

Eugene Y. Turin (SB # 342413)  
MC GUIRE LAW, P.C.  
1089 Willowcreek Road, Suite 200  
San Diego, CA 92131  
Tel: (312) 893-7002 Ex. 3  
Fax: 312-275-7895  
eturin@mcgpc.com

*Counsel for Plaintiff and the Putative Class Members*

**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA  
SAN FRANCISCO DIVISION**

DARIUS H. JAMES, individually and  
on behalf of similarly situated  
individuals, ) Case No.  
Plaintiff, ) **COMPLAINT**  
v. ) **CLASS ACTION**  
CEREBRAS SYSTEMS INC., a ) **DEMAND FOR JURY TRIAL**  
Delaware corporation, )  
Defendant. )

1 Plaintiff Darius H. James (“Plaintiff”), on behalf of himself and all others  
2 similarly situated, bring this class action complaint (“Complaint”) against Defendant  
3 Cerebras Systems Inc. (“Cerebras” or “Defendant”).

4 **OVERVIEW**

5 1. Artificial intelligence (“AI”) refers to software engineered to mimic  
6 human-like reasoning and inference through algorithmic processes, typically  
7 leveraging statistical methods.

8 2. Large language models (“LLMs”) are AI software programs designed to  
9 reply to user prompts with natural-sounding text outputs. Cerebras-GPT is a series of  
10 LLMs trained, created, and then released by Cerebras in April 2023.

11 3. While the traditional coding process involves human coders inputting  
12 explicit instructions, an LLM is instead trained by processing vast quantities of text  
13 from diverse sources (a “pre-training dataset”), learning statistical patterns and  
14 associations within that data, and encoding those abstract representations into a vast  
15 array of numerical values known as parameters. The goal is to enable the model to  
16 learn general language patterns, grammar, factual knowledge, and contextual  
17 relationships. This results in a versatile base model that can understand and generate  
18 human-like text.

19 4. The pre-training dataset used by Cerebras to pre-train its LLMs used  
20 public domain, licensed, and, crucially, unlicensed copyrighted materials.

21 5. Plaintiff and Class members are authors. They own registered copyrights  
22 in certain books (the “Infringed Works”) that were included in the pre-training dataset  
23 that Cerebras pirated and copied from the internet to train its Cerebras-GPT models.  
24 Plaintiff and Class members never authorized Cerebras to copy, store, and use their  
25 copyrighted works as pre-training materials. Cerebras copied, and thus infringed on  
26 these copyrighted works multiple times to train its Cerebras-GPT large language  
27 models.

1       6. Through the acts described in further detail below, Defendant has  
2 infringed on Plaintiff's copyrighted works, and it continues to do so by continuing to  
3 store, copy, use, and process the pre-training datasets containing copies of Plaintiff's  
4 and the putative Class's Infringed Works.

## JURISDICTION AND VENUE

6       7.     This Court has subject-matter jurisdiction under 28 U.S.C. § 1331  
7 because this case arises under the Copyright Act (17 U.S.C. § 501).

8. Jurisdiction and venue are proper in this judicial district under 28 U.S.C.  
§ 1331 because Cerebras is headquartered in this District. Cerebras created the  
Cerebras-GPT models and distributes them as open-sourced models while using them  
to demonstrate the capabilities of their wafer-scale AI hardware's efficiency.  
Therefore, a substantial part of the events giving rise to the claim occurred in this  
District. A substantial portion of the affected interstate trade and commerce was  
carried out in this District. Defendant has transacted business, maintained substantial  
contacts, and/or committed overt acts in furtherance of the illegal scheme and  
conspiracy throughout the United States, including in this District. Defendant's  
conduct has had the intended and foreseeable effect of causing injury to persons  
residing in, located in, or doing business throughout the United States, including in  
this District.

20       9. Under Civil Local Rule 3-2(c), assignment of this case to the San  
21 Francisco Division is proper because this case pertains to intellectual-property rights,  
22 which is a district-wide case category under General Order No. 44, and therefore  
23 venue is proper in any courthouse in this District.

## PARTIES

25       10. Plaintiff Darius H. James is an author and performance artist who resides  
26 in Connecticut. He is the author of two American publications, *Negrophobia*  
27 and *That's Blaxploitation!: Roots of the Baadassss 'Tude*, and two bilingual German

1 publications, *Voodoo Stew* and *Froggie Chocolates' Christmas Eve*.

2 11. A list of Plaintiff's registered copyrights owned are attached hereto as  
 3 Exhibit A.

4 12. Defendant Cerebras is a Delaware corporation with its principal place of  
 5 business at 1237 E Arques Avenue, Sunnyvale, CA 94085.

6 13. The unlawful acts alleged against the Defendant in this Complaint were  
 7 authorized, ordered, or performed by the Defendant's respective officers, agents,  
 8 employees, representatives, or shareholders while actively engaged in the  
 9 management, direction, or control of the Defendant's business or affairs.

10 14. Various persons or firms not named as defendants may have participated  
 11 as co-conspirators in the violations alleged herein and may have performed acts and  
 12 made statements in furtherance thereof.

13 **FACTUAL ALLEGATIONS**

14 15. Cerebras is an AI hardware company founded in 2015 that originally  
 15 focused on wafer-scale processors for accelerating deep learning workloads and has  
 16 since expanded to become an AI inference and training platform. Organizations across  
 17 fields like medical research, cryptography, energy, and agentic AI use Defendant's  
 18 CS-2 and CS-3 systems to build on-premise supercomputers, while developers and  
 19 enterprises have access to pay-as-you-go cloud offerings as well.

20 16. LLMs are trained by first ingesting a large pre-training dataset in a pre-  
 21 training process. Creating this pre-training dataset necessarily involves creating or  
 22 copying an enormous quantity of textual works.

23 17. During pre-training, the LLM processes each textual work in the training  
 24 dataset to learn statistical patterns and associations within it. The LLM adjusts its  
 25 parameters through optimization techniques to get progressively better at predicting  
 26 sequences in the data, capturing general linguistic structures rather than specific  
 27 expressions. The results of this learning process are encoded in a large set of numbers

1 called parameters stored within the model. These parameters are derived from the  
 2 entire pre-training dataset. For instance, one of the Cerebras-GPT series LLMs stores  
 3 13 billion parameters.

4       18. Thus, the ability of an LLM to generate human-like textual outputs is  
 5 directly based on it having ingested large pre-training datasets.

6       19. The Pile is a pre-training dataset assembled by EleutherAI for use in  
 7 training AI models. In December 2020, EleutherAI introduced this dataset in a paper  
 8 called “*The Pile: An 800GB Dataset of Diverse Text for Language Modeling*”.<sup>1</sup>

9       20. The paper provides a description of the Books3 dataset contained within  
 10 The Pile:

11  
 12       Books3 is a dataset of books derived from a copy of the contents of the  
 13 Bibliotik private tracker ... Bibliotik consists of a mix of fiction and  
 14 nonfiction books and is almost an order of magnitude larger than our  
 15 next largest book dataset (BookCorpus2). We included Bibliotik  
 16 because books are invaluable for long-range context modeling research  
 17 and coherent storytelling.<sup>2</sup>

18       21. Bibliotik is one of a number of notorious pirate websites that also  
 19 includes Library Genesis (aka LibGen), Z-Library (aka B-ok), Sci-Hub, and Anna’s  
 20 Archive. These pirate libraries have long been of interest to the AI-training  
 21 community because they host and distribute vast quantities of unauthorized  
 22 copyrighted material, including books. For that reason, copying these libraries also  
 23 violates the U.S. Copyright Act.

24       22. Shawn Presser, an individual credited with assembling the Books3  
 25 dataset, has confirmed in public statements that it represents “all of Bibliotik” and  
 26 contains approximately 196,640 books.

---

27<sup>1</sup> <https://arxiv.org/abs/2101.00027>

28<sup>2</sup> *Id.* at 3-4.

1       23. EleutherAI’s website encourages the public to download The Pile dataset  
 2 from a website known as “The-Eye.” Anyone who downloads The Pile from The-Eye  
 3 is therefore also downloading a copy of Books3.

4       24. The Pile is also available for download from the “Hugging Face”  
 5 website. Before October 2023, the Books3 subset of The Pile was available for  
 6 download from Hugging Face as a standalone dataset. But in October 2023, the  
 7 Books3 dataset was removed with a message that it “is defunct and no longer  
 8 accessible due to reported copyright infringement.”<sup>3</sup> Thus, anyone who downloaded  
 9 The Pile dataset(s) from Hugging Face, or any other platform prior to October of 2023  
 10 also downloaded the Books3 dataset.

11       25. Plaintiff’s copyrighted books are among the works in the Books3 dataset.

12       26. In order to pre-train its LLM model(s), Cerebras downloaded and stored  
 13 copies of The Pile containing Books3, which contained copyrighted works including  
 14 those copyrighted by Plaintiff and Class members, and used the dataset to develop  
 15 and train its Cerebras-GPT LLMs.

16       27. In April 2023, Cerebras released its Cerebras-GPT series of LLMs which  
 17 had been pre-trained on The Pile, including Books3.

18       28. Cerebras’ own research confirms that its LLM(s) were pre-trained on  
 19 The Pile. *See Cerebras-GPT: Open Computation-Optimal Language Models Trained on*  
 20 *the Cerebras Wafer-Scale Cluster* (Attached hereto as Exhibit B) (Section 2.2, p. 2).

21       29. Cerebras further admits that it “pre-train[s] models on the Pile dataset,  
 22 which consists of data from 22 data sources, including Common Crawl, PubMed  
 23 Central, Books3, OpenWebText2, Github, and arXiv (Gao et al., 2020).” (Ex. B,  
 24 Section 2.2, p. 3.); *see also* (*Id.* at p. 2) (“[w]e train Cerebras-GPT . . . models . . . on  
 25 the Pile dataset”).

---

26       3

27       

---

 28 [https://web.archive.org/web/20231127101818/https://huggingface.co/datasets/the\\_pile\\_books3](https://web.archive.org/web/20231127101818/https://huggingface.co/datasets/the_pile_books3)

1       30. Because Defendant acquired these datasets before October 2023,  
2 Defendant copied, stored, and used versions of The Pile containing Books3 to develop  
3 and train its Cerebras-GPT models.

4       31. Thus, Cerebras directly infringed on Plaintiff's and Class Members'  
5 copyrighted works on a massive scale. Cerebras downloaded and copied these books  
6 and the Infringed Works, without authorization from, or after providing compensation  
7 to, their authors. Cerebras then continued copying and storing the datasets and used  
8 them to develop and train its Cerebras-GPT LLMs.

9       32. Certain books written by Plaintiff are part of Books3—including the  
10 Infringed Works—and thus Cerebras necessarily trained its models on one or more  
11 copies of the Infringed Works, thereby directly infringing the copyrights of the  
12 Plaintiff and the members of the putative class.

## CLASS ALLEGATIONS

14       33. The “Class Period” as defined in this Complaint begins on at least  
15 October 30, 2022 and runs through the present. Because Plaintiff does not yet know  
16 when the unlawful conduct alleged herein began, but believes, on information and  
17 belief, that the conduct likely began prior to October 30, 2022, Plaintiff reserves the  
18 right to amend the Class Period to comport with the facts and evidence uncovered  
19 during further investigation or through discovery.

20       34. Plaintiff seeks certification of the following Class pursuant to Federal  
21 Rules of Civil Procedure 23(a), 23(b)(2), and 23(b)(3):

23 All persons or entities domiciled in the United States that own a United  
24 States copyright in any work that was copied, stored, or used as pre-  
training data by Defendant during the Class Period.

26       35. Plaintiff will fairly and adequately represent and protect the interests of  
27 the other members of the Class. Plaintiff has retained counsel with substantial

1 experience in prosecuting complex litigation and class actions. Plaintiff and his  
 2 counsel are committed to vigorously prosecuting this action on behalf of the other  
 3 members of the Class, and have the financial resources to do so. Neither Plaintiff nor  
 4 his counsel have any interest adverse to those of the other members of the Class.

5       36. Absent a class action, most members of the Class would find the cost of  
 6 litigation their claims to be prohibitive and would have no effective remedy. The class  
 7 treatment of common questions of law and fact is also superior to multiple individual  
 8 actions or piecemeal litigation in that it conserves the resources of the courts and the  
 9 litigants and promotes consistency and efficiency of adjudication.

10       37. Defendant has acted and failed to act on grounds generally applicable to  
 11 Plaintiff and the other members of the Class, requiring the Court's imposition of  
 12 uniform relief to ensure compatible standards of conduct toward the members of the  
 13 Class, and making injunctive or corresponding declaratory relief appropriate for the  
 14 Class as a whole.

15       38. The factual and legal bases of Defendant's liability to Plaintiff and to the  
 16 other members of the Class are the same, resulting in injury to Plaintiff and to all of  
 17 the other members of the Class. Plaintiff and the other members of the Class have all  
 18 suffered harm and damages as a result of Defendant's unlawful and wrongful conduct.

19       39. There are many questions of law and fact common to the claims of  
 20 Plaintiff and the other members of the Class, and those questions predominate over  
 21 any questions that may affect individual members of the Class. Common questions  
 22 for the Class include, but are not limited to, the following:

- 23       a. Whether Defendant violated the copyrights of Plaintiff and the Class by  
 24           obtaining and creating copies of Plaintiff's Infringed Works with the  
 25           intent to use the Infringed Works for commercial benefit;
- 26       b. Whether Defendant did use the Infringed Works of Plaintiff and the  
 27           Class for commercial benefit;

- c. Whether Defendant violated the copyrights of Plaintiff and the Class by using illicitly obtained copies of Plaintiff's Infringed Works to train Defendant's AI models; and
- d. Whether Defendant caused further infringement of the Infringed Works by distributing the AI models under an open license.

**FIRST CAUSE OF ACTION  
Direct Copyright Infringement,  
(17 U.S.C. § 501)  
(On Behalf of Plaintiff and the Class)**

40. Plaintiff repeats the allegations contained in the foregoing paragraphs as if fully set forth herein.

41. Plaintiff, as the owner of the registered copyrights hold the exclusive rights to those books under 17 U.S.C. § 106.

42. In order to supply enough data for pre-training of the Cerebras-GPT models, Cerebras downloaded and stored copies of the Pile, which included Books3. Books3 includes the Plaintiff's copyrighted works. Cerebras made multiples copies of the Books3 dataset (and thus Plaintiff and the Class members' copyrighted works) for pre-training its AI models.

43. Neither Plaintiff nor Class members authorized Cerebras to make copies of, make derivative works, publicly display copies (or derivative works), or distribute copies (or derivative works) of their copyrighted works. The U.S. Copyright Act bestows all the aforementioned rights only on Plaintiff and the Class members.

44. By copying, storing, processing, reproducing, and using the datasets containing copies of Plaintiff's Infringed Works, Defendant has directly infringed on Plaintiff's exclusive rights in his copyrighted works.

45. By copying, storing, processing, and reproducing the Cerebras-GPT LLMs trained on Plaintiff's copyrighted works, Defendant has directly infringed on Plaintiff's exclusive rights in his copyrighted works.

46. Defendant repeatedly copied, stored, and used the Infringed Works without Plaintiff's and members of the Class's permission in violation of their exclusive rights under the Copyright Act.

47. By and through the actions alleged above, Defendant has infringed and will continue to infringe on Plaintiff's copyrights.

48. Plaintiff has been injured by Defendant's acts of direct copyright infringement. Plaintiff is entitled to statutory damages, actual damages, restitution of profits, and all appropriate legal and equitable relief.

## PRAYER FOR RELIEF

WHEREFORE, Plaintiff, individually and on behalf of all others similarly situated, seeks the following against Defendant:

- a. An order certifying the Class, naming Plaintiff as Class Representative, and naming Plaintiff's attorneys as Class Counsel to represent the Class;
- b. An order declaring that Defendant's conduct violates 17 U.S.C. § 501;
- c. An award of statutory and other damages under 17 U.S.C. § 504 for violations of the copyrights of Plaintiff and the Class by Defendant;
- d. Reasonable attorneys' fees and reimbursement of costs under 17 U.S.C. § 505 or otherwise;
- e. A declaration that such infringement is willful;
- f. Destruction or other reasonable disposition of all copies Defendant made or used in violation of the exclusive rights of Plaintiff and the Class, under 17 U.S.C. § 503(b);
- g. Pre- and post-judgment interest on the damages awarded to Plaintiff and the Class, and that such interest be awarded at the highest legal rate from and after the date this class action complaint is first served on Defendant;
- h. Further relief for Plaintiff and the Class as the Court deems may be appropriate.

## JURY TRIAL DEMAND

Plaintiff demands a trial by jury on all causes of action and issues so triable.

DATED: October 30, 2025 Respectfully submitted,

DARIUS H. JAMES, individually and on behalf of  
similarly situated individuals

By: /s/ Eugene Y. Turin  
Plaintiff's Attorney

Eugene Y. Turin (SB # 342413)  
David L. Gerbie (*pro hac vice* forthcoming)  
Jordan R. Frysinger (*pro hac vice* forthcoming)  
MC GUIRE LAW, P.C.  
1089 Willowcreek Road, Suite 200  
San Diego, CA 92131  
Tel: (312) 893-7002 Ex. 3  
Fax: 312-275-7895  
eturin@mcgpc.com  
dgerbie@mcgpc.com  
jfrysinger@mcgpc.com

*Counsel for Plaintiff and the Putative Class*

# **EXHIBIT A**

---

Registration Number / Date: TX0003372564 / 1992-08-05

Title: *Negrophobia: an urban parable: a novel* / Darius James.

Copyright Claimant:

Darius James

USCO Catalog Link:

[https://publicrecords.copyright.gov/detailed-record/voyager\\_14732620](https://publicrecords.copyright.gov/detailed-record/voyager_14732620)

---

Registration Number / Date: TX0004193939 / 1996-01-03

Title: *That's blaxploitation!: roots of the baadassss 'tude* / Darius James.

Copyright Claimant: Darius James

USCO Catalog Link:

[https://publicrecords.copyright.gov/detailed-record/voyager\\_15167875](https://publicrecords.copyright.gov/detailed-record/voyager_15167875)

## **EXHIBIT B**



# Cerebras-GPT: Open Compute-Optimal Language Models Trained on the Cerebras Wafer-Scale Cluster

Nolan Dey, Gurpreet Gosal, Zhiming (Charles) Chen, Hemant Khachane, William Marshall, Ribhu Pathria, Marvin Tom, Joel Hestness

Cerebras Systems [{nolan,joel}@cerebras.net](mailto:{nolan,joel}@cerebras.net)

## Abstract

We study recent research advances that improve large language models through efficient pre-training and scaling, and open datasets and tools. We combine these advances to introduce Cerebras-GPT, a family of open compute-optimal language models scaled from 111M to 13B parameters. We train Cerebras-GPT models on the Eleuther Pile dataset following DeepMind Chinchilla scaling rules for efficient pre-training (highest accuracy for a given compute budget). We characterize the predictable power-law scaling and compare Cerebras-GPT with other publicly-available models to show all Cerebras-GPT models have state-of-the-art training efficiency on both pre-training and downstream objectives. We describe our learnings including how Maximal Update Parameterization ( $\mu$ P) can further improve large model scaling, improving accuracy and hyperparameter predictability at scale. We release our pre-trained models and code, making this paper the first open and reproducible work comparing compute-optimal model scaling to models trained on fixed dataset sizes. Cerebras-GPT models are available on HuggingFace: <https://huggingface.co/cerebras>.

## 1 Introduction

Recent research in large language models (LLMs) shows important advances that can improve LLM quality and efficiency. Scaling law studies show predictable and significant improvements in model performance by increasing model and dataset size (Hestness et al., 2017; Kaplan et al., 2020). Language models can also be improved just by training on more data (Hoffmann et al., 2022; Touvron et al., 2023). Recent works, such as Maximal Update Parameterization ( $\mu$ P), also show techniques to improve training stability and performance as models scale up (e.g., Bachlechner et al. (2020); Yang et al. (2021)).

Concurrently with these advances, the research community has trained and released many open-source models. Models like GPT-J, GPT-NeoX, OPT, and Pythia have each held state-of-the-art accuracy for open source models for their size, and these models can be tested and used simply by downloading the pre-trained weights (Wang & Komatsuzaki, 2021; Black et al., 2022; Zhang et al., 2022; Biderman et al., 2023). While these models are important contributions, they have not aimed to be compute-efficient. The research community needs more reproducible scaling efforts that can guide collective decisions about training large foundation models in a compute-efficient way.

We introduce Cerebras-GPT, our open effort to combine recent LLM efficient scaling techniques to produce compute-optimal pre-trained models and corresponding scaling laws. Cerebras-GPT is a family of GPT-3-like models that we scale from 111M to 13B parameters. We train them on the open-source dataset, the Pile (Gao et al., 2020), following DeepMind’s Chinchilla scaling rules (Hoffmann et al., 2022). Cerebras-GPT models show state-of-the-art training efficiency when targeting both upstream Pile evaluations as well as a suite of downstream tasks. Our largest model shows state-of-the-art performance on pre-training and most downstream tasks compared to other comparably-sized public models. We also characterize some of the training stability challenges when scaling Cerebras-GPT. We address the challenges by training models with  $\mu$ P, which shows further accuracy improvements and hyperparameter predictability.

Cerebras-GPT models form the compute-optimal Pareto frontier for both pre-training and popular downstream objectives. Figure 1 shows the upstream Pile frontiers compared to contemporary works. We characterize the Pareto frontiers with scaling laws that can be used to predict the benefits of further model and dataset scaling efforts. We also observe and discuss that future open efforts should consider aggregate compute budget (both pre-training and expected inferences) when deciding the appropriate balance of model size and pre-training dataset size.

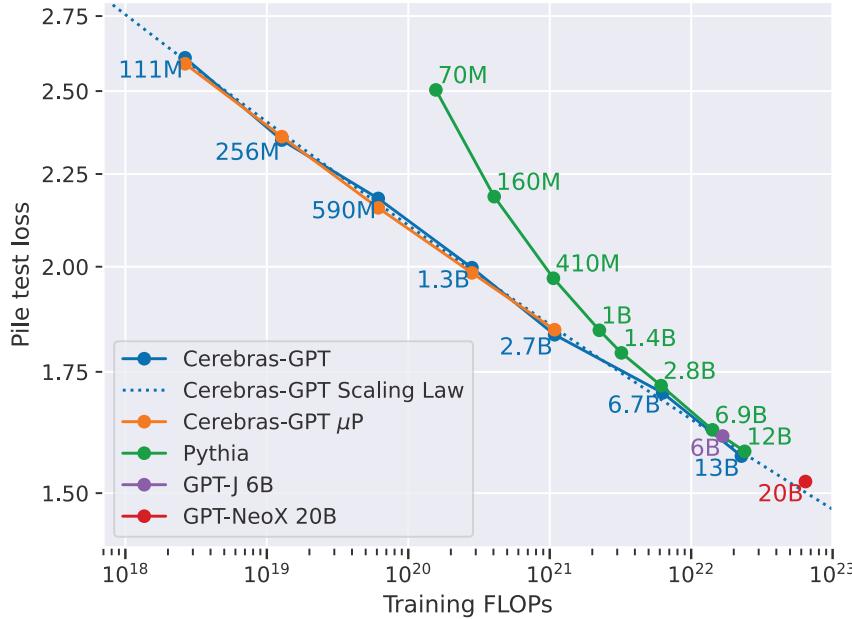


Figure 1: Pile test set loss given pre-training FLOPs for Cerebras-GPT, GPT-J, GPT-NeoX, and Pythia.

Overall, the contributions of this work are as follows:

- We train Cerebras-GPT compute-optimal models scaled from 111M to 13B parameters on the Pile dataset following Chinchilla scaling rules to collect compute-efficient scaling laws.
- We show that these models provide state-of-the-art pre-training efficiency on both pre-training and downstream objectives compared to other open models—the first such open effort.
- We provide detailed instructions to reproduce our results, including the use of  $\mu$ P to improve training stability and transfer hyperparameters as models scale up.
- We document our experience training these models on the Andromeda AI Cluster, comprising 16 Cerebras CS-2 systems, and we describe the simplicity of scaling models and performance.

Finally, we aim to enable the research community to consume these results. We release our pre-trained models and code, and we share details about our training process here, so the community can use and reproduce our results. Pre-trained models are available on HuggingFace: <https://huggingface.co/cerebras>. Source code is available in the Cerebras Modelzoo: <https://github.com/Cerebras/modelzoo>. We hope these models will be a valuable addition for the open-source community.

## 2 Methodology

In this section, we describe the details of the models we trained, including hyperparameters used at each model scale and details about how we obtain and use the Pile dataset. We also motivate the need for techniques to stabilize scaling, and we describe how we use Maximal Update Parameterization ( $\mu$ P).

## 2.1 Model Architecture

Cerebras-GPT models have a GPT-3-like architecture, an autoregressive transformer decoder model (Brown et al., 2020). The main difference is that unlike GPT-3, which uses alternating dense and sparse-banded attention, we use dense attention in all decoder blocks. We select model dimensions to either follow aspect ratio  $\sim 80$  ( $d_{\text{model}}/n_{\text{layers}}$ ) or the same shape as GPT-3 models. All models are trained with a maximum sequence length of 2048 tokens. Table 1 lists the specific model dimensions for each model size. Our formula for the number of parameters is provided in appendix E.

Table 1: Cerebras-GPT model architecture and training algorithm details

Parameters	Model Dimensions				Total tokens	Batch Size (tokens)	Learning Rate (LR)	LR Decay Type
	$d_{\text{model}}$	$n_{\text{layers}}$	$d_{\text{head}}$	$d_{\text{ffn}}$				
111M	768	10	64	3072	2.2B	246K	6.0E-04	Linear
256M	1088	14	64	4352	5.1B	541K	6.0E-04	Linear
590M	1536	18	128	6144	11.8B	541K	2.0E-04	Linear
1.3B	2048	24	128	8192	26.3B	1.08M	2.0E-04	Cosine
2.7B	2560	32	80	10240	53.0B	1.08M	2.0E-04	Cosine
6.7B	4096	32	128	16384	133.2B	2.13M	1.2E-04	Linear
13B	5120	40	128	20480	257.1B	1.47M → 2.21M	1.2E-04	Cosine
111M + $\mu$ P	768	10	64	3072	2.2B	246K	6.0E-03	Linear
256M + $\mu$ P	1088	14	64	4352	5.1B	541K	6.0E-03	Linear
590M + $\mu$ P	1536	18	128	6144	11.8B	541K	6.0E-03	Linear
1.3B + $\mu$ P	2048	24	128	8192	26.3B	1.08M	6.0E-03	Linear
2.7B + $\mu$ P	2560	32	80	10240	53.0B	1.08M	6.0E-03	Linear

## 2.2 Pre-training Corpus

We pre-train models on the Pile dataset, which consists of data from 22 data sources, including Common Crawl, PubMed Central, Books3, OpenWebText2, Github, and arXiv (Gao et al., 2020). We use the dataset splits for train, test, and validation sets provided in the Pile configuration. We tokenize the corpora with byte-pair encoding and the GPT-2 vocabulary of size 50257 (Sennrich et al., 2016; Radford et al., 2019). We do not perform deduplication of Pile but believe that deduplication could further improve our results. We include more details about Pile and dataset pre-processing in Appendix A.1.

To evaluate pre-training, we compare Cerebras-GPT models to several publicly-available models using cross-entropy loss on the Pile test set. To ensure fair comparisons, we run evaluation ourselves on all checkpoints rather than using published numbers, though in most cases, our evaluations match prior works. For models that use different vocabularies, we correct cross-entropy back to the equivalent value with the GPT-2 vocabulary based on the number of tokens in each dataset.

## 2.3 Model Training

We train models using the following training configurations. We use the AdamW optimizer (Loshchilov & Hutter, 2017) with  $(\text{beta1}, \text{beta2}) = (0.9, 0.95)$ . We set epsilon to  $1e-8$  for small models and to  $1e-9$  for 6.7B and 13B parameter models. We use weight decay of 0.1 for all models. We do not use dropout for pre-training. For all runs, we use gradient norm clipping of 1.0.

We use learning rates and batch sizes consistent with prior works, as listed in Table 1. We find that linear learning rate decay tends to perform better than cosine decay, so we use it in most of our pre-training runs. With either decay type, we warm up learning rate linearly over 375M tokens and then decay to 10% of the maximum learning rate. The table also includes batch sizing. For the 13B parameter model, we train with a batch size of 720 sequences of length 2048 tokens for the first 84B tokens. At that point, we observed the gap between validation and train loss growing, indicating that the gradient noise was growing, so we increased the batch size to 1080 sequences for the rest of training.

To scale Cerebras-GPT model training in a compute-efficient way, we follow the DeepMind Chinchilla scaling methodology outlined in (Hoffmann et al., 2022). Specifically, we test and find that models trained with roughly 20 tokens per parameter offer the most compute-efficient pre-training, consistent with the Chinchilla results. We believe this paper is the first open effort to estimate the compute-efficient tokens per parameter for the Pile dataset. Our results Section 3.1 characterizes the effect of training with more tokens per parameter, and we include further test results in Appendix D.

Finally, we train models using both FP16 mixed precision and bfloat16 precision (Micikevicius et al., 2018; Abadi et al., 2016). Overall, we find bfloat16 to be more stable due to its extra exponent range, so we use it for all Cerebras-GPT models that we release. We include further discussion of precision in Appendix A.2.

## 2.4 Standard (SP) and Maximal Update Parameterization ( $\mu$ P)

**Standard Parameterization (SP):** We configure our main Cerebras-GPT models with the common standard parameterization (SP) approach. In SP, model weights are initialized from normal distributions with constant standard deviation or standard deviation based on the shape of each layer (Glorot & Bengio, 2010). We initialize embedding and hidden layer weights with a truncated normal distribution with standard deviation  $\sigma = 0.02$ . An exception is that we use a standard deviation of  $\sigma = 0.02/\sqrt{2 \cdot n_{\text{layers}}}$  for the last layer inside each residual network, following the GPT-2 initialization (Radford et al., 2019).

Unfortunately, the SP approach does not account for potential inter-layer interactions and resulting training dynamics that arise when scaling to very large models. As SP models scale, they tend to become unstable as weight and activation values bump up against the limits of the numerical representations used to train them. For large models, unstable training can cause very costly restarts and researchers might not have budget for extensive hyperparameter tuning.

**Maximal Update Parameterization ( $\mu$ P):** To address these issues, we also train a set of Cerebras-GPT models with Maximal Update Parameterization ( $\mu$ P) (Yang et al., 2021).  $\mu$ P controls initialization, layer-wise learning rates, and activation magnitudes to ensure analytically stable training independent of a model’s layer widths. In addition to improving training stability,  $\mu$ P also improves the transferability of training hyperparameters from smaller to larger scale models, a technique called  $\mu$ Transfer.  $\mu$ Transfer permits directly using the same settings for some optimizer hyperparameters, most notably the learning rate.

We train a set of Cerebras-GPT models using  $\mu$ P. We follow the  $\mu$ Transfer approach by first tuning hyperparameters for a small, 40M parameter  $\mu$ P model. Then, we transfer the hyperparameters along our  $\mu$ P scaling law up to a 2.7B parameter model.  $\mu$ P requires small changes to our baseline Cerebras-GPT models, including adding element-wise activation tensor scaling, adjusting initializers for affected layers, and adding layer-wise learning rates scaling to certain layers. We discuss the benefits we see with  $\mu$ P in Section 3.3. Refer to Appendix G for our tips to implement  $\mu$ P and our hyperparameter tuning notes.

## 3 Results

In this section, we show pre-training and downstream evaluations of Cerebras-GPT models, scaled from 111M to 13B parameters, and we compare against recent related works. We characterize the compute-efficient Pareto frontier for pre-training models on the Pile dataset and show that models on this frontier are also competitive on downstream tasks. We believe this is the first study to release a compute-optimal scaling law for pre-training on the Pile dataset that is openly reproducible by the community.

We show that Cerebras-GPT models define the state-of-the-art compute-optimal Pareto frontier on both pre-training and downstream objectives. Further, our largest model with 13B parameters shows improved accuracy on most downstream tasks compared to other comparably-sized publicly-available models<sup>1</sup>. We also train Cerebras-GPT models configured using  $\mu$ P. We show that  $\mu$ P enables direct hyperparameter transfer from smaller to larger models and improves the compute-optimal frontier loss by 0.4%.

<sup>1</sup>We believe the LLaMa 13B model is better than Cerebras-GPT on downstream tasks because it was trained for 4 $\times$  more tokens, but were unable to get access to test the model ourselves.

### 3.1 Pre-training Results

We scaled and pre-trained Cerebras-GPT models from 111M–13B parameters on the Pile dataset. We compare the Pile test set loss<sup>2</sup> for Cerebras-GPT models against other publicly available pre-trained models, GPT-J, GPT-NeoX, and Pythia (Wang & Komatsuzaki, 2021; Black et al., 2022; Biderman et al., 2023). We believe these models to be fair comparisons either because the models were trained directly on Pile or on similarly-prepared datasets.

Figure 2 plots pre-training efficiency (values also listed in Table 2). The horizontal axis plots floating-point operations (FLOPs) spent during pre-training (log scale), and the vertical axis plots Pile test loss (log scale)<sup>3</sup>. Across all model scales, Cerebras-GPT sets the efficiency frontier, largely because models were pre-trained with 20 tokens per parameter, consistent with findings in the Chinchilla paper. Other public models use more tokens per parameter, requiring more FLOPs to achieve similar loss.

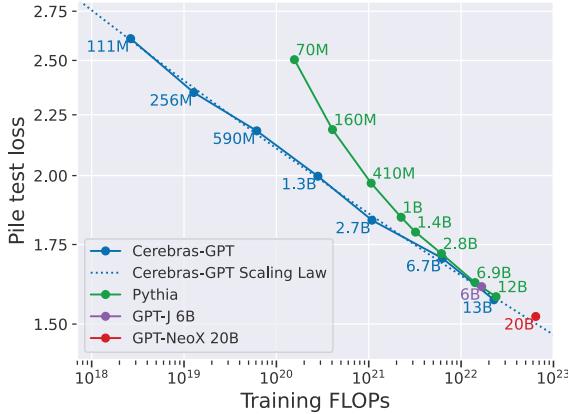


Figure 2: Pile test set loss given pre-training FLOPs for Cerebras-GPT, GPT-J, GPT-NeoX, and Pythia.

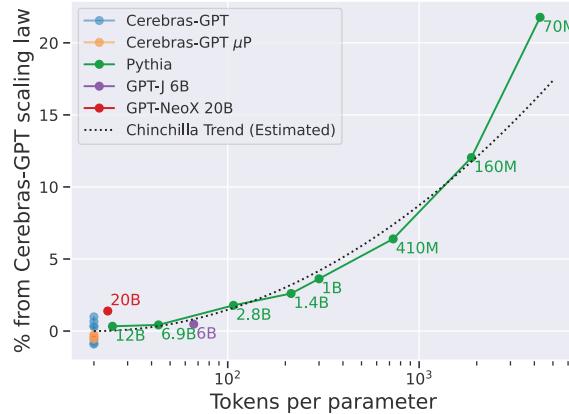


Figure 3: Percent loss degradation from Cerebras-GPT compute-optimal scaling law.

There are a couple notable observations from Figure 2. First, the scaling law for Cerebras-GPT models extrapolates accurately to larger model scales. We estimated the 13B model loss using a similar scaling law from models up to 6.7B parameters, and the 13B model trained to within 0.5% of projected loss. Extending the existing scaling law shows that if we budgeted to train a model with FLOPs equivalent to GPT-NeoX 20B, we would expect the Cerebras-GPT model loss to be  $\sim 1.2\%$  better than GPT-NeoX 20B. For future reference, we include the compute-optimal frontier scaling law here ( $f$  is compute FLOPs to loss,  $\mathcal{L}$ ):

$$\mathcal{L}(f) = (f/5.984e22)^{-0.0737} + 0.5066 \quad (1)$$

Second, increasing tokens per parameter above 20 leads to smoothly degraded loss for the FLOP budget. Pythia models are each trained using 299.9B tokens from the Pile. As model size increases, tokens per parameter decreases reciprocally, and losses move closer to the compute-optimal frontier. The largest Pythia model at 12B parameters is trained with 25.3 tokens per parameter and is just 0.3% loss above the Cerebras-GPT scaling law.

The loss gap from the compute-optimal frontier appears to be predictable in terms of tokens per parameter. In Figure 3, we plot the percentage loss increase compared to the Cerebras-GPT frontier as a function of tokens per parameter. Here, Cerebras-GPT models cluster at 20 tokens per parameter, and Pythia results show the smooth curve away from the frontier for more tokens per parameter. We also include an estimate of the Chinchilla loss degradation from curve fitting data in their plots (Hoffmann et al. (2022), Figure 3). These results confirm the estimate that compute optimal pre-training on the Pile should use roughly 20

<sup>2</sup>All Cerebras-GPT development and hyperparameter tuning was evaluated using the Pile validation set.

<sup>3</sup>Pile test loss is crossentropy in nats/token. We correct all crossentropy results for different vocabularies to be comparable to the GPT-2 vocabulary.

tokens per parameter, a striking consistency with the Chinchilla results on the MassiveText dataset. Further tokens per parameter tests are in Appendix D.

### 3.2 Downstream Results

We evaluate Cerebras-GPT and publicly-available models on a suite of seven common sense reasoning tasks in both the zero-shot and five-shot settings using the EleutherAI evaluation harness (Gao et al., 2021). In particular, we evaluate models on the tasks HellaSwag, PIQA, WinoGrande, Lambda, ARC (both the easy and challenge versions), and OpenBookQA (Zellers et al., 2019; Bisk et al., 2020; Sakaguchi et al., 2021; Paperno et al., 2016; Clark et al., 2018; Mihaylov et al., 2018). We include more detail about these tasks in Appendix B. In addition to models we evaluate on upstream Pile, we add downstream results for OPT models (Zhang et al., 2022), which were trained on a broader dataset but still using 300B pre-training tokens.

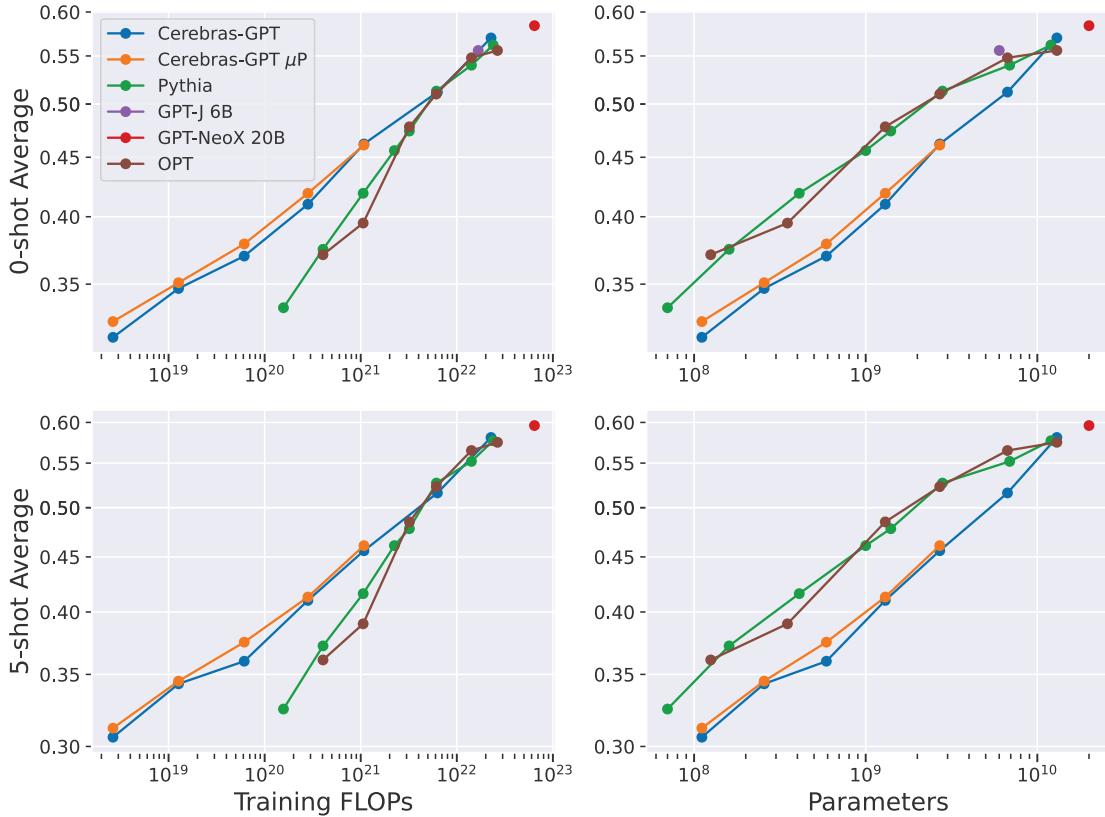


Figure 4: Average zero- and five-shot downstream task accuracy plotted against FLOPs (left) and parameters (right). Higher accuracy is better. Individual tasks are plotted in Figures 9 and 11.

Like the pre-training results, Cerebras-GPT models form the compute-optimal Pareto frontier for downstream tasks as well. Figure 4 summarizes the average downstream task results for both zero- and five-shot evaluations<sup>4</sup> comparing Cerebras-GPT to GPT-J, GPT-NeoX, and Pythia. As Pythia and OPT models grow close to the 20 tokens per parameter count, they approach the Cerebras-GPT frontier FLOPs-to-accuracy. Here again, the Cerebras-GPT 13B model shows the best average downstream result for models of comparable size.

Figure 4 also plots downstream averages against model size in parameters (right column). For each model size smaller than 13B parameters, GPT-J, OPT, and Pythia models show significantly better downstream accuracy than Cerebras-GPT models, as expected. The Pythia and OPT accuracy frontiers deflect from straight

<sup>4</sup>Here, we report accuracy result from each model predictions using token-level probability, consistent with reported results in the GPT-NeoX paper. We report additional accuracy measures in Appendix C.2.

lines (power-laws in log-log-scale), whereas Cerebras-GPT frontiers continue, indicating that downstream accuracy is predictable by model size for models trained with fixed tokens-per-parameter. The Cerebras-GPT trend suggests these models would be competitive with GPT-NeoX 20B if scaled to that size.

Table 2: Zero-shot downstream task results for large publicly-available models. Full results in Table 8.

Model	Pre-training (↓)		Downstream task accuracy (↑)								
	Training FLOPs	Pile test xent	Hella-Swag	PIQA	Wino-Grande	Lambada	ARC-e	ARC-c	Open-BookQA	Downstream Avg.	
OPT	2.7B	6.1e21	-	<b>0.458</b>	<b>0.738</b>	0.610	0.637	0.609	0.268	<b>0.250</b>	0.510
Pythia	2.8B	6.1e21	<b>1.720</b>	0.451	0.737	<b>0.612</b>	<b>0.654</b>	<b>0.629</b>	<b>0.288</b>	0.220	<b>0.513</b>
Cerebras-GPT	2.7B	<b>1.1e21</b>	1.834	0.386	0.701	0.559	0.567	0.571	0.246	0.206	0.462
GPT-J	6.1B	1.7e22	<b>1.613</b>	<b>0.518</b>	0.752	0.640	<b>0.683</b>	<b>0.670</b>	<b>0.340</b>	<b>0.288</b>	<b>0.556</b>
OPT	6.7B	1.4e22	-	0.505	<b>0.763</b>	<b>0.654</b>	0.677	0.656	0.307	0.276	0.548
Pythia	6.9B	1.4e22	1.626	0.482	0.746	0.611	0.679	0.669	0.323	0.270	0.540
Cerebras-GPT	6.7B	<b>6.3e21</b>	1.704	0.447	0.739	0.602	0.636	0.643	0.282	0.238	0.512
OPT	13B	2.7e22	-	<b>0.524</b>	0.759	<b>0.651</b>	0.687	0.671	0.329	0.270	0.556
Pythia	12B	2.4e22	1.582	0.505	0.761	0.645	<b>0.705</b>	0.700	0.336	0.284	0.562
Cerebras-GPT	13B	<b>2.3e22</b>	<b>1.572</b>	0.513	<b>0.766</b>	0.646	0.696	<b>0.714</b>	<b>0.367</b>	<b>0.286</b>	<b>0.570</b>
GPT-NeoX	20B	<b>6.4e22</b>	<b>1.519</b>	<b>0.535</b>	<b>0.779</b>	<b>0.661</b>	<b>0.720</b>	<b>0.723</b>	<b>0.380</b>	<b>0.290</b>	<b>0.584</b>
	2.8B	6.1e21	1.724	0.466	0.743	0.612	0.672	0.662	0.299	0.232	0.526
Pythia	6.9B	1.4e22	1.644	0.488	0.756	0.636	0.695	0.667	0.320	0.252	0.545
Pile-dedup	12B	2.4e22	1.601	0.516	0.761	0.639	0.712	0.697	0.341	0.280	0.564

Finally, Table 2 shows more detailed downstream task comparisons for large publicly-available models, grouped into comparable sizes. We bold the results that are the best for each task and model size group. Each model family has at least one model that is best for some tasks. In this table, we also include results for Pythia models trained on a deduplicated version of the Pile. We separated these results, since they may not be directly comparable to others above, which were trained using the same or similar dataset preparation. As expected from the deduplication process, Pythia models show more difficulty generalizing to the pre-training Pile test loss task than other open models, which might have seen duplicated data during training. However, the Pythia Pile-dedup models typically improve accuracy on downstream tasks (1.8% on average), indicating the potential benefits of deduplication.

### 3.3 Maximal Update Parameterization ( $\mu$ P) and $\mu$ Transfer

As we scaled the Cerebras-GPT models with standard parameterization (SP) along our scaling law, we experienced challenges predicting appropriate hyperparameters, and these models show substantial variance around their common scaling law. To address these challenges, we also test  $\mu$ P and  $\mu$ Transfer tuned hyperparameters to 111M–2.7B parameter Cerebras-GPT models. Across model sizes, our  $\mu$ P models exhibit an average of 0.43% improved Pile test loss and 1.7% higher average downstream task accuracy compared to our SP models. Here, we also show that  $\mu$ P performance scales more predictably, enabling more accurate performance extrapolation.

As we scaled up models with SP, we found model training could become unstable when configured with hyperparameters used in other prior works. At different model size scales, the numerical characteristics of different layers can training instability. These instabilities can lead the practitioner to adjust prior hyperparameters in an effort to work around the issues<sup>5</sup>. However, moving away from known good configurations can lead to costly tuning efforts and blocked scaling progress. By shifting to  $\mu$ P, we find more stable training dynamics—key metrics like weight and gradient norms behave similarly at different scales.

We see the benefits of  $\mu$ P readily as we scale. First, after tuning hyperparameters with a small 40M parameter model, we were able to use the same learning rate hyperparameters for all model scales, as we noted in Table 1.  $\mu$ P features were the only changes we made to these models, so scaling was very simple.

<sup>5</sup>Appendix A.2 describes example stability challenges, such as FP16 mixed precision training causing numerical underflows.

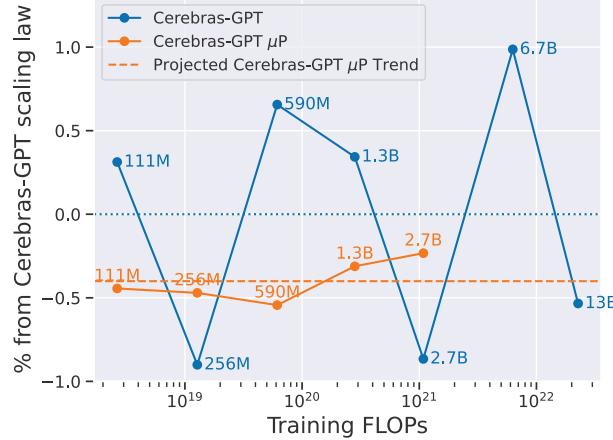


Figure 5: Percentage loss increase relative to Cerebras-GPT scaling law plotted against training FLOPs.

Second, models show significantly more predictable scaling. Figure 5 plots the percentage loss increase for each SP and μP model relative to the SP scaling law (negative values are improved loss). μP models show an average of 0.43% better Pile test loss compared to the Cerebras-GPT SP scaling law fit. Further, μP models show substantially lower variance with just 0.04% standard deviation relative to the SP scaling law, while SP models show deviation 0.66% ( $\sim 16\times$  more noisy). For perspective, the run-to-run standard deviation in loss when using different initialization and data random seeds is around 0.35%.

Table 3: Pile pre-training test loss and zero-shot downstream task results for μP and SP models.

Model	Pre-train Pile test xent (↓)	Downstream task accuracy (↑)							Down- stream Average	
		Hella- Swag	PIQA	Wino- Grande	Lambada	ARC-e	ARC-c	Open- BookQA		
Cerebras-GPT	111M	2.608	<b>0.268</b>	0.594	0.488	0.194	0.380	0.166	0.118	0.315
Cerebras-GPT + μP	111M	<b>2.588</b>	<b>0.268</b>	<b>0.598</b>	<b>0.519</b>	<b>0.204</b>	<b>0.390</b>	<b>0.176</b>	<b>0.124</b>	<b>0.325</b>
Cerebras-GPT	256M	<b>2.349</b>	<b>0.274</b>	0.613	<b>0.511</b>	<b>0.293</b>	0.410	0.170	<b>0.158</b>	0.347
Cerebras-GPT + μP	256M	2.359	<b>0.274</b>	<b>0.617</b>	0.505	0.287	<b>0.427</b>	<b>0.194</b>	0.156	<b>0.351</b>
Cerebras-GPT	590M	2.181	0.291	0.627	0.498	<b>0.366</b>	0.464	0.190	0.158	0.370
Cerebras-GPT + μP	590M	<b>2.155</b>	<b>0.295</b>	<b>0.644</b>	<b>0.517</b>	0.362	<b>0.470</b>	<b>0.194</b>	<b>0.172</b>	<b>0.379</b>
Cerebras-GPT	1.3B	1.997	0.325	0.664	<b>0.521</b>	0.462	0.508	<b>0.224</b>	0.166	0.410
Cerebras-GPT + μP	1.3B	<b>1.984</b>	<b>0.334</b>	<b>0.682</b>	0.512	<b>0.471</b>	<b>0.515</b>	0.223	<b>0.196</b>	<b>0.419</b>
Cerebras-GPT	2.7B	<b>1.834</b>	0.386	<b>0.701</b>	<b>0.559</b>	<b>0.567</b>	<b>0.571</b>	<b>0.246</b>	0.206	<b>0.462</b>
Cerebras-GPT + μP	2.7B	1.846	<b>0.388</b>	0.697	0.557	0.558	0.569	0.241	<b>0.218</b>	0.461

In addition to its pre-training advantages, μP also improves downstream capabilities of these models. In the previous Figure 4, we plotted downstream results for μP models, where we see improved accuracy and distinctively smoother scaling than SP models. Table 3 also lists these zero-shot downstream results for SP and μP models. In particular, μP models show a 1.7% relative improvement in downstream tasks on average. These results are robust across model scales besides the 2.7B parameter model. We believe that we were just lucky when choosing the SP 2.7B model hyperparameters such that it performs significantly better than the SP Pile scaling law. Despite the SP model’s upstream advantage, however, the 2.7B + μP model still performs as well on downstream tasks on average.

#### 4 Trading Off Training and Inference FLOPs

Up to this point, our analysis has focused on compute-optimal pre-training, where compute cost is proportional to the square of the model’s size, because we train models to a constant number of tokens per parameter. However, recent work has started to also consider model inference costs, showing smaller models trained on more tokens can still significantly improve loss (Hoffmann et al., 2022; Touvron et al., 2023). At

inference time, the compute cost is proportional to the model’s size and number of inferences. Thus, smaller models will have an overall inference cost advantage proportional to their size.

We propose a technique to identify training+inference compute-optimal frontiers that practitioners can use to estimate how they should pre-train their models when considering inference deployment costs. Specifically, we define a compute cost metric equal to pre-training FLOPs added to the model’s inference cost and the expected number of inference tokens. Here,  $F$  is the total compute cost,  $f$  represents FLOPs costs for full pre-training and per-token inference,  $n_{\text{infer\_tokens}}$  is the number of expected inference tokens for the given model, and  $p$  is the model’s parameter count<sup>6</sup>:

$$F = f_{\text{pre-train\_total}} + n_{\text{infer\_tokens}} \cdot f_{\text{infer/token}} \\ \propto \mathcal{O}(p^2) + n_{\text{infer\_tokens}} \cdot \mathcal{O}(p) \quad (2)$$

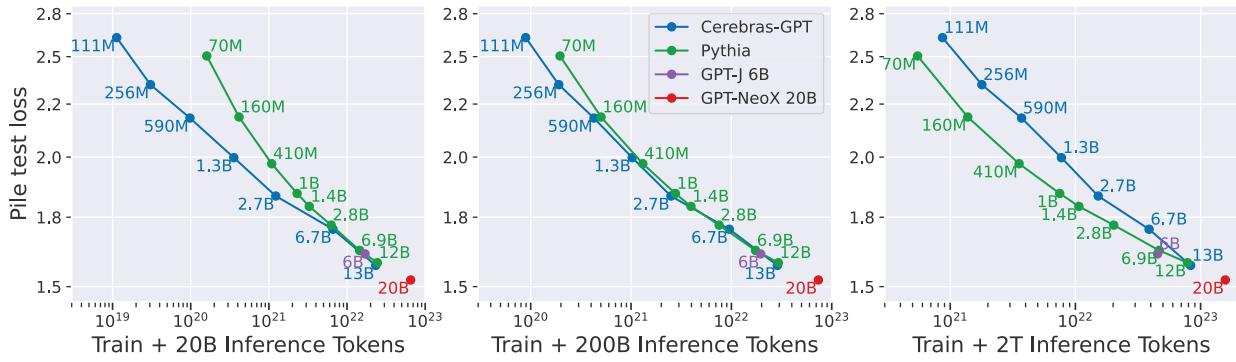


Figure 6: Pile test loss when accounting for both pre-training and expected inference FLOPs. Plots account for 20B (left), 200B (middle), 2T (right) tokens inference.

With this formulation, we can estimate the number of model inferences before the total compute budget matches models trained on fewer or more tokens. Figure 6 plots a comparison of total pre-train + inference compute cost for Cerebras-GPT, GPT-J, GPT-NeoX, and Pythia models assuming either 20B, 200B, or 2T inference tokens. These results show that most Cerebras-GPT models would provide better Pile test-loss-per-compute-FLOP than Pythia models until all models reach roughly 200B inference tokens. Since this total compute metric forms a continuum trade-off, models pre-trained on some number of tokens in between the Cerebras-GPT and Pythia frontiers are likely to achieve better loss for the same total compute budget.

Following this formulation, organizations and governments can better assess the total costs when budgeting large-scale training runs. Specifically, if a model is to be trained in a pre-training compute-*inefficient* way using too many data samples, that model may need to be used in a very large number of inferences before the training compute cost can be amortized and well-justified. Similar analysis can be applied to monetary, energy, or carbon footprint costs as well. We encourage the community to consider these total costs when training future models.

## 5 Cerebras Stack

To collect our compute-efficient LLM scaling laws, we run all studies on the Cerebras Wafer-Scale Cluster named “Andromeda”, which contains 16 Cerebras CS-2 systems. As far as we are aware, this is the first scaling laws study performed on Cerebras systems, which are capable of simple large-scale model training and high-performance scale-out to many systems. In this section, we describe the Andromeda AI Supercomputer, and the Cerebras software platform (CSoft) that we use for scaling and training. We show that Andromeda performance scales linearly up to the full 16 CS-2s, and we describe the simplicity of training models for this study.

<sup>6</sup>Note that the big- $\mathcal{O}$  order relations here could incorporate constant factors to account for model compression, quantization, or other techniques that decrease the relative inference costs.

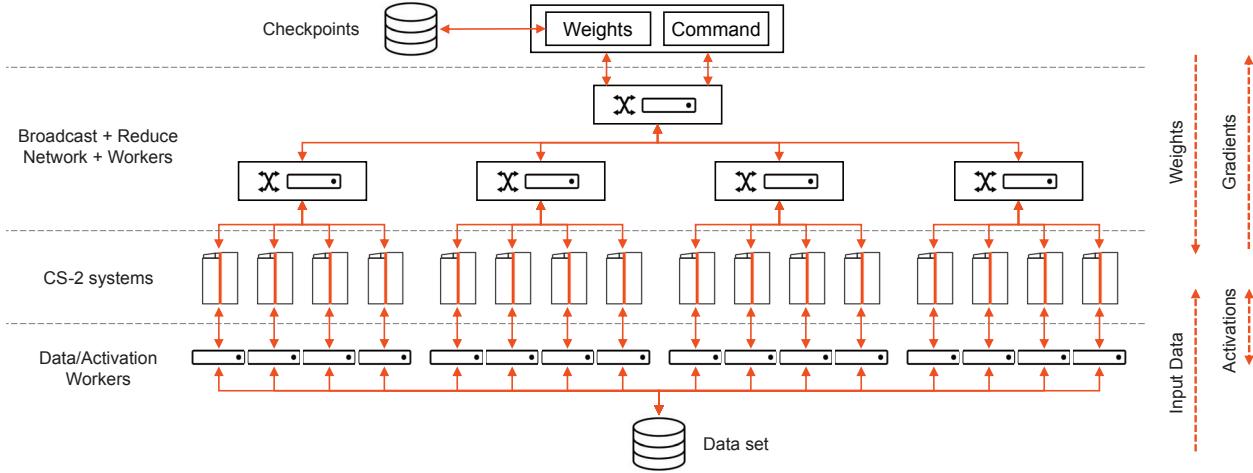


Figure 7: Andromeda AI Supercomputer: logical architecture of the Cerebras Wafer-Scale Cluster.

### 5.1 Andromeda AI Supercomputer

Andromeda is a Cerebras Wafer-Scale Cluster composed of 16 CS-2 systems. Figure 7 shows the architecture of Andromeda, which aligns well with the large-scale parallel nature of deep learning training. Each CS-2 system contains a Cerebras Wafer-Scale Engine (WSE-2) processor, which has 40 GB of high bandwidth SRAM and compute capability of 7.5 PetaFLOP/s half precision peak throughput. The WSE-2's processing cores are specifically designed to perform all compute operations required for deep learning models. Overall, Andromeda has peak throughput of 120 PFLOP/s from these CS-2s.

Weights and command servers drive the CS-2's computation by broadcasting the weights and control instructions through a broadcast + reduce tree network. This same network collects and reduces gradients from the CS-2s for each training step. When weights servers receive the reduced gradients, they perform the optimizer step and update model weights. They also save and restore model checkpoints to/from disk.

Activation workers act as servers to handle input data and activations. Each worker reads an independent shard of the dataset from disk and creates subbatches to send to a corresponding CS-2 for training. In cases where models must be trained using activation checkpointing, the CS-2 can evict the activations to the corresponding activation worker, which can later refill the activation on the CS-2 when needed.

### 5.2 CSoft Platform and Weight Streaming Mode

Andromeda runs deep learning applications through the Cerebras Software Platform (CSoft). For this study, we write and train models in both Tensorflow and PyTorch (reported results are with PyTorch), and CSoft compiles and orchestrates running these models on the hardware. In this process, CSoft automatically selects things like data parallel subbatch sizing and gradient accumulation, activation recomputation and checkpointing, and appropriate data layouts and kernel configurations for high performance.

The logical data flow in Figure 7 is called the Weight Streaming mode, because weight servers stream the weights to the CS-2s and collect gradients on each training step. This execution mode permits training models of size only limited by the memory capacity of weight servers, and we have tested the ability to train beyond the full GPT-3 175B parameter model with no changes outside of model configurations.

The Weight Streaming design stands in contrast with existing accelerator execution modes. Recent trends in large language model training typically require parallelizing training across tens to thousands of accelerator devices, such as GPUs. These efforts require complicated combinations of data and model parallelism (e.g., (Smith et al., 2022)). Models must be carefully divided to fit into memory close to the devices to achieve high throughput at relatively small per-device batch sizes. Weight Streaming permits moving the weights to the wafer and gradients from the wafer—achieving solid performance at small per-system batch sizes—without the need for model parallelism.

## Our Language Model Scaling Experience

We find CSoft Weight Streaming to be significantly easier to develop and scale models than existing accelerator approaches. First, we were able to run each Cerebras-GPT model and even larger models for many training steps on a single CS-2 system. This capability made it easy to quickly test that features of our model and dataset loader implementations would work well even for very large models. Second, the cluster’s near-linear performance scaling meant that we could accurately estimate total training time for each run as we scaled to more CS-2 systems. Finally, it was easy to configure these large-scale runs; Scaling to many CS-2 systems requires changing only the number of systems on which to train, and CSoft automatically chooses the data parallel configurations for us.

### 5.3 Performance Scalability

Andromeda provides near-linear performance scaling up to the full 16 CS-2s. We show performance (training speed) scaling from our initial model tests, followed by performance scaling results from our actual training runs. First, as Andromeda came online, we tested performance using a weak scaling approach: As we increased the number of systems, we increase the batch size proportionally (here, batch size is number of sequences of length 2048). We ran 100 training steps for each configuration and take an average training step time over the 100 steps. Table 4 shows the weak scaling performance relative to 1 CS-2. Andromeda achieves linear scaling within 9% for all model sizes and CS-2 system counts.

Table 4: Andromeda weak scaling tests show linear performance scaling up to 16 CS-2s

Model	Sequence Length	Per CS-2 Batch Size	Performance relative to 1 CS-2			
			2 CS-2s	4 CS-2s	8 CS-2s	16 CS-2s
GPT-3 XL 1.3B	2,048	121	1.99x	3.94x	7.87x	15.50x
GPT-3 XL 1.3B	10,000	33	1.99x	3.97x	7.95x	15.87x
GPT-3 2.7B	2,048	121	1.98x	3.91x	7.86x	15.62x
GPT-3 6.7B	2,048	85	1.99x	3.89x	7.91x	15.45x
GPT-3 20B	2,048	50	1.92x	3.75x	7.93x	15.32x
GPT-J 6B	2,048	65	1.97x	3.65x	7.69x	14.52x
GPT-NeoX 20B	2,048	50	1.98x	3.92x	8.05x	15.45x

We also show that Andromeda achieves high utilization even when strong scaling on batch size. We choose to scale out the fixed batch sizes from our training runs across different numbers of Andromeda systems. When running on fewer CS-2s, if the per-CS-2 batch size requires too much memory to fit in each WSE-2’s on-wafer SRAM, the software stack automatically selects a smaller per-CS-2 batch size and accumulates gradients up to the user’s chosen batch size. Table 5 lists the relative performance compared to running on a single CS-2. These results show consistent performance scalability for the batch sizes commonly chosen for these models.

Table 5: Strong scaling performance for batch sizes used to train larger models. To get to the user’s full batch size, CSoft uses data parallelism across systems and gradient accumulation.

Model	Batch Size	Performance relative to 1 CS-2 (per CS-2 batch)			
		1 CS-2	2 CS-2s	4 CS-2s	8 CS-2s
1.3B	528	1.0x (132)	1.99x (132)	3.97x (132)	7.10x (66)
2.7B	528	1.0x (88)	1.99x (88)	3.77x (66)	7.43x (66)
6.7B	1,040	1.0x (65)	1.99x (65)	3.97x (65)	7.90x (65)
13B	1,040	1.0x (65)	1.99x (65)	3.95x (65)	7.84x (65)

Finally, as we increased model sizes along our scaling law, we tested and compared the cluster’s FLOP/s utilization for each training run. Table 6 lists Andromeda’s utilization relative to the 111M parameter model running on one CS-2. Performance deviates by less than 8% at all model scales. In addition to robust scaling across many machines, these results indicate consistent performance across a range of model and batch sizes.

Table 6: Andromeda FLOP/s utilization relative to 1 CS-2 training the 111M parameter model. Here, larger values mean higher utilization.

Model	Batch Size	Number of CS-2s	Per CS-2 Batch Size	Relative Utilization
111M	120	1	120	1.00
256M	264	1	264	1.00
590M	264	1	264	0.92
1.3B	528	4	132	0.96
2.7B	528	4	132	0.96
6.7B	1040	16	65	1.05
13B	1080	12	45	1.02

## 6 Related Work

Early deep learning scaling law studies show that when scaling dataset and model size, loss improves predictably (Hestness et al., 2017; Kaplan et al., 2020). These studies indicate generally that scaling could give substantial modeling improvements. From this observation, many organizations scaled to train the largest possible models on their available infrastructure: GPT-3 175B (Brown et al., 2020), Jurassic-1 178B (Lieber et al., 2021), Gopher 280B (Rae et al., 2022), HyperCLOVA 82B (Kim et al., 2021), Ernie 3.0 Titan 260B (Wang et al., 2021), Yuan 1.0 (Wu et al., 2021), PanGu- $\alpha$  (Zeng et al., 2021), Megatron-Turing NLG 530B (Smith et al., 2022), PaLM 540B (Chowdhery et al., 2022), and LaMDA 137B (Thoppilan et al., 2022). These models show significant performance improvement on many downstream tasks compared to prior language models. However, these studies typically only scale model size without scaling the dataset size as suggested by the early works, often training on roughly 300B tokens. Further, these models could only be trained by select organizations with large compute clusters, and the datasets and resulting pre-trained models have not been released publicly for analysis by the research community.

The research community *has* released large datasets and pre-trained models—typically much smaller than the largest models above but still quite valuable—and we have noted many of them previously: GPT-J, GPT-Neo, GPT-NeoX, OPT, and Pythia. Another notable work that releases dataset and model is the Big Science collaborative effort to train BLOOM 176B (Scao et al., 2022; 2023). These studies, datasets, and models enable the community to test, compare, and use large language models they would otherwise not have access to or compute budget to train.

In 2022, studies started revisiting early scaling works to note that although model size scaling improves performance, consistently scaling the dataset size is still critical to get the best possible models. Hoffmann et al. (2022) show that for compute-optimal pre-training, the dataset size should grow linearly with transformer model size in parameters, and they scaled training up to a 70B parameter model on 1.4T tokens. The dataset and models are not publicly-available, so our work aims to reproduce these results to offer to the community an open and reproducible scaling law. Recently, the LLaMa paper (Touvron et al., 2023) also reproduces large models trained on large open datasets to improve pre-training. Although these models show strong performance, most are trained in a compute-inefficient way by training on larger datasets than would be compute-optimal for the given model sizes. LLaMa models are available through request. The resulting models trained in these works perform better than prior larger models trained on smaller datasets.

Large language model training is prone to instability, and it is very costly when large model training runs fail due to instability. Various techniques have been developed to control training dynamics and train models stably (Glorot & Bengio, 2010; Yang & Schoenholz, 2017; Schoenholz et al., 2017; Yang & Schoenholz, 2018; Zhang et al., 2019; Bachlechner et al., 2020; Huang et al., 2020; Liu et al., 2020; Li et al., 2022).  $\mu$ P is the first comprehensive method to analytically control width-related training instabilities and allow optimal hyperparameters of small models to be the same as optimal hyperparameters for very large models. We find that the comprehensive nature of  $\mu$ P simplifies our training efforts, so we feel it is useful to share our experience and encourage the community to use it rather than considering combinations of other techniques.

## 7 Limitations

In this work, we train well-established model architectures to create foundation models, but we did not explore recent architectural features, downstream task tuning procedures, or dataset cleaning approaches used in contemporary works. Model features worth exploring in future work include position embeddings, such as RoPE (Su et al., 2022) and ALiBi (Press et al., 2022), and activation functions, like SwiGLU (Shazeer, 2020). There are also training paradigms worth exploring, such as denoising pre-training objectives (Tay et al., 2023) and instruction fine-tuning (Ouyang et al., 2022). Finally, we expect that further dataset cleaning can further improve pre-trained models. For instance, our testing in Appendix C.2 shows that the Pythia models improve downstream task accuracy when trained on a deduplicated version of the Pile.

We have not yet tested Cerebras-GPT models extensively in downstream tasks or in real application settings. Specifically, we have not tested for factual accuracy, profanity, toxicity, or other socially undesirable text generation. We do evaluate the bias of our Cerebras-GPT models using the CrowS-Pairs dataset in Appendix C.4. Further safety-related testing, mitigations, and output curation should be applied to our pre-trained models before presenting results to users. Please refer to the model card in the Appendix, Table 7.

## 8 Conclusion

In this paper, we introduce Cerebras-GPT, a family of open models scaled from 111M to 13B parameters and pre-trained in a compute-optimal way on the Pile dataset. These models show state-of-the-art pre-training efficiency on pre-training and downstream objectives when compared to other open-source models. We believe this is the first such open effort, and we provide detailed instructions to reproduce our results and we release our pre-trained model checkpoints<sup>7</sup>. We combine this scaling with  $\mu$ P, a comprehensive technique to improve large model stability, and we show it further improves our scaling results. We document our experience training these models on the Andromeda AI Cluster, comprising 16 Cerebras CS-2 systems, and we describe the simplicity of scaling models and performance.

## References

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Gregory S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems, 2016. URL <http://arxiv.org/abs/1603.04467>.

Thomas Bachlechner, Bodhisattwa Prasad Majumder, Huanru Henry Mao, Garrison W. Cottrell, and Julian McAuley. ReZero is All You Need: Fast Convergence at Large Depth, 2020. URL <https://arxiv.org/abs/2003.04887>.

Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. Pythia: A Suite for Analyzing Large Language Models Across Training and Scaling, 2023. URL <https://arxiv.org/abs/2304.01373>.

Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. PIQA: Reasoning about Physical Commonsense in Natural Language. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.

Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, et al. GPT-NeoX-20B: An Open-Source Autoregressive Language Model. In *Proceedings of the ACL Workshop on Challenges & Perspectives in Creating Large Language Models*, 2022.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, 2020.

<sup>7</sup>Pre-trained models are available on HuggingFace: <https://huggingface.co/cerebras>. Source code is available in the Cerebras Modelzoo: <https://github.com/Cerebras/modelzoo>.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. PaLM: Scaling Language Modeling with Pathways, 2022. URL <https://arxiv.org/abs/2204.02311>.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge, 2018. URL <https://arxiv.org/abs/1803.05457>.

Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, et al. Scaling Vision Transformers to 22 Billion Parameters, 2023. URL <https://arxiv.org/abs/2302.05442>.

Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The Pile: An 800GB Dataset of Diverse Text for Language Modeling, 2020. URL <https://arxiv.org/abs/2101.00027>.

Leo Gao, Jonathan Tow, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Kyle McDonell, Niklas Muennighoff, et al. A Framework for Few-shot Language Model Evaluation, 2021.

Xavier Glorot and Yoshua Bengio. Understanding the Difficulty of Training Deep Feedforward Neural Networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (PMLR)*, 2010.

Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md. Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. Deep Learning Scaling is Predictable, Empirically, 2017. URL <https://arxiv.org/abs/1712.00409>.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. An Empirical Analysis of Compute-optimal Large Language Model Training. In *The Conference on Neural Information Processing Systems (NeurIPS)*, 2022.

Xiao Shi Huang, Felipe Perez, Jimmy Ba, and Maksims Volkovs. Improving Transformer Optimization Through Better Initialization. In *Proceedings of the 37th International Conference on Machine Learning*, 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. URL <https://arxiv.org/abs/2001.08361>.

Boseop Kim, HyoungSeok Kim, Sang-Woo Lee, Gichang Lee, Donghyun Kwak, Jeon Dong Hyeon, Sunghyun Park, Sungju Kim, Seonhoon Kim, Dongpil Seo, et al. What Changes Can Large-scale Language Models Bring? Intensive Study on HyperCLOVA: Billions-scale Korean Generative Pretrained Transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021.

Conglong Li, Minjia Zhang, and Yuxiong He. The Stability-Efficiency Dilemma: Investigating Sequence Length Warmup for Training GPT Models. In *Advances in Neural Information Processing Systems*, 2022.

Opher Lieber, Or Sharir, Barak Lenz, and Yoav Shoham. Jurassic-1: Technical Details And Evaluation, 2021. URL [https://uploads-ssl.webflow.com/60fd4503684b466578c0d307/61138924626a6981ee09caf6\\_jurassic\\_tech\\_paper.pdf](https://uploads-ssl.webflow.com/60fd4503684b466578c0d307/61138924626a6981ee09caf6_jurassic_tech_paper.pdf).

Liyuan Liu, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, and Jiawei Han. Understanding the Difficulty of Training Transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.

Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*, 2017.

Sam McCandlish, Jared Kaplan, Dario Amodei, and OpenAI Dota Team. An Empirical Model of Large-Batch Training, 2018. URL <https://arxiv.org/abs/1812.06162>.

Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. Mixed Precision Training. In *International Conference on Learning Representations*, 2018.

Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a Suit of Armor Conduct Electricity? A New Dataset for Open Book Question Answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.

Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 2019.

Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. CrowdPairs: A Challenge Dataset for Measuring Social Biases in Masked Language Models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training Language Models to Follow Instructions With Human Feedback, 2022. URL <https://arxiv.org/abs/2203.02155>.

Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. The LAMBADA dataset: Word Prediction Requiring a Broad Discourse Context. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2016.

Ofir Press, Noah Smith, and Mike Lewis. Train Short, Test Long: Attention with Linear Biases Enables Input Length Extrapolation. In *International Conference on Learning Representations*, 2022.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language Models are Unsupervised Multitask Learners, 2019. URL <https://d4mucfpksyv.wcloudfront.net/better-language-models/language-models.pdf>.

Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling Language Models: Methods, Analysis & Insights from Training Gopher, 2022. URL <https://arxiv.org/abs/2112.11446>.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. WinoGrande: An Adversarial Winograd Schema Challenge at Scale. *Communications of the ACM*, 2021.

Teven Le Scao, Thomas Wang, Daniel Hesslow, Lucile Saulnier, Stas Bekman, M Saiful Bari, Stella Biderman, Hady Elsahar, Niklas Muennighoff, Jason Phang, et al. What Language Model to Train if You Have One Million GPU Hours?, 2022. URL <https://arxiv.org/abs/2210.15424>.

Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. BLOOM: A 176B-Parameter Open-Access Multilingual Language Model, 2023. URL <https://arxiv.org/abs/2211.05100>.

Samuel S. Schoenholz, Justin Gilmer, Surya Ganguli, and Jascha Sohl-Dickstein. Deep Information Propagation. In *International Conference on Learning Representations*, 2017.

Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2016.

Chris Shallue, Jaehoon Lee, Joseph Antognini, Jascha Sohl-dickstein, Roy Frostig, and George Dahl. Measuring the Effects of Data Parallelism on Neural Network Training. *Journal of Machine Learning Research (JMLR)*, 2018.

Noam Shazeer. GLU Variants Improve Transformer, 2020. URL <https://arxiv.org/abs/2002.05202>.

Shaden Smith, Mostafa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, A Large-Scale Generative Language Model, 2022. URL <https://arxiv.org/abs/2201.11990>.

Robyn Speer. ftfy, 2019. Version 5.5.

Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. RoFormer: Enhanced Transformer with Rotary Position Embedding, 2022. URL <https://arxiv.org/abs/2104.09864>.

Yi Tay, Mostafa Dehghani, Vinh Q. Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Siamak Shakeri, Dara Bahri, Tal Schuster, et al. UL2: Unifying Language Learning Paradigms, 2023. URL <http://arxiv.org/abs/2205.05131>.

Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. LaMDA: Language Models for Dialog Applications, 2022. URL <https://arxiv.org/abs/2201.08239>.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. LLaMA: Open and Efficient Foundation Language Models, 2023. URL <https://arxiv.org/abs/2302.13971>.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *Advances in Neural Information Processing Systems*, 2017.

Ben Wang and Aran Komatsuzaki. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>, 2021.

Shuohuan Wang, Yu Sun, Yang Xiang, Zhihua Wu, Siyu Ding, Weibao Gong, Shikun Feng, Junyuan Shang, Yanbin Zhao, Chao Pang, et al. ERNIE 3.0 Titan: Exploring Larger-scale Knowledge Enhanced Pre-training for Language Understanding and Generation, 2021. URL <https://arxiv.org/abs/2112.12731>.

Shaohua Wu, Xudong Zhao, Tong Yu, Rongguo Zhang, Chong Shen, Hongli Liu, Feng Li, Hong Zhu, Jiangang Luo, Liang Xu, and Xuanwei Zhang. Yuan 1.0: Large-Scale Pre-trained Language Model in Zero-Shot and Few-Shot Learning. *ArXiv*, abs/2110.04725, 2021. URL <https://arxiv.org/abs/2110.04725>.

Greg Yang and Sam Schoenholz. Mean Field Residual Networks: On the Edge of Chaos. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017.

Greg Yang and Sam S. Schoenholz. Deep Mean Field Theory: Layerwise Variance and Width Variation as Methods to Control Gradient Explosion, 2018. URL <https://openreview.net/forum?id=rJGY8GbR->.

Greg Yang, Edward Hu, Igor Babuschkin, Szymon Sidor, Xiaodong Liu, David Farhi, Nick Ryder, Jakub Pachocki, Weizhu Chen, and Jianfeng Gao. Tuning Large Neural Networks via Zero-Shot Hyperparameter Transfer. In *Advances in Neural Information Processing Systems*, 2021.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. HellaSwag: Can a Machine Really Finish Your Sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019.

Wei Zeng, Xiaozhe Ren, Teng Su, Hui Wang, Yi Liao, Zhiwei Wang, Xin Jiang, ZhenZhang Yang, Kaisheng Wang, Xiaoda Zhang, et al. PanGu- $\alpha$ : Large-scale Autoregressive Pretrained Chinese Language Models with Auto-parallel Computation, 2021. URL <https://arxiv.org/abs/2104.12369>.

Hongyi Zhang, Yann N. Dauphin, and Tengyu Ma. Residual Learning Without Normalization via Better Initialization. In *International Conference on Learning Representations*, 2019.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuhui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. OPT: Open Pre-trained Transformer Language Models, 2022. URL <https://arxiv.org/abs/2205.01068>.

## Appendix

### Model Card

Table 7 shows the model card for the largest Cerebras-GPT model following guide in (Mitchell et al., 2019).

Table 7: Cerebras-GPT 13B Parameter Model Card

<b>Release details</b>
<ul style="list-style-type: none"> <li><b>Organization:</b> Cerebras Systems</li> <li><b>Model date:</b> March 2023</li> <li><b>Model type:</b> Autoregressive Transformer Language Model (more details in Section 2.3)</li> <li><b>Feedback on the model:</b> Nolan Dey and Joel Hestness, {nolan, joel}@cerebras.net</li> </ul>
<b>Model details</b>
<ul style="list-style-type: none"> <li><b>Model architecture:</b> Cerebras-GPT 13B is an autoregressive transformer decoder-only model with 13 billion parameters. The architecture is similar to GPT-2 and GPT-3. More details in Section 2.1.</li> <li><b>Hidden size:</b> 5,120</li> <li><b>Number of layers:</b> 40</li> <li><b>Head size:</b> 128</li> <li><b>Filter size:</b> 20,480</li> <li><b>Context (sequence) length:</b> 2,048</li> <li><b>Initialization:</b> Model is trained from randomly initialized weights. The base variant uses standard parameterization initialization (see Section 2.4).</li> <li><b>Release license:</b> Apache 2.0</li> </ul>
<b>Data Overview</b>
<ul style="list-style-type: none"> <li><b>Training data:</b> Cerebras-GPT is trained on the Pile dataset (Gao et al., 2020)</li> <li><b>Pre-processing:</b> Pile was cleaned using ftfy library to normalize text, and then filtered using scripts provided by Eleuther. Then, data was tokenized with byte-pair encoding using the GPT-2 vocabulary.</li> <li><b>Evaluation data:</b> Upstream (pre-training) evaluations were completed using the Pile validation and test set splits. Downstream evaluations were performed on standardized tests. Cloze and completion tasks: LAMBADA, HellaSwag. Common Sense Reasoning tasks: PIQA, ARC, OpenBookQA. Winograd schema type tasks: Wino-grande. Downstream evaluations were performed using the Eleuther lm-eval-harness (Gao et al., 2021).</li> <li><b>Motivation:</b> Evaluation tasks were chosen to closely match related works and cover a broad cross-section of task types.</li> </ul>
<b>Usage</b>
<ul style="list-style-type: none"> <li><b>Primary intended uses:</b> The primary intended use is to further research into large language models. Model can be used as a foundation model for NLP, applications, ethics, and alignment research.</li> <li><b>Primary intended users:</b> Researchers who are working to improve LLMs and practitioners who are looking for reference implementations, training setups, hyperparameters, or pre-trained models.</li> <li><b>Limitations:</b> Due to financial and compute budgets, Cerebras-GPT models were only trained and evaluated following the approaches described in this document.</li> <li><b>Out-of-scope uses:</b> Further safety-related testing and mitigations should be applied before using the Cerebras-GPT model family in production downstream applications.</li> </ul>
<b>Metrics</b>
<ul style="list-style-type: none"> <li><b>Model performance measures:</b> Model is evaluated using text prediction cross-entropy on upstream tasks and text generation accuracy on downstream tasks. Results are compared against many publicly available large language models. Details can be found in Section 3.</li> <li><b>Uncertainty and variability:</b> Model is not evaluated for prediction uncertainty or calibration. Due to restricted compute budget, variability analysis was only performed for small variants of Cerebras-GPT models using multiple runs from different random initializations and data loader seeds to assess variance in task performance.</li> </ul>
<b>Ethical considerations</b>
<ul style="list-style-type: none"> <li><b>Data:</b> The Pile dataset has been thoroughly analyzed from various ethical standpoints, and the dataset is known to contain content considered toxic, gender biased, pejorative, racially sensitive, etc. Please refer to Pile dataset references.</li> <li><b>Human life:</b> The outputs from this model may or may not align with human values. The risk needs to be thoroughly investigated before deploying this model in a production environment where it can directly impact human life.</li> <li><b>Risks and harms:</b> There can be distributional bias in the Pile dataset that can manifest in various forms in the downstream model deployment. There are other risks associated with large language models such as amplifying social stereotypes, memorizing training data, or revealing private or secure information.</li> <li><b>Mitigations:</b> Only mitigations in standard Pile dataset pre-processing were employed when pre-training Cerebras-GPT.</li> </ul>
<b>Factors</b>
<ul style="list-style-type: none"> <li><b>Evaluation factors:</b> Cerebras-GPT was evaluated for various bias factors using the Crows-Pairs dataset task. Details are in Appendix C.4.</li> </ul>
<b>Implementation infrastructure</b>
<ul style="list-style-type: none"> <li><b>Hardware:</b> Andromeda AI Supercomputer: Cerebras Wafer-Scale Cluster with 16 Cerebras CS-2 systems</li> <li><b>Software:</b> PyTorch, Cerebras Software Platform (CSoft) release 1.8</li> </ul>

## Cerebras-GPT Open-Source References

We release our pre-trained models and code, so the community can use and reproduce our results. Pre-trained models are available on HuggingFace: <https://huggingface.co/cerebras>. We are initially releasing seven Cerebras-GPT models with 111M, 256M, 590M, 1.3B, 2.7B, 6.7B, and 13B parameters trained with standard parameterization (SP). These models are released under Apache 2.0 license, which permits commercial and non-commercial use. Source code is available in the Cerebras Modelzoo: <https://github.com/Cerebras/modelzoo>. We hope these models will be a valuable addition for the open-source community.

## Author Contributions and Acknowledgements

We would like to acknowledge the contributions of those who helped in preparation of this manuscript.

**Experimental planning and strategy:** Nolan Dey, Joel Hestness

**Model training:** Zhiming (Charles) Chen, Hemant Khachane, Ribhu Pathria, Gurpreet Gosal

**Dataloader development and dataset preparation:** Gurpreet Gosal

**Numerical configuration and validation:** Joel Hestness, Hemant Khachane, Gurpreet Gosal

**Upstream loss comparisons:** Gurpreet Gosal, Charles Chen

**Downstream task comparisons:** William Marshall

**Manuscript preparation:** Nolan Dey, Joel Hestness, Gurpreet Gosal, William Marshall

**Overall project leadership:** Joel Hestness, Marvin Tom

**Overall technical leadership:** Joel Hestness

In addition, we would like to thank others who helped in the preparation of this work. Bowen Yang and Faisal Al-Khateeb helped prepare the Pile dataset. We are also thankful for helpful feedback on the manuscript provided by Sean Lie, Anshul Samar, and Vithu Thangarasa. In general, we would also like to acknowledge the contributions of the many Cerebras engineers who made this work possible.

## A Methods Details

### A.1 Pile Dataset Preprocessing

We preprocess Pile using tools and instructions provided by Eleuther and the community. We clean the raw text data sources using the ftfy library to normalize text, including cleaning corrupted unicode (Speer, 2019). Our tokenized version of the Pile training set contains roughly 371B tokens (validation 380M, test 371M), similar to results reported in the GPT-NeoX paper (Black et al., 2022). The resulting tokenized dataset files contain contiguous samples from the raw text. For the best model generalization, we find it critical to shuffle samples across all training set documents, rather than shuffling within a window of even a few thousand documents. So, we also shuffle the training dataset across all documents as a final preprocessing step. This dataset-wide shuffling improves validation loss by 0.7-1.5% compared to aggressive shuffling settings over sets of contiguous documents as we tested with our dataloaders.

The Pile dataset has been thoroughly analyzed from various ethical standpoints, and the dataset is known to contain content considered toxic, gender biased, pejorative, racially sensitive, etc. Please refer to Pile dataset references for further information.

### A.2 Ensuring Stable Training

As we scaled up models to larger sizes, we encountered and resolved a few issues that improve training stability. We share some details here in hopes they assist others in their scaling efforts.

**Mixed Precision Training:** Initially, we trained models using FP16 mixed precision, a technique that carries model weights and activations in IEEE half precision floating-point (FP16) while performing dot-products and reductions in single precision 32-bit (FP32). This approach ensures that reductions maintain precision, while taking advantage of the smaller 16-bit data format for storing activations. Because FP16 has a significantly reduced exponent range compared to FP32, models need to be trained with loss scaling, an

approach that multiplies the gradients by a large positive value before back-propagation, and then divides out this multiplier just before applying the calculated gradients to the weights in the optimizer step. A dynamic approach to loss scaling sets the scale value by periodically testing larger values to see the largest scale such that the gradients do not overflow FP16.

We found that for models larger than roughly 1.3B parameters (hidden size 2048), loss scaling alone was not sufficient to ensure stable model training. As model weights grow during training, the gradients through layers like softmaxes can become eccentric, leading to large gradient values that tend to overflow. These large values push down the maximum allowed loss scale and cause other gradients to be very small. Very small gradient values have a tendency to underflow in FP16. Underflow can cause weights to receive either no gradient or low-precision, eccentric gradients, which can further exacerbate dynamic loss scale and underflow.

**Underflows and Weight Growth:** We detect underflows by observing any significant increase in the number of identically zero values in tensors as they go through cast operations from FP32 to FP16. Specifically, we find that attention layer softmax gradients are particularly susceptible to underflow. To fix this issue, we recommend carrying gradients in FP32 from the softmax back through the corresponding query-key dot-product and when calculating the gradients for the query and key projection weights and biases. We have tested various open-source mixed precision attention implementations that suffer this same issue.

We also find specific layers to be most susceptible to eccentric gradients caused by underflow. In the attention layers, the bias weights of the keys projection, specifically, have expected value close to zero early in training. If gradients to these weights partially underflow, the remaining gradients will be eccentric and large relative to their expectation. These K bias weights will tend to grow very quickly under these circumstances. We detect this issue by inspecting weight growth—measuring the weight standard deviation and norms—over many training steps compared to an implementation that uses FP32.

**Switching to bfloat16:** Another approach to avoid underflows is to use a larger exponent range for activation and gradient tensors. Brain floating-point (bfloat16) is a numerical format introduced by Google Brain and used in various hardware platforms to improve half precision floating point range. Specifically, bfloat16 has 8 bits of exponent compared to 5 bits for FP16. Typical bfloat16 model training implementations still use FP32 for intermediate values (mixed precision) in reduction operations to ensure mantissa precision.

Bfloat16 eliminates the need for dynamic loss scaling that is used with mixed precision, because the exponent range significantly reduces the likelihood of underflows. We find that although bfloat16 does not completely eliminate low-precision training dynamics concerns, it does significantly improve training stability, so we use bfloat16 for all final models that we train in this paper and release publicly. We find that our experience with bfloat16 training stability is consistent with prior works.

**Setting Adam Epsilon:** When gradients for a set of weights are small, using a relatively large Adam epsilon value can cause weights to grow slowly. This might be an appealing approach in the presence of large weight growth among weights that are expected to be small. However, a large Adam epsilon can cause very poor weights resolution and degrade model quality. Given the Adam update at step  $t$  on weights  $\theta$ :

$$\theta_t = \theta_{t-1} - \gamma m_t / (\sqrt{v_t} + \epsilon) \quad (3)$$

Here,  $m_t$  is the momentum, a running average of the gradient, and  $v_t$  is the velocity, a running average of the squared gradient. When gradients to a weight are small (e.g., in the case of K bias weights growth above),  $v_t$  will tend to be very small, because it is squared. In this case,  $\epsilon$  needs to be chosen to be small relative to each  $\sqrt{v_t}$ , or the Adam update denominator will be large, causing the weight updates to be small. As a rule-of-thumb, we find  $\epsilon$  should be less than  $\sqrt{\mu_v}/1000$ , where  $\mu_v$  is the mean of the velocity state weights, to ensure models do not suffer from stagnant weight growth. This analysis is how we choose to lower epsilon for our 6.7B and 13B parameter models.

## B Downstream Task Details

We evaluate our models on the following six downstream tasks in both the zero-shot and the few-shot setting. Here, we briefly describe each of the tasks: HellaSwag, PIQA, WinoGrande, Lambada, ARC, and OpenBookQA.

1. **HellaSwag** is a dataset of multiple choice questions aimed to test a model's common sense reasoning abilities (Zellers et al., 2019). For example,

A woman is outside with a bucket and a dog. The dog is running around trying to avoid a bath. She...

- A. rinses the bucket off with soap and blow dry the dog's head.
- B. uses a hose to keep it from getting soapy.
- C. gets the dog wet, then it runs away again.
- D. gets into a bath tub with the dog.

The authors of the dataset adversely select examples such that they are difficult for language models while still trivial for humans (with reported greater than 95% accuracy).

2. **PIQA** tests a model's common sense reasoning about the physical world by posing a prompt and two potential completions (Bisk et al., 2020). For example

```
[Goal] Make an outdoor pillow
[Sol1] Blow into a tin can and tie with rubber band
[Sol2] Blow into a trash bag and tie with rubber band
```

The model must choose which of the two continuations is more likely to follow from the prompt. Human performance on this dataset is approximately 95%.

3. **WinoGrande** consists of a set of pronoun resolution problems (Sakaguchi et al., 2021). Samples are constructed as pairs of similar sentences, each with a pronoun referring to a noun earlier in the sentence. The task is to predict which noun the pronoun refers to. For example, in the sample

- a. The trophy doesn't fit into the brown suitcase because it's too large.
- b. The trophy doesn't fit into the brown suitcase because it's too small.

in sentence (a), "it's" refers to "trophy", while in sentence (b), changing a single context word modifies the meaning of the sentence such that "it's" now refers to "suitcase".

4. **Lambada** is a word prediction task that tests a model's ability to understand text, with a particular emphasis on global context (Paperno et al., 2016). For example

Context: They tuned, discussed for a moment, then struck up a lively jig. Everyone joined in, turning the courtyard into an even more chaotic scene, people now dancing in circles, swinging and spinning in circles, everyone making up their own dance steps. I felt my feet tapping, my body wanting to move.  
 Target sentence: Aside from writing, I've always loved \_\_\_\_.  
 Target word: dancing

There are two versions of the Lambada dataset. The original version is that which was published in (Paperno et al., 2016). However, researchers more commonly use version of the dataset with slightly different formatting that was created by Radford et al. in order to evaluate their GPT-2 model (Radford et al., 2019). In our evaluations we use the latter version, referred to as "lambada\_openai" in the Eleuther eval harness (Gao et al., 2021).

5. **ARC** tests a model's ability to answer multiple choice science questions (Clark et al., 2018). For example

Which property of a mineral can be determined just by looking at it?  
 (A) luster [correct] (B) mass (C) weight (D) hardness

This dataset is split into an "easy" set and a "challenge" set where samples are selected for the challenge set if they are answered incorrectly by word co-occurrence and retrieval based algorithms.

6. **OpenBookQA** is a multiple choice common sense question answering dataset (Mihaylov et al., 2018). One example question from this dataset is

What is the most likely to be an effect of acid rain on an aquatic environment?

- increase in plant growth
- increase in fish population
- decrease in plant life
- cleaner and clearer water

## C Additional Results

### C.1 Pre-training Losses Throughout Training

In Figure 8, we show the intermediate Pile test losses achieved throughout training for Pythia and Cerebras-GPT models. For all model sizes and compute budgets, Cerebras-GPT models tend to have similar trajectory when approaching their final results along the scaling law. In contrast, Pythia models trained for more tokens per parameter follow less efficient trajectory, trending away from the scaling law and indicating their over-training. For Pythia models trained closer to 20 tokens per parameter, the trajectories align more closely with Cerebras-GPT models.

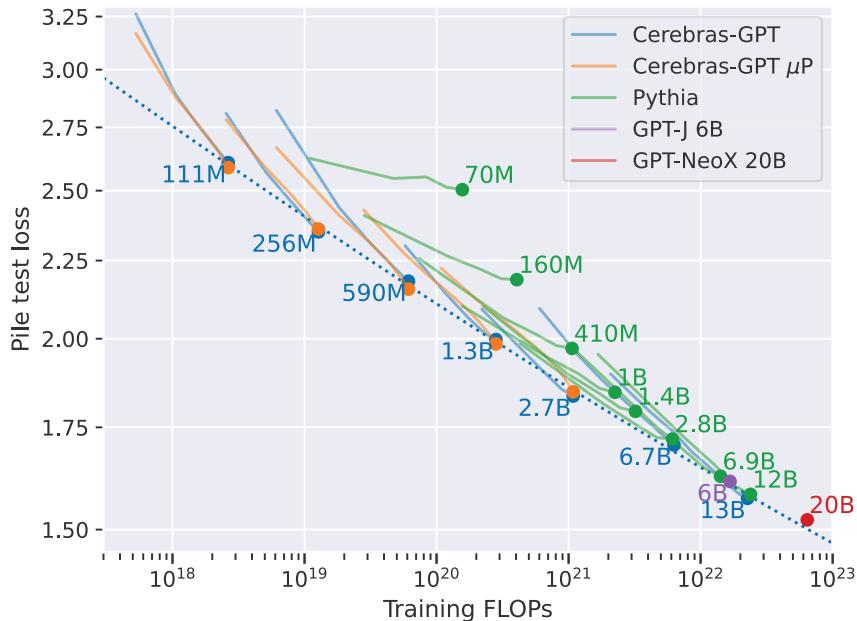


Figure 8: Pile test set loss given pre-training FLOPs throughout training for Cerebras-GPT and Pythia.

### C.2 Complete Downstream Task Testing

For completeness, we include all downstream task results we collected for this study. Table 8 includes upstream Pile evaluations and all downstream zero-shot tasks for models GPT-J, GPT-NeoX, OPT, and Pythia, as well as Cerebras-GPT. Similarly, Table 9 shows the few-shot (five-shot) results for all models. Full downstream results are plotted in Figures 9 and 11.

Some prior works also use a different methods to select model predictions when evaluating the model’s accuracy on some downstream tasks. Specifically, there are two commonly used techniques to select a model’s prediction. First, the model can predict the probability of an output (continuation) sequence given a context sequence. Here, the selection criteria would be to select the continuation with maximum probability. We use this maximum probability approach in all prior results in the paper to be consistent with results in the GPT-NeoX paper. The second approach is to normalize the model’s predicted probability in the log domain by the length of the continuation, and choose the continuation with the smallest length-normalized negative

log-likelihood (NLL) ( $\text{argmin}_i (-\ln(p_i)/|c_i|)$ , where  $p_i$  is the model’s predicted probability of continuation sequence  $c_i$ , and  $|c_i|$  is the length of that sequence). This approach will tend to favor longer continuations with moderate probability, which might be preferred for some tasks. For comparison against prior works that report minimum length-normalized NLL, we report Cerebras-GPT results in Tables 10 and 10.

### C.3 Differences Between Cerebras-GPT and Other Models

The Cerebras-GPT 13B model improves over other publicly-available models of comparable size. This is a surprising result given that the creators of these other models modified the original GPT-2/3 architecture intending to improve convergence and training efficiency. There are many confounders that could contribute to Cerebras-GPT’s advantages, but we briefly list known differences here to give an idea of the space of possible opportunities for future study.

- GPT-J, GPT-NeoX, and Pythia models use rotary positional embeddings, which show modest loss/accuracy improvements and ability to extend to longer sequence lengths (Su et al., 2022). Cerebras-GPT uses standard trainable positional embeddings.
- Some GPT-J variants disable bias weights for fully-connected layers in transformer attention blocks. Other studies explain that disabling biases can increase accelerator utilization without loss degradation (Chowdhery et al., 2022; Dehghani et al., 2023). We believe this approach might also improve training stability issues caused by key projection bias weights growth as we describe in Appendix A.2. We have not tested the effects on loss/accuracy from disabling these bias weights.
- GPT-J, GPT-NeoX, and Pythia use a parallel structure for attention and feed forward layers (Black et al., 2022). This residual architecture has been reported to cause degradation in model performance at similar scales of models, so it is typically only adopted to increase accelerator utilization (Chowdhery et al., 2022). OPT and Cerebras-GPT models use the standard GPT-2 transformer block, which orders attention sequentially before the feed forward layers.
- GPT-NeoX and Pythia models use vocabulary and tokenization designed specifically for the Pile dataset (Black et al., 2022). The resulting vocabulary is different in a few ways from the GPT-2/3 vocabulary. GPT-J, OPT, and Cerebras-GPT models use the GPT-2/3 vocabulary and tokenizer.
- Pythia models also include those that were trained on a deduplicated version of Pile (Biderman et al., 2023). These models show an average of 1.2% advantage on downstream tasks. This indicates further opportunity to improve models with further dataset curation.
- OPT are trained on a dataset combining the datasets used for RoBERTa, the PushShift.io Reddit dataset, and Pile, along with their own dataset pre-processing (Zhang et al., 2022).

Table 8: Pile pre-training test loss and zero-shot downstream task results for publicly available models.

Model	Pre-training (↓)		Downstream task accuracy (↑)							Down-stream Average	
	Training FLOPs	Pile test xent	Hella-Swag	PIQA	Wino-Grande	Lambada	ARC-e	ARC-c	Open-BookQA		
GPT-J	6.1B	1.7e22	1.613	0.518	0.752	0.640	0.683	0.670	0.340	0.288	0.556
GPT-NeoX	20B	6.4e22	1.519	0.535	0.779	0.661	0.720	0.723	0.380	0.290	0.584
OPT	125M	4.1e20	-	0.292	0.630	0.503	0.379	0.435	0.189	0.166	0.371
	350M	1.1e21	-	0.320	0.644	0.523	0.452	0.440	0.207	0.176	0.395
	1.3B	3.2e21	-	0.415	0.717	0.595	0.579	0.570	0.234	0.234	0.478
	2.7B	6.1e21	-	0.458	0.738	0.610	0.637	0.609	0.268	0.250	0.510
	6.7B	1.4e22	-	0.505	0.763	0.654	0.677	0.656	0.307	0.276	0.548
	13B	2.7e22	-	0.524	0.759	0.651	0.687	0.671	0.329	0.270	0.556
Pythia	70M	1.6e20	2.504	0.270	0.590	0.491	0.259	0.413	0.185	0.132	0.334
	160M	4.1e20	2.186	0.293	0.627	0.519	0.389	0.452	0.181	0.160	0.375
	410M	1.1e21	1.971	0.333	0.668	0.530	0.505	0.504	0.213	0.178	0.419
	1B	2.2e21	1.845	0.376	0.705	0.545	0.566	0.559	0.243	0.196	0.456
	1.4B	3.2e21	1.793	0.398	0.711	0.565	0.604	0.576	0.256	0.204	0.474
	2.8B	6.1e21	1.720	0.451	0.737	0.612	0.654	0.629	0.288	0.220	0.513
	6.9B	1.4e22	1.626	0.482	0.746	0.611	0.679	0.669	0.323	0.270	0.540
	12B	2.4e22	1.582	0.505	0.761	0.645	0.705	0.700	0.336	0.284	0.562
	70M	1.6e20	2.549	0.273	0.607	0.526	0.257	0.404	0.175	0.136	0.340
	160M	4.1e20	2.204	0.294	0.632	0.509	0.370	0.451	0.204	0.172	0.376
	410M	1.1e21	1.989	0.341	0.668	0.534	0.514	0.519	0.206	0.180	0.423
	1B	2.2e21	1.858	0.387	0.712	0.546	0.585	0.568	0.241	0.212	0.464
Pile-dedup	1.4B	3.2e21	1.889	0.403	0.729	0.561	0.610	0.582	0.265	0.198	0.478
	2.8B	6.1e21	1.724	0.466	0.743	0.612	0.672	0.662	0.299	0.232	0.526
	6.9B	1.4e22	1.644	0.488	0.756	0.636	0.695	0.667	0.320	0.252	0.545
	12B	2.4e22	1.601	0.516	0.761	0.639	0.712	0.697	0.341	0.280	0.564
	111M	2.6e18	2.608	0.268	0.594	0.488	0.194	0.380	0.166	0.118	0.315
	256M	1.3e19	2.349	0.274	0.613	0.511	0.293	0.410	0.170	0.158	0.347
Cerebras-GPT	590M	6.1e19	2.181	0.291	0.627	0.498	0.366	0.464	0.190	0.158	0.370
	1.3B	2.8e20	1.997	0.325	0.664	0.521	0.462	0.508	0.224	0.166	0.410
	2.7B	1.1e21	1.834	0.386	0.701	0.559	0.567	0.571	0.246	0.206	0.462
	6.7B	6.3e21	1.704	0.447	0.739	0.602	0.636	0.643	0.282	0.238	0.512
	13B	2.3e22	1.572	0.513	0.766	0.646	0.696	0.714	0.367	0.286	0.570
	111M	2.6e18	2.588	0.268	0.598	0.519	0.204	0.390	0.176	0.124	0.325
Cerebras-GPT + μP	256M	1.3e19	2.359	0.274	0.617	0.505	0.287	0.427	0.194	0.156	0.351
	590M	6.1e19	2.155	0.295	0.644	0.517	0.362	0.470	0.194	0.172	0.379
	1.3B	2.8e20	1.984	0.334	0.682	0.512	0.471	0.515	0.223	0.196	0.419
	2.7B	1.1e21	1.846	0.388	0.697	0.557	0.558	0.569	0.241	0.218	0.461

Table 9: Five-shot downstream task accuracy results. Higher accuracy is better.

Model		Hella-Swag	PIQA	Wino-Grande	Lambada	ARC-e	ARC-c	Open-BookQA	Down-stream Avg.
GPT-J	6.1B	0.494	0.756	0.660	0.662	0.705	0.360	0.310	0.564
GPT-NeoX	20B	0.538	0.774	0.683	0.698	0.746	0.410	0.326	0.596
OPT	125M	0.289	0.628	0.520	0.303	0.426	0.197	0.166	0.361
	350M	0.321	0.647	0.521	0.384	0.464	0.208	0.184	0.390
	1.3B	0.413	0.726	0.597	0.553	0.604	0.273	0.230	0.485
	2.7B	0.458	0.749	0.616	0.603	0.651	0.305	0.276	0.523
	6.7B	0.505	0.773	0.663	0.660	0.692	0.340	0.318	0.565
	13B	0.524	0.763	0.684	0.678	0.714	0.358	0.306	0.575
	70M	0.269	0.589	0.491	0.192	0.399	0.184	0.148	0.325
Pythia	160M	0.292	0.631	0.515	0.329	0.469	0.205	0.164	0.372
	410M	0.333	0.669	0.522	0.448	0.526	0.229	0.188	0.416
	1B	0.374	0.709	0.562	0.514	0.596	0.265	0.206	0.461
	1.4B	0.398	0.712	0.573	0.553	0.622	0.274	0.214	0.478
	2.8B	0.448	0.738	0.621	0.629	0.673	0.328	0.254	0.527
	6.9B	0.478	0.750	0.646	0.641	0.699	0.355	0.296	0.552
	12B	0.506	0.759	0.662	0.673	0.731	0.383	0.322	0.577
	70M	0.272	0.604	0.519	0.192	0.403	0.177	0.152	0.331
	160M	0.294	0.639	0.507	0.309	0.472	0.215	0.178	0.373
	410M	0.339	0.673	0.513	0.456	0.537	0.232	0.190	0.420
Pythia	1B	0.384	0.710	0.552	0.529	0.588	0.259	0.226	0.464
	1.4B	0.400	0.730	0.566	0.565	0.617	0.283	0.232	0.485
	2.8B	0.463	0.758	0.609	0.637	0.681	0.327	0.282	0.537
	6.9B	0.492	0.762	0.637	0.671	0.705	0.344	0.308	0.560
	12B	0.516	0.765	0.678	0.696	0.728	0.386	0.326	0.585
	111M	0.267	0.588	0.475	0.158	0.356	0.166	0.136	0.306
	256M	0.278	0.606	0.522	0.225	0.422	0.183	0.164	0.343
Cerebras-GPT	590M	0.291	0.634	0.479	0.281	0.475	0.206	0.152	0.360
	1.3B	0.326	0.668	0.536	0.395	0.529	0.241	0.174	0.410
	2.7B	0.382	0.697	0.543	0.487	0.590	0.267	0.224	0.456
	6.7B	0.444	0.736	0.590	0.591	0.667	0.314	0.270	0.516
	13B	0.514	0.768	0.674	0.655	0.743	0.398	0.318	0.581
	111M	0.268	0.581	0.520	0.146	0.368	0.175	0.124	0.312
	256M	0.278	0.619	0.534	0.220	0.415	0.193	0.154	0.345
Cerebras-GPT + μP	590M	0.298	0.652	0.515	0.301	0.479	0.206	0.174	0.375
	1.3B	0.329	0.672	0.513	0.396	0.531	0.235	0.212	0.413
	2.7B	0.382	0.704	0.560	0.510	0.595	0.267	0.210	0.461

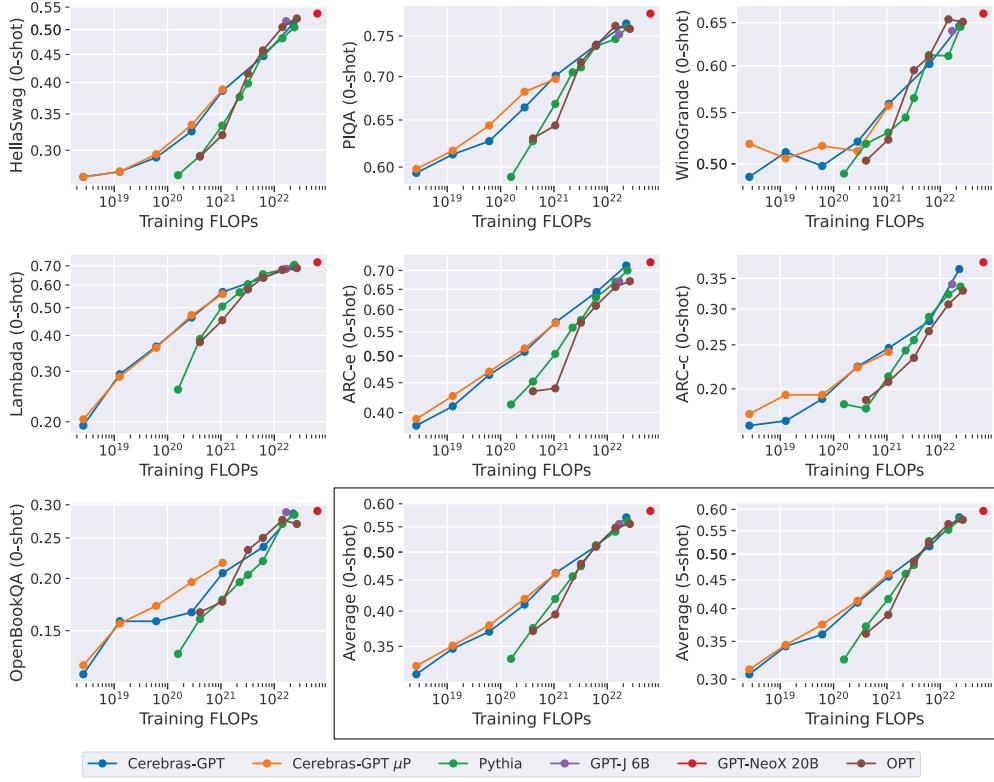


Figure 9: Individual zero-shot downstream task accuracy, zero-shot average, and five-shot average, plotted against training FLOPs.

Figure 10: Zero-shot downstream task accuracy using length-normalized NLL selection criteria.

Model		HellaSwag	PIQA	ARC-e	ARC-c	OpenBookQA
Cerebras-GPT	111M	0.271	0.581	0.348	0.206	0.278
	256M	0.286	0.614	0.376	0.218	0.254
	590M	0.324	0.629	0.412	0.235	0.280
	1.3B	0.384	0.666	0.458	0.250	0.290
	2.7B	0.488	0.707	0.525	0.273	0.318
	6.7B	0.589	0.740	0.579	0.312	0.366
Cerebras-GPT + $\mu$ P	111M	0.276	0.598	0.344	0.223	0.260
	256M	0.287	0.618	0.376	0.225	0.258
	590M	0.333	0.637	0.411	0.237	0.270
	1.3B	0.400	0.670	0.460	0.247	0.298
	2.7B	0.493	0.704	0.495	0.287	0.332

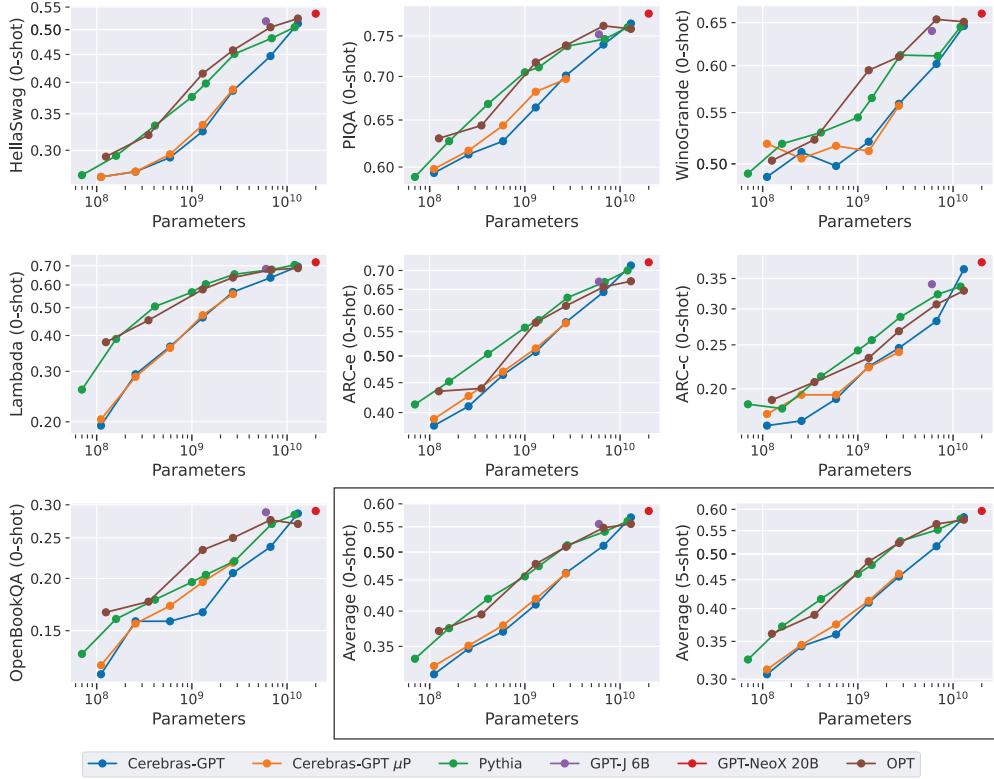


Figure 11: Individual zero-shot downstream task accuracy, zero-shot average, and five-shot average, plotted against model parameters.

Table 10: Five-shot downstream task accuracy using length-normalized NLL selection criteria.

Model		HellaSwag	PIQA	ARC-e	ARC-c	OpenBookQA
Cerebras-GPT	111M	0.270	0.582	0.350	0.208	0.252
	256M	0.291	0.607	0.391	0.211	0.266
	590M	0.324	0.622	0.449	0.229	0.270
	1.3B	0.387	0.665	0.512	0.266	0.286
	2.7B	0.488	0.699	0.576	0.285	0.326
	6.7B	0.589	0.744	0.668	0.343	0.370
	13B	0.694	0.774	0.747	0.433	0.420
Cerebras-GPT + μP	111M	0.275	0.583	0.343	0.202	0.268
	256M	0.288	0.614	0.391	0.225	0.262
	590M	0.337	0.646	0.449	0.227	0.270
	1.3B	0.401	0.670	0.516	0.263	0.294
	2.7B	0.492	0.701	0.591	0.294	0.322

#### C.4 Bias

Language models carry with them the risk of causing harm through the propagation of bias, toxicity, and other negative traits found in their training data. Accordingly, it is important to test models for such biases. We evaluate our models on the CrowS-Pairs dataset (Nangia et al., 2020), which measures bias across nine different categories. In Table 11, we compare bias measurements for our family of models to Pythia 70M–12B, as well as three well regarded baselines: GPT-3 175B (Brown et al., 2020), OPT 175B (Zhang et al., 2022), and LLaMA 65B (Touvron et al., 2023).

Table 11: Analyzing levels of bias using the CrowS-Pairs dataset, comparing GPT-3 175B, OPT 175B, Pythia, and LLaMA 65B models to Cerebras-GPT. Higher values correspond to higher bias.

Model	Race/ Color	Socio- economic status	Gender	Age	Religion	Disabil- ity	Sexual orienta- tion	Nation- ality	Physical appear- ance	Average	
GPT-3	175B	64.7	73.8	62.6	64.4	73.3	76.7	76.2	61.6	74.6	67.2
OPT	175B	68.6	76.2	65.7	67.8	68.6	76.7	78.6	62.9	76.2	69.5
Pythia	70M	51.2	63.7	58.1	54.9	64.0	66.2	79.6	41.2	66.7	56.6
	160M	49.8	56.8	55.9	56.0	72.1	66.2	73.1	45.4	63.9	55.8
	410M	53.7	62.6	61.9	63.7	65.8	72.3	81.7	52.3	62.5	60.1
	1B	54.1	65.8	62.2	63.7	72.1	73.8	78.5	52.3	63.9	61.1
	1.4B	51.8	65.8	63.4	62.6	76.6	72.3	82.8	57.4	68.1	61.8
	2.8B	53.7	66.3	63.1	63.7	78.4	78.5	83.9	55.1	73.6	62.9
	6.9B	55.5	72.6	66.6	72.5	80.2	72.3	84.9	56.5	76.4	65.6
	12B	55.9	68.4	63.4	68.1	75.7	72.3	83.9	57.9	73.6	64.0
	LLaMA	65B	57.0	71.5	70.6	70.1	79.0	66.7	81.0	64.2	77.8
Cerebras-GPT	111M	41.3	69.5	55.6	42.9	64.9	60.0	78.5	43.5	61.1	52.9
	256M	52.8	63.2	57.8	53.8	60.4	67.7	80.6	44.4	61.1	56.9
	590M	46.9	62.6	58.1	59.3	79.3	64.6	79.6	47.7	66.7	57.2
	1.3B	50.6	60.5	58.1	61.5	73.0	69.2	73.1	45.4	72.2	57.6
	2.7B	53.7	65.8	60.3	64.8	76.6	67.7	78.5	52.8	69.4	61.1
	6.7B	54.1	65.3	64.4	65.9	80.2	72.3	86.0	59.7	73.6	63.9
	13B	55.1	72.1	67.5	73.6	81.1	73.8	78.5	59.7	75.0	65.7

The Cerebras-GPT models exhibit less bias on average than any of the larger model baselines. However, Cerebras-GPT 13B does show bias greater than GPT-3, OPT, or LLaMa on six of the nine bias categories, indicating that compute-efficient pre-training is not immune to large bias.

When observing bias levels across Cerebras-GPT or Pythia models, we see that biases tend to grow with model size. The Cerebras-GPT models tend to show a larger range of bias values over the growing model sizes (e.g., gender), while Pythia models sometimes show similar bias across model sizes (e.g., disability). This suggests that models trained on a fixed dataset size may be likely to extract similar levels of bias regardless of model size. On the other hand, more compute-efficient training (smaller datasets for smaller models) might mitigate some bias issues compared to models trained on more data.

Finally, we note that when comparing Cerebras-GPT and Pythia models trained with similar compute budgets, Pythia models tend to have slightly lower bias. In particular, the Cerebras-GPT models 1.3B, 2.7B, 6.7B, and 13B use similar compute to Pythia models 160M, 410M, 2.8B, and 12B, respectively. These Cerebras-GPT models show roughly 1–2% higher bias on average. This suggests that bias is more efficiently extracted from the training data when using a compute-optimal pre-training setup, possibly due to the larger model sizes.

Overall, we recommend further bias evaluation and mitigations for Cerebras-GPT and larger models if deploying them in production settings.

#### D Additional Tokens-per-parameter Experiments

Here, we give more evidence that 20 tokens-per-parameter is nearly compute-optimal when pre-training GPT-like models on the Pile dataset.

### Estimated Chinchilla Losses

First, in Figure 3 in Section 3, we include a curve that estimates the Chinchilla loss degradation when changing the number of tokens per parameter for which a model is pre-trained. To get that estimate, we start by fitting a curve to points we estimate in the Chinchilla paper plot (Figure 3), which shows loss for different model sizes and tokens trained with fixed compute budgets. Given our approach, our estimates are likely to introduce error, and we do not know the true functional form of their fits. However, we pull points from three different FLOP levels, and we validate our curve fit has low error for a fourth held-out FLOP level. This result seems surprising that large changes in model size and training tokens do not appear to have a large effect on the expected degradation from computationally-inefficient training, but it might indicate another invariant to training scale.

We use the Chinchilla curve fit to estimate the proportional loss degradation,  $\Delta\mathcal{L}$ , when changing tokens-per-parameter,  $\tau$  (this is the Chinchilla trend plotted in Figure 3):

$$\Delta\mathcal{L}(\tau) = 0.023 \cdot \ln(\sqrt{20/\tau})^2 \quad (4)$$

This degradation formulation shows good agreement with our tests. Further, Chinchilla models were trained on MassiveText, while our models and Pythia models were trained on the Pile, suggesting the two datasets have significant commonality in their scaling characteristics when training on more than 20 tokens per parameter.

### Our Tokens-per-parameter Experiments

Figure 12 (left) plots the loss degradation (%) from our Cerebras-GPT compute-efficient scaling law for different tokens per parameter (similar to Figure 3), and we add our small scale experiments around 20 tokens per parameter for model sizes 111M, 256M, and 590M parameters. The right plot in Figure 12 zooms in on the region from 15 to 60 tokens per parameter. The right plot shows that losses are quite stable between 20 and 40 tokens per parameter and our empirical best results are between 20 and 30 for each of the three models. We note that the variance in loss for runs at this scale is roughly 0.35%, so most loss values here are also within expected run-to-run variance. Based on these results and the strong agreement between Chinchilla and Pythia results, we were comfortable to conclude that 20 tokens-per-parameter is nearly compute-optimal for these models trained the Pile dataset.

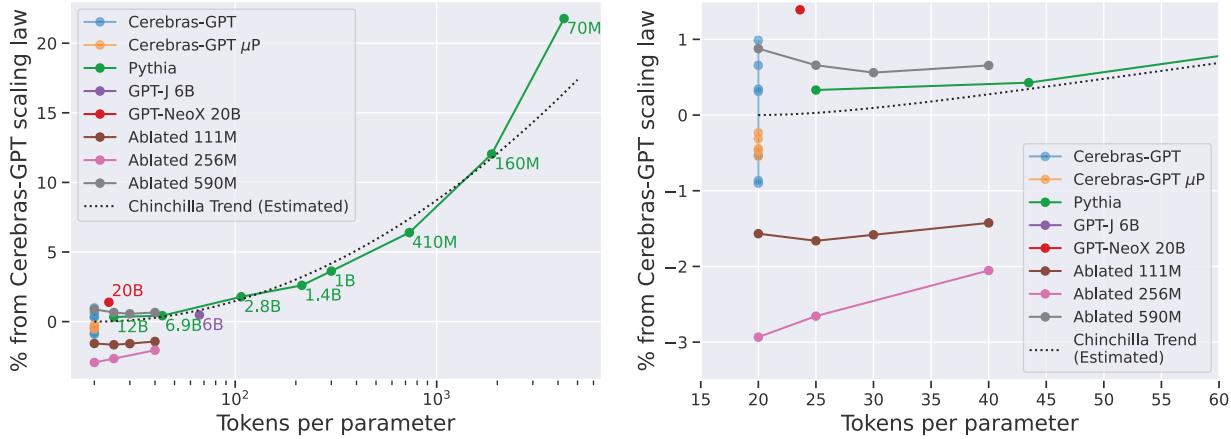


Figure 12: Percent loss degradation compared to the Cerebras-GPT compute-efficient scaling law for varying tokens per parameter.

## E Number of Model Parameters

Table 12 shows the formula we use to calculate parameter counts for Cerebras-GPT models.

Table 12: Python code to calculate parameter count in Cerebras-GPT models

---

```
def get_n_params(vocab_size, d_model, num_layers, seq_length=2048):
    embedding = vocab_size * d_model + d_model * seq_length
    ln1 = 2 * d_model
    attn = 4 * (d_model**2 + d_model)
    ln2 = 2 * d_model
    ffn = 8 * d_model**2 + 5 * d_model
    encoder = num_layers * (ln1 + attn + ln2 + ffn)
    final_ln = 2 * d_model
    n_params = embedding + encoder + final_ln
    return n_params
```

---

## F Number of Training FLOPs

We calculate the number of training FLOPs with a formula similar to Chinchilla, but with two modifications. First, we account for the dot product between  $\text{softmax}(QK^T)$  and  $V$ . Second, we account for the fact that embedding layers do not need to calculate a delta gradient for earlier layers. Code for FLOPs count is in Table 13. We consider this FLOP calculation to be a measure of the algorithmic calculations required for forward and backward gradient steps of training, or “Algorithmic FLOPs”. This formula does not include any additional FLOPs for things like activation checkpointing and recomputation or specifics related to software or hardware implementation.

Table 13: Python code to calculate FLOPs per sequence in Cerebras-GPT models

---

```
def get_flops_per_seq(vocab_size, d_model, num_layers, key_size, seq_len, inference=False):
    num_heads = d_model // key_size
    embeddings = 2 * seq_len * vocab_size * d_model
    position_embeddings = 2 * d_model * seq_len

    kqv_proj = num_layers * 2 * 3 * seq_len * d_model * (key_size * num_heads)
    kq_logits = num_layers * 2 * (seq_len**2) * (key_size * num_heads)
    softmax = num_layers * 3 * (key_size * num_heads) * (seq_len**2)
    softmax_q_red = num_layers * (seq_len**2) * (key_size * num_heads)
    final_linear = num_layers * 2 * seq_len * (key_size * num_heads) * d_model
    sm_v_dot = num_layers * 2 * (seq_len**2) * (key_size * num_heads)
    attention_blocks = kqv_proj + kq_logits + softmax + softmax_q_red + sm_v_dot + final_linear

    dense_blocks = num_layers * 16 * seq_len * (d_model**2)
    final_logits = 2 * seq_len * d_model * vocab_size

    # Layer norms: 7 FLOPs/activation, 2 LNs per decoder block
    layer_norm_flops = num_layers * 2 * 7 * (seq_len * d_model)

    # GeLU: estimate 20 FLOPs/activation, applied in FFN with 4x hidden dim
    gelu_flops = num_layers * 20 * 4 * (seq_len * d_model)

    total_flops_per_step = embeddings + position_embeddings + layer_norm_flops + attention_blocks +
        dense_blocks + final_logits + gelu_flops
    inference_flops_per_step = total_flops_per_step

    # Account for backward pass too
    total_flops_per_step *= 3

    # Embeddings don't need to pass a delta back
    total_flops_per_step -= embeddings
    total_flops_per_step -= position_embeddings

    if inference:
        return inference_flops_per_step
    else:
        return total_flops_per_step
```

---

## G Additional $\mu$ P Details

We spent time orienting ourselves to use  $\mu$ P, so this section describes that process for other practitioners.

### G.1 $\mu$ P Implementation

In Table 14 we detail all the changes required to implement  $\mu$ P in GPT-like models, and we describe these changes in more detail here.

- $\mu$ P adds a tunable embedding output activation multiplier,  $m_{\text{emb}}$ , which is multiplied by the sum of token and position embeddings. This multiplier controls relative activation and gradient magnitudes between the embeddings layers and the transformer backbone.
- Similarly, to control the relative gradient magnitude between embeddings and the transformer backbone, the model’s output logits activation tensor (pre-softmax) is multiplied by  $1/m_{\text{width}}$  (when using shared embedding and output weights).
- To control the expected magnitude of activations in the transformer blocks,  $\mu$ P scales the initialization variance for each fully-connected layer’s weights by  $1/m_{\text{width}}$ .
- To control the relative weight magnitudes throughout training,  $\mu$ P scales the learning rate of each fully-connected layer’s weights by  $1/m_{\text{width}}$ .
- Under the assumption that query and key projections have significant alignment later in training,  $\mu$ P scales the key-query dot-product activations by  $1/d_{\text{head}}$ , rather than using the  $1/\sqrt{d_{\text{head}}}$  scaling originally proposed in Vaswani et al. (2017).
- The next subsection describes how we tune the three transferrable  $\mu$ P hyperparameters: the base learning rate ( $\eta_{\text{base}}$ ), the base initialization standard deviation ( $\sigma_{\text{base}}$ ), and the embedding output multiplier ( $m_{\text{emb}}$ ).

### G.2 $\mu$ P Hyperparameter Search Details

We tune  $\mu$ P hyperparameters on a 40M parameter proxy model, and  $\mu$ Transfer those hyperparameters to our models with 111M–2.7B parameters. Figure 13 shows the results of a 200 sample random hyperparameter search on a 40M parameter proxy model ( $d_{\text{model}} = d_{\text{model,base}} = 256$ ,  $n_{\text{layers}} = 32$ ,  $d_{\text{head}} = 128$ ) trained on 600M tokens with a batch size of 131k tokens. From this sweep we obtained the following tuned hyperparameters for our  $\mu$ P models:  $\eta_{\text{base}} = 6e-3$ ,  $\sigma_{\text{base}} = 0.08$ ,  $m_{\text{emb}} = 10$ . These hyperparameters also closely match those used by Yang et al. (2021).

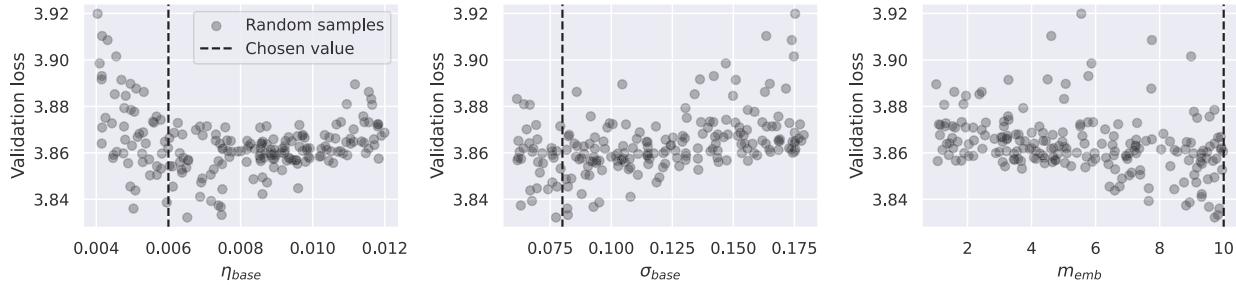


Figure 13: Results from random hyperparameter search on 40M parameter  $\mu$ P proxy model.

### G.3 Advice for Practitioners

#### Use Large Enough Batch Sizes For Each Model Size

The  $\mu$ P paper suggests small proxy models should be trained with batch sizes similar to the larger target model to which you would like to  $\mu$ Transfer the hyperparameters. We find, however, that choice of proxy

Table 14: Cheat Sheet: All implementation details required to compare SP and  $\mu$ P.

	Standard Parameterization (SP)	Maximal Update ( $\mu$ P)
<b>Variables</b>		
$W$	A multiplicative or fully-connected weights tensor	
$b$	A bias weights tensor	
$X, Y$	Activation tensors: layer input, output, respectively	
$d_{\text{model},\text{base}}$	Proxy (base) model's layer width	
$d_{\text{model}}$	Width of each layer	
$d_{\text{head}}$	Size of each attention head	
embed	Combined token + position embedding function	
$\eta_{\text{base}}$	The base learning rate (LR): Maximum in training schedule	
$\sigma_{\text{base}}$	The base initialization standard deviation	
$m_{\text{width}}$	—	Layer width multiplier: $d_{\text{model}}/d_{\text{model},\text{base}}$
$m_{\text{emb}}$	—	Embedding output multiplier
<b>Empirically Tuned Values</b>		
$d_{\text{model},\text{base}}$	—	256
$\eta_{\text{base}}$	Must tune for each model size	6e-3
$\sigma_{\text{base}}$	0.02	0.08
$m_{\text{emb}}$	—	10.0
<b>Formulas</b>		
Embedding initializer	$W_{\text{emb}} \sim N_{\text{trunc}}(0, \sigma_{\text{base}}^2)$	
Embedding LR	$\eta_{\text{emb}} = \eta_{\text{base}}$	
Embedding output	$Y_{\text{emb}} = \text{embed}(X)$	$Y_{\text{emb}} = m_{\text{emb}} \cdot \text{embed}(X)$
LN initializer	$W_{\gamma} \sim 1, b_{\beta} \sim 0$	
LN LR	$\eta_{\text{LN}} = \eta_{\text{base}}$	
Bias initializer	$b \sim 0$	
Bias LR	$\eta_b = \eta_{\text{base}}$	
MHA equation	$\text{softmax} \left( \frac{Q^T K}{\sqrt{d_{\text{head}}}} \right) V$	$\text{softmax} \left( \frac{Q^T K}{d_{\text{head}}} \right) V$
QKV weights initializer	$W_{\text{qkv}} \sim N_{\text{trunc}}(0, \sigma_{\text{base}}^2)$	$W_{\text{qkv}} \sim N_{\text{trunc}}(0, \sigma_{\text{base}}^2/m_{\text{width}})$
QKV weights LR	$\eta_{\text{qkv}} = \eta_{\text{base}}$	$\eta_{\text{qkv}} = \eta_{\text{base}}/m_{\text{width}}$
O weights initializer	$W_o \sim N_{\text{trunc}}(0, \frac{\sigma_{\text{base}}^2}{2 \cdot n_{\text{layers}}})$	$W_o \sim N_{\text{trunc}}(0, \frac{\sigma_{\text{base}}^2}{2m_{\text{width}} \cdot n_{\text{layers}}})$
O weights LR	$\eta_o = \eta_{\text{base}}$	$\eta_o = \eta_{\text{base}}/m_{\text{width}}$
ffn1 weights initializer	$W_{\text{ffn1}} \sim N_{\text{trunc}}(0, \sigma_{\text{base}}^2)$	$W_{\text{ffn1}} \sim N_{\text{trunc}}(0, \sigma_{\text{base}}^2/m_{\text{width}})$
ffn1 weights LR	$\eta_{\text{ffn1}} = \eta_{\text{base}}$	$\eta_{\text{ffn1}} = \eta_{\text{base}}/m_{\text{width}}$
ffn2 weights initializer	$W_{\text{ffn2}} \sim N_{\text{trunc}}(0, \frac{\sigma_{\text{base}}^2}{2 \cdot n_{\text{layers}}})$	$W_{\text{ffn2}} \sim N_{\text{trunc}}(0, \frac{\sigma_{\text{base}}^2}{2m_{\text{width}} \cdot n_{\text{layers}}})$
ffn2 weights LR	$\eta_{\text{ffn2}} = \eta_{\text{base}}$	$\eta_{\text{ffn2}} = \eta_{\text{base}}/m_{\text{width}}$
Output logits multiplier	$Y_{\text{logits}} = W_{\text{unemb}} X$	$Y_{\text{logits}} = W_{\text{unemb}} X / m_{\text{width}}$

model batch size can be more flexible as long as it is large enough. When tuning on a small model, Yang and Hu choose batch size to be quite large such that it is an appropriate batch size for the largest models to which they  $\mu$ Transfer those hyperparameters. We find that proxy model tuning does *not* need to be performed on a batch size appropriate for the largest model. Rather, batch size must be sufficiently large such that the proxy model’s gradient dynamics are likely to be consistent with the larger models. More specifically, setting the proxy model’s batch size larger than its critical batch size (dictated by gradient noise (McCandlish et al., 2018)) is sufficient to get good hyperparameter transferrability. Then, batch size should be scaled appropriately as model size scales.

### Simple Batch Size + Learning Rate Scaling with $\mu$ P

More precisely, we find that  $\mu$ P learning rate transfers as long as each model size is trained with a batch size roughly consistent with or larger than the critical batch size. The closer the batch size is to the critical batch size for a given model, the better the loss will be when using the  $\mu$ Transferred learning rate. Further, when training models with a batch size smaller than the critical batch size, learning rate should be reduced linearly proportional to the reduction in batch size—consistent with the findings in (Shallue et al., 2018; Yang et al., 2021).