
TRANSFORMER MODELS: AN INTRODUCTION AND CATALOG

Xavier Amatriain
 Los Gatos, CA 95032
 xavier@amatriain.net

February 16, 2023

ABSTRACT

I have a terrible memory for names. In the past few years we have seen the meteoric appearance of dozens of models of the Transformer family, all of which have funny, but not self-explanatory, names. The goal of this post is to offer a short and simple catalog and classification of the most popular Transformer models. In other words, I needed a Transformers cheat-sheet and could not find a good enough one online, so I thought I'd write my own. I hope it can be useful to you too.

Contents

1	Introduction: What are Transformers	3
1.1	Encoder/Decoder architecture	3
1.2	Attention	5
1.3	What are Transformers used for and why are they so popular	5
1.4	RLHF	6
1.5	Diffusion	6
2	The Transformers catalog	8
2.1	Features of a Transformer	8
2.1.1	Pretraining Architecture	8
2.1.2	Pretraining Task	8
2.1.3	Application	9
2.2	Catalog table	9
2.3	Family Tree	9
2.4	Chronological timeline	9
2.5	Catalog List	9
2.5.1	ALBERT	9
2.5.2	AlphaFold	9
2.5.3	Anthropic Assistant	12
2.5.4	BART	13
2.5.5	BERT	13

2.5.6	Big Bird	13
2.5.7	BlenderBot3	14
2.5.8	BLOOM	14
2.5.9	ChatGPT	14
2.5.10	Chinchilla	15
2.5.11	CLIP	15
2.5.12	CM3	15
2.5.13	CTRL	16
2.5.14	DALL-E	16
2.5.15	DALL-E 2	16
2.5.16	Decision Transformers	17
2.5.17	DialoGPT	17
2.5.18	DistilBERT	17
2.5.19	DQ-BART	18
2.5.20	ELECTRA	18
2.5.21	ERNIE	18
2.5.22	Flamingo	19
2.5.23	Gato	19
2.5.24	GLaM	19
2.5.25	GLIDE	20
2.5.26	Global Context ViT	20
2.5.27	Gopher	20
2.5.28	GopherCite	21
2.5.29	GPT	21
2.5.30	GPT-2	21
2.5.31	GPT-3	22
2.5.32	GPT-3.5	22
2.5.33	InstructGPT	22
2.5.34	GPT-Neo	23
2.5.35	GPT-NeoX-20B	23
2.5.36	HTML	23
2.5.37	Imagen	23
2.5.38	Jurassic-1	24
2.5.39	LAMDA	24
2.5.40	mBART	24
2.5.41	Megatron	25
2.5.42	Minerva	25
2.5.43	MT-NLG (Megatron TouringNLG)	25
2.5.44	OPT	26

2.5.45	PalM	26
2.5.46	Pegasus	26
2.5.47	RoBERTa	27
2.5.48	SeeKer	27
2.5.49	Sparrow	27
2.5.50	StableDiffusion	28
2.5.51	Swin Transformer	28
2.5.52	Switch	28
2.5.53	T5	29
2.5.54	Trajectory Transformers	29
2.5.55	Transformer XL	29
2.5.56	Turing-NLG	30
2.5.57	ViT	30
2.5.58	Wu Dao 2.0	30
2.5.59	XLNet	31
2.5.60	XLNet	31
3	Further reading	31

1 Introduction: What are Transformers

Transformers are a class of deep learning models that are defined by some architectural traits. They were first introduced in the now famous "Attention is All you Need" paper by Google researchers in 2017 [1] (the paper has accumulated a whopping 38k citations in only 5 years) and associated blog post¹.

The Transformer architecture is a specific instance of the encoder-decoder models[2]² that had become popular just over the 2–3 years prior. Up until that point however, attention was just one of the mechanisms used by these models, which were mostly based on LSTM (Long Short Term Memory)[3] and other RNN (Recurrent Neural Networks)[4] variations. The key insight of the Transformers paper was that, as the title implies, attention could be used as the only mechanism to derive dependencies between input and output.

It is beyond the scope of this blog to go into all the details of the Transformer architecture. For that, I will refer you to the original paper above or to the wonderful The Illustrated Transformer³ post. That being said, we will briefly describe the most important aspects since we will be referring to them in the catalog below. Let's start with the basic architectural diagram from the original paper, and describe some of the components.

1.1 Encoder/Decoder architecture

A generic encoder/decoder architecture (see Figure 1) is made up of two models. The encoder takes the input and encodes it into a fixed-length vector. The decoder takes that vector and decodes it into the output sequence. The encoder and decoder are jointly trained to minimize the conditional log-likelihood. Once trained the encoder/decoder can generate an output given an input sequence or can score a pair of input/output sequences.

In the case of the original Transformer architecture, both encoder and decoder had 6 identical layers. In each of those 6 layers the Encoder has two sub layers: a multi-head attention layer, and a simple feed forward network. Each sublayer has a residual connection and a layer normalization. The output size of the Encoder is 512. The Decoder adds a third sublayer, which is another multi-head attention layer over the output of the Encoder. Besides, the other multi-head layer in the decoder is masked to prevent attention to subsequent positions.

¹<https://ai.googleblog.com/2017/08/transformer-novel-neural-network.html>

²<https://machinelearningmastery.com/encoder-decoder-long-short-term-memory-networks/>

³<https://jalammar.github.io/illustrated-transformer/>

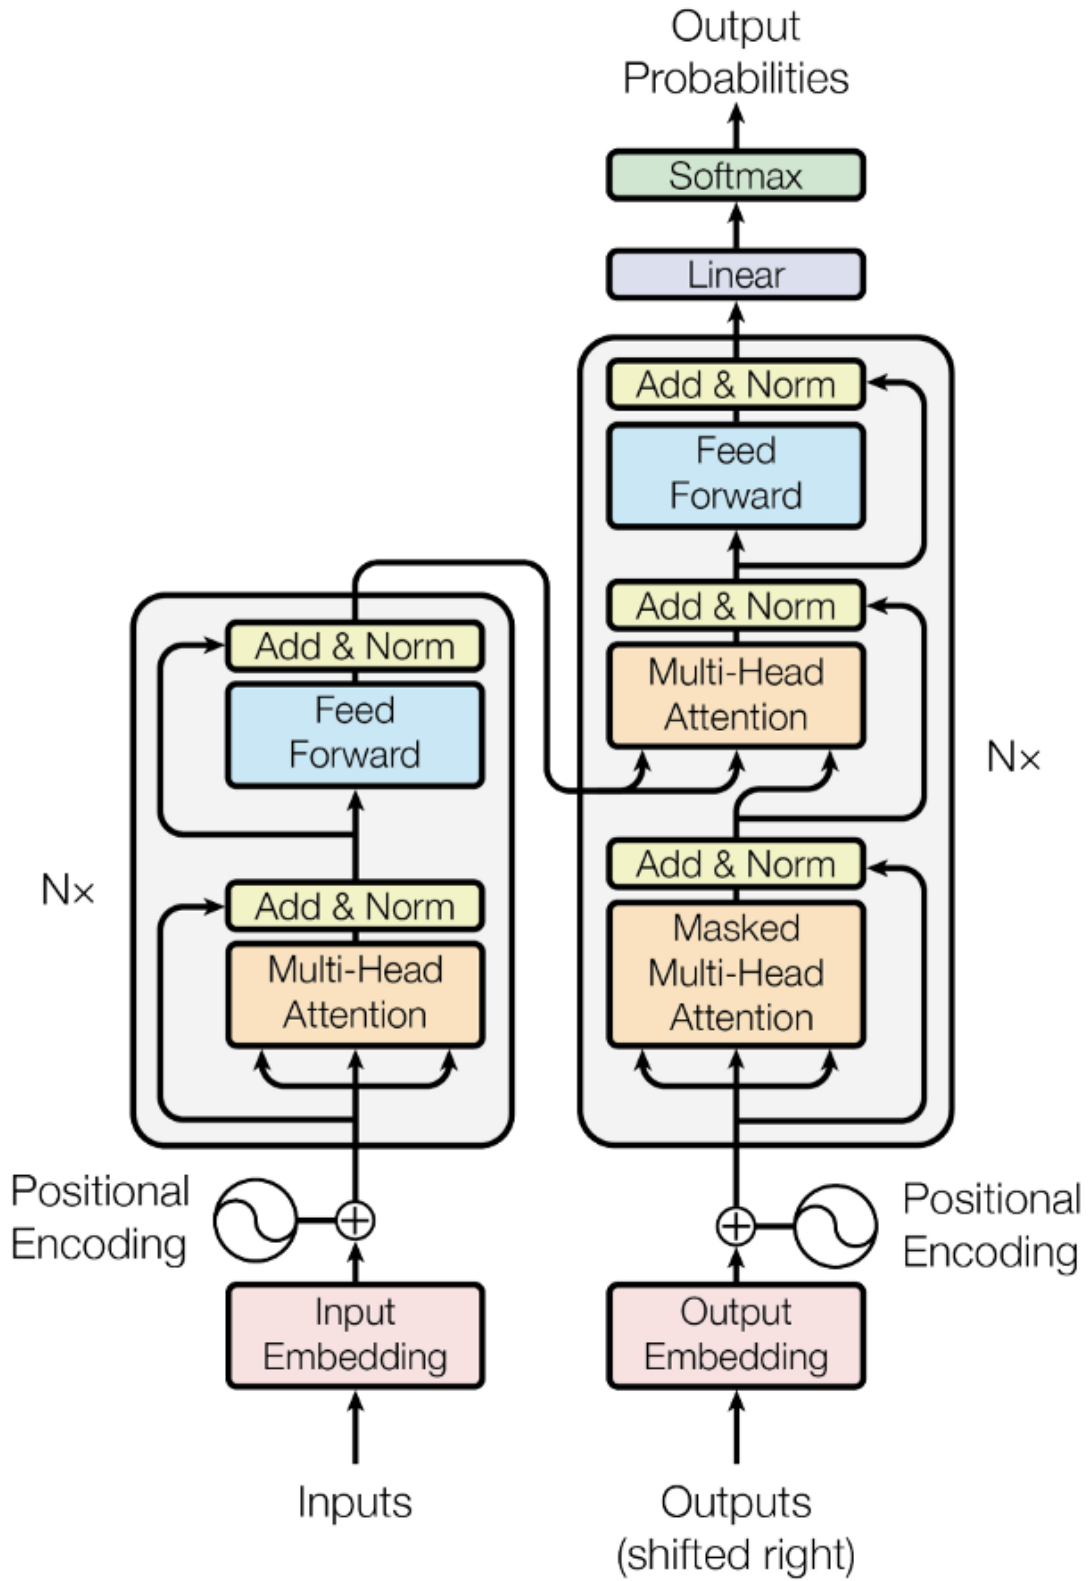


Figure 1: Transformer Architecture from [1]

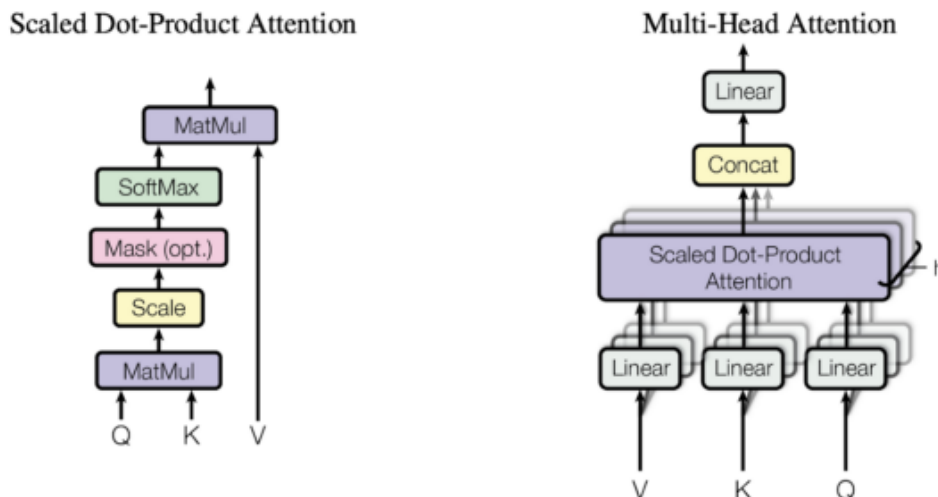


Figure 2: (left) Scaled Dot-Product Attention. (right) Multi-Head Attention consists of several attention layers running in parallel.

Figure 2: The Attention Mechanism

1.2 Attention

It is clear from the description above that the only “exotic” elements of the model architecture are the multi-headed attention, but, as described above, that is where the whole power of the model lies! So, what is attention anyway? An attention function is a mapping between a query and a set of key-value pairs to an output. The output is computed as a weighted sum of the values, where the weight assigned to each value is computed by a compatibility function of the query with the corresponding key. Transformers use multi-headed attention, which is a parallel computation of a specific attention function called scaled dot-product attention. I will refer you again to the The Illustrated Transformer⁴ post for many more details on how the attention mechanism works, but will reproduce the diagram from the original paper in Figure 2 so you get the main idea

There are several advantages of attention layers over recurrent and convolutional networks, the two most important being their lower computational complexity and their higher connectivity, especially useful for learning long-term dependencies in sequences.

1.3 What are Transformers used for and why are they so popular

The original transformer was designed for language translation, particularly from English to German. But, already the original paper showed that the architecture generalized well to other language tasks. This particular trend became quickly noticed by the research community. Over the next few months most of the leaderboards for any language-related ML task became completely dominated by some version of the transformer architecture (see for example the well known SQUAD leaderboard⁵ for question answer where all models at the top are ensembles of Transformers).

One of the key reasons Transformers were able to so quickly take over most NLP leaderboards is their ability to quickly adapt to other tasks, a.k.a. Transfer learning. Pretrained Transformer models can adapt extremely easily and quickly to tasks they have not been trained on, and that has huge advantages. As an ML practitioner, you no longer need to train a large model on a huge dataset. All you need to do is re-use the pretrained model on your task, maybe just slightly

⁴<https://jalamar.github.io/illustrated-transformer/>

⁵<https://rajpurkar.github.io/SQuAD-explorer/>

adapting it with a much smaller data set. A specific technique used to adapt pretrained models to a different task is the so-called fine tuning⁶.

It turns out that the capability of Transformers to adapt to other tasks is so great, that, while they were initially developed for language related tasks, they quickly became useful for other tasks ranging from vision[5] or audio and music⁷ applications all the way to playing chess[6] or doing math[7].

Of course all these applications would have not been possible if it wasn't because of the myriad of tools that made them readily available to anyone that could write a few lines of code. Not only were Transformers quickly integrated into the main AI frameworks (namely Pytorch⁸ and TF⁹), but they even enabled the creation of an entire company around them. Huggingface¹⁰, a startup that has raised over \$ 60M to this day, is almost entirely built around the idea of commercializing their open source Transformers library¹¹.

Last but not least, I would be remiss if I did not mention the impact of GPT-3[8] on the popularization of Transformers. GPT-3 is a Transformer model introduced by OpenAI in May 2020 as a follow up to their earlier GPT and GPT-2. The company made a big splash by introducing the model in a preprint[8] in which they claimed that the model was so powerful that they were not in a position to release it to the world. Since then, the model has not only been released, but also commercialized through a very large partnership¹² between OpenAI and Microsoft. GPT-3 powers over 300 different applications¹³, and is the foundation for OpenAI's commercial strategy (which is a lot to say for a company that has received over \$ 1B in funding).

1.4 RLHF

Reinforcement Learning from Human Feedback (or Preferences) aka RLHF (or RLHP) has become a huge addition to the AI toolkit as of lately. The concept was introduced already in 2017 in the paper ["Deep reinforcement learning from human preferences"](<https://arxiv.org/abs/1706.03741>). More recently though, it has been applied to ChatGPT and similar dialog agents like BlenderBot3 or Sparrow. The idea is pretty simple though: Once a language model is pretrained, we can generate different responses to a dialog and have Humans rank the results. We can use those ranking (aka preferences or feedback) to train a reward, in the reinforcement learning context (see Figure 3). You can read much more in these two wonderful posts by Huggingface¹⁴ or Weights and Bias¹⁵.

1.5 Diffusion

Diffusion models have become the new SOTA in image generation, clearly pushing aside the previous approaches such as GANs (Generative Adversarial Networks). What are diffusion models? They are a class of latent variable models trained variational inference. What this means in practice is that we train a deep neural network to denoise images blurred with some sort of noise function. Networks that are trained this way are in fact learning the latent space of what those images represent (see figure 4).

Diffusion models have relation with other generative models like the famous [Generative Adversarial Networks (GAN)]¹⁶, which they have mostly replaced in many applications and, particularly with (denoising) Autoencoders. Some authors¹⁷ will go as far as saying that Diffusion models are just a specific instance of autoencoders. However, they also admit that the small differences do transform their application, from the latent representation of autoencoders to the pure generative nature of Diffusion models.

⁶<https://huggingface.co/docs/transformers/training>

⁷<https://magenta.tensorflow.org/music-transformer>

⁸https://pytorch.org/tutorials/beginner/transformer_tutorial.html

⁹<https://www.tensorflow.org/text/tutorials/transformer>

¹⁰<https://huggingface.co/docs>

¹¹<https://github.com/huggingface/transformers>

¹²<https://openai.com/blog/openai-licenses-gpt-3-technology-to-microsoft/>

¹³<https://openai.com/blog/gpt-3-apps/>

¹⁴<https://huggingface.co/blog/rlhf>

¹⁵<https://wandb.ai/ayush-thakur/RLHF/reports/Understanding-Reinforcement-Learning-from-Human-Feedback-RLHF-Part-1>

¹⁶https://en.wikipedia.org/wiki/Generative_adversarial_network

¹⁷<https://benanne.github.io/2022/01/31/diffusion.html>

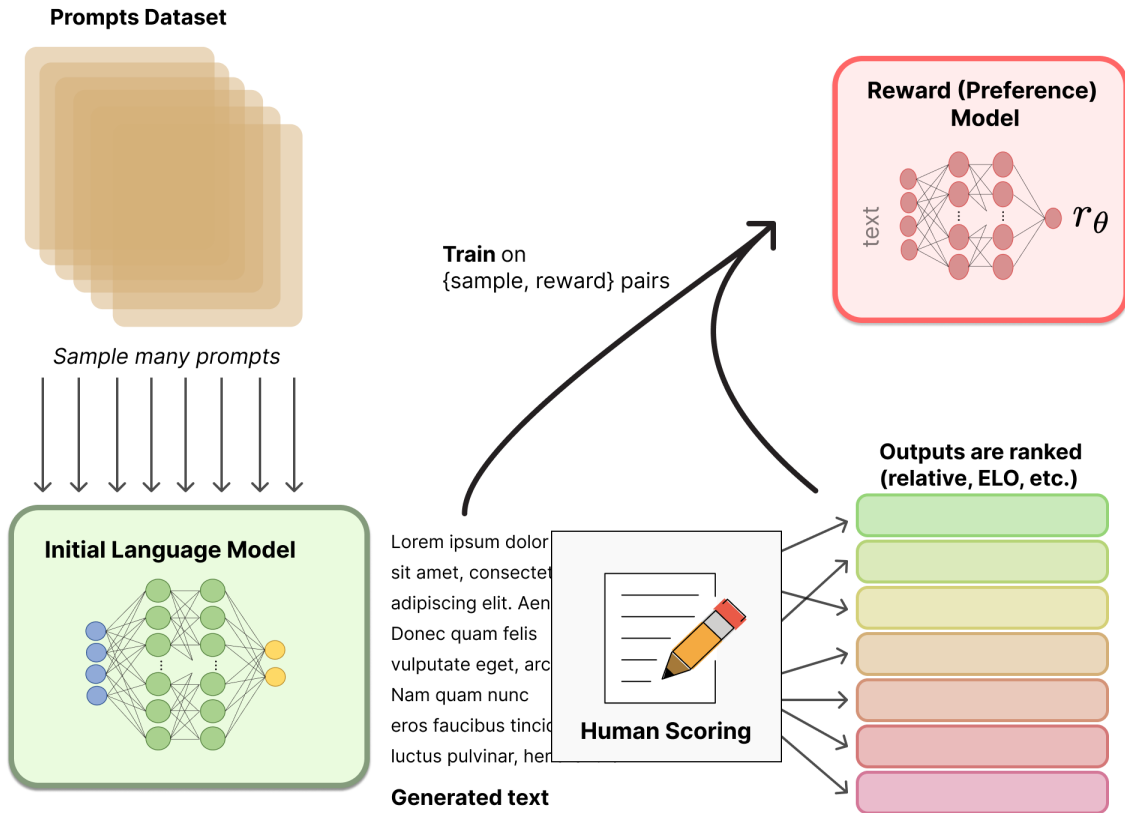


Figure 3: Reinforcement Learning with Human Feedback. From HuggingFace’s RLHF blog post at <https://huggingface.co/blog/rlhf>.

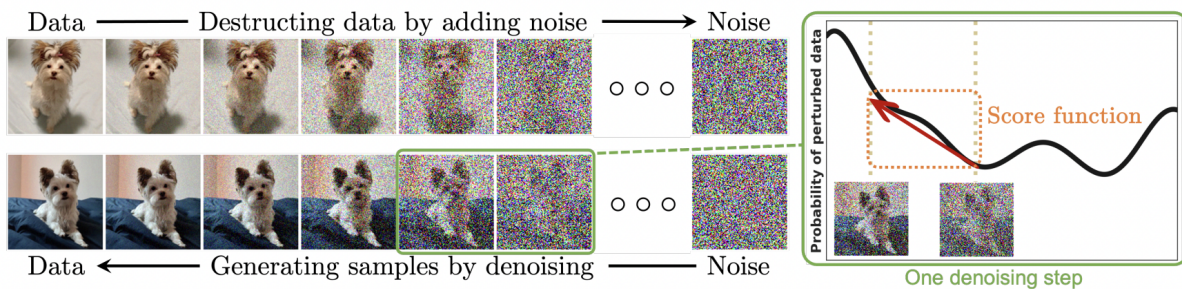


Fig. 2. Diffusion models smoothly perturb data by adding noise, then reverse this process to generate new data from noise. Each denoising step in the reverse process typically requires estimating the score function (see the illustrative figure on the right), which is a gradient pointing to the directions of data with higher likelihood and less noise.

Figure 4: Probabilistic diffusion model architecture from “Diffusion Models: A Comprehensive Survey of Methods and Applications” [9]

2 The Transformers catalog

Note: For all the models available in Huggingface, I decided to directly link to the page in the documentation since they do a fantastic job of offering a consistent format and links to everything else you might need, including the original papers. Only a few of the models are not included in Huggingface. For those, I try to include a link to their github if available or blog post if not. For all, I also include bibliographic reference.

2.1 Features of a Transformer

So hopefully by now you understand what Transformer models are, and why they are so popular and impactful. In this section I will introduce a catalog of the most important Transformer models that have been developed to this day. I will categorize each model according to the following properties: Pretraining Architecture, Pretraining Task, Compression, Application, Year, and Number of Parameters. Let's briefly define each of those:

2.1.1 Pretraining Architecture

We described the Transformer architecture as being made up of an Encoder and a Decoder, and that is true for the original Transformer. However, since then, different advances have been made that have revealed that in some cases it is beneficial to use only the encoder, only the decoder, or both.

Encoder Pretraining These models, which are also called bi-directional or auto-encoding, only use the encoder during pretraining, which is usually accomplished by masking words in the input sentence and training the model to reconstruct. At each stage during pretraining, attention layers can access all the input words. This family of models are most useful for tasks that require understanding complete sentences such as sentence classification or extractive question answering.

Decoder Pretraining Decoder models, often called auto-regressive, use only the decoder during a pretraining that is usually designed so the model is forced to predict the next word. The attention layers can only access the words positioned before a given word in the sentence. They are best suited for tasks involving text generation.

Transformer (Encoder-Decoder) Pretraining Encoder-decoder models, also called sequence-to-sequence, use both parts of the Transformer architecture. Attention layers of the encoder can access all the words in the input, while those of the decoder can only access the words positioned before a given word in the input. The pretraining can be done using the objectives of encoder or decoder models, but usually involves something a bit more complex. These models are best suited for tasks revolving around generating new sentences depending on a given input, such as summarization, translation, or generative question answering.

2.1.2 Pretraining Task

When training a model we need to define a task for the model to learn on. Some of the typical tasks, such as predicting the next word or learning to reconstruct masked words were already mentioned above. "Pre-trained Models for Natural Language Processing: A Survey"[10] includes a pretty comprehensive taxonomy of pretraining tasks, all of which can be considered self-supervised:

1. **Language Modeling (LM):** Predict next token (in the case of unidirectional LM) or previous and next token (in the case of bidirectional LM)
2. **Masked Language Modeling (MLM):** mask out some tokens from the input sentences and then trains the model to predict the masked tokens by the rest of the tokens
3. **Permuted Language Modeling (PLM):** same as LM but on a random permutation of input sequences. A permutation is randomly sampled from all possible permutations. Then some of the tokens are chosen as the target, and the model is trained to predict these targets.
4. **Denosing Autoencoder (DAE):** take a partially corrupted input (e.g. Randomly sampling tokens from the input and replacing them with "[MASK]" elements. randomly deleting tokens from the input, or shuffling sentences in random order) and aim to recover the original undistorted input.
5. **Contrastive Learning (CTL):** A score function for text pairs is learned by assuming some observed pairs of text that are more semantically similar than randomly sampled text. It includes:
 - **Deep InfoMax (DIM):** maximize mutual information between an image representation and local regions of the image;

- **Replaced Token Detection (RTD):** predict whether a token is replaced given its surroundings;
- **Next Sentence Prediction (NSP):** train the model to distinguish whether two input sentences are continuous segments from the training corpus; and
- **Sentence Order Prediction (SOP):** Similar to NSP, but uses two consecutive segments as positive examples, and the same segments but with their order swapped as negative examples

2.1.3 Application

Here we will note what are the main practical applications of the Transformer model. Most of these applications will be in the language domain (e.g. question answering, sentiment analysis, or entity recognition). However, as mentioned before, some Transformer models have also found applications well beyond NLP and are also included in the catalog.

2.2 Catalog table

Figure 5 is a screenshot of a large table where I have tabulated all the models. If you are interested in the table, access it directly in ¹⁸

2.3 Family Tree

The diagram in figure 6 is a simple view that highlights the different families of transformers and how they relate to each other.

2.4 Chronological timeline

Another interesting perspective of the catalog is to see it as a chronological timeline. In Figure 7 you will find all the transformers in the catalog sorted by their date of publication. In this first visualization, the Y-axis is only used to cluster transformers of related heritage/family.

In Figure 8, the Y-axis represents model size in millions of parameters. You won't be able to see all the models in the catalog since many fall right on the same time and size, so please refer to the previous image for that.

2.5 Catalog List

Finally, here is the full list view that might be easier to follow along in some cases:

2.5.1 ALBERT

- **Reference:**¹⁹[11]
- **Family:** BERT
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** MLM/NSP
- **Extension:** Compressed version of BERT using parameter sharing, which is much more efficient given the same number of parameters
- **Application:** Same as BERT
- **Date (of first known publication):** 09/2019
- **Num. Params:** b12M, Large = 18M, XLarge = 60M*
- **Corpus:** Same as BERT
- **Lab:** Google

2.5.2 AlphaFold

- **Reference:**²⁰[12]

¹⁸<https://docs.google.com/spreadsheets/d/1ltyrAB6BL29c0v2fSpNqnnq2vbX8UrH147d7FkIf6t4>

¹⁹https://huggingface.co/docs/transformers/model_doc/albert

²⁰<https://www.deepmind.com/publications/highly-accurate-protein-structure-prediction-with-alphafold>

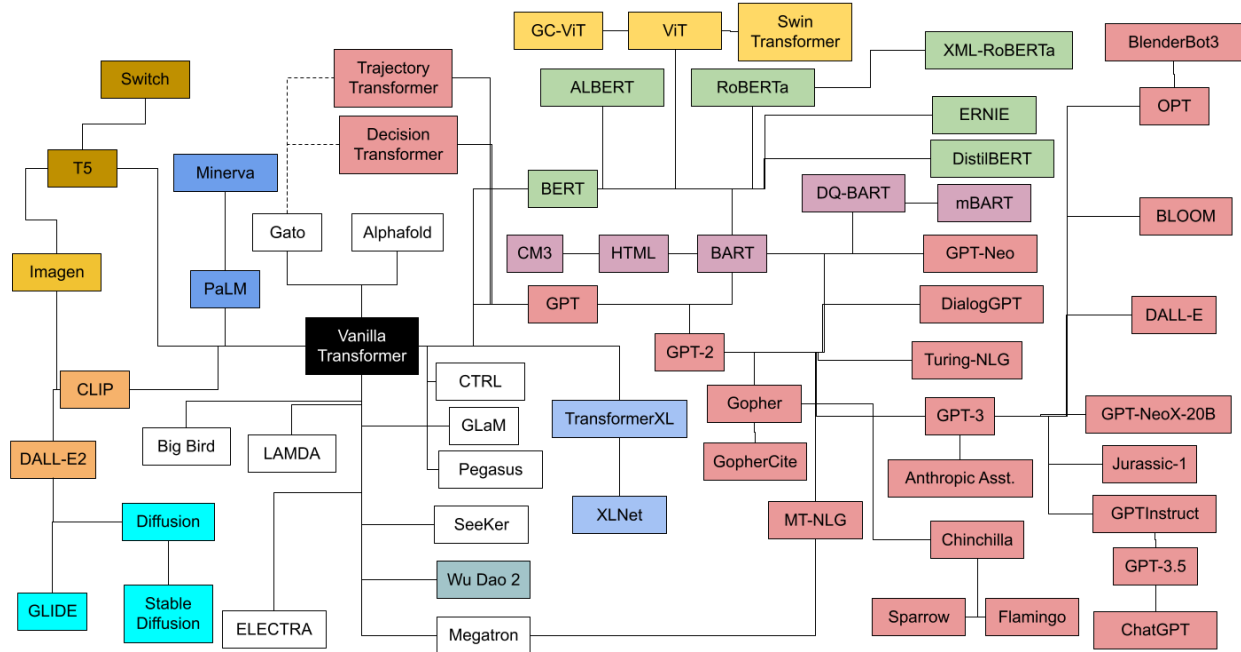


Figure 6:

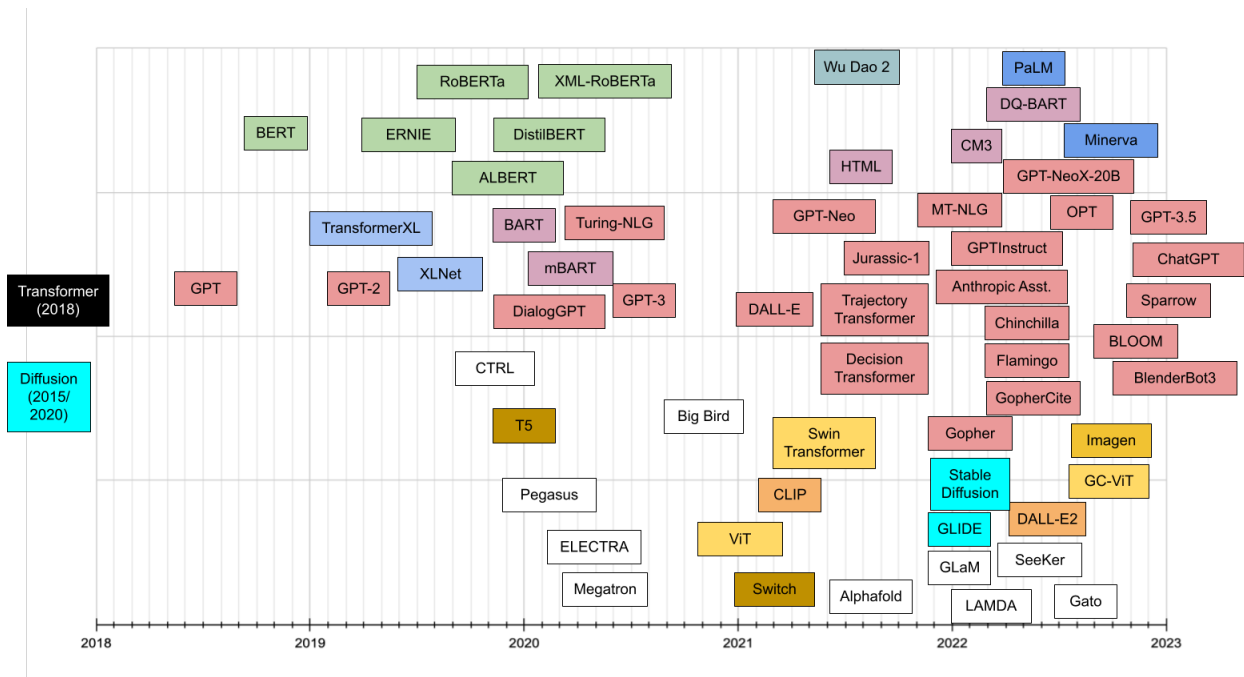


Figure 7: Transformer timeline. Colors describe Transformer family.




Figure 8: Transformer timeline. On the vertical axis, number of parameters. Colors describe Transformer family.

- **Family:** SE(3) Transformer²¹
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** Protein folding prediction*ion of BERT using parameter sharing, which is much more efficient given the same number of parameters
- **Extension:**The original Alphafold used a BERT-style transformer. The details of Alphafold’s Transformer are not known, but it is believed it is an extension of the SE(3)-Tranfomer, a 3-D equivariant Transformer (see this blog post²²)
- **Application:** Same as BERT
- **Date (of first known publication):** 09/2019
- **Num. Params:**b12M, Large = 18M, XLarge = 60M*
- **Corpus:**Same as BERT
- **Lab:**Google

2.5.3 Anthropic Assistant

- **Reference:**²³ see also²⁴ [13, 14]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** Protein folding prediction*ion of BERT using parameter sharing, which is much more efficient given the same number of parameters
- **Extension:**These models do not introduce novelties at the architecture/pretraining level and they are based on GPT-3 but rather focuses on how to improve alignment through fine-tuning and prompting. Note that the Anthropic Assistant includes several models optimized for different tasks. Latest versions of this work focus on the benefits of RLHF

²¹<https://arxiv.org/abs/2006.10503>

²²<https://fabianfuchsm1.github.io/alphafold2/>

²³<https://arxiv.org/abs/2112.00861>

²⁴<https://arxiv.org/abs/2204.05862>

- **Application:** Different models with different applications from general dialog to code assistant.
- **Date (of first known publication):** 12/2021
- **Num. Params:** 10M to 52B
- **Corpus:** 400B tokens from filtered Common Crawl and Books. They also create several Dialogue Preference datasets for the RLHF training.
- **Lab:** Anthropic

2.5.4 BART

- **Reference:**²⁵[15]
- **Family:** BERT for encoder, GPT for Decoder
- **Pretraining Architecture:** Encoder/Decoder
- **Pretraining Task:** DAE
- **Extension:** It can be seen as a generalization of BERT and GPT in that it combines ideas from both in the encoder and decoder
- **Application:** Mostly text generation but also some text understanding tasks*
- **Date (of first known publication):** 10/2019*
- **Num. Params:** 10 % more than BERT
- **Corpus:** Same as RoBERTa (160Gb of news, books, stories)
- **Lab:** Facebook

2.5.5 BERT

- **Reference:**²⁶[16]
- **Family:** BERT
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** MLM/NSP
- **Extension:** It can be seen as a generalization of BERT and GPT in that it combines ideas from both in the encoder and decoder
- **Application:** General Language Understanding and Question Answering. Many other language applications followed
- **Date (of first known publication):** 10/2018
- **Num. Params:** Base = 110M, Large = 340MT
- **Corpus:** Toronto Book Corpus and Wikipedia (3.3B Tokens)
- **Lab:** Google

2.5.6 Big Bird

- **Reference:**²⁷[17]
- **Family:** BERT
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** MLM
- **Extension:** Big Bird can extend other architectures such as BERT, Pegasus, or RoBERTa by using a sparse attention mechanism that eliminates the quadratic dependency thus making it more suitable for longer sequences
- **Application:** Particularly well suited for longer sequences, not only in text but also e.g. in genomics
- **Date (of first known publication):** 07/2020

²⁵https://huggingface.co/docs/transformers/model_doc/bart

²⁶https://huggingface.co/docs/transformers/model_doc/bert

²⁷https://huggingface.co/docs/transformers/model_doc/big_bird

- **Num. Params:** Depends on the overall architecture
- **Corpus:** Books, CC-News, Stories and Wikipedia)
- **Lab:** Google

2.5.7 BlenderBot3

- **Reference:**²⁸[18]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** BlenderBot 3 is based on a pre-trained OPT. It adds features needed for a dialog agent such as long-term memory or the ability to search the internet. It is also fine-tuned for some specific tasks given human feedback on them.
- **Application:** Same as GPT-3
- **Date (of first known publication):** 08/2022
- **Num. Params:** 175B
- **Corpus:** 180B tokens = RoBERTa + the Pile + PushShift.io Reddit
- **Lab:** Facebook

2.5.8 BLOOM

- **Reference:**²⁹
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Main difference to GPT-3 is that it uses full attention instead of sparse attention
- **Application:** Same as GPT-3
- **Date (of first known publication):** 07/2022
- **Num. Params:** 176B
- **Corpus:** 366B tokens (1.5 TB of text data) multilingual dataset
- **Lab:** Big Science/Huggingface

2.5.9 ChatGPT

- **Reference:**³⁰
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** ChatGPT takes a GPT3.5 (aka GPT3 Davinci-003) pretrained model and uses RLHF to finetune the model mostly like described in InstructGPT but with slight differences in the data collection. ChatGPT is also more than a model since it includes extensions for Memory Store and retrieval similar to BlenderBot3
- **Application:** Dialog agents
- **Date (of first known publication):** 10/2022
- **Num. Params:** Same as GPT3
- **Corpus:** Same as GPT3 + datasets generated for RLHF
- **Lab:** OpenAI

²⁸<https://arxiv.org/abs/2208.03188>

²⁹https://huggingface.co/docs/transformers/model_doc/bloom

³⁰<https://openai.com/blog/chatgpt/>

2.5.10 Chinchilla

- **Reference:**³¹[19]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Same as Gopher but with optimizations to reduce model size and therefore training/inference time with equal or superior performance
- **Application:** Same as Gopher/GPT3
- **Date (of first known publication):** 03/2022
- **Num. Params:**70B
- **Corpus:** Massive Text
- **Lab:** Deepmind

2.5.11 CLIP

- **Reference:**³²[20]
- **Family:** CLIP (Also using Resnet, ViT, and vanilla transformer for text)
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** predict which of the $N \times N$ possible (image, text) pairings across a batch actually occurred
- **Extension:** Combines Resnet and ViT for the visual encoding with Transformer for the Textual encoder
- **Application:** Image/object classification
- **Date (of first known publication):** 02/2021
- **Num. Params:**?
- **Corpus:** WIT (WebImageText) - 400 million text,image pairs
- **Lab:** OpenAI

2.5.12 CM3

- **Reference:**³³[21]
- **Family:** HTML
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** Causality-masked LM
- **Extension:** This is somewhat similar to HTML in its use of structured training data. However, it is a different architecture and uses causal masking
- **Application:** Multimodal language model with the ability to do structured prompting
- **Date (of first known publication):** 01/2022
- **Num. Params:**13B (largest)
- **Corpus:** CC-News, English Wikipedia
- **Lab:** Facebook

³¹<https://arxiv.org/abs/2203.15556>

³²https://huggingface.co/docs/transformers/model_doc/clip

³³<https://arxiv.org/abs/2201.07520>

2.5.13 CTRL

- **Reference:**³⁴[22]
- **Family:**
- **Pretraining Architecture:** Decoder
- **Pretraining Task:**
- **Extension:** model can generate text conditioned on control codes that specify domain, style, topics, dates, entities, relationships between entities, plot points, and task-related behavior
- **Application:** Controllable text generation
- **Date (of first known publication):** xx
- **Num. Params:**1.63B
- **Corpus:** 140 GB of text including: Wikipedia (En, De, Es, Fr), Project Gutenberg, 45 subreddits, OpenWeb-Text2, Amazon Reviews, Europarl and UN data from WMT, question-answer pairs from ELI5, and the MRQA shared task3, which includes the Stanford Question Answering Dataset, NewsQA, TriviaQA, SearchQA, HotpotQA , and Natural Questions
- **Lab:** Salesforce

2.5.14 DALL-E

- **Reference:**³⁵[23]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** Caption prediction
- **Extension:** A differential variational auto-encoder is used to learn the visual codebook. The transformer is a variation of GPT-3
- **Application:** Text to image
- **Date (of first known publication):** 01/2021
- **Num. Params:**12B
- **Corpus:** 250 million text-images pairs from the internet
- **Lab:** OpenAI

2.5.15 DALL-E 2

- **Reference:**³⁶[24]
- **Family:** CLIP, GLIDE
- **Pretraining Architecture:** Encoder/Decoder
- **Pretraining Task:** Caption prediction
- **Extension:** Combines CLIP encoder and Diffusion decoder similar to GLIDE
- **Application:** Text to image
- **Date (of first known publication):** 04/2022
- **Num. Params:**3.5B
- **Corpus:** Combination of the DALL-E and CLIP datasets
- **Lab:** OpenAI

³⁴https://huggingface.co/docs/transformers/model_doc/ctrl

³⁵<https://openai.com/blog/dall-e/>

³⁶<https://openai.com/dall-e-2/>

2.5.16 Decision Transformers

- **Reference:**³⁷[25]
- **Family:** GPT, Control Transformers” (not per se a family, but grouping here those transformers that try to model more general control, RL-like, tasks)
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** Next action prediction
- **Extension:** Decision transformers use a GPT architecture and extend it by encoding trajectories in a way that they can be learned by an auto-regressive task
- **Application:** General RL (reinforcement learning tasks)
- **Date (of first known publication):** 06/2021
- **Num. Params:** Same as GPT
- **Corpus:** Different corpus for different experiments
- **Lab:** Google/UC Berkeley/Facebook

2.5.17 DialoGPT

- **Reference:**³⁸[26]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** GPT-2 architecture trained on dialog data
- **Application:** Text generation in dialog settings
- **Date (of first known publication):** 10/2019
- **Num. Params:** 1.5B
- **Corpus:** 140M Reddit conversations
- **Lab:** Microsoft

2.5.18 DistilBERT

- **Reference:**³⁹[27]
- **Family:** BERT
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** MLM/NSP
- **Extension:** Compressed version of BERT using distillation, which is much more efficient given the same number of parameters
- **Application:** Same as BERT
- **Date (of first known publication):** 10/2019
- **Num. Params:** 66M
- **Corpus:** Same as BERT
- **Lab:** Huggingface

³⁷<https://arxiv.org/abs/2106.01345>

³⁸https://huggingface.co/docs/transformers/model_doc/dialogpt

³⁹https://huggingface.co/docs/transformers/model_doc/distilbert

2.5.19 DQ-BART

- **Reference:**⁴⁰[28]
- **Family:** BART
- **Pretraining Architecture:** Encoder/Decoder
- **Pretraining Task:** DAE
- **Extension:** Adds quantization and distillation to a BART model to improve performance and model size
- **Application:** Text generation and understanding
- **Date (of first known publication):** 03/2022
- **Num. Params:** Up to 30x reduction in parameters compared to standard BART
- **Corpus:** CNN/DM, XSUM, ELI5, WMT16 En-Ro (1M tokens)
- **Lab:** Amazon

2.5.20 ELECTRA

- **Reference:**⁴¹[29]
- **Family:**
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** RTD
- **Extension:** Same as BERT
- **Application:** 03/2020
- **Date (of first known publication):** xx
- **Num. Params:** Base = 110M, Large = 330M
- **Corpus:** Same as BERT except for Large with is same as XLNet
- **Lab:** Stanford/Google

2.5.21 ERNIE

- **Reference:**⁴²[30]
- **Family:** BERT
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** MLM
- **Extension:** Uses BERT for Encoder architecture, but stacks and aggregates two of them for text and entities. This architecture could be understood as BERT for text + knowledge graphs
- **Application:** Knowledge intensive related tasks that might benefit from knowledge graphs or entities such as entity recognition
- **Date (of first known publication):** 05/2019
- **Num. Params:** 114M
- **Corpus:** English Wikipedia + Wikidata for entities (note that they initialize model to original BERT parameter values)
- **Lab:** Various Chinese institutions

⁴⁰<https://arxiv.org/abs/2203.11239>

⁴¹https://huggingface.co/docs/transformers/model_doc/electra

⁴²<https://arxiv.org/abs/1905.07129>

2.5.22 Flamingo

- **Reference:**⁴³[31]
- **Family:** Chinchilla
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** Log likelihood of text given some visual input
- **Extension:** It uses a frozen textual language model (like Chinchilla) conditioned on the visual representation, which is encoded from a Normalizer-Free ResNet
- **Application:** Text to image
- **Date (of first known publication):** 04/2022
- **Num. Params:**80B (largest)
- **Corpus:** MultiModal MassiveWeb (M3W): 185 million images and 182 GB text + a number of text paired with image datasets: ALIGN + LTIP (Long Text & Image Pairs) = 312 million images, and VTP (Video & Text Pairs) = 27 million short videos (approximately 22 seconds on average)
- **Lab:** Deepmind

2.5.23 Gato

- **Reference:**⁴⁴[32]
- **Family:** “Control Transformers” (not per se a family, but grouping here those transformers that try to model more general control, RL-like, tasks)
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** MLM (where tokens are either text or agent actions)
- **Extension:** The standard decoder-only transformer architecture is preceded by an embedding layer that can embed text and images, plus add position encodings to add spatial information when applicable.
- **Application:** Gato presents a generalizable agent that can be used beyond text to tasks such as playing Atari or controlling a robot arm.
- **Date (of first known publication):** 05/2022
- **Num. Params:**1.2B
- **Corpus:** 1.5T tokens including standard text (e.g. MassiveText), vision (e.g. ALIGN), and simulation environments (e.g. ALE Atari, or RGB Stacking Real Robot)
- **Lab:** Deepmind

2.5.24 GLaM

- **Reference:**⁴⁵[33]
- **Family:** Transformer
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** GLaM introduces a Mixture of 64 Experts to increase parameter count and generalization properties in a somewhat standard decoder-only. Transformer architecture. Only two experts get activated at a time per token, which makes the model also more efficient in training and inference.
- **Application:** General language modeling
- **Date (of first known publication):** 12/2021
- **Num. Params:**1.2T across 64 experts, but only 96B get activated for inference
- **Corpus:** 1.6T tokens including web pages filtered by Wikipedia and books for quality
- **Lab:** Google

⁴³<https://arxiv.org/abs/2204.14198>

⁴⁴<https://www.deepmind.com/publications/a-generalist-agent>

⁴⁵<https://ai.googleblog.com/2021/12/more-efficient-in-context-learning-with.html>

2.5.25 GLIDE

- **Reference:**⁴⁶[34]
- **Family:** Diffusion models
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** Caption prediction
- **Extension:** GLIDE can be seen as an extension of the ADM (Ablated Diffusion Model) by the same authors. However, ADM is not per se a transformer architecture although it does resemble one in some of the configurations the authors use. Given that ADM is by the same authors and was quickly followed up by GLIDE, I think it is fair to consider GLIDE as the first of its kind.
- **Application:** Text to image
- **Date (of first known publication):** 12/2021
- **Num. Params:**3.5B diffusion model (2.3B for visual encoding, 1.2B for textual) + 1.5B for model for upsampling
- **Corpus:** Same as DALL-E
- **Lab:** OpenAI

2.5.26 Global Context ViT

- **Reference:**⁴⁷[35]
- **Family:** ViT
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** Image classification
- **Extension:** hierarchical ViT architecture consisting of local and global self-attention modules
- **Application:** Image generation
- **Date (of first known publication):** 06/2022
- **Num. Params:**90M
- **Corpus:** Imagenet-1K and other task dependent datasets
- **Lab:** NVidia

2.5.27 Gopher

- **Reference:**⁴⁸[36]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Same as GPT-2 but use RSNorm instead of LayerNorm and relative positional encoding rather than absolute
- **Application:** Mostly Language Modeling and NLU, but also extensible like GPT
- **Date (of first known publication):** 12/2021
- **Num. Params:**280B
- **Corpus:** Massive Text (2.35 billion documents, or about 10.5 TB of text including Massive Web, Books, Github, News, C4, and Wikipedia.
- **Lab:** Deepmind

⁴⁶<https://arxiv.org/abs/2112.10741>

⁴⁷<https://arxiv.org/abs/2206.09959>

⁴⁸<https://www.deepmind.com/blog/language-modelling-at-scale-gopher-ethical-considerations-and-retrieval>

2.5.28 GopherCite

- **Reference:**⁴⁹[37]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** GopherCite is based on Gopher but adds a step using RLHP (Reinforcement Learning from Human Preferences) to learn whether not only a response is plausible but also supported
- **Application:** Dialog systems, Q&A, general language generation tasks
- **Date (of first known publication):** 03/2022
- **Num. Params:**280B
- **Corpus:** Same as Gopher plus specific dataset generated in the RLHP process
- **Lab:** Deepmind

2.5.29 GPT

- **Reference:**⁵⁰[38]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:**
- **Application:** Text generation, but adaptable to many other NLP tasks when fine tuned.
- **Date (of first known publication):** 06/2018
- **Num. Params:**117M
- **Corpus:** Unsupervised Pretraining on BookCorpus dataset. Supervised Finetuning on several task-specific datasets including SNLI, RACE, Quora. . .
- **Lab:** OpenAI

2.5.30 GPT-2

- **Reference:**⁵¹[39]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Minor extensions to the GPT architecture (e.g. layer normalization moved to the input of each sub-layer, or increased context size from 512 to 1024)
- **Application:** Text generation, but adaptable to many other NLP tasks when fine tuned.
- **Date (of first known publication):** 02/2019
- **Num. Params:**1.5B
- **Corpus:** 8 million web pages (40 GB). 10X GPT . WebText dataset is created by crawling all links at Reddit with at least 3 Karma points.
- **Lab:** OpenAI

⁴⁹<https://arxiv.org/abs/2203.11147>

⁵⁰https://huggingface.co/docs/transformers/model_doc/openai-gpt

⁵¹https://huggingface.co/docs/transformers/model_doc/gpt2

2.5.31 GPT-3

- **Reference:**⁵²[8]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Same as GPT-2 with the only addition of alternating dense and locally banded sparse attention patterns, inspired by the Sparse Transformer
- **Application:** Initially text generation, but has over time been used for a large range of applications in areas such as code generation, but also image and audio generation
- **Date (of first known publication):** 05/2020
- **Num. Params:** 175 B
- **Corpus:** 500B tokens including CommonCrawl (410B), WebText2 (19B), Books1 (12B), Books2 (55B), and Wikipedia (3B)
- **Lab:** OpenAI

2.5.32 GPT-3.5

- **Reference:**⁵³
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** The GPT3.5 series includes a number of models like Davinci-003. They are basically versions of the InstructGPT model. See [here](https://scale.com/blog/gpt-3-davinci-003-comparison) for details on the comparison of the performance to older GPT3 models.
- **Application:** Dialog and general language, but there is a code specific model too
- **Date (of first known publication):** 10/2022
- **Num. Params:** 175B
- **Corpus:** Same as InstructGPT
- **Lab:** OpenAI

2.5.33 InstructGPT

- **Reference:**⁵⁴[40]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** GPTInstruct starts off with a pretrained GPT3 model and adds reward modeling through reinforcement learning after a supervised finetuning
- **Application:** Knowledge-intensive dialog or language tasks
- **Date (of first known publication):** 01/2022
- **Num. Params:** Same as GPT3
- **Corpus:** Same as GPT3 for pretraining, but finetuned and optimized using labeler data and prompts
- **Lab:** OpenAI

⁵²<https://github.com/openai/gpt-3>

⁵³<https://beta.openai.com/docs/model-index-for-researchers>

⁵⁴<https://openai.com/blog/instruction-following/>

2.5.34 GPT-Neo

- **Reference:**⁵⁵
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Similar to GPT-2 but uses local attention in every other layer with a window size of 256 tokens
- **Application:** Text generation, but adaptable to many other NLP tasks when fine tuned
- **Date (of first known publication):** 03/2021
- **Num. Params:** 5B, 2.7B (XL)
- **Corpus:** Pile — 840 GB open source text dataset that combines 22 pre existing datasets
- **Lab:** EleutherAI

2.5.35 GPT-NeoX-20B

- **Reference:**⁵⁶[41]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Similar to GPT-3 with rotary encoders instead of positional, parallel attention and feed forward layers, different initialization, and all dense layers instead of alternate dense/sparse
- **Application:** same as GPT-3
- **Date (of first known publication):** 04/2022
- **Num. Params:**20B
- **Corpus:** Pile — 840 GB open source text dataset that combines 22 pre existing datasets
- **Lab:** EleutherAI

2.5.36 HTML

- **Reference:**⁵⁷[42]
- **Family:** BART
- **Pretraining Architecture:** Encoder/Decoder
- **Pretraining Task:** DAE
- **Extension:** As opposed to BART, they don't do sentence shuffling
- **Application:** General purpose language model that allows structured HTML prompting
- **Date (of first known publication):** 07/2021
- **Num. Params:**400M
- **Corpus:** 23TB of simplified HTML extracted from CommonCrawl
- **Lab:** Facebook

2.5.37 Imagen

- **Reference:**⁵⁸[43]
- **Family:** T5, CLIP, Diffusion models
- **Pretraining Architecture:** T5 (or CLIP or BERT) for frozen text encoder + U-net architecture for cascaded diffusion models for text to image

⁵⁵https://huggingface.co/docs/transformers/model_doc/gpt_neo

⁵⁶<https://arxiv.org/abs/2204.06745>

⁵⁷<https://arxiv.org/abs/2107.06955>

⁵⁸<https://imagen.research.google/>

- **Pretraining Task:** image/text pair prediction
- **Extension:** Imagen adds a few extensions to the U-net diffusion architecture (pooled embedding vector, cross attention over text embeddings, and Layer Normalizations)
- **Application:** Text to image
- **Date (of first known publication):** 06/2022
- **Num. Params:**2B
- **Corpus:** a combination of internal datasets, with 460M image-text pairs, and the publicly available Laion dataset, with 400M image-text pairs
- **Lab:** Google

2.5.38 Jurassic-1

- **Reference:**⁵⁹[44]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Very similar to GPT-3, but far more parameters and improved training efficiency mostly because of the improved tokenizer. Also, different ratio of depth to breadth
- **Application:** Similar to GPT-3
- **Date (of first known publication):** 09/2021
- **Num. Params:**178B (Jumbo), 7.5B (Large)
- **Corpus:** 300B tokens (same as GPT-3)
- **Lab:** AI21

2.5.39 LAMDA

- **Reference:**⁶⁰[45]
- **Family:** Transformer
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** LAMDA focuses on how to improve safety, quality, and groundedness using different fine-tuning strategies
- **Application:** General language modeling
- **Date (of first known publication):** 01/2022
- **Num. Params:**137B
- **Corpus:** 1.56T words from public dialog data and other public web documents
- **Lab:** Google

2.5.40 mBART

- **Reference:**⁶¹[46]
- **Family:** BART
- **Pretraining Architecture:** Encoder/Decoder
- **Pretraining Task:** DAE
- **Extension:**

⁵⁹https://uploads-ssl.webflow.com/60fd4503684b466578c0d307/61138924626a6981ee09caf6_jurassic_tech_paper.pdf

⁶⁰<https://ai.googleblog.com/2022/01/lamda-towards-safe-grounded-and-high.html>

⁶¹https://huggingface.co/docs/transformers/model_doc/mbart

- **Application:** Translation
- **Date (of first known publication):** 01/2020
- **Num. Params:** Same as BART
- **Corpus:**
- **Lab:** CC25 Corpus includes 25 monolingual corpuses in different languages. Largest corpuses are English (300 GB) and Russian (280GB)

2.5.41 Megatron

- **Reference:**⁶²[47]
- **Family:** GPT/BERT/T5
- **Pretraining Architecture:** Encoder or Decoder, depending on the base model
- **Pretraining Task:** Same as base model
- **Extension:** Megatron is a family of models that extend previously known architectures (namely GPT-2 and BERT originally, but also T5 more recently) by introducing model parallelism primitives. In the case of BERT, the authors also replace the next sentence prediction head with sentence order prediction and use whole word n-gram masking.
- **Application:** Same as base model
- **Date (of first known publication):** 03/2020
- **Num. Params:** 8.3B (GPT-like), 3.9B (BERT-like)
- **Corpus:** Original paper uses an aggregate dataset consisting of Wikipedia, CC-Stories, RealNews, and OpenWebtext
- **Lab:** NVidia

2.5.42 Minerva

- **Reference:**⁶³[48]
- **Family:** PaLM
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Extends PaLM by fine-tuning on the mathematical dataset
- **Application:** Mathematical reasoning
- **Date (of first known publication):** 06/2022
- **Num. Params:** 540B
- **Corpus:** Same as PaLM + 118GB dataset of scientific papers from the arXiv preprint server and web pages that contain mathematical expressions using LaTeX, MathJax, or other mathematical typesetting formats
- **Lab:** Google

2.5.43 MT-NLG (Megatron TuringNLG)

- **Reference:**⁶⁴[49]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Uses parallelization similar to Megatron to train a LM double the size of GPT-3
- **Application:** Language generation and others (similar to GPT-3)
- **Date (of first known publication):** 10/2021

⁶²<https://github.com/NVIDIA/Megatron-LM>

⁶³<https://ai.googleblog.com/2022/06/minerva-solving-quantitative-reasoning.html>

⁶⁴<https://developer.nvidia.com/blog/using-deepspeed-and-megatron-to-train-megatron-turing-nlg-530b-the-worlds-lar>

- **Num. Params:**530B
- **Corpus:** The Pile⁶⁵ (800GB dataset) + 2 Common Crawl snapshots
- **Lab:** NVidia

2.5.44 OPT

- **Reference:**⁶⁶[50]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Basically same architecture as GPT-3 but with some training improvements introduced in Megatron-LM
- **Application:** Same as GPT-3
- **Date (of first known publication):** 05/2022
- **Num. Params:** 175B (and other smaller versions)
- **Corpus:** 180B tokens = RoBERTa + the Pile + PushShift.io Reddit
- **Lab:** Facebook

2.5.45 PalM

- **Reference:**⁶⁷[51]
- **Family:** Transformer
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Palm uses a typical decoder-only transformer architecture, but adds quite a few extensions: SwiGLU activations, parallel layers, multi-query attention, RoPE embeddings, Shared Input-Output Embeddings, no biases, and a 256k SentencePiece vocabulary generated from the training data.
- **Application:** PalM is designed as a general purpose language model with applicability to hundreds of different language tasks
- **Date (of first known publication):** 04/2022
- **Num. Params:**540B
- **Corpus:** 780B tokens from filtered webpages, books, Wikipedia, news articles, source code, and social media conversations. Code includes 24 programming languages.
- **Lab:** Google

2.5.46 Pegasus

- **Reference:**⁶⁸[52]
- **Family:** Transformer
- **Pretraining Architecture:** Encoder/Decoder
- **Pretraining Task:** DAE (more concretely GSG) and MLM
- **Extension:** Extends vanilla Transformer by using a different pretraining task (GSG: Gap Sentence Generation) that is better suited for summarization
- **Application:** Summarization
- **Date (of first known publication):** 12/2019
- **Num. Params:** Base = 223M, Large = 568M
- **Corpus:** C4 (750GB) + HugeNews (3.8 TB)
- **Lab:** UCL/Google

⁶⁵<https://arxiv.org/abs/2101.00027>

⁶⁶<https://ai.facebook.com/blog/democratizing-access-to-large-scale-language-models-with-opt-175b/>

⁶⁷<https://ai.googleblog.com/2022/04/pathways-language-model-palm-scaling-to.html>

⁶⁸https://huggingface.co/docs/transformers/model_doc/pegasus

2.5.47 RoBERTa

- **Reference:**⁶⁹[53]
- **Family:** BERT
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** MLM (Dynamic)
- **Extension:** Extension of BERT with optimized training procedure and more data
- **Application:** Same as BERT
- **Date (of first known publication):** 07/2019
- **Num. Params:** 356M
- **Corpus:** Same as BERT + CC News + OpenWebText + Stories (33B Tokens)
- **Lab:** UW/Google

2.5.48 SeeKer

- **Reference:**⁷⁰[54]
- **Family:** GPT (but can extend any family)
- **Pretraining Architecture:** Encoder/decoder or decoder only, depending on the base model it’s extending
- **Pretraining Task:** Depends on the base model
- **Extension:** SeeKer is an extension that can be applied to any Transformer architecture by introducing “search”, “knowledge”, and “response” modules that are introduced during pretraining
- **Application:** Same as base models
- **Date (of first known publication):** 03/2022
- **Num. Params:** Depends on the base model
- **Corpus:** Same as base model
- **Lab:** Facebook

2.5.49 Sparrow

- **Reference:**⁷¹[55]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Starts from the Chinchilla 70B model but adds RLHF (Reinforcement Learning with Human Feedback). It also adds inline evidence a la GopherCite
- **Application:** Dialog agents and general language generation applications like Q&A
- **Date (of first known publication):** 09/2022
- **Num. Params:** 70B
- **Corpus:** Same as Chinchilla + interactive data gathering with human annotators during the RLHF process
- **Lab:** Deepmind

⁶⁹https://huggingface.co/docs/transformers/model_doc/roberta

⁷⁰<https://parl.ai/projects/seeker/>

⁷¹<https://arxiv.org/abs/2209.14375>

2.5.50 StableDiffusion

- **Reference:**⁷²[56]
- **Family:** Diffusion
- **Pretraining Architecture:** Encoder/Decoder
- **Pretraining Task:** Caption prediction
- **Extension:** Stable diffusion is basically the Latent Diffusion model developed by LMU Munich researchers + some learnings on conditional diffusion from DALL-e and Imagen
- **Application:** Text to image
- **Date (of first known publication):** 12/2021
- **Num. Params:** 890M (although there are different, smaller, variants)
- **Corpus:** LAION-5B, a publicly available dataset derived from Common Crawl
- **Lab:** LMU Munich + Stability.ai + Eleuther.ai

2.5.51 Swin Transformer

- **Reference:**⁷³[57]
- **Family:** ViT
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** Same as ViT
- **Extension:** Extends ViT by replacing the standard multi-head self attention (MSA) module by a module based on shifted windows (Swin) allowing ViT-like architectures to generalize to higher resolution images
- **Application:** Image (object detection, image classification..)
- **Date (of first known publication):** 03/2021
- **Num. Params:** 29M-197M
- **Corpus:** Imagenet and Imagenet-22k
- **Lab:** Facebook

2.5.52 Switch

- **Reference:**⁷⁴[58]
- **Family:** T5
- **Pretraining Architecture:** Encoder/Decoder
- **Pretraining Task:** DAE
- **Extension:** Goal to increase parameter count while keeping FLOP operations constant by using efficient routing of MoE (Mixture of Experts)
- **Application:** General language tasks (e.g. question answering)
- **Date (of first known publication):** 01/2021
- **Num. Params:** 1T
- **Corpus:** Colossal Clean Crawled Corpus
- **Lab:** Google

⁷²<https://huggingface.co/CompVis/stable-diffusion>

⁷³<https://github.com/microsoft/Swin-Transformer>

⁷⁴<https://arxiv.org/abs/2101.03961>

2.5.53 T5

- **Reference:**⁷⁵[59]
- **Family:**
- **Pretraining Architecture:** Encoder/Decoder
- **Pretraining Task:** DAE
- **Extension:** Same as original Transformer with some additions such as relative positional embeddings like Transformer XL
- **Application:** General language tasks including machine translation, question answering, abstractive summarization, and text classification
- **Date (of first known publication):** 10/2019
- **Num. Params:** 11 B (up to)
- **Corpus:** Colossal Clean Crawled Corpus (C4) — Cleaned up version of the Common Crawl dataset — 750 GB
- **Lab:** Google

2.5.54 Trajectory Transformers

- **Reference:**⁷⁶[60]
- **Family:** GPT, Control Transformers” (not per se a family, but grouping here those transformers that try to model more general control, RL-like, tasks)
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** predict most likely sequence
- **Extension:** Similarly to the Decision transformers, the main extension introduced by Trajectory Transformers is a way to encode a trajectory (state, actions, rewards)
- **Application:** General RL (reinforcement learning tasks)
- **Date (of first known publication):** 06/2021
- **Num. Params:** Smaller architecture than GPT
- **Corpus:** D4RL dataset and other RL datasets depending on the task at hand
- **Lab:** UC Berkeley

2.5.55 Transformer XL

- **Reference:**⁷⁷[61]
- **Family:**
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Relative positioned embeddings enable longer-context attention when compared to vanilla Transformer model
- **Application:** General language tasks
- **Date (of first known publication):** 01/2019
- **Num. Params:** 151M
- **Corpus:** Different training datasets depending on experiments, but baseline is Wikitext-103
- **Lab:** CMU/Google

⁷⁵https://huggingface.co/docs/transformers/model_doc/t5

⁷⁶<https://arxiv.org/abs/2106.02039>

⁷⁷https://huggingface.co/docs/transformers/model_doc/transfo-xl

2.5.56 Turing-NLG

- **Reference:**⁷⁸[62]
- **Family:** GPT
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** LM
- **Extension:** Optimized version of GPT2 with optimal hyperparameters and software/hardware platform to improve training
- **Application:** Same as GPT-2/3
- **Date (of first known publication):** 02/2020
- **Num. Params:** 17B originally, up to 530B more recently
- **Corpus:** Highest quality subset from The Pile + 2 CC snapshots (339B tokens)
- **Lab:** Microsoft

2.5.57 ViT

- **Reference:**⁷⁹[63]
- **Family:** BERT
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** Image classification
- **Extension:** Extension of BERT architecture to train on patches of images
- **Application:** Image classification
- **Date (of first known publication):** 10/2020
- **Num. Params:** 86M(Base) to 632M (Huge)
- **Corpus:** From standard Imagenet to JFT-300M (large inhouse dataset)
- **Lab:** Google

2.5.58 Wu Dao 2.0

- **Reference:**⁸⁰
- **Family:** GLM (General Language Model)
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** Autoregressive blank infilling
- **Extension:** Similar to GPT in that it uses a Decoder/autoregressive architecture but applies a different pretraining task proposed in the GLM family of models. Besides, Wu Dao uses a Fast Mixture of Experts <https://github.com/laekov/fastmoe> approach to scale training to trillions of parameters
- **Application:** Language and multimodal (particularly image)
- **Date (of first known publication):** 06/2021
- **Num. Params:** 1.75T
- **Corpus:** ?
- **Lab:** Beijing Academy of Artificial Intelligence

⁷⁸<https://www.microsoft.com/en-us/research/blog/turing-nlg-a-17-billion-parameter-language-model-by-microsoft/>

⁷⁹https://huggingface.co/docs/transformers/model_doc/vit

⁸⁰https://en.wikipedia.org/wiki/Wu_Dao

2.5.59 XLM-RoBERTa

- **Reference:**⁸¹[64]
- **Family:** RoBERTa
- **Pretraining Architecture:** Encoder
- **Pretraining Task:** MLM (Dynamic)
- **Extension:** An extension of RoBERTa that introduces small parameter tuning insights in the context of multilingual applications
- **Application:** Translation and other cross-lingual language tasks
- **Date (of first known publication):** 10/2019
- **Num. Params:** Base = 270M, Large = 550M
- **Corpus:** Cleaned Common Crawl in 100 languages
- **Lab:** Facebook

2.5.60 XLNet

- **Reference:**⁸²[65]
- **Family:** Transformer XL
- **Pretraining Architecture:** Decoder
- **Pretraining Task:** PLM
- **Extension:** This model basically adapts Transformer XL architecture to permutation-based LM
- **Application:** General language tasks
- **Date (of first known publication):** 05/2019
- **Num. Params:** Base=117M, Large=360M
- **Corpus:** Same as BERT + Giga5 (16GB text) + and aggressively filtered ClueWeb 2012-B (19GB), Common Crawl (110 GB)
- **Lab:** CMU/Google

3 Further reading

Most of the following references have already been mentioned in the post. However, it is worth listing them here in case you need more details:

- The Huggingface Transformers documentation⁸³ and course is extremely good and comprehensive. I have used myself in this post, and I can't recommend enough as a natural follow up to what you will find here.
- A survey of transformers⁸⁴[66] includes a 40 page long survey with over 170 references and a full blown taxonomy.
- Pre-trained Models for Natural Language Processing: A Survey[10] is also a very comprehensive survey that includes many of the pretrained models with a particular focus on NLP

References

- [1] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [2] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.

⁸¹https://huggingface.co/docs/transformers/model_doc/xlm-roberta

⁸²https://huggingface.co/docs/transformers/model_doc/xlnet

⁸³<https://huggingface.co/course/chapter1/1?fw=pt>

⁸⁴<https://arxiv.org/abs/2106.04554>

QIU XP, et al. *Pre-trained Models for Natural Language Processing: A Survey* March (2020)

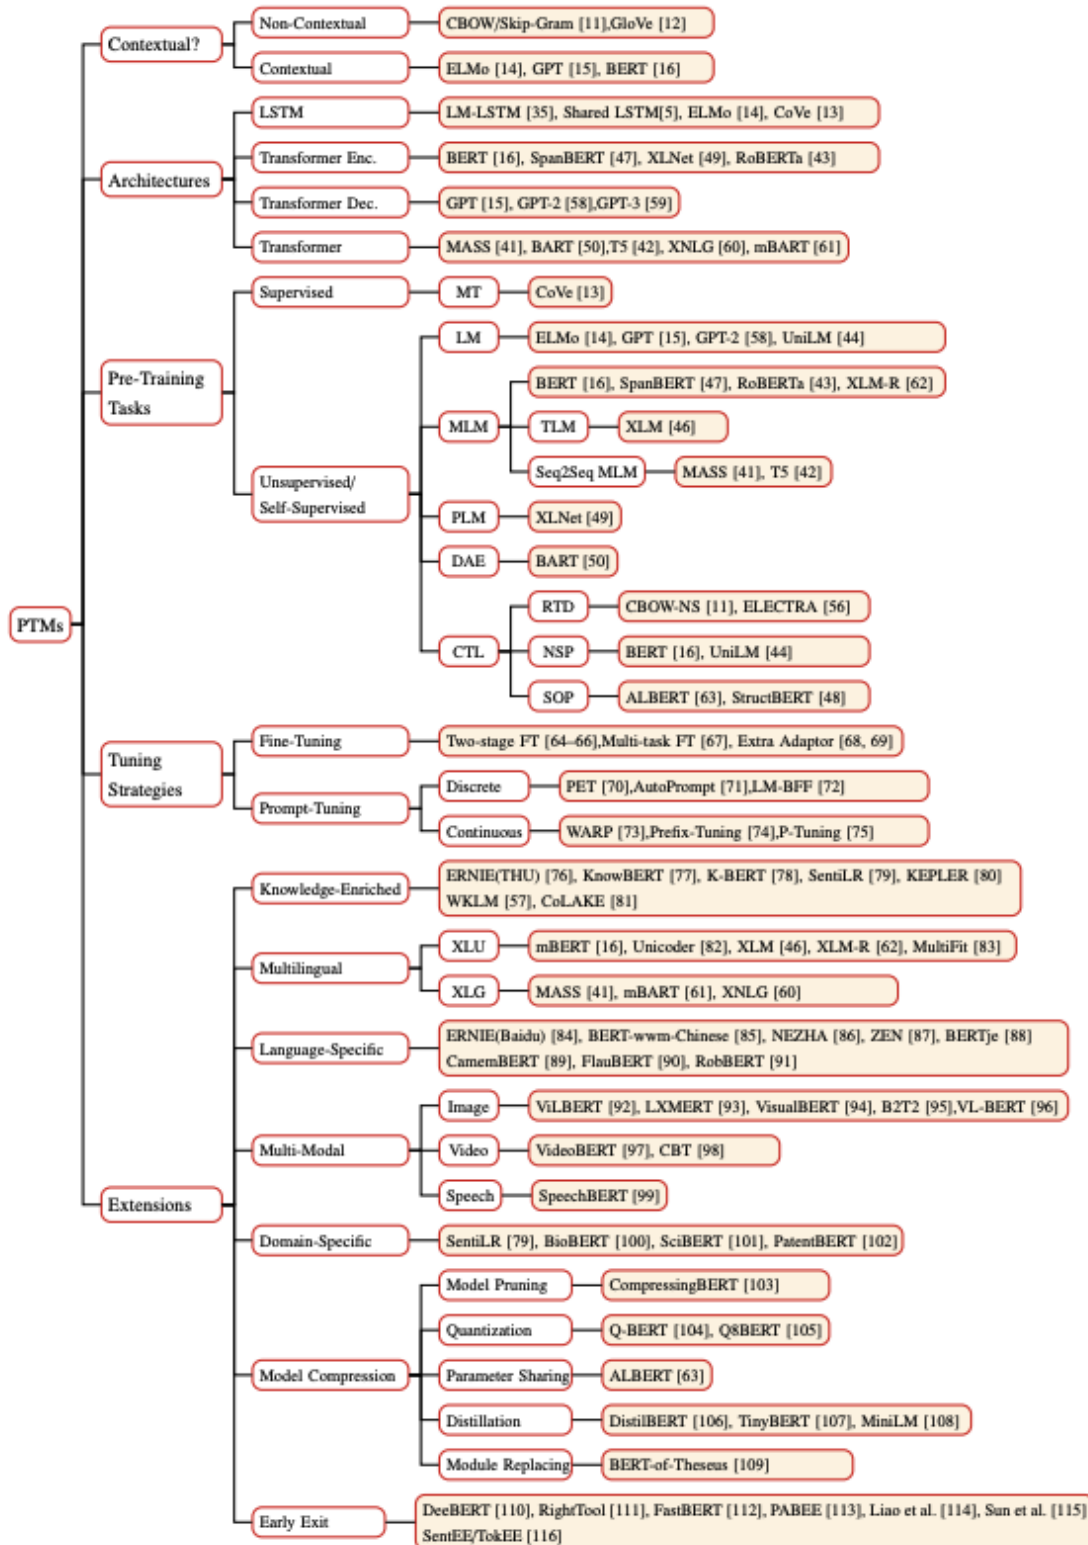


Figure 3: Taxonomy of PTMs with Representative Examples

Figure 9: Transformer taxonomy from [10]

- [3] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [4] Tomas Mikolov, Martin Karafiát, Lukáš Burget, Jan Cernocký, and Sanjeev Khudanpur. Recurrent neural network based language model. In *Interspeech*, volume 2, pages 1045–1048. Makuhari, 2010.
- [5] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM computing surveys (CSUR)*, 54(10s):1–41, 2022.
- [6] David Noever, Matt Ciolino, and Josh Kalin. The chess transformer: Mastering play using generative language models. *arXiv preprint arXiv:2008.04057*, 2020.
- [7] Kimia Noorbakhsh, Modar Sulaiman, Mahdi Sharifi, Kallol Roy, and Pooyan Jamshidi. Pretrained language models are symbolic mathematics solvers too! *arXiv preprint arXiv:2110.03501*, 2021.
- [8] 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. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [9] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Yingxia Shao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796*, 2022.
- [10] Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, 63(10):1872–1897, 2020.
- [11] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.
- [12] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- [13] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- [14] Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, et al. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*, 2021.
- [15] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [17] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. *Advances in neural information processing systems*, 33:17283–17297, 2020.
- [18] Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, et al. Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage. *arXiv preprint arXiv:2208.03188*, 2022.
- [19] 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.
- [20] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [21] Armen Aghajanyan, Bernie Huang, Candace Ross, Vladimir Karpukhin, Hu Xu, Naman Goyal, Dmytro Okhonko, Mandar Joshi, Gargi Ghosh, Mike Lewis, et al. Cm3: A causal masked multimodal model of the internet. *arXiv preprint arXiv:2201.07520*, 2022.
- [22] Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*, 2019.
- [23] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021.

- [24] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [25] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- [26] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. Dialogpt: Large-scale generative pre-training for conversational response generation. *arXiv preprint arXiv:1911.00536*, 2019.
- [27] 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.
- [28] Zheng Li, Zijian Wang, Ming Tan, Ramesh Nallapati, Parminder Bhatia, Andrew Arnold, Bing Xiang, and Dan Roth. Dq-bart: Efficient sequence-to-sequence model via joint distillation and quantization. *arXiv preprint arXiv:2203.11239*, 2022.
- [29] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*, 2020.
- [30] Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. Ernie: Enhanced language representation with informative entities. *arXiv preprint arXiv:1905.07129*, 2019.
- [31] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022.
- [32] Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gomez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Gimenez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, et al. A generalist agent. *arXiv preprint arXiv:2205.06175*, 2022.
- [33] Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, pages 5547–5569. PMLR, 2022.
- [34] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.
- [35] Ali Hatamizadeh, Hongxu Yin, Jan Kautz, and Pavlo Molchanov. Global context vision transformers. *arXiv preprint arXiv:2206.09959*, 2022.
- [36] 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. *arXiv preprint arXiv:2112.11446*, 2021.
- [37] Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick, Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, et al. Teaching language models to support answers with verified quotes. *arXiv preprint arXiv:2203.11147*, 2022.
- [38] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- [39] 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.
- [40] 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. *arXiv preprint arXiv:2203.02155*, 2022.
- [41] Sid 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. *arXiv preprint arXiv:2204.06745*, 2022.
- [42] Armen Aghajanyan, Dmytro Okhonko, Mike Lewis, Mandar Joshi, Hu Xu, Gargi Ghosh, and Luke Zettlemoyer. Htlm: Hyper-text pre-training and prompting of language models. *arXiv preprint arXiv:2107.06955*, 2021.
- [43] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022.

- [44] Opher Lieber, Or Sharir, Barak Lenz, and Yoav Shoham. Jurassic-1: Technical details and evaluation. *White Paper: AI21 Labs*, 1, 2021.
- [45] 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. *arXiv preprint arXiv:2201.08239*, 2022.
- [46] Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742, 2020.
- [47] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- [48] Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative reasoning problems with language models. *arXiv preprint arXiv:2206.14858*, 2022.
- [49] Shaden Smith, Mostofa Patwary, Brandon Norrick, 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. *arXiv preprint arXiv:2201.11990*, 2022.
- [50] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- [51] 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. *arXiv preprint arXiv:2204.02311*, 2022.
- [52] Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR, 2020.
- [53] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [54] Kurt Shuster, Mojtaba Komeili, Leonard Adolphs, Stephen Roller, Arthur Szlam, and Jason Weston. Language models that seek for knowledge: Modular search & generation for dialogue and prompt completion. *arXiv preprint arXiv:2203.13224*, 2022.
- [55] Amelia Glaese, Nat McAleese, Maja Trębacz, John Aslanides, Vlad Firoiu, Timo Ewalds, Maribeth Rauh, Laura Weidinger, Martin Chadwick, Phoebe Thacker, et al. Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*, 2022.
- [56] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [57] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- [58] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *J. Mach. Learn. Res.*, 23:1–40, 2021.
- [59] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- [60] Michael Janner, Qiyang Li, and Sergey Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34:1273–1286, 2021.
- [61] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.
- [62] Corby Rosset. Turing-nlg: A 17-billion-parameter language model by microsoft. *Microsoft Blog*, 1(2), 2020.

- [63] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [64] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*, 2019.
- [65] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.
- [66] Tianyang Lin, Yuxin Wang, Xiangyang Liu, and Xipeng Qiu. A survey of transformers. *AI Open*, 2022.