

# Training Transformers Together

**Alexander Borzunov\***

**Max Ryabinin\***

*HSE University, Yandex*

BORZUNOV.ALEXANDER@GMAIL.COM

MRYABININ0@GMAIL.COM

**Tim Dettmers\***

*University of Washington*

DETTMERS@CS.WASHINGTON.EDU

**Quentin Lhoest\***

**Lucile Saulnier\***

*Hugging Face*

QUENTIN@HUGGINGFACE.CO

LUCILE@HUGGINGFACE.CO

**Michael Diskin**

*HSE University, Yandex*

MICHAEL.S.DISKIN@GMAIL.COM

**Yacine Jernite**

**Thomas Wolf**

*Hugging Face*

YACINE@HUGGINGFACE.CO

THOMAS@HUGGINGFACE.CO

**Editors:** Douwe Kiela, Marco Ciccone, Barbara Caputo

## Abstract

The infrastructure necessary for training state-of-the-art models is becoming overly expensive, which makes training such models affordable only to large corporations and institutions. Recent work proposes several methods for training such models collaboratively, i.e., by pooling together hardware from many independent parties and training a shared model over the Internet. In this demonstration, we collaboratively trained a text-to-image transformer similar to OpenAI DALL-E. We invited the viewers to join the ongoing training run, showing them instructions on how to contribute using the available hardware. We explained how to address the engineering challenges associated with such a training run (slow communication, limited memory, uneven performance between devices, and security concerns) and discussed how the viewers can set up collaborative training runs themselves. Finally, we show that the resulting model generates images of reasonable quality on a number of prompts.

**Keywords:** distributed training, volunteer computing, transformers, text-to-image, memory efficiency, communication efficiency, heterogeneous hardware, security

## 1. Introduction

Training state-of-the-art deep learning models is becoming ever more computationally demanding. One infamous example of this trend is transformers (Vaswani et al., 2017), a popular architecture widely used in NLP (Devlin et al., 2019; Liu et al., 2019; Brown et al., 2020), speech processing (Gulati et al., 2020; Li et al., 2019), and computer vision (Dosovitskiy et al., 2020; Touvron et al., 2021; Caron et al., 2021). Transformers benefit from having billions of parameters (Brown et al., 2020; Kaplan et al., 2020; Ott et al., 2018) and large-batch training (Popel and Bojar, 2018), which makes them dependent on large-scale training infrastructure (Narayanan et al., 2021; Shoeybi et al., 2019; Lepikhin et al., 2020).

---

\* Equal contribution.

Unfortunately, this kind of infrastructure can be prohibitively expensive, whether one buys the hardware or rents cloud resources (Turner; Li, 2020). As a result, most researchers simply cannot afford to conduct the necessary experiments to develop their ideas, which ultimately slows down scientific progress.

To make large-scale deep learning more accessible, recent work proposes to train these models collaboratively, i.e., to pool together the hardware from many independent parties and train a shared model over the Internet (Pascutto and Linscott, 2019; Ryabinin and Gusev, 2020; Kijispongse et al., 2018; Atre et al., 2021; Diskin et al., 2021). Such work proposes general distributed algorithms for training on many devices with uneven compute capability and reliability. However, to make them practical, one must overcome several engineering challenges, such as slow communication, limited memory, and security concerns.

In this demonstration, we collaboratively trained a text-to-image transformer similar to DALL-E (Ramesh et al., 2021). Our contributions are the following:

- We modify the DALL-E model, making it suitable for training over the Internet using the method from Diskin et al. (2021) and the `hivemind` library (hivemind, 2020). We set up the infrastructure for such a training run and publish the training results.
- We provide a webpage<sup>1</sup> explaining how to join the ongoing training run, address challenges related to collaborative training runs (slow communication, low memory budget, support of heterogeneous devices), and set up such a training run by yourself.
- We provide an interactive “calculator” that shows the memory consumed by different models in case of using various memory-efficiency techniques. Also, we present a tutorial on setting up dataset streaming and model compression using the `datasets` and `bitsandbytes` libraries (Lhoest et al., 2021; Dettmers et al., 2021).

## 2. Demonstration Contents

### 2.1. Main webpage

The central part of our demonstration is a webpage where people can explore the demonstration materials. The webpage describes the motivation behind collaborative training projects, the method for efficient training from Diskin et al. (2021), and the ongoing collaborative training of our adapted version of DALL-E (see Section 3). Here, we also show a plot of the training objective and the number of active participants.

Next, we provide instructions on how to join the training run using free cloud providers or their own GPU. This involves (1) joining a specific Hugging Face organization, where we can authenticate the users and measure their contribution, and (2) running a Jupyter notebook (Kluyver et al., 2016) with the training code. Our intention was that the user can explore our collaborative training environment through active participation while at the same time reading the detailed explanations of how it works. Here, we also provide the link to the interactive dashboard which shows the statistics and the leaderboard of contributors and provides further information about the training run, such as model checkpoints uploaded to the Model Hub, notebooks for inference, and links to the source code.

Then, we proceed to discuss the engineering challenges of collaborative training runs:

---

1. See <https://training-transformers-together.github.io>

- **Communication efficiency.** Most distributed training algorithms are designed for the networks inside HPC clusters with a 10–100 Gbit/s bandwidth. However, typical Internet connections are orders of magnitude slower (10–100 Mbit/s). To make training over the Internet practical, one can reduce the communication costs using large-batch training (You et al., 2020), gradient compression (Dettmers, 2015; Lin et al., 2018; Vogels et al., 2019; Tang et al., 2021), parameter sharing (Lan et al., 2020; Xue et al., 2021), and overlapping computation with communication (Ren et al., 2021).
- **Uneven device performance.** Traditional data-parallel training waits for the slowest device on every batch. Diskin et al. (2021) allow the devices to process different numbers of samples for a batch, while keeping the guarantees of synchronous training.
- **Memory efficiency.** Distributed training requires either storing all parameters and optimizer statistics on each participant, which is challenging in the case of low-end hardware, or using model parallelism which introduces another level of complexity. Fortunately, the first option is often viable if we reduce the memory consumption with 8-bit optimizers (Dettmers et al., 2021), by offloading the statistics to CPU, with gradient checkpointing or parameter sharing (Lan et al., 2020; Xue et al., 2021).
- **Dataset streaming.** Participants often cannot store or even download the whole dataset, since datasets used for pretraining transformers may contain hundreds of gigabytes of data. To address that, one can use dataset streaming tools, such as the `datasets` library (Lhoest et al., 2021).
- **Security.** Crucially, the participants only exchange tensors and never send code to be executed on each other’s computers. Since a malicious participant also could influence the training outcome by sending wrong tensors, we should either authenticate participants, as described in Diskin et al. (2021), and/or use gradient aggregation techniques robust to outliers (Karimireddy et al., 2020; Gorbunov et al., 2021).

Finally, we provide a recipe on how to combine all that and set up a new collaborative training run using the `hivemind` library (hivemind, 2020).

## 2.2. Memory calculator

The demonstration webpage includes an interactive “calculator” showing the benefits of various memory-efficiency techniques and their combinations. It can compute the consumption of RAM and GPU memory for BERT (Devlin et al., 2019), T5 (Raffel et al., 2020), GPT-2 (Radford et al., 2019), GPT-3 (Brown et al., 2020), GPT-J (Wang and Komatsuzaki, 2021), and DALL-E (Ramesh et al., 2021) in case of using 8-bit optimizers, offloading the optimizer statistics to CPU, using gradient checkpointing and parameter sharing.

## 2.3. Tutorial on memory-efficiency techniques

The demonstration webpage refers to a tutorial on setting up dataset streaming with the `datasets` library (Lhoest et al., 2021) and model compression with the `bitsandbytes` library (Dettmers et al., 2021). The goal of the tutorial is to fine-tune the GPT-2 Large model (Radford et al., 2019) on the C4 dataset (Raffel et al., 2020) using only a low-end GPU, which is possible with the 8-bit Adam optimizer.

### 3. Collaborative Training Run

#### 3.1. Model

For the practical example of a collaborative training run, we chose to train a text-to-image transformer similar to DALL-E (Ramesh et al., 2021), based on the code from Wang (2021). Specifically, we used a decoder-only transformer with 1024 hidden units and 64 layers, each of which uses 16 attention heads with a per-head state size of 64 ( $\approx 1.1$ B parameters in total). We alternated the attention masks as in the original paper, i.e., repeated “row, column, row, row” masks until the last layer, which had the convolutional mask.

To improve communication and memory efficiency, we tied weights of all “row, column, row, row” layer groups (Lan et al., 2020) and tied the input and output embeddings (Press and Wolf, 2016), so the model uses  $\approx 8$ x fewer parameters (but the same amount of compute). We also used reversible layers (Brügger et al., 2019) to reduce memory usage and rotary embeddings (Su et al., 2021) to improve training stability.

We replaced dVAE with VQ-GAN (Esser et al., 2021), since it has a smaller reconstruction error. We used the checkpoint with  $f=8$  and the codebook size 8192. Finally, we used CLIP ViT/B-32 (Radford et al., 2021) to choose the best 4 out of 128 generated images.

#### 3.2. Dataset

We trained the model on the first 100 million image-text pairs from LAION-400M (Schuhmann et al., 2021). We skipped  $\approx 10\%$  images due to short captions, extreme aspect ratios, and NSFW labels.

Before training, we preprocessed all images with VQGAN and uploaded the VQGAN codes and captions, both compressed with Brotli (Alakuijala et al., 2018), to the Hugging Face Dataset Hub (Lhoest et al., 2021). During training, we streamed the compressed codes instead of the original images, thus consuming  $\approx 18$ x less bandwidth.

#### 3.3. Training procedure

We followed the distributed training procedure from Diskin et al. (2021) and used the 8-bit LAMB optimizer (You et al., 2020; Dettmers et al., 2021) offloaded to CPU. We used the linear training schedule with 31250 steps (the first 10% is the warm-up) and the peak learning rate of  $2.5 \cdot 10^{-3}$ . While exchanging gradients and parameters, we used the 8-bit quantization (Dettmers, 2015) for tensors with  $\geq 2^{16}$  elements and the 16-bit precision for other tensors. Unlike the original paper, we did not use PowerSGD (Vogels et al., 2019).

#### 3.4. Results

The training run lasted for 2.5 months and passed  $\approx 80\%$  of the training schedule. Besides the authors, 37 volunteers have contributed for at least 10 minutes (see Appendix A).

During inference, we note that limiting sampling to top 256 logits or top logits whose probability sums up to  $p = 0.75$  greatly improves the image quality. The final model generates realistic images for some prompts but fails to draw correct shapes for the others, while using the appropriate image style, textures, and colors (see Appendix B). We attribute that to the fact that our model is too small to remember the full diversity of images in LAION-400M. Still, the model can generalize to the concepts not present in the dataset.

## References

- Jyrki Alakuijala, Andrea Farruggia, Paolo Ferragina, Eugene Kliuchnikov, Robert Obryk, Zoltan Szabadka, and Lode Vandevenne. Brotli: A general-purpose data compressor. *ACM Transactions on Information Systems (TOIS)*, 37(1):1–30, 2018.
- Medha Atre, Birendra Jha, and Ashwini Rao. Distributed deep learning using volunteer computing-like paradigm, 2021.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Robin Brügger, Christian F. Baumgartner, and Ender Konukoglu. A partially reversible u-net for memory-efficient volumetric image segmentation. *arXiv:1906.06148*, 2019.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *arXiv preprint arXiv:2104.14294*, 2021.
- Tim Dettmers. 8-bit approximations for parallelism in deep learning. *ICLR*, 2015.
- Tim Dettmers, Mike Lewis, Sam Shleifer, and Luke Zettlemoyer. 8-bit optimizers via block-wise quantization, 2021.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019.
- Michael Diskin, Alexey Bukhtiyarov, Max Ryabinin, Lucile Saulnier, Quentin Lhoest, Anton Sinitsin, Dmitriy Popov, Dmitry Pyркиn, Maxim Kashirin, Alexander Borzunov, Albert Villanova del Moral, Denis Mazur, Ilia Kobelev, Yacine Jernite, Thomas Wolf, and Gennady Pekhimenko. Distributed deep learning in open collaborations. *CoRR*, abs/2106.10207, 2021. URL <https://arxiv.org/abs/2106.10207>.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12873–12883, 2021.
- Eduard Gorbunov, Alexander Borzunov, Michael Diskin, and Max Ryabinin. Secure distributed training at scale. *arXiv preprint arXiv:2106.11257*, 2021.
- Anmol Gulati, Chung-Cheng Chiu, James Qin, Jiahui Yu, Niki Parmar, Ruoming Pang, Shibo Wang, Wei Han, Yonghui Wu, Yu Zhang, and Zhengdong Zhang, editors. *Conformer: Convolution-augmented Transformer for Speech Recognition*, 2020.

- hivemind. Hivemind: a Library for Decentralized Deep Learning. <https://github.com/learning-at-home/hivemind>, 2020.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models, 2020.
- Sai Praneeth Karimireddy, Lie He, and Martin Jaggi. Learning from history for byzantine robust optimization. *arXiv preprint arXiv:2012.10333v1*, 2020.
- Ekasit Kijsipongse, Apivadee Piyatumrong, and Suriya U-ruekolan. A hybrid gpu cluster and volunteer computing platform for scalable deep learning. *The Journal of Supercomputing*, 04 2018. doi: 10.1007/s11227-018-2375-9.
- Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian E Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica B Hamrick, Jason Grout, Sylvain Corlay, et al. *Jupyter Notebooks-a publishing format for reproducible computational workflows.*, volume 2016. 2016.
- Zhen-Zhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*, 2020.
- Dmitry Lepikhin, H. Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Y. Huang, M. Krikun, Noam Shazeer, and Z. Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. *ArXiv*, abs/2006.16668, 2020.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. Datasets: A community library for natural language processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. URL <https://aclanthology.org/2021.emnlp-demo.21>.
- Chuan Li. Demystifying gpt-3 language model: A technical overview, 2020. "<https://lambdalabs.com/blog/demystifying-gpt-3>".
- N. Li, Shujie Liu, Yanqing Liu, Sheng Zhao, and Ming Liu. Neural speech synthesis with transformer network. In *AAAI*, 2019.
- Yujun Lin, Song Han, Huizi Mao, Yu Wang, and William J Dally. Deep Gradient Compression: Reducing the communication bandwidth for distributed training. In *The International Conference on Learning Representations*, 2018.



- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692, 2019.
- Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, et al. Efficient large-scale language model training on gpu clusters. *arXiv preprint arXiv:2104.04473*, 2021.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. Scaling neural machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 1–9, Brussels, Belgium, October 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-6301. URL <https://www.aclweb.org/anthology/W18-6301>.
- Gian-Carlo Pascutto and Gary Linscott. Leela chess zero, 2019. URL <http://lczero.org/>.
- M. Popel and Ondrej Bojar. Training tips for the transformer model. *The Prague Bulletin of Mathematical Linguistics*, 110:43 – 70, 2018.
- Ofir Press and Lior Wolf. Using the output embedding to improve language models. *arXiv preprint arXiv:1608.05859*, 2016.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- Colin Raffel, Noam M. Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, W. Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *ArXiv*, abs/1910.10683, 2020.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021.
- Jie Ren, Samyam Rajbhandari, Reza Yazdani Aminabadi, Olatunji Ruwase, Shuangyan Yang, Minjia Zhang, Dong Li, and Yuxiong He. Zero-offload: Democratizing billion-scale model training, 2021.
- Max Ryabinin and Anton Gusev. Towards crowdsourced training of large neural networks using decentralized mixture-of-experts. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 3659–3672. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/25ddc0f8c9d3e22e03d3076f98d83cb2-Paper.pdf>.

- Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using gpu model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*, 2021.
- Hanlin Tang, Shaoduo Gan, Ammar Ahmad Awan, Samyam Rajbhandari, Conglong Li, Xiangru Lian, Ji Liu, Ce Zhang, and Yuxiong He. 1-bit adam: Communication efficient large-scale training with adam’s convergence speed. In *International Conference on Machine Learning*, pages 10118–10129. PMLR, 2021.
- Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021.
- Elliot Turner. Estimate of GPT-3 training cost based on public cloud GPU/TPU cost models, from Elliot Turner’s personal page (accessed on May 29, 2020).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>.
- Thijs Vogels, Sai Praneeth Karimireddy, and Martin Jaggi. Powersgd: Practical low-rank gradient compression for distributed optimization. *Advances in Neural Information Processing Systems*, 32, 2019.
- Ben Wang and Aran Komatsuzaki. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>, May 2021.
- Phil Wang. DALLE-pytorch. Implementation / replication of DALL-E, OpenAI’s Text to Image Transformer, in Pytorch. <https://github.com/lucidrains/DALLE-pytorch>, 2021.
- Fuzhao Xue, Ziji Shi, Futao Wei, Yuxuan Lou, Yong Liu, and Yang You. Go wider instead of deeper. *arXiv preprint arXiv:2107.11817*, 2021.
- Yang You, Jing Li, Sashank Reddi, Jonathan Hseu, Sanjiv Kumar, Srinadh Bhojanapalli, Xiaodan Song, James Demmel, Kurt Keutzer, and Cho-Jui Hsieh. Large batch optimization for deep learning: Training bert in 76 minutes, 2020.