup>+</sup>19] is an open challenge organized as part of the ImageCLEF 2019 initiative [IMP<sup>+</sup>19]. The associated dataset belongs to the Visual Question Answering (VQA) problem domain. with a focus on radiology images. Despite its small size (a training set of 3,200 images with 12,792 question-answer pairs, divided into four C1–C4 categories) the dataset is scattered, noisy and heavily biased, thus we decided to use it as a test bed for PTP, as we could encounter all the challenges that one must deal with in practical scenarios.



Figure 6: Multi-stage training of the Supporting Facts Network. Green color denotes components pretrained on external datasets, whreas blue indicates parts of the pipeline pre-trained on Image-CLEF.

Summarizing our efforts, within a month we have build more than 50 diverse pipelines, developed dozens of components and made 4 entries to the leaderboard by submitting the files generated straight in PTP. Our final model, called **Supporting Facts Network** (SFN) [KRS<sup>+</sup>19], shared knowledge between upstream and downstream tasks through the use of a coarse-grained categorizer combined with five task-specific, fine-grained classifiers. Training of the final SFN architecture (the whole final pipeline had 53 components!) was a complex procedure, with several transfer learning steps (represented as red arrows in Fig. 6): (0) using pre-trained "Word Embeddings" (GloVe.6B.50d) and "Feature Map Extractor" (ResNet-50), (1) pre-training of "Question Categorizer" on C1,C2,C3 and C4 categories, (2) pre-training "Input Fusion" on C1,C2 and C3 categories, (3) loading and freezing the "Question Categorizer", loading and fine-tuning the "Input Fusion" and (4) training jointly the final SFN architecture using 5 losses (one for each of the fine-grained classifier). Despite we haven't won the challenge (ranked 7), we found PTP extremely useful, supporting all the abovementioned operations and enabling us to rapidly prototype new solutions and ideas.

## References

- [AAL<sup>+</sup>15] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings* of the IEEE international conference on computer vision, pages 2425–2433, 2015.
- [ABC<sup>+</sup>16] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. TensorFlow: a system for large-scale machine learning. In OSDI, volume 16, pages 265–283, 2016.
- [BAM18] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):423–443, 2018.
- [BCB14] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [BHD<sup>+</sup>19] Asma Ben Abacha, Sadid A. Hasan, Vivek V. Datla, Joey Liu, Dina Demner-Fushman, and Henning Müller. VQA-Med: Overview of the medical visual question answering task at imageclef 2019. In *CLEF 2019 Working Notes*, CEUR Workshop Proceedings, Lugano, Switzerland, September 09-12 2019. CEUR-WS.org < http://ceur – ws.org/ >.
- [BLP<sup>+</sup>12] Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, James Bergstra, Ian Goodfellow, Arnaud Bergeron, Nicolas Bouchard, David Warde-Farley, and Yoshua Bengio. Theano: new features and speed improvements. *arXiv preprint arXiv:1211.5590*, 2012.
  - [Car97] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- [CBM02] Ronan Collobert, Samy Bengio, and Johnny Mariéthoz. Torch: a modular machine learning software library. Technical report, Idiap, 2002.
- [DCK18] Mikyas T Desta, Larry Chen, and Tomasz Kornuta. Object-based reasoning in VQA. arXiv preprint arXiv:1801.09718, 2018.
- [DKG<sup>+</sup>17] Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. Visual dialog. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 326–335, 2017.
- [FPY<sup>+</sup>16] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, 2016.
- [GGN<sup>+</sup>17] Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. Allennlp: A deep semantic natural language processing platform. 2017.
- [GMH13] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In 2013 IEEE international conference on acoustics, speech and signal processing, pages 6645–6649. IEEE, 2013.
- [GRG<sup>+</sup>18] Ross Girshick, Ilija Radosavovic, Georgia Gkioxari, Piotr Dollár, and Kaiming He. Detectron. https://github.com/facebookresearch/detectron, 2018.
- [HGDG17] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In Proceedings of the IEEE international conference on computer vision, pages 2961– 2969, 2017.
  - [HM18] Drew A Hudson and Christopher D Manning. Compositional attention networks for machine reasoning. In International Conference on Learning Representations (ICLR), 2018.

- [HRS<sup>+</sup>17] Jonathan Huang, Vivek Rathod, Chen Sun, Menglong Zhu, Anoop Korattikara, Alireza Fathi, Ian Fischer, Zbigniew Wojna, Yang Song, Sergio Guadarrama, and Kevin. Speed/accuracy trade-offs for modern convolutional object detectors. In *Proceedings* of the IEEE conference on computer vision and pattern recognition, pages 7310–7311, 2017.
  - [Hut18] Matthew Hutson. Artificial intelligence faces reproducibility crisis. *Science*, 359(6377):725–726, 2018.
- [IMP<sup>+</sup>19] Bogdan Ionescu, Henning Müller, Renaud Péteri, Yashin Dicente Cid, Vitali Liauchuk, Vassili Kovalev, Dzmitri Klimuk, Aleh Tarasau, Asma Ben Abacha, Sadid A. Hasan, Vivek Datla, Joey Liu, Dina Demner-Fushman, Duc-Tien Dang-Nguyen, Luca Piras, Michael Riegler, Minh-Triet Tran, Mathias Lux, Cathal Gurrin, Obioma Pelka, Christoph M. Friedrich, Alba García Seco de Herrera, Narciso Garcia, Ergina Kavallieratou, Carlos Roberto del Blanco, Carlos Cuevas Rodríguez, Nikos Vasillopoulos, Konstantinos Karampidis, Jon Chamberlain, Adrian Clark, and Antonio Campello. ImageCLEF 2019: Multimedia retrieval in medicine, lifelogging, security and nature. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, Proceedings of the 10th International Conference of the CLEF Association (CLEF 2019), Lugano, Switzerland, September 9-12 2019. LNCS Lecture Notes in Computer Science, Springer.
- [JNC<sup>+</sup>18] Yu Jiang, Vivek Natarajan, Xinlei Chen, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Pythia v0. 1: the winning entry to the vqa challenge 2018. *arXiv preprint arXiv:1807.09956*, 2018.
  - [KE17] Vahid Kazemi and Ali Elqursh. Show, ask, attend, and answer: A strong baseline for visual question answering. *arXiv preprint arXiv:1704.03162*, 2017.
- [KFF15] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137, 2015.
- [KGG<sup>+</sup>18] Oleksii Kuchaiev, Boris Ginsburg, Igor Gitman, Vitaly Lavrukhin, Jason Li, Huyen Nguyen, Carl Case, and Paulius Micikevicius. Mixed-precision training for nlp and speech recognition with openseq2seq. arXiv preprint arXiv:1805.10387, 2018.
- [KLN<sup>+</sup>19] Oleksii Kuchaiev, Jason Li, Huyen Nguyen, Oleksii Hrinchuk, Ryan Leary, Boris Ginsburg, Samuel Kriman, Stanislav Beliaev, Vitaly Lavrukhin, Jack Cook, Patrice Castonguay, Mariya Popova, Jocelyn Huang, and Jonathan M. Cohen. NeMo: a toolkit for building ai applications using neural modules, 2019.
- [KMM<sup>+</sup>18] Tomasz Kornuta, Vincent Marois, Ryan L McAvoy, Younes Bouhadjar, Alexis Asseman, Vincent Albouy, TS Jayram, and Ahmet S Ozcan. Accelerating machine learning research with mi-prometheus. In *NeurIPS Workshop on Machine Learning Open Source Software (MLOSS)*, volume 2018, 2018.
- [KOL<sup>+</sup>17] Jin-Hwa Kim, Kyoung-Woon On, Woosang Lim, Jeonghee Kim, Jung-Woo Ha, and Byoung-Tak Zhang. Hadamard product for low-rank bilinear pooling. In *ICLR*, 2017.
- [KRS<sup>+</sup>19] Tomasz Kornuta, Deepta Rajan, Chaitanya Shivade, Alexis Asseman, and Ahmet S Ozcan. Leveraging medical visual question answering with supporting facts. *arXiv preprint arXiv:1905.12008*, 2019.
- [KSH12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [LBH15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep Learning. *Nature*, 521(7553):436, 2015.
- [LMB<sup>+</sup>14] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.

- [MD18] Mateusz Malinowski and Carl Doersch. The visual qa devil in the details: The impact of early fusion and batch norm on clevr. *arXiv preprint arXiv:1809.04482*, 2018.
- [MF14] Mateusz Malinowski and Mario Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *Advances in neural information processing systems*, pages 1682–1690, 2014.
- [MFF<sup>+</sup>17] Alexander H Miller, Will Feng, Adam Fisch, Jiasen Lu, Dhruv Batra, Antoine Bordes, Devi Parikh, and Jason Weston. Parlai: A dialog research software platform. arXiv preprint arXiv:1705.06476, 2017.
- [MJA<sup>+</sup>18] Vincent Marois, TS Jayram, Vincent Albouy, Tomasz Kornuta, Younes Bouhadjar, and Ahmet S Ozcan. On transfer learning using a MAC model variant. In *NeurIPS'18 Visually-Grounded Interaction and Language (ViGIL) Workshop*, 2018.
- [MKK19] Aditya Mogadala, Marimuthu Kalimuthu, and Dietrich Klakow. Trends in integration of vision and language research: A survey of tasks, datasets, and methods. *arXiv* preprint arXiv:1907.09358, 2019.
- [Mol19] Piero Molino. Ludwig: a type-based declarative deep learning toolbox. *To appear*, 2019.
- [PGC<sup>+</sup>17] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. 2017.
  - [PY09] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions* on knowledge and data engineering, 22(10):1345–1359, 2009.
- [RDGF16] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
  - [SK10] Jason Sanders and Edward Kandrot. *CUDA by example: an introduction to generalpurpose GPU programming.* Addison-Wesley Professional, 2010.
- [SNJ<sup>+</sup>18] Amanpreet Singh, Vivek Natarajan, Yu Jiang, Xinlei Chen, Meet Shah, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Pythia-a platform for vision & language research. In *SysML Workshop, NeurIPS*, volume 2018, 2018.
- [SRB<sup>+</sup>17] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Tim Lillicrap. A simple neural network module for relational reasoning. In Advances in neural information processing systems, pages 4967–4976, 2017.
- [TAHvdH18] Damien Teney, Peter Anderson, Xiaodong He, and Anton van den Hengel. Tips and tricks for visual question answering: Learnings from the 2017 challenge. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 4223–4232, 2018.
  - [TOHC15] Seiya Tokui, Kenta Oono, Shohei Hido, and Justin Clayton. Chainer: a next-generation open source framework for deep learning. In *Proceedings of workshop on machine learning systems (LearningSys) in the twenty-ninth annual conference on neural information processing systems (NIPS)*, volume 5, pages 1–6, 2015.
  - [VBB<sup>+</sup>18] Ashish Vaswani, Samy Bengio, Eugene Brevdo, Francois Chollet, Aidan N Gomez, Stephan Gouws, Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, et al. Tensor2tensor for neural machine translation. arXiv preprint arXiv:1803.07416, 2018.
  - [WCB14] Jason Weston, Sumit Chopra, and Antoine Bordes. Memory networks. *arXiv preprint arXiv:1410.3916*, 2014.