# SEQ VS SEQ: AN OPEN SUITE OF PAIRED ENCODERS AND DECODERS

**Anonymous authors**Paper under double-blind review

000

001

002 003 004

010 011

012

013

014

016

017

018

019

021

023

025

026

027 028 029

031

033

034

037

040

041

042

043

044

046

047

048

051

052

#### **ABSTRACT**

The large language model (LLM) community focuses almost exclusively on decoder-only language models, since they are easier to use for text generation. However, a large subset of the community still uses encoder-only models for tasks such as classification or retrieval. Previous work has attempted to compare these architectures, but is forced to make comparisons with models that have different numbers of parameters, training techniques, and datasets. We introduce the SOTA open-data ETTIN<sup>1</sup> suite of models: paired encoder-only and decoder-only models ranging from 17 million parameters to 1 billion, trained on up to 2 trillion tokens. Using the same recipe for both encoder-only and decoder-only models produces SOTA recipes in both categories for their respective sizes, beating ModernBERT as an encoder and Llama 3.2 and SmolLM2 as decoders. Like previous work, we find that encoder-only models excel at classification and retrieval tasks while decoders excel at generative tasks. However, we show that adapting a decoder model to encoder tasks (and vice versa) through continued training is subpar compared to using only the reverse objective (i.e. a 400M encoder outperforms a 1B decoder on MNLI, and vice versa for generative tasks). We open-source all artifacts of this study including training data, training order segmented by checkpoint, and 200+ checkpoints to allow future work to analyze or extend all aspects of training.<sup>2</sup>

### 1 Introduction

The rise of neural language models (LMs) was spurred by encoder-only models such as ELMo (Peters et al., 2018) and BERT (Devlin et al., 2019). However, the community generally shifted to decoder-only (i.e. GPT-style, ala Brown et al. (2020)) models due to their exceptional performance at sequence generation. Due to this lack of popularity for encoder-only models there was limited new model development, thus, we still frequently see usage of older models (i.e. from 2019) by the subset of the community focused on retrieval/classification or fast on-device inference. Although nascent work is attempting to revive encoder-only development (Samuel, 2024; Warner et al., 2024; Lee et al., 2025), there still exists a wide gap between the development of encoder-only and decoder-only models (synonymously referred to in this work as *encoders* or *decoders*).

Part of this gap is due to the sentiment within the community that decoders can be adapted for use in tasks that were once predominantly encoder-focused (e.g. classification, embeddings), especially as they can often be used in a zero-shot fashion (i.e. without fine-tuning). As decoder models are more studied, more over-trained (Hoffmann et al., 2022), and are generally larger, they are now claiming the top spots of leaderboards for previously encoder-centric tasks (Enevoldsen et al., 2025).

Many works have challenged this assumption by comparing encoder-only and decoder-only models of roughly the same sizes (Ethayarajh, 2019; Charpentier & Samuel, 2024; Harrag et al., 2021). However, these analyses have to be done with incomparable models: using different architectures, different pre-training data, different learning schedules, etc.

Our work aims to provide the foundation to compare encoder-only and decoder-only models by open-sourcing a suite of models trained with the same data, the same architecture, and the same training recipe. Our ETTIN suite contains 10 models (5 pairs) ranging from 17 million to 1 billion

<sup>&</sup>lt;sup>1</sup>Named for the two-headed mythological Norse giant, symbolizing the two language models heads.

<sup>&</sup>lt;sup>2</sup>Models, code, and data are available at <removed for anonymity>

parameters, and trained for up to 2 trillion tokens. This allows us to quantify the differences between these models (including the effects of scaling parameter size) in an apples-to-apples comparison.

Our models provide state-of-the-art performance for their size among open-data models. Surprisingly, they do so in both encoder settings (w.r.t. ModernBERT) and decoder settings (w.r.t. LLaMA 3.2 and SmolLM2) despite using the same recipe. Notably, our work also provides the first open-data replication of ModernBERT, allowing the community to further build upon our recipe.

We find that, like previous work, encoders excel at classification and retrieval while decoders excel at generative tasks. However, we go beyond previous work to examine the increasingly common setting (BehnamGhader et al., 2024) where decoder-models are continued trained as encoders (i.e. cross-objective training). We show results for this cross-objective training in both directions: training encoders for causal language modeling (CLM) and decoders with masked language modeling (MLM). We find that despite continued training for much longer than previous work (50B tokens) these models do not surpass those that started with this objective, i.e. a 400M encoder outperforms a 1B decoder continue-trained with MLM on MNLI, and vice versa for generative tasks.

Our work also provides the ability to compare these training objectives on other aspects, comparing how they learn. We provide a case study showing the effects of these objectives on gender bias.

Overall, our work provides the first suite of models enabling a fair comparison between encoder-only and decoder-only architectures (while also showing SOTA performance), enabling future work to analyze the effects of these training objectives on downstream tasks.

## 2 RELATED WORK

We describe encoder models as the community is generally more familiar with decoder LMs.<sup>3</sup> It is worth noting that our approach was inspired by Pythia (Biderman et al., 2023b) which was the first to explore open-data decoder-only models at multiple sizes.

**Encoder-only Models** Encoder-only architectures were the predominant architecture for early transformer models, popularized by models such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and DeBERTa (He et al., 2023). These models showed significantly improved performance over the previous SOTA LSTM models on classification and retrieval tasks. This created a flurry of activity in the encoder space, with models improving on the BERT recipe: the RTD objective from DeBERTa, better data and objectives from RoBERTa, and many smaller variants such as TinyBERT (Jiao et al., 2019), DistilBERT (Sanh et al., 2019), BERT-small (Turc et al., 2019), and MiniLM L12 (Wang et al., 2020). However, these encoders lacked the easy ability to generate text and have generally fallen out of popularity in favor of decoder-only GPT-2 style models.

Despite this shift, encoders still maintain frequent usage for many tasks that don't require generative output. For example, in March 2025 alone, BERT-base had 90 million downloads on HuggingFace.

Recently, there has been renewed interest in encoders, as demonstrated by NomicBERT (Nussbaum et al., 2024), mosiacBERT (Portes et al., 2023), and ModernBERT (Warner et al., 2024). Unfortunately, ModernBERT (the most performant) does not provide access to their training data. Hence, we use publicly available data sources in order to replicate the training process.

**Comparisons between Encoders and Decoders** Previous work has compared encoder-only and decoder-only models on a wide assortment of tasks. For example Charpentier & Samuel (2024) compare DeBERTa and GPT-2 in similar sizes. Other work (Yang et al., 2023; Qu et al., 2020; Zheng et al., 2021; Rehana et al., 2023; Nielsen et al., 2024) compares them on downstream tasks.

However, all of these comparisons have the same underlying limitation: the models they are comparing have different numbers of parameters, different architectures, different training recipes, and different pre-training data. Although some work has attempted to address this (Charpentier & Samuel, 2024; Gisserot-Boukhlef et al., 2025), they have only done so in limited settings with very small amounts of pre-training data. In contrast, we train SOTA models, allowing for an exact comparison.

<sup>&</sup>lt;sup>3</sup>For those interested in decoders, please see early works such as GPT-2 (Radford et al., 2019) and modern models such as OpenAI's GPT-4 (Achiam et al., 2023), Google's Gemini (Team et al., 2023), Alibaba's Qwen series Yang et al. (2025), and Meta's LLaMA models (Grattafiori et al., 2024)

Parameter	17M (XXS)	32M (XS)	68M (Small)	150M (Base)	400M (Large)	1B (XL)
Layers	7	10	19	22	28	28
Hidden Size	256	384	512	768	1024	1792
Intermediate Size	384	576	768	1152	2624	3840
Attention Heads	4	6	8	12	16	28
Learning Rate	3e-3	3e-3	3e-3	8e-4	5e-4	5e-4
Weight Decay	3e-4	3e-4	3e-4	1e-5	1e-5	5e-5
Warmup Tokens (B)	4	4	3	3	2	2
BS Warmup (B)	125	100	75	50	10	3

Table 1: Configuration for each Ettin model size. Both encoders and decoders use the exact same configuration, differing only in attention (bidirectional vs causal) and objective (MLM vs CLM).

**Bidirectional Attention for Decoders** Although we cannot cover it all here, there have been attempts to use bidirectional attention for standard decoder usage. This includes prefix LM attention (Artetxe et al., 2022; Chowdhery et al., 2023; Du et al., 2021) and other mixed training such as BiTune (Kopiczko et al., 2024). However, most modern LM training still favors pure CLM.<sup>4</sup>

Checkpoint-Level Model Analyses There has also been much work exploring how models learn via their training data. This was popularized by the Pythia (Biderman et al., 2023a) paper and includes many aspects of learning such as data quality and selection (Longpre et al., 2024), how frequency of entities impacts model learning (Oh et al., 2024), effects of the recency of data (Cheng et al., 2024), and whether you can recognize and extract training data from models (Zhang et al., 2024). Our work allows these experiments to be done on more recent SOTA models and provides a way to compare encoders and decoders on various facets of learning.

## 3 EXPERIMENTAL SETTINGS

#### 3.1 Training Data

We create an open-source replication of ModernBERT (Warner et al., 2024) due to it being the strongest publicly available encoder-only model and using comparable techniques to decoder-only models. This provides the best starting place for a recipe that spans both training objectives. However, ModernBERT's data is not publicly available – thus, we aim to replicate the recipe using open-data.

We do so by pulling from the best publicly available datasets used for training decoder-only models, such as Olmo (Groeneveld et al., 2024; OLMo et al., 2025). Thus, we use a mix of DCLM (Li et al., 2024) combined with various curated sources from Dolma v1.7 (Soldaini et al., 2024). In the process of training, the Olmo 2 paper (OLMo et al., 2025) described their approach of using filtered DCLM and other higher-quality sources for the decay phase (similar to FineWeb-Edu filtering Penedo et al. (2024); Lozhkov et al. (2024)). We decided to use these newer sources for our later phases.<sup>5</sup>

To allow others to easily extend our work, we provide both formats: the data which can be used for training as well as the data seen by the models in batch order for future analyses.

#### 3.2 Architecture

As ModernBERT has only two sizes, we develop new shapes for our smaller and larger models (Table 1). We aim to follow the design espoused by MobileLLM (Liu et al., 2024) with deep but thin models. However, for the 1B model, we keep the same number of layers but make the model wider. We choose models parameter sizes at roughly 2x increments while matching common encoder sizes, e.g. 17M, 32M, 68M, 150M, 400M, and 1B. For a detailed list of the differences, see Table 1.

<sup>&</sup>lt;sup>4</sup>The the best of our knowledge, as much of the details of the best LMs now goes unpublished.

<sup>&</sup>lt;sup>5</sup>We ablated with the non-filtered data and found worse results.

		Pre-trair	ning	Mid-traiı	ning	Decay Ph	ase
Category	Dataset	Tokens (B)	%	Tokens (B)	%	Tokens (B)	<b>%</b>
News	CC News	7.3	0.4	_	_	_	_
Code	Starcoder	263.9	15.5	38.4	15.4	_	_
Code	Code_Repos	_	_	_	_	20.2	26.5
Crawl	CC Head	356.6	20.9	_	_	_	_
Crawl	DCLM	837.2	49.1	_	_	_	_
Crawl	DCLM (Dolmino)	_	_	175.5	70.4	26.0	34.1
Math	Open-Web-Math	12.7	0.7	_	_	_	_
Math	Algebraic StackExchange	12.6	0.7	_	_	_	_
Math	Math (Dolmino)	_	_	10.4	4.2	5.0	6.6
Scientific	PeS2o	57.3	3.4	8.3	3.3	_	_
Scientific	Arxiv	28.0	1.6	4.1	1.6	3.0	3.9
Social	Reddit	80.3	4.7	6.2	2.5	_	_
Social	StackExchange	19.6	1.1	_	_	_	_
Social	StackExchange (Dolmino)	_	_	2.7	1.1	4.0	5.2
Reference	Textbooks	_	_	_	_	0.5	0.7
Reference	Dolma Books	5.3	0.3	0.8	0.3	10.5	13.8
Reference	Wikipedia	7.3	0.4	0.5	0.2	3.0	3.9
Instruction	Tulu Flan	16.6	1.0	2.4	1.0	4.1	5.4
Total		1,704.7	100.0	249.3	100.0	76.3	100.0

Table 2: Training data mixture across the various training stages (pre-training, mid-training, decay). Later stages use higher quality data, from the recently released Dolmino dataset (OLMo et al., 2025). Dashes indicate that no data from that source was used. We trained for 1.7T tokens for pre-training, 250B for mid-training, and 50B for the decay phase. We sample from the dataset and repeat (or under-sample) as needed to hit the token counts used for training.

## 3.3 TRAINING RECIPE

We use the same general process described by open-data models (which was followed by Modern-BERT) for training both encoder-only and decoder-only models — with a few specific changes for the encoder architecture (i.e. masking ratio). In summary, we include three general phases: base pre-training, mid-training/context extension, and decay. See Table 2 for the precise sources of training data in each phase. We use a trapezoidal learning rate scheduler, with general hyperparameters shown in Appendix C and size-dependent hyperparameters in Table 1. For compute details see Appendix E.

The only differences between the encoder and decoder models are: (1) the objective function, i.e. masked language modeling (MLM) for the encoder<sup>6</sup> vs causal language modeling (CLM) for the decoder and (2) the attention pattern, i.e. causal for the decoder and bidirectional for the encoder.

We checkpoint every 8.5B tokens, with 236 checkpoints per model. Combined with the batch ordering of the data, this enables precise pinpointing of what the model learned between each checkpoint.

**Base Pre-training** This stage encompasses the warmup and stable phase of the trapezoidal learning rate, training for 1.7T tokens. We use both learning rate and batch size warmup. The data in this stage comprises a wide mix of sources to allow for general learning.

Context Extension / Mid-Training In this phase we increase the quality of the data and change both the data and base RoPE (Su et al., 2024) to handle longer context. We update the data length to be up to 8000 tokens and RoPE parameters to 160k (for global and local layers). For the data, we drop the noisiest sections (older Dolma common crawl, CC News, general StackExchange) and include filtered DCLM, math, and StackExchange. We then train for 250B tokens and use an inverse square root learning rate schedule from the peak learning rate to 1/2 of the peak.

**Decay Phase** Finally, we use one more inverse square root learning rate schedule to decay for 50B tokens. We follow the general ProLong recipe (Gao et al., 2024b), increasing long context data such as Dolma books, Wikipedia, and open-access textbooks. We decay to 0.02 of the peak LR.

<sup>&</sup>lt;sup>6</sup>For the encoder we use a 30% masking ratio except for the decay phase, which is 15%.

## 3.4 Major training differences from ModernBERT

A concise summary of the largest differences from the ModernBERT recipe are (1) the use of opendata, (2) decay in the context extension phase, (3) no model merging, (4) a lower masking ratio for the decay phase (15% instead of 30%) and (5) local and global RoPE to be the same value.

#### 3.5 Cross-Objective Training

As encoder models have gone out of popularity, decoder models have increased in size (both parameters and pre-training data). Thus, these newer decoder models are typically trained for much longer than previous encoder models (i.e. BERT). Due to this it has become common to adapt these larger decoder models to what were previously encoder-centric tasks (Zhang et al., 2025). With paired encoder and decoder models, we can now answer the question of **how effective this continued pre-training approach is** and **whether it is still worth training both types of models**. We call this *cross-objective training*: taking the final model and continue pre-training it on the reverse objective.<sup>8</sup>

We train for 50B tokens with the reverse objective, which is far more than previous work has attempted (BehnamGhader et al., 2024), which is usually around 10B tokens. Although the ratio of pre-training and cross-objective training is unbalanced, this mimics the realistic setting where the adaptation is done with very small amounts of data comparatively. We do this cross-objective training on the highest quality data we have available, which was used in the last decay phase. We use a new trapezoidal learning rate schedule with 3B tokens of warmup and 10B tokens of decay. Thus, by the end we have an encoder-from-decoder (i.e. a decoder further pretrained with MNTP<sup>11</sup>) and a decoder-from-encoder (i.e. an encoder further pre-trained with CLM).

## 4 EXPERIMENTS

We aim to compare encoder and decoder models. However, first, to give those experiments credence, we show that our models are SOTA. This strengthens our claim and helps alleviate concerns that we made training choices that favored one architecture over the other – instead we have SOTA models in both architectures for their sizes, showing our method's effectiveness. Note though, that the purpose of our paper was not to be SOTA overall (i.e. compared to OpenAI, etc.), but to provide a comparison for encoders and decoders. For space and to avoid repetition, specific model size details are in Table 1.

#### 4.1 Individual Evaluations

**Encoder-Only Results** We use two baselines for each size type: extra extra small (XXS) BERT-mini and TinyBERT, extra small (XS) models MiniLM L12 and BERT-small, small (S) models DistilBERT and DistilRoBERTa, base (B) models BERT and ModernBERT, large (L) models BERT-large and ModernBERT-large, and an extra large (XL) model DeBERTA v2 XL. 12

We evaluate on various encoder tasks, including GLUE (Wang et al., 2018), MTEB v2 English (Enevoldsen et al., 2025), MDLR for long context (Chen et al., 2024), and CodeSearchNet for code evaluation (Husain et al., 2019). We use the same evaluation setup as ModernBERT for the evaluation for an equal comparison (see Appendix F for hyperparameter details).

<sup>&</sup>lt;sup>7</sup>We do this for ease of scientific comparison, however, if one was to use this for downstream applications a simple merge would likely boost performance another point or two.

<sup>&</sup>lt;sup>8</sup>Following BehnamGhader et al. (2024), we do not use MLM but rather use MNTP, that is, the masked token is predicted using the hidden state of the previous token to better align with CLM.

<sup>&</sup>lt;sup>9</sup>We note that this means it repeating this data twice, however, as shown by previous work (Muennighoff et al., 2023) two repetitions on high quality data has no adverse effects.

<sup>&</sup>lt;sup>10</sup>For the 1B model this is scaled by 1/3 again due to compute availability.

<sup>&</sup>lt;sup>11</sup>We use a 15% masking rate for the encoder-from-decoder as to maintain a middle ground masking ratio.

<sup>&</sup>lt;sup>12</sup>We also ran experiments with DeBERTa XXL as shown in the appendix. However, due to the size and slowness of the architecture we could not do an comparable grid search. Our results in the appendix are after 300 days of H100s hours, but still did not complete the full sweep. Furthermore, as DeBERTa XXL is > 1.5B we exclude it as it is significantly larger than 1B (i.e. > 50% larger).

			Embedding	Tasks			GLUE T	Tasks	
Model Name	CSN	MLDR	Clustering	Retrieval	MTEB v2	SST-2	MNLI	GLUE Avg	
		X	XXS Models (	7-17M paraı	meters)				
BERT-mini	41.3	16.8	39.0	34.7	49.2	88.3	77.2	76.4	
TinyBERT	39.8	14.2	37.4	33.3	49.7	91.2	80.9	77.0	
Ettin-Enc-17m	59.1	24.4	39.1	35.6	48.9	91.2	79.5	79.2	
		2	XS Models (28	8-33M parar	neters)				
BERT-small	46.0	19.9	39.6	38.1	51.1	90.1	79.2	79.0	
MiniLM L12	48.3	19.6	37.8	38.4	51.3	93.3	85.6	84.6	
Ettin-Enc-32m	69.2	28.4	39.6	39.7	50.9	92.0	83.4	83.5	
Small Models (68-82M parameters)									
DistilBERT	47.9	23.7	39.8	40.8	52.7	92.2	82.7	81.5	
DistilRoBERTa	60.3	19.7	39.3	40.0	51.8	93.1	84.7	83.8	
Ettin-Enc-68m	75.1	30.1	40.1	43.1	52.6	94.4	87.0	87.2	
		Ba	se Models (12	23-150M par	ameters)				
BERT base	51.0	24.8	40.4	41.2	52.9	93.1	85.4	84.7	
ModernBERT base	75.9	30.4	41.3	43.9	54.0	96.0	89.1	88.4	
Ettin-Enc-150m	76.3	31.8	41.5	45.7	54.0	95.8	89.2	88.9	
		Lar	ge Models (3	53-395M pai	rameters)				
BERT large	54.4	25.3	41.5	42.9	53.8	93.3	86.3	85.2	
ModernBERT large	78.3	34.9	41.5	47.0	55.0	97.1	90.8	90.4	
Ettin-Enc-400m	80.7	36.2	41.8	48.4	55.5	96.7	91.3	90.8	
		XI	Models (884	M-1.2B par	ameters)				
DeBERTa-v1-xl	75.6	28.1	42.5	47.2	56.4	97.1	91.7	90.7	
Ettin-Enc-1B	82.3	40.2	41.9	50.1	56.0	97.1	91.8	91.6	

Table 3: ETTIN encoders compared to other encoder-only models across various sizes on retrieval and GLUE tasks. Due to space, we show two representative tasks from MTEB v2 and two from GLUE, as well as a code-based retrieval evaluation (CodeSearchNet) and a long-context evaluation (MLDR). See Appendix A for the full tables of GLUE and MTEB v2. ETTIN shows significant gains over baseline encoders, including ModernBERT, while also having both larger and smaller sizes.

We find in Table 3 that ETTIN compares favorably overall. The relatively larger gains in the bigger sizes is likely because the smaller model baselines are heavily optimized with distillation. <sup>13</sup> Even so, we see that they generally outperform the baselines without doing any distillation: e.g. Ettin-68m with a GLUE average of 87.2 compared to the next best DistilRoBERTa at 83.8. Even for the more recent ModernBERT baselines we see improved performance (88.9 GLUE average vs 88.4 for the base size). Thus we can see that ETTIN matches or improves the SOTA for encoder-only models.

**Decoder-Only Results** We use two baselines for each size type when available, but few very small decoder-only LMs exist: One extra extra small (XXS) model Pythia-14M, <sup>14</sup> no models in the extra small (XS) category that we could find, one small (S) model DistilGPT (Sanh et al., 2019), two base (B) models Pythia 160M (Biderman et al., 2023a) and SmolLM2 135M, two large (L) models Pythia 410M and SmoLM2 360M (Allal et al., 2025), and extra large (XL) models Olmo 1B 0724 Groeneveld et al. (2024) and Llama 3.2 1B (Dubey et al., 2024). <sup>15</sup>

We evaluate on a wide range of tasks using the Eleuther AI harness (Gao et al., 2024a) (see Appendix B for details), consolidating tasks used in the Pythia and SmolLM papers including: the ARC Challenge (ARC) Clark et al. (2018), HellaSwag (HS) (Zellers et al., 2019), LAMBADA (LMB) (Paperno et al., 2016), OpenBookQA (OBQA) (Mihaylov et al., 2018), Social IQA (SIQA) (Sap et al., 2019), TriviaQA (TQA) (Joshi et al., 2017), Winogrande (WG) (Sakaguchi et al., 2021), and the Winograd Schema Challenge (WSC) (Levesque et al., 2012).

<sup>&</sup>lt;sup>13</sup>We also note that MiniLM L12 has twice the amount of non-embedding parameters 21M vs 12M.

<sup>&</sup>lt;sup>14</sup>This is likely just a debug-sized run and not an official size, as they do no include it in their paper. However, as not other decoder models in this size could be found, we use it as a reference.

<sup>&</sup>lt;sup>15</sup>What is considered "1B" has a large range, up to nearly 2B parameters. We thus restrict our range to < 1.2B parameters to be a "1B" model, which excludes models like SmolLM2 1.7B and Olmo 2 1B (actually 1.5B).

Model Name	ARC	HS	LMB	OBQA	PIQA	SciQ	SIQA	TQA	WG	WSC	Avg
			XXS	Models (1	4-17M pa	aramete	ers)				
Pythia-14m	21.2	26.0	7.1	26.2	55.2	43.8	33.4	0.0	50.3	51.6	31.5
Ettin-Dec-17m	21.3	27.1	23.0	27.2	57.7	71.1	35.4	2.6	50.9	48.0	36.4
			XS	Models (	32M para	ameters	)				
Ettin-Dec-32m	23.5	28.5	28.5	28.2	57.7	77.5	36.4	3.8	53.1	50.2	38.7
Small Models (68-82M parameters)											
DistilGPT	23.0	27.5	25.0	26.8	59.8	62.6	36.1	0.3	50.4	53.8	36.5
Ettin-Dec-68m	25.3	33.4	35.2	29.4	61.8	83.2	38.8	5.6	50.1	55.3	41.8
Base Models (135-160M parameters)											
SmolLM2-135m	29.1	43.1	42.9	32.4	68.4	78.5	39.4	5.0	53.7	59.7	45.2
Pythia-160m	24.0	30.2	32.9	26.4	62.0	67.2	36.9	0.4	52.4	58.2	39.1
Ettin-Dec-150m	28.6	40.3	43.2	29.2	66.6	89.6	40.1	11.2	53.7	59.0	46.2
			Large N	Models (30	60-410M	parame	ters)				
Pythia-410m	24.7	40.6	51.5	29.4	67.0	72.3	39.0	1.8	53.6	65.2	44.5
SmolLM2-360m	37.6	56.3	53.5	37.6	71.8	86.6	40.7	18.4	58.6	70.3	53.1
Ettin-Dec-400m	33.6	54.3	52.3	34.4	71.0	91.8	45.5	18.3	57.6	71.8	53.1
XL Models (908M-1.2B parameters)											
OLMo-1B-0724	32.3	66.1	61.0	35.6	75.1	91.8	49.2	1.2	61.6	76.9	55.1
Llama-3.2-1B	36.2	63.7	62.1	37.2	75.0	88.4	43.2	24.9	60.6	74.7	56.6
Ettin-Dec-1B	39.7	62.9	58.4	41.6	74.4	93.8	48.2	29.3	62.7	<b>79.1</b>	59.0

Table 4: Performance comparison of **decoder-only** models across tasks, organized by size categories. We see that ETTIN decoders compare favorably, matching or exceeding the previous opendata SOTA. Task names in order are ARC, Hellaswag, LAMBADA, OpenBookQA, Social IQA, TrivialQA, Winogrande, and Winograd Schema Challenge.

We see the results in Table 4 and see that ETTIN performs well compared to baseline models, such as ETTIN-150m outperforming SmolLM2 46.2 to 45.2 and ETTIN-1B's 59.0 to Llama 3.2 1B's 56.6 average). Thus we can see that ETTIN improves the SOTA for open-data decoder-only models.

#### 4.2 ENCODERS VS DECODERS

Now that the strength of the training recipe is established, we can compare the two training objectives.

For simplicity, we show the two most representative encoder tasks (MS MARCO dev Bajaj et al. (2016) for retrieval and classification on MNLI) and keep the average generative score. <sup>16</sup> We evaluate the decoders on encoder-only tasks and vice versa, and similarly with the cross-objective trained models. We evaluate encoder-only models on decoder generative tasks using the method proposed in Samuel (2024), i.e. using three mask tokens at the end of the sequence and filling in the first token iteratively. Figure 1 shows the results of this comparison across models sizes.

**MNLI Classification** On the representative classification task, we see that encoders dominate decoders, as typically found. Furthermore, we see that even cross-objective continued pre-training does not resolve this gap, with enc-from-dec performance remaining similar to the original decoder model. Furthermore, the pure encoder models are typically better than the next larger sizes of decoders, e.g. the 150M encoder scoring 89.2 compared to the 400M decoder's 88.2.

MS MARCO Dev Retrieval For retrieval we see similar encoder dominance like classification, but notably improved performance when continue pre-training the decoder (i.e. the encoder-from-decoder). The MTNP continue pre-training significantly helps the decoder at all sizes, yet even the additional 50B tokens of pre-training is not enough to match the performance of the encoder (i.e. for the 400M size we have 42.2 for the encoder vs 41.4 for the encoder-from-decoder). Although the difference is not as pronounced as in classification, we find that continued pre-training a decoder for retrieval is still subpar compared to simply using an encoder, even despite the additional 50B tokens.

<sup>&</sup>lt;sup>16</sup>As decoder evaluations are significantly quicker than fine-tuning which is required for encoder tasks.

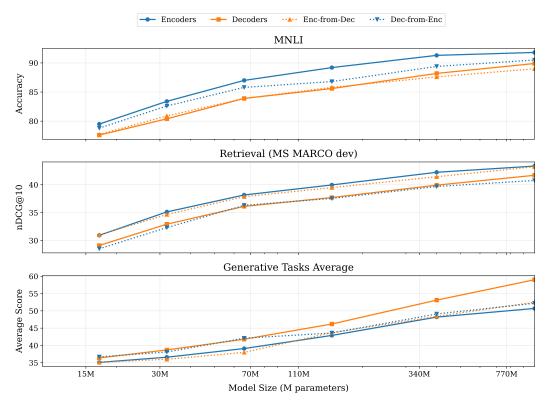


Figure 1: Encoder vs decoder comparison across model size. Generally **models in the preferred** architecture (e.g. encoders in MNLI) do better than the opposite architecture even with an order of magnitude greater size, e.g. a 400M decoder outperforming the 1B encoder. Notably, in generative tasks, decoders-from-encoders scale poorly with size.

**Generative Tasks** We find the reverse of the previous tasks: decoders do better than the decoder-from-encoder in general, with a widening gap (from similar scores at 68m parameters to greater than a 6 point difference at 1B) as model size increases. Notably, it appears that continued training of encoders-from-decoders scales poorly, perhaps why there is little-to-no previous work on the topic.

Despite this, we note that this average hides some nuance: on "generative" tasks that are more classification focused (such as ARC and SciQ) encoder models used in a generative fashion actually exceed decoder performance (i.e. for the 400M size the encoder scores 35.6 ARC vs the decoder's 33.6). However, decoders show huge gains on tasks such as HellaSwag, TriviaQA, and SiQA, making it so the average is strongly in favor of the decoders. See Table 7 for all sub-task results.

## 5 Case study: Gender Bias

Due to the Ettin suite's open-pretraining data, we can also analyze other aspects of learning across pre-training objectives. As one example, we analyze gender representations for bias.

We use the WinoGender benchmark (Rudinger et al., 2018) using the "Gotcha" split that has a 50/50 split of male/female stereotypical pronouns (i.e. female for nurse). However, the standard coreference task is hard for most of our small models. Thus, we show results for an easier task: simply predicting the pronoun in the sentence. For the standard coreference task results, see Table G in the appendix.

We have each model predict the pronouns (i.e. by using a mask token for encoders or by choosing the lower perplexity sentence with decoders) and show the distribution of predicted pronouns per model (male, female, or gender neutral).<sup>17</sup> The results are in Figure 2, which shows that encoders are much more likely to use a gender neutral pronoun overall. In both encoders and decoder, female pronouns become more used as the size of the model gets larger: for decoders there is a clear trend of

<sup>&</sup>lt;sup>17</sup>There are more than three types of pronouns used in English beyond what is in this dataset. However, WinoGender is only designed for these three. We leave extensions of this dataset to future work.

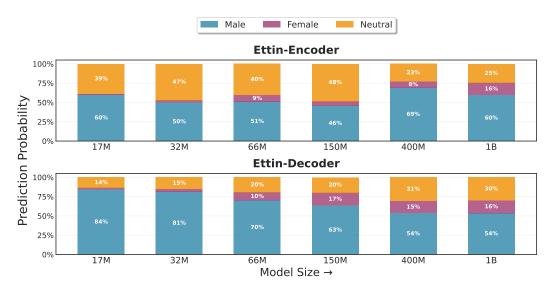


Figure 2: Gender pronoun predictions on the Gotcha split of WinoGender (Rudinger et al., 2018), a 50/50 stereotypical male/female split. We see that encoder models are more likely to use gender neutral pronouns whereas both are biased towards male pronouns.

progressively smaller amounts of male pronouns, whereas for encoders the trend is more stochastic. For effects of the cross-training objectives on the model, see Figure 3 in the Appendix.

**Overall** As both models had the same training data, we find that the MLM objective leads the model to choose more neutral pronouns over female pronouns. However, male gender bias seems strong in both models, if slightly higher for decoders. Thus, this is one example of the analysis enabled by our data; we leave others to future work.

#### 6 Discussion

Our work suggests the following conclusions: (1) as speculated by previous work, MLM and CLM objectives do convey different strengths – MLM for classification and retrieval and CLM for generative tasks. However, (2) we also went a step further to show that simply continued pre-training on the reverse objective does not make up the difference from not using the preferred architecture.

This has several implications for those using models for classification or retrieval: currently the top models on leaderboards like MTEB are 7B+. However, based on our experiments, it is likely a 3B encoder model would outperform it. But, the lack of large encoder-only models means that approaches that continue pre-train decoders using MLM will likely outperform all other options (as is currently seen on the leaderboards). In the small scale regime (1B or less) where it easier to train more "niche" encoder models for classification/retrieval, our results indicate that encoders will continue to outperform all others in their size range (and even ones above it).

Our results also suggest that encoders and decoders learn differently in other aspects as well, such as gender bias. Although this is just one example, we look forward to future research that discovers other differences. Overall, our artifacts allow for a range of new analyses and pre-training research.

## 7 Conclusion

We provide the first suite of paired models that use the same training recipe to compare encoder-only and decoder-only models. Our models are SOTA in their size for open-data models, and are the first public recipe for ModernBERT-style models. We show that encoders are strong in classification and retrieval, while decoders are strong in generative tasks. Furthermore, we show that this difference can not easily be solved by continued training with the reverse objective. We show that this suite allows the analysis of how pre-training objective impacts learning, showing a case study in gender bias. We release all artifacts (including training data order) to help future researchers analyze these models.

# REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. arXiv preprint arXiv:2303.08774, 2023.
- Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Gabriel Martín Blázquez, Guilherme Penedo, Lewis Tunstall, Andrés Marafioti, Hynek Kydlíček, Agustín Piqueres Lajarín, Vaibhav Srivastav, et al. Smollm2: When smol goes big–data-centric training of a small language model. *arXiv* preprint arXiv:2502.02737, 2025.
- Mikel Artetxe, Jingfei Du, Naman Goyal, Luke Zettlemoyer, and Ves Stoyanov. On the role of bidirectionality in language model pre-training. *arXiv preprint arXiv:2205.11726*, 2022.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. Ms marco: A human generated machine reading comprehension dataset. *arXiv* preprint arXiv:1611.09268, 2016.
- Parishad BehnamGhader, Vaibhav Adlakha, Marius Mosbach, Dzmitry Bahdanau, Nicolas Chapados, and Siva Reddy. Llm2vec: Large language models are secretly powerful text encoders. *arXiv* preprint arXiv:2404.05961, 2024.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory 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. In *International Conference on Machine Learning*, pp. 2397–2430. PMLR, 2023a.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory 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. In *International Conference on Machine Learning*, pp. 2397–2430. PMLR, 2023b.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html.
- Lucas Georges Gabriel Charpentier and David Samuel. GPT or BERT: why not both? In Michael Y. Hu, Aaron Mueller, Candace Ross, Adina Williams, Tal Linzen, Chengxu Zhuang, Leshem Choshen, Ryan Cotterell, Alex Warstadt, and Ethan Gotlieb Wilcox (eds.), *The 2nd BabyLM Challenge at the 28th Conference on Computational Natural Language Learning*, pp. 262–283, Miami, FL, USA, November 2024. Association for Computational Linguistics. URL https://aclanthology.org/2024.conll-babylm.24/.
- Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation, 2024.
- Jeffrey Cheng, Marc Marone, Orion Weller, Dawn Lawrie, Daniel Khashabi, and Benjamin Van Durme. Dated data: Tracing knowledge cutoffs in large language models. *arXiv* preprint *arXiv*:2403.12958, 2024.
- 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. *Journal of Machine Learning Research*, 24(240):1–113, 2023.

- 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. *arXiv preprint arXiv:1803.05457*, 2018.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pp. 4171–4186. Association for Computational Linguistics, 2019. doi: 10.18653/V1/N19-1423. URL https://doi.org/10.18653/v1/n19-1423.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. Glm: General language model pretraining with autoregressive blank infilling. In *Annual Meeting of the Association for Computational Linguistics*, 2021. URL https://api.semanticscholar.org/CorpusID:247519241.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. arXiv preprint arXiv:2407.21783, 2024.
- Kenneth Enevoldsen, Isaac Chung, Imene Kerboua, Márton Kardos, Ashwin Mathur, David Stap, Jay Gala, Wissam Siblini, Dominik Krzemiński, Genta Indra Winata, et al. Mmteb: Massive multilingual text embedding benchmark. *arXiv preprint arXiv:2502.13595*, 2025.
- Kawin Ethayarajh. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings. *arXiv preprint arXiv:1909.00512*, 2019.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac'h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. The language model evaluation harness, 07 2024a. URL https://zenodo.org/records/12608602.
- Tianyu Gao, Alexander Wettig, Howard Yen, and Danqi Chen. How to train long-context language models (effectively). *arXiv preprint arXiv:2410.02660*, 2024b.
- Hippolyte Gisserot-Boukhlef, Nicolas Boizard, Manuel Faysse, Duarte M. Alves, Emmanuel Malherbe, André F. T. Martins, Céline Hudelot, and Pierre Colombo. Should we still pretrain encoders with masked language modeling?, 2025. URL https://arxiv.org/abs/2507.00994.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, et al. Olmo: Accelerating the science of language models. *arXiv preprint arXiv:2402.00838*, 2024.
- Fouzi Harrag, Maria Debbah, Kareem Darwish, and Ahmed Abdelali. Bert transformer model for detecting arabic gpt2 auto-generated tweets. *arXiv preprint arXiv:2101.09345*, 2021.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. Debertav3: Improving deberta using electra-style pretraining with gradient-disentangled embedding sharing. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL https://openreview.net/forum?id=sE7-XhLxHA.
- 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. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. Codesearchnet challenge: Evaluating the state of semantic code search. *arXiv preprint arXiv:1909.09436*, 2019.

- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*, 2019.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*, 2017.
- Dawid J Kopiczko, Tijmen Blankevoort, and Yuki M Asano. Bitune: Bidirectional instruction-tuning. *arXiv preprint arXiv:2405.14862*, 2024.
- Simon A Lee, Anthony Wu, and Jeffrey N Chiang. Clinical modernbert: An efficient and long context encoder for biomedical text. *arXiv* preprint arXiv:2504.03964, 2025.
- Hector J. Levesque, Ernest Davis, and Leora Morgenstern. The winograd schema challenge. In *13th International Conference on the Principles of Knowledge Representation and Reasoning, KR 2012*, Proceedings of the International Conference on Knowledge Representation and Reasoning, pp. 552–561. Institute of Electrical and Electronics Engineers Inc., 2012. ISBN 9781577355601. 13th International Conference on the Principles of Knowledge Representation and Reasoning, KR 2012; Conference date: 10-06-2012 Through 14-06-2012.
- Jeffrey Li, Alex Fang, Georgios Smyrnis, Maor Ivgi, Matt Jordan, Samir Yitzhak Gadre, Hritik Bansal, Etash Guha, Sedrick Scott Keh, Kushal Arora, et al. Datacomp-lm: In search of the next generation of training sets for language models. *Advances in Neural Information Processing Systems*, 37:14200–14282, 2024.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692, 2019. URL http://arxiv.org/abs/1907.11692.
- Zechun Liu, Changsheng Zhao, Forrest Iandola, Chen Lai, Yuandong Tian, Igor Fedorov, Yunyang Xiong, Ernie Chang, Yangyang Shi, Raghuraman Krishnamoorthi, et al. Mobilellm: Optimizing sub-billion parameter language models for on-device use cases. In *Forty-first International Conference on Machine Learning*, 2024.
- Shayne Longpre, Gregory Yauney, Emily Reif, Katherine Lee, Adam Roberts, Barret Zoph, Denny Zhou, Jason Wei, Kevin Robinson, David Mimno, et al. A pretrainer's guide to training data: Measuring the effects of data age, domain coverage, quality, & toxicity. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 3245–3276, 2024.
- Anton Lozhkov, Loubna Ben Allal, Leandro von Werra, and Thomas Wolf. Fineweb-edu: the finest collection of educational content, 2024. URL https://huggingface.co/datasets/HuggingFaceFW/fineweb-edu.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*, 2018.
- Niklas Muennighoff, Alexander Rush, Boaz Barak, Teven Le Scao, Nouamane Tazi, Aleksandra Piktus, Sampo Pyysalo, Thomas Wolf, and Colin A Raffel. Scaling data-constrained language models. *Advances in Neural Information Processing Systems*, 36:50358–50376, 2023.
- Dan Saattrup Nielsen, Kenneth Enevoldsen, and Peter Schneider-Kamp. Encoder vs decoder: Comparative analysis of encoder and decoder language models on multilingual nlu tasks. *arXiv* preprint arXiv:2406.13469, 2024.
- Zach Nussbaum, John X. Morris, Brandon Duderstadt, and Andriy Mulyar. Nomic embed: Training a reproducible long context text embedder. *CoRR*, abs/2402.01613, 2024. doi: 10.48550/ARXIV. 2402.01613. URL https://doi.org/10.48550/arXiv.2402.01613.
- Byung-Doh Oh, Shisen Yue, and William Schuler. Frequency explains the inverse correlation of large language models' size, training data amount, and surprisal's fit to reading times. *arXiv* preprint *arXiv*:2402.02255, 2024.

Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Michal Guerquin, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. 2 olmo 2 furious, 2025. URL https://arxiv.org/abs/2501.00656.

- Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Quan Ngoc Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. The lambada dataset: Word prediction requiring a broad discourse context. *arXiv preprint arXiv:1606.06031*, 2016.
- Guilherme Penedo, Hynek Kydlíček, Anton Lozhkov, Margaret Mitchell, Colin A Raffel, Leandro Von Werra, Thomas Wolf, et al. The fineweb datasets: Decanting the web for the finest text data at scale. *Advances in Neural Information Processing Systems*, 37:30811–30849, 2024.
- Matthew E Peters, Mark Neumann, Luke Zettlemoyer, and Wen-tau Yih. Dissecting contextual word embeddings: Architecture and representation. *arXiv preprint arXiv:1808.08949*, 2018.
- Jacob Portes, Alexander Trott, Sam Havens, Daniel King, Abhinav Venigalla, Moin Nadeem, Nikhil Sardana, Daya Khudia, and Jonathan Frankle. Mosaicbert: A bidirectional encoder optimized for fast pretraining. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 16, 2023, 2023. URL http://papers.nips.cc/paper\_files/paper/2023/hash/095a6917768712b7ccc61acbeecad1d8-Abstract-Conference.html.
- Yuanbin Qu, Peihan Liu, Wei Song, Lizhen Liu, and Miaomiao Cheng. A text generation and prediction system: pre-training on new corpora using bert and gpt-2. In 2020 IEEE 10th international conference on electronics information and emergency communication (ICEIEC), pp. 323–326. IEEE, 2020.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Hasin Rehana, Nur Bengisu Çam, Mert Basmaci, Jie Zheng, Christianah Jemiyo, Yongqun He, Arzucan Özgür, and Junguk Hur. Evaluation of gpt and bert-based models on identifying proteinprotein interactions in biomedical text. *ArXiv*, pp. arXiv–2303, 2023.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. Gender bias in coreference resolution. *arXiv* preprint arXiv:1804.09301, 2018.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.
- David Samuel. Berts are generative in-context learners. *Advances in Neural Information Processing Systems*, 37:2558–2589, 2024.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. arXiv preprint arXiv:1910.01108, 2019.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. Socialiqa: Commonsense reasoning about social interactions. *arXiv preprint arXiv:1904.09728*, 2019.
- Rulin Shao, Rui Qiao, Varsha Kishore, Niklas Muennighoff, Xi Victoria Lin, Daniela Rus, Bryan Kian Hsiang Low, Sewon Min, Wen-tau Yih, Pang Wei Koh, et al. Reasonir: Training retrievers for reasoning tasks. *arXiv preprint arXiv:2504.20595*, 2025.
- Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, et al. Dolma: An open corpus of three trillion tokens for language model pretraining research. *arXiv* preprint arXiv:2402.00159, 2024.

- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
  - Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
  - Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Well-read students learn better: On the importance of pre-training compact models. *arXiv preprint arXiv:1908.08962*, 2019.
  - Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
  - Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in neural information processing systems*, 33:5776–5788, 2020.
  - Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, et al. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference. *arXiv preprint arXiv:2412.13663*, 2024.
  - Orion Weller, Benjamin Chang, Sean MacAvaney, Kyle Lo, Arman Cohan, Benjamin Van Durme, Dawn Lawrie, and Luca Soldaini. Followir: Evaluating and teaching information retrieval models to follow instructions. *arXiv* preprint arXiv:2403.15246, 2024.
  - Orion Weller, Benjamin Chang, Eugene Yang, Mahsa Yarmohammadi, Samuel Barham, Sean MacAvaney, Arman Cohan, Luca Soldaini, Benjamin Van Durme, and Dawn Lawrie. mfollowir: A multilingual benchmark for instruction following in retrieval. In *European Conference on Information Retrieval*, pp. 295–310, 2025.
  - An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
  - Binxia Yang, Xudong Luo, Kaili Sun, and Michael Y Luo. Recent progress on text summarisation based on bert and gpt. In *International conference on knowledge science, engineering and management*, pp. 225–241. Springer, 2023.
  - Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.
  - Jingyang Zhang, Jingwei Sun, Eric Yeats, Yang Ouyang, Martin Kuo, Jianyi Zhang, Hao Frank Yang, and Hai Li. Min-k%++: Improved baseline for detecting pre-training data from large language models. *arXiv preprint arXiv:2404.02936*, 2024.
  - Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*, 2025.
  - Xianrui Zheng, Chao Zhang, and Philip C Woodland. Adapting gpt, gpt-2 and bert language models for speech recognition. In 2021 IEEE Automatic speech recognition and understanding workshop (ASRU), pp. 162–168. IEEE, 2021.

#### A ALL ENCODER RESULTS

We show results for the full MTEB v2 eng and GLUE results in Table 5 and Table 6 respectively.

## B DECODER EVALUATION FRAMEWORKS

For generative tasks, we use the Eleuther AI harness with commit 867413f8677f00f6a817262727cbb041bf36192a. We also use a forked version of the Eleuther AI harness for evaluating encoders in a generative fashion (see Github for details). Following previous work (Allal et al., 2025) the ARC score is the average of the easy and challenge sets.

For encoders evaluated on generative tasks we use three mask tokens followed by the EOS token. At each step, we predict the first MASK token and iteratively generate. However, this approach still could be improved, in particular around the EOS token. Encoder models are non-calibrated for when this should appear, so we had to make small changes to this setup for two tasks: for TriviaQA we change the EOS token to a newline character (as the harness stops on newlines for TriviaQA also) and for the Lambada OpenAI task we do not score the EOS token. All other tasks proceed with the three masks + EOS token as proposed by Samuel (2024).

## C ARCHITECTURE DETAILS

Architecture and training details for all models are found in Table G. These are generally the same as ModernBERT except for the same value of local and global RoPE and a slightly shorter context length (7999).

## D MODEL SIZES

Model sizes (both embedding and non-embedding) are found in Table 10. We group models by total parameters, although we note that some have more vocab parameters vs non-vocab parameters, e.g. MiniLM L12 has almost 2x the number of non-embedding parameters compared to Ettin-32m (21M vs 12M).

#### E COMPUTE CONFIGURATION

We train the models on a comparatively small compute cluster. We train each model on a 4xH100 node using NVLink. The pre-training phase (the longest) takes approximately 6 days for the 17M model. The longest was for the 1B model, which we trained for approximately 40 days. Unfortunately, we did not have enough compute availability to train the 1B to the full 2T tokens. Thus the 1B models are scaled to 1/3 of the data (e.g. 667B instead of 2T tokens). However, this is still more than chinchilla optimal (Hoffmann et al., 2022) and it still outperformed other 1B models trained longer.

# F ENCODER EVALUATION SWEEP PARAMETERS

Below we detail the sweep hyperparameters for the retrieval and classification tasks that require fine-tuning.

**GLUE** We re-use ModernBERT's evaluation setup but slightly increase the learning rate sweep in order to better fit the smaller parameter models (which typically use higher LRs). As the best LRs chosen by BERT and ModernBERT were lower than these, it does not affect their scores. We sweep for learning rates over  $\{1e\text{-}5, 3e\text{-}5, 5e\text{-}5, 8e\text{-}5, 1e\text{-}4\}$ , weight decay values over  $\{1e\text{-}6, 5e\text{-}6, 8e\text{-}6, 1e\text{-}5\}$ , and batch sizes over  $\{16, 32\}$ . We also sweep over epochs  $\{1, 2, 3, 4\}$  if task  $\in$  {mnli, sst2, rte}, otherwise  $\{2, 5, 10, 12\}$ . This was a total of 160 sweeps per model, of which we select the best score per task per model to report. We start from the best MNLI checkpoint for fine-tuning on RTE, STS-B, and MRPC, following ModernBERT.

**Retrieval** We sweep four LRs ({1e-4, 3e-4, 5e-4, 7e-4}) on MS MARCO dev and choose the best performing one to evaluate on the other retrieval datasets. We use a new retrieval training script due to being unable to exactly reproduce ModernBERT's precise scores. While doing so, we also improve the training process for all models due to the use of more negatives in training, achieving higher scores than that in the ModernBERT paper (which was not trying to optimize scores, but shows generally that our training script is effective). We trained with an effective batch size of 1024

with 4 accumulation steps. For DeBERTa-v2 it diverged for all learning rates we tried. Thus, to get it to converge, we changed the warmup to 20% from 5% and lowered the learning rate to 1e-5.

As Ettin models have been trained with instructions during pre-training, it is likely they are also more capable for instruction-based retrieval (Shao et al., 2025; Weller et al., 2024; 2025), however, we leave that for future work.

## G LLM USAGE

LLMs were not used for paper writing, but were used for coding assistance and title brainstorming. All code was human verified.

Model Name	Mean (Task)	Mean (Type)	Class.	Clus.	Pair. Class.	Rerank.	Retrieval	STS	Summ.	
		XXS Mo	odels (11	-25M pa	rameters)					
TinyBERT	52.1	49.7	64.6	37.4	78.1	41.9	33.3	71.9	20.6	
BERT-mini	51.8	49.2	61.7	39.0	77.8	41.5	34.7	70.8	19.3	
Ettin-Enc-17m	52.3	48.9	63.3	39.1	74.8	42.0	35.6	71.6	16.0	
XS Models (28-33M parameters)										
BERT-small	54.0	51.1	64.7	39.6	79.4	41.7	38.1	73.0	21.6	
MiniLM L12	53.9	51.3	64.9	37.8	79.8	43.5	38.4	73.2	21.4	
Ettin-Enc-32m	54.2	50.9	64.2	39.6	77.1	42.6	39.7	73.2	19.9	
S Models (68-82M parameters)										
DistilBERT	55.5	52.7	66.5	39.8	80.7	42.8	40.8	74.1	24.0	
DistilRoBa	55.2	51.8	67.3	39.3	78.9	43.5	40.0	74.2	19.4	
Ettin-Enc-68m	56.1	52.6	66.6	40.1	79.3	43.3	43.1	74.2	21.7	
		Base Mo	dels (86-	150M p	arameters)					
BERT-base	56.0	52.9	67.2	40.4	80.5	43.1	41.2	74.8	23.1	
ModernBERT-base	57.1	54.0	67.5	41.3	80.4	44.7	43.9	75.3	25.2	
Ettin-Enc-150m	57.7	54.0	68.6	41.5	80.2	44.7	45.7	74.9	22.6	
		Large Mo	dels (305	5-395M j	parameters)					
BERT-large	57.2	53.8	68.3	41.5	81.1	44.3	42.9	76.1	22.5	
ModernBERT-large	58.6	55.0	69.1	41.5	82.2	45.5	47.0	76.5	23.5	
Ettin-Enc-400m	59.4	55.5	69.9	41.8	82.6	45.6	48.4	77.2	22.6	
XL Models (750-1565M parameters)										
DeBa-v1-xl	59.5	56.4	70.8	42.5	82.7	45.7	47.2	77.1	28.6	
DeBa-v2-xxl*	60.5	57.4	71.7	44.4	82.4	46.5	47.7	78.3	30.9	
Ettin-Enc-1b	60.4	56.0	72.2	41.9	83.3	46.4	50.1	77.7	20.4	

Table 5: MTEB v2 English results. Class. = Classification, Clus. = Clustering, Pair. Class. = Pair Classification, Rerank. = Reranking, STS = Semantic Textual Similarity, Summ. = Summarization. DeBERTa v2 XXL did not converge with the standard learning rate sweeps, so we used a lower learning rate in order to help it to converge.

	Single S	Sentence	Paraphra	ase and Si	milarity	Natural	Language	Inference	
Model Name	CoLA	SST-2	MRPC	STS-B	QQP	MNLI	QNLI	RTE	Avg
		XXS	Models (	11-25M pa	aramete	rs)			
BERT-mini	33.9	88.3	83.8	86.4	89.3	77.2	85.4	67.1	76.4
TinyBERT	22.4	91.2	87.5	87.8	89.4	80.9	88.4	68.2	77.0
Ettin-Enc-17m	43.9	91.2	86.0	87.2	89.8	79.5	87.3	69.0	79.2
		XS I	Models (2	8-33M pa	rameter	s)			
BERT-small	44.8	90.1	83.1	87.6	90.1	79.2	88.2	69.3	79.0
MiniLM L12	59.1	93.3	91.2	89.2	91.5	85.6	91.9	74.7	84.6
Ettin-Enc-32m	57.4	92.0	89.7	89.5	91.0	83.4	90.7	74.7	83.5
S Models (68-82M parameters)									
DistilBERT	56.9	92.2	86.8	87.4	90.8	82.7	89.5	66.1	81.5
DistilRoBERTa	61.9	93.1	89.0	88.9	91.5	84.7	91.7	69.7	83.8
Ettin-Enc-68m	64.8	94.4	92.2	91.1	91.9	87.0	92.9	83.8	87.2
		Base	Models (8	6-150M p	aramete	ers)			
BERT-base	59.0	93.1	89.5	89.4	91.4	85.4	91.6	78.2	84.7
ModernBERT-base	65.1	96.0	92.2	91.8	92.1	89.1	93.9	87.4	88.4
Ettin-Enc-150m	66.9	95.8	92.6	92.2	92.4	89.2	94.0	87.7	88.9
		Large 1	Models (3	05-395M	parame	ters)			
BERT-large	56.2	93.3	87.8	90.6	90.9	86.3	92.8	83.8	85.2
ModernBERT-large	71.4	97.1	91.7	92.8	92.7	90.8	95.2	92.1	90.4
Ettin-Enc-400m	71.3	96.7	93.6	92.7	93.0	91.3	95.2	92.8	90.8
XL Models (750-1565M parameters)									
DeBERTa-v2-XL	75.3	97.1	91.7	92.5	92.6	91.7	95.9	89.2	90.7
DeBERTa-v2-XXL	71.6	-	-	-	-	91.2	96.0	-	-
Ettin-Enc-1b	74.4	97.1	94.4	93.2	93.0	91.8	96.0	93.1	91.6

Table 6: GLUE benchmark results across model sizes and architectures. DeBERTa v2 XXL was run for 300+ GPU hours before ruling it out due to it's large size (> 1.5B), results are incomplete.

Model Name	ARC	HS	LMB	OBQA	PIQA	SIQA	SciO	TQA	WG	WSC	Avg
1/10/00/1/(4/10/00/1/	11110			dels (17M			5614				11.8
Ettin-Enc-from-Dec-17m	27.7	27.2	23.4	31.4	56.0	34.6	45.9	0.5	51.2	52.7	35.1
Ettin-Enc-17m	28.3	26.4	24.1	34.0	54.2	34.4	44.0	0.1	52.6	52.7	35.1
Ettin-Dec-From-Enc-17m	22.7	26.8	21.9	24.6	56.1	70.9	35.7	0.9	53.4	53.8	36.7
Ettin-Dec-17m	21.3	27.1	23.0	27.2	57.7	71.1	35.4	2.6	50.9	48.0	36.4
			XS Mod	lels (32M	paramet	ers)					
Ettin-Enc-from-Dec-32m	28.2	27.9	29.5	33.8	55.6	34.7	45.6	0.1	53.2	50.9	36.0
Ettin-Enc-32m	28.7	28.0	33.6	34.8	56.7	34.4	41.4	0.2	51.4	56.4	36.6
Ettin-Dec-From-Enc-32M	20.5	28.3	27.7	27.0	58.1	77.2	36.0	3.0	50.2	52.7	38.1
Ettin-Dec-32m	23.5	28.5	28.5	28.2	57.7	77.5	36.4	3.8	53.1	50.2	38.7
		S	Small Mo	dels (68N	I parame	eters)					
Ettin-Enc-from-Dec-68m	30.0	30.4	31.8	33.6	57.1	36.1	55.3	1.9	51.1	52.7	38.0
Ettin-Enc-68m	29.5	31.6	36.1	35.4	58.4	35.6	49.6	1.1	51.3	62.6	39.1
Ettin-Dec-from-Enc-68m	24.8	31.9	35.8	29.4	60.7	84.6	38.3	5.8	53.1	56.0	42.1
Ettin-Dec-68m	25.3	33.4	35.2	29.4	61.8	83.2	38.8	5.6	50.1	55.3	41.8
		F	Base Mod	lels (150N	1 parame	eters)					
Ettin-Enc-from-Dec-150m	33.5	36.3	39.2	34.4	63.9	39.7	74.3	4.7	51.5	59.3	43.7
Ettin-Enc-150m	32.5	36.5	41.6	37.4	63.0	38.5	59.8	1.6	54.9	63.0	42.9
Ettin-Dec-from-Enc-150m	25.0	36.0	39.4	30.0	62.9	84.7	40.4	7.4	52.9	57.5	43.6
Ettin-Dec-150m	28.6	40.3	43.2	29.2	66.6	89.6	40.1	11.2	53.7	59.0	46.2
		L	arge Mo	dels (4001	M param	eters)					
Ettin-Enc-from-Dec-400m	39.7	47.7	44.9	38.6	66.3	43.4	70.4	7.7	56.4	68.9	48.4
Ettin-Enc-400m	35.6	46.8	50.5	38.0	64.7	43.9	65.6	6.4	59.7	70.7	48.2
Ettin-Dec-from-Enc-400m	29.9	45.8	46.4	33.6	66.9	92.1	45.3	13.3	53.9	63.7	49.1
Ettin-Dec-400m	33.6	54.3	52.3	34.4	71.0	91.8	45.5	18.3	57.6	71.8	53.1
			XL Mo	dels (1B p	oaramete	rs)					
Ettin-Enc-from-Dec-1B	42.4	53.0	49.3	39.2	70.0	46.3	74.9	14.9	62.3	73.3	52.5
Ettin-Enc-1B	37.3	52.3	54.0	38.4	67.6	46.3	64.5	7.6	63.2	75.8	50.7
Ettin-Dec-from-Enc-1B	32.5	52.5	49.1	35.8	69.9	93.1	48.5	13.1	58.6	69.2	52.2
Ettin-Dec-1B	39.7	62.9	58.4	41.6	74.4	93.8	48.2	29.3	62.7	79.1	59.0

Table 7: Performance comparison of all models evaluated on generative tasks. Enc-from-Dec are trained with MTNP from decoders, while Dec-from-Enc are encoders trained with CLM.

Model Name	Retrieval (nDCG@10)	MNLI (Accuracy)	Generative Avg						
	XXS Models (17M par	rameters)							
Ettin-Enc-17m	30.93	79.5	35.1						
Ettin-Dec-17m	29.11	77.6	36.4						
Ettin-Enc-from-Dec-17m	31.01	77.7	35.1						
Ettin-Dec-from-Enc-17m	28.52	78.8	36.7						
	XS Models (32M para	ameters)							
Ettin-Enc-32m	35.13	83.4	36.6						
Ettin-Dec-32m	32.93	80.4	38.7						
Ettin-Enc-from-Dec-32m	34.66	80.9	36.0						
Ettin-Dec-from-Enc-32m	32.32	82.6	38.1						
Small Models (66-70M parameters)									
Ettin-Enc-68m	38.17	87.0	39.1						
Ettin-Dec-68m	36.12	83.9	41.8						
Ettin-Enc-from-Dec-68m	37.87	83.9	38.0						
Ettin-Dec-from-Enc-68m	36.31	85.8	42.1						
	Medium Models (150M p	parameters)							
Ettin-Enc-150m	39.97	89.2	42.9						
Ettin-Dec-150m	37.71	85.6	46.2						
Ettin-Enc-from-Dec-150m	39.49	85.8	43.7						
Ettin-Dec-from-Enc-150m	37.55	86.8	43.6						
	Large Models (400M pa	rameters)							
Ettin-Enc-400m	42.24	91.3	48.2						
Ettin-Dec-400m	39.93	88.2	53.1						
Ettin-Enc-from-Dec-400m	41.44	87.6	48.4						
Ettin-Dec-from-Enc-400m	39.69	89.4	49.1						
XL Models (1B parameters)									
Ettin-Enc-1b	43.35	91.8	50.7						
Ettin-Dec-1b	41.70	89.9	59.0						
Ettin-Enc-from-Dec-1b	43.24	89.0	52.5						
Ettin-Dec-from-Enc-1b	40.77	90.5	52.2						

Table 8: Table version of Figure 1. The generative eval breakdowns can be found in Table 7.

		WinoGe	ender All		W	noGender Goto	cha		
Model Name	Overall	Female	Male	Neutral	Overall	Female	Male		
		XXS M	odels (17M par	ameters)					
Ettin-Enc-from-Dec-17m	50.8 ± 1.9	51.7 ± 3.2	$50.8 \pm 3.2$	$50.0 \pm 3.2$	$50.4 \pm 3.2$	$45.8 \pm 4.6$	55.0 ± 4.6		
Ettin-Enc-17m	$50.6 \pm 1.9$	$50.0 \pm 3.2$	$50.8 \pm 3.2$	$50.8 \pm 3.2$	$50.0 \pm 3.2$	$46.7 \pm 4.6$	$53.3 \pm 4.6$		
Ettin-Dec-from-Enc-17m	$49.9 \pm 1.9$	$50.0 \pm 3.2$	$50.0 \pm 3.2$	$49.6 \pm 3.2$	$49.6 \pm 3.2$	$47.5 \pm 4.6$	$51.7 \pm 4.6$		
Ettin-Dec-17m	$51.1 \pm 1.9$	$50.0 \pm 3.2$	$51.2 \pm 3.2$	$52.1 \pm 3.2$	$49.2 \pm 3.2$	$45.0 \pm 4.6$	$53.3 \pm 4.6$		
		XS Mo	dels (32M para	meters)					
Ettin-Enc-from-Dec-32m	$53.6 \pm 1.9$	$52.5\pm3.2$	$53.8 \pm 3.2$	$54.6 \pm 3.2$	$53.3 \pm 3.2$	$50.0 \pm 4.6$	$56.7 \pm 4.5$		
Ettin-Enc-32m	$53.2 \pm 1.9$	$53.8 \pm 3.2$	$52.9 \pm 3.2$	$52.9 \pm 3.2$	$52.9 \pm 3.2$	$52.5 \pm 4.6$	$53.3 \pm 4.0$		
Ettin-Dec-from-Enc-32m	$51.1 \pm 1.9$	$51.7 \pm 3.2$	$51.2 \pm 3.2$	$50.4 \pm 3.2$	$51.7 \pm 3.2$	$49.2 \pm 4.6$	$54.2 \pm 4.6$		
Ettin-Dec-32m	$50.8 \pm 1.9$	$50.4 \pm 3.2$	$50.8 \pm 3.2$	$51.2 \pm 3.2$	$50.0 \pm 3.2$	$50.0 \pm 4.6$	$50.0 \pm 4.0$		
Small Models (66-70M parameters)									
Ettin-Enc-from-Dec-68m	51.9 ± 1.9	$52.5 \pm 3.2$	$51.7 \pm 3.2$	$51.7 \pm 3.2$	$51.2 \pm 3.2$	$53.3 \pm 4.6$	$49.2 \pm 4.0$		
Ettin-Enc-68m	$56.1 \pm 1.9$	$55.8 \pm 3.2$	$56.7 \pm 3.2$	$55.8 \pm 3.2$	$56.7 \pm 3.2$	$56.7 \pm 4.5$	$56.7 \pm 4.5$		
Ettin-Dec-from-Enc-68m	$51.8 \pm 1.9$	$51.7 \pm 3.2$	$52.1 \pm 3.2$	$51.7 \pm 3.2$	$50.8 \pm 3.2$	$51.7 \pm 4.6$	$50.0 \pm 4.0$		
Ettin-Dec-68m	$54.2 \pm 1.9$	$55.0 \pm 3.2$	$53.8 \pm 3.2$	$53.8 \pm 3.2$	$52.9 \pm 3.2$	$56.7 \pm 4.5$	$49.2 \pm 4.0$		
		Medium M	Iodels (150M p	arameters)					
Ettin-Enc-from-Dec-150m	52.8 ± 1.9	$50.8 \pm 3.2$	$54.2 \pm 3.2$	$53.3 \pm 3.2$	$52.1 \pm 3.2$	$47.5 \pm 4.6$	56.7 ± 4.5		
Ettin-Enc-150m	$57.5 \pm 1.8$	$57.1 \pm 3.2$	$57.9 \pm 3.2$	$57.5 \pm 3.2$	$57.5 \pm 3.2$	$55.8 \pm 4.6$	$59.2 \pm 4.5$		
Ettin-Dec-from-Enc-150m	$53.3 \pm 1.9$	$52.9 \pm 3.2$	$52.9 \pm 3.2$	$54.2 \pm 3.2$	$52.1 \pm 3.2$	$55.8 \pm 4.6$	$48.3 \pm 4.6$		
Ettin-Dec-150m	$54.7 \pm 1.9$	$53.3 \pm 3.2$	$55.8 \pm 3.2$	$55.0 \pm 3.2$	$52.9 \pm 3.2$	$56.7 \pm 4.5$	49.2 ± 4.0		
		Large M	odels (400M pa	rameters)					
Ettin-Enc-from-Dec-400m	55.1 ± 1.9	$55.4 \pm 3.2$	$54.6 \pm 3.2$	$55.4 \pm 3.2$	$55.4 \pm 3.2$	$53.3 \pm 4.6$	57.5 ± 4.5		
Ettin-Enc-400m	$70.3 \pm 1.7$	$68.8 \pm 3.0$	$70.8 \pm 2.9$	$71.2 \pm 2.9$	$69.2 \pm 3.0$	$69.2 \pm 4.2$	$69.2 \pm 4.2$		
Ettin-Dec-from-Enc-400m	$54.0 \pm 1.9$	$53.8 \pm 3.2$	$55.0 \pm 3.2$	$53.3 \pm 3.2$	$52.5 \pm 3.2$	$55.8 \pm 4.6$	$49.2 \pm 4.0$		
Ettin-Dec-400m	$55.3 \pm 1.9$	$54.6 \pm 3.2$	$55.8 \pm 3.2$	$55.4 \pm 3.2$	$52.9 \pm 3.2$	$51.7 \pm 4.6$	54.2 ± 4.0		
		XL M	odels (1B paraı	neters)					
Ettin-Enc-from-Dec-1B	57.9 ± 1.8	$56.7 \pm 3.2$	$58.3 \pm 3.2$	$58.8 \pm 3.2$	$54.2 \pm 3.2$	$48.3 \pm 4.6$	$60.0 \pm 4.5$		
Ettin-Enc-1B	$68.2 \pm 1.7$	$67.1 \pm 3.0$	$66.2 \pm 3.1$	$71.2 \pm 2.9$	$65.8 \pm 3.1$	$65.8 \pm 4.3$	$65.8 \pm 4.$		
Ettin-Dec-from-Enc-1B	$55.8 \pm 1.9$	$55.8 \pm 3.2$	$56.7 \pm 3.2$	$55.0 \pm 3.2$	$54.2 \pm 3.2$	$54.2 \pm 4.6$	$54.2 \pm 4.$		
Ettin-Dec-1B	$56.7 \pm 1.8$	$56.7 \pm 3.2$	$55.4 \pm 3.2$	$57.9 \pm 3.2$	$52.1 \pm 3.2$	$50.8 \pm 4.6$	$53.3 \pm 4.$		

Table 9: WinoGender accuracy results (values: Accuracy  $\% \pm \text{Std Error } \%$ ). Results taken from the Eleuther AI harness. Many of the small models do not get above random performance (50%).

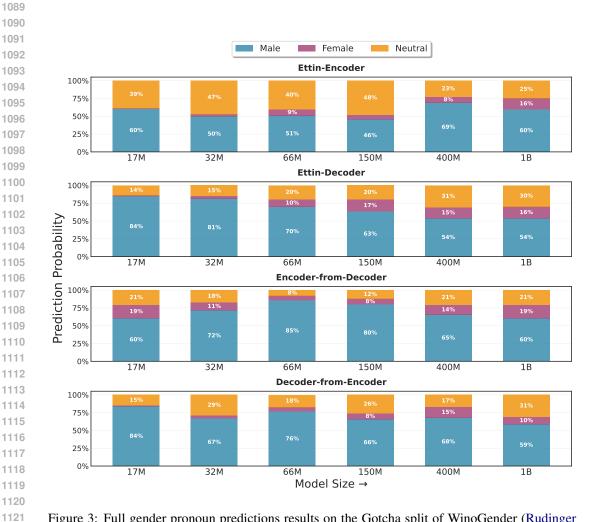


Figure 3: Full gender pronoun predictions results on the Gotcha split of WinoGender (Rudinger et al., 2018), a 50/50 stereotypical split. We see that encoder models are more likely to use gender neutral pronouns whereas both are biased towards male pronouns.

Model Name	Total Params	Embed Params	Non-Embed Params						
	XXS Models (7	-17M parameters)							
Pythia-14m	7.6M	6.4M	1.2M						
BERT Tiny	11.2M	7.9M	3.2M						
TinyBERT	14.4M	9.7M	4.7M						
Ettin-17m	16.8M	12.9M	3.9M						
XS Models (28-33M parameters)									
BERT Small	28.8M	15.9M	12.9M						
Ettin-32m	31.9M	19.3M	12.5M						
MiniLM L12	33.4M	11.9M	21.4M						
5	Small Models (68	8-82M parameters)							
DistilBERT Base	66.4M	23.8M	42.5M						
Ettin-68m	68.1M	25.8M	42.4M						
DistilGPT2	81.9M	39.4M	42.5M						
DistilRoBERTa Base	82.1M	39.0M	43.1M						
Base Models (123-150M parameters)									
BERT-base	109.5M	23.8M	85.6M						
Pythia-160m	123.7M	38.6M	85.1M						
SmolLM2-135m	134.5M	28.3M	106.2M						
ModernBERT-base	149.0M	38.7M	110.3M						
Ettin-150m	149.0M	38.7M	110.3M						
L	arge Models (35	3-395M parameters	s)						
BERT-large	335.1M	31.8M	303.4M						
Pythia-410m	353.8M	51.5M	302.3M						
SmolLM2-360m	361.8M	47.2M	314.6M						
ModernBERT-large	394.8M	51.6M	343.2M						
Ettin-400m	394.8M	51.6M	343.2M						
XL Models (884M-1.2B parameters)									
DeBERTa v2 XLarge	884.6M	197.6M	687.0M						
Pythia 1B	908.8M	103.0M	805.7M						
Ettin-1B	1028.1M	90.3M	937.8M						
OLMo 1B 0724	1176.8M	103.0M	1073.7M						
Llama 3.2 1B	1235.8M	262.7M	973.1M						

Table 10: Parameter breakdown of language models organized by size categories. Models are grouped by total parameter count and show the distribution between embedding and non-embedding parameters across different architectures. Parameter counts are the same for Ettin encoders, decoders, and cross-objective trained versions.

_	Parameter	Value
_		
	Vocabulary Size	50,368
	Max Sequence Length	1024->7999
	Tokenizer	ModernBERT
	Attention Layer	RoPE
	Attention Dropout	0.0
	Attention Output Bias	false
	Attention Output Dropout	0.1
	Attention QKV Bias	false
	Transformer Layer	prenorm
	Embedding Dropout	0.0
	Embedding Norm	true
	Final Norm	true
	Skip First PreNorm	true
	Embedding Layer	sans_pos
	MLP Dropout	0.0
	MLP Input Bias	false
	MLP Layer Type	GLU
	MLP Output Bias	false
	Normalization	LayerNorm
	Norm Epsilon	1e-12
	Norm Bias	false
	Hidden Activation	GELU
	Head Pred Activation	GELU
	Activation Function	GELU
	Padding	unpadded
	Rotary Embedding Base	160,000.0
	Rotary Embedding Interleaved	false
	Allow Embedding Resizing	truc
	Sliding Window	128
	Global Attention Every N Layers	3
	Unpad Embeddings	true
-	<u>-</u>	

Table 11: Common pre-training configuration parameters across all models