

# Sequential Modeling Enables Scalable Learning for Large Vision Models

Yutong Bai<sup>1,2,\*</sup> Xinyang Geng<sup>1,\*</sup> Karttikeya Mangalam<sup>1</sup> Amir Bar<sup>1</sup>  
 Alan L. Yuille<sup>2</sup> Trevor Darrell<sup>1</sup> Jitendra Malik<sup>1</sup> Alexei A. Efros<sup>1</sup>

<sup>1</sup>UC Berkeley (BAIR)

<sup>2</sup>Johns Hopkins University

## Abstract

*We introduce a novel sequential modeling approach which enables learning a Large Vision Model (LVM) without making use of any linguistic data. To do this, we define a common format, “visual sentences”, in which we can represent raw images and videos as well as annotated data sources such as semantic segmentations and depth reconstructions without needing any meta-knowledge beyond the pixels. Once this wide variety of visual data (comprising 420 billion tokens) is represented as sequences, the model can be trained to minimize a cross-entropy loss for next token prediction. By training across various scales of model architecture and data diversity, we provide empirical evidence that our models scale effectively. Many different vision tasks can be solved by designing suitable visual prompts at test time.*

## 1. Introduction

Large language models (LLMs) such as GPT [10] and LLaMA [80] have taken the world by storm. What would it take to build a Large Vision Model (LVM)? From the animal world, we know that visual competences are not dependent on language. In particular, many experiments have shown that the visual world of non-human primates is remarkably similar to that of humans. So while the space of vision-language models such as LLaVA [53] is interesting and worthwhile to pursue, in this paper we seek an answer to a different question – how far can we go from pixels alone?

The key features of contemporary LLMs that we seek to emulate in LVMs are: 1) scaling in the presence of big data, and 2) flexible specification of tasks through prompting (in-context learning). How do we achieve this? As usual, there are three main components that must be specified:

**Data:** We want to exploit all the remarkable diversity in visual data. First of all, just raw unannotated images and videos. Next, we want to exploit the variety of annotated visual data sources that have been produced over the last

couple of decades – semantic segmentations, depth reconstructions, keypoints, multiple views of 3D objects, among others. We define a common format, “visual sentences”, in which to represent these different annotations without needing any meta-knowledge beyond the pixels. The total size of our training dataset is 1.64 billion images/frames.

**Architecture:** We use a large transformer architecture (3 billion parameters) trained on visual data represented as sequence of tokens, using a learned tokenizer that maps each image to a string of 256 vector-quantized tokens.

**Loss function:** We draw inspiration from the natural language community, where masked token modeling has given way to sequential autoregressive prediction. Once images/videos/annotated images can all be represented as sequences, we can train the model to minimize the cross-entropy loss for predicting the next token.

With this extremely simple design, we demonstrate some noteworthy behaviors:

- Appropriate scaling behavior as one increases model size and data size.
- Many different vision tasks can now be “solved” by designing suitable prompts at test time. While the results don’t show as high performance as bespoke, specifically-trained models, the fact that so many tasks are all addressed by a single vision model is quite encouraging.
- We see a clear benefit of the amount of unsupervised data on the performance on various standard vision tasks.
- We see a hint of an ability for general visual reasoning – handling out-of-distribution data, and performing novel tasks. But further investigation is needed.

## 2. Related Work

**Pretrained Vision Models.** The value of using pretrained models (such as ImageNet-pretrained AlexNet [45]) has been demonstrated as far back as 2015 in R-CNN [34], and it has since become standard practice in computer vision. Self-supervised pretraining was proposed as a way to vastly increase the amount of data available for pretraining [16, 25, 37, 61, 62, 99]. Unfortunately, this was not very successful, likely because the CNN-based architectures

\*Equal Contribution. Webpage: <https://yutongbai.com/lvm.html>.

of that time did not have enough capacity to absorb the data. With the introduction of Transformers [84], which have much higher capacity, researchers revisited self-supervised pretraining, and showed that transformer-based masked image reconstruction approaches, such as BEiT [7], MAE [38], SimMIM [91], MaskGIT [13] perform vastly better than their CNN-based counterparts [62]. Yet, despite their recent successes, current pretrained vision-only models have had trouble scaling up to the really large datasets, such as LAION [72].

**Multi-task Learning and In-context Learning.** From the classic one-model-per-task setups, computer vision is slowly moving toward a single model performing multiple different tasks. Various multi-task learning approaches [24, 40, 43, 73, 97] exist but they are typically limited to a fixed, pre-defined number of tasks. More recently, methods inspired by in-context learning in LLMs forgo any notion of tasks and instead let the model infer the task directly from the input prompt. For example, Visual Prompting [8, 87] takes in a task input/output example pair and a query image at test time, concatenates them into a single 2-by-2 image, and uses inpainting to generate the desired output. But, since the inpainting is performed using a variant of MAE [38], the same problems with scaling are inherited by these approaches.

**Auto-regressive Visual Models.** The idea of using auto-regressive models for synthesizing visual data goes back at least 70 years. Inspired by Shannon’s use of  $N$ -grams to synthesize language [74, 75], a number of works, starting with Attneave’s seminal 1954 paper [5], applied this idea to sequentially synthesizing pixels [28, 31, 39, 64], image patches [27], video frames [69], and motion capture data [4, 44, 48]. As deep models became popular, newer works replaced  $N$ -grams with RNNs or CNNs for pixel [81, 82] and video [66] synthesis. Most recently, transformer-based autoregressive visual generation methods have been proposed [15, 29, 94, 96], and, combined with language, have demonstrated impressive image synthesis results, e.g. Parti [95].

### 3. Data

“Data! Data! Data! I can’t make bricks without clay!”

SHERLOCK HOLMES

The key requirement of any Large Pre-trained Model is that it must be trained on vast amounts of data. For language models, very large and very diverse datasets are fairly easy to obtain. For instance, the popular Common Crawl repository [1] contains 250 billion web pages spanning the entire Web, is extremely diverse, and includes “natural demonstrations” like language translations, question answering, etc. In computer vision, we are still very far from having a data source of comparable size and diversity. One of the

central contributions of our work is the first step toward curating such a dataset that we call *Unified Vision Dataset v1 (UVDv1)*. To assemble it, we leverage many different sources of visual data: (1) unlabelled images, (2) images with visual annotations, (3) unlabelled videos, (4) videos with visual annotations, and (5) 3D synthetic objects. The unlabeled images, which represent over 80% of our data, capture a huge cross-section of our visual world, and provide the required diversity, at the cost of lower quality. Images with annotations have a much more constrained distribution, but are usually of higher quality. Video data is even more constrained (typically, to human-centric activities), but is an invaluable source of temporal data. Renderings of 3D synthetic objects are the lowest in diversity but can provide valuable hints about the behavior of 3D structures. Importantly, UVDv1 is a purely visual dataset, with no non-visual meta-data (e.g. text) included. All together, UVDv1 contains 1.64 billion images.

Another important difference from Large Language Models is that language data has a natural, unified one-dimensional structure for all the data – a stream of text. Unfortunately, this is not the case for visual data, with different sources all having different structures. In this work we propose *visual sentence* as a unified unit of visual data, which enables us to train scalable models from a diverse set sources. A visual sentence is simply a sequence containing one or more images followed by an end-of-sentence (EOS) token. Figure 1 shows how the various data sources are partitioned into visual sentences. In particular:

**Single images.** A single image itself represents the simplest form of a visual sentence – { image, EOS}. We use the filtered subset of 1.49 billion images [88] from the LAION 5B [71] dataset. This is by far the largest part of our data, comprising 88.5%.

**Image sequences.** A sequence of images is a natural form of visual sentence. We create such sequences by sourcing video data from a wide range of existing datasets [11, 12, 21, 35, 36, 46, 50, 51, 55, 57–59, 63, 68, 76–78, 92, 93]. Visual sentences of 16 frames are formed by randomly sampling the videos at three different strides (10, 20, and 30). In addition, we utilize synthetic 3D objects from the Objaverse Dataset [22] to generate object-centric multi-view sequences for a variety of objects. For each object, we sample one radius length between the object center and the camera from 1.5 to 2.2, and sample one constant elevation from  $-45$  degrees to  $45$  degrees, then traverse different views of the object by changing the azimuth with a step length of  $15$  degrees and render 24 views. We rendered 42000 such sequences in total for training and 8000 for testing. Finally, we can also represent images belonging to the same semantic category as being (part of) a sequence. We use categories from ImageNet, concatenating together groups of images (2,4,8, or 16) from the same category into a 16-image long

# Visual Sentences

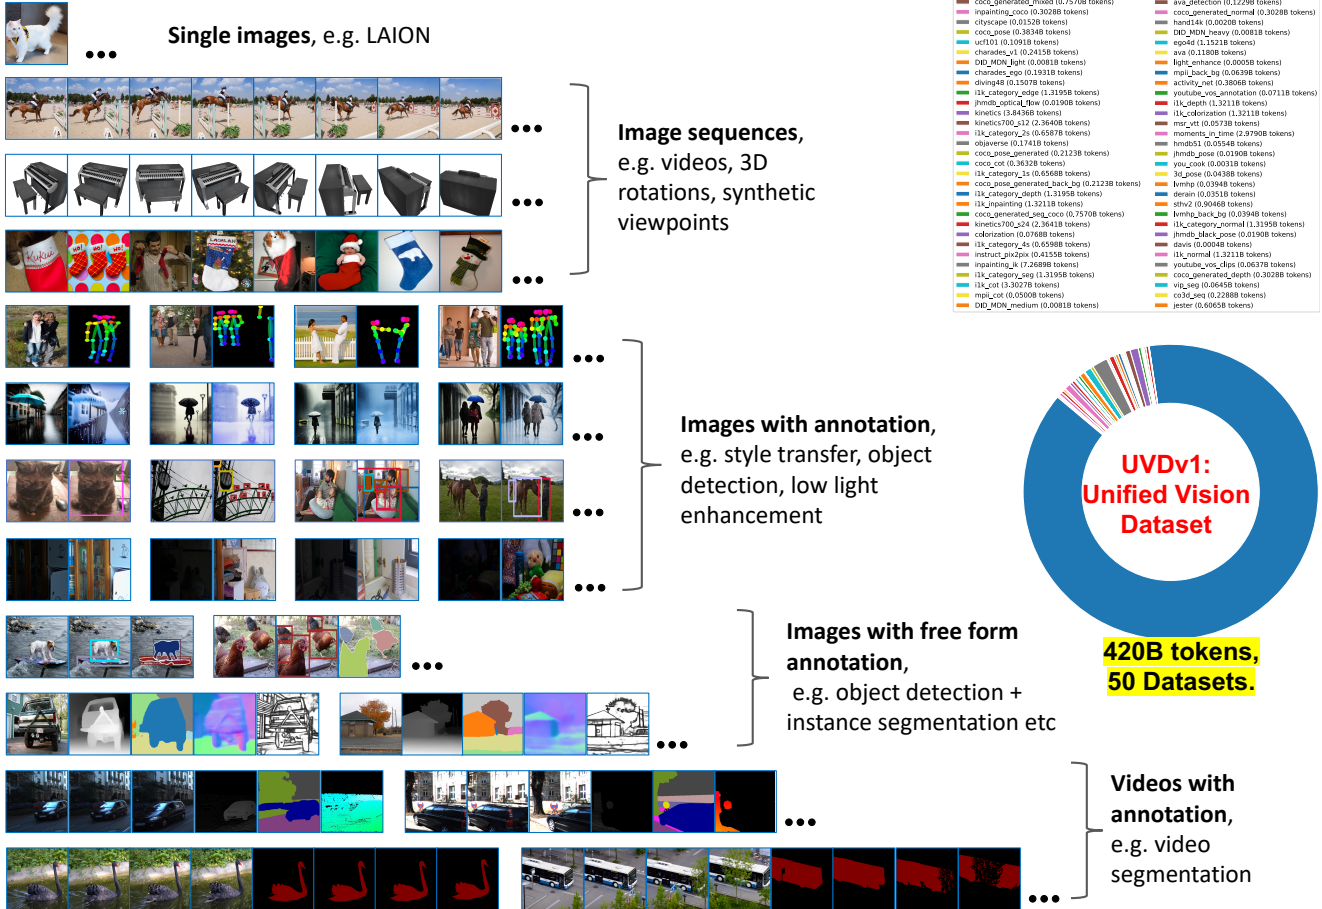


Figure 1. **Visual sentences** allow us to format diverse vision data into the unified structure of image sequences.

visual sentences.

**Images with annotations.** To handle different types of image annotations in a uniform way, we choose to represent all annotations as images. Some data types, e.g. semantic segmentation maps [100], edge maps [79], depth [65] and normal images [6], are already represented this way. For others we apply tailored methods for each specific annotation type: 1) Object Detection: We create annotations by overlaying a color-coded bounding box around each object, following the methodology in [14]; 2) Human Pose: Human skeletons are rendered in pixel space, adhering to the OpenPose format, utilizing MMPose [19]; 3) Depth Estimation, Surface Normal, and Edge Detection: given ImageNet and COCO images, we generate annotations in line with the protocols from [54]. 3) Style Transfer [9], De-rain [98], De-noise [85], Low Light Enhancement [89], and Stereo Datasets [33]: These are all represented as image pairs (e.g. input/output). 4) Colorization: We convert ImageNet images to greyscale, producing image pairs. 5) Inpainting: The process involves randomly adding black-colored boxes

in images to simulate corruption, resulting in image pairs. For all the above annotation types, we can create visual sentences by concatenating 8 image pairs of the same annotation type into a 16-image visual sentence. For datasets containing  $k$  different annotations for the same image we use a different approach: for each set of  $1 + k$  images (input plus  $k$  annotations), we randomly select  $m$  elements, where  $m \leq k + 1 \leq 16$ . These  $m$ -tuples are then concatenated to form visual sequences.

**Image sequences with annotations.** When converting annotated video data (VIPSeg [56], Hand14K [30], AVA [59], JHMDB [41]) to visual sentences, we apply two complementary strategies. The first is similar to how we treat image data with paired annotations: each visual sentence is constructed by concatenating frames with their annotations – {frame1,annot1,frame2,annot2,...}. The second method involves grouping multiple frames followed by their corresponding annotations – {frame1,frame2,annot1,annot2,...}.

We present a detailed summary of all the data sources, annotation type and data statistics of UVDv1 in the Appendix.



## 4. Approach

In this section, we describe the design of our autoregressive Large Vision Model. Unlike text data, which naturally exhibits discrete sequential structure, it is not straightforward to model image pixels in visual sentences. In this work, we take a two-stage approach: 1) train a large visual tokenizer (which operates on individual images) to convert each image into a sequence of visual tokens; 2) train an autoregressive transformer model on visual sentences, each represented as a sequence of tokens. We summarize our approach in Figure 2.

### 4.1. Image Tokenization

While the visual sentences exhibit a sequence structure between consecutive images, we don't have such natural sequence structure within an image. Therefore, in order to apply a transformer model to images, prior works typically do one of the following: either divide the image into patches in scan-line order, and treat that as a sequence [26], or use a pre-trained image tokenizer, such as VQVAE [83] or VQGAN [29], to cluster image features into a grid of discrete tokens, which, again, are turned into a sequence in scan-line order. We adopt the latter approach since the discrete categorical output from a model naturally forms a probabilistic distribution that one can easily sample from, enabling flexible conditional generation of new images within a visual sentence.

Specifically, we employ semantic tokens generated by a VQGAN model, a concept introduced by Esser et al [29]. This framework consists of an encoding and a decoding mechanism, featuring a quantization layer that assigns input images to a sequence of discrete tokens from an established codebook. Our encoders and decoders are constructed purely with convolutional layers. The encoder is equipped with several downsampling modules to contract the spatial dimension of the input, whereas the decoder is fitted with an equivalent series of upsampling modules to restore the image to its initial size. For a given image, our VQGAN tokenizer produces 256 discrete tokens.

It is important to note that our tokenizer operates on individual images independently, rather than on the entire visual sentence at once. This independence allows us to decouple the tokenizer training from the downstream Transformer model so that the tokenizer can be trained on a dataset of single images without having to consider the distribution of visual sentences.

**Implementation Details:** We adopt an off-the-shelf VQGAN architecture from Chang et al. [13]. We follow the exact configuration in Chang et al. [13], which uses a downsampling factor of  $f = 16$  and codebook size 8192. This means that for an image of size  $256 \times 256$ , our VQGAN tokenizer produces  $16 \times 16 = 256$  tokens where each can take 8192 different values. We found that using the results of an ImageNet pre-trained tokenizer did not generalize well

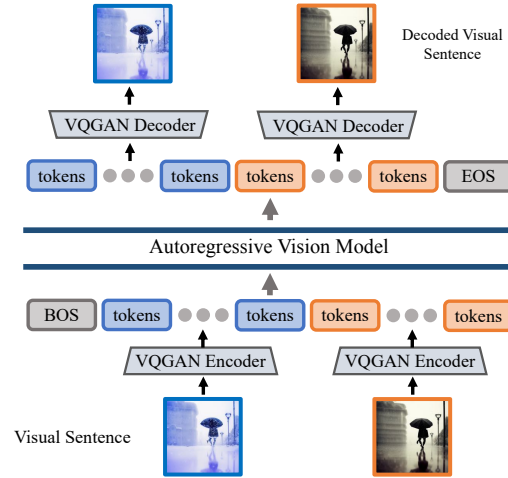


Figure 2. **Architecture of LVM.** We first convert individual images from a visual sentence into discrete tokens using a VQGAN encoder. The resulting tokens from all images are then concatenated into a 1D sequence, and fed into an autoregressive Transformer model to predict the next token in the sequence. The predicted visual tokens are decoded into images using the VQGAN decoder.

beyond ImageNet images. Therefore, we train our own tokenizer on a 1.5B subset of the LAION 5B dataset [71].

### 4.2. Sequence Modeling of Visual Sentences

After converting images into discrete tokens with VQGAN, we treat our visual sentence as a unified sequence by concatenating the discrete tokens from multiple images into a 1D sequence. Importantly, all visual sentences are treated equally – we do not make use of any special tokens to indicate particular tasks or formats. We train a causal Transformer model with the next token prediction objective using a cross-entropy loss, similar to the standard approach for language models [10]. Training the model the same way on all visual sentences enables the model to infer the relation between images from context instead of from task- or format-specific tokens. This gives the model an opportunity to generalize to other, unseen visual sentence structures.

**Implementation Details:** After tokenizing each image in a visual sentence into 256 tokens, we concatenate them to form a 1D sequence of tokens. On top of the sequences of visual tokens, our Transformer model is virtually the same as an autoregressive language model, so we adopt the Transformer architecture of LLaMA [80], a popular open-source language model with widely available implementations. We use a context length of 4096 tokens, which can fit 16 images under our VQGAN tokenizer. Similar to language models, we add a [BOS] (begin of sentence) token to the beginning of each visual sentence and an [EOS] (end of sentence) token to the end, and use sequence concatenation [18] during training time to improve efficiency. We train our model on our entire UVDv1 dataset (420 billion tokens) using one epoch (simple epoch training is standard in language models to avoid poten-



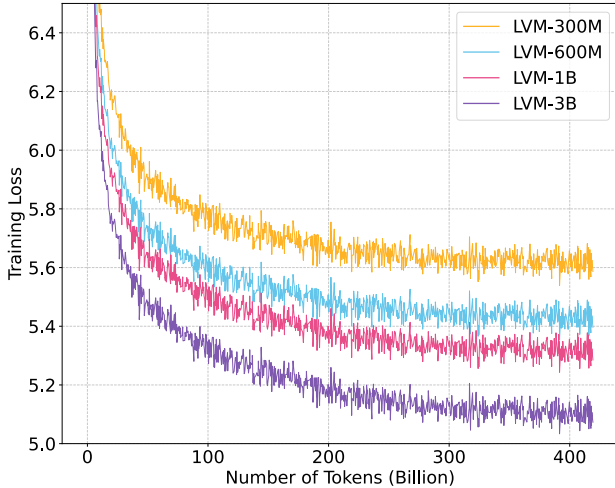


Figure 3. **Training loss for the 300M, 600M, 1B, and 3B models.** All models are trained on 420B tokens (1.64B images). The training scales well with model sizes. Zoom in for more details: notice how the shape of the 4 curves is very similar. This is the consequence of the deterministic training pipeline, so the loss fluctuations are determined by the intrinsic difficulty of the data stream.

tial overfitting). We train 4 models with different numbers of parameters: 300 million, 600 million, 1 billion and 3 billion, following the same training configurations. We provide the detailed training hyperparameters in Appendix 5.5.

### 4.3. Inference by Visual Prompting

Since the autoregressive Transformer in our model outputs a probability distribution of the next token conditioned on previous tokens, we can easily sample from this distribution to generate new visual tokens that complete a visual sentence. To use the model for downstream tasks, one can construct a partial visual sentence that defines a task at test time, and apply the model to generate the output. This is similar to in-context learning in language models [10] or visual prompting in computer vision [8, 39].

## 5. Experimental Results and Analysis

In this section, we evaluate the scaling abilities of our trained model, as well as its ability to understand and answer a range of diverse prompted tasks.

### 5.1. Scalability

We investigate the scaling behavior of our model in terms of the training loss and downstream task performance as we increase the model size as well as the number of tokens seen during training.

**Training loss.** We first inspect the training loss of LVM with different parameter sizes, which we present in Figure 3. Since all our models are trained for only one epoch on the dataset, the model sees a given data sample just once, and therefore the training loss at any point during training is very similar to the validation loss. One can observe that as

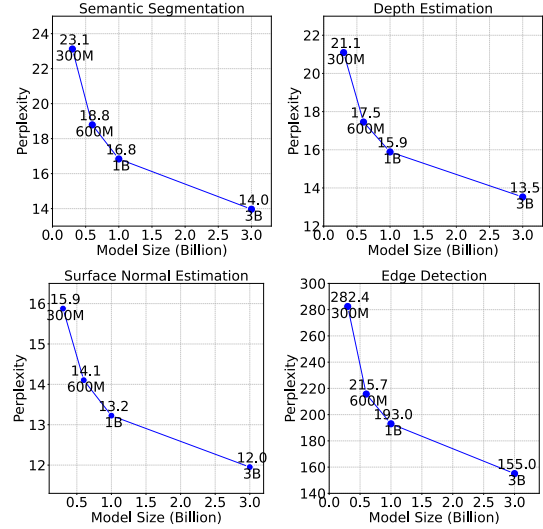


Figure 4. **Larger LVMs perform better on downstream tasks.** We evaluate LVMs of varying sizes on 4 different downstream tasks, following the 5 shot setting on the ImageNet validation set and report the perplexity. We find that perplexity decreases with larger models across all tasks, indicating strong scalability.

training progresses: 1) the training loss (perplexity) of the models, regardless of their size, continues to decrease; 2) as we increase the size of the model (parameter count), the loss decreases faster. These observations indicate that LVM shows strong scalability behavior with both larger models and more data.

**Scalability on downstream benchmarks.** While the LVM overall loss scales well during training, there is no guarantee that the better overall model would also perform better on a given specific downstream task. Therefore, we evaluate different sizes of models on 4 downstream tasks: semantic segmentation, depth estimation, surface normal estimation, and edge detection. We evaluate these tasks on the ImageNet validation set and generate all the annotations using the corresponding method described in Sec. 3. For each task, we give 5 pairs consisting of the inputs and corresponding ground-truth annotations as well as the query image as input prompt and evaluate the perplexity of the ground-truth annotation under our model’s prediction of the next 256 tokens (one image). We report the results in Figure 4. We see that larger models indeed attain lower perplexity across all tasks, showcasing that our scalable overall performance does transfer to a range of downstream tasks.

**Dataset ablation.** While LVM attains better performance with larger models and more data, it is natural to ask whether each data component we collect in UVDv1 helps. To answer this question, we conduct an ablation study on our dataset by training several 3B models on subsets of our dataset, and compare their performances on downstream tasks. We use the same 4 downstream tasks and settings as before and present the results in Figure 5. We observe that each




Figure 5. We evaluate the perplexity of 4 models trained on different subsets of our datasets on multiple tasks using the ImageNet validation set. All models are 3B parameters and all evaluations are conducted in the 5-shot setting. We can see that the model benefits from each single images, videos and annotations, demonstrating the importance of our training dataset diversity.

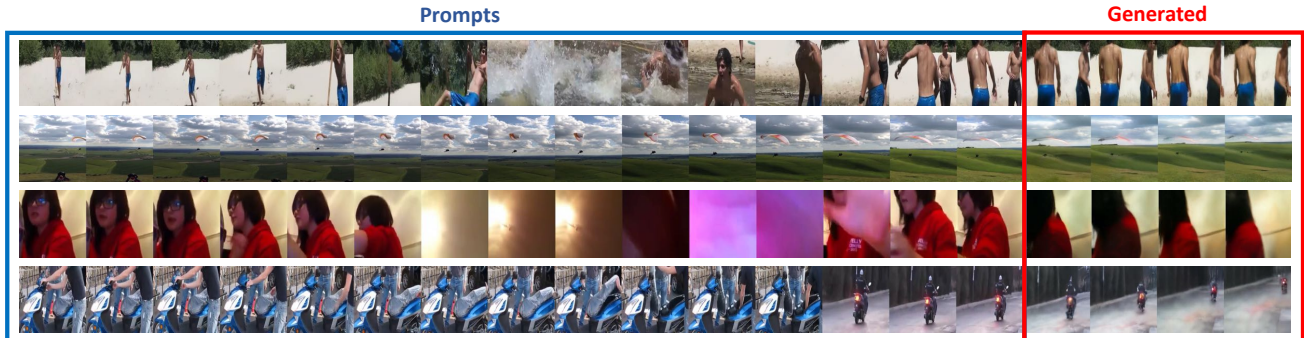


Figure 6. **Frame predictions.** LVM predicts the next frames (marked in red) given previous video frames as prompt. The results reveal that the LVM can predict the video frames while considering dynamic objects and camera motion.

data component contributes positively to the downstream tasks. LVM not only benefits from larger data, but also improves with more diversity in the dataset, which includes both annotated and unsupervised image and video data.

## 5.2. Sequential Prompting

We begin with the most intuitive and straightforward approach to visually prompt the LVM: sequential reasoning. Here the prompt construction is very simple: we present the model with a sequence of 7 images and ask it to predict the next image (256 tokens).

**Video frame prediction.** The most direct task for sequential prompting is video prediction. Figure 6 presents several next frame prediction examples, prompted by sequences from the Kinetics-700 validation set. At the top, 7 frame prompts (blue border) are followed by the predicted frame (red border). We observe a certain degree of inferential ability regarding spatial positioning, viewpoint and object understanding. Perplexity of prediction on Kinetics val set is 49.8. The last 4 rows show predictions with longer context (15 frames) and a longer prediction (4 frames). See Figures 16 to 21 in the Appendix for many more examples.

**Rotation and Category prediction.** The same type of simple sequential prompting can be used in other ways as well. For example, Figure 15 shows how prompting the model with a sequence of 3D rotations of a synthetic object around an arbitrary axis allows it to predict further rotation. Or we can think of a list of items of a given category as a sequence and predict other ideas in that same category, as shown in Figure 14. Note that, while the system was trained

on groups of images from the same ImageNet category, here the prompt consists of sketches, which have not been seen in any annotated data.

**Context length analysis.** Next we ask how much temporal context is required to accurately predict the subsequent frame? We assessed the model’s frame generation perplexity when prompted with a context of varying lengths (1 to 15 frames). On the Kinetics-700 val set, we see a clear improvement in perplexity from 1 to 11 frames after which it stabilizes (from 62.1  $\rightarrow$  48.4).

## 5.3. Analogy Prompting

Our study progresses by evaluating a more complex prompting structure, which we call ‘Analogy Prompting’. This method challenges the model to comprehend analogies of arbitrary length and complexity, thereby testing its advanced interpretative abilities.

**Qualitative Results.** Figure 7 shows a sampling of qualitative results with analogy prompting on a number of tasks. The prompts consist of a sequence of 14 images giving examples of various tasks, followed by a 15th query image. Given each prompt, the next image predicted is the result. The top part of the figure shows several example prompts defining tasks that were part of the training set (but these actual images were never seen at training). The bottom part of the figure demonstrates generalization to tasks never shown at training. See Appendix for many more qualitative examples.

**Unseen Tasks and Dataset.** We present the results for keypoint detection on Pascal 3D+ [90], evaluated using the standard Percentage of Correct Keypoints (PCK) metric with




Figure 7. **In and out of distribution prompting examples.** Every row is a prompt that contains a sequence of images interleaved with annotations, followed by a query. The last image is predicted by the model (marked in red). The last 5 rows show examples where the query image is out of distribution (painting, sketch, etc) for the task it was trained for.

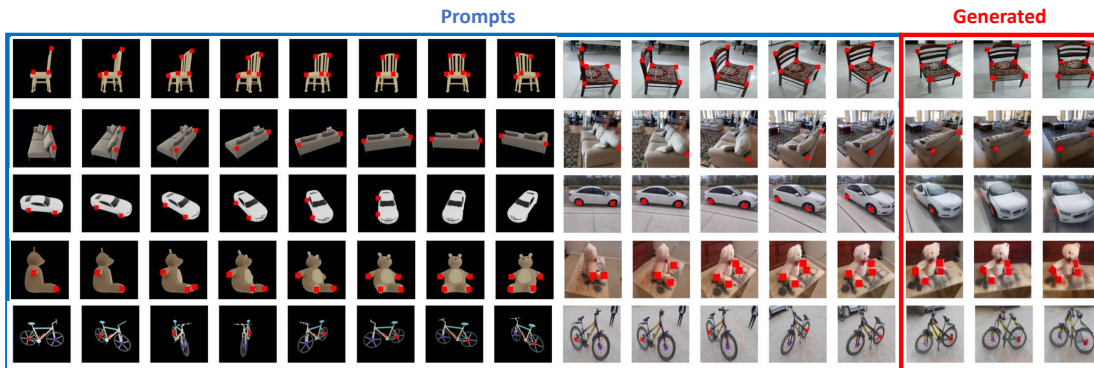


Figure 8. **Task Compositionality.** Examples of prompts that combine two different tasks – object rotation and keypoint tracking.

Model	Foreground Segmentation $\uparrow$				Single Object Detection $\uparrow$				Colorization $\downarrow$	
	Split 0	Split 1	Split 2	Split 3	Split 1	Split 2	Split 3	Split 4	MSE	LPIPS
MAE (IN-1k)	1.92	6.76	3.85	4.57	1.37	1.98	1.62	1.62	1.13	0.87
MAE-VQGAN (IN-1k)	2.22	7.07	5.48	6.28	3.34	3.21	2.80	2.80	3.31	0.75
MAE (CVF)	17.42	25.70	18.64	16.53	5.49	4.98	5.24	5.84	<b>0.43</b>	0.55
MAE-VQGAN (CVF)	27.83	30.44	26.15	24.25	24.19	25.20	25.36	25.23	0.67	<b>0.40</b>
Ours	<b>48.94</b>	<b>51.29</b>	<b>47.66</b>	<b>50.82</b>	<b>48.25</b>	<b>49.60</b>	<b>50.08</b>	<b>48.92</b>	0.51	0.46

Table 1. **Comparison with Visual Prompting [8].** For Foreground Segmentation and Single Object Detection, we report the *mIOU* score. For Colorization, we report the *MSE* and *LPIPS*.

a of threshold 0.1. Remarkably, LVM achieves a PCK of 81.2 without training on this dataset, demonstrating impressive generalization capabilities. In comparison, we show

some existing task-specific model: StackedHourglass [60] scores 68.0 PCK, MSS-Net [42] achieves 68.9 PCK, and StarMap [101] registers 78.6 PCK.

**Comparison with Visual Prompting.** The closest approach to ours that also allows for defining arbitrary tasks is Visual Prompting [8]. In Table 1, we compare various visual prompting models on few-shot segmentation, object detection, and colorization tasks. Note that our sequential LVM beats previous approaches on almost all tasks.



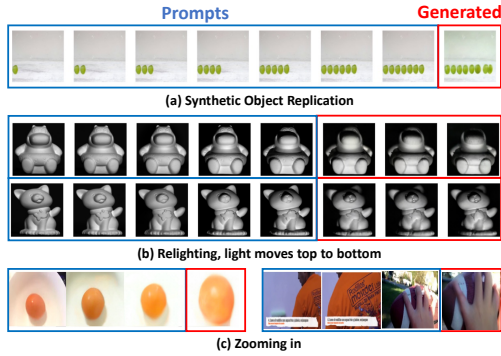


Figure 9. **Miscellaneous Prompting.** A variety of simple vision tasks, such as object replication (top), relighting (middle), and zooming in (bottom), can be simply specified via a suitably chosen visual sentence prompt that expresses the task to the LVM.




Figure 10. **What comes next?** Tasks that are not always easily describable in language.




Figure 11. **“Sparks of AGI?”** 😊 We prompt LVM with a masked reasoning visual sentence to infer the solution for non-verbal reasoning questions which are prevalent in IQ tests (masked image, second from the right). Left are successful cases and right is failure. We find that the model often infers and applies the abstract visual pattern correctly. So, we graciously hand over to you, our gentle reader, the task of pondering whether our modest LVM also exhibits the much-vaunted ‘Sparks of AGI’.

**Task Compositing.** Figure 8 demonstrates compositing several tasks together within a single prompt. Here, we demonstrate the rotation task together with the novel key-point correspondence task and ask the model to continue the pattern. The model can successfully combine these two at test time, demonstrating some degree of compositionality.

#### 5.4. Miscellaneous Prompts

Here we try to see how far we can push our model by offering it various prompts it has not seen before. Figure 9 shows a few such prompts that happened to work reasonably well. Figure 10 shows some prompts which are not easily describable by words – these are the type of tasks where LVMs might eventually outshine LLMs.

In Figure 11, we show initial qualitative results on a typical visual reasoning question as found on non-verbal human IQ tests (Raven’s Progressive Matrices [67]). With considerable squinting, one could imagine the LVM having a

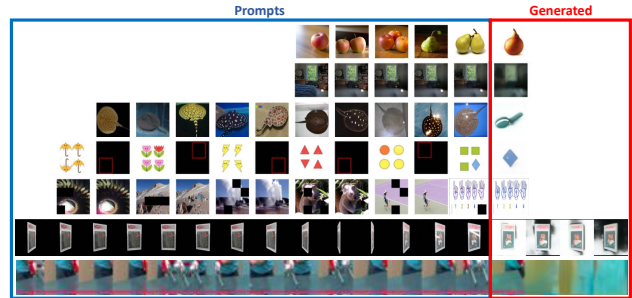


Figure 12. **Failure Cases.** This figure illustrates seven examples of failure cases: (1) Task confusion: the model interprets the counting task as style transfer, resulting in a pear-shaped apple. (2) Task entanglement: the model perceives the prompt as high frequencies from the entire image, rather than just from some parts of image (objects). (3) Wrong instance: the model correctly identifies rotation but mistakenly generates a brush instead of a manta. (4) Outlier detection versus generation: rather than detecting an outlier, the model directly generates it. (5) Digits: the model fails to represent digit sequences. (6) Tokenizer: poor performance with some out-of-distribution synthetic data generation (wrong background). (7) Sequence degeneration: sometimes frame predictions go astray.

latent ability for grasping abstract visual patterns and applying them to extrapolate the shown visual sequence. We also perform perplexity analysis on 10 classic Raven 5-way multiple-choice Matrices in Table 2, choosing the answer with lowest perplexity under our model. It shows that our results are better than chance. To further understand this process, we generate 900 new Raven-like synthetic tests to study individual factors: color, shape, and size following [8].

	Raven IQ matrices	shape	color	size
chance	20%	33%	33%	50%
Ours	30%	45%	42%	94%

Table 2. Performance on 10 Raven’s Progressive Matrices and 900 synthetic reasoning tasks by lowest perplexity under our model.

#### 5.5. Failure Cases

Figure 12 shows some typical failure cases of the current model. One common element, the use of visual prompt to define a task is often under-constrained (more so than in language, since images are very high-dimensional), or the requested task might be beyond the capabilities of the current system. Other, more mundane failures involve issues with the tokenizer and lack of high-quality video training data.

**Acknowledgements:** We are grateful to many friends and colleagues for discussions and comments on this work, including Yossi Gandelsman, Aleksander Holynski, Angjoo Kanazawa, Qianqian Wang, Sophia Koepke, Quoc Le, Chen Liang, Zichao Li, Ekin D. Cubuk, Assaf Shocher, Amir Zamir, Carl Vondrick, Ludwig Schmidt, Aviral Kumar, Xiaolong Wang and Miki Rubinstein. We thank the support from Google during Yutong’s part-time internship. The work has been supported, in part, by ONR MURI N0001 4-22-1-2773, N00014-21-1-2801, and N00014-21-1-2812, ERC HOLI, Apple Graduate Fellowship to Yutong Bai, and compute donation via the Google TPU Research Cloud. Darrell’s group is supported in part by the US DoD, BAIR&BDD industrial consortia.

## References

- [1] Common crawl repository. <https://commoncrawl.org/>.
- [2] Abdelrahman Abdelhamed, Stephen Lin, and Michael S Brown. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1692–1700, 2018.
- [3] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [4] Okan Arıkan and D. A. Forsyth. Interactive motion generation from examples. *ACM Trans. Graph.*, 21(3):483–490, 2002.
- [5] Fred Attneave. Some informational aspects of visual perception. *Psychological review*, 61(3):183, 1954.
- [6] Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. Estimating and exploiting the aleatoric uncertainty in surface normal estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13137–13146, 2021.
- [7] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. In *International Conference on Learning Representations*, 2021.
- [8] Amir Bar, Yossi Gandelsman, Trevor Darrell, Amir Globerson, and Alexei Efros. Visual prompting via image inpainting. *Advances in Neural Information Processing Systems*, 35:25005–25017, 2022.
- [9] Tim Brooks, Aleksander Holynski, and Alexei A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, 2023.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [11] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015.
- [12] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019.
- [13] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11315–11325, 2022.
- [14] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019.
- [15] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *International Conference on Machine Learning*, pages 1691–1703. PMLR, 2020.
- [16] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15750–15758, 2021.
- [17] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. 2022.
- [18] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.
- [19] MMPose Contributors. Openmmlab pose estimation toolbox and benchmark. <https://github.com/open-mmlab/mmpose>, 2020.
- [20] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [21] Pradipto Das, Chenliang Xu, Richard F Doell, and Jason J Corso. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2634–2641, 2013.
- [22] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023.
- [23] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [24] Carl Doersch and Andrew Zisserman. Multi-task self-supervised visual learning. In *Proceedings of the IEEE international conference on computer vision*, pages 2051–2060, 2017.
- [25] 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 (ICCV)*, 2015.
- [26] 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. In *International Conference on Learning Representations*, 2020.
- [27] Alexei A. Efros and William T. Freeman. Image quilting for

- texture synthesis and transfer. In *Proceedings of SIGGRAPH 2001*, 2001.
- [28] Alexei A. Efros and Thomas K. Leung. Texture synthesis by non-parametric sampling. In *IEEE International Conference on Computer Vision*, pages 1033–1038, Corfu, Greece, 1999.
- [29] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12873–12883, 2021.
- [30] Alireza Fathi, Xiaofeng Ren, and James M Rehg. Learning to recognize objects in egocentric activities. In *CVPR 2011*, pages 3281–3288. IEEE, 2011.
- [31] David Donovan Garber. *Computational Models for Texture Analysis and Texture Synthesis*. PhD thesis, University of Southern California, Image Processing Institute, 1981.
- [32] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012.
- [33] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [34] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Region-based convolutional networks for accurate object detection and segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 38(1):142–158, 2015.
- [35] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The” something something” video database for learning and evaluating visual common sense. In *Proceedings of the IEEE international conference on computer vision*, pages 5842–5850, 2017.
- [36] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022.
- [37] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- [38] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022.
- [39] Aaron Hertzmann, Charles E Jacobs, Nuria Oliver, Brian Curless, and David H Salesin. Image analogies. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 327–340, 2001.
- [40] Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppula, Daniel Zoran, Andrew Brock, Evan Shelhamer, et al. Perceiver io: A general architecture for structured inputs & outputs. In *International Conference on Learning Representations*, 2021.
- [41] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M. J. Black. Towards understanding action recognition. In *International Conf. on Computer Vision (ICCV)*, pages 3192–3199, 2013.
- [42] Lipeng Ke, Ming-Ching Chang, Honggang Qi, and Siwei Lyu. Multi-scale structure-aware network for human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 713–728, 2018.
- [43] Iasonas Kokkinos. Ubernet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6129–6138, 2017.
- [44] Lucas Kovar, Michael Gleicher, and Frederic Pighin. Motion Graphs. In *Proceedings of SIGGRAPH '02*, San Antonio, TX, 2002.
- [45] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [46] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *2011 International conference on computer vision*, pages 2556–2563. IEEE, 2011.
- [47] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J Black, and Peter V Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6050–6059, 2017.
- [48] Jehee Lee, Jinxiang Chai, Paul Reitsma, Jessica K. Hodgins, and Nancy Pollard. Interactive control of avatars animated with human motion data. *ACM Transactions on Graphics (TOG)*, 21(3):491 – 500, 2002.
- [49] Jianshu Li, Jian Zhao, Yunchao Wei, Congyan Lang, Yidong Li, Terence Sim, Shuicheng Yan, and Jiashi Feng. Multiple-human parsing in the wild. *arXiv preprint arXiv:1705.07206*, 2017.
- [50] Yingwei Li, Yi Li, and Nuno Vasconcelos. Resound: Towards action recognition without representation bias. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 513–528, 2018.
- [51] Yixuan Li, Lei Chen, Runyu He, Zhenzhi Wang, Gangshan Wu, and Limin Wang. Multisports: A multi-person video dataset of spatio-temporally localized sports actions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13536–13545, 2021.
- [52] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.



- [53] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024.
- [54] Shikun Liu, Linxi Fan, Edward Johns, Zhiding Yu, Chaowei Xiao, and Anima Anandkumar. Prism: A vision-language model with an ensemble of experts. *arXiv preprint arXiv:2303.02506*, 2023.
- [55] Joanna Materzynska, Guillaume Berger, Ingo Bax, and Roland Memisevic. The jester dataset: A large-scale video dataset of human gestures. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*, pages 0–0, 2019.
- [56] Jiayu Miao, Xiaohan Wang, Yu Wu, Wei Li, Xu Zhang, Yunchao Wei, and Yi Yang. Large-scale video panoptic segmentation in the wild: A benchmark. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [57] Mathew Monfort, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, Tom Yan, Lisa Brown, Quanfu Fan, Dan Gutfreund, Carl Vondrick, et al. Moments in time dataset: one million videos for event understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(2):502–508, 2019.
- [58] Mathew Monfort, Bowen Pan, Kandan Ramakrishnan, Alex Andonian, Barry A McNamara, Alex Lascelles, Quanfu Fan, Dan Gutfreund, Rogério Schmidt Feris, and Aude Oliva. Multi-moments in time: Learning and interpreting models for multi-action video understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9434–9445, 2021.
- [59] Naila Murray, Luca Marchesotti, and Florent Perronnin. Ava: A large-scale database for aesthetic visual analysis. In *2012 IEEE conference on computer vision and pattern recognition*, pages 2408–2415. IEEE, 2012.
- [60] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 483–499. Springer, 2016.
- [61] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pages 69–84. Springer, 2016.
- [62] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.
- [63] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 724–732, 2016.
- [64] Kris Popat and Rosalind W. Picard. Novel cluster-based probability model for texture synthesis, classification, and compression. In *Proc. SPIE Visual Comm. and Image Processing*, 1993.
- [65] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12179–12188, 2021.
- [66] MarcAurelio Ranzato, Arthur Szlam, Joan Bruna, Michael Mathieu, Ronan Collobert, and Sumit Chopra. Video (language) modeling: a baseline for generative models of natural videos. *arXiv preprint arXiv:1412.6604*, 2014.
- [67] John C Raven. Standardization of progressive matrices, 1938. *British Journal of Medical Psychology*, 1941.
- [68] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10901–10911, 2021.
- [69] Arno Schodl, Richard Szeliski, David H. Salesin, and Irfan Essa. Video textures. In *Proceedings of SIGGRAPH '00*, pages 489–498, 2000.
- [70] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.
- [71] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: An open large-scale dataset for training next generation image-text models, 2022.
- [72] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022.
- [73] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, 31, 2018.
- [74] Claude E. Shannon. A mathematical theory of communication. *Bell Sys. Tech. Journal*, 27, 1948.
- [75] Claude E Shannon. Prediction and entropy of printed english. *Bell system technical journal*, 30(1):50–64, 1951.
- [76] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 510–526. Springer, 2016.
- [77] Gunnar A Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. Charades-ego: A large-scale dataset of paired third and first person videos. *arXiv preprint arXiv:1804.09626*, 2018.
- [78] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.

- [79] X. Soria, E. Riba, and A. Sappa. Dense extreme inception network: Towards a robust cnn model for edge detection. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1912–1921, Los Alamitos, CA, USA, 2020. IEEE Computer Society.
- [80] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [81] Aaron van den Oord, Nal Kalchbrenner, Lasse Espeholt, koray kavukcuoglu, Oriol Vinyals, and Alex Graves. Conditional image generation with pixelcnn decoders. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2016.
- [82] Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 1747–1756, New York, New York, USA, 2016. PMLR.
- [83] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- [84] 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.
- [85] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th International Conference on Machine Learning*, page 1096–1103, 2008.
- [86] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems*, pages 10506–10518, 2019.
- [87] Xinlong Wang, Wen Wang, Yue Cao, Chunhua Shen, and Tiejun Huang. Images speak in images: A generalist painter for in-context visual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6830–6839, 2023.
- [88] Ryan Webster, Julien Rabin, Loic Simon, and Frederic Jurie. On the de-duplication of laion-2b, 2023.
- [89] Chen Wei, Wenjing Wang, Wenhao Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560*, 2018.
- [90] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2014.
- [91] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simsim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9653–9663, 2022.
- [92] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5288–5296, 2016.
- [93] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*, 2018.
- [94] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved vqgan. In *International Conference on Learning Representations*, 2021.
- [95] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. Scaling autoregressive models for content-rich text-to-image generation, 2022.
- [96] Yingchen Yu, Fangneng Zhan, Rongliang Wu, Jianxiong Pan, Kaiwen Cui, Shijian Lu, Feiying Ma, Xuansong Xie, and Chunyan Miao. Diverse image inpainting with bidirectional and autoregressive transformers. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 69–78, 2021.
- [97] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3712–3722, 2018.
- [98] He Zhang and Vishal M Patel. Density-aware single image de-raining using a multi-stream dense network. In *CVPR*, 2018.
- [99] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016.
- [100] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127(3):302–321, 2019.
- [101] Xingyi Zhou, Arjun Karapur, Linjie Luo, and Qixing Huang. Starmap for category-agnostic keypoint and viewpoint estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 318–334, 2018.