

Beyond Instructional Videos: Probing for More Diverse Visual-Textual Grounding on YouTube

Jack Hessel

Allen Institute for AI

jackh@allenai.org

Zhenhai Zhu Bo Pang Radu Soricut

Google Research

{zhenhai, bopang, rsoricut}@google.com

Abstract

Pretraining from unlabelled web videos has quickly become the de-facto means of achieving high performance on many video understanding tasks. Features are learned via prediction of *grounded* relationships between visual content and automatic speech recognition (ASR) tokens. However, prior pretraining work has been limited to only instructional videos; *a priori*, we expect this domain to be relatively “easy:” speakers in instructional videos will often reference the literal objects/actions being depicted. We ask: can similar models be trained on more diverse video corpora? And, if so, what types of videos are “grounded” and what types are not? We fit a representative pretraining model to the diverse YouTube8M dataset, and study its success and failure cases. We find that visual-textual grounding is indeed possible across previously unexplored video categories, and that pretraining on a more diverse set results in representations that generalize to both non-instructional and instructional domains.

1 Introduction

Self-supervised pretraining approaches have recently been adapted to web videos (Sun et al., 2019a,b; Miech et al., 2019, 2020; Zhu and Yang, 2020; Amrani et al., 2020); the resulting models have achieved state-of-the-art performance on a wide range of video understanding tasks, e.g., dense caption generation, action localization, etc.

In general, the pretraining step requires a large, unlabelled corpus of web videos. The training objective aligns visual content (i.e., video segments) with automatic speech recognition (ASR) tokens, and the resulting representations are fine-tuned for downstream tasks. The assumption underlying this family of approaches is that, in the pretraining corpus, spoken words have *some* consistent,

grounded relationship with the temporally corresponding visual content.

However, in contrast to the highly diverse corpora utilized for text-based pretraining (Wikipedia, Common Crawl, etc.), pretraining for web videos (so far) has been limited to instructional videos. This domain restriction is motivated by the commonly accepted notion that “procedural knowledge tends to be inherently multimodal” (Malmaud et al., 2015). We *expect* that the semantic information in video frames and ASR tokens is readily correlated in instructional videos. But corpus diversity brings significant benefits: in the text-only case, models can effectively represent diverse real-world entities (Roberts et al., 2020) precisely because pretraining is not restricted to, e.g., only fictional stories (Zhu et al., 2015).

In search of more general representations, our main question is: **does video-ASR pretraining “work” for more diverse pretraining corpora?** Are certain categories of non-instructional videos “groundable,” thus enabling diverse representation learning? Or are some types too difficult, only acting as training noise? We conclude that: 1) grounding is indeed possible in a wide range of yet-to-be-computationally-exploited YouTube video categories, e.g., walk-throughs, vehicles, tech reviews, etc., with some harder than others; 2) transferable representations can be successfully learned by training on a more diverse set, which may provide more versatility.

2 Related Work

ASR is known to be a useful signal source in various instructional video understanding tasks (Gupta et al., 2017; Huang et al., 2017; Huang* et al., 2018; Moriya et al., 2019), e.g., action detection/classification (Yu et al., 2014; Alayrac et al., 2017; Chang et al., 2019; Kuehne et al., 2019),

segmentation/captioning (Sener et al., 2015), and instruction alignment (Malmaud et al., 2015; Alayrac et al., 2016). A number of multimodal instructional video datasets have been proposed (Wang et al., 2019; Tang et al., 2019; Sanabria et al., 2018). A notable recent example of work addressing a non-instructional video corpus is Ignat et al. (2019), who analyze grounded-ness in lifestyle vlogs. Fouhey et al. (2018) highlight the difference between keyword search vs. implicitly mining action data of interest from a broader corpus (e.g., Bregler (1997); Gu et al. (2018)).

Operational grounding. Our work builds upon prior operational notions of grounding: if an algorithm is able to consistently predict specific visual-textual relationships, then that relationship is said to be “grounded” (Lu et al., 2008; Berg et al., 2010; Parikh and Grauman, 2011; Hill and Korhonen, 2014; Hessel et al., 2018). Yanai and Barnard (2005), for example, examine an image+text corpus and rank “substrings of text by how well their occurrence can be predicted from visual features.” One shortcoming of any model-based operationalization of “grounding” is that only *positive* instances of groundedness can be identified: if one model fails to ground something, perhaps a better model could have.

3 Video-ASR pretraining + our model

Recent work in designing pretraining objectives: 1) assumes that ASR tokens have, on average, some correspondence to temporally co-occurring video frames within the same video; and 2) ignores clips that lack ASR. We consider a model that encapsulates both of these assumptions.¹ The model is a slight simplification of Miech et al. (2019), where a joint embedding for the visual content and ASR tokens is learned. While more sophisticated methods based on self-attention models have since been examined (e.g., Zhu and Yang (2020)), joint embedding models are still performant and offer greater interpretability, thus enabling our later error analyses.

Model details. The similarity between clip i and ASR caption j , $s_{i,j}$, is estimated by computing the cosine similarity between their corresponding embeddings in the joint space. Joint embedding models are parameterized using gated, multi-layer

¹While more complicated models are possible, our goal is to conduct an error analysis of a simple, representative model, not to necessarily achieve state-of-the-art results.

feedforward networks. The visual features we use as input are: frame-wise 2D Inception-v1 pre-trained for object detection (Szegedy et al., 2015; Sun et al., 2017) and 3D CNN S3D-G features pretrained for action recognition (Xie et al., 2018; Kay et al., 2017). The language feature inputs are 300 dimensional vectors per word-type; these are fine-tuned during the training process. Max pooling is used for both token embeddings and for per-frame visual features to achieve a single visual and textual embedding for each clip.²

During training, temporally corresponding (clip, ASR caption) pairs are sampled (“Positive” cases). For each positive case, a set of mismatched “Negative” cases is also sampled both from other videos and from the same video in equal proportion. In contrast to Miech et al. (2019), we control for clip length, and sample temporally *fixed-length* segments. In initial experiments with variable-length segments, we found that our models were capable of “cheating” the grounding task by aligning longer (and shorter, respectively) clips with longer (and shorter) ASR captions, largely ignoring content. Thus, this simplifying choice makes our error analysis significantly more straightforward, and results in minimal performance change. We use 5 second segments, but results are similar with 10 or 30 second windows (see Appendix E). To generate segments, we initially randomly sample 256 per video before discarding ones that have no temporally-accompanying ASR. Segments may overlap, though results are similar without overlaps (see Appendix D). The following hinge loss is minimized for margin δ :

$$\sum_{i,j \in \mathcal{P}, \mathcal{N}} \max(0, \delta + s_{i,j} - s_{i,i}) + \max(0, \delta + s_{j,i} - s_{i,i}) \quad (1)$$

We trained with Adam (Kingma and Ba, 2015), a learning rate of .001, and set $\delta = .1$, but didn’t undertake significant hyperparameter optimization. We terminate training after 300K steps.

CrossTask replication. To verify that our model simplifications didn’t significantly hinder performance, we replicated key experiments from Miech et al. (2019). In particular, we sought to gather the pretraining corpus they used, HowTo100M, which consists of 1.22M videos. Because of, e.g., users deleting videos, we were able to gather features for only 87% of the original set, 1.06M videos.

²When training on YouTube-600K the vocabulary size is 61K.

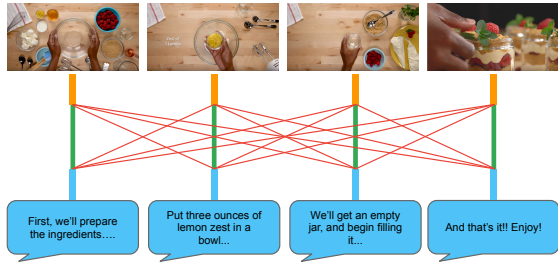


Figure 1: Intra-video AUC metric: the model scores all possible links between clips and ASR captions within a single video; the model is rewarded for assigning higher similarity to **temporally-aligned segments** versus **mismatched ones**.

We verify the performance of our model using the CrossTask localization task (Zhukov et al., 2019). While we defer details to the original paper, the goal of CrossTask is to temporally localize a set of procedural steps for a task in an unlabelled/unsegmented video depicting that task. An algorithm’s performance is evaluated with a recall metric (higher is better). We follow the evaluation procedure given in Miech et al. (2019), except instead of embedding each frame individually, we embed a sliding 5-second window of video clips.

Our simplified model trained on less data performs comparably to Miech et al. (2019)’s. We achieve 32.6 recall, while they report 33.6 recall; for reference, a supervised upper-bound without pre-training achieves 31.6 recall (full results and more details are in Appendix B).

Measuring visual-textual alignment. Viewed through the lens of link prediction between truly co-occurring (clip, ASR) pairs, Eq. 1 can be seen as a differentiable approximation of AUC (Rendle et al., 2009). Thus, we propose to operationally measure the groundedness using *intra-video* AUC: a single score is assigned to each video, rewarding the model if it is able to successfully align temporal pairs within the same video (and penalizing it if not). Fig. 1 presents a visualization of this method. One notable advantage of AUC versus other link prediction metrics is that it is insensitive to the label distribution: shorter videos are not systematically assigned higher scores simply because there are fewer incorrect links.

4 A More Diverse Corpus

YouTube-600K. YouTube8M (Abu-El-Haija et al., 2016) is a dataset of 6.1M YouTube

Category	% videos	Category	% videos
BBC	74.1%	Pachinko	0.4%
President (USA)	71.7%	Jumbotron	0.4%
Hair conditioner	71.1%	Chipmunk	0.3%
Madden NFL	69.3%	Taiko no Tatsujin	0.2%
Wig	67.9%	Yo-kai Watch	0.1%
Magic (card game)	67.9%	Zee Bangla	0.1%
Booster pack	67.4%	Karaoke box	0.1%
Raw foodism	66.5%	Wangan Midnight	0.1%
NBA 2K14	65.2%	Caporales	0.0%
Silver	65.2%	Military band	0.0%

(a) Most ASR

(b) Least ASR

Table 1: Categories of YouTube8M with the highest and lowest availability of English ASR (minimum 1K videos); corpus mean = 17%.

videos,³ where each video is labeled across 3K categories, ranging from “cooking” to “games” to “nature.” It is among the largest and most diverse publicly available dataset of YouTube videos. Due to user deletions and videos without detected spoken words, we are able to collect ASR via the YouTube API for 1.4M (29%) videos; we further filtered to 817K videos tagged with English ASR.⁴ There is an extremely wide variance of ASR availability per category, e.g., 74% of “BBC” videos (a category that generally contains news videos by the broadcaster) have ASR, whereas almost no “Military band” videos do (Table 1). While the percentage of ASR-available videos is higher in many instructional video categories, e.g., “cooking” at 31%, “cosmetics” at 44%, etc., many non-instructional categories on YouTube have ASR available (e.g., “silver” at 65%; mostly videos about coins). Maintaining the train / validation split of the original data release yields 639K training videos (henceforth referred to as YouTube-600K) and 167K validation-set videos.

Human annotation of “Is-it-instructional” While a qualitative examination of YouTube8M reveals clear topical and stylistic diversity compared to domain-restricted corpora, we quantitatively verify that YouTube8M does not consist of mostly instructional videos.

We sample 6.8K videos with English ASR from the validation set for human labeling. Each video is shown to three paid annotators, who each provide a Yes/No answer to the question: “Does this

³v3 of the dataset is smaller than v1/v2, due to videos becoming unavailable over time and other refinements.

⁴We expect that non-English videos will similarly be an excellent source of visual-textual grounding training data, particularly for under-resourced languages. We focused on English to simplify our error analyses. But in future work, we expect to not impose such a limitation.

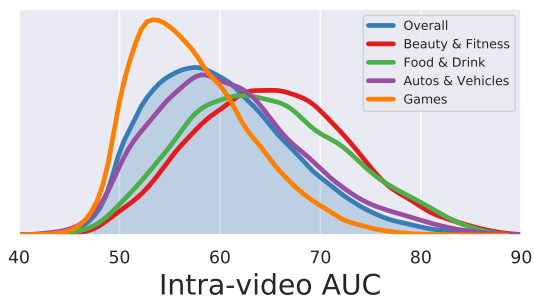


Figure 2: Distribution of intra-video AUC scores, grouped by meta category, compared to the overall distribution.

video focus on real-world human actions accompanied by procedural language that explains what is happening on screen in reasonable detail?” Note that our definition of “instructional” intends to include the usual “how-to” videos, but also attempts to capture a more general notion of “instructional-ness”. For instance, an un-boxing video where parts of a product are *taken out and assembled* along with corresponding narration should receive “Yes”, whereas a video showing only a product from different angles should receive “No”, due to a lack of narrated human actions.

After a pilot study with a few iterations over the guidelines and examples, the annotators reach high agreement: in 96% of cases, all three judges are unanimous. From these annotations, we estimate that around 74% of the videos in the YouTube-600K corpus are **not** instructional, even for the generalized notion of “instructional-ness.” For reference, [Miech et al. \(2019\)](#) conduct an analysis of 100 videos from HowTo100M (constructed with the intention to focus on how-to videos) and estimate that 71% are instructional.

The annotated *i3-video* corpus ([is-it-instructional-video](#)) is available for download.⁵ One potential use-case: consider an automated tool designed exclusively for use on instructional videos. A classifier trained on our labelled corpus could be used to determine if applying the automated tool is appropriate or not for an unlabelled input video.

Which categories are easiest/hardest? We train our model on YouTube-600K, and compute intra-video AUC for each of the 178K validation videos. First, we average all videos labeled with a particular category to produce a per-category AUC score. The performance in a vast majority of categories

⁵<https://github.com/google-research-datasets/i3-video>

Domain	Example Categories (AUC)
Vehicles	Crossover SUV (70); Sedan (69); Minivan (69); Station wagon (68)
Walkthroughs	Hotel Suite (71); Apartment (69); Dining room (68); Living room (68)
Advertising	Advertising (68); Television advertisement (66); Infomercial (63)
Tech Reviews	CNET (66), Netbook (65), Asus (63)
Toys	iPhone 5S (64), MacBook (64)
Appliances	Funko (66); Monster High (64); Figurine (64)
Places	Action Figure (64)
	Home appliance (65); Washing machine (64); Kitchen stove (64)
	Greenhouse (62) University (61); Amusement park (59)

Table 2: Domains with high operational groundability.

is above the 50 AUC (random) baseline, and ranges from 51 (“Mixtape”) to 76 (“Muffin”). To make sure that the model is not succeeding simply because a category happened to be frequent in the dataset, we note the correlation between category AUC and category frequency is essentially zero ($\rho = .02, p > .58$). This suggests that at least *some* aspect of most categories of videos can be visual-textually grounded.

We next coarsely aggregate the YouTube8M categories into meta-categories, e.g., “Food and Drink.”⁶ The AUC distribution of 4 popular meta categories relative to the overall AUC distribution is given in Fig. 2. In general, the grounding succeeds most readily on makeup/hair videos (e.g., “Eye liner” AUC = 74, “Updo” AUC = 68, etc.) and cooking videos (e.g., “Vegetarian cuisine” AUC = 71), domains that have been previously used in video grounding work. Besides these already-studied domains, other high-scoring category types emerge (Table 2). Conversely, some categories are more difficult for the model, e.g., video game categories like “RuneScape” AUC = 54 and “First-person shooter” AUC = 55; speakers in these videos often reference diverse topics unrelated to the game itself. Non-video-game categories can also be difficult, e.g., “Unidentified flying object” AUC = 56, “Dashcam” AUC = 54.

We next ask: *are instructional videos indeed easier to operationally ground?* While human judgements of instructional-ness and intra-video AUC are positively correlated $\rho = .20$ ($p \ll 0$), the low magnitude of this correlation provides additional empirical confirmation that other types of videos are also promising.

Within-category observations. To this point,

⁶These meta-categories are called “verticals,” and are released with YouTube8M.




Figure 3: Correlation between per-segment AUC scores and segment timing within video (left column) and the number of tokens in a segment (right column)

we have identified broad categories of YouTube videos that are more groundable than others. However, it is not yet clear *why*, e.g., the algorithm gets 64 AUC on “Action Figure,” or 55 AUC on “Call of Duty” (a first-person shooter game). We now define a *segment-level* AUC metric, analogous to the intra-video AUC metric previously defined: it quantifies how readily *individual ASR captions* are temporally localized by the model within the same video (see Menon and Elkan (2011) for a description of different AUC variants).

Before examining the relationship between content-based features and segment-level AUC, contextual factors must be considered. Fig. 3 illustrates clear relationships 1) between ASR caption placement within a video and segment AUC (segments at the very beginning and very end of videos tend to be easier); and 2) between the number of tokens in an ASR caption and segment AUC. For “Action Figure”, ASR segments with more words are easier (this is the case with most categories), but for “Call of Duty”, the opposite is true.

After controlling for contextual variables, we train OLS regression models to predict segment AUC from lexical unigram features, while controlling for timing/length features. Lexical features add predictive capacity ($p \ll .01$, F-test). While we find some patterns predictive of segment AUC for both categories, e.g., intro/outro-language (e.g., “hey”, “welcome”, “peace”), we also observe topical patterns, e.g., several unigrams associated with specific action figure body parts (“knee”, “shoulder”, “joint”, etc.) are positively associated with segment AUC.

5 Implications for Pretraining

While we’ve thusfar shown that self-grounding is possible for a diverse set of domains, do we *gain* anything by training on a more diverse corpus? Or do difficult-to-ground videos introduce noise and degrade representations for downstream tasks?

We compare two versions of our model: one with parameters learned from training on the diverse YouTube-600K corpus (M_{Diverse}), and one with parameters learned from a domain-specific corpus of 1M instructional videos ($M_{\text{Instructional}}$).

First, we evaluate each model’s capacity to localize instructional steps on the CrossTask (Zhukov et al., 2019) dataset. M_{Diverse} performs admirably, even with significant domain mismatch and fewer pretraining videos: recall drops by only 15% (32.6 \rightarrow 27.4) when swapping from $M_{\text{Instructional}}$ to M_{Diverse} .

We next evaluate each model’s performance on the same-video clip alignment task over a diverse set of videos: the sample of 6.8K human-annotated videos from the YouTube8M validation set. In terms of intra-video AUC, M_{Diverse} outperforms $M_{\text{Instructional}}$ on 59% of videos. If we split the data across the “Is-it-instructional” human judgements and compare the two models in each subset, $M_{\text{Instructional}}$ “wins” in 57% of the instructional videos, whereas M_{Diverse} “wins” in 65% of non-instructional cases.

In short: both models achieve reasonable performance under instructional vs. non-instructional train/test domain mismatch. Taken together, this is a promising result for future pretraining work with more diverse corpora: at least for these evaluations, good performance on an instructional video grounding task is still possible under domain shift. And while the comparison of intra-video AUC is not necessarily definitive, it suggests that diverse corpora may provide more versatility, and we look forward to exploring this further in future work.

6 Conclusion

Peeking through the lens of a joint embedding model, we probe into learning visual-textual grounding over a more diverse corpus of YouTube videos vs. prior work. We find that learning visual-textual grounding is possible across many yet-to-be-explored categories of YouTube videos, and that it’s possible to learn generalizable representations from a more diverse video set.

Acknowledgements. In addition to the anonymous reviewers, the authors would like to thank Soravit Changpinyo, Sebastian Goodman, Bryan Seybold, Chen Sun and Ashish Thapliyal for insightful discussions and implementation help. JH completed this work during an internship at Google.

References

- Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. 2016. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*.
- Jean-Baptiste Alayrac, Piotr Bojanowski, Nishant Agrawal, Josef Sivic, Ivan Laptev, and Simon Lacoste-Julien. 2016. Unsupervised learning from narrated instruction videos. In *CVPR*.
- Jean-Baptiste Alayrac, Ivan Laptev, Josef Sivic, and Simon Lacoste-Julien. 2017. Joint discovery of object states and manipulation actions. In *ICCV*.
- Elad Amrani, Rami Ben-Ari, Daniel Rotman, and Alex Bronstein. 2020. Noise estimation using density estimation for self-supervised multimodal learning. *arXiv preprint arXiv:2003.03186*.
- Tamara L Berg, Alexander C Berg, and Jonathan Shih. 2010. Automatic attribute discovery and characterization from noisy web data. In *ECCV*.
- Christoph Bregler. 1997. Learning and recognizing human dynamics in video sequences. In *Computer Society Conference on Computer Vision and Pattern Recognition*.
- Chien-Yi Chang, De-An Huang, Yanan Sui, Li Fei-Fei, and Juan Carlos Niebles. 2019. D3tw: Discriminative differentiable dynamic time warping for weakly supervised action alignment and segmentation. In *CVPR*.
- David F Fouhey, Wei-cheng Kuo, Alexei A Efros, and Jitendra Malik. 2018. From lifestyle vlogs to everyday interactions. In *CVPR*.
- Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, and Jitendra Malik. 2018. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*.
- Abhinav Gupta, Yajie Miao, Leonardo Neves, and Florian Metze. 2017. Visual features for context-aware speech recognition. In *ICASSP*.
- Jack Hessel, David Mimno, and Lillian Lee. 2018. Quantifying the visual concreteness of words and topics in multimodal datasets. In *NAACL*.
- Felix Hill and Anna Korhonen. 2014. Learning abstract concept embeddings from multi-modal data: Since you probably cant see what I mean. In *EMNLP*.
- De-An Huang*, Shyamal Buch*, Lucio Dery, Animesh Garg, Li Fei-Fei, and Juan Carlos Niebles. 2018. Finding “it”: Weakly-supervised, reference-aware visual grounding in instructional videos. In *CVPR*.
- De-An Huang, Joseph J Lim, Li Fei-Fei, and Juan Carlos Niebles. 2017. Unsupervised visual-linguistic reference resolution in instructional videos. In *CVPR*.
- Oana Ignat, Laura Burdick, Jia Deng, and Rada Mihalcea. 2019. Identifying visible actions in lifestyle vlogs. In *ACL*.
- Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. 2017. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Hilde Kuehne, Ahsan Iqbal, Alexander Richard, and Juergen Gall. 2019. Mining youtube-a dataset for learning fine-grained action concepts from webly supervised video data. *arXiv preprint arXiv:1906.01012*.
- Yijuan Lu, Lei Zhang, Qi Tian, and Wei-Ying Ma. 2008. What are the high-level concepts with small semantic gaps? In *CVPR*.
- Jonathan Malmaud, Jonathan Huang, Vivek Rathod, Nick Johnston, Andrew Rabinovich, and Kevin Murphy. 2015. What’s cookin’? interpreting cooking videos using text, speech and vision. In *NAACL*.
- Aditya Krishna Menon and Charles Elkan. 2011. Link prediction via matrix factorization. In *ECML-PKDD*.
- Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. 2020. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*.
- Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. 2019. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *ICCV*.
- Yasufumi Moriya, Ramon Sanabria, Florian Metze, and Gareth JF Jones. 2019. Grounding object detections with transcriptions. *arXiv preprint arXiv:1906.06147*.
- Devi Parikh and Kristen Grauman. 2011. Interactively building a discriminative vocabulary of nameable attributes. In *CVPR*.

- Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. Bpr: Bayesian personalized ranking from implicit feedback. In *UAI*.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? *arXiv preprint arXiv:2002.08910*.
- Ramon Sanabria, Ozan Caglayan, Shruti Palaskar, Desmond Elliott, Loïc Barrault, Lucia Specia, and Florian Metze. 2018. How2: a large-scale dataset for multimodal language understanding. In *NeurIPS Workshops*.
- Ozan Sener, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. 2015. Unsupervised semantic parsing of video collections. In *ICCV*.
- Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. 2019a. Contrastive bidirectional transformer for temporal representation learning. *arXiv preprint arXiv:1906.05743*.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019b. VideoBERT: A joint model for video and language representation learning. In *ICCV*.
- Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *ICCV*.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *CVPR*.
- Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou. 2019. COIN: A large-scale dataset for comprehensive instructional video analysis. In *CVPR*.
- Weiyang Wang, Yongcheng Wang, Shizhe Chen, and Qin Jin. 2019. Youmakeup: A large-scale domain-specific multimodal dataset for fine-grained semantic comprehension. In *EMNLP*.
- Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. 2018. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*.
- Keiji Yanai and Kobus Barnard. 2005. Image region entropy: a measure of visualness of web images associated with one concept. In *ACM MM*.
- Shou-I Yu, Lu Jiang, and Alexander Hauptmann. 2014. Instructional videos for unsupervised harvesting and learning of action examples. In *ACM MM*.
- Linchao Zhu and Yi Yang. 2020. Actbert: Learning global-local video-text representations. In *CVPR*.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *ICCV*.
- Dimitri Zhukov, Jean-Baptiste Alayrac, Ramazan Gokberk Cinbis, David Fouhey, Ivan Laptev, and Josef Sivic. 2019. Cross-task weakly supervised learning from instructional videos. In *CVPR*.

A Additional Model Details

We adapt [Miech et al. \(2019\)](#)’s joint embedding model that pre-trains by aligning ASR tokens with corresponding video frames. The main difference between our implementation and theirs is how we generated (ASR, caption) pairs. While we considered generating clips according to their methodology, we ran into two problems. First, in early experiments, we found that the interpretability our error analysis was significantly impacted by varying clip length. For example: we were worried that it might not be consistent to compare the model’s ability to temporally ground a 1s clip vs. a 15 second clip. There was also high correlation between caption length and temporal clip duration, which further complicated interpretation. Sampling clips of uniform duration solved these problems.

Second, [Miech et al. \(2019\)](#)’s temporal segmentation was generated by relying on the *scrolling timing* of the ASR tokens on the YouTube, i.e., the time that YouTube decides to generate a linebreak, removing a line of caption from the screen. Via manual inspection, we found that scrolling time was temporally unreliable, e.g., the time in which ASR captions scroll on YouTube often differs significantly from when particular words were said. Instead, we sample 256 candidate 5 second segments uniformly at random from the video, and then discard segments that have no corresponding ASR.

Additional visual feature details. For 2D features, we sample frames at 1FPS from all of the videos in our corpus, resize frames to be 256 by 256, and pass them through Inception-v1 ([Szegedy et al., 2015](#)) pretrained on JFT ([Sun et al., 2017](#)). For 3D convolutional networks, we follow a similar procedure to ([Sun et al., 2019](#)), sample frames at 30FPS, aggregate frames into one second non-overlapping clips of 1 second each, and run an S3D-G ([Xie et al., 2018](#)) network that is pretrained on the Kinetics action recognition dataset ([Kay et al., 2017](#)). Both 2D and 3D features are L2 normalized. The result of this process is a 2524-D feature vector for each second of video in our corpus.

B Comparison to HowTo100M

The full per-task recall comparisons are given in Table 3. Our results, like those of [Miech et al. \(2019\)](#), use [Zhukov et al. \(2019\)](#)’s dynamic programming postprocessing method. We found that

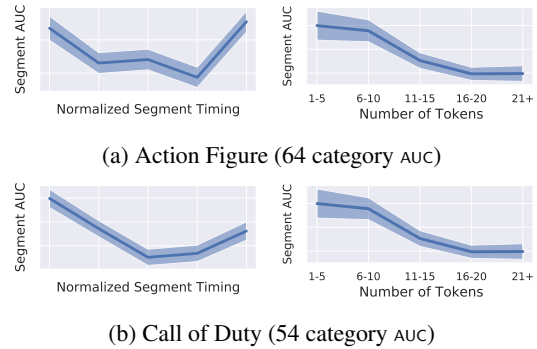


Figure 4: Correlation between per-segment AUC scores and segment timing within video (left column) and the number of tokens in a segment (right column) for the 140K checkpoint (compare to the 300K checkpoint in the main paper).

it usually resulted in a small performance increase.

Our simplified model performs only slightly worse (3%) than [Miech et al. \(2019\)](#)’s. While we argue that our model is certainly still representative, there are several reasons why this gap might exist. For example, there may be a regularizing effect when the model is allowed to view clips of varying length. Furthermore, our feature set was different; we used different (but comparable) base neural networks for feature extraction. Also, our model is trained on less data due to authors deleting their videos. Finally — we didn’t tune the training hyperparameters for our model/implementation, e.g., hinge size, learning rate, batch size, etc.

C Stability of results to checkpoint

To ensure the results related to intra-video AUC were insensitive to the particular choice of model checkpoint, we re-did the experiments in §4 using a version of our model checkpointed at 140K iterations vs. the 300K presented in the main paper; these experiments were conducted over 21K dev videos instead of the full 167K dev videos presented in the main paper. Figures+tables in that section were consistent with the presented results, and the qualitative observations about the “Action Figure” category held (see Figure 4 for replicated figures).

D Stability of results to overlapping windows

When we generate our windows at training and testing time to compute intra-video AUC, given that

	Make Kimchi Rice	Pickle Cucumber	Make Banana Ice Cream	Grill Steak	Jack Up Car	Make Jello Shots	Change Tire	Make Lemonade	Add Oil to Cur	Make Latte	Build Shelves	Make Taco Salad	Make French Toast	Make Insh Coffee	Make Strawberry Cake	Make Pancakes	Make Meringue	Make Fish Curry	Average	
Zhukov et al. (2019)	15.6	10.6	7.5	14.2	9.3	11.8	17.3	13.1	6.4	12.9	27.2	9.2	15.7	8.6	16.3	13.0	23.2	7.4	13.3	
Supervised upper-bound (Zhukov et al., 2019)	19.1	25.3	38.0	37.5	25.7	28.2	54.3	25.8	18.3	31.2	47.7	12.0	39.5	23.4	30.9	41.1	53.4	17.3	31.6	
HowTo100M (1.2M videos) → Crosstask (Miech et al., 2019)	33.5	27.1	36.6	37.9	24.1	35.6	32.7	35.1	30.7	28.5	43.2	19.8	34.7	33.6	40.4	41.6	41.9	27.4	33.6	
↳ only 600K instructional videos																				32.6
↳ only 200K instructional videos																				31.1
Our HowTo100M (1.06M videos) → Crosstask	24.5	30.0	39.9	32.0	27.0	37.2	33.6	33.5	24.4	27.7	44.7	19.1	32.9	31.7	35.3	46.6	43.4	22.9	32.6	
Our YouTube-600K (639K videos; 166K instr) → Crosstask	21.5	24.7	35.2	26.2	19.6	29.5	25.8	30.1	20.9	22.9	32.7	18.4	26.7	27.2	31.0	43.3	37.7	20.5	27.4	

Table 3: Comparison between our simplified model and Miech et al. (2019)’s model on CrossTask, and the effect of pretraining the model on YouTube-600K instead of HowTo100M. Note that HowTo100M → Crosstask results are pre-trained on less data when compared to the original works due to video deletion.

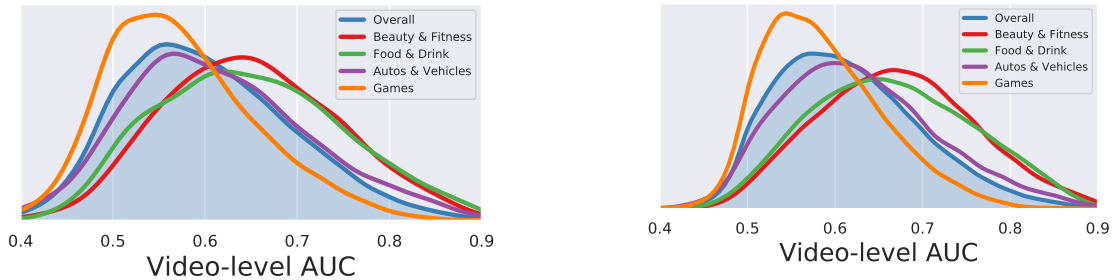


Figure 5: Meta-category AUC with window size of 5, but test-time windows are sampled without overlap.


we sample 256 candidates per video (and then filter out clips without associated temporal ASR), the windows frequently overlap. We ran additional experiments to ensure that our results held when we sampled non-overlapping clips at testing time.

We computed an alternate version of the intra-video AUC results using a 140K training iteration checkpoint. Instead of sampling 256 segments at testing time, we only (randomly) sample up to 10 segments, but force them to be non-overlapping. There are some videos that are discarded in this process. For 1/500 videos (or so) we cannot sample non-overlapping segments. However, among the majority of videos for which the sampling is successful, the Spearman correlation with the category-level results allowing for overlap is $\rho = .98$. Figure 5 reproduces the meta-category plot from the main paper, but with test-time segments sampled without overlap.

E Additional Window Sizes

The results presented in the main paper use a temporal window size of five seconds. We were curious as to the stability of our observations with respect to the choice of this window size. While changing the window size, to an extent, changes

(a) Window size = 10



(b) Window size = 30

Figure 6: Meta-category AUC using models trained with alternate window sizes of $t = 10, 30$ seconds. The main paper results are with $t = 5$ second windows.

the nature of the task, we ran with window size 10s and window size 30s to measure the stability of the results.

Category-level AUC. We computed the intra-video AUC values for the models trained with alternate window sizes. Largely, while the individual AUC values may change by a point or two, the relative pattern stayed the same. The per-category Spearman correlation with the results for the $t = 5$ windows was high: $\rho = .99$ for $t = 10$ and $\rho = .96$ for $t = 30$ (both correlations are true with $p \ll .001$). It should be noted that there are marginal differences in the development set videos used for the experiments at $t = 5, 10, 30s$: because experiments were run at different times, due

Domain	Example Categories (AUC)
Vehicles	Crossover SUV (73); Sedan (72); Minivan (72); Station wagon (71)
Walkthroughs	Hotel Suite (74); Apartment (72); Dining room (71); Living room (70)
Advertising	Advertising (69); Television advertisement (67); Infomercial (64)
Tech Reviews	CNET (67), Netbook (68), Asus (66)
Toys	iPhone 5S (66), MacBook (66)
Appliances	Funko (68); Monster High (66); Figurine (66)
Places	Action Figure (66)
	Home appliance (67); Washing machine (66); Kitchen stove (66)
	Greenhouse (63) University (63); Amusement park (60)

Table 4: High groundability categories with $t = 10$ seconds.

to video deletion, 2% fewer videos were available for later experiments. Figure 6 recreates the meta-category plot from the main paper with the different window sizes; Tables 4 and 5 report the same category-level AUC values as the table in the main paper. Overall, the results are very similar.

Transfer learning. In the main paper, we explored the differences between models trained on YouTube-600K vs. HowTo100M. Here, we conducted the same YouTube8M self-grounding experiments as described in the main paper (varying the training set) with $w = 10$ instead of $w = 5$. The results are very similar. In terms of intra-video AUC, M_{Diverse} (trained on YouTube-600K) outperforms $M_{\text{Instructional}}$ (trained on HowTo100M) on 58% of videos (compared to 59% for $w = 5$ in the main paper). If we split the data across the “Is-it-instructional” human judgments and compare the two models in each subset, $M_{\text{Instructional}}$ “wins” in 57% of the instructional videos (compared to 57% for $w = 5$ in the main paper), whereas M_{Diverse} “wins” in 63% of non-instructional cases (compared to 65% for $w = 5$ in the main paper).

F Additional Reproducibility Information

The models were trained and evaluated on a mix of TPU and GPU clusters. Depending on the particular hardware, the training process takes roughly 1 or 2 days to get 300K training iterations. While we make no particular claim about parameter/time efficiency, the number of parameters of our models are similar to HowTo100M’s, i.e., roughly 50M. The runtime of our model is relatively fast — on GPUs, batches of hundreds of videos can be processed in seconds. For hyperparameters not

Domain	Example Categories (AUC)
Vehicles	Crossover SUV (72); Sedan (72); Minivan (71); Station wagon (71)
Walkthroughs	Hotel Suite (74); Apartment (71); Dining room (72); Living room (71)
Advertising	Advertising (67); Television advertisement (65); Infomercial (61)
Tech Reviews	CNET (65), Netbook (67), Asus (65)
Toys	iPhone 5S (64), MacBook (65)
Appliances	Funko (66); Monster High (65); Figurine (65)
Places	Action Figure (66)
	Home appliance (65); Washing machine (64); Kitchen stove (64)
	Greenhouse (61) University (60); Amusement park (59)

Table 5: High groundability categories with $t = 30$ seconds.

specifically described, we mirror the choices made in the HowTo100M public repo.⁷ For evaluating CrossTask, we wrote our own recall computing code that mirrors the setup publicly released by the authors;⁸ this includes several ad-hoc decisions, e.g., computing the floor/ceiling of temporal annotations in CrossTask to form the window, and only counting a “hit” when the predicted time is strictly *less* than the ceiling (rather than less than or equal).

References

- Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. 2017. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*.
- Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. 2019. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *ICCV*.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. VideoBERT: A joint model for video and language representation learning. In *ICCV*.
- Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *ICCV*.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *CVPR*.

⁷<https://github.com/antoine77340/howto100m>

⁸<https://github.com/DmZhukov/CrossTask>

Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. 2018. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*.

Dimitri Zhukov, Jean-Baptiste Alayrac, Ramazan Gokberk Cinbis, David Fouhey, Ivan Laptev, and Josef Sivic. 2019. Cross-task weakly supervised learning from instructional videos. In *CVPR*.