

Scaling Laws for Acoustic Models

Jasha Droppo, Oguz Elibol

Amazon Alexa

{drojasha, oelibol}@amazon.com

Abstract

There is a recent trend in machine learning to increase model quality by growing models to sizes previously thought to be unreasonable. Recent work has shown that autoregressive generative models with cross-entropy objective functions exhibit smooth power-law relationships, or scaling laws, that predict model quality from model size, training set size, and the available compute budget. These scaling laws allow one to choose nearly optimal hyper-parameters given constraints on available training data, model parameter count, or training computation budget. In this paper, we demonstrate that acoustic models trained with an auto-predictive coding loss behave as if they are subject to similar scaling laws. We extend previous work to jointly predict loss due to model size, to training set size, and to the inherent “irreducible loss” of the task. We find that the scaling laws accurately match model performance over two orders of magnitude in both model size and training set size, and make predictions about the limits of model performance.

Index Terms: speech recognition, acoustic modeling

1. Introduction

In the field of language modeling, it has been shown that generative models have predictable and favorable scaling properties with respect to training set size, model size, and available compute budget[1, 2]. It has also been shown that, as the models become larger and more accurate, less task-specific fine-tuning is necessary to achieve state-of-the-art results[3].

Consequently, the field of natural language processing is in the middle of an “arms race” to build the largest and most powerful models. It started with smaller models like GPT[4] and Bert[5], which had hundreds of millions of parameters. Today, state-of-the-art models like SWITCH-C[6] have more than 1.5 trillion parameters.

Meanwhile, in the field of acoustic modeling of speech, there has been growing interest in improving automatic speech recognition by using bulk untranscribed audio for unsupervised pre-training. Successful techniques use either autoregressive predictive coding (APC) [7, 8, 9] or contrastive predictive coding (CPC) [10, 11, 12] to pre-train an acoustic model from a large corpus of unlabeled speech. Once the pre-training is complete, the models are used to generate features for downstream modeling tasks, such as automatic speech recognition.

If acoustic models trained with predictive coding exhibit similar scaling behavior as the autoregressive generative models studied in [2], then we can use this analysis to discover best practices for scaling and training acoustic models that are much larger and more accurate than today’s state-of-the-art. By accurately predicting the amount of data, number of parameters, or length of compute, we eliminate these factors from potentially expensive hyperparameter tuning on the models we build. By understanding whether our model’s performance (development-set loss) is being limited by a lack of data or by the model size,

we can make intelligent decisions about how to improve them.

This paper tests this premise and develops a generalization of previous work that accounts for the interplay between limited training data, limited model parameters, and the development-set loss of a converged model.

We find that these acoustic models do exhibit smooth power-law relationships with respect to the training constraints, that training to convergence is inefficient, that data requirements scale sub-linearly with model size, and that larger models are more efficient learners as was shown for language models in [1].

This paper is organized as follows. In Section 2, we describe the acoustic models that we test, our training procedure, and our parameter scaling strategy. In Section 3, we review previously published results on model scaling, and fit them to measurements of fully trained acoustic models. Section 4 highlights the advantage of not training to convergence, and or conclusions are presented in Section 5.

2. Experimental Setup

The two models evaluated in this work share the same basic auto-predictive coding (APC) design, shown in Figure 1. Both models share the same encoder and prediction structures, with different context modules.




Figure 1: APC model structure.

In this work, the encoder module is a dense affine embedding layer that projects each input feature vector into the layer width of the context module, without additional frame-stacking or decimation.

The context module is a sequence-to-sequence model that converts the encoded input sequence into a sequence of context vectors. We experiment with two different designs for the context model: the LSTM and the Transformer. To maintain causality, the LSTM are uni-directional and the Transformer uses masking to prevent the network from using future frames.

The prediction module uses the output of the context module to predict and reconstruct ten consecutive input feature vectors. It does this with ten separate prediction networks. Each one consists of a linear projection to 512 units, a ReLU nonlinearity, and a linear projection to the input data dimension. The model is trained to minimize the sum of the L1 loss between each of the ten prediction networks and the ten contiguous input feature vectors. Following the design in [13], the first of

these feature vector targets is identical to the context vector’s latest input.

The difference between this APC design and a more traditional autoencoder is that in an autoencoder, the network distills and recreates the input, and in an APC design, the network predicts the future from the past.

The LSTM context module is a stack of uni-directional LSTM layers. Each of the L layers incorporates layer normalization, and every layer but the first employs a residual structure. The complexity of this context module, both in terms of number of parameters N and the number of multiplications per frame of input data performed during inference M are shown in Table 1. The values are a function of u , the number of units in the LSTM.

Table 1: *LSTM layer parameter count N and inference multiplication count M .*

Computation	N	M
Input, Output, Forget Gates	$3u(2u + 1)$	$3u(2u + 1)$
Cell Update	$u(2u + 1)$	$u(2u + 1)$
Layer Norm	$2u$	u
Total	$4u(2u + 1)$	$u(8u + 5)$

The Transformer context model is a stack of Transformer encoder layers. Each layer comprises layer norm, masked multi-headed self-attention, a second layer norm, and a then a point-wise feed-forward network. Both the self-attention and the feed-forward network employ a residual connection. The number of parameters N and the number of multiplications per input frame during inference M are given in Table 2. Unlike the LSTM, the computational complexity of the transformer-based context module is dependent on both the number of units in the model u and the attention component’s context length n_c .

Table 2: *Transformer encoder layer parameter count N and forward-pass multiplication count M .*

Computation	N	M
Two Layer Norms	$4u$	$2u$
Att. Embedding	$3u(u + 1)$	$3u(u + 1)$
Self Attention	0	$2n_{\text{ctx}}u$
Att. Projection	$u(u + 1)$	$u(u + 1)$
Feed-Forward	$u(8u + 5)$	$u(8u + 5)$
Total	$u(12u + 13)$	$u(12u + 2n_{\text{ctx}} + 11)$

The LSTM exhibits $O(u^2)$ complexity, regardless of the input data. The Transformer exhibits $O(u^2)$ complexity for short context lengths, and $O(u)$ for long context lengths. Specifically, when $n_c < 6u$, the term $12u^2$ dominates the Transformer multiplication count, and when $n_c > 6u$, the term $2un_c$ becomes more important.

2.1. Training

To train the models, we used an Adam optimizer with a warm hold decay learning rate scheduler. The learning rate was ramped up from 1×10^{-4} to 2×10^{-4} over 3,000 steps, held there until 50,000 steps were complete, and then exponentially decayed to hit 1×10^{-5} at 150,000 steps. Each step consisted of a 64-example mini-batch of speech sequences, with a feature dimension of 64 and roughly bucketed sequence lengths. Loss

function values were computed on a held-out development set every 25,000 steps during training.

All acoustic data used in this paper was drawn from a 23 thousand hour corpus of untranscribed, de-identified, far-field, English voice command and voice query speech collected from home environments. This data is presented to the network as a series of log-Mel frequency filterbank feature vectors, at a rate of 100 vectors per second of audio. Although this data is not publicly available, the authors believe that the phenomena described in this paper should apply to any similar set of speech recordings.

2.2. Model Scaling Strategy

Although one could scale model size by independently choosing the number of layers n_{layer} and number of units u in the context module, we chose to scale the models by coupling the two hyper-parameters with a fixed aspect ratio.

Preliminary experimentation indicated that for a given number of parameters, the LSTM model performs best with a constant aspect ratio of $u = 256n_{\text{layer}}$. Similarly, the transformer model performs best with an aspect ratio of $u = 64n_{\text{layer}}$. These ratios are used exclusively throughout this paper.

3. Scaling of Converged Models

This section describes the relationship between the number of model parameters N , the amount of training data D , and the model loss of a fully trained model $L(N, D)$.

3.1. Limited Data or Limited Model Parameters

It was empirically demonstrated in [1] and [2] that, for several generative modeling tasks, the loss of autoregressive Transformer models behave predictably when their performance is limited by either the number of model parameters N or the amount of training data D . Specifically, loss is predictable according to a power-law plus constant formulation as in Eq. 1.

$$L(D) = L_\infty + \left(\frac{D_C}{D}\right)^{\alpha_D}, \quad L(N) = L_\infty + \left(\frac{N_C}{N}\right)^{\alpha_N} \quad (1)$$

The parameter D_C represents the critical value of D , where the contribution of limited data to the loss function is equal to 1.0. The parameter α_D controls the relationship between increased data and reduced loss. The parameters N_C and α_N perform similar functions in relating N to $L(N)$. The constant L_∞ is the “irreducible loss” that is independent of either D or N .

To demonstrate the effect of limited D on loss, we trained several models that were likely to be performance-limited by D . Five instances of 11-layer Transformer models were on differently sized subsets of the training data. The training data sizes D were chosen to be exponentially spaced from 180 hours to 11,500 hours. The loss values for these models are marked with dots in Figure 2. The parametric $L(D)$ from Eq. 1 is shown with constants $D_C = 7.350 \times 10^{-23}$ hours, $\alpha_D = 0.01946$, and $L_\infty = 0.316$.

The small value for D_C indicates that almost no data is necessary to bring the objective function below 1.0. This makes sense for our data, where if the model did nothing but predict the mean spectral value, the objective function would already be less than 1.0.

To demonstrate the effect of limited N on loss, we trained several models that were likely to be performance-limited by N . Five models were trained on 23 thousand hours of data, with




Figure 2: When model performance is limited only by available training data the relationship between training data and loss is linear on a log-log plot.

model size restricted to 2, 3, 5, 7, and 11 layers. The resulting loss values are marked with dots in Figure 3, together with the parametric $L(N)$ with $N_C = 9.410 \times 10^{-25}$ parameters, $\alpha_N = 0.01601$, and $L_\infty = 0.316$. Note that the value used for irreducible loss here in $L(N)$ is identical to the one used for $L(D)$. The small value for N_C indicates that loss values less than 1.0 are achievable with zero parameters in the context network.




Figure 3: When model performance is limited by model size, the relationship between parameter count and loss is linear on a log-log plot.

Together, these relationships describe how to improve model performance when it is being limited either by D or N .

At less than about one thousand hours of training data, the 11-layer transformer is data-limited. With the estimated value of α_D , a five percent reduction in loss requires a 14.0-fold increase in training data D .

At less than about two million parameters, the models trained on a large amount of data are parameter-limited. With the estimated value of α_N , a five percent reduction in loss requires a 24.6-fold increase in model size N .

The ratio of α_N to α_D describes how quickly D should increase to support an increase in N . If D scales more slowly, then the model performance will eventually be limited by D and the increase in N will be wasted. For our data, $\frac{\alpha_N}{\alpha_D} < 1$, which means data requirements scale sub-linearly with model parameter count. Specifically, with every doubling of N , D should increase by a factor of 1.77.

If our only constraint is model size N , we should choose D large enough so that its contribution to the loss is at least one order of magnitude smaller than $L(N)$. For our data, this implies a lower limit of $D > 0.0436N^{0.8230}$.

3.2. Limited Data and Limited Model Size

In [1], the authors demonstrated that when limited by both data and model size, but with unlimited computational power, the two previous relationships can be combined into a single scaling law, Eq. 2. Unfortunately, this equation was developed for language models where L_∞ can be ignored. The authors of [2] did not present an analogous relationship that incorporates the concept of irreducible loss.

$$L(N, D) = \left[\left(\frac{N_C}{N} \right)^{\frac{\alpha_N}{\alpha_D}} + \left(\frac{D_C}{D} \right)^{\alpha_D} \right]^{\alpha_D} \quad (2)$$

To adapt this relationship to account for irreducible loss, we introduce the following generalization, which is equivalent to Eq. 2 with the addition of an “irreducible loss” term and a slight change of variables.

$$L(N, D) = \left[(L_\infty)^{\frac{1}{\alpha}} + \left(\frac{N_C}{N} \right)^{\frac{\alpha_N}{\alpha}} + \left(\frac{D_C}{D} \right)^{\frac{\alpha_D}{\alpha}} \right]^{\alpha} \quad (3)$$

To test this scaling law’s ability to describe model performance, we trained twenty-one Transformer based models. We used a variety of n_{layer} from 2 to 11, a fixed aspect ratio of $u = 64n_{\text{layer}}$, and training data set sizes from 134 hours to 23 thousand hours of speech. Each model was trained until the development set loss began to increase, or until 1.5 million parameter updates were complete.




Figure 4: The converged loss of several Transformer models with 2 through 11 layers, plotted as a function of training data D . Each marker represents a different experiment across n_{layer} and D . Lines are drawn for each model size representing the prediction of Eq. 3. The diagonal dashed line represents a performance frontier imposed by limited data, and the horizontal dotted line represents the irreducible loss.

The loss values for these models are represented by markers in Figure 4. The performance barrier due to limited training set size is shown as a diagonal dashed line. Regardless of model size, we do not expect any model will be able to surpass this limit. The performance barrier due to irreducible loss is represented as a horizontal dotted line. If Eq. 3 is correct, no model can be better than this limit, regardless of model size or training data set size.

The smooth curves in Figure 4 are computed using Eq. 3 with the same constants reported above with the addition of $\alpha = 0.01363$. It is clear that the model of Eq. 3 is a good match to our data.

4. Efficient Scaling Properties

Both [1] and [2] showed that for a broad class of models, training to convergence is a computationally inefficient way to maximize model performance. Instead, they demonstrated that for a given loss, the most efficient way to train a model to that loss is to train a larger model using fewer model update steps. We repeat this analysis for acoustic models.

Figure 5 shows the evolution of development-set loss during training for both LSTM and Transformer based models. Each curve represents an experiment with a different N . Each experiment is stopped after 300,000 model updates, which represents less than one full pass over the training data. For each point, the x-coordinate is the approximate number of multiplications and additions executed during training, following the conventions established in [1].

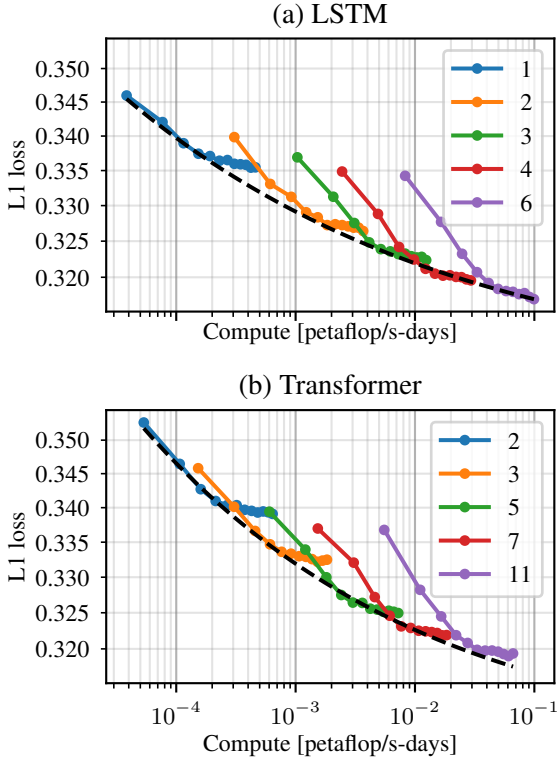


Figure 5: Development set loss for both LSTM and Transformer models for models with the indicated number of layers. The dashed line represents the computationally efficient frontier defined in Eq. 4.

Regardless of model size, no experiment surpasses the power-law plus constant fit shown as a dotted line. This line is the compute-efficient frontier, and has a similar form to Eq. 1.

$$L(C) = L_{\infty} + \left(\frac{C_C}{C}\right)^{\alpha_C} \quad (4)$$

When a model reaches $L(C)$, it means that a different model with enough capacity, but with fewer parameters, would need more computation and more data to reach the same loss value. Alternatively, a model with more parameters would need more computation and less data to reach the same loss value.

Where curves for two experiments meet, it is an indication that the same amount of compute can reach the given loss value

through two different methods. One can either use more parameters and fewer data, or use fewer parameters and more data.

The constant L_{∞} is 0.306 in both figures. This represents a shared asymptote between the LSTM and Transformer systems, which will never be surpassed, regardless of the computational or data budget. The fact that the same asymptote applies to both systems hints that irreducible loss is indeed a fundamental property of the data and not the model. Additionally, this constant is similar to the value found in Section 3.1. The authors suspect that the constants should be identical, but our precision in measuring it is limited.

The LSTM models exhibit a compute-efficient frontier with a slope of -0.167. A doubling of computation yields a 10.9% reduction in objective function. A halving of objective function would come with a 63.5 fold increase in computation.

The slope of the compute-efficient frontier for Transformer models is -0.197. When computation is increased by a factor of r , then the reducible loss will be changed by a factor of $r^{-0.197}$. At that rate, a doubling of computation yields a 12.7% reduction in objective function. A halving of objective function would come with a 33.7 fold increase in computation.

The difference in slope between the LSTM and Transformer experiments indicate that the Transformer architecture makes more efficient use of increased model parameters and increased training data. Although LSTM is superior to transformer at smaller model sizes, as the model size grows, and these trends continue, the transformer will eventually be more efficient.

Finally, the experimental data show that larger models learn more quickly from the same amount of data. Each of the points plotted in Figure 5 represent the consumption of an additional 25 thousand minibatches of training data. At the first point, second, or third, each model has processed the same data, but the larger models have achieved better accuracy on the held-out development set.

5. Conclusion

In this paper, we demonstrated that the effect of limited model parameters and limited training data on the quality of APC acoustic models follows a power-law given by Eq. 3. This extends previous results in language modeling and recurrent generative modeling. The relationship can be used to predict the amount of data needed to support a given model size, or to predict the proper model size for a given amount of data.

For our data, we have shown that the irreducible loss L_{∞} is the same, whether the model uses Transformer or LSTM components. Furthermore, under two different power-law relationships, the L_{∞} is similar. Together, this suggests that it is a fundamental property of the task and data, but not the model. Although speculative, authors believe that the irreducible loss derived from such scaling studies can translate into a measure of information content, or quality of the training data. One study that can be performed, beyond the scope of this work, is to add controlled amount of noise and study the effect on the irreducible loss. Future work should tie L_{∞} to the intrinsic difficulty in modeling a set of acoustic data, similar to perplexity measurements in language modeling tasks.

This paper has implicitly assumed that improvements in L1 loss are desirable. Future work will relate L1 loss to performance on a downstream acoustic modeling task.

6. References

- [1] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, “Scaling Laws for Neural Language Models,” *arXiv e-prints*, p. arXiv:2001.08361, Jan. 2020.
- [2] T. Henighan, J. Kaplan, M. Katz, M. Chen, C. Hesse, J. Jackson, H. Jun, T. B. Brown, P. Dhariwal, S. Gray, C. Hallacy, B. Mann, A. Radford, A. Ramesh, N. Ryder, D. M. Ziegler, J. Schulman, D. Amodei, and S. McCandlish, “Scaling Laws for Autoregressive Generative Modeling,” *arXiv e-prints*, p. arXiv:2010.14701, Oct. 2020.
- [3] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, “Language models are few-shot learners,” 2020.
- [4] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, “Improving language understanding with unsupervised learning,” OpenAI, Tech. Rep., 2018.
- [5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” *arXiv e-prints*, p. arXiv:1810.04805, Oct. 2018.
- [6] W. Fedus, B. Zoph, and N. Shazeer, “Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity,” *arXiv e-prints*, p. arXiv:2101.03961, Jan. 2021.
- [7] S. Ling and Y. Liu, “DeCoAR 2.0: Deep Contextualized Acoustic Representations with Vector Quantization,” *arXiv e-prints*, p. arXiv:2012.06659, Dec. 2020.
- [8] Y. Chung and J. Glass, “Generative pre-training for speech with autoregressive predictive coding,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 3497–3501.
- [9] Y.-A. Chung, H. Tang, and J. Glass, “Vector-quantized autoregressive predictive coding,” October 2020, pp. 3760–3764.
- [10] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” 2020.
- [11] A. van den Oord, Y. Li, and O. Vinyals, “Representation Learning with Contrastive Predictive Coding,” *arXiv e-prints*, p. arXiv:1807.03748, Jul. 2018.
- [12] Q. Xu, A. Baevski, T. Likhomanenko, P. Tomasello, A. Conneau, R. Collobert, G. Synnaeve, and M. Auli, “Self-training and Pre-training are Complementary for Speech Recognition,” *arXiv e-prints*, p. arXiv:2010.11430, Oct. 2020.
- [13] S. Ling, Y. Liu, J. Salazar, and K. Kirchhoff, “Deep contextualized acoustic representations for semi-supervised speech recognition,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6429–6433.