
Learning Visual Prior via Generative Pre-Training

Anonymous Author(s)

Affiliation
Address
email

1 1 Appendix

2 1.1 Examples of Training Sequences

- 3 Here, we give some examples of various types of training sequences on different datasets:

Human Pose (COCO):

key point; multiple instances; large; 1; 18; person; [a 190 120 b 266 146 c 318 143 d 385 232 e 338 269 f 214 150 g 0 0 h 0 0 i 312 280 j 365 296 k 359 420 l 258 283 m 194 344 n 301 383 o 197 100 p 181 103 q 234 84 r 0 0]

Human Pose (CrowdPose):

key point; multiple instances; large; 2; 14; person, person; [a 312 201 b 306 200 c 311 232 d 269 214 e 298 257 f 231 206 g 296 275 h 307 275 i 251 244 j 271 235 k 274 292 l 283 295 m 304 153 n 310 191] [a 179 247 b 165 245 c 164 313 d 160 315 e 221 316 f 207 279 g 155 343 h 144 366 i 242 337 j 240 367 k 210 431 l 300 418 m 172 176 n 177 227] key point; multiple instances; large; 2; 14; person, person; [a 240 178 b 304 168 c 228 239 d 0 0 e 261 236 f 0 0 g 251 296 h 289 296 i 0 0 j 0 0 k 0 0 l 0 0 m 261 92 n 272 156] [a 314 160 b 363 158 c 274 232 d 356 264 e 224 260 f 271 263 g 298 315 h 341 324 i 0 0 j 332 442 k 0 0 l 0 0 m 287 64 n 333 133]

Instance Mask:

mask; multiple instances; medium; 1; 0; clock; [m0 224 291 m1 226 299 m2 227 306 m3 228 313 m4 233 320 m5 238 325 m6 245 329 m7 252 332 m8 259 334 m9 266 335 m10 274 333 m11 281 330 m12 288 327 m13 293 323 m14 299 318 m15 303 312 m16 305 305 m17 307 298 m18 310 291 m19 308 284 m20 307 276 m21 303 269 m22 299 263 m23 295 257 m24 288 254 m25 280 251 m26 273 250 m27 266 249 m28 259 249 m29 252 251 m30 246 256 m31 240 260 m32 235 265 m33 229 270 m34 227 277 m35 225 284]

Object Centric Bounding-Box:

box; object centric; large; 1; 0; castle; [xmin 236 ymin 142 xmax 413 ymax 232]

4

5 1.2 Implementation Details

All experimental evaluations were conducted on eight NVIDIA Tesla V100-32GB GPUs using PyTorch. In order to include special words, we created a new vocabulary containing a total of 30,769 words based on a standard vocabulary. To optimize computational efficiency and memory utilization, we utilized the DeepSpeed framework. To serialize visual locations, we first resized the long side of each image to a length of 512 pixels and then shifted the image content to the center by padding the short side to a length of 512 pixels. As a result, the number of bins m was set to 512. The flag of **[Size]** indicates the average area of all instances in the image and we set the flag according to the rule:

$$\begin{cases} \text{"small"} & \text{average area} < 32^2 \\ \text{"medium"} & 32^2 \leq \text{average area} < 96^2 \\ \text{"large"} & \text{average area} \geq 96^2 \end{cases} .$$

- 6 We omitted person instances with fewer than five keypoints. To enable continuous generation, we designed and
7 trained models based on the prompt format (b). Specifically, VISORGPT[†] (a&b) and VISORGPT (a&b) were
8 trained using the same number of sequences as VISORGPT[†] (a) and VISORGPT (a), respectively. The only
9 difference is that we randomly utilized prompt format (a) or (b) to construct each training sequence.

10 During the evaluation stage, we set the maximum sequence length of our model (VISORGPT) to 256 tokens to
 11 ensure efficient inference. In the ablation studies, we added special words only to the [Coordinate] term, and we
 12 reported the average KL divergence between the location and shape priors learned by VISORGPT and those in
 13 the real world. Since training large-scale language models is time- and resource-consuming, we trained only
 14 three types of VISORGPT with respect to GPT-2 (base, medium, large) with a maximum token length of 256 in
 15 50,000 iterations on COCO (Box) data.

16 1.3 Evaluation Details

17 To estimate discrete visual prior from VISORGPT, we infer a series of sequences via prompting as below:

Code in Python:
 18 `f"box; multiple instances; random.choice(['small', 'medium', 'large']);
 random.randint(2, 10); 0; category name,"`

19 To ensure that each category in a given dataset is sufficiently represented in the sequence data used for estimating
 20 the visual prior, we specify a minimum number of sequences in which each category must appear. Table 1
 provides an overview of the predicted sequences that are used for evaluation.

Table 1: Details about the predicted sequences for evaluation.

Datasets	#Categories	#Predicted Seq.	Min #Seq. Per Category
Open Images (Box)	600	48,000	~80
Objects365 (Box)	365	29,200	~80
COCO (Box)	80	6,400	~80

21
 22 In our study, we adopt the Kullback-Leibler divergence to quantify the similarity between two given discrete
 23 distributions. Specifically, let p and q denote the estimated probabilistic priors derived from the real-world data
 24 and the VISORGPT, respectively. The degree of similarity between these two distributions can be computed as:

$$\text{KL}(p||q) = p \log(p/q). \quad (1)$$

25 1.4 Visualization

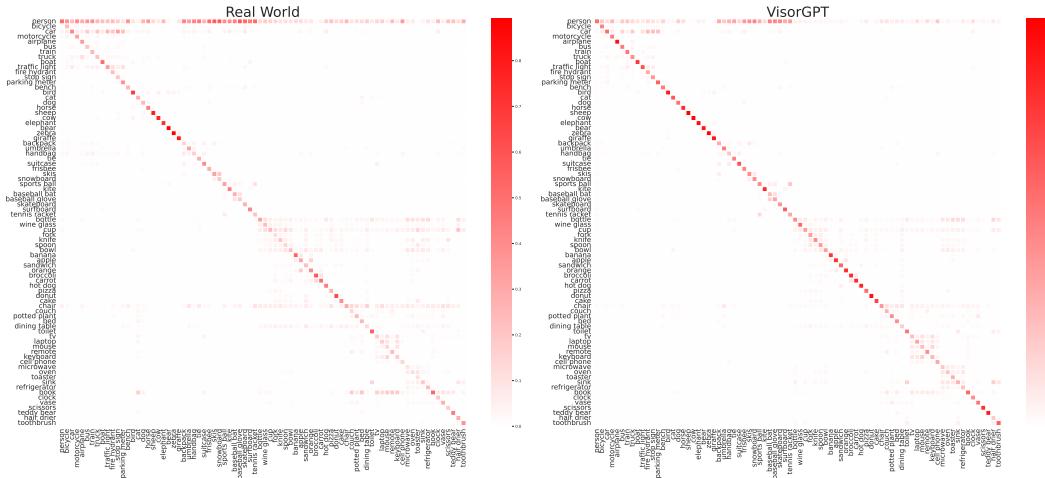


Figure 1: Relation among 80 categories on COCO.

26 **Relation Prior of COCO.** Fig. 1 illustrates the comparison between the real and learned relation prior among
 27 80 categories on the COCO dataset. As can be observed, there is a high degree of similarity between the two
 28 relation matrices.

29 **More Visual Comparison.** We provide more comparison of visual prior between the real world and one learned
 30 by our VISORGPT and failure cases on COCO dataset in Fig. 2.

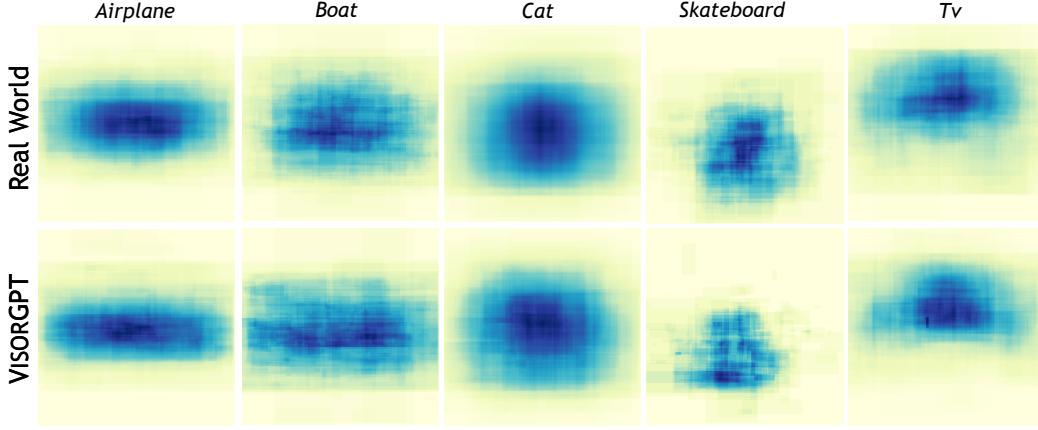
31 **Continuous Generation.** Fig. 3 presents a set of examples showcasing continuous generation based on
 32 the current scene. Notably, in each row, the proposed VISORGPT is able to successfully complete a scene

33 that involves many individuals annotated with 14/18 keypoints or objects with bounding boxes, based on the
34 information provided in the corresponding scene depicted in the previous columns.

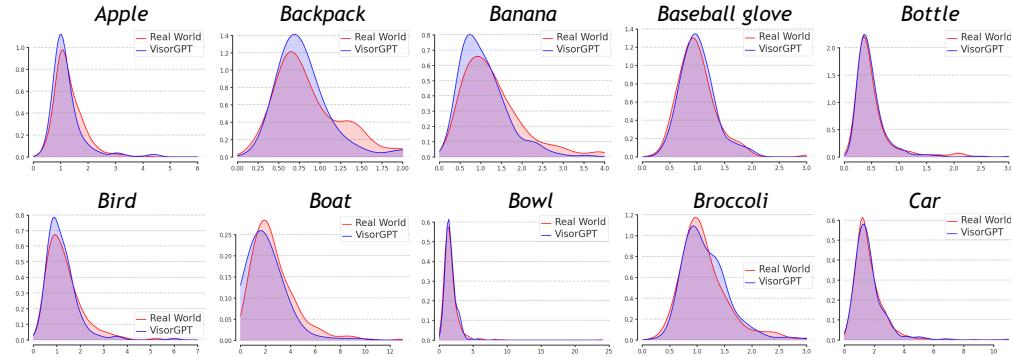
35 Figs. 4 and 5 present more visualization results.

36 **1.5 Broader Impact**

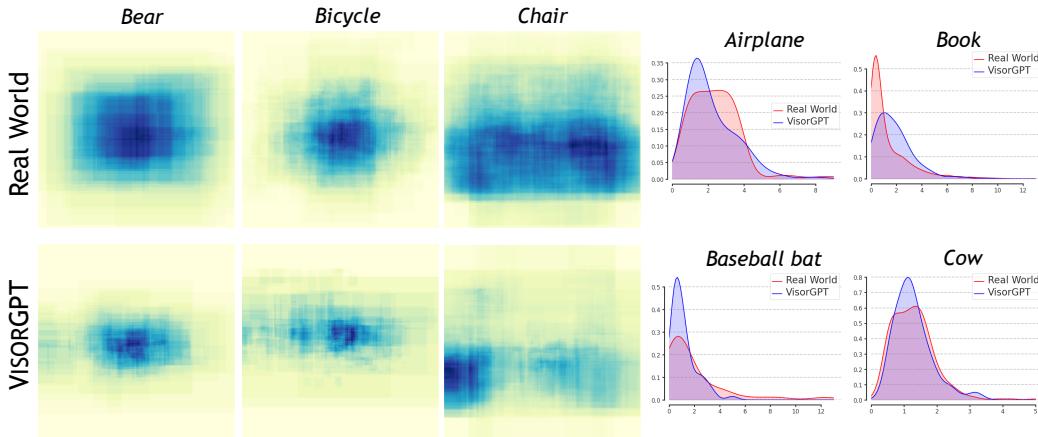
37 One of the key advantages of our VISORGPT is that it has learned the visual prior that can be used to sample
38 various types of sequences. This allows for a high degree of customization in terms of the spatial conditions
39 of bounding boxes, human poses, and instance masks, from many aspects such as object size, number of
40 instances, and classes. In this way, the generated spatial conditions can be used to continuously synthesize paired
41 image-box/pose/mask data such that we can potentially train more generalized visual intelligence models that
42 are capable of handling a wider range of scenarios.



(a) Comparison of location prior



(b) Comparison of shape prior



(b) Failure cases

Figure 2: Comparison of visual prior between the real world and one learned by VISORGPT on COCO dataset.

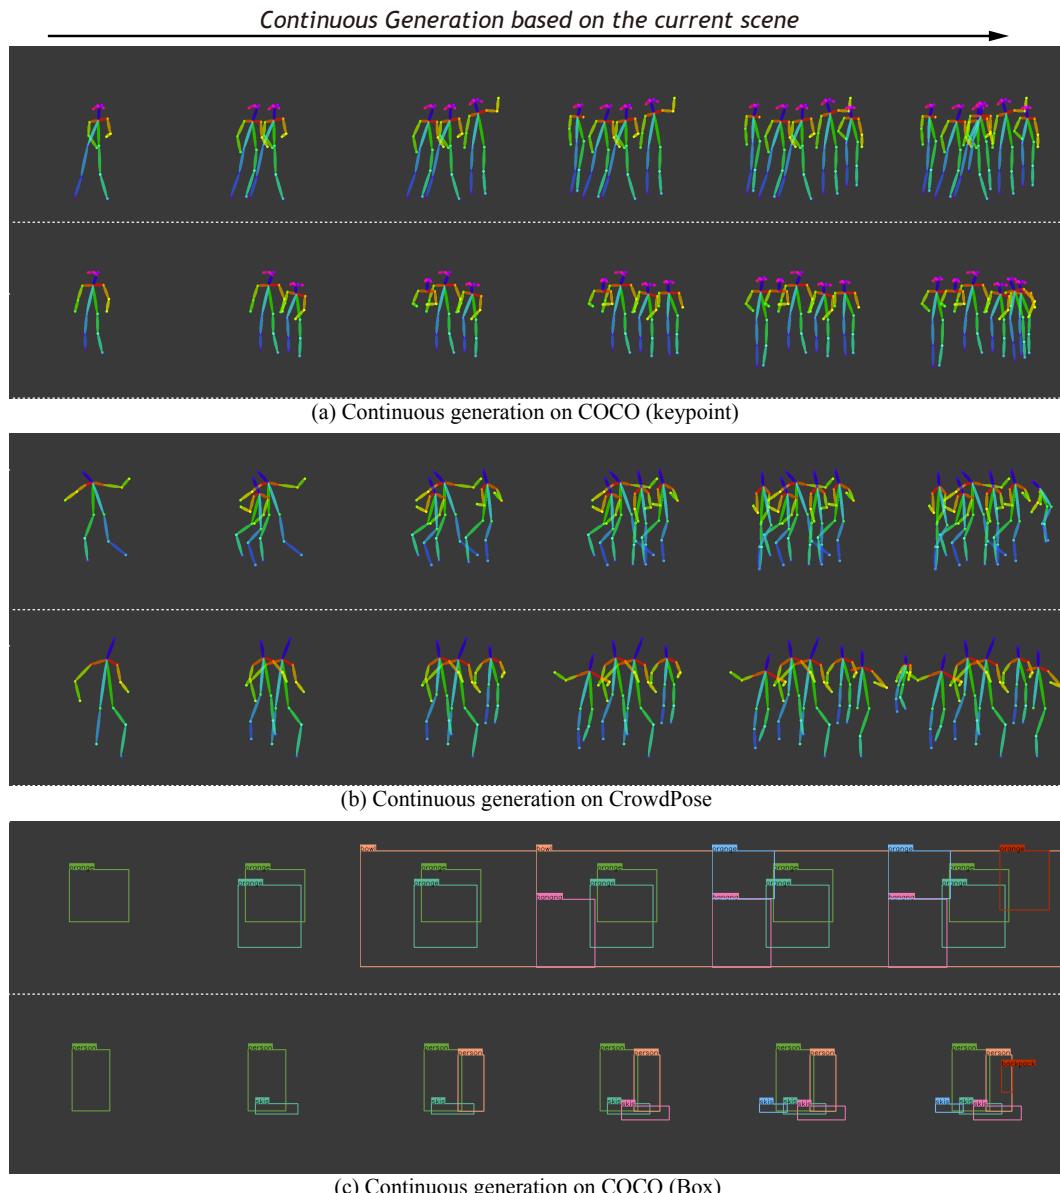


Figure 3: Examples of continual generation.



Figure 4: Examples of input prompts, output sequences, decoded results and synthetic images.



Figure 5: Examples of input prompts, output sequences, decoded results, and synthetic images.