

# MERLOT RESERVE: Neural Script Knowledge through Vision and Language and Sound

Rowan Zellers<sup>1</sup> Jiasen Lu<sup>2</sup> Ximing Lu<sup>1,2</sup> Youngjae Yu<sup>2</sup> Yanpeng Zhao<sup>3</sup>  
 Mohammadreza Salehi<sup>1</sup> Aditya Kusupati<sup>1</sup> Jack Hessel<sup>2</sup> Ali Farhadi<sup>1</sup> Yejin Choi<sup>1,2</sup>

<sup>1</sup>Paul G. Allen School of Computer Science & Engineering, University of Washington

<sup>2</sup>Allen Institute for Artificial Intelligence

<sup>3</sup>Institute for Language, Cognition and Computation, University of Edinburgh

[rowanzellers.com/merlotreserve](http://rowanzellers.com/merlotreserve)

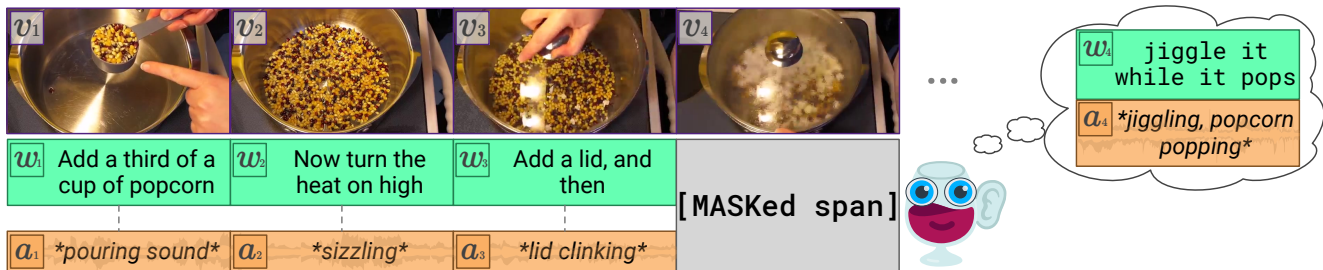


Figure 1: MERLOT RESERVE learns *multimodal neural script knowledge* representations of video – jointly reasoning over video frames, text, and audio. Our model is pretrained to predict which snippet of text (and audio) might be hidden by the MASK. This task enables it to perform well on a variety of vision-and-language tasks, in both zero-shot and finetuned settings.

## Abstract

As humans, we navigate the world through all our senses, using perceptual input from each one to correct the others. We introduce MERLOT RESERVE, a model that represents videos jointly over time – through a new training objective that learns from audio, subtitles, and video frames. Given a video, we replace snippets of text and audio with a MASK token; the model learns by choosing the correct masked-out snippet. Our objective learns faster than alternatives, and performs well at scale: we pretrain on 20 million YouTube videos.

Empirical results show that MERLOT RESERVE learns strong representations about videos through all constituent modalities. When finetuned, it sets a new state-of-the-art on both VCR and TVQA, outperforming prior work by 5% and 7% respectively. Ablations show that both tasks benefit from audio pretraining – even VCR, a QA task centered around images (without sound). Moreover, our objective enables out-of-the-box prediction, revealing strong multimodal commonsense understanding. In a fully zero-shot setting, our model obtains competitive results on four video understanding tasks, even outperforming supervised approaches on the recently proposed Situated Reasoning (STAR) benchmark.

We analyze why incorporating audio leads to better vision-language representations, suggesting significant opportunities for future research. We conclude by discussing ethical and societal implications of multimodal pretraining.

## 1. Introduction

The world around us is dynamic. We experience and learn from it using all of our senses, reasoning over them temporally through *multimodal script knowledge* [98, 126]. Consider Figure 1, which depicts someone cooking popcorn. From the images and dialogue alone, we might be able to imagine what *sounds* of the scene are: the process might begin with raw kernels scattering in an empty, metallic pot, and end with the dynamic ‘pops’ of popcorn expanding, along with the jiggling of a metal around the stove.

Predicting this sound is an instance of *learning from reentry*: where time-locked correlations enable one modality to educate others. Reentry has been hypothesized by developmental psychologists to be crucial for how we as humans learn visual and world knowledge, much of it without need for an explicit teacher [88, 35, 19, 100]. Yet, we ask – can we build machines that likewise learn vision, language, and sound *together*? And can this paradigm enable learning *neural script knowledge*, that transfers to language-and-vision tasks, *even those without sound*?

In this work, we study these questions, and find that the answers are ‘yes.’ We introduce a new model that learns self-supervised representations of videos, through all their modalities (audio, subtitles, vision). We dub our model MERLOT RESERVE<sup>1</sup>, henceforth RESERVE for short.

<sup>1</sup>Short for **M**ultimodal **E**vent **R**epresentation **L**earning **O**ver **T**ime, with

Our model differs from past work that learns from audio-image pairs [53, 70], from subtitled videos [105, 126], or from static images with literal descriptions [106, 20, 91]. Instead, we learn joint representations from *all modalities of a video*, using each modality to teach others. We do this at scale, training on over 20 million YouTube videos.

We introduce a new *contrastive masked span* learning objective to learn script knowledge across modalities. It generalizes and outperforms a variety of previously proposed approaches (e.g. [29, 106, 91, 126]), while enabling audio to be used as signal. The idea is outlined in Figure 1: the model must figure out which span of text (or audio) was MASKED out of a video sequence. We combine our objective with a second contrastive learning approach, tailored to learning *visual recognition* from scratch: the model must also match each video frame to a contextualized representation of the video’s transcript [126]. Through ablations, we show that our framework enables rapid pretraining of a model and readily scales to ‘large’ transformer sizes (of 644M parameters).

Experimental results show that 🎧RESERVE learns powerful representations, useful even for tasks posed over only a subset of the studied modalities. For example, when finetuned on Visual Commonsense Reasoning [124] (a vision+language task with no audio), it sets a new state-of-the-art, outperforming models trained on supervised image-caption pairs by **over 5%**. It does even better on video tasks: fine-tuning without audio, it outperforms prior work on TVQA [74] by a margin of **over 7%** (and given TVQA audio, performance increases even further). These performance improvements do not come at the expense of efficiency: our largest model uses one-fifths the FLOPs of a VisualBERT.

🎧RESERVE also performs well in zero-shot settings. We evaluate on four diverse benchmarks: Situated Reasoning (STAR) [119], EPIC-Kitchens [26], LSMDC-FiB [95], and MSR-VTT QA [120]. These benchmarks require visual reasoning with respective emphasis on *temporality*, *future prediction*, and both *social* and *physical understanding*. With no fine-tuning or supervision, our model obtains competitive performance on each. Of note, it nearly doubles [122]’s SoTA zero-shot accuracy on MSR-VTT QA, and it outperforms supervised approaches (like ClipBERT [73]) on STAR.

Finally, we investigate *why*, and *on which training instances* audio-powered multimodal pretraining particularly helps. For instance, predicting audio rewards models for recognizing *dynamic state changes* (like cooked popcorn) and *human communication dynamics* (what are people’s emotions and towards whom). Our model progressively learns these phenomena as pretraining progresses. These signals are often orthogonal to what snippets of text provide, which motivates learning from both modalities.

In summary, our key contributions are the following:

- a. 🎧RESERVE, a model for multimodal script knowledge, fusing vision, audio, and text.
- b. A new contrastive span matching objective, enabling our model to learn from text *and audio* self-supervision.
- c. Experiments, ablations, and analysis, that demonstrate strong multimodal video representations.

Overall, the results suggest that learning representations from *all modalities* – in a time-locked, reentrant manner – is a promising direction, and one that has significant space for future work. We release code and model checkpoints at [rowanzellers.com/merlotreserve](https://rowanzellers.com/merlotreserve).

## 2. Related Work

Our work brings together two active lines of research.

**Joint representations of multiple modalities.** Many language-and-vision tasks benefit from *early fusion* of different modalities [7]. A family of ‘VisualBERT’ models have been proposed for this: typically, these use a supervised object detector image encoder backbone, and pretrain on images paired with literal captions [106, 76, 80, 20, 123, 73]. Cross-modal interactions are learned in part through a *masked language modeling* (mask LM) objective [29], where subwords are replaced with ‘MASK’, and models independently predict each subword conditioned on both images and unmasked tokens.<sup>2</sup> Extensions to this paradigm include audio in the captioning setting [61] – for instance, OPT [78] was trained on image-caption pairs narrated by a text-to-speech system; yet like the other approaches, this still only learns (from audio) at the level of *literal captions*.

Perhaps closest to our work is MERLOT [126], which learns a joint vision-text model from web videos with automatic speech recognition (ASR). Through a combination of objectives (including a variant of mask LM), MERLOT established strong results on a variety of video QA benchmarks when finetuned. However, it lacks audio: it is limited to learning from (and representing) video frames paired with subtitles. Our proposed 🎧RESERVE, which represents and learns from audio, outperforms MERLOT.

**Co-supervision between modalities.** A common pitfall when training a joint multimodal model is that complex *inter-modal* interactions can be ignored during learning, in favor of simpler *intra-modal* interactions [50, 23, 58]. For example, when using the aforementioned mask LM objective, models often *ignore visual input completely* in favor of text-text interactions [13]; this issue is magnified when training on videos with ASR, as spoken language often contains redundant words [126].

As a result, a line of recent work has learned independent modality-specific encoders coupled with objectives that can’t

<sup>2</sup>Recent papers propose exceptions, like generating masked-out spans [21] or text [77, 116], but it is unclear whether they can outperform the ‘VisualBERT’s on vision-language tasks like VCR [124].

be shortcutted with simple intra-modal patterns. Models like CLIP learn image classification by matching images with their captions, contrastively [130, 91, 62]. Recent work has explored this paradigm for matching video frames with their transcripts [121], with their audio signal [96, 114], or both [3, 2]; these works likewise perform well on single-modality tasks like audio classification and activity recognition. These independent encoders can be combined through late fusion [96], yet late fusion is strictly less expressive than our proposed joint encoding (early fusion) approach.

**Our work** combines both lines of research. We train a joint model using all modalities in a video – vision, text, and audio – and use a new objective to enable *co-supervision* between modalities. For instance, we can learn a strong model of vision-and-text in part by using *audio as supervision*: when the audio is of ‘MASKed-out’ regions of the video, models cannot easily ‘cheat’ by learning intra-modal vision-vision or text-text interactions (Figure 1). We turn to this next.

### 3. Model: 🎧RESERVE

In this section, we present 🎧RESERVE, including: our model architecture (3.1), new pretraining objectives (3.2), and pretraining video dataset (3.3). At a high level, 🎧RESERVE represents a video by fusing its constituent modalities (vision, audio, and text from transcribed speech) together, and over time. These representations enable both finetuned and zero-shot downstream applications.

More formally, we split a video  $\mathcal{V}$  into a sequence of non-overlapping segments in time  $\{s_t\}$ . Each segment has:

- a. A frame  $v_t$ , from the middle of the segment,
- b. The ASR tokens  $w_t$  spoken during the segment,
- c. The audio  $a_t$  of the segment.

Segments default to 5 seconds in length; we discuss details of how we split videos into segments in Appendix C.

As the text  $w_t$  was automatically transcribed by a model given audio  $a_t$ , it is reasonable to assume that it contains strictly less information content.<sup>3</sup> Thus, for each segment  $s_t$ , we provide models with exactly one of text *or* audio. We will further *mask out* portions of the text and audio during pretraining, to challenge models to recover what is missing.

#### 3.1. Model architecture

An overview of 🎧RESERVE is shown in Figure 2. Each modality is first encoded independently by a separate encoder (a Transformer [110] for images/audio; a BPE lookup table

<sup>3</sup>Despite being derived from the audio, pretraining with text is still paramount: 1) in §3.2 we discuss how jointly modeling audio+text prevents models from shortcutting pretraining objectives via surface correlations; 2) in §4.2 we show that incorporating both transcripts and audio during fine-tuning improves performance; and 3) a textual interface to the model is required for downstream vision+language with textual inputs.

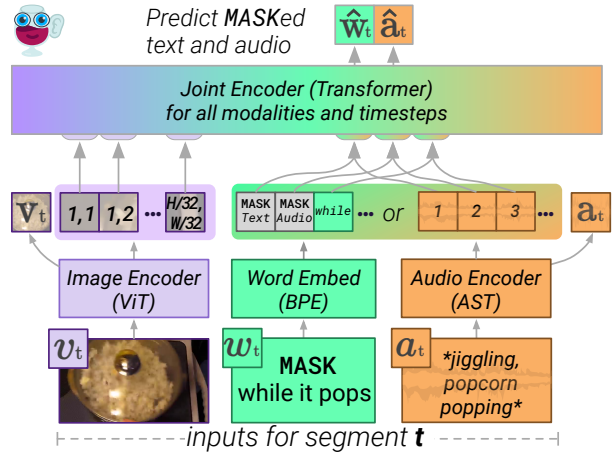


Figure 2: 🎧RESERVE architecture. We provide sequence-level representations of video frames, and *either* words or audio, to a joint encoder. The joint encoder contextualizes over modalities and timesteps, to predict what is behind MASK for audio  $\hat{a}_t$  and text  $\hat{w}_t$ . We supervise these predictions with independently encoded targets:  $a_t$  from the audio encoder, and  $w_t$  from a separate text encoder (not shown).

for text). Then, the representations will be combined with a joint encoder. All encoders are bidirectional.

**Image encoder.** We first encode each frame independently with a Vision Transformer (ViT; [34]), which represents an image of dimensions  $H \times W$  as a sequence of  $16 \times 16$ -pixel patches. We conserve memory in the joint encoder by then performing a  $2 \times 2$  pooling of the feature map, giving us a final  $H/32 \times W/32$  sequence of dimension  $d_h$ .<sup>4</sup>

**Audio encoder.** We split the audio in each segment  $a_t$  into three equal-sized *subsegments*, for compatibility with the lengths at which we mask text (Appendix C). Each subsegment is encoded independently using an Audio Spectrogram Transformer (AST; [46]), which treats each subsegment as an  $M \times T$  image to be ViT-encoded (where  $M=64$  is the Mel filterbank size, and  $T=30$  is time). We treat it as a one-dimensional sequence by using a patch size of  $M \times 2$ . We attention-pool the resulting feature map with a window size of 5, giving a final feature map of  $T/10=6$ .

**Text encoder.** We provide text to our model by BPE-encoding it, embedding the result with a learned table [99].

**Joint encoder.** Finally, we jointly encode all modalities (over all input video segments) using a bidirectional Transformer. We use a linear projection of the final layer’s hidden states for all objectives (e.g.  $\hat{w}_t$  and  $\hat{a}_t$ ).

**Contrastive targets.** Our contrastive pretraining objectives (3.2) require *independently-encoded targets* for each modality.<sup>5</sup> Doing this is straightforward for the image and

<sup>4</sup>We pool through self-attention: the query is the average of the vectors in the  $2 \times 2$  grid; the keys and values are learned projections of the vectors.

<sup>5</sup>Otherwise, a model can trivially cheat by coordinating predictions [6].




Figure 3: Contrastive span training. Given a video with all modalities temporally aligned, we MASK out a region of text and audio. The model must maximize its similarity *only* to an independent encoding of the text  $\mathbf{w}_t$  and audio  $\mathbf{a}_t$ .

audio encoders: we add a CLS to their respective inputs, and extract the final hidden state  $\mathbf{v}_t$  or  $\mathbf{a}_t$  at that position.<sup>6</sup> For text, we learn a separate bidirectional Transformer *span encoder*, which computes targets  $\mathbf{w}_t$  from a CLS and (shared) embedded tokens of a candidate text span. This will also enable zero-shot prediction (4.3).

**Architecture sizes.** We consider two model sizes in this work, which we pretrain from random initialization:

1. 🌐 RESERVE-B, with a hidden size of 768, a 12-layer ViT-B/16 image encoder, and a 12-layer joint encoder.
2. 🌐 RESERVE-L, with a hidden size of 1024, a 24-layer ViT-L/16 image encoder, and a 24-layer joint encoder.

We always use a 12-layer audio encoder, and a 4-layer text span encoder. We pretrain on 192x320 images, and adapt the resolution to 288x512 widescreen images for downstream tasks. To simplify adapting to variable input sizes, we eschew learned position embeddings in favor of a multi-dimensional rotary position encoding [104, 10]. Details and hyperparameters are in Appendix B.

### 3.2. Contrastive Span Training

We introduce *contrastive span* training, which enables learning across and between the three modalities. As shown in Figure 3, the model is given a sequence of video segments. For each one, we include the video frame, and then three ‘subsegments’, each that are either text *or* audio (never both, due to redundancy of text given audio). We train by replacing 25% of these text and audio subsegments with a special MASK token. The model must match the representation atop the MASK *only* with an independent encoding of its span.

Our approach combines past success at matching images to their captions [91, 62] along with ‘VisualBERT’-style prediction of independent tokens [106, 20] – though, crucially,

<sup>6</sup>The CLS hidden states are not involved in the spatial pooling, and likewise are not given to the joint encoder.

we predict representations at a higher-level semantic unit than individual tokens. Our approach also enables the model to learn from both audio and text, while discouraging *memorization* of raw perceptual input, or tokens – which can harm representation quality [112].

Formally, we minimize the cross entropy between the MASKed prediction  $\hat{\mathbf{w}}_t$  and its corresponding phrase representation  $\mathbf{w}_t$ , versus others in the batch  $\mathcal{W}$ :

$$\mathcal{L}_{\text{mask} \rightarrow \text{text}} = \frac{1}{|\mathcal{W}|} \sum_{\mathbf{w}_t \in \mathcal{W}} \left( \log \frac{\exp(\sigma \hat{\mathbf{w}}_t \cdot \mathbf{w}_t)}{\sum_{\mathbf{w} \in \mathcal{W}} \exp(\sigma \hat{\mathbf{w}}_t \cdot \mathbf{w})} \right). \quad (1)$$

We first  $L^2$ -normalize  $\mathbf{w}$  and  $\hat{\mathbf{w}}$ , and scale their dot product with a parameter  $\sigma$  [91].<sup>7</sup> We then add this to its transposed version  $\mathcal{L}_{\text{text} \rightarrow \text{mask}}$ , giving us our text-based loss  $\mathcal{L}_{\text{text}}$ . Analogously, we define  $\mathcal{L}_{\text{audio}}$  for audio, between the MASKed prediction  $\hat{\mathbf{a}}_t$  and its target  $\mathbf{a}_t$ , versus others  $\mathbf{a}$  in the batch.

In addition to these masked text and audio objectives, we simultaneously train the model to match video frames with a contextualized encoding of the transcript. In MERLOT [126], this objective was found to be critical for learning visual recognition from self-supervised videos. Here, the joint encoder encodes the entire video’s transcript at once, extracting a single hidden representation per segment  $\hat{\mathbf{v}}_t$ . We use the same contrastive setup as Equation 1 to maximize the similarity of these vectors with the corresponding  $\mathbf{v}_t$  vectors from the frames, giving us a symmetric frame-based loss  $\mathcal{L}_{\text{frame}}$ . The final loss is the sum of the component losses:

$$\mathcal{L} = \mathcal{L}_{\text{text}} + \mathcal{L}_{\text{audio}} + \mathcal{L}_{\text{frame}}. \quad (2)$$

**Avoiding shortcut learning.** Early on, we observed that training a model to predict a *perceptual* modality (like audio or vision) given input from *the same modality*, led to shortcut learning – a low training loss, but poor representations. We hypothesize that this setup encourages models to learn imperceptible features, like the exact model of the microphone, or the chromatic aberration of the camera lens [33]. We avoid this, while still using audio as a target, by simultaneously training on two kinds of masked videos:

- Audio only as target.** We provide only text and video frames. The model produces representations of both *audio* and *text* that fill in MASKed blanks. It thus learns from audio supervision, without intra-modality shortcuts.
- Audio as input.** We provide the model video frames, and audio (or subtitles) at each segment. Because the model is given audio as an input somewhere, the model only produces representations for MASKed *text*.

There is one more data issue that must be addressed to avoid shortcuts. YouTube’s captions are not perfectly time-aligned with the underlying audio. During our initial

<sup>7</sup>Following past work, we optimize  $\sigma$  and clip it at 100, which enables the model to ‘warm-up’ its emphasis placed on hard negatives [91, 113].

exploration, models took ready advantage of this shortcut: for instance, predicting an audio span based on what adjacent (overlapping) words sound like. We introduce a masking algorithm to resolve this; details in Appendix C.

**Pretraining setup.** We train 🌐RESERVE on TPU v3-512 accelerators; training takes 5 days for 🌐RESERVE-B, and 16 days for 🌐RESERVE-L. We made pretraining more efficient through several algorithmic and implementation improvements, listed in full in Appendix B.5. Of note, we provide the joint encoder at most  $N=8$  frames at a time, and incorporate extra contrastive text negatives from written text.

We use a batch size of 1024 videos, each with  $N=16$  segments. The transcript frame matching task uses all 16384 images from the batch. Likewise, the contrastive span task uses all 49152 audio clips and text candidates in the batch. We use AdamW [68, 79] to minimize Equation 2, details and hyperparameters are in Appendix B.5.

### 3.3. Pretraining Dataset

Recent prior work on static images that demonstrates empirical improvements by increasing dataset size – all the way up to JFT-3B [69, 34, 91, 128]. The same pattern emerges in videos: prior work that has shown promising empirical improvements not only by scaling to 6 million videos/180M frames [126], but also by collecting a diverse set (i.e., going beyond instructional videos [59]).

To this end, we introduce a new training dataset of 20 million English-subtitled YouTube videos, and 1 billion frames, called YT-Temporal-1B. At the same time, we take steps to protect user privacy, directing scraping towards public, large, and monetized channels. We detail our collection, preprocessing, and release strategy in Appendix E.

## 4. Experiments

In this section, we show that a finetuned 🌐RESERVE obtains state-of-the-art results on both VCR [124] as well as TVQA [74], suggesting strong multimodal script knowledge understanding. We then show that our model has strong zero-shot capability, over four challenging zero-shot tasks.

### 4.1. Visual Commonsense Reasoning

We evaluate 🌐RESERVE first through finetuning on VCR [124], a challenging multiple choice task that is commonly used as a benchmark for vision-and-language models. Most competitive models for VCR learn exclusively from images paired with captions, often with supervised visual representations (like from an object detector). To the best of our knowledge, the only exception is MERLOT [126], which uses YouTube video frames and text as part of pretraining; no VCR model to date was pretrained on audio.

**VCR Task.** A model is given an image from a movie, and a question. The model must choose the correct answer

given four multiple choice options ( $Q \rightarrow A$ ); it then is given four *rationales* justifying the answer, and it must choose the correct one ( $QA \rightarrow R$ ). The results are combined with a  $Q \rightarrow AR$  metric, where a model must choose the right answer *and then* the right rationale, to get the question ‘correct.’

**Finetuning approach.** We finetune models using [126]’s approach: ‘drawing on’ VCR’s detection tags directly onto the image, and jointly finetuning on  $Q \rightarrow A$  and  $QA \rightarrow R$ . For both subproblems, we consider four sequences (one per answer or rationale option). Each sequence has pooled image features, the question (and, in  $QA \rightarrow R$ , the answer); and then the option. We pool a hidden representation from a MASK inserted after the text. Using a newly-initialized  $d_h$  vector, we extract a logit for that candidate, and use cross-entropy to optimize that score (details in Appendix D.1.1.)

#### 4.1.1 Ablations: contrastive learning with audio helps.

While we present our final, state-of-the-art VCR performance in 4.1.2, we first use the corpus for an ablation study, empirically justifying design decisions of our pretraining setup. We use the same architecture and data throughout, allowing apples-to-apples comparison between modeling decisions.<sup>8</sup>

**Contrastive Span helps for Vision+Text modeling.** We start by comparing pretraining objectives for learning from YouTube ASR and video alone:

- Mask LM.** This objective trains a bidirectional model by having it *independently* predict masked-out tokens. We make this baseline as strong as possible by using SpanBERT-style masking [63], where text spans are masked out (identical to our *contrastive spans*). Each span  $w$  is replaced by a MASK token. We predict each of its subwords  $w_i$  independently: like [63], we concatenate the MASK’s hidden state with a position embedding for index  $i$ , pass the result through a two-layer MLP, and use tied embedding weights to predict  $w_i$ .
- VirTex [27].** In this objective, we likewise mask text subsegments and extract their hidden states. The difference is that we sequentially predict tokens  $w_i \in w$ , using a left-to-right language model (LM) with the same architecture details as our proposed span encoder. The LM is given the pooled hidden state, and uses tied embeddings to predict new tokens  $w_{i+1}$  conditioned on previous  $w_{<i}$ .

Results are in Table 1. Versus these widely used baselines, our contrastive span objective boosts performance by over 2%, after one epoch of pretraining only on vision and text. We hypothesize that its faster learning is caused by encouraging models to learn concept-level span representations; this might not happen when predicting tokens individually [22].

<sup>8</sup>We use the transcript-frame matching objective throughout, as MERLOT [126]’s ablations demonstrate that it is indispensable when training vision representations from random initialization. Our starting ‘Mask LM’ configuration is similar to MERLOT’s, and we improve from there.

Configuration for one epoch of pretraining	VCR val Q→A (%)
V+T Mask LM [29, 106, 126]	67.2
VirTex-style [27]	67.8
🗣️ Contrastive Span	<b>69.7</b>
V+T+A 🗣️ Audio as target	70.4
🗣️ Audio as input and target	<b>70.7</b>
Audio as input and target, but sans strict localization	70.6
🗣️ RESERVE-B	<b>71.9</b>

Table 1: **Ablation study** of our contrastive span objective. It outperforms prior work even without audio; with an additional 1% boost when audio is used. Our full setup, adding written text, improves another 1%. A 🗣️ denotes a component of our full model.

**Audio pretraining helps**, even for the audio-less VCR:

- d. **Audio as target.** Here, the model is only given video frames and ASR text as input. In addition to performing contrastive-span pretraining over the missing text spans, it does the same over the (held-out) audio span (Equation 2). This boosts VCR accuracy by 0.7%.
- e. **Audio as input and target.** The model does the above (for video+text input sequences), and simultaneously is given video+text+audio sequences, wherein it must predict missing text. This boosts accuracy by 1% in total.
- f. **Sans strict localization.** We evaluate the importance of our strict localization in time. Here, in addition to correct subsegments at the *true* position  $t$  as a correct match, we count adjacent MASKed out regions as well. An extreme version of this was proposed by [48], where a positive match can be of any two frames in a video. Yet even in our conservative implementation, performance drops slightly, suggesting localization helps.

Putting these all together, we find that contrastive span pretraining outperforms mask LM, with improved performance when audio is used **both as input and target**. For our flagship model, we report results in Table 1 on simultaneously training on web-text sequences as well (Appendix C.4), this improves performance by an additional 1%.

#### 4.1.2 VCR Results

Encouraged by these results, we train our models for 10 epochs on YT-Temporal-1B. Figure 4 demonstrates that finetuned VCR performance tracks with the number of pretraining epochs, as well as the validation loss.<sup>9</sup>

<sup>9</sup>The plot suggests that if we pretrained longer, VCR performance might continue to increase. While this could be true, an additional confounding factor is the learning rate decay schedule. With access to compute beyond our current capacity, future work would be well-suited to consider this and other pre-training modifications.




Figure 4: **Pretraining progress:** performance on contrastive-span validation accuracy. Pretraining 🗣️ RESERVE-B for 9 more epochs boosts performance by 5%; L by 8%.

Model	VCR (test)		
	Q→A	QA→R	Q→AR
Caption/ObjDet-based			
ERNIE-ViL-Large [123]	79.2	83.5	66.3
Villa-Large [39]	78.9	83.8	65.7
UNITER-Large [20]	77.3	80.8	62.8
Villa-Base [39]	76.4	79.1	60.6
ViLBERT [80]	73.3	74.6	54.8
B2T2 [4]	72.6	75.7	55.0
VisualBERT [76]	71.6	73.2	52.4
Video-based			
MERLOT [126]	80.6	80.4	65.1
🗣️ RESERVE-B	79.3	78.7	62.6
🗣️ RESERVE-L	<b>84.0</b>	<b>84.9</b>	<b>72.0</b>

Table 2: 🗣️ RESERVE gets **state-of-the-art leaderboard performance on VCR**. We compare it with the largest submitted single models, including image-caption models that utilize heavy manual supervision (e.g. object detections and captions).

Finally, in Table 2, we compare 🗣️ RESERVE against the largest published models from the VCR leaderboard. Of note, 🗣️ RESERVE-L outperforms all prior work, by **over 5%** on Q→AR metric. It outperforms even large ensembles (e.g. 15 ERNIE-Large’s) submitted by industry [123], though we do not show these on this table to focus on only single models.

**Efficiency.** The accuracy increase of 🗣️ RESERVE is not simply due to compute.<sup>10</sup> In fact, our 🗣️ RESERVE-L requires *one-fifth the FLOPs* of detector-based systems, like UNITER-Large [20] (Appendix B.3). Moreover, because 🗣️ RESERVE-L uses a pure ViT backbone versus MERLOT’s ViT-ResNet hybrid, it uses fewer FLOPs than MERLOT, while scoring 7% higher. Meanwhile, 🗣️ RESERVE-B outperforms ‘base’ detector-based models, while using *less than one-tenth their FLOPs*.

In terms of parameter count, 🗣️ RESERVE-B is comparable to prior work. On VCR, including the vision stack, 🗣️ RESERVE-B has 200M finetunable parameters and performs similarly to the 378M parameter UNITER-Large. 🗣️ RESERVE-L has 644M parameters.

#### 4.2. Finetuning on TVQA

We use TVQA [74] to evaluate our model’s capacity to transfer to video understanding tasks when finetuned. In TVQA, models are given a video, a question, and five answer choices. The scenes come from American TV shows, and depict characters interacting with each other through dialogue – which past work represents through subtitles. We will show that providing audio to 🗣️ RESERVE improves performance.

**Audio-Subtitle Finetuning.** To evaluate how much audio can help for TVQA, we finetune 🗣️ RESERVE jointly between

<sup>10</sup>Here, we use FLOPs as our key efficiency metric, as they are a critical bottleneck in model scaling [65, 34, 128]. On the other hand, we argue that parameter count can be misleading – for instance, many Transformer parameters can be tied together with minimal performance loss [71].

Model	TVQA (acc; %)	
	Val	Test
Human [74]	–	89.4
MERLOT [126]	78.7	78.4
MMFT-BERT [109]	73.5	72.8
Kim et al [67]	76.2	76.1
RESERVE-B (subtitles)	81.3	–
RESERVE-L (subtitles)	<b>85.6</b>	<b>84.8</b>
RESERVE-B (audio)	82.5	–
RESERVE-L (audio)	<b>85.9</b>	<b>85.6</b>
RESERVE-B	83.1	82.7
RESERVE-L	<b>86.5</b>	<b>86.1</b>

Table 3: RESERVE gets state-of-the-art results on TVQA by a margin of **over 7%**, over methods from prior work – that do not make use of audio information.

the ‘vision and subtitles’ and ‘vision and audio’ settings. Like on VCR, we consider one sequence per candidate: each contains video frame features, the question, the answer candidate, and a MASK token (from where we pool a hidden representation). During training, each sequence is duplicated: we provide one sequence with *subtitles* from the video, and for the other, we use *audio*. This lets us train a single model, and then test how it will do *given subtitles*, *given audio*, or *given both* (by averaging the two softmax predictions)

**Results.** We show TVQA results in Table 3. Without audio so as to enable fair comparison, our RESERVE-B outperforms all prior work with a margin of 2.6%. Audio performs even better than subtitles. And, combining subtitle-only and audio-only predictions performs best of all, improving over 4% versus the prior state-of-the-art, MERLOT (and in turn over other models). The same pattern holds (with additional performance gains) as model size increases: RESERVE-L improves over prior work by **7.6%**.

### 4.3. Zero-Shot Experiments

Next, we show that our model exhibits strong zero-shot performance for a variety of downstream tasks. Our zero-shot interface is enabled by our *contrastive span objective*. For QA tasks that require predicting an option from a label space of short phrases, we encode this label space as vectors, and predict the closest phrase to a MASKED input. We consider:

- i. **Situated Reasoning (STAR)** [119]. This task requires the model to reason over short situations in videos, covering four axes: interaction, sequence, prediction, and feasibility. The model is given a video, a templated question, and 4 answer choices. We convert templated questions into literal statements (which are more similar to YouTube dialogue); the label space is the set of four options.
- ii. **Action Anticipation in Epic Kitchens** [26]. Here, the goal is to predict *future actions* given a video clip, which requires reasoning temporally over an actor’s motivations

Model	Situated Reasoning (STAR)					EPIC-Kitchens			LSMDC	MSR-VTT QA	
	Interaction	Sequence	Prediction	Feasibility	Overall	(val class-mean R@5; %)			(F1B test %)	(test acc %)	
Supervised SoTA											
			ClipBERT [73]							MERLOT [126]	
Random	25.0	25.0	25.0	25.0	25.0	6.2	2.3	0.1	0.1	0.1	0.5
CLIP (ViT-B/16) [91]	39.8	40.5	35.5	36.0	38.0	16.5	12.8	2.3	2.0	3.0	11.9
CLIP (RN50x16) [91]	39.9	41.7	36.5	<b>37.0</b>	38.7	13.4	14.5	2.1	2.3	2.3	9.7
Just Ask (ZS) [122]										2.9	8.8
RESERVE-B	44.4	40.1	38.1	35.0	39.4	17.9	15.6	2.7	26.1	3.7	10.8
RESERVE-L	42.6	41.1	37.4	32.2	38.3	15.6	19.3	4.5	26.7	4.4	11.5
RESERVE-B (+audio)	<b>44.8</b>	<b>42.4</b>	<b>38.8</b>	36.2	<b>40.5</b>	20.9	17.5	3.7	29.1	4.0	12.0
RESERVE-L (+audio)	43.9	42.6	37.6	33.6	39.4	<b>23.2</b>	<b>23.7</b>	<b>4.8</b>	<b>31.0</b>	<b>5.8</b>	<b>13.6</b>

Table 4: Zero shot results. On STAR, RESERVE obtains state-of-the-art results, outperforming finetuned video models. It performs well on EPIC-Kitchens (verb and noun forecasting), along with LSMDC, despite their long-tail distributions. On MSR-VTT QA, it outperforms past work on weakly-supervised video QA. Further, it outperforms CLIP (that cannot handle dynamic situations), and benefits from audio when given.

and intentions. The dataset has a long tail of rare action combinations, making zero-shot inference challenging (since we do not assume access to this prior). As such, prior work [45, 38] trains on the provided in-domain training set. To adapt RESERVE to this task, we provide it a single MASK token as text input, and use as our label space of all combinations of verbs and nouns in the vocabulary (e.g. ‘cook apple, cook avocado’, etc.).

- iii. **LSMDC** [81, 95]. Models are given a video clip, along with a video description (with a MASK to be filled in). We compare it with the vocabulary used in prior work [126].
- iv. **MSR-VTT QA** [120]. This is an open-ended video QA task about what is literally happening in a web video. We use GPT3 [16], prompted with a dozen (unlabelled) questions, to reword the questions into statements with MASKS. This introduces some errors, but minimizes domain shift. We use [122]’s label space, with 1k options.

For these tasks, we use  $N=8$  video segments (dilating time when appropriate), and provide audio input when possible. Details and prompts are in Appendix D. We compare against both finetuned and zeroshot models, including running CLIP [91] on all tasks. CLIP is a strong model for zero-shot classification, particularly when *encyclopedic knowledge about images* is helpful; our comparisons showcase where multimodal script knowledge helps.

**Results.** Table 4 shows our model performs competitively:

- i. On STAR, it obtains state-of-the-art results, with performance gain when audio is included. Interestingly, RESERVE-B outperforms its larger variant; we hypothesize that this is due to limited prompt searching around question templates. We qualitatively observed that RESERVE-L sometimes excludes topically correct options if they sound grammatically strange (to it).
- ii. On EPIC-Kitchens, our model obtains strong results at correctly anticipating the verb and noun - despite the

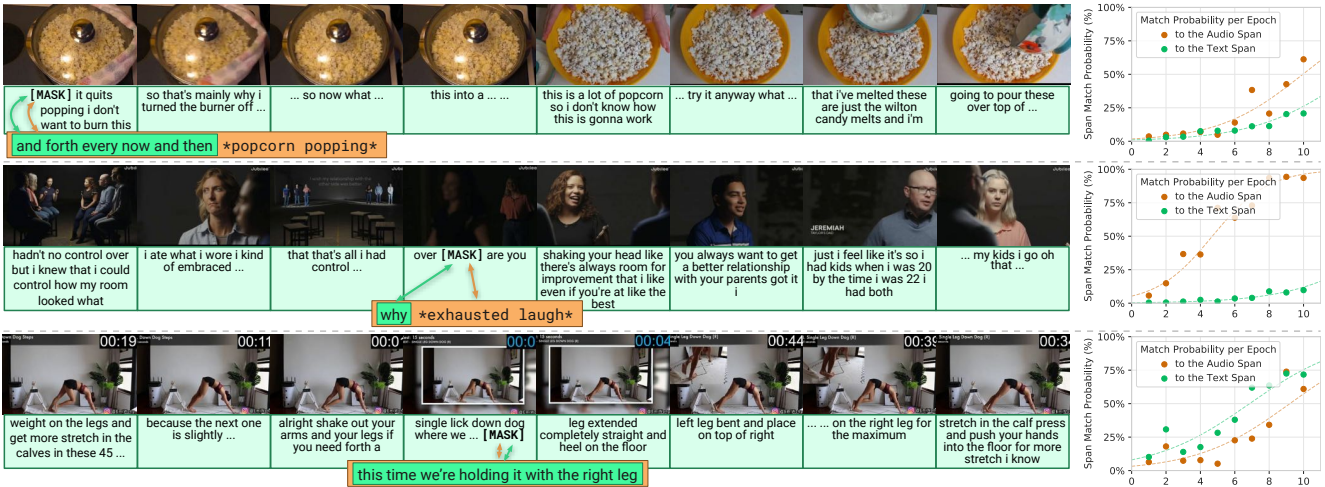


Figure 5: **Exploring MASKed audio self-supervision.** Shown are example videos from our validation set, with predictions from 🧠RESERVE-B. During pretraining, our model progressively learns to pick up on audio-specific clues. It seems to recognize physical dynamics of *cooking popcorn*, matching the first row to its MASKed audio. Likewise, it seems to use social reasoning to match the second row to its audio. Both of these clues are orthogonal to subtitles provide.

heavy-tailed nature of both distributions. It is worse on getting both right (‘action’), we suspect that this might be due to priors (motifs) between noun and verb [127]. These are easy to learn given access to training data, but we exclude these as we consider the zero-shot task.

- iii. On LSMDC, our model obtains strong results at filling-in-the-blank, likewise despite a heavy (unseen) frequency bias. Notably, it outperforms CLIP significantly, with CLIP often preferring templates that use visually-relevant words, even if they don’t make sense as a whole.<sup>11</sup>
- iv. Finally, our model performs well on MSR-VTT QA, outperforming past work that directly rewords subtitled instructional videos into video QA instances [122].

## 5. Qualitative Analysis: Why does audio help?

What can 🧠RESERVE learn from both text *and* audio? Three validation set examples are shown in Figure 5. The model is given the displayed text and video frames, and must match the MASK to the correct missing text and audio span (out of 48k total in the batch). The plots show 🧠RESERVE-B’s probability of correctly identifying the correct audio or text span, as it progresses through 10 epochs of pretraining.

**Audio’s supervisory signal.** In the first two rows of Figure 5, audio provides orthogonal supervision to text:

1. The first row is from the recurring ‘popcorn’ example, and the MASKed audio contains the distinctive sound of popcorn pops slowing, as most kernels have already finished. By the final epoch, 🧠RESERVE-B selects this specific auditory cue with 60% probability, over others

(including from adjacent frames, with different distinct stages of popping). Here, sound provides signal for learning joint vision-text understanding of the situation. This cannot be captured by predicting the MASKed text ‘and forth every now and then,’ as evidenced by the lower match-probability.

2. The second row contains only the text ‘why,’ with the audio providing greatly more information — a female-presenting speaker (shown in the next frame) laughs, astonished that the child (in the frame afterwards) might want a better relationship with their parents.
3. In the third row, matching performance is similar between modalities, possibly as the yogi is narrating over a (muted) video recording, and not adding much information.

**Role of text.** Text is still a crucial complement to audio, in terms of the supervision it provides. Consider the second row: 🧠RESERVE-B learns to match the audio almost perfectly (perhaps reasoning that the speaker is shown in the next frame, and is laughing). In later epochs, its text-match probability increases: knowing that a ‘why’ question is likely to be asked is a valid *social* inference to make about this (tense) situation.

**Learning through multimodal reentry.** In this work, we proposed a *reentrant* learning paradigm: predicting a modality like audio given fused input from vision and text. This paradigm contrasts with today’s standard practice of learning primarily through language (either from language alone, or from images paired with captions).

Indeed, we argue that learning from audio *through reentry* goes beyond offering orthogonal signal to models. Our pretraining task forces 🧠RESERVE to abstract away from intramodal perceptual features (like pixels or tokens), towards learning a more complete understanding of the world

<sup>11</sup>For instance, on a picture of a smiling mailman ready to leave, CLIP chooses ‘the mailman smiles off,’ versus ‘the mailman takes off.’



as viewed from the held-out modality. Developmental psychologists have hypothesized that human children adopt such a learning strategy, learning connections not just between vision, language, and audio but also touch, smell, and taste [35, 100]. Our work suggests it has potential for machines as well.

## 6. Conclusion, Limitations, Broader Impact

We introduced 🗣️ RESERVE, which learns representations jointly through sound, language, and vision, guided through a new pretraining objective. Our model performs well in both finetuned and zero-shot settings, yet there are still limitations:

- a. Our model only learns to reason over 40-second long video chunks at once (8 ‘segments’).
- b. Our setup learns only from videos with English ASR subtitles. It would not work for the majority of the world’s languages, where ASR systems today do worse.
- c. Our model can only *match*, not *generate*, text and audio.

Still, we foresee several potential broader societal impacts of our work. An expanded version of this technology might someday serve as a component of assistive technology for improved video captioning, e.g., to improve accessibility for low vision and blind users [75] or for people who are d/Deaf [47]. Yet, the same technology can have impacts that we authors consider to be negative, including surveillance, or applications that hegemonize social biases that are widespread on YouTube. We discuss these further in Appendix A: key dimensions include respecting user privacy during dataset collection, exploring biases in YouTube data, dual use, and energy consumption. We discuss our plan to *release our model and data* for research use so others can critically study this approach to learning script knowledge.

## Acknowledgements

We thank Jae Sung Park, Oren Etzioni, Gabriel Ilharco, and Mitchell Wortsman for feedback on this work. Thanks also to Zak Stone and the Google Cloud TPU team for providing access to the TPU machines used for conducting experiments. Thanks to James Bradbury and Skye Wanderman-Milne for help with JAX on TPUs. Thanks to the AI2 ReVIZ team, including Jon Borchardt and M Kusold, for help with the demo. This work was funded by DARPA MCS program through NIWC Pacific (N66001-19-2-4031), and the Allen Institute for AI. Last, but not least, thanks to the YouTubers whose work and creativity helps machines to learn about the multimodal world.

## References

[1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. Youtube-8m: A large-scale video

classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016. 15

[2] Hassan Akbari, Linagzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. VATT: transformers for multimodal self-supervised learning from raw video, audio and text. *arXiv preprint arXiv:2104.11178*, 2021. 3

[3] Jean-Baptiste Alayrac, Adrià Recasens, Rosalia Schneider, Relja Arandjelović, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks. *arXiv preprint arXiv:2006.16228*, 2020. 3

[4] Chris Alberti, Jeffrey Ling, Michael Collins, and David Reitter. Fusion of detected objects in text for visual question answering. *arXiv preprint arXiv:1908.05054*, 2019. 6

[5] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018. 19, 20

[6] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *arXiv preprint arXiv:1906.00910*, 2019. 3

[7] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443, 2018. 2

[8] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623, 2021. 15, 17

[9] Emily M. Bender and Alexander Koller. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online, July 2020. Association for Computational Linguistics. 17

[10] Stella Biderman, Sid Black, Charles Foster, Leo Gao, Eric Hallahan, Horace He, Ben Wang, and Phil Wang. Rotary embeddings: A relative revolution, 2021. [Online; accessed ]. 4, 18

[11] Abeba Birhane, Vinay Uday Prabhu, and Emmanuel Kahembwe. Multimodal datasets: misogyny, pornography, and malignant stereotypes. *arXiv preprint arXiv:2110.01963*, 2021. 17

[12] Sophie Bishop. Anxiety, panic and self-optimization: Inequalities and the youtube algorithm. *Convergence*, 24(1):69–84, 2018. 17

[13] Yonatan Bitton, Gabriel Stanovsky, Michael Elhadad, and Roy Schwartz. Data efficient masked language modeling for vision and language. *arXiv preprint arXiv:2109.02040*, 2021. 2

[14] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. 15

[15] Luke Breitfeller, Emily Ahn, David Jurgens, and Yulia Tsvetkov. Finding microaggressions in the wild: A case

- for locating elusive phenomena in social media posts. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1664–1674, 2019. 17
- [16] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020. 7, 15, 17
- [17] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 67–74. IEEE, 2018. 15, 16
- [18] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. *arXiv preprint arXiv:2012.07805*, 2020. 15
- [19] Robin S Chapman. Children’s language learning: An interactionist perspective. *The Journal of Child Psychology and Psychiatry and Allied Disciplines*, 41(1):33–54, 2000. 1
- [20] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020. 2, 4, 6, 18, 19
- [21] Jaemin Cho, Jie Lei, Hao Tan, and Mohit Bansal. Unifying vision-and-language tasks via text generation. In *ICML*, 2021. 2
- [22] Kyunghyun Cho. Tweet. "important dependences between the image features and words/phrases in the description could be explained away by the dependencies among words/phrases". 5
- [23] Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. Don’t take the easy way out: Ensemble based methods for avoiding known dataset biases. In *EMNLP*, pages 4069–4082, Hong Kong, China, Nov. 2019. Association for Computational Linguistics. 2
- [24] Matthew Crain. The limits of transparency: Data brokers and commodification. *new media & society*, 20(1):88–104, 2018. 15
- [25] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018. 25
- [26] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision (IJCV)*, 2021. 2, 7, 25
- [27] Karan Desai and Justin Johnson. Virtex: Learning visual representations from textual annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11162–11173, 2021. 5, 6
- [28] Sunipa Dev, Masoud Monajatipoor, Anaelia Ovalle, Arjun Subramonian, Jeff M Phillips, and Kai-Wei Chang. Harms of gender exclusivity and challenges in non-binary representation in language technologies. *arXiv preprint arXiv:2108.12084*, 2021. 17
- [29] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2, 6
- [30] Travis L Dixon. Crime news and racialized beliefs: Understanding the relationship between local news viewing and perceptions of african americans and crime. *Journal of Communication*, 58(1):106–125, 2008. 27
- [31] Travis L Dixon and Daniel Linz. Overrepresentation and underrepresentation of african americans and latinos as lawbreakers on television news. *Journal of communication*, 50(2):131–154, 2000. 27
- [32] Jesse Dodge, Maarten Sap, Ana Marasovic, William Agnew, Gabriel Ilharco, Dirk Groeneveld, and Matt Gardner. Documenting the english colossal clean crawled corpus. *CoRR*, abs/2104.08758, 2021. 17
- [33] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015. 4
- [34] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3, 5, 6, 20
- [35] Gerald M Edelman. Neural darwinism: selection and reentrant signaling in higher brain function. *Neuron*, 10(2):115–125, 1993. 1, 9
- [36] David F Fouhey, Wei-cheng Kuo, Alexei A Efros, and Jitendra Malik. From lifestyle vlogs to everyday interactions. In *CVPR*, 2018. 28
- [37] Christian Fuchs. An alternative view of privacy on facebook. *Information*, 2(1):140–165, 2011. 15
- [38] Antonino Furnari and Giovanni Maria Farinella. Rolling-unrolling lstms for action anticipation from first-person video. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2020. 7, 25, 26
- [39] Zhe Gan, Yen-Chun Chen, Linjie Li, Chen Zhu, Yu Cheng, and Jingjing Liu. Large-scale adversarial training for vision-and-language representation learning. *arXiv preprint arXiv:2006.06195*, 2020. 6
- [40] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020. 24
- [41] Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith. Realtotoxicityprompts: Evaluating neural toxic degeneration in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, 2020. 17

- [42] Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *Proc. IEEE ICASSP 2017*, New Orleans, LA, 2017. 29, 30
- [43] Tarleton Gillespie. Content moderation, ai, and the question of scale. *Big Data & Society*, 7(2):2053951720943234, 2020. 17
- [44] Franklin D Gilliam Jr, Shanto Iyengar, Adam Simon, and Oliver Wright. Crime in black and white: The violent, scary world of local news. *Harvard International Journal of press/politics*, 1(3):6–23, 1996. 27
- [45] Rohit Girdhar and Kristen Grauman. Anticipative video transformer. *arXiv preprint arXiv:2106.02036*, 2021. 7, 25, 26
- [46] Yuan Gong, Yu-An Chung, and James Glass. Ast: Audio spectrogram transformer. *arXiv preprint arXiv:2104.01778*, 2021. 3, 19
- [47] Steven M Goodman, Ping Liu, Dhruv Jain, Emma J McDonnell, Jon E Froehlich, and Leah Findlater. Toward user-driven sound recognizer personalization with people who are d/deaf or hard of hearing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(2):1–23, 2021. 9
- [48] Daniel Gordon, Kiana Ehsani, Dieter Fox, and Ali Farhadi. Watching the world go by: Representation learning from unlabeled videos. *arXiv preprint arXiv:2003.07990*, 2020. 6
- [49] Jonathan Gordon and Benjamin Van Durme. Reporting bias and knowledge acquisition. In *Proceedings of the 2013 workshop on Automated knowledge base construction*, pages 25–30. ACM, 2013. 17
- [50] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in vqa matter: Elevating the role of image understanding in visual question answering. In *CVPR*, volume 1, page 9, 2017. 2
- [51] Ben Green. Good” isn’t good enough. In *Proceedings of the AI for Social Good workshop at NeurIPS*, 2019. 17
- [52] Daniel Griffin and Jae Lim. Signal estimation from modified short-time fourier transform. *IEEE Transactions on acoustics, speech, and signal processing*, 32(2):236–243, 1984. 19
- [53] Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel. Audioclip: Extending clip to image, text and audio. *arXiv preprint arXiv:2106.13043*, 2021. 2, 29, 30
- [54] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M Branham. Gender recognition or gender reductionism? the social implications of embedded gender recognition systems. In *Proceedings of the 2018 chi conference on human factors in computing systems*, pages 1–13, 2018. 17
- [55] Donna Haraway. Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist studies*, 14(3):575–599, 1988. 17
- [56] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 19
- [57] Don Heider. *White news: Why local news programs don’t cover people of color*. Routledge, 2014. 27
- [58] Jack Hessel and Lillian Lee. Does my multimodal model learn cross-modal interactions? it’s harder to tell than you might think! In *EMNLP*, 2020. 2
- [59] Jack Hessel, Bo Pang, Zhenhai Zhu, and Radu Soricut. A case study on combining ASR and visual features for generating instructional video captions. In *CoNLL*, Nov. 2019. 5
- [60] Dirk Hovy and Shrimai Prabhumoye. Five sources of bias in natural language processing. *Language and Linguistics Compass*, 15(8):e12432, 2021. 15
- [61] Gabriel Ilharco, Yuan Zhang, and Jason Baldridge. Large-scale representation learning from visually grounded untranscribed speech. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 55–65, 2019. 2
- [62] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. *arXiv preprint arXiv:2102.05918*, 2021. 3, 4
- [63] Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77, 2020. 5, 21
- [64] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. “my data just goes everywhere:” user mental models of the internet and implications for privacy and security. In *Eleventh Symposium On Usable Privacy and Security (SOUPS) 2015*, pages 39–52, 2015. 15, 27
- [65] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. 6, 20
- [66] Os Keyes, Zoë Hitzig, and Mwenza Blell. Truth from the machine: artificial intelligence and the materialization of identity. *Interdisciplinary Science Reviews*, 46(1-2):158–175, 2021. 17
- [67] Seonhoon Kim, Seohyeong Jeong, Eunbyul Kim, Inho Kang, and Nojun Kwak. Self-supervised pre-training and contrastive representation learning for multiple-choice video qa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13171–13179, 2021. 7
- [68] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. 5
- [69] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 491–507. Springer, 2020. 5
- [70] Jatin Lamba, Jayaprakash Akula, Rishabh Dabral, Preethi Jyothi, Ganesh Ramakrishnan, et al. Cross-modal learning for audio-visual video parsing. *arXiv preprint arXiv:2104.04598*, 2021. 2
- [71] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite

- bert for self-supervised learning of language representations. In *International Conference on Learning Representations*, 2019. 6, 20
- [72] Cheolhyoung Lee, Kyunghyun Cho, and Wanmo Kang. Mixout: Effective regularization to finetune large-scale pre-trained language models. In *International Conference on Learning Representations*, 2019. 25
- [73] Jie Lei, Linjie Li, Luwei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7331–7341, 2021. 2, 7
- [74] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L Berg. Tvqa: Localized, compositional video question answering. In *EMNLP*, 2018. 2, 5, 6, 7
- [75] Marco Leo, G Medioni, M Trivedi, Takeo Kanade, and Giovanni Maria Farinella. Computer vision for assistive technologies. *Computer Vision and Image Understanding*, 154:1–15, 2017. 9
- [76] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 2, 6
- [77] Xudong Lin, Gedas Bertasius, Jue Wang, Shih-Fu Chang, Devi Parikh, and Lorenzo Torresani. Vx2text: End-to-end learning of video-based text generation from multimodal inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7005–7015, 2021. 2
- [78] Jing Liu, Xinxin Zhu, Fei Liu, Longteng Guo, Zijia Zhao, Mingzhen Sun, Weining Wang, Jinqiao Wang, and Hanqing Lu. Opt: Omni-perception pre-trainer for cross-modal understanding and generation. *arXiv preprint arXiv:2107.00249*, 2021. 2
- [79] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [80] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViL-BERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, pages 13–23, 2019. 2, 6
- [81] Tegan Maharaj, Nicolas Ballas, Anna Rohrbach, Aaron C Courville, and Christopher Joseph Pal. A dataset and exploration of models for understanding video data through fill-in-the-blank question-answering. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 7, 26, 27
- [82] Alice E Marwick and danah boyd. Networked privacy: How teenagers negotiate context in social media. *New media & society*, 16(7):1051–1067, 2014. 15
- [83] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *ICCV*, 2019. 15
- [84] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *ICCV*, 2019. 28
- [85] Heather Molyneaux, Susan O’Donnell, Kerri Gibson, Janice Singer, et al. Exploring the gender divide on youtube: An analysis of the creation and reception of vlogs. *American Communication Journal*, 10(2):1–14, 2008. 17
- [86] Arsha Nagrani, Joon Son Chung, and Andrew Zisserman. Voxceleb: a large-scale speaker identification dataset. *arXiv preprint arXiv:1706.08612*, 2017. 15, 16, 29
- [87] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*, 2021. 18
- [88] Jean Piaget and Margaret Trans Cook. The origins of intelligence in children. 1952. 1
- [89] Karol J. Piczak. ESC: Dataset for Environmental Sound Classification. In *Proceedings of the 23rd Annual ACM Conference on Multimedia*, pages 1015–1018. ACM Press, 2011. 29
- [90] Yael Pritch, Sarit Ratovitch, Avishai Hendel, and Shmuel Peleg. Clustered synopsis of surveillance video. In *2009 Sixth IEEE international conference on advanced video and signal based surveillance*, pages 195–200. IEEE, 2009. 17
- [91] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021. 2, 3, 4, 5, 7, 15, 16, 17, 24, 25, 26, 29, 30
- [92] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21:1–67, 2020. 15, 17, 21, 28
- [93] Micah Rajunov and A Scott Duane. *Nonbinary: Memoirs of Gender and Identity*. Columbia University Press, 2019. 17
- [94] Manoel Horta Ribeiro, Raphael Ottoni, Robert West, Virgílio AF Almeida, and Wagner Meira Jr. Auditing radicalization pathways on youtube. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 131–141, 2020. 17
- [95] Anna Rohrbach, Atousa Torabi, Marcus Rohrbach, Niket Tandon, Chris Pal, Hugo Larochelle, Aaron Courville, and Bernt Schiele. Movie description. *International Journal of Computer Vision*, 2017. 2, 7, 26
- [96] Andrew Rouditchenko, Angie Boggust, David Harwath, Brian Chen, Dhiraj Joshi, Samuel Thomas, Kartik Audhkhasi, Hilde Kuehne, Rameswar Panda, Rogerio Feris, et al. Avinet: Learning audio-visual language representations from instructional videos. *arXiv preprint arXiv:2006.09199*, 2020. 3
- [97] Justin Salamon, Christopher Jacoby, and Juan Pablo Bello. A dataset and taxonomy for urban sound research. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 1041–1044, 2014. 29
- [98] Roger C. Schank and Robert P. Abelson. Scripts, plans, and knowledge. In *Proceedings of the 4th International Joint*

- Conference on Artificial Intelligence - Volume 1, IJCAI'75*, pages 151–157, San Francisco, CA, USA, 1975. Morgan Kaufmann Publishers Inc. [1](#)
- [99] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, 2016. [3](#)
- [100] Linda Smith and Michael Gasser. The development of embodied cognition: Six lessons from babies. *Artificial life*, 11(1-2):13–29, 2005. [1](#), [9](#)
- [101] Nick Srnicek. *Platform capitalism*. John Wiley & Sons, 2017. [17](#)
- [102] Michael Strangelove. *Watching YouTube*. University of Toronto press, 2020. [15](#), [17](#)
- [103] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in nlp. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, 2019. [18](#)
- [104] Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*, 2021. [4](#), [18](#)
- [105] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. VideoBERT: A joint model for video and language representation learning. In *ICCV*, 2019. [2](#)
- [106] Hao Tan and Mohit Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *EMNLP*, 2019. [2](#), [4](#), [6](#), [18](#)
- [107] Rachael Tatman. Gender and dialect bias in youtube’s automatic captions. *EACL 2017*, page 53, 2017. [17](#)
- [108] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Herve Jegou. Fixing the train-test resolution discrepancy. *Advances in Neural Information Processing Systems*, 32:8252–8262, 2019. [19](#)
- [109] Aisha Urooj, Amir Mazaheri, Mubarak Shah, et al. Mmftbert: Multimodal fusion transformer with bert encodings for visual question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 4648–4660, 2020. [7](#)
- [110] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [3](#)
- [111] Aurélie Villard, Alan Lelah, and Daniel Brissaud. Drawing a chip environmental profile: environmental indicators for the semiconductor industry. *Journal of Cleaner Production*, 86:98–109, 2015. [18](#)
- [112] Jacob Walker, Carl Doersch, Abhinav Gupta, and Martial Hebert. An uncertain future: Forecasting from static images using variational autoencoders. In *European Conference on Computer Vision*, pages 835–851. Springer, 2016. [4](#)
- [113] Feng Wang and Huaping Liu. Understanding the behaviour of contrastive loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2495–2504, 2021. [4](#)
- [114] Luyu Wang, Pauline Luc, Adria Recasens, Jean-Baptiste Alayrac, and Aaron van den Oord. Multimodal self-supervised learning of general audio representations. *arXiv preprint arXiv:2104.12807*, 2021. [3](#)
- [115] Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, et al. Tacotron: Towards end-to-end speech synthesis. *arXiv preprint arXiv:1703.10135*, 2017. [19](#)
- [116] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint arXiv:2108.10904*, 2021. [2](#)
- [117] Zeerak Waseem, Smarika Lulz, Joachim Bingel, and Isabelle Augenstein. Disembodied machine learning: On the illusion of objectivity in nlp. *arXiv preprint arXiv:2101.11974*, 2021. [17](#)
- [118] Georg Wiese, Dirk Weissenborn, and Mariana Neves. Neural domain adaptation for biomedical question answering. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 281–289, 2017. [25](#)
- [119] Bo Wu, Shoubin Yu, Zhenfang Chen, Joshua B. Tenenbaum, and Chuang Gan. Star: A benchmark for situated reasoning in real-world videos. *2021 Conference on Neural Information Processing Systems*, 2021. [2](#), [7](#), [26](#)
- [120] Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video question answering via gradually refined attention over appearance and motion. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1645–1653, 2017. [2](#), [7](#), [27](#)
- [121] Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, and Florian Metze Luke Zettlemoyer Christoph Feichtenhofer. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *arXiv preprint arXiv:2109.14084*, 2021. [3](#)
- [122] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just ask: Learning to answer questions from millions of narrated videos. *arXiv preprint arXiv:2012.00451*, 2020. [2](#), [7](#), [8](#), [27](#)
- [123] Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. Ernie-vil: Knowledge enhanced vision-language representations through scene graph. *arXiv preprint arXiv:2006.16934*, 2020. [2](#), [6](#)
- [124] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6720–6731, 2019. [2](#), [5](#), [19](#)
- [125] Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. Defending against neural fake news. In *Advances in Neural Information Processing Systems 32*, 2019. [15](#), [17](#)
- [126] Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. Merlot: Multimodal neural script knowledge models. *arXiv preprint arXiv:2106.02636*, 2021. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [15](#), [17](#), [18](#), [19](#), [21](#), [25](#), [27](#), [28](#), [29](#)
- [127] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In

*Conference on Computer Vision and Pattern Recognition*, 2018. 8

- [128] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers, 2021. 5, 6, 18
- [129] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. *arXiv preprint arXiv:2101.00529*, 2021. 18
- [130] Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D Manning, and Curtis P Langlotz. Contrastive learning of medical visual representations from paired images and text. *arXiv preprint arXiv:2010.00747*, 2020. 3
- [131] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2989, 2017. 17
- [132] Shoshana Zuboff. Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30(1):75–89, 2015. 17

## Abstract

We provide the following materials in the appendix:

- A full broader impact statement (Section A)
- Details about our model architecture (Section B)
- Details about how we provide video data into the model, including how we align the modalities and perform the masking (Section C)
- Details about how we adapted our model to downstream tasks (Section D)
- Details about how we collected data (Section E)
- Additional experiments (Section F)

## A. Broader Impact Statement

In this paper, we have presented a model for learning multimodal neural script knowledge, through incorporation of audio as a first-class citizen alongside text and video frames. We argue that academic study of this learning paradigm is important, in part because it relates to how we as humans understand the world. We as humans process situations by perceiving through multiple modalities and interpreting the result holistically.

At the same time, the work and methodology that we outlined risks dual use. Like other large machine learning systems pretrained on web data, our system may reproduce harmful social biases present in its training data. While a variety of past work has studied risks of *language-only* pretraining [125, 14, 8, 60], the video-centric pretraining that we explore in our work might have different benefits and risks. We discuss these below, along with how we worked to mitigate them through our work.

### A.1. Privacy.

A significant risk with training on data at YouTube scale is protecting user privacy. We took several proactive steps to ensure this, that in turn build off prior work and community norms [1, 83, 126]:

- a. We plan to release only video IDs for download, following prior work [1, 83]. Thus, when a user deletes a video off of YouTube, it becomes removed from YT-Temporal-1B as well, giving content creators a right to opt out of all uses of their videos.
- b. Building off of past work [126], we directed our data collection towards *public* and *monetized channels*. These channels are identifiable insofar as they contain more subscribers, and more videos. They include companies that have official accounts, including journalism outlets like the *New York Times* and *Vox*. They also include individuals for whom making public YouTube videos is their full time job. In either case, our use videos in question for research purposes can be seen as *fair use*.

Model	Accuracy (%)		
	Voice	Image+Voice	Image
🗨️ RESERVE-L	10.8	9.6	10.7
CLIP ViT-B/16 [91]			<b>86.0</b>

Table 5: Zero-shot person (face/voice) recognition accuracy on VoxCeleb2 [86] and VGGFace2 [17], using different modalities. While 🗨️ RESERVE can perform person recognition from several modalities, its performance is much lower than the recognition-optimized CLIP model in the image-to-name setting. We hypothesize that this is due to a similarity between this setting and CLIP’s pretraining data – news articles often include celebrity images, paired with their names.

**Framing of privacy.** Privacy is a nuanced topic with many societally, culturally, and generationally-specific interpretations. We took inspiration from Marwick and Boyd [82]’s framework of *networked privacy*, which posits that users posting public videos might *encode* private information – enough so that their intended viewership (friends, possibly) can catch the gist, but not enough so as to leak private details like phone numbers to the world.

Through the lens of networked privacy, we see key differences between studying videos on a moderated platform, versus NLP work that trains models from the open web (e.g. [92, 16]). When YouTube users upload videos, they seem to understand details of its privacy policy, beyond consenting to it [64]. YouTubers typically upload their own videos [102]. These details differ from text on the open web. Today, ‘data brokers’ post private details (like phone numbers) to the web for profit [24]; concerningly, a study on language models suggests that models are vulnerable at *memorizing* this private information [18].

It is worth examining our research through other framings of privacy as well. For example, internet platforms profit off of user data, whereas users do not share equally in these profits [37]. For this, and for the other reasons mentioned, we aim to release our model only for research-based use.

#### A.1.1 Empirical study: can 🗨️ RESERVE identify individual celebrities?

Inspired by work studying language model memorization of private information [18], we wish to empirically probe 🗨️ RESERVE’s ability to recognize individuals. Our goal during model development was **not** to optimize for this ability. Instead, our goal was to study models for *multimodal script knowledge* (what people might be doing in a situation over time, and why) instead of long-tailed *visual recognition* (including who those individuals are). These goals might trade off – for instance, our training data only has individuals’

names when they are mentioned in ASR subtitles, a pairing that might be significantly noisier than images and text on the open web.

We study this capacity on the VoxCeleb2 and VGGFace2 datasets [86, 17], where we created a test set of 120 celebrities, with 100 samples of each. We study these datasets not to promote them, but to establish a **conservative upper-bound** for the capacity of a model to recognize non-celebrities. We hypothesize that if 🗨️RESERVE struggles to select the right celebrity out of 120 predefined options, it would struggle much more at identifying random people (where the set of candidate names is much greater). We test this hypothesis over three zero-shot settings:

- 1. Voice to name.** Given an audio clip sampled for a celebrity, we encode it with our model’s audio encoder. We provide our model’s joint encoder the text ‘the sound of MASK’, followed by the encoded audio. A blank image is provided. We extract the representation on top of the MASK, and choose the most similar celebrity name.
- 2. Image+voice to name.** Here, we adopt the same format as ‘Audio to name,’ except we additionally encode an image of the celebrity’s face in question.
- 3. Image to name.** Here, 🗨️RESERVE encodes an image of the celebrity in question, and we provide it with text ‘A picture of MASK.’ No audio is provided. Using our model’s joint encoder, we select the closest encoded celebrity name, out of all options.

We use this format to compare to a CLIP model, which was trained on web images with captions [91]. For the CLIP comparison, we use it to encode each image, and for all considered celebrity names, the sentence ‘A picture of  $\{\$ \{name\}$ ’. We choose the closest encoded sentence to the encoded image.

We show our results in Table 5. In all modes, our model is less than 11% accurate at recognizing celebrities. Curiously, the accuracy *drops* given both the image and the voice, suggesting that the way we fused a celebrity’s image and voice together might be outside the model’s training distribution. These results are significantly lower than CLIP’s 86% accuracy at classifying a person from their image.

In Figure 6, we investigate more into which celebrities our model is best at recognizing. Only a few celebrities are reliably classified; these tend to be very famous celebrities like Oprah Winfrey and Justin Bieber. Several sports players are recognized well (including LeBron James and Roger Federer), which could imply that our model learned their identities from watching sports replays or commentary. Most other celebrities are hardly recognized, whereas CLIP does well across the board.

**Results summary.** Together, these results show that while models like CLIP focus on encyclopedic knowledge that results in strong zero-shot person recognition accuracy,




Figure 6: VoxCeleb2 results per-celebrity, comparing 🗨️RESERVE-L versus CLIP ViT-B/32 in the same ‘image-text’ setting. Our model reliably recognizes A-list celebrities like Oprah Winfrey, very famous musicians (Justin Bieber) and sports players (LeBron James). However, it struggles on every other celebrity, particularly compared with CLIP. This suggests that our model primarily learns **semantic** as opposed to **recognition-level** encyclopedic knowledge.



🤖 RESERVE is not as effective as other models in memorizing particular celebrities— and, thus, perhaps not as effective as memorizing particular *non-celebrities*. These results suggest that 🤖 RESERVE’s objectives and data might make it *less of a concern* to release privacy-wise, versus models trained on web images with captions.

As the rest of the paper emphasizes however, 🤖 RESERVE performs well on tasks with temporal understanding and commonsense reasoning as the primary goal. On a broader level, these results suggest that it is possible to learn strong models about temporal reasoning *without* person-level memorization, though more work is needed.

## A.2. Biases in (pre)training data.

The ‘knowledge’ that our model learns should be viewed as situated within YouTube [66], which has numerous biases (that we will discuss next). Past work has made similar observations for language model pretraining on the open web[8]. One of the root causes of such bias is learning objectives that encourage memorization of surface level cooccurrences, rather than truly causal factors [55, 9, 117]. Though it is possible that in the very long term, a paradigm of *grounded learning* might resolve some of these issues, the objectives in this work still likely reify biases that exist in the YouTube data.

**Platform biases.** Unlike many other pretraining efforts, that scrape data from the open internet (e.g. [92, 16, 91]) which directly leads to toxic biases (e.g. [41, 32, 11]); we trained our model on YouTube, which is a moderated platform [101]. Though the content moderation might perhaps reduce overtly ‘toxic’ content, social media platforms like YouTube still contain harmful microaggressions [15], and alt-lite to alt-right content [94]. Additionally, it should be mentioned that the content moderation on YouTube disproportionately filters out minoritized voices [43]. Thus, despite us not using any word-based ‘blocklist,’ our model’s pretraining data is still biased [32]. Even without videos being explicitly removed, the ‘YouTube algorithm’ incentivizes the production of certain types of content over others [12, 102]; e.g. people’s roles in YouTube videos tend to be highly gendered [85], which might bias situation understanding [131].

**Bias amplification.** In this work, we pretrained a model primarily on ASR text, which is itself produced by another model. The automatic captions in YouTube are known to suffer from gender bias [107], which our model (like neural models generally) might in turn *amplify* [131]. The transcriptions on YouTube are also likely poor at handling important identity markers, like pronouns. Already, text-only models like BERT struggle with pronouns like they/them and zi/zir; our reliance on ASR text makes our corpus likely worse in this regard [28]. While past work, namely MERLOT [126], ‘cleaned’ this ASR text – through another large language model – we opted not to do so for this work

due to computational expense. Though in that work, the ASR-denoisification was found to boost performance in VCR, it seems unlikely that it would solve this core issue of model bias.

## A.3. Dual use.

Learning connections between video, audio, and text – though an important area of study as we have argued – can be used for undesirable applications, beyond what we have outlined under ‘biases.’ We outline and discuss a few below.

**Generating fake content.** A concern for pretrained models is that they can generate fake content, that could be used by ‘bad’ actors for their ends [125]. It should be noted that our model cannot explicitly ‘generate’ text, audio, or vision in a direct sense. Nonetheless, however, it is possible that a finetuned or expanded version of this model could be used for that purpose – and that our model would be *more helpful* to such an actor versus them training their own (perhaps larger) model from scratch.

**Surveillance.** Our model might contain representations that enable it to be used in surveillance applications. As we note in Appendix A.1.1, our model’s low performance on person recognition suggests that it might perform poorly recognition-focused applications. Still, one possibility is that a *neural script knowledge* could ‘summarize’ surveillance videos in some form (like identifying an activity of interest), without identifying the person(s).

We suspect (but cannot definitively prove) that the *reporting bias* of the YouTube data that it was trained on might make it poor for such a surveillance-focused task [49]. Namely, most surveillance videos are sparse in nature – finding an activity of interest is like finding a ‘needle in a haystack’ [90]. Though, some surveillance videos are inevitably posted on YouTube and then captioned, these disproportionately contain *interesting events* (like somebody’s car crashing into a house). It is not clear whether our system could be easily adapted to such a sparse problem; the amount of work required suggests that it might be out-of-scope at least for low-skill actors. On the other hand, this broad research agenda, and perhaps all of computer vision for that matter, might enable large actors to do just that [132]; which might not be addressable through purely technical solutions [51].

**Harmful outcomes if deployed.** Beyond the *biases* that our system possesses, some applications of our system – if deployed in production – could cause harm, particularly to groups already harmed by AI systems. Of note, linking someone’s voice with their appearance is not always a good thing [93]. Likely some of the key features that our model learns – though we did not teach it this explicitly – involve recognizing gender, and this is harmful especially to transgender individuals [54].

## A.4. Energy consumption.

Our model cost a lot amount of energy to pretrain [103]; roughly 3 weeks of time on a TPU v3-512. The total carbon footprint of our work was a net 8.23 tons of CO<sub>2</sub> equivalent, which is roughly 4.5% of the emissions of a jet plane flying round-trip from San Francisco to New York.<sup>12</sup>

At the same time, it is possible that our model could save energy overall, when shared with researchers who build off of our system. Indeed, RESERVE-B uses less energy than MERLOT [126] (due to a smaller vision backbone, and smaller image sizes), MERLOT in turn is more efficient than past work which used expensive detector-based backbones (e.g. [106, 20, 129]), that are made more expensive because some of their computational primitives (like non-maximum suppression) are difficult to make efficient on-device.

## A.5. Synthesis.

With these risks in mind, we release our video IDs, as well as RESERVE’s checkpoints, exclusively for research use. We believe that at this point in time, we as a field lack full knowledge of the privacy, bias, and dual-use risks of video-based models – though, we hope that our analysis in this section provides a good starting point. For instance, while the objectives that we have studied were designed to promote learning general neural script knowledge above *encyclopedic memorization*, they have not yet been tested in all possible cases. By opening our models to the research community, we hope to promote fundamental work in uncovering both promising aspects of these systems, alongside examining their risks. We hope to contribute to these lines of research as well.

## B. Model implementation details

In this section, we discuss at a more in-depth, technical level, how we implement certain aspects of RESERVE, and other details (like its runtime in FLOPs).

### B.1. Rotary position encoding

As mentioned in the main text, we use a rotary position encoding to model the relative location of input sequences [104, 10]. We chose this primarily because we did not want to use absolute (additive) position embeddings, which would have to be added to the inputs of each encoder, and possibly

<sup>12</sup>CO<sub>2</sub> Calculation. It is also important to consider the *location* where these TPUs are located, as the renewables portion at each datacenter is not equal [87]. Our TPUs were in the ‘europe-west4’ region, which uses on average 60% carbon-free energy, and a Grid Carbon intensity of 0.410 kgCO<sub>2</sub>eq / kWh. A single TPU v3 processor (with 8 cores over 2 chips) has a power average of 283 W, so after performing the math from [87], our training cost 20,000 kWh. This gives us a net 8.23 tons of CO<sub>2</sub> equivalent. It should be mentioned that this figure only covers the *electricity usage* given the chips (and the datacenter), not the raw materials involved in making these chips (which is significant [111]).

at multiple levels in the hierarchy (e.g. for the joint encoder, the video segment index  $t$  would be needed as well).

The rotary encoding uses no parameters, and instead uses a kernel trick to allow the model to recover relative distances between key and query elements in a Transformer’s attention head. This can be seen as ‘rotating’ pairs of elements; we apply the rotation to only the first *half* of each 64-dimensional head, and the second half is kept as is.

**Multidimensional coordinates.** We treat each token as having a 4-dimensional position of  $(h, w, \ell, t)$ , corresponding to the  $h, w$  coordinates in the image, the position  $\ell$  in the text-sequence, and the segment index  $t$ . If a dimension is irrelevant to a modality (like  $h, w$  for text), we set it to 0. Thus, for our various encoders, we use the following coordinate schemes:

- Video Frame Encoder (ViT): just the  $h, w$  coordinates of the image; so  $(h, w, 0, 0)$ .
- Audio Encoder: Only the 1-D position  $\ell$  of the patch in the spectrogram:  $(0, 0, \ell, 0)$ .
- Text Span Encoder: Only the 1-D position  $\ell$  of the token in the input:  $(0, 0, \ell, 0)$ .
- Joint encoder: Here, we use all coordinates. Inputs from the video frame encoder have coordinates  $(h, w, 0, t)$ , where  $t$  is their segment index. The text and (pooled) audio inputs are merged, and they each have coordinates  $(0, 0, \ell, t)$ , where  $\ell$  here is the absolute position in the entire sequence (across segments).

As part of our implementation, we normalize the rotary coordinates.  $h, w$  are scaled to be in the range  $[-1/2, 1/2]$ , such that text is implicitly ‘in the center’ of the image. Likewise,  $\ell$  and  $t$  are scaled to be in the range of  $[0, 1]$ . The positions are used to compute relative distances, by using a kernel trick to rotate coordinates in the keys and values of each  $d_h$ -sized Transformer attention head.

### B.2. Sequence lengths

We briefly remark on the sequence lengths used by parts of the model.

- Video Frame Encoder (ViT): Most YouTube videos are widescreen (16x9). We thus used a widescreen resolution for our video frame encoder. It takes in patches of size 16x16, and we used a layout of 12 patches (in height) by 20 patches (in width). This corresponds to **192x320**. Among other factors that are important are ensuring that TPUs do not excessively pad the sequence length [128]. The sequence length is 241 in this case, as there is a CLS token, and it gets padded to 256.

As we note in the main text, afterwards we apply attention pooling in a 2x2 grid (ignoring the CLS token here), which gives us a  $H/32$  by  $W/32$  grid for the joint encoder (6 x 10).

- b. **Audio Encoder.** Our model independently encodes each 1.6 second of audio. We do this through spectrograms. Each window involves 1536 samples at a sample rate of 22500 Hz, and there are 588 samples ‘hops’ between windows. We chose these hyperparameters largely around efficiency. We found that the Discrete Fourier Transform is fastest if the window size is close to a multiple of 2. We used a small number of mel spectrogram bins (64) because we found that at that threshold, we could reconstruct the original sequence at an acceptable level using the Griffin-Lim algorithm, [52] which itself might be a *lower bound* on quality as neural methods trained for this purpose have been shown to do better [115].

In our implementation, we compute the spectrogram for an entire video segment (5 seconds) at once; this is of size 64 mel bins by 192 windows. During pretraining, we perform what is effectively a ‘random crop’ over the spectrogram: we extract three sequential 64x60 sub-spectrograms, for each audio subsegment. We constrain them to not overlap, which means that 12 (random) windows are held out.

We note that our Audio Encoder AST is quite different from the one proposed by [46]. Though it operates over spectrograms, we opted for a linear ‘1-dimensional’ layout rather than a two-dimensional (image-like) one. We also did not pretrain our audio encoder on any supervised data (they used ImageNet and found, perhaps surprisingly, that it helped initialize the model). We used a patch size of 64 mel bins by 2 windows; the resulting (1D) sequence is of size 30. After adding a CLS token, the result is a sequence of length **31**.

As we note in the main text, we apply attention pooling afterwards (for all elements except the CLS token), pooling by a factor of five to resize the length-30 sequence to a length of 6 ‘audio tokens.’

- c. **Text Span Encoder:** We operate on spans that are at most of length 15, with an additional CLS token. Its length is thus **16**.
- d. **Joint encoder.** Let  $L$  be the number of text or pooled audio tokens given to the model per segment, on average; we set  $L=20$ . Let  $T$  be the number of video segments. Then, the joint model’s sequence length is  $T \times (L + W/32 \times H/32)$ . Our total sequence length was thus **640**.

To better adapt our model to downstream tasks – particularly single-image tasks like VCR [124], where past work tends to use a resolution much higher than 192x320, after pretraining, we performed FixRes pretraining (for one epoch on RESERVE-B, and one half epoch on RESERVE-L [108]).<sup>13</sup> Here, we trained the model on larger images – simultaneously on 288x512 widescreen images (18 patches

<sup>13</sup>We had intended to do a full epoch for RESERVE-L, but our job got preempted, and the loss seemed to have already converged.

Model	GFlops, from		Total	VCR
	Image Encoder	Joint Encoder		Q→AR Acc(%)
UNITER-Base[20]	1766	28	1794	58.2
UNITER-Large[20]	1767	99	1867	62.8
MERLOT [126]	236	67	303	65.1
RESERVE-B	99	46	146	62.6
RESERVE-L	176	165	341	<b>71.5</b>

Table 6: Efficiency metrics of our model versus others, measured in terms of (giga) floating point operations required to process a single image, question, and answer candidate on VCR. We compare with the overall VCR performance on the combined Q→AR metric. Our RESERVE family of models are significantly more efficient than prior work, with RESERVE-L being roughly on par with MERLOT [126] in terms of FLOPs, yet improving accuracy by over 6%.

by 32 patches), and on 384x384 square images (24 patches on each side). The joint encoder, correspondingly, uses a sequence length of 1312.

During 10 epochs of pretraining, we used a cosine decay of the learning rate down to 0.02 its maximum. During FixRes pretraining afterwards, we warmed up the learning rate to 0.02x its peak, over the first 1/5th of an epoch, and afterwards used a cosine schedule to anneal it towards 0.

### B.3. Efficiency metrics of our model

In Table 6, we report efficiency metrics of RESERVE, versus others. We calculate these metrics in the context of scoring a single VCR question and answer candidate. This requires encoding one image, and using 128 tokens for each question and answer combined (for all models). We compare against a UNITER [20], which is a representative VisualBERT style model, along with MERLOT [126]. Our models are far more efficient in terms of FLOPs, with RESERVE-L being roughly on par with MERLOT, yet outperforming it by 6% in terms of VCR accuracy. We discuss key differences below:

- a. **UNITER.** We note that UNITER, like other VisualBERT models, uses a supervised object detection backbone [5]. This processes images using a ResNet 101 model [56], at a resolution of 600x800; the final ResNet ‘C4’ block is applied densely over the entire image to obtain object-detection potentials everywhere in the image. Both factors greatly increase the FLOPs count.

When computing UNITER’s FLOPs count, we exclude operations like non-max suppression, which is an operation that is difficult to implement (and thus whose FLOP count might vary significantly depending on implementation). Our FLOPs count is thus a lower-bound. 36

detection regions are extracted, which is why the ‘joint encoder’ for UNITER is smaller than the equivalents for MERLOT and 🌈RESERVE.

- b. MERLOT. This model has two key differences versus our 🌈RESERVE. First, it uses a larger image resolution for VCR: 384x704, versus our 288x512. Second, it uses a hybrid ViT-ResNet50 backbone for encoding images. The backbone here is lighter weight than the object detection backbone of UNITER (in particular, the final ‘C4’ block is removed), and thus, as shown in Table 6, though it uses more FLOPs than does our 🌈RESERVE-L, it uses far fewer FLOPs than UNITER.

We choose flops as our primary comparison metric as past work shows that it is one of the key factors in model scaling [65, 34]. Parameters are arguably more fungible. For instance, in text-only representation learning, ALBERT [71] demonstrates that it is possible to *tie parameters* together at all layers of a BERT-like transformer, reducing parameters by an order of magnitude (while not modifying compute), with a minimal performance drop. We did not do this for this work, as we wanted to use a more ‘vanilla’ Transformer architecture; however, it suggests that representation learning models with hundreds of millions of parameters might be *FLOPs bound* as opposed to *parameter-bound*.

Nonetheless, UNITER-Base has 154 million parameters, though some are frozen (86 million from their Transformer, 23 million from the word embedding layer, and then 44 million from their object detector [5]). UNITER-Large has 378 million parameters (303 from their Transformer, 31 million from word embeddings, and 44 million from the same object detector. Meanwhile, MERLOT has 223M parameters. Versus our 🌈RESERVE-B, 14 million extra parameters are due to a larger vocabulary, and 10 million parameters are due to a ResNet50 encoder – but these parameters have a disproportionate impact in FLOPs count.

#### B.4. Full model hyperparameters

In Table 7, we present full hyperparameters for our model. Among other details, we used AdamW as our optimizer, with  $\beta_2 = 0.98$  and  $\epsilon = 1e - 6$ . We increased the learning rate linearly to its peak value (4e-4 for 🌈RESERVE-B, 3e-4 for 🌈RESERVE-L) over 3750 steps ( $\frac{1}{20}$ th of an epoch). Our number of warmup steps is lower than many other pretraining work; we note that all of our contrastive objectives involve learning a  $\sigma$  parameter, which functions as a secondary ‘warmup.’

We did not use gradient clipping. We trained and evaluated in 16-bit bfloat16 precision wherever we could – casting all gradients to that precision as well, and saving the AdamW running mean and variance to be 16-bit as well. A few times during pretraining 🌈RESERVE-L, we found that some values in gradients would be NaN. We addressed this by always setting NaN values to be 0. This seemed to address

	Base	Large	
Audio size	Sample rate	22050 Hz	
	FFT hop length	588 samples	
	FFT window size	1536	
	Mel bins	64	
	Subsegment length	60 hops, ( $\approx 1.6$ sec)	
	Patch size	64 mels $\times$ 2hops	
Pooling ratio	5		
Final size	6 tokens		
Image	ViT patch size	16	
	Pretraining size	192 $\times$ 320	
	Res-adaptation size	288 $\times$ 512 and 384 $\times$ 384	
	Pooling window	2 $\times$ 2	
Text	Max. span length	15	
	Mean span length	5.5	
Joint sizes	$N$ video segments	16	
	video segment groups	2 (each with 8 segments)	
	Pretraining seq. length	640 (160 text&pooled audio; 480 pooled vision)	
	Res-adapted seq. length	1312 (160 text&pooled audio; 1152 pooled vision)	
Batch sizes	Videos	1024	
	# Frames (for matching)	16384	
	Masking rate	25% (of subsegments)	
	Text spans	49152	
Audio spans	49152		
architecture	Hidden size	768	1024
	Num attention heads	12	16
	Size per head	64	
	Rotary size (per head)	32	
	Vision num layers	12	24
	Audio num layers	12	
	Text-span num layers	4	
	Joint num layers	12	24
optimizer	Peak learning rate	4e-4	3e-4
	Weight decay	0.1	
	AdamW $\beta_2$	0.98	
	AdamW $\epsilon$	1e-6	
	Warmup steps	3750	
	Training steps	750k (+ 75k for res. adaptation)	
	Training epochs	10 (+ 1 for res. adaptation)	
$\sigma$ Maximum scale	100.0		
Pretraining compute	TPU v3-512 for 16 days	TPU v3-512 for 5 days	

Table 7: Architecture details, and pretraining hyperparameters, for both model sizes.

	Base	Large
VCR	Batch Size	32
	Training Epochs	5
	Image Size	288×512
	Learning Rates Tried	1e-5, 2e-5, 3e-5
	Learning Rate	2e-5
TVQA	Batch Size	32
	Training Epochs	3
	Image Size	288×512
	Learning Rates Tried	5e-6, 1e-5
	Learning Rate	5e-6

Table 8: Hyperparameters for finetuning on downstream tasks.

the symptoms of training instability – though sometimes the training loss would spike to roughly around the same loss as random initialization, it always converged back to *slightly better* than it was before the spike. We are not currently sure why this happens.

### B.5. Speed improvements during pretraining

We made several high-level algorithmic and engineering implementations to our implementation, which made pretraining run faster, and that we discuss here.

**Duplicated video copies.** As mentioned in the main text, we create two copies per each video – allowing us to learn separately *how to handle audio as an input* as well as *how to learn from audio*. We chose this in part because copying a video *does not* increase the total compute required by a factor of two. Instead:

1. We encode the underlying video frames and audio clips only once (for the two video copies), and then duplicate the encodings; this is far more efficient than encoding them both separately from scratch.
2. For the two video copies, we sampled two disjoint sets of masks (for which audio and text subsegments are replaced with MASK) at a 25% rate. This increases the pool of negative samples for contrastive learning, again increasing training efficiency.

**Reducing memory usage.** The memory usage of our Transformer implementation scales quadratically with sequence length, which could pose a problem since we operate on sequences of videos. We split the video into two groups of 8 segments, and encode these each group separately by the joint encoder.

**Vectorization.** We vectorize all joint transformer inputs together into a single call. During this vectorization, we also encode the transcript (for the transcript-frame matching objective).

We note that this vectorization is incompatible with the Mask LM variant proposed by MERLOT [126]. In this variant, which the authors called ‘attention masking,’ two transformer calls must happen *sequentially* – **first**, a language only encoder must encode the inputs and mark down (what is presumably) visually-grounded tokens; **second**, these tokens are masked for the joint encoder. We found that such an objective was unnecessary when pretraining under our contrastive learning approach, which in turn enabled more efficient pretraining.

We discuss the exact pretraining data formatting technique that we used in the next section.

## C. Pretraining Data Formatting: alignment and masking

In this section, we discuss how we turn a video  $\mathcal{V}$  into a (masked) list of segments  $\{s_t\}$  for pretraining.

Recall that each segment contains a video frame  $v_t$ , ASR tokens  $w_t$ , and audio  $a_t$ . We generate the list of segments by iterating through the video with a 5-second sliding window.<sup>14</sup>

**Audio and text subsegments for masking.** We want audio to be used in part as a *target* for contrastive prediction. However, during early exploration we found that 5 seconds of audio could correspond to many BPE tokens; roughly 15 on average. We use past work in language modeling as a guide [63, 92] and wanted an average span length of around 5 tokens. To get this, we split each audio segment into three equal *subsegments*, each with a duration of 1.66 seconds. We can then perform masked language modeling at the *aligned subsegment level*, where we mask out the text corresponding to an audio subsegment, and have the model (contrastively) predict the masked-out span of text, as well as the corresponding span of audio. We use a masking rate of 25%, which means that a quarter of the subsegments will be corrupted and replaced by a MASK token.

In theory, splitting the videos into (masked) segments ought to be straightforward. However, the key challenge that we ran into is that **the YouTube caption timing information is unreliable**. Problems might arise when we perform pretraining with both audio and text, on misaligned data. Suppose the model is given audio in segment  $s_{t-1}$  that ends with somebody saying the word ‘pasta.’ If the alignment between audio and text is off, the model might be able to cheat the desired task by simply predicting the word ‘pasta’ for segment  $s_t$  – thereby turning the challenging masked-prediction task into an easier speech recognition task; we discuss this in more detail in Appendix C.1.

<sup>14</sup>Sometimes there are long ‘pauses’ in videos where nothing gets said. When this happens – if two segments in a row have fewer than 8 BPE tokens – we merge them 90% of the time, in effect ‘fast-forwarding’ the audio and still extracting a frame from the middle. We do this at most twice, so the total length is at most 15 seconds here (in effect, a ‘playback’ rate of 1x, 2x, or 3x). In roughly 90% of cases, the segments are 5 seconds of length.

One way of addressing the timing issue would be to run our own ASR model over all videos, but we chose not to do this due to computational expense. Instead, we adopted two complementary strategies. First, we trained a lightweight regressor to refine the timing information (C.2); second, we mask audio and text conservatively, to minimize alignment errors (C.3). Finally, we discuss how we combine everything efficiently (in a vectorized way) in C.4.

### C.1. YouTube Caption Timings

YouTube provides automatically generated captions for accessibility purposes, which include timing information on each word. In the subtitle encoding that we used (vtt), each word  $w$  contains a single timestamp  $t$  which corresponds to when the word should flash on-screen. The timings are *mostly* accurate, but we found two key issues:

- a. First, they show up on average roughly 0.1 seconds before each word is spoken, which we suspect might be for usability purposes (perhaps so that while the viewer is reading the caption, they hear the word).
- b. Second, with a single timestamp  $t$  for each word, it is difficult to infer about pauses. For each word  $w$ , we can use its timestamp  $t$ , and the timestamps of adjacent words, to loosely infer an interval  $[t'_s, t'_e]$  around when the word is said. However, the interval is not tight. We can only infer that the word is being actively spoken for some subinterval  $[t_s, t_e]$  such that  $t'_s \leq t_s \leq t_e \leq t'_e$ .<sup>15</sup>

This can lead to high absolute error (in terms of a difference between timesteps), when pauses occur. For example, suppose a speaker says a word, and then pauses. The interval given by the subtitles,  $[t'_s, t'_e]$ , might be rather large (possibly a few seconds), even though the actual word was spoken for a fraction of that time.

### C.2. Refining timing information

We trained a simple multilayer perceptron regressor to correct the timing information of YouTube transcripts. For data, we used 2000 videos with transcripts from YT-Temporal-180M, and also used Google Cloud’s (highest quality, paid) ASR service to transcribe them. After aligning the words for these transcripts, this gave us tuples of the YouTube ASR word  $w$ , its provided interval  $[t'_s, t'_e]$ , and the ‘ground truth’ interval  $[t_s, t_e]$ .<sup>16</sup> Our modeling objective was then to predict the desired offsets with respect to the provided interval:  $\delta_s = t_s - t'_s$  and  $\delta_e = t_e - t'_e$ . We took a feature based approach.

<sup>15</sup>Note that this is compounded with the first problem, the ground truth interval  $[t_s, t_e]$  might not be fully contained in the provided interval  $[t'_s, t'_e]$  due to ‘captions being shown before audio’, the error here is typically small though (0.1 seconds).

<sup>16</sup>When matching YouTube ASR to Google Cloud’s ASR, we skipped words without an ‘exact-match’ alignment, as well as words that were over 0.25 seconds apart (i.e., where either  $\delta_s > 0.25$  or  $\delta_e > 0.25$ ).

For each input  $(w, t'_s, t'_e)$ , we used as features:

- i. the length of  $w$  in characters,
- ii. the length of  $w$  in BPE tokens,
- iii. whether  $w$  is uppercase or not,
- iv. the number of vowels in  $w$ ,
- v. the number of punctuation characters in  $w$ ,
- vi. the value of  $t'_e - t'_s$ .

We provided these features as input to the model, as well as the corresponding features for the next word, and the previous word. We z-normalized all features and used a two-layer multilayer perceptron, with a hidden size of 32 and RELU activations. We used a tanh activation at the end to bound the regression. The final predictions for  $\delta_s$  (analogously for  $\delta_e$ ) were then given by the following equation:

$$\delta_s = c \tanh(\mathbf{w} \cdot \mathbf{h} + b_1) + b_2 \quad (3)$$

where  $\mathbf{h}$  is the hidden state, and with learnable parameters  $c$ ,  $\mathbf{w}$ ,  $b_1$ , and  $b_2$ . The learned bounds mean that, no matter what the input, the model will never predict an offset of above  $c + b_2$  (of which it learned for both parameters  $c \approx 0.2$  and  $b_2 \approx 0.11$ , so the offsets can never be above 0.3 seconds). We trained our lightweight regression model using an L<sup>1</sup> loss, and used it to correct the timing on all of the transcripts.

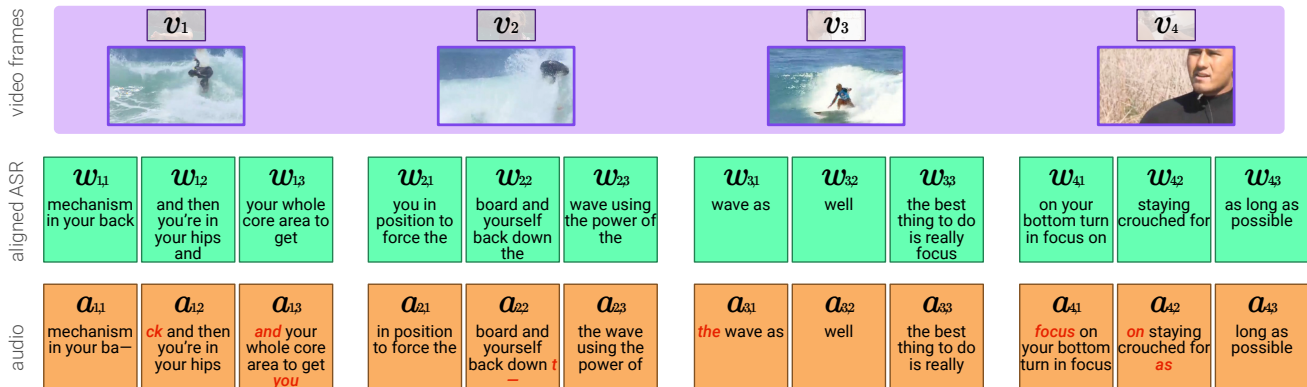
### C.3. Handling worst-case scenarios in masking, when alignment isn’t perfect

The regressor that we described reduces the average timing error of a transcript, as a preprocessing step, but it is not perfect. Thankfully, however, we find that most of the remaining alignment errors are *single* words that are slightly misaligned. For instance, for three words  $w_t, w_{t+1}, w_{t+2}$ , the audio corresponding to the time interval around  $w_t$  might contain sound from  $w_{t+1}$  being spoken, but rarely  $w_{t+2}$ . We suspect this is primarily due to the difficulty inferring pauses: by definition, no other word can be said in a pause, so the errors are local.

We present a high level approach for masking audio and text, that in turn addresses these alignment issues (making it difficult for models to cheat). A diagram is in Figure 7.

Recall that in our framework, we only either go from ‘vision and text  $\rightarrow$  text and audio’ (VT $\rightarrow$ TA), or, ‘vision, text, and audio  $\rightarrow$  text’ (VTA $\rightarrow$ T). One of the reasons we did this is to avoid allowing a model to cheat by performing speaker identification (or even ‘microphone identification’), which might be feasible if audio was given to the joint model as input. We can handle the two cases separately:

- a. **Vision and text  $\rightarrow$  text and audio** (VT $\rightarrow$ TA). Here, the text as input (to the joint encoder) might overlap with the audio we are trying to predict. Our solution here is thus to **donate nearby tokens** from the predicted span, to the input. Let the span that we are trying to predict (and that



The audio has **alignment errors** versus the aligned ASR text. So...




Figure 7: An overview of our masking strategy for dealing with sequences of video frames, ASR, and audio. We have noisy timing information for each word, so we can align the ASR text with audio spans of 1.6 seconds each, using three sub-segments of audio and text for each video frame. However, there exist **alignment errors** between the ASR and audio sub-segments – certain words (and sub-words) have phonemes that are in the wrong segment (like ‘back’ in  $w_{1,1}$  is only partially said in the first sub-segment; the ‘k’ sound is said in the second. When audio is only a target, we address these by ‘donating’ tokens to predicted spans. When audio is only provided as input, we address this by sandwiching ‘mask’ tokens between text input (so alignment does not ‘bleed’ over).

we will ‘mask out’) have a start time of  $t_s$  and an ending time of  $t_e$ . If the final token in the previous text span, if any, has a timestamp of greater than  $t_s - 0.125$ , we move it to the predicted span; likewise, if the first token in the next text span has a timestamp of less than  $t_e + 0.125$ , we move it to the predicted span as well.

- b. **Vision, text, and audio**  $\rightarrow$  **text** (VTA $\rightarrow$ T). In this prediction task, models are given information from all modalities as input, and must predict masked-out text spans. Note that models are only given a single ‘speech’ modality – either text, or audio – at each timestep. What this means is that we can carefully choose *which input subsegments* to turn into ‘audio subsegments,’ and which to turn into ‘text subsegments.’ Our strategy is, given a masked out subsegment, to turn 80% of adjacent subsegments into ‘text subsegments.’

We give an illustration of this in Figure 7, part 2. Here the word ‘focus’ is part of  $a_{4,1}$  but also  $w_{3,3}$ ). This

might make  $w_{3,3}$ ) overly easy to predict, if we gave the model  $a_{4,1}$  as input. Our solution is thus to give the model text from  $w_{3,2}$ ) and from  $w_{4,1}$ ) as input; we are guaranteed that there is no misalignment overlap here between input and prediction spans. All of the other subsegments (not adjacent to one of the 25% that we mask out) will be provided as audio.

#### C.4. Putting it all together, along with web text

Finally, we discuss how we combine the various masking approaches into the prediction tasks outlined in the main text.

Each video has  $N = 16$  video segments, and three sub-segments of audio or text spans per segment. We consider two sub-problems for this video sequence:

- i. in VT $\rightarrow$ TA, vision and text are provided as input, and the model must predict masked-out text and audio. These are done on top of separately-encoded MASK tokens and MASKAUDIO tokens, to enable the model to learn different

predictions for each modality over two separate transformer ‘columns.’

- ii. In  $VTA \rightarrow T$ , vision, text and audio are provided as input, and models must predict masked-out text. Here, we use the term ‘predict’ as a shorthand for our contrastive objective – in which a model must match a context (a jointly-encoded MASK) to the exact missing span in question, where many negative contexts and spans are provided.

We use a masking rate of 25% for audio and text subsegments, and there are 3 subsegments per segment. This means that a single video instance gives us  $48 \times 0.25 = 12$  masked-out spans of text, for each of  $VT \rightarrow TA$  and  $VTA \rightarrow T$ , so 24 in total (as we use disjoint masked-out subsegments). Likewise, it gives us 12 masked-out spans of audio. If we scaled these to the whole batch of 1024 videos, we would have 12k audio span options and 24k text span options. This might suffice, but scaling up the pool of candidates boosts performance in a contrastive setting, as suggested from prior work (e.g. [91]), and as our ablations (Table 1) support as well. Thus, we do the following:

- a. **Text candidates.** We scale up the text candidates by simultaneously training the model on web text, from The Pile [40]. The joint encoder – which can handle pooled video, pooled audio, and BPE-encoded text – is simultaneously given a sequence of web text, for each video that we have. By performing the span-contrastive objective with this piece of web text as well, we can not only teach the model about written (as opposed to spoken) language, but we can scale up the set of text candidates as well.

Let each web-text sequence be of length  $L$ . We first divide it into fake regions that ‘look like’ the text subsegments in length. We do this by calculating the empirical length distribution of the text subsegments, and then using this (categorical) distribution to sample a sequence of sub-segment lengths  $\ell_1, \dots, \ell_K$ .<sup>17</sup> We clip the sampled sequence, such that  $\sum_i \ell_i = L$ .

Next, we mask the fake subsegments. During pretraining, we use text sequences of length  $L = 800$ , but a model sequence length of only 640. Because we are masking spans and not individual tokens, the text sequences ‘shrink’ when we mask them. We extract exactly 38 masked-out spans, which corresponds to around 25% of total text.

Finally, we combine the target spans that we took from the webtext sequence, with the target spans from the video. We note that sometimes – especially in a video – text spans might be empty. Not every 1.6 second slice of a video has someone speaking. We thus try to not

<sup>17</sup>The empirical distribution for each length, in order from a length of 1 to 15, is [0.03, 0.05, 0.08, 0.11, 0.13, 0.13, 0.12, 0.10, 0.07, 0.05, 0.03, 0.02, 0.01, 0.006, 0.003].

use these empty spans in our contrastive objective. For each video (which is paired with text for implementation reasons) we select the ‘best’ 48 text spans out of the (38+24) options – penalizing empty spans, and choosing spans from videos 4x as often.

These ‘best 48’ text spans, as well as the pooled contexts that they were paired with, will be used in the contrastive objective. Aggregating over the entire batch of 1024 videos (and 1024 web text sequences), this gives us 49152 text spans as candidates, for the all-pairs symmetric softmax between text spans and contexts.


- b. **Audio candidates.** For each video, we note that we have exactly 12 pooled MASKAUDIO tokens, where the model is trying to predict the corresponding audio span. One option would be to just use those 12 corresponding audio spans as the targets, aggregate these over the batch, and do a symmetric-cross-entropy loss.

However, we can do even better *for free*. Note that for the  $VTA \rightarrow T$  direction, we might have to encode many of the audio spans *anyways*, using the lower level audio encoder (which simultaneously extracts a CLS representation and a sequence-level pooled representation). To simplify implementation, we encode *all 48 audio spans* per video. We can use these audio spans as candidates.

Thus, we do the following when computing the loss over audio prediction. We aggregate all 12288 *contexts* from the MASKAUDIO tokens in the batch, and we aggregate all 49152 candidate audio spans. We perform an all-pairs dot product between these two sets, and use it to compute a symmetric cross-entropy loss over both directions. We did not encounter any trouble using the same temperature for both directions (even though for one direction, there are 12288 options, and for the other, there are 49152).

The combination of these design decisions provide more ‘hard negatives’ for the model during training. We also found that they worked well to reduce wasted computation on a TPU. For each video, the joint transformer uses one  $L = 640$  length sequence for transcript-frame matching, two length- $L$  sequences for the  $VT \rightarrow TA$  direction (as we break it up into two groups of 8 frames each), two length  $L$  sequences for the  $VTA \rightarrow T$  direction, and finally one length- $L$  sequence of text. These sequences can all be vectorized together, and the total batch size is  $6 \times$  the number of videos. This is helpful because using an even-numbered batch size reduces wasted computation on a TPU.

## D. Downstream Task Implementation Details

In this section, we present information for how we adapted  RESERVE on downstream tasks.



## D.1. Setup for finetuned Tasks

For adapting 🗯️RESERVE in a finetuned setting, we take the following approach. We always trained the model for either 3 or 5, depending on the dataset size. We always used a batch size of 32. We use a linear warmup of the learning rate over the first half of the first epoch, with a linear decay thereafter to 0. To find the learning rate, we did a small grid search generally centered around 1e-5; details in Table 7.

When finetuning (and pretraining), we did not use any dropout to make implementation simpler. Instead, as a way to apply regularization, we used the same  $L_2$  penalty as in pretraining (a weight decay of 0.1), but with respect to the *pretrained* weights. This idea was used in [118] among other works, and although it often tends to underperform dropout [72], it is simple to implement. We also randomly resized images (or video frames) to between 100% and 110% of the network’s size. We suspect that additional work in applying other regularization techniques (e.g. AutoAugment [25]) could improve our results.

### D.1.1 Visual Commonsense Reasoning

We randomly flip images left or right, so long as there is no instance of the word ‘left’ or ‘right’ in the question, answer, or rationale candidates.

We also followed past work in ‘drawing on’ the provided detection tags to the image [126]. These are unambiguous references to entities that are then referred to in the question, answer, and rationale. For example, text might reference a ‘person1’, which corresponds to an image region. When drawing on these detection tags, we do so in a deterministic way – for example, ‘person1’ always gets the same box color. We determine the box color by hashing the object’s ID (in this case, ‘person1’) and using that to determine the hue. The model learns the connection between boxes with different hues, and the names, during finetuning.

### D.1.2 TVQA

TVQA provides models with a video, a question, and five answer candidates; we represent this as five distinct sequences for the model to score (one per candidate). The version of TVQA that we used also gives models annotations for the *time region* in the video that is being referenced. It is not clear that only using this region would provide enough context to be able to understand what is going on – enough to answer correctly. Thus, for each question, we extract 35 seconds of video around the provided time region. We then provided the model with two numbers corresponding to the time region, relative to the cropped time interval. For example, if the provided timestamp annotation is  $[t_0, t_1]$ , we

use the following region:

$$t_c = \frac{(t_0 + t_1)}{2} \tag{4}$$

$$t_s = t_c - 17.5 \tag{5}$$

$$t_e = t_c + 17.5 \tag{6}$$

The location of  $[t_0, t_1]$  in relative coordinates is then:

$$t_0^r = \frac{t_0 - t_s}{t_e - t_s} \tag{7}$$

$$t_1^r = \frac{t_1 - t_s}{t_e - t_s} \tag{8}$$

We provide models with  $t_0^r$  and  $t_1^r$ , multiplied by 100 and casted to an integer. Thus, an example TVQA instance might look like:

1 to 28 What is Janice Holding on to after Chandler sends Joey to his room? Chandler’s tie.  
MASK[subtitles or audio]

This text input corresponds to the first ‘segment’ of a video; to it we append subtitles (or audio representations) from seven segments from the provided TVQA video (with accompanying frames).

## D.2. Setup and prompting for Zero-shot tasks

Here, we discuss how we set up various tasks for 🗯️RESERVE in a fully zero-shot setting. In addition to evaluating 🗯️RESERVE, we also evaluate CLIP [91] in the same zero-shot setting. CLIP is not pretrained on videos, and it cannot jointly encode text. For each task, we construct CLIP’s label space by taking our prompt and substituting in each possible answer option. We average together the logits over all frames, and take a softmax, giving us a distribution over the task-specific label space.

### D.2.1 Zero-shot Action Anticipation on EPIC-Kitchens

We study the task of action anticipation from the EPIC-Kitchens dataset [26], a large egocentric video dataset with 700 unscripted and untrimmed videos of cooking activities. In action anticipation, a model must predict a *future action* that comes  $\tau_a$  seconds after a given video clip. The observed segments are of arbitrary length; we follow prior work [26] and set  $\tau_a = 1$ .

The model tries to choose the correct *noun* and *verb* that happens next, given a list of predefined options for each. We report results on each category using the class-mean top-5 recall.

**Zero-shot inference approach.** We directly evaluate the pretrained 🗯️RESERVE on action anticipation to verify the knowledge learned during pre-training. All prior work reported on the official leaderboard use supervision from the in-domain training set, which we do not use at all [45, 38].

Model	Overall			Unseen Kitchen			Tail Classes		
	Verb	Noun	Act	Verb	Noun	Act	Verb	Noun	Act
RULSTM [38]	27.8	30.8	14.0	<b>28.8</b>	<b>27.2</b>	<b>14.2</b>	19.8	22.0	11.1
AVT+ (TSN) [45]	25.5	31.8	14.8	25.5	23.6	11.5	18.5	25.8	12.6
AVT+ [45]	<b>28.2</b>	<b>32.0</b>	<b>15.9</b>	19.5	23.9	11.9	<b>21.1</b>	<b>25.8</b>	<b>14.1</b>
Chance	6.4	2.0	0.2	14.4	2.9	0.5	1.6	0.2	0.1
CLIP (ViT-B/16) [91]	13.3	14.5	2.0	12.3	8.4	2.1	14.3	14.3	1.7
CLIP (RN50x16) [91]	16.5	12.8	2.2	13.4	7.0	1.2	17.1	12.6	2.5
Validation									
🧠 RESERVE-B	17.9	15.6	2.7	11.0	15.7	4.4	18.0	12.7	2.0
🧠 RESERVE-L	15.6	19.3	4.5	14.1	18.4	3.4	14.7	18.5	4.4
🧠 RESERVE-B (+audio)	20.9	17.5	3.7	15.5	20.1	4.3	20.7	14.5	3.2
🧠 RESERVE-L (+audio)	<b>23.2</b>	<b>23.7</b>	<b>4.8</b>	<b>20.3</b>	<b>21.0</b>	<b>5.9</b>	<b>22.7</b>	<b>21.6</b>	<b>4.0</b>
Test									
RULSTM [38]	25.3	26.7	11.2	19.4	26.9	<b>9.7</b>	17.6	16.0	7.9
AVT+ [45]	<b>25.6</b>	<b>28.8</b>	<b>12.6</b>	20.9	22.3	8.8	19.0	22.0	<b>10.1</b>
🧠 RESERVE-L (+audio)	24.0	25.5	5.8	<b>22.7</b>	<b>26.4</b>	7.0	<b>23.7</b>	<b>24.2</b>	4.7

Table 9: 🧠 RESERVE gets competitive results on EPIC Kitchen Action Anticipation challenge with zero-shot, over methods from prior work.

For each action segment, we sample at most  $N = 8$  image frames and their associated audio, with fixed time interval  $t = 2.0$  preceding it and ending  $\tau_a$  seconds before the start of the action. We append a MASK token as the sole text input (at the last frame, after audio is optionally included).<sup>18</sup> We create short phrases out of all candidate nouns and verbs, and use that as our label space to simultaneously predict them both. We compute the score for each verb and noun independently by averaging their scores, over all labels for which they appear.

**Results.** We show the full zero-shot action anticipation results in Table 9. We also show our results on the test set here for our best performing model (🧠 RESERVE-L, with audio provided). It gets competitive results on verb and noun prediction – with only 1.6% and 3.3% lower compared to the challenge winner method AVT+ [45], which is fully supervised and use additional object-level annotations. On Unseen Kitchen and Tail Classes, our model **outperforms AVT+** on noun and verb. Overall, audio significantly improves the results – 🧠 RESERVE-L (+audio) outperforms 🧠 RESERVE-L with an average 3.0%, which suggests that it is useful for this task.

## D.2.2 Zero-shot Situated Reasoning

Next, we evaluate on situated reasoning (STAR) [119] which requires the model to capture the knowledge from surrounding situations and perform reasoning accordingly. STAR dataset includes four types of questions, including interaction, sequence, prediction, and feasibility. A model is given a video clip, a templated question, and 4 answer choices.

**Zero-shot inference approach.** For each video clip, we sample  $N = 8$  image frames uniformly from the video, we

<sup>18</sup>We were unable to find a better text based prompt than this, as we found that they often biased the model towards linguistically relevant words; however, we suspect that such a prompt does exist.

also optionally include the video’s sound.

To reduce domain shift between YouTube data – where people don’t typically ask visual questions, and where ASR typically does not insert question marks – we convert the question-answer pair into a statement. We did so using the question-answer templates provided by the author, with the answer replaced by a MASK. For example, “*Q: What did the person do with the bottle? – A: Put down.*” will be converted to “*The person MASK the bottle.*”.

We put the converted statement into the first frame and use the four candidate answers as a unique label space (that differs from example to example). Like with EPIC-Kitchens, we also evaluate how much audio can help by masking the audio inputs.

**Results.** We show our zero-shot STAR results in Table 4 in the main text. Our base model outperforms all supervised prior work by 3.7%. The model with audio performs better, with average 1.1% improvement. Interestingly, 🧠 RESERVE-L is worse than 🧠 RESERVE-B, we suspect the reason is 🧠 RESERVE-L is sensitive to grammar details. Given the previous example, we note that while ‘Put down’ is a valid answer that might make sense both semantically and syntactically, a different answer ‘pick up’ might be flagged by some English speakers as being ungrammatical: the instantiated template would then be ‘the person pick up the bottle.’ We noticed instances of the larger model paying greater attention to these syntax-level details, even though they were not the focus of the task. It does suggest, however, that additional prompting (or label space augmentation) could resolve these issues and increase performance even further.

## D.2.3 Zero-shot LSMDC

We evaluate our model on Movie Fill-in-the-Blank [95, 81] task, which based on descriptive audio description for the visually impaired. Given a movie clip and an aligned description with a blank in it, the task is to fill in the blank with the correct word. Following [81], we report prediction accuracy in test set of 30,354 examples from 10K movie clips.

**Zero-shot Inference approach.** We sample  $N = 8$  video segments uniformly over the movie clip, and extract the audio and middle frame of each segment. We replace the ‘blank’ token in each description with a MASK token, and provide it (as text-based input) to the model at its final segment. For the other segments, we optionally provide the model with audio; for all segments, we provide the associated image frame. We use the vocabulary set in the LSMDC dataset as our label space (for what the ‘missing word’ might be).

**Results.** Our results are shown in Table 4 in the main text. Our model obtains 31% when audio is included, which

outperforms human text-only performance (30.2 %) [81], predicted by human annotators. A supervised LSTM obtains 34.4% in this text-only setting [81] which suggests that there is a certain textual bias in this task, which our model cannot learn (as it is zero-shot). This also suggests that state-of-the-art supervised models exploit patterns in this vocabulary distribution.

Without such an advantage, our model performs well, outperforming CLIP (2%) by a large margin. This suggests that jointly reasoning over both the visual situation, and the linguistic context of the provided sentence, is helpful for zero-shot performance on LSMDC fill-in-the-blank.

### D.2.4 Zero-shot MSRVTQA

Finally, we evaluate our model on MSR VTT-QA, a question-answering task over videos [120]. We provide a model with  $N = 8$  video segments sampled uniformly from the video clip, and extract an image from each one. For the first seven segments, we optionally include audio extracted from that point; at the last segment, we insert a converted version of the question, along with a MASK. We compare the similarity of that hidden state to the top 2000 most common answers, similar to past work [126].

Similar to STAR, we convert the questions into statements to minimize drift away from the pretraining distribution. We use GPT3 prompted with several examples for this. Our exact prompt is the following:

```

Input:  what is a car being driven through?
Output: a car is being driven through _ .

Input:  who are running across screen?
Output: _ are running across screen.

Input:  when is a girl performing?
Output: a girl is performing at _ .

Input:  what is a cartoon doing?
Output: a cartoon is _ .

Input:  how many women talk in a bedroom?
Output: _ women talk in a bedroom.

Input:  what a man playing while dancing with others?
Output: a man is playing _ while dancing with others.

Input:  where is a flag hoisted?
Output: a flag is hoisted in _ .

Input:  who talks to another man on the couch?
Output: _ talks to another man on the couch.

Input:  what does a teenage girl try to get at a public restroom?
Output: a teenage girl tries to get _ at a public restroom.

Input:  when do the models walk as the audience watches?
Output: the models walk as the audience watches at _ .

Input:  what shows a person killing animals in a green forest?
Output: _ shows a person killing animals in a green forest.

Input:  who does a man ask to go on a date?
Output: a man asks _ to go on a date.

Input:  what are three people sitting on?

```

```

Output: three people are sitting on _ .
Input:  ${question}
Output:

```

Then, given a new question  $\text{\${question}}$ , GPT3 generates a converted output, wherein we can replace its underscore with a MASK. GPT3 works well at this conversion, though sometimes it generates a sentence where inserting the ‘correct answer’ feels grammatically strange. For example, the question ‘how many women talk in a bedroom?’ suggests any integer might be a reasonable answer. On the other hand, ‘\_ women talk in a bedroom’ implies that ‘one’ is not a valid answer (since ‘women’ is plural). We note that the errors caused by this conversion technique are specific to English grammar, and so if such a question-conversion approach was done in other languages, there could be more (or less) errors that directly result.

Our results are shown in Table 4. Of note, our model through automatic question-conversion outperforms Just Ask [122], which performs an analogous (supervised-guided) question conversion on all its YouTube transcripts, before pretraining. Our model also outperforms CLIP, which cannot naturally handle dynamic situations.

## E. Dataset Collection

In this section, we discussed how we curated data for YT-Temporal-1B. We had several goals in mind. We wanted to use only public-facing data, which motivated our choice of YouTube as it is a public platform *that users understand is public* [64]. We wanted to use this platform to examine to what extent we can learn multimodal neural script knowledge from web data alone.

Our data collection strategy in this work was informed by past work, notably MERLOT [126]. That paper found that increasing the diversity and scale of a video corpus both allowed for better learned representations. At the same time, the data collected by MERLOT (YT-Temporal-180M) has issues. Of note, the authors’ scraping strategies – to prioritize *monetized content* – also led to a lot of U.S. local news being in that corpus (roughly 30% of all data). Local news might be problematic to learn from, particularly in that quantity, due to its numerous biases (e.g. racist coverage on ‘crime’ [44, 31, 30, 57]). Our goal was to expand the dataset in both diversity and size to 20 million videos, while having *less local news* and without scraping private content.

**High level approach.** We adopt a similar dataset collection strategy as in MERLOT [126]. In the first phase, we identify a candidate set of videos ID to download. In the second phase, we open each video ID in YouTube and apply several filtering steps that go from inexpensive to expensive. The filtering steps allow us to exit early and possibly avoid downloading the video if the video seems unsuitable for our purpose from the title, description, and captions alone.

### E.1. Candidate video IDs

For MERLOT’s YT-Temporal-180M, the bulk of the video IDs were identified by applying breadth-first-search on YouTube channels from HowTo100M [84] and VLOG [36]. Each channel often links to other channels, and given a channel it is inexpensive to obtain a list of all its videos using the youtube-dl Python package.

In this paper, we considered numerous approaches to search for diverse, visually grounded videos. We ended up using an approach where we used YouTube’s recommended videos algorithm to suggest similar videos to YT-Temporal-180M. We went through all non-news and non-sports videos YT-Temporal-180M, and opened each video up in YouTube. For each other video that YouTube recommended, we retrieved its channel ID – giving us access to not just that video, but all other videos. This approach yielded 2 million channels, with 200 million videos among them.

### E.2. Filtering video IDs by channel

Given this (large) list of channels, each with many videos, we took steps to filter it further. We used the python c1d3 library to remove channels whose titles might not be in English. We then finetuned, and used, a language model to identify channels likely to have visually grounded videos, which we describe next.

In more detail, we selected 2000 videos, and asked workers on Mechanical Turk to rate their level of groundedness, their genre, and whether they had explicit content or not. The questions we asked are shown in Figure 8. We annotated 2k videos under this schema, and trained a model to predict the annotations given video metadata.

For model training, we used a slightly different setting to what we gave the crowdworkers. We trained a model to predict the labels, given a formatted list of 5 video titles from the same channel. During training, we made the weak-supervision assumption that all videos from a channel have exactly the same rating (as the video we annotated). This enabled us to collect 84k examples from our 2k annotations. The model we chose was T5-base model [92], which generates the labels left-to-right in text form (and which we converted automatically to a structured representation).

We then used this model to identify channels that seem especially promising. For each channel with at least 5 videos, we randomly sampled 8 sets of length-5 videos, and used the finetuned T5 model to classify them. We filtered out any channel that had at least 25% of likely non-English or irrelevant-English videos, any channel that had at least 25% of slideshows, and any channel that likely had racist or sexist content.

One side benefit of this model is that it allowed us to estimate our videos’ genre breakdown before downloading them. We found 1% Gaming videos, 11% News videos, 20% How-To videos, 20% ‘chatting’ videos, 5% sports videos, 5%

\${VIDEO}
Q1. How would you describe the role of English speech in the video?
a. This video doesn’t have spoken English, or if it does, it’s irrelevant to what’s going on in the video.
b. This video has English speech that describes, or adds onto, the visual content.
Q2. Select at least one genres of the video:
a. Gaming
b. News
c. How-to
d. Chatting
e. Sports
f. Music
g. Movies / Drama
h. Documentary
i. Miscellaneous
Q3. Select if any of the following are true:
a. A variety of objects are interacted with.
b. A variety of actions are performed.
c. A variety of scenes are performed.
d. This video is a slideshow.
e. This video contains racist or sexist content..

Figure 8: Video annotation. We had workers on Mechanical Turk annotate 2000 videos in our dataset with this questionnaire, allowing us to then train a model to identify suitable channels for our purpose.

Music videos, 3% Movies/Drama videos, 4% Documentary videos, and 31% Miscellaneous. The Gaming videos were then filtered out.

We used the classification model to create a budget for how many videos to download from each channel; with the aim to download more videos from likely more-grounded channels. Using the answers to Q3 (from Figure 8), we gave each channel 1 point for likely having ‘a variety of objects’, 2 points for ‘a variety of actions’, and 0.5 points for ‘a variety of scenes.’ We subtracted 3 points if it was likely to be a slideshow. (Likely-racist or sexist channels were already filtered out.) We then z-normalized and softmaxed the channel scores, and used the result as the channel-level budgets. Any channel with an aggregate ‘interestingness’ score of 1 standard deviation above the mean would then have a budget of 8x larger than the mean. We clipped the channel-level budgets to include at most 500 videos per channel.

This process (finally!) gave us 30 million YouTube video IDs that were likely to be high-quality.

### E.3. Filtering videos from their metadata

Last, we filtered and downloaded these videos using a filtering approach similar to [126]. We first retrieved the




Figure 9: An image prompt used in zero-shot audio classification. Here, “the sound of” is always inserted, and the word “birds” is one of the labels in ESC50 [89]. We consider one image prompt for each label in ESC50 (or whichever dataset we are using).

video metadata and used it to filter out ‘gaming’ videos. We then retrieved the video’s transcript, and filtered out any video without a ‘dense’ span of spoken words – defined as an interval of 30 seconds where at least 50 words are spoken. Additionally, we used the Python package `cld3` to filter out any transcript with a probability of less than 80% of being English. Last, we used a hidden feature in the YouTube API to download four thumbnails of the video. Using the image classification model from [126], we filtered out videos whose four thumbnails had an average cosine similarity of above 85%, or that contained fewer than 1 object from COCO.

Unlike [126], we did not use a sequence-to-sequence model to ‘translate’ spoken text to text that appears more stylistically like written English (i.e., by adding capitalization and punctuation, and removing filler words).

## F. Additional Experiments and Exploration

In this section, we briefly include additional experiments, showcasing our model’s performance on specific tasks that do not necessarily require multimodal script knowledge.

### F.1. Zero-shot Audio classification

We evaluate 🗯️RESERVE on the task of zero-shot audio classification, to study to what extent its learned audio representations can directly predict text-based labels. We conduct this evaluation on environmental sounds from ESC50 [89], urban sounds from US8K [97], and (as part of the privacy-minded exploration in Appendix A) celebrity voices from VoxCeleb2 [86].

We consider the format where we encode an audio input into a CLS level representation, and retrieve the most-similar label given a set of encoded options. We encode the audio input with our encoder, which takes in as input audio clips of length at most 1.6 seconds. For shorter audio clips (like many sounds in ESC50), we repeat them in time until their length is at least 1.6 seconds. For longer audio clips, we encode multiple CLS representations and then average the

Model	Prompting	Accuracy (%)		
		ESC50	US8K	VoxCeleb2
AudioClip		68.6	68.8	
🗯️RESERVE-L	Text-only.	41.6	60.2	10.8
	Image-only.	42.8	54.3	<b>13.3</b>
	Image and text.	<b>52.2</b>	<b>62.3</b>	9.6

Table 10: Zero-shot audio classification accuracies (%) on ESC50 [89], US8K [97], and VoxCeleb2 [86]. We compare our model with AudioClip [53], which was pretrained on supervised data from AudioSet [42]. Our 🗯️RESERVE performs well across the board, especially when given *both the image and the text* as a prompt – demonstrating its OCR capability.

resulting vectors.

We consider the following ways to encode the labels:

- Text-only.** Inspired by the prompt ‘a photo of’, which is used in CLIP’s zero-shot image classification task [91], we give 🗯️RESERVE’s joint encoder a blank image, with associated tokens `the sound of ${label}`. We do this once for each label, giving us a single ‘target’ vector for each possible label in the dataset.
- Image-only.** Inspired by YouTube videos of sound effects<sup>19</sup>, we created image-only prompts that suggest a sound (of the target class) is playing in the background. An example is shown in Figure 9. We encode each image with our joint encoder, and do this once for each label. We note that for VoxCeleb2, we use face images of celebrities rather than this image-based prompt, due to our interest in exploring whether models can perform person-level recognition due to the privacy issue (Appendix A.1.1).
- Image and text.** Here, we combine both of the above options: encoding one input for each label, using both the image and text prompt.

For each prompt, we append the token ‘MASKAUDIO’ and extract the hidden state from there, as our final representation for that label.

We present our results in Table 10. The results show, possibly surprisingly, that 🗯️RESERVE can perform optical character recognition over image prompts like Figure 9 – given just the image, its accuracy on ESC50 is higher than given just text. Its accuracy on ESC50 and US8K improves further when given both an image and text.

These results are slightly different for VoxCeleb2, which emphasizes long-tail recognition of people – something that might be more encyclopedic than semantic, and that we did not wish to optimize in this work. There, when given an image of a celebrity’s face, it demonstrates some capacity

<sup>19</sup>For instance, [youtu.be/VmgKryu4\\_k](https://youtu.be/VmgKryu4_k).

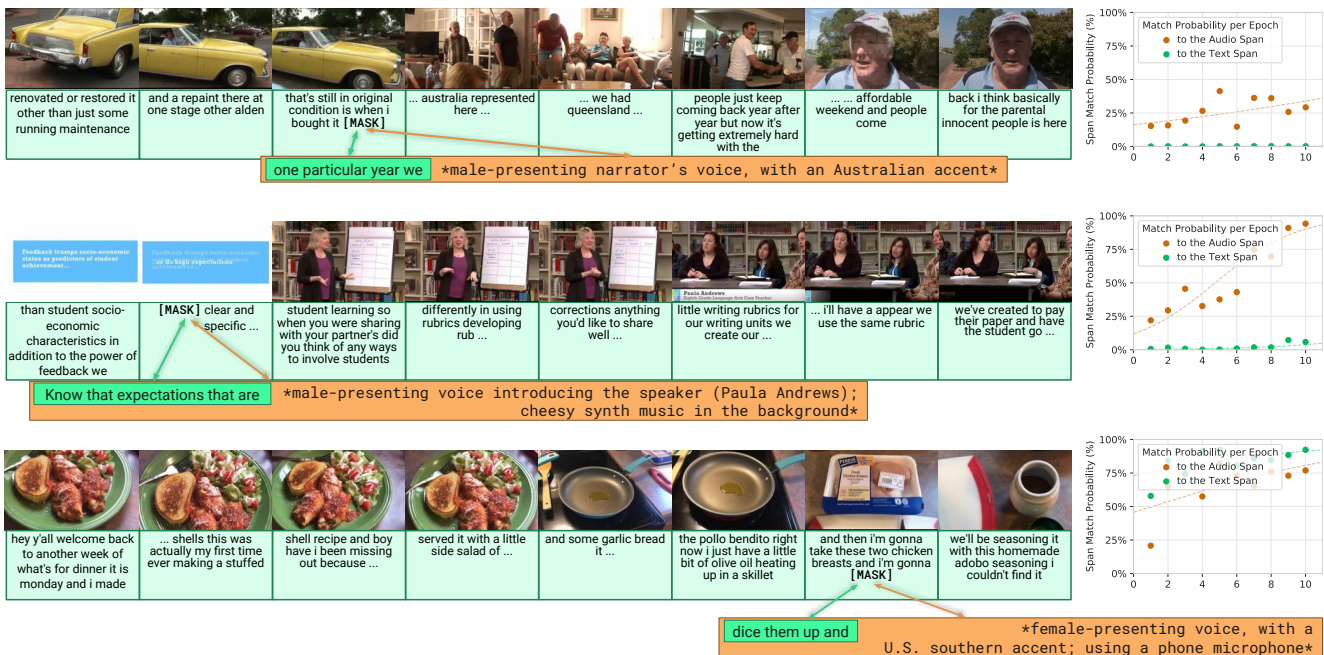


Figure 10: MASKed audio self-supervision on different examples. Similar to Figure 5, we show predictions from 🗨️RESERVE-B over the course of pretraining. Match performance increases over time. The audio prediction in the first row is perhaps made easier by the speaker’s australian accent. The audio prediction in the second row is perhaps easier due to the lecture-video setting. In the third row, both audio and text span prediction improves, with text being slightly favored in the end. This might be in part because of the truncation we do on audio (Section C.3) – the audio span is shorter than the text span of ‘dice them up and’ so as to not leak information, making prediction more challenging.

at linking it with one of their audio clips – a capacity that *decreases* if prompted with additional text. We suspect that this is due to interpreting the given text as spoken, for example, *Justin Bieber himself saying ‘the sound of Justin Bieber.’* On all celebrities, 🗨️RESERVE struggles versus recognition-focused models like CLIP [91] (Appendix A.1.1).

Overall, our model displays strong audio understanding ability. In comparison, AudioCLIP [53] (which is supervised on human-annotated labels from AudioSet [42]), performs 16% higher on ESC50, and 6.4% higher on US8K.

## F.2. Additional Qualitative Analysis

In Figure 10, we include an additional figure of examples, of the same format as Figure 5. The examples are chosen randomly – not by how much 🗨️RESERVE improved at retrieving their audio or text spans over the course of training.