



Improving Joint Speech-Text Representations Without Alignment

Cal Peysers^{1,2}, Zhong Meng², Ke Hu², Rohit Prabhavalkar², Andrew Rosenberg², Tara N. Sainath²,
Michael Picheny¹, Kyunghyun Cho¹

¹Center for Data Science, New York University, New York City, USA

²Google Inc., U.S.A

cpeysers@google.com

Abstract

The last year has seen astonishing progress in text-prompted image generation premised on the idea of a cross-modal representation space in which the text and image domains are represented jointly. In ASR, this idea has found application as joint speech-text encoders that can scale to the capacities of very large parameter models by being trained on both unpaired speech and text. While these methods show promise, they have required special treatment of the sequence-length mismatch inherent in speech and text, either by up-sampling heuristics or an explicit alignment model. In this work, we offer evidence that joint speech-text encoders naturally achieve consistent representations across modalities by disregarding sequence length, and argue that consistency losses could forgive length differences and simply assume the best alignment. We show that such a loss improves downstream WER in both a large-parameter monolingual and multilingual system.

1. Introduction

The power of very large models trained on vast unsupervised corpora in a single modality has become increasingly clear. This has been demonstrated in the text domain where language models have achieved unprecedented zero-shot capabilities [1, 2], as well as in the audio domain, in which a single model has been shown to be adaptable to a surprisingly wide array of acoustic tasks [3, 4]. These successes have given rise to the question of how to apply these methods for problems involving two modalities, which historically have depended on manually paired data.

One very promising solution to this problem is to train a large encoder on both modalities, such that either modality may be provided as an unpaired example, but which learns to map paired examples to similar points in representation space. In the image/text domain, such a representation has proved achievable and capable of attaining state-of-the-art performance on many image and text comprehension tasks in a single model [5, 6].

In the audio/text domain, joint speech and text models have been utilized for a wide range of tasks [7, 8, 9]. In speech recognition, the past few years has seen a trend toward models with a joint speech and text encoder to allow pretraining on unpaired speech and text data [10, 11, 12, 13]. However, speech recognition presents the particular challenge of two sequence modalities, one of which (speech) is typically represented by a much longer sequence than the other (text). This complicates the task of representing both modalities in the same embedding space, since we cannot make a direct, frame-wise comparison of an encoder’s speech representation to its text representation.

This complication has largely been handled either by up-sampling or an explicit alignment model. Fixed up-sampling of the text inputs has been applied successfully for ASR in [13]

and SLU in [14], proving that an approximate alignment is sufficient for learning a joint representation. On the other hand, [15] addresses the problem with a separately-trained alignment model that aims for perfect alignment. In [12], it’s shown that such an alignment model permits the use of “consistency” regularization in which the encoder’s outputs on corresponding speech and text are compared frame-wise and pushed together in representation space. [12] goes on to show that “consistency” regularization yields a more closely joined representation space leading to better WER.

Consistency regularization itself follows naturally from the literature on generative models. Systems like autoencoders applied to augmented data (e.g. [16]) explicitly push representations of matched examples together, while contrastive systems like [17] do the same implicitly. The success of the same idea in speech using an explicit alignment begs the question of if the same can be done with an implicit alignment; that is, without knowing the particular alignment between speech and text.

In this paper, we ask if consistency regularization may be applied using the implicit alignments learned in upsampling systems like [13] to achieve the performance improvements seen with the explicit alignments in [12]. To this end, we develop an algorithm inspired by dynamic time warping [18] that finds the *best possible alignment* between an encoder’s representation of a paired speech and text example. We measure the quality of this *best alignment* in a system without an explicit alignment model and show that that it is not only learned during training but in fact improves at deeper layers of the network. Inspired by the improvements shown in [15] and [12], we then show that by changing the criteria of the consistency regularization to encourage consistency under *some* alignment, instead of a direct frame-wise comparison, we can achieve robust WER improvements against strong, semi-supervised baselines in both a monolingual and multilingual setting, all without any learned alignment model. Our results suggest that enforcing consistency in cross-modal representations can be done by simply forgiving misalignment.

The rest of this paper proceeds as follows. Section 2 specifies our setup for joint speech/text modeling and consistency regularization, and the details of our best-alignment algorithm. Section 3 specifies details of our data and training. Section 4 presents our analysis of the best alignment in an unregularized model and the results of optimizing that alignment with our consistency loss. We conclude in Section 5.

2. Methods

In this section we present our setup for semi-supervised ASR by joint speech/text modeling, for which we mostly follow [13].

We then present our proposed best-alignment algorithm and define a corresponding consistency loss inspired by [15].

2.1. Model Architecture

Figure 1 gives our model architecture. Essentially, we perform supervised ASR with streaming and non-streaming decoders, where the encoder is split into “audio-only”/“text-only” and “shared” components to permit text injection. The simultaneous ASR and text-injection tasks give rise to a joint representation in the shared encoder. Specifically, given a corpus of supervised examples $(x, y) \in \mathcal{S}$ and an unpaired text corpus $y \in \mathcal{U}$, our model contains the following components:

- E_a : The audio encoder, which embeds audio features x .
- E_t : The text encoder, which embeds text features y .
- E_s^{stream} : The shared streaming encoder, which may consume either $E_a(x)$ or $E_t(y)$ and maps them to a joint representation. Since this encoder is “streaming”, it only receives past acoustic frames.
- $E_s^{\text{full-context}}$: The full-context encoder, which consumes the outputs of E_s^{stream} and which is given forward acoustic frames.
- D^{stream} : The streaming decoder, which consumes the outputs of E_s^{stream} and emits streaming ASR hypothesis.
- $E_s^{\text{non-streaming}}$: The non-streaming decoder, which consumes the outputs of $E_s^{\text{full-context}}$ and emits non-streaming ASR hypothesis.

Our model is trained simultaneously on two tasks: ASR, and masked text reconstruction. For ASR, audio is passed into audio encoder, and hypothesis are compared against ground truth text with the conventional cross-entropy loss. Masked text reconstruction makes use of unpaired text data. A mask (15% of the transcript) is applied to a phonemic representation of text, which is then passed into the text encoder. The hypothesis is compared to the masked portion of the input again with a cross-entropy loss.




Figure 1: Our architecture for semi-supervised ASR, adapted from [13] and [12].

2.2. Consistency Loss

Consider a paired example (x, y) , where $x = (x_0, \dots, x_n)$ are speech inputs and $y = (y_0, \dots, y_m)$ are text inputs and where $n > m$. Let us define the shared representations of audio and text as

$$R_a = E_s(E_a(x)) \quad R_t = E_s(E_t(y))$$

where E_s can represent either the application of only E_s^{stream} (for a streaming representation) or E_s^{stream} followed by $E_s^{\text{non-streaming}}$ (for a non-streaming representation). A “consistency loss”, as developed in [15] and [12], is some loss

$\mathcal{L}^{\text{consistency}}(R_a, R_t)$ that measures the similarity of the two representations.

Since the audio x and the text y are sequences of different lengths, we require some notion of an alignment to define a meaningful consistency loss. By alignment, we mean a specific up-sampling of y such that each audio frame $x[i]$ will correspond to some text frame $y[j]$. With this in mind, we define an alignment $\mathcal{A} = (a_0, a_1, \dots, a_n)$ as a list of indexes into y , such that for any audio frame i , $x[i]$ corresponds to $y[a_i]$ in the alignment. We will also add the constraint that $a_i \leq a_{i+1}$ for all i . That is, we constrain \mathcal{A} to be monotonically increasing, so that so that sequential audio frames may not correspond to text backwards.

This formulation is one of many conceivable ways to define an “alignment” and we’ve chosen it for the practicality it offers in efficiently computing the best alignment (see Section 2.3 below). We note that in this formulation, each audio frame is considered exactly once, while each text frame can be repeated or skipped over entirely.

With this definition in mind, we define the consistency loss for a given alignment as

$$\mathcal{L}_{\mathcal{A}}^{\text{consistency}}(R_a, R_t) = \sum_{x=0}^n \frac{\mathcal{L}^{\text{frame}}(R_a[x], R_t[\mathcal{A}[x]])}{n}$$

where $\mathcal{L}^{\text{frame}}$ is some frame-wise similarity measure (in this work, we use L2). That is, $\mathcal{L}_{\mathcal{A}}^{\text{consistency}}(R_a, R_t)$ gives the average frame-wise similarity between R_a and the specific up-sampling of R_t given by \mathcal{A} .

The setups in [15] and [12] use such a consistency loss successfully, taking \mathcal{A} from a neural alignment model. We propose, as an alternative, to optimize the consistency over the *best possible* alignment:

$$\mathcal{L}^{\text{consistency}}(R_a, R_t) = \min_{\mathcal{A}} \mathcal{L}_{\mathcal{A}}^{\text{consistency}}(R_a, R_t)$$

In order to train with such a loss, we require an efficient algorithm to compute the best alignment.

2.3. Computing the Best Alignment

Dynamic time warping [18] relies on an inductive rule in order to define a recursive algorithm to match two sequences based on a cost function. We do the same, specifying the cost:

$$\mathcal{C}(i, j) = \min_{\mathcal{A}} \mathcal{L}^{\text{consistency}}(R_a[:i], R_t[:j])$$

That is, the cost $\mathcal{C}(i, j)$ gives the consistency loss under the best alignment between the prefix of the audio representation up to the index i and the prefix of the text representation up to the index j . We may then specify a inductive rule:

$$\mathcal{C}(i, j) = \min_{k \leq j} [\mathcal{C}(i, k-1) + \mathcal{L}^{\text{frame}}(R_a[i], R_t[k])]$$

That is, the best alignment for the prefixes $R_a[:i]$ and $R_t[:j]$ aligns the previous $i-1$ audio frames to some shorter prefix $R_t[:k]$, and then appends to it the specific alignment of $R_a[i]$ to $R_t[k]$.

We may back out the indexes of the best alignment from this computation. This rule gives rise to a dynamic programming algorithm for finding the best alignment in $\mathcal{O}(nm^2)$ time and memory.

We note that the minimization across all alignments precludes differentiation of the alignment-finding. Instead, we compute the best alignment during forward-propagation, and then differentiate $\mathcal{L}^{\text{frame}}$ as applied to the aligned frames. That is, we use the pass-through approximation of the gradient:

$$\frac{\partial \mathcal{L}^{\text{consistency}}(R_a, R_t)}{\partial \theta} \approx \frac{\partial \mathcal{L}_{\mathcal{A}^*}^{\text{consistency}}(R_a, R_t)}{\partial \theta}$$

where

$$\mathcal{A}^* = \arg \min_{\mathcal{A}} \mathcal{L}_{\mathcal{A}}^{\text{consistency}}(R_a, R_t)$$

3. Experiments

In this section, we provide details of our model settings and data.

3.1. Model Settings

We specify component’s parameterizations according to the list in Section 2.1:

- E_a : The audio encoder consists of a single conformer [19] layer with 8 attention heads and dimension 2048. The audio encoder consumes 128 dimensional log-mels spanning 32ms each and spaced apart by 10ms. We then stack each frame with the frame before it and the two frames after it to yield 512 dimensional representations. Finally, we subsample by taking each third frame, yielding a final frame rate of 30ms.
- E_t : The text encoder consists of an embedding projection followed by a conformer layer. As in [13], we find it necessary to supply the text encoder with phonemic representations of text transcripts. We then continue to follow [13] by repeating each phoneme twice as an alignment heuristic.
- E_s^{stream} : The shared streaming encoder consists of five conformer layers, with layer-norm applied at the end.
- $E_s^{\text{full-context}}$: The full-context shared encoder consists of nine additional conformer layers, with layer-norm applied at the end.
- D^{stream} : The streaming decoder is a HAT decoder [20] in which both the prediction and joint layers have dimension 640.
- $E_s^{\text{non-streaming}}$: The non-streaming decoder, is identical to the streaming decoder.

Together, our model contains about 165M parameters. Training is done in two phases. First, the audio encoder, joint encoders, and decoders are all trained on paired data for 800k steps with a batch size of 2048. The text encoder is then added and the model is further trained with equally weighted supervised and unsupervised loss as described in Section 2.1, with the best alignment loss from Section 2.3 optionally included. The model trains in this manner for 100k further steps with a batch size of 2048 for both the supervised and unsupervised data.

All models are implemented in Tensorflow, with the best alignment algorithm itself implemented as a CPU kernel. We find that the addition of the best alignment computation does not significantly increase training time over the baseline model.

3.2. Datasets

Text-injection methods in ASR have historically been applied in two broad settings. Strong baselines are often fine-tuned

with very large text corpora to improve performance on difficult words. Alternatively, text-injection may be used for models trained on limited supervised data may be used to improve the internal language model and get closer to a viable system. With these two settings in mind, we study the best alignment loss in two setups:

- A large English corpus consisting of about 200k hours of supervised speech, together with an unsupervised text dataset of about 200B examples.
- A multilingual corpus consisting of the following eleven languages: English (**En**), French (**Fr**), Spanish (**Sp**), Arabic (**Ar**), Portuguese (**Po**), German (**De**), Russian (**Ru**), Hindi (**Hi**), Italian(**It**), Mandarin, and Japanese.

This setting involves no unsupervised text, with the MLM objective applied instead to the supervised transcripts. The dataset consists of about 140M paired examples.

Bolded abbreviations are given above for languages for which we are able to report WER in 1. For simplicity with the large number of test sets, we report only non-streaming WER from this model.

All datasets are anonymized and human transcribed.

4. Results

In this section, we seek to demonstrate that even without consistency regularization, our model learns an alignment between paired speech and text examples. We then seek to show that optimizing this alignment with our proposed best-alignment consistency regularization improves WER.

4.1. Best-Alignment in an Unregularized Model

For this analysis, we use our baseline model from the monolingual setup as described in Section 3.1. Our objective is to measure $\mathcal{L}^{\text{consistency}}$ on a small set of random development examples for R_s and R_t taken at each of the first five conformer layers of the streaming joint encoder. We interpret a lower value for $\mathcal{L}^{\text{consistency}}$ as reflecting a stronger implicit alignment between speech and text.

For each layer l we sample 2000 random pairs of audio and text embeddings and compute the mean μ_l and standard deviation σ_l^2 of the distribution of distances between pairs. We then compare two alignments: the naive frame-wise alignment and our computed best alignment. For each of these alignments \mathcal{A} , we report:

$$\mu_l - \frac{\mathcal{L}_{\mathcal{A}}^{\text{consistency}}(R_s, R_t)}{\sigma_l^2}$$

That is, we report the consistency in terms of the number standard deviations away from the mean, such that a result of 0 suggests that the alignment is no better than random and a result below 0 suggests that the alignment is stronger than random.

Table 2 presents these measurements. We see that while the consistency of the frame-wise alignment is close to that of the random alignment, the best alignment is considerably better than random. Furthermore, the quality of the best alignment improves steadily as we progress deeper into the model. That is, while text and speech are not modeled jointly at the frame level, there is *some* alignment for which paired speech and text

	En	Fr	Sp	Ar	Po	De	Ru	Hi	It
M_0	9.1	10.6	6.4	12.6	7.9	14.8	13.0	19.7	10.3
M_10	8.5	10.4	5.8	11.8	7.7	13.4	12.5	19.4	9.8
M_1	8.5	10.5	6.1	11.9	8.1	13.9	12.7	19.3	10.0
M_0.1	8.6	10.3	6.2	12.1	8.0	13.9	12.9	19.5	9.9
M_0.01	8.8	10.5	6.3	12.2	7.9	14.0	13.0	19.6	10.3

Table 1: Evaluation Results for the Multilingual Setting.

Table 2: Consistency of the linear and best alignments at layers of the shared encoder.

Layer	Frame-wise Alignment	Best Alignment
1	-0.06	-1.47
2	-0.23	-2.15
3	-0.29	-2.61
4	-0.37	-2.67
5	-0.49	-3.06

are mapped to similar points in the embedding space, and this alignment improves with the depth of the network.

To illustrate the presence of this alignment, we visualize the relationship between shared encoder’s final representations of the speech and text from a single test example. Figure 2a plots the distance between each pair of frames in the embeddings, and demonstrates that is indeed a single alignment with low distance. Figure 2b shows how the best alignment algorithm recovers this trajectory.

	Main	Noisy		Main	Noisy
E_0	5.40	8.75	E_0	7.99	13.33
E_10	5.37	8.70	E_10	8.21	13.08
E_1	5.35	8.42	E_1	8.07	13.00
E_0.1	5.27	8.77	E_0.1	7.90	12.63
E_0.01	5.32	8.54	E_0.01	7.94	12.74

(a) Non-Streaming


(b) Streaming

Table 3: Evaluation Results for the English-Only Setting.


4.2. Consistency Regularization Results

We present results for the best-alignment loss at different interpolation weights and for both of the settings specified in Section 3.2. For ease of reading, we specify each experiment by a letter and a number. The letter is either **E** for the English-only setting or **M** for the multilingual setting. The number is the interpolation weight of the best-alignment loss as a percentage. For example, **E_0** is the baseline English-only model with unregularized semisupervised finetuning, while **M_0.1** is a multilingual model with the best alignment loss interpolated during finetuning at 0.1 percent.

Table 3 gives our results in the high-resource, English-only setting. There, we see small but consistent WER improvements with the best-alignment loss, although we note the necessity of selecting the correct interpolation weight. Table 1 gives our results in the multilingual setting, where we see larger improve-



(a) Distances



(b) Best Alignment

Figure 2: Visualizations of embedding distances (a) and the best alignment (b) between an audio embedding on the horizontal axis and the corresponding text embedding on the vertical axis. Darker points in (a) represent pairs of audio and text frames with nearby embeddings, and yellow points in (b) represent pairs in the recovered best alignment.

ments. We believe that the strength of the method in the multilingual setting is due to the increased difficulty of the problem and the smaller dataset leaving more room for the model to improve.

5. Conclusions

We’ve shown that a semi-supervised speech/text encoder learns a joint representation of the two modalities that can be observed by choosing the best alignment. We’ve exploited that fact to enforce domain consistency with an extra loss term which optimizes the modality match for the best alignment. We show consistent improvements over an unregularized joint model across multiple setups without adding any parameters.

6. References

- [1] 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," in *Advances in Neural Information Processing Systems*, 2020.
- [2] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, P. Schuh, K. Shi, S. Tsvyashchenko, J. Maynez, A. Rao, P. Barnes, Y. Tay, N. Shazeer, V. Prabhakaran, E. Reif, N. Du, B. Hutchinson, R. Pope, J. Bradbury, J. Austin, M. Isard, G. Gur-Ari, P. Yin, T. Duke, A. Levskaya, S. Ghemawat, S. Dev, H. Michalewski, X. Garcia, V. Misra, K. Robinson, L. Fedus, D. Zhou, D. Ippolito, D. Luan, H. Lim, B. Zoph, A. Spiridonov, R. Sepassi, D. Dohan, S. Agrawal, M. Omernick, A. M. Dai, T. S. Pillai, M. Pellat, A. Lewkowycz, E. Moreira, R. Child, O. Polozov, K. Lee, Z. Zhou, X. Wang, B. Saeta, M. Diaz, O. Firat, M. Catasta, J. Wei, K. Meier-Hellstern, D. Eck, J. Dean, S. Petrov, and N. Fiedel, "Palm: Scaling language modeling with pathways," 2022.
- [3] S. Yang, P. Chi, Y. Chuang, C. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G. Lin, T. Huang, W. Tseng, K. Lee, D. Liu, Z. Huang, S. Dong, S. Li, S. Watanabe, A. Mohamed, and H. Lee, "SUPERB: speech processing universal performance benchmark," in *INTERSPEECH*, 2021.
- [4] Z. Borsos, R. Marinier, D. Vincent, E. Kharitonov, O. Pietquin, M. Sharifi, O. Teboul, D. Grangier, M. Tagliasacchi, and N. Zeghidour, "Audiolm: a language modeling approach to audio generation," 2022.
- [5] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, Y. Hasson, K. Lenc, A. Mensch, K. Millican, M. Reynolds, R. Ring, E. Rutherford, S. Cabi, T. Han, Z. Gong, S. Samangooei, M. Monteiro, J. Menick, S. Borgeaud, A. Brock, A. Nematzadeh, S. Sharifzadeh, M. Binkowski, R. Barreira, O. Vinyals, A. Zisserman, and K. Simonyan, "Flamingo: a visual language model for few-shot learning," in *Advances in Neural Information Processing Systems*, 2022.
- [6] J. Cho, J. Lei, H. Tan, and M. Bansal, "Unifying vision-and-language tasks via text generation," in *International Conference on Machine Learning*, 2021.
- [7] A. Renduchintala, S. Ding, M. Wiesner, and S. Watanabe, "Multi-modal data augmentation for end-to-end ASR," in *INTERSPEECH*, 2018.
- [8] Y. Huang, H. Kuo, S. Thomas, Z. Kons, K. Audhkhasi, B. Kingsbury, R. Hoory, and M. Picheny, "Leveraging unpaired text data for training end-to-end speech-to-intent systems," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2020.
- [9] S. Mariooryad, M. Shannon, S. Ma, T. Bagby, D. Kao, D. Stanton, E. Battenberg, and R. Skerry-Ryan, "Learning the joint distribution of two sequences using little or no paired data," 2022.
- [10] Y. Tang, J. M. Pino, C. Wang, X. Ma, and D. Genzel, "A general multi-task learning framework to leverage text data for speech to text tasks," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2021.
- [11] A. Bapna, C. Cherry, Y. Zhang, Y. Jia, M. Johnson, Y. Cheng, S. Khanuja, J. Riesa, and A. Conneau, "mislam: Massively multilingual joint pre-training for speech and text," 2020.
- [12] Z. Chen, Y. Zhang, A. Rosenberg, B. Ramabhadran, P. Moreno, A. Bapna, and H. Zen, "Maestro: Matched speech text representations through modality matching," in *INTERSPEECH*, 2022.
- [13] T. N. Sainath, R. Prabhavalkar, A. Bapna, Y. Zhang, Z. Huo, Z. Chen, B. Li, W. Wang, and T. Strohman, "Joist: A joint speech and text streaming model for asr," in *IEEE Spoken Language Technology Workshop*, 2022.
- [14] S. Thomas, H.-K. J. Kuo, B. Kingsbury, and G. Saon, "Towards reducing the need for speech training data to build spoken language understanding systems," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022.
- [15] Z. Chen, Y. Zhang, A. Rosenberg, B. Ramabhadran, P. Moreno, and G. Wang, "Tts4pretrain 2.0: Advancing the use of text and speech in asr pretraining with consistency and contrastive losses," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022.
- [16] C. Chadebec, E. Thibeau-Sutre, N. Burgos, and S. Allasonniere, "Data augmentation in high dimensional low sample size setting using a geometry-based variational autoencoder," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [17] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International Conference on Machine Learning*, 2020.
- [18] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," in *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 1978.
- [19] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," in *INTERSPEECH*, 2020.
- [20] E. Variani, D. Rybach, C. Allauzen, and M. Riley, "Hybrid autoregressive transducer (hat)," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2020.