

PoseScript: Linking 3D Human Poses and Natural Language

Ginger Delmas^{1,2}, Philippe Weinzaepfel², Thomas Lucas², Francesc Moreno-Noguer¹, Grégory Rogez²

¹ Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain

² NAVER LABS Europe

¹{gdelmas, fmoreno}@iri.upc.edu, ²{name.surname}@naverlabs.com

Abstract—Natural language plays a critical role in many computer vision applications, such as image captioning, visual question answering, and cross-modal retrieval, to provide fine-grained semantic information. Unfortunately, while human pose is key to human understanding, current 3D human pose datasets lack detailed language descriptions. To address this issue, we have introduced the PoseScript dataset. This dataset pairs more than six thousand 3D human poses from AMASS with rich human-annotated descriptions of the body parts and their spatial relationships. Additionally, to increase the size of the dataset to a scale that is compatible with data-hungry learning algorithms, we have proposed an elaborate captioning process that generates automatic synthetic descriptions in natural language from given 3D keypoints. This process extracts low-level pose information, known as “posecodes”, using a set of simple but generic rules on the 3D keypoints. These posecodes are then combined into higher level textual descriptions using syntactic rules. With automatic annotations, the amount of available data significantly scales up (100k), making it possible to effectively pretrain deep models for finetuning on human captions. To showcase the potential of annotated poses, we present three multi-modal learning tasks that utilize the PoseScript dataset. Firstly, we develop a pipeline that maps 3D poses and textual descriptions into a joint embedding space, allowing for cross-modal retrieval of relevant poses from large-scale datasets. Secondly, we establish a baseline for a text-conditioned model generating 3D poses. Thirdly, we present a learned process for generating pose descriptions. These applications demonstrate the versatility and usefulness of annotated poses in various tasks and pave the way for future research in the field. The dataset is available at <https://europe.naverlabs.com/research/computer-vision/posescript/>.

Index Terms—3D human pose, natural language, multi-modal learning, description, generation, retrieval, captioning.

I. INTRODUCTION

THE pose has the head down, ultimately touching the floor, with the weight of the body on the palms and the feet. The arms are stretched straight forward, shoulder width apart; the feet are a foot apart, the legs are straight, and the hips are raised as high as possible¹. The text above describes the downward dog yoga pose¹, and a reader can easily visualize such a pose from this natural language description. Being able to automatically map natural language descriptions and accurate 3D human poses would open the door to a number of applications such

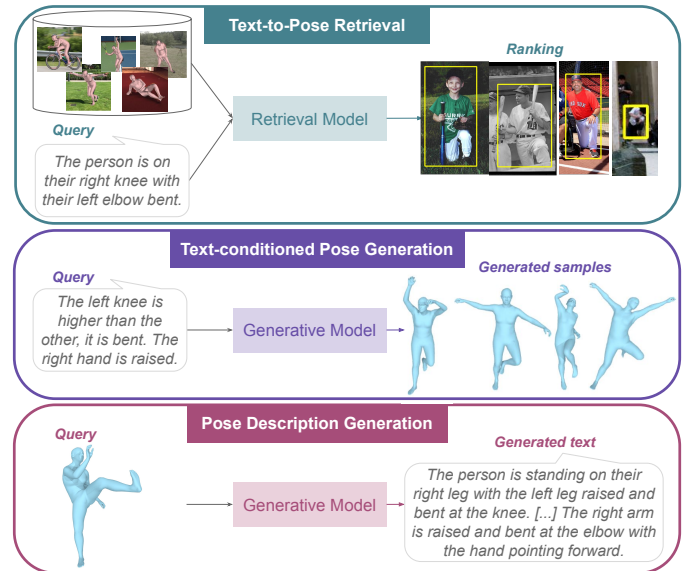


Fig. 1. Illustration of three multi-modal learning applications that can be implemented using PoseScript. The top figure illustrates text-to-pose retrieval where the goal is to retrieve poses in a large-scale database given a text query. This can be applied to databases of images with associated SMPL fits. The middle figure shows an example of text-conditioned pose generation. The bottom figure presents the generation of pose descriptions.

as helping image annotation when the deployment of Motion Capture (MoCap) systems is not practical; performing *pose-based* semantic searches in large-scale datasets (see Figure 1 top), which are currently only based on high-level metadata such as the action being performed [1]–[3]; complex pose or motion data generation in digital animation (see Figure 1 middle); or teaching posture skills to visually impaired [4] (see Figure 1 bottom).

While the problem of combining language and images or videos has attracted significant attention [5]–[8], in particular with the impressive results obtained by the recent multimodal neural networks CLIP [9] and DALL-E [10], the problem of linking text and 3D geometry is only starting to grow. There have been a few recent attempts at mapping text to rigid 3D shapes [11], [12], and at using natural language for 3D object localization [13] or 3D object differentiation [14]. More recently, Fieraru *et al.* [15] introduce AIFit, an approach

¹https://en.wikipedia.org/wiki/Downward_Dog_Pose




Fig. 2. Examples of pose descriptions from PoseScript, produced by human annotators (left) and by our automatic captioning pipeline (right).

to automatically generate human-interpretable feedback on the difference between a reference and a target motion. There have also been a number of attempts to model humans using various forms of text. Attributes have been used for instance to model body shape [16] and face images [17]. Other approaches [18]–[21] leverage textual descriptions to generate motion, but without fine-grained control of the body limbs. More related to our work, Pavlakos *et al.* [22] exploit the relation between two joints along the depth dimension, and Pons-Moll *et al.* [23] describe 3D human poses through a series of *posebits*, which are binary indicators for different types of questions such as ‘Is the right hand above the hips?’. However, these types of Boolean assertions have limited expressivity and remain far from the natural language descriptions a human would use.

In this paper, we propose to map 3D human poses with both free-form and arbitrarily complex structural descriptions, in natural language, of the body parts and their spatial relationships. To that end, we first introduce the multi-modal *PoseScript* dataset, which consists of captions written by human annotators for about 6,000 poses from the AMASS dataset [2]. To scale up this dataset, we additionally propose an automatic captioning pipeline for human-centric poses that makes it possible to annotate thousands of human poses in a few minutes. Our pipeline is built on (a) low-level information obtained via an extension of posebits [23] to finer-grained categorical relations of the different body parts (e.g. ‘the knees are slightly/relatively/completely bent’), units that we refer to as *posecodes*, and on (b) higher-level concepts that come either from the action labels annotated by the BABEL dataset [3], or combinations of posecodes. We define rules to select and aggregate posecodes using linguistic aggregation principles, and convert them into sentences to produce textual descriptions. As a result, we are able to automatically extract human-like captions for a normalized input 3D pose. Importantly, since the process is randomized, we can generate several descriptions per pose, as different human annotators would do. We used this procedure to describe 100,000 poses extracted from the AMASS dataset. Figure 2 shows examples of human-written and automatic captions.

Using the PoseScript dataset, we propose to tackle three multi-modal learning tasks, see Figure 1. The first is a cross-modal retrieval task where the goal is to retrieve from a database the 3D poses that are most similar to a given text

query; this can also be applied to RGB images by associating them with 3D human fits. The second task consists in generating diverse human poses conditioned on a textual description. The third aims to generate a textual description from a provided 3D pose. In all cases, our experiments demonstrate that it is beneficial to pretrain models using the automatic captions before finetuning them on real captions. In summary, our contributions are threefold:

- We introduce the PoseScript dataset (Section III). It associates human poses and structural descriptions in natural language, either obtained through human-written annotations or using our automatic captioning pipeline.
- We then study the task of text-to-pose retrieval (Section IV).
- We present the task of text-conditioned pose generation (Section V).
- We finally focus on the task of pose description generation (Section VI).

A preliminary version of this work was presented in [24], in which we already introduced the PoseScript dataset and baselines for the two first tasks. In this work, we extend our contribution by increasing the size of the PoseScript dataset, studying the pose description generation problem and providing additional analysis. In particular, we study the impact of different aspects of our captioning pipeline and of the number of automatically generated captions used for pretraining. We further improved our models using the additional data, and transformers [25] for our text encoder.

II. RELATED WORK

Text for humans in images. Some previous works have used attributes as semantic-level representation to edit body shapes [16], human clothing [26] and images of faces [17]. In contrast, we leverage natural language, which has the advantage of being unconstrained and more flexible. While others also do so for videos of faces [27] or images of clothed humans in modeling poses [28], our approach focuses on diverse 3D body poses and pose semantics. Closer to our work, [29], [30] focus on generating human 2D poses, SMPL parameters or even images from captions. However, they use MS Coco [31] captions, which are generally simple image-level statements on the activity performed by the human, and which sometimes

relate to the interaction with other elements from the scene, *e.g.* ‘A soccer player is running while the ball is in the air’. By contrast, we focus on fine-grained detailed captions about the pose only. FixMyPose [32] provides manually annotated captions about the difference between human poses in two synthetic images. These captions also include information about the objects in the environment, *e.g.* ‘carpet’ or ‘door’. Similarly, AIFit [15] proposes to automatically generate text about the discrepancies between a reference motion and a performed one, based on differences of angles and positions. Our approach, instead, is focused on describing one single pose without relying on any other visual element.

Text for human motion. We deal with static poses, whereas several existing methods have mainly studied 3D action (sequence) recognition or text-based 2D [19] or 3D motion synthesis. These approaches either condition their model on action labels [20], [33], [34], or descriptions in natural language [18], [21], [35]–[44]. Alternatively, other works leverage appearance information [45], [46]. Yet, even if motion descriptions effectively constrain *sequences* of poses, they do not specifically inform about individual poses. What if an animation studio looks for a sequence of 3D body poses where ‘the man is running with his hands on his hips’? The model used by the artists to initialize the animation should have a deep understanding of the relations between the body parts. To this end, it is important to learn about specific pose semantics, beyond global pose sequence semantics.

Pose semantic representations. Our captioning generation process relies on posecodes that capture relevant information about the pose semantics. Posecodes are drawn inspiration on posebits [23], where images showing a human are annotated with various binary indicators. This data is then used to reduce ambiguities in 3D pose estimation. Conversely, we automatically extract posecodes from normalized 3D poses in order to generate descriptions in natural language. Ordinal depth [22] can be seen as a special case of posebits, focusing on the depth relationship between two joints. The authors exploit such annotations on training images to improve a human mesh recovery model by adding extra constraints. Poselets [47] can also be seen as another way to extract discriminative pose information, but they are not easily interpreted. In contrast to these representations, we propose to generate pose descriptions in natural language, which have the advantage (a) of being a very intuitive way to communicate ideas, and (b) of providing greater flexibility. More recently, TIPS [48] introduce structural descriptions for poses from DeepFashion [49] images. These are collected by selecting states for some body parts among predefined lists. While the underlying process of our proposed captioning pipeline is similar, we deal with 3D poses instead of images, and devise an automatic approach. Besides, we do not limit to fashion and focus on a large variety of poses. Eventually, we collect actual free-form text descriptions from human annotators.

In summary, our proposed PoseScript dataset differs from existing datasets in that it focuses on single 3D poses instead of motion [50], which are diverse and not restricted to modeling poses [48]. Moreover, it provides direct descriptions in natural




Fig. 3. **Origin of the selected poses.** The top bar plot shows the proportion of sequences that are eventually used in PoseScript with respect to available sequences in AMASS. A sequence is ‘used’ if it provided at least one pose to PoseScript. The bottom bar plot shows the distribution of the PoseScript poses over the AMASS sub-datasets.

language instead of simple action labels [1], [3], [33], [51], [52], binary relations [22], [23] or modifying texts [15], [32].

III. THE POSESCRIPT DATASET

The PoseScript dataset is composed of static 3D human poses, together with fine-grained semantic annotations in natural language. We provide **Human-written** annotated descriptions (PoseScript-H), and further increase the amount of data with **Automatically generated captions** (PoseScript-A). The pose selection strategy is presented in Section III-A, the crowd-sourced data collection process in Section III-B, and the automatic captioning pipeline in Section III-C. Finally, aggregated statistics over the PoseScript dataset are reported in Section III-D.

A. Pose selection

The PoseScript dataset contains a total of 100,000 human poses sampled from 14,096 AMASS [2] sequences. To ensure diversity, we excluded the first and last 25 frames of each sequence, which contained initialization poses such as T-poses. Additionally, we sampled only one pose from every 25 frames to avoid redundant poses that were too similar to each other. To further maximize variability, we utilized a farthest-point sampling algorithm. First, we normalized the joint positions of each pose using the neutral body model with default shape

coefficients and global orientation set to 0. Then, we randomly selected one pose from the dataset and iteratively added the pose with the highest MPJE (mean per-joint error) to the set of already-selected poses. This process continued until we obtained the desired number of poses with maximum variability.

Figure 3 presents the AMASS sub-datasets from which the 100,000 selected poses were chosen. Notably, it appears that all sequences of DanceDB and MPI_Limits available in AMASS are used in PoseScript; and that most of the poses in PoseScript actually come from DanceDB (28%), CMU (20%) and BioMotionLab (13%). This is because these sub-datasets exhibit a greater variability of poses compared to others in AMASS.

B. Dataset collection

We collect human-written captions for 3D human poses extracted from the AMASS dataset [2], using Amazon Mechanical Turk² (AMT), a crowd-sourced annotation platform.

Annotation interface. The interface, depicted in Figure 4, first presents the annotators with the mesh of a human pose (in blue), and a slider for controlling the viewpoint. The task is to write the description of the blue pose, using directions relative to the subject (the ‘left’ is the body’s ‘left’), inter-body parts indications (e.g. ‘the right hand is on the hip’), common pose references (e.g. ‘in a headstand’) and analogies. In a second step, to encourage discriminative captions, we additionally display 3 discriminator poses (in gray), which are semantically close to the pose to annotate. The workers are then instructed to refine their description such that it fits only the blue pose. While this interface was first designed to be one-step (with all poses shown at once), we found that annotators would sometimes just provide enough information about the blue pose to distinguish it from the displayed gray poses, but not enough information to fully describe it in detail. This two-step design is an attempt at limiting this phenomenon, in order to obtain both complete and precise descriptions. Some PoseScript-H examples are shown in Figure 2 (left).

Pose discriminators are selected in PoseScript. They should be similar to the target pose, as measured by the distance between pose embeddings obtained with an early version of our retrieval model. They are also required to have at least 15 different posecode categorizations. This ensures that the selected pose discriminators share some semantic similarities with the pose to be annotated while having sufficient differences to be easily distinguished by the annotators.

Annotators qualifications. The annotation task was initially made available to workers living in English-speaking countries, who got at least 5000 of their previous assignments approved, and have an approval rate over 95%. We manually reviewed close to 1000 annotations, based on the following criteria: the description is ‘complete’ (i.e., nearly all the body parts are described), there is no left/right confusion, no distance metric (as these are not scalable with the body size), no subjective comments, few spelling errors and good

Fig. 4. Interface presented to the AMT annotators in order to collect discriminative descriptions of the blue pose following a two-step process.

grammar. Based on these evaluations, we further qualified 41 workers and made the assignments available to them only; only some of their annotations, selected at random, would then be manually reviewed. 61 other workers submitted great annotations, but did not complete enough to get qualified.

Pricing. The time to complete a HIT was estimated to be 2-3 minutes. Each HIT was rewarded \$0.50, based on the minimum wage in California for 2022. We additionally paid \$2 bonus to qualified annotators for every 50 annotations.

Semantic analysis. We report in Table I the results of a semantic analysis carried out on 115 annotations. It emerges that one challenging aspect of the dataset is the implicit side description of some body parts: the deduction of the corresponding side involves reasoning about previously described body parts and hierarchical relations between them. In the course of this study, we moreover measured that the annotations describe in average 6.2 different body parts: those vary across annotations, and multiple indications can be given to detail the position of one body part. This hints at the level of details in the annotations.

C. Automatic captioning pipeline

We now describe the process used to generate synthetic textual descriptions for 3D human poses. As depicted in Figure 5, it relies on the extraction, selection and aggregation of elementary pieces of pose information, called *posecodes*, that are eventually converted into sentences to produce a description.

The process takes 3D keypoint coordinates of human-centric poses as input. These are inferred with the SMPL-H body model [53] using the default shape coefficients and a normalized global orientation along the y-axis.

1. Posecode extraction. A posecode describes a relation between a specific set of joints. We capture five kinds of

²<https://www.mturk.com>

TABLE I
SEMANTIC ANALYSIS ON 115 POSESCRIPT ANNOTATIONS.

Property	Proportion	Examples
Egocentric relations	86%	"above their head" / "looking at their right arm" / "at about shin level"
Analogies	24%	"in a circular position" / "almost in a seated position" / "like a hug"
Implicit side description	41%	"right leg stretched out forward, Ø heel on the ground"

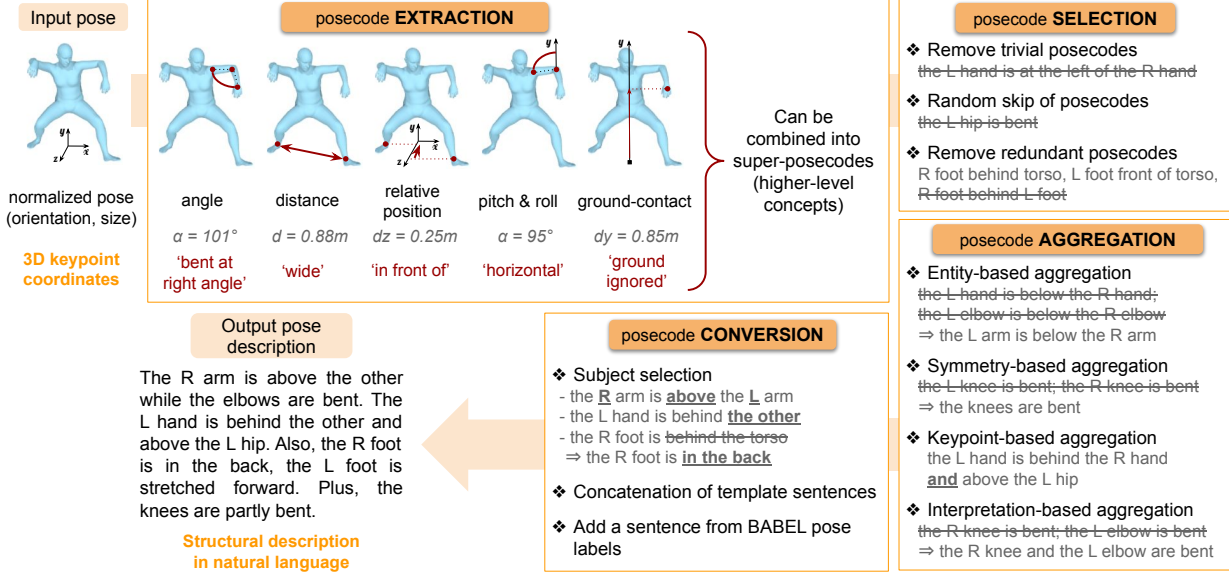


Fig. 5. **Overview of our captioning pipeline.** Given a normalized 3D pose, we use posecodes to extract semantic pose information. These posecodes are then selected, merged or combined (when relevant) before being converted into a structural pose description in natural language. Letters 'L' and 'R' stand for 'left' and 'right' respectively.

elementary relations: angles, distances and relative positions (as in [23]), but also pitch, roll and ground-contacts.

- *Angle posecodes* describe how a body part 'bends' at a given joint, e.g. the left elbow. Depending on the angle, the posecode is assigned one of the following attributes: 'straight', 'slightly bent', 'partially bent', 'bent at right angle', 'almost completely bent' and 'completely bent'.

- *Distance posecodes* categorize the L2-distance between two keypoints (e.g. the two hands) into 'close', 'shoulder width apart', 'spread' or 'wide' apart.

- *Posecodes on relative position* compute the difference between two keypoints along a given axis. The possible categories are, for the x -axis: 'at the right of', ' x -ignored', 'at the left of'; for the y -axis: 'below', ' y -ignored', 'above'; and for the z -axis: 'behind', ' z -ignored' and 'in front of'. In particular, comparing the x -coordinate of the left and right hands allows to infer if they are crossed (i.e., the left hand is 'at the right' of the right hand). The 'ignored' interpretations are ambiguous configurations which will not be described.

- *Pitch & roll posecodes* assess the verticality or horizontality of a body part defined by two keypoints (e.g. the left knee and hip together define the left thigh). A body part is 'vertical' if it is approximately orthogonal to the y -hyperplane, and 'horizontal' if it is in it. Other configurations are 'pitch-roll-ignored'.

- *Ground-contact posecodes*, used for intermediate computation only, denote whether a keypoint is 'on the ground' (i.e., vertically close to the keypoint of minimal height in the body, considered as the ground) or 'ground-ignored'.

Handling ambiguity in posecode categorization. Posecode categorizations are obtained using predefined thresholds. As these values are inherently subjective, we randomize the binning step by also defining a noise level applied to the measured angles and distances values before thresholding.

Higher-level concepts. We also define a few *super-posecodes* to extract higher-level pose concepts. These posecodes are binary (they either apply or not to a given pose configuration), and are expressed from elementary posecodes. For instance, the super-posecode 'kneeling' can be defined as having both knees 'on the ground' and 'completely bent'.

2. Posecode selection aims at selecting an interesting subset of posecodes among those extracted, to obtain a concise yet discriminative description. First, we remove trivial settings (e.g. 'the left hand is at the left of the right hand'). Next, based on a statistical study over the whole set of poses, we randomly skip a few non-essential – i.e., non-trivial but non highly discriminative – posecodes, to account for natural human oversights. We also set highly-discriminative posecodes as unskippable. Finally, we remove redundant posecodes based on statistically frequent pairs and triplets of posecodes, and transitive relations between body parts. Details are provided in the supplementary material.




Fig. 7. **Overview of the training scheme of the retrieval model.** The input pose and caption are first separately fed into a pose encoder and a text encoder, respectively, to map them into a joint embedding space. The loss function is designed to encourage the pose embedding y_i and its corresponding caption embedding x_i to be close to each other in this shared latent space, while also pulling them apart from the features of other poses in the same training batch (e.g., y_k and y_l).

pose encoder $\phi(\cdot)$ to encode the captions and poses, respectively. Let $x = \theta(c) \in \mathbb{R}^d$ and $y = \phi(p) \in \mathbb{R}^d$ be the L_2 -normalized representations of a caption c and a pose p in the joint embedding space, as shown in Figure 7. The similarity score between the two modalities is computed based on the distance between their respective embeddings.

Encoders. The caption is tokenized and then embedded using either a bi-GRU [55] on top of GloVe word embeddings [56] or a transformer [25] on top of the frozen pretrained DistilBERT [57] word embeddings. The pose is encoded as a matrix of size (52, 3), consisting in the rotation of the SMPL-H body joints in axis-angle representation. It is flattened and fed to the pose encoder, chosen as the VPoser encoder [58]. An added ReLU and final projection layer produce an embedding of the same size d as the text encoding.

Training. Given a batch of B training pairs (x_i, y_i) , we use the Batch-Based Classification (BBC) loss which is common in cross-modal retrieval [59]:

$$\mathcal{L}_{\text{BBC}} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(\gamma \sigma(x_i, y_i))}{\sum_j \exp(\gamma \sigma(x_i, y_j))}, \quad (1)$$

where γ is a learnable temperature parameter and σ is the cosine similarity function $\sigma(x, y) = x^\top y / (\|x\|_2 \times \|y\|_2)$.

Evaluation protocol. Text-to-pose retrieval is evaluated by ranking the whole set of poses for each of the query texts. We then compute the recall@K ($R@K$), which is the proportion of query texts for which the corresponding pose is ranked in the top- K retrieved poses. We proceed similarly to evaluate pose-to-text retrieval. We use $K = 1, 5, 10$ and additionally report the mean recall (mRecall) as the average over all recall@K values from both retrieval directions.

Quantitative results. We report results on the test set of PoseScript in Table II, both on automatic and human-written captions. Our model trained on automatic captions obtains a mean recall of 72.8%, with a $R@1$ close to 50% and a $R@10$ above 80% on automatic captions. However, the performance degrades on human captions, as many words from the richer human vocabulary are unseen during training on automatic captions. When trained on human captions, the model obtains a higher – but still rather low – performance. Using human

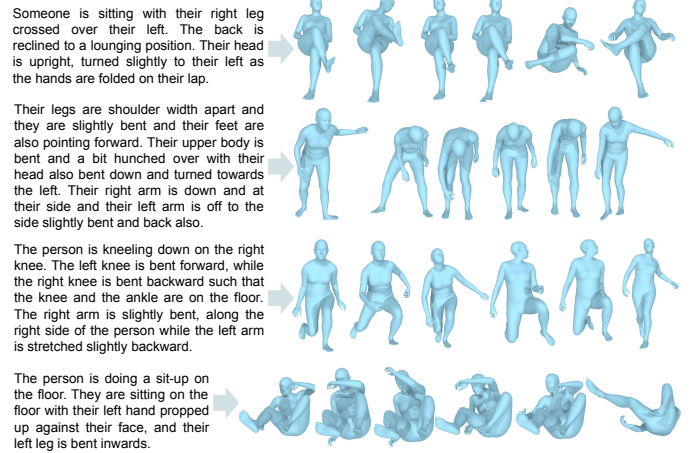


Fig. 8. **Text-to-pose retrieval results** for human-written captions from PoseScript. Directions such as 'left' and 'right' are relative to the body.

captions to finetune the initial model trained on automatic ones brings an improvement of almost a factor 2, with a mean recall (resp. $R@10$ for text-to-pose) of 40.9% (resp. 57.9%) compared to 23.0% (resp. 35.7%) when training from scratch. This experiment clearly shows the benefit of using the automatic captioning pipeline to scale-up the PoseScript dataset. In particular, this suggests that the model is able to derive new concepts in human-written captions from non-trivial combination of existing posecodes in automatic captions. The last two rows show further improvement when using a transformer-based text encoder and applying data augmentation by mirroring the poses (i.e., switching *left* and *right* side words in the texts).

Qualitative retrieval results. Examples of text-to-pose retrieval results are presented in Figure 8. It appears that the model is able to encode several pose concepts concurrently and to distinguish between the left and right body parts.

Retrieval in image databases. MS Coco [31] is one of several real-world datasets that have been used for human mesh recovery. We resort to the 74,834 pseudo-ground-truth SMPL fits provided by EFT [60], on which we apply our text-to-pose retrieval model trained with PoseScript. We then retrieve 3D poses among this MS Coco-EFT set, and display the corresponding images with the associated bounding box around the human body. Results are shown in Figure 9. We observe that overall, the constraints specified in the query text are satisfied in the images. Retrieval is based on the poses and not on the context, hence the third image of the first row where the pose is close to an actual kneeling one. This shows one application of a retrieval model trained on the PoseScript dataset and applied to a third modality: specific pose retrieval in images. Our model can be applied to any dataset of images containing humans, as long as SMPL fits are also available.

V. APPLICATION TO TEXT-CONDITIONED POSE GENERATION

We next study the problem of *text-conditioned human pose generation*, i.e., generating possible matching poses for a given text query. Our proposed model is based on Variational Auto-Encoders (VAEs) [61].

TABLE II

TEXT-TO-POSE AND POSE-TO-TEXT RETRIEVAL RESULTS ON THE TEST SPLIT OF THE POSESCRIPT DATASET. FOR HUMAN-WRITTEN CAPTIONS (POSESCRIPT-H), WE EVALUATE MODELS TRAINED ON EACH SPECIFIC CAPTION SET ALONE, AND ONE PRETRAINED ON AUTOMATIC CAPTIONS (POSESCRIPT-A) THEN FINETUNED (FT) ON HUMAN CAPTIONS. UNLESS SPECIFIED OTHERWISE, MODELS ALL HAVE THE GLoVe-BIGRU CONFIGURATION. RESULTS ARE AVERAGED OVER 3 RUNS.

	mRecall \uparrow	pose-to-text			text-to-pose		
		$R@1\uparrow$	$R@5\uparrow$	$R@10\uparrow$	$R@1\uparrow$	$R@5\uparrow$	$R@10\uparrow$
<i>test on PoseScript-A (19,990 samples)</i> trained on PoseScript-A	72.8 \pm 0.4	47.2 \pm 0.5	78.5 \pm 0.3	87.1 \pm 0.2	52.4 \pm 0.8	82.0 \pm 0.5	89.3 \pm 0.4
<i>test on PoseScript-H (1234 samples)</i> trained on PoseScript-A	5.9 \pm 0.4	2.3 \pm 0.4	6.9 \pm 0.7	11.3 \pm 0.6	1.4 \pm 0.1	5.1 \pm 0.3	8.6 \pm 0.4
trained on PoseScript-H	23.0 \pm 0.6	8.9 \pm 0.3	24.4 \pm 1.5	34.8 \pm 0.5	9.3 \pm 1.0	24.6 \pm 0.1	35.7 \pm 0.9
trained on PoseScript-A, FT on PoseScript-H	40.9 \pm 0.1	19.8 \pm 0.4	44.9 \pm 0.7	56.2 \pm 0.7	19.9 \pm 0.6	46.5 \pm 0.1	57.9 \pm 0.3
using the transformer text encoder	43.3 \pm 0.6	21.0 \pm 0.5	48.2 \pm 1.2	60.4 \pm 1.0	21.7 \pm 1.3	48.1 \pm 0.5	60.6 \pm 0.5
with mirroring augmentation	45.3 \pm 0.4	22.3 \pm 1.6	50.1 \pm 1.0	62.9 \pm 0.7	22.1 \pm 0.4	51.4 \pm 0.3	63.1 \pm 0.9




Fig. 9. **Retrieval results in image databases.** We use our text-to-pose retrieval model trained on human captions from PoseScript to retrieve 3D poses from SMPL fits on MS Coco, for some given text queries. We display the corresponding pictures and bounding boxes for the top retrieved poses.




Fig. 10. **Overview of the text-conditioned generative model.** During training, it follows a VAE but where the latent distribution \mathcal{N}_p from the pose encoder has a KL divergence term with the prior distribution \mathcal{N}_c given by the text encoder. At test time, the sample z is drawn from \mathcal{N}_c .

Training. Our goal is to generate a pose \hat{p} given its caption c . To this end, we train a conditional VAE model that takes a tuple (p, c) composed of a pose p and its caption c at training time. Figure 10 gives an overview of our model. A pose encoder maps the pose p to a posterior over latent variables by producing the mean $\mu(p)$ and variance $\Sigma(p)$ of a normal distribution $\mathcal{N}_p = \mathcal{N}(\cdot | \mu(p), \Sigma(p))$. Another encoder

is used to obtain a prior distribution \mathcal{N}_c , independent of p but conditioned on c . A latent variable $z \sim \mathcal{N}_p$ is sampled from \mathcal{N}_p and decoded into a reconstructed pose \hat{p} . The training loss combines a reconstruction term $\mathcal{L}_R(p, \hat{p})$ between the original and reconstructed poses, p and \hat{p} , and a regularization term, the Kullback-Leibler (KL) divergence between \mathcal{N}_p and \mathcal{N}_c :

$$\mathcal{L} = \mathcal{L}_R(p, \hat{p}) + \mathcal{L}_{KL}(\mathcal{N}_p, \mathcal{N}_c). \quad (2)$$

We also experiment with an additional loss term \mathcal{L}_{reg} denoting $\mathcal{L}_{KL}(\mathcal{N}_p, \mathcal{N}(\cdot | 0, I)) + \mathcal{L}_{KL}(\mathcal{N}_c, \mathcal{N}(\cdot | 0, I))$: KL divergences between the posterior (resp. the prior) and the standard Gaussian $\mathcal{N}_0 = \mathcal{N}(\cdot | 0, I)$. These can be seen as other regularizers and they also allow to sample poses from the model without conditioning on captions. We treat the variance of the decoder as a learned constant [62] and use a negative log likelihood (nll) as reconstruction loss, either from a Gaussian – which corresponds to an L2 loss and a learned variance term – or a Laplacian density, which corresponds to an L1

TABLE III
EVALUATION OF THE TEXT-CONDITIONED GENERATIVE MODEL ON POSESCRIPT-A FOR A MODEL WITHOUT OR WITH \mathcal{L}_{reg} (TOP) AND ON POSESCRIPT-H WITHOUT OR WITH PRETRAINING ON POSESCRIPT-A (BOTTOM). UNLESS SPECIFIED OTHERWISE, MODELS ALL HAVE THE GLOVE-BIGRU CONFIGURATION. RESULTS ARE AVERAGED OVER 3 RUNS. THE VARIABILITY OF R/G (RESP. G/R) mRECALL IS DUE TO THE RANDOM SELECTION OF A GENERATED POSE SAMPLE AT TEST (RESP. TRAINING) TIME. FOR COMPARISON, THE mRECALL WHEN TRAINING AND TESTING ON REAL POSES IS 72.8 WITH POSESCRIPT-A AND 45.3 ON POSESCRIPT-H.

	FID↓	ELBO jts↑	ELBO vert.↑	ELBO rot.↑	mRecall R/G↑	mRecall G/R↑
<i>test on PoseScript-A</i>						
without \mathcal{L}_{reg}	0.12 ± 0.02	1.76 ± 0.00	2.05 ± 0.01	1.06 ± 0.02	22.7 ± 1.4	41.5 ± 4.5
with \mathcal{L}_{reg}	0.07 ± 0.00	1.78 ± 0.00	2.09 ± 0.01	1.10 ± 0.01	25.1 ± 0.2	40.0 ± 3.0
<i>test on PoseScript-H, for the model with \mathcal{L}_{reg}</i>						
without pretraining	0.29 ± 0.04	0.98 ± 0.02	1.41 ± 0.02	0.53 ± 0.02	5.2 ± 0.5	20.5 ± 3.4
with pretraining	0.04 ± 0.01	1.39 ± 0.01	1.75 ± 0.01	0.87 ± 0.01	19.5 ± 2.6	35.1 ± 2.5
using the transformer text encoder	0.03 ± 0.00	1.40 ± 0.00	1.74 ± 0.01	0.85 ± 0.02	35.6 ± 1.2	44.5 ± 0.1
with mirroring augmentation	0.03 ± 0.00	1.41 ± 0.00	1.75 ± 0.01	0.86 ± 0.01	37.5 ± 0.5	45.2 ± 1.1

loss. Following VPoser, we use SMPL(-H) inputs with the axis-angle representation, and output joint rotations with the continuous 6D representation of [63]. Our reconstruction loss $\mathcal{L}_R(p, \hat{p})$ is a sum of the reconstruction losses between the rotation matrices – evaluated with a Gaussian log-likelihood – the position of the joints and the position of the vertices, both evaluated with a Laplacian log-likelihood.

Text-conditioned generation. At test time, a caption c is encoded into \mathcal{N}_c , from which z is sampled and decoded into a generated pose \hat{p} .

Evaluation metrics. We evaluate sample quality following the principle of the Fréchet inception distance: we compare the distributions of features extracted using our retrieval model (see Section IV), using real test poses and poses generated from test captions. This is denoted FID with an abuse of notation. We also report the mean-recall of retrieval models trained on real poses and evaluated on generated poses (mR R/G), and vice-versa (mR G/R). Both metrics are sensitive to sample quality: the retrieval model will fail if the data is unrealistic. The second metric is also sensitive to diversity: missing parts of the data distribution hinder the retrieval model trained on samples. Finally, we report the Evidence Lower Bound (ELBO) computed on joints, vertices or rotation matrices, normalized by the target dimension.

Results. We present quantitative results in Table III. We first find that adding the extra-regularization loss \mathcal{L}_{reg} to the model trained and evaluated on automatic captions is slightly helpful. Also, it is convenient to sample poses without any conditioning. We keep this configuration and evaluate it when (a) training on human captions and (b) pretraining on automatic captions then finetuning on human captions. Pretraining improves greatly all metrics, showing that it helps to yield realistic and diverse samples. The transformer-based text encoder improves over the one based on GloVe and biGRU, achieving substantially better performance in terms of mRecall. This suggests that the transformer-based text encoder has a finer understanding of the pose semantics. Slight improvement is yielded by the mirroring augmentation. We display generated samples in Figure 11; the poses are realistic and generally correspond to the query. There are some

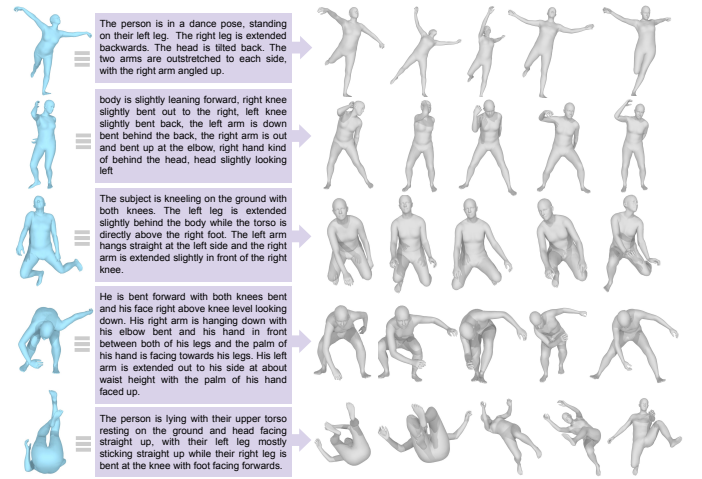


Fig. 11. **Examples of generated samples.** We show several generated samples (in grey) obtained for the human-written captions presented in the middle. For reference, we also show in blue the pose for which this annotation was originally collected.

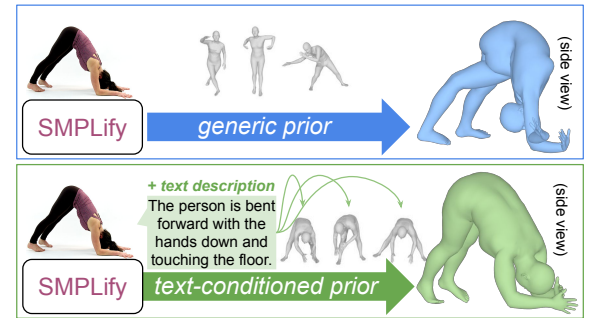


Fig. 12. **Example of potential application to SMPL fitting in images.** Using the text-conditioned pose prior (bottom) yields a more accurate 3D pose than a generic pose prior (top) when running the optimization-based SMPL fitting method SMPLify.

variations, especially when the text allows it, for instance with the height of the right leg in the top example or the distance between the legs in the fourth row. Failure cases can happen; in particular rare words like ‘lying’ in the bottom row lead to higher variance in the generated samples; some of them are nevertheless close to the reference.




Fig. 13. **Overview of the pose description generative model.** Pose information is injected in the text transformer via cross-attention. The cross-entropy loss between the output probability distribution over the vocabulary and the shifted target trains the model to recursively predict the next word starting from the *BOS* token.

Application to SMPL fitting in image. We showcase the potential of leveraging text data for 3D tasks on a challenging example from SMPLify [64], in Figure 12. We use our text-conditional prior instead of the generic VPoser prior [58] to initialize to a pose closer to the ground truth and to better guide the in-the-loop optimization. This helps to avoid bad local minima traps.

VI. APPLICATION TO POSE DESCRIPTION GENERATION

We now introduce our *learned* approach to generate pose descriptions in natural language. Conversely to the process presented in Section III-C, this one does not rely on heuristics nor template structures. It is trained on the human-written captions, resulting in generated texts that are more concise, with improved formulation and more high-level concepts. Note that this model does not debate the pipeline from Section III-C: just like the others, it benefits tremendously from the pretraining on the automatic captions the pipeline produces.

We use an auto-regressive model, that starts from the *BOS* (beginning-of-sequence) token and produces iteratively each new token given all those previously generated, and the pose conditioning; see Figure 13.

Training. The model is provided with the tokenized caption $c_{1:L}$ of L tokens. It embeds each token, adds positional encoding then feeds the result to a transformer, which accounts for the pose conditioning thanks to the cross-attention mechanisms. The use of a causal attention mask prevents the model from attending tokens from $l+1$ to L when dealing with token l . Finally, the model outputs a probability distribution $w_{1:L}$ over the vocabulary, where w_l corresponds to the probability $p(\cdot|c_{1:l})$. The cross-entropy loss between $w_{1:L-1}$ and $c_{2:L}$ maximizes $p(c_{l+1}|c_{1:l})$, thus training the model to predict the next token $l+1$ from previous ones $c_{1:l}$.

Inference. Given the *BOS* token and the input pose, the description is decoded iteratively in a greedy fashion by maximizing the likelihood at each step l , until the special token *EOS* is decoded. At step l , token l is decoded as the one that maximizes the output w_l ; it is then appended to the tokens 1 to $l-1$ decoded earlier, so as to predict token $l+1$ in the next pass.

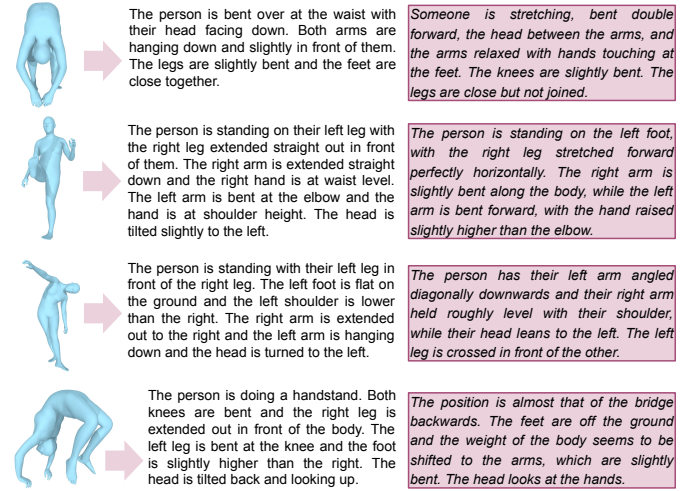


Fig. 14. **Examples of model-generated descriptions** for the poses on the left. For reference, the annotated texts are presented in the colored boxes.

Evaluation. We report the standard natural language processing (NLP) metrics BLEU-4 [65], Rouge-L [66] and METEOR [67]. Following TM2T [41], we use our retrieval model from Section IV to measure the recall at $K \in \{1, 2, 3\}$ (top-k R-precision) when ranking, for a query pose, the corresponding text generated by our model among 31 randomly sampled generated texts yielded for other poses. Additionally, we report reconstruction metrics (MPJE, MPVE and geodesic distance on the joint rotations) obtained by comparing the input pose with the one generated by our model from Section V, when given the generated text. While the NLP metrics measure the common n-grams between the reference text and the generated one, the other metrics evaluate the semantic content of the generated description. Indeed, an insufficiently detailed text could not help generate nor retrieve back the input pose.

Results. Again, we notice from Table IV that pretraining on the automatic captions leads to substantially better descriptions, and that the mirroring augmentation helps. We note that R-Precision and reconstruction metrics rely on trained models, and their biased understanding of the data (*i.e.*, some concepts may be ill-encoded). This could explain why the produced texts appear to yield better results than the original ones.

Examples of descriptions generated by our model are presented in Figure 14. It appears that the model is able to produce meaningful descriptions, with egocentric relations and high-level concepts (*e.g.* handstand). However, it sometimes hallucinate (leg position in the last example), or it struggles to understand the pose as a whole, especially in rare cases like upside-down poses (the head would indeed be looking up if the body was not bent backwards that much).

VII. CHARACTERISTICS OF AUTOMATIC CAPTIONS

In this section, we aim to study the impact of different aspects of the automatic captioning pipeline. To this end, we generate 6 different kinds of captions per pose, each with different characteristics: we generate all the captions using the same pipeline, which is presented in Section III, and disable some steps of the process to produce the different versions.

TABLE IV

CAPTION GENERATION RESULTS. THE TOP BLOCK SHOWS SOME REFERENCE MEASURES WHILE THE LOWER BLOCK EVALUATES THE GENERATED TEXTS FOR POSESCRIPT-H. RESULTS ARE AVERAGED OVER 3 RUNS.

	R-Precision \uparrow			NLP \uparrow			Reconstruction \downarrow (<i>best of 30</i>)		
	R@1	R@2	R@3	BLEU-4	ROUGE-L	METEOR	MPJE	MPVE	Geodesic
random text from PoseScript-H	3.00	5.83	8.18	7.07	23.98	25.96	419	328	12.51
matching text from PoseScript-A	44.49	57.62	66.13	24.52	35.15	43.20	203	170	9.13
matching text from PoseScript-H	81.60	87.60	91.25	100.00	100.00	100.00	203	168	8.93
without pretraining	24.39 \pm 2.34	36.01 \pm 2.22	43.87 \pm 1.78	11.35 \pm 0.13	31.32 \pm 0.22	30.88 \pm 0.19	324 \pm 1	254 \pm 2	10.74 \pm 0.04
with pretraining	88.09 \pm 0.64	93.33 \pm 0.33	95.54 \pm 0.07	13.07 \pm 0.11	33.86 \pm 0.16	32.95 \pm 0.20	203 \pm 1	169 \pm 0	8.80 \pm 0.02
with mirroring augmentation	88.98 \pm 1.51	94.46 \pm 0.53	96.14 \pm 0.34	13.22 \pm 0.17	34.07 \pm 0.04	33.09 \pm 0.21	200 \pm 2	166 \pm 2	8.75 \pm 0.07

TABLE V

SUMMARY OF THE AUTOMATIC CAPTION VERSIONS. \checkmark SYMBOLS INDICATE WHEN CHARACTERISTICS APPLY TO EACH CAPTION VERSION. ALL MODELS WERE TRAINED ON A POOL OF 3 CAPTIONS PER POSE (MULTIPLICITY). MEAN RECALL RESULTS ARE AVERAGED OVER 3 RUNS OF MODELS TRAINED WITH THE BI-GRU CONFIGURATION.

Version (pretraining)	Multiplicity	Random skip	Implicitness	Auxiliary labels	Ripple effect	mRecall
N1	$\times 3$	-	-	-	-	37.5 \pm 1.6
N2	$\times 3$	\checkmark	-	-	-	38.8 \pm 1.3
N3	$\times 3$	-	\checkmark	-	-	39.8 \pm 0.5
N4	$\times 3$	-	-	(w/o dancing label)	-	37.6 \pm 0.7
N4d	$\times 3$	-	-	(w/ dancing label)	-	38.7 \pm 1.9
N5	$\times 3$	-	-	-	\checkmark	37.1 \pm 2.0
PoseScript-A	$\times 1$	\checkmark	-	-	-	40.9 \pm 0.1
	$\times 1$	\checkmark	\checkmark	(w/ dancing label)	-	
	$\times 1$	\checkmark	\checkmark	-	\checkmark	

Specifically, steps that were deactivated include: (1) randomly skipping eligible posecodes for description; (2) aggregating posecodes (“implicitness”), omitting support keypoints (e.g. ‘the right foot is behind the torso’ does not turn into ‘the right foot is in the back’ when this step is deactivated) and randomly referring to a body part by a substitute word (e.g. ‘it’/‘they’, ‘the other’); (3) adding a sentence constructed from high-level pose annotations given by BABEL [3]; and (4) removing redundant posecodes based on ripple effect rules.

Among all 100k poses of PoseScript, only 36,317 are annotated in BABEL and may benefit from an additional sentence in their automatic description. As 28% of PoseScript poses come from DanceDB, which was not annotated in BABEL, we additionally assign the ‘dancing’ label to those DanceDB-originated poses, for one variant of the automatic captions that already leverages BABEL auxiliary annotations. This results in 64,758 poses benefiting from an auxiliary label.

Table V summarises the characteristics of the 6 captions versions introduced in this section (N1 to N5), along with those of the captions versions composing PoseScript-A (last row), which is used in the rest of this article.

We pretrain 6 different retrieval models (GloVe+biGRU configuration), one with each caption version N1 to N5. Each model is trained on a pool of 3 generated captions from the same version. Next, we finetune the retrieval models on PoseScript-H. We report the mean recall on the test set of PoseScript-H in the last column of Table V.

Due to relatively high variability, it is hard to tell which aspects of the automatic captioning pipeline are the most important. However, it appears clearly that implicitness (i.e., posecode aggregations) make for automatic captions that are




Fig. 15. Training with different amounts of data. Results are reported on the test split of PoseScript-H₂. The green curve corresponds to the size of automatic data used for pretraining, while the violet curve corresponds to the size of human-written data used for direct training.

closer to human-written ones, as leveraging captions N3 for pretraining achieve the best results. Eventually, the best performance is obtained by pretraining the model on a pool of 3 captions from different versions (last row). This last pool of captions, namely PoseScript-A, is the one used in the experiments for all other sections.

VIII. SIZE OF THE TRAINING DATA

In this section, we study the impact of the number of automatic captions used for pretraining. To this end, we define data subsets of different sizes with automatic (10k, 20k and 100k) or human-written annotations (3.9k and 6.3k). We train one retrieval model (GloVe+biGRU configuration) for each of the PoseScript-A subsets, and finetune it on PoseScript-H. Besides, we train retrieval models directly on each human-written data subsets. We compare results on the test set of PoseScript-H in Figure 15.

We observe better results when leveraging larger amounts of automatic captions at pretraining time, especially when the number of automatic captions exceeds significantly the number of human-written descriptions.

IX. IMPLEMENTATION DETAILS

We follow VPoser [58] for the pose encoder and decoder architectures (except that we use the 52 joints of SMPL-H [53]). GloVe word embeddings are 300-dimensional. We use a one-layer bidirectional GRU with 512-dimensional hidden state features. Our transformer-based [25] text encoder uses frozen DistilBERT [57] word embeddings of dimension 768, which are then passed to a ReLU and projected into a 512-dimensional space. We next apply a cosine positional encoding, and feed the result to a transformer composed of 4

layers, 4 heads, GELU activations, and feed-forward networks of size 1024. The final embedding of the text sequence is obtained by average-pooling of the output. We use the same encoders for retrieval and generative tasks. Our text decoder is a transformer with the same configuration as for the text encoder, except that it has 8 heads; both the word embeddings and the latent space are of size 512. Models are optimized with Adam [68]. We use an initial loss temperature of $\gamma = 10$ for the retrieval model, and a learning rate coefficient of 0.1 for the pose auto-encoder in the pose generative model at finetuning time. See Table VI for details.

TABLE VI

IMPLEMENTATION DETAILS. d IS THE POSE EMBEDDING SIZE, lr STANDS FOR 'LEARNING RATE'. THE SEVERAL VALUES ARE FOR TRAINING ON POSESCRIPT-A AND POSESCRIPT-H RESPECTIVELY.

model	epochs	batch size	d	lr init	lr or weight decay
retrieval	1000 / 300	512 / 32	512	$2 \cdot 10^{-4}$	0.5 every 400 / 40 epochs
pose generative	2000 / 1000	128	32	$10^{-4} / 10^{-5}$	10^{-4}
text generative	3000 / 2000	128	32 \rightarrow 512	$10^{-4} / 10^{-5}$	10^{-4}

X. DISCUSSION AND CONCLUSION

We introduced PoseScript, the first dataset to map 3D human poses and descriptions in natural language. We provided multi-modal applications to text-to-pose retrieval, to text-conditioned human pose generation and pose description generation. Pre-training on the automatic texts improves performance notably (by a factor 2), systematically over all three studied tasks.

Limitations. The accuracy of our models depends in high part on the training data. For instance, our models struggle due to the limited amount of self-contact or upside-down poses. An alternative to collecting more human-written texts would be to design specific posecodes describing such pose configurations. Another general observation is that our models strive to produce results meeting *all* the text requirements, as the PoseScript descriptions are very rich and complex.

Future works. The PoseScript dataset could be extended to account for multi-people interactions. One could also leverage knowledge from large multi-modal models (e.g. text-to-image) to help filling in the gaps of the collected data in some aspects (e.g. activity concepts). One could further explore the use of a text-based pose prior (i.e. with body semantics awareness) for other applications, e.g. action recognition.

ACKNOWLEDGMENTS



This work is supported by the Spanish government with the project MoHuCo PID2020-120049RB-I00, and by NAVER LABS Europe under technology transfer contract 'Text4Pose'.

REFERENCES

- [1] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu, "Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments," *IEEE Trans. PAMI*, 2014.
- [2] N. Mahmood, N. Ghorbani, N. F. Troje, G. Pons-Moll, and M. J. Black, "AMASS: Archive of motion capture as surface shapes," in *ICCV*, 2019.
- [3] A. R. Punnakal, A. Chandrasekaran, N. Athanasiou, A. Quiros-Ramirez, and M. J. Black, "BABEL: Bodies, action and behavior with english labels," in *CVPR*, 2021.

- [4] C. Suveren-Erdogan and S. Suveren, "Teaching of basic posture skills in visually impaired individuals and its implementation under aggravated conditions," *Journal of Education and Learning*, 2018.
- [5] A. Karpathy and L. Fei-Fei, "Deep visual-semantic alignments for generating image descriptions," in *CVPR*, 2015.
- [6] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, "Show and tell: A neural image caption generator," *CVPR*, 2015.
- [7] S. Li, T. Xiao, H. Li, W. Yang, and X. Wang, "Identity-aware textual-visual matching with latent co-attention," in *ICCV*, 2017.
- [8] F. Feng, X. Wang, and R. Li, "Cross-modal retrieval with correspondence autoencoder," in *ACMMM*, 2014.
- [9] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *ICML*, 2021.
- [10] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, "Zero-shot text-to-image generation," in *ICML*, 2021.
- [11] K. Chen, C. B. Choy, M. Savva, A. X. Chang, T. Funkhouser, and S. Savarese, "Text2shape: Generating shapes from natural language by learning joint embeddings," in *ACCV*, 2018.
- [12] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall, "Dreamfusion: Text-to-3d using 2d diffusion," *International Conference on Learning Representations*, 2023.
- [13] D. Z. Chen, A. X. Chang, and M. Nießner, "Scanrefer: 3D object localization in rgb-d scans using natural language," in *ECCV*, 2020.
- [14] P. Achlioptas, J. Fan, R. Hawkins, N. Goodman, and L. J. Guibas, "Shapeglot: Learning language for shape differentiation," in *ICCV*, 2019.
- [15] M. Fieraru, M. Zanfir, S. C. Pirlea, V. Olaru, and C. Sminchisescu, "AIFit: Automatic 3D human-interpretable feedback models for fitness training," in *CVPR*, 2021.
- [16] S. Streuber, M. A. Quiros-Ramirez, M. Q. Hill, C. A. Hahn, S. Zuffi, A. O'Toole, and M. J. Black, "Body talk: Crowdshaping realistic 3D avatars with words," *ACM TOG*, 2016.
- [17] Y. Jiang, Z. Huang, X. Pan, C. C. Loy, and Z. Liu, "Talk-to-edit: Fine-grained facial editing via dialog," in *ICCV*, 2021.
- [18] A. Ghosh, N. Cheema, C. Oguz, C. Theobalt, and P. Slusallek, "Synthesis of compositional animations from textual descriptions," in *ICCV*, 2021.
- [19] H. Ahn, T. Ha, Y. Choi, H. Yoo, and S. Oh, "Text2action: Generative adversarial synthesis from language to action," in *ICRA*, 2018.
- [20] M. Petrovich, M. J. Black, and G. Varol, "Action-conditioned 3D human motion synthesis with transformer VAE," in *ICCV*, 2021.
- [21] C. Ahuja and L.-P. Morency, "Language2pose: Natural language grounded pose forecasting," *3DV*, 2019.
- [22] G. Pavlakos, X. Zhou, and K. Daniilidis, "Ordinal depth supervision for 3D human pose estimation," in *CVPR*, 2018.
- [23] G. Pons-Moll, D. J. Fleet, and B. Rosenhahn, "Posebits for monocular human pose estimation," in *CVPR*, 2014.
- [24] Delmas, Ginger and Weinzaepfel, Philippe and Lucas, Thomas and Moreno-Noguer, Francesc and Rogez, Grégory, "PoseScript: 3D Human Poses from Natural Language," in *ECCV*, 2022.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *NeurIPS*, 2017.
- [26] J. Fu, S. Li, Y. Jiang, K.-Y. Lin, C. Qian, C.-C. Loy, W. Wu, and Z. Liu, "Stylegan-human: A data-centric odyssey of human generation," *ECCV*, 2022.
- [27] J. Yu, H. Zhu, L. Jiang, C. C. Loy, W. Cai, and W. Wu, "CelebV-Text: A large-scale facial text-video dataset," in *CVPR*, 2023.
- [28] Y. Jiang, S. Yang, H. Qiu, W. Wu, C. C. Loy, and Z. Liu, "Text2human: Text-driven controllable human image generation," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 4, pp. 1–11, 2022.
- [29] Y. Zhang, R. Briq, J. Tanke, and J. Gall, "Adversarial synthesis of human pose from text," in *GCPR*, 2020.
- [30] R. Briq, P. Kochar, and J. Gall, "Towards better adversarial synthesis of human images from text," *arXiv preprint arXiv:2107.01869*, 2021.
- [31] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *ECCV*, 2014.
- [32] H. Kim, A. Zala, G. Burri, and M. Bansal, "FixMyPose: Pose correctional captioning and retrieval," in *AAAI*, 2021.
- [33] C. Guo, X. Zuo, S. Wang, S. Zou, Q. Sun, A. Deng, M. Gong, and L. Cheng, "Action2motion: Conditioned generation of 3D human motions," in *ACMMM*, 2020.
- [34] T. Lucas, F. Baradel, P. Weinzaepfel, and G. Rogez, "PoseGPT: Quantizing human motion for large scale generative modeling," in *ECCV*, 2022.

- [35] M. Plappert, C. Mandery, and T. Asfour, "Learning a bidirectional mapping between human whole-body motion and natural language using deep recurrent neural networks," *Robotics Auton. Syst.*, 2018.
- [36] A. S. Lin, L. Wu, R. Corona, K. W. H. Tai, Q. Huang, and R. J. Mooney, "Generating animated videos of human activities from natural language descriptions," in *NeurIPS workshops*, 2018.
- [37] T. Yamada, H. Matsunaga, and T. Ogata, "Paired recurrent autoencoders for bidirectional translation between robot actions and linguistic descriptions," *IEEE RAL*, 2018.
- [38] M. Petrovich, M. J. Black, and G. Varol, "Temos: Generating diverse human motions from textual descriptions," in *ECCV*, 2022.
- [39] C. Guo, S. Zou, X. Zuo, S. Wang, W. Ji, X. Li, and L. Cheng, "Generating diverse and natural 3d human motions from text," in *CVPR*, 2022.
- [40] G. Tevet, B. Gordon, A. Hertz, A. H. Bermano, and D. Cohen-Or, "Motionclip: Exposing human motion generation to clip space," in *ECCV*, 2022.
- [41] C. Guo, X. Zuo, S. Wang, and L. Cheng, "Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts," in *ECCV*, 2022.
- [42] J. Kim, J. Kim, and S. Choi, "Flame: Free-form language-based motion synthesis & editing," *AAAI*, 2023.
- [43] G. Tevet, S. Raab, B. Gordon, Y. Shafir, D. Cohen-or, and A. H. Bermano, "Human motion diffusion model," in *The Eleventh International Conference on Learning Representations*, 2023. [Online]. Available: <https://openreview.net/forum?id=SJKSyO2jwu>
- [44] M. Zhang, Z. Cai, L. Pan, F. Hong, X. Guo, L. Yang, and Z. Liu, "Motiondiffuse: Text-driven human motion generation with diffusion model," *arXiv preprint arXiv:2208.15001*, 2022.
- [45] F. Hong, M. Zhang, L. Pan, Z. Cai, L. Yang, and Z. Liu, "Avatarclip: Zero-shot text-driven generation and animation of 3d avatars," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 4, pp. 1–19, 2022.
- [46] K. Youwang, K. Ji-Yeon, and T.-H. Oh, "Clip-actor: Text-driven recommendation and stylization for animating human meshes," in *ECCV*, 2022.
- [47] L. Bourdev and J. Malik, "Poselets: Body part detectors trained using 3D human pose annotations," in *ICCV*, 2009.
- [48] P. Roy, S. Ghosh, S. Bhattacharya, U. Pal, and M. Blumenstein, "Tips: Text-induced pose synthesis," in *The European Conference on Computer Vision (ECCV)*, October 2022.
- [49] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang, "Deepfashion: Powering robust clothes recognition and retrieval with rich annotations," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1096–1104.
- [50] M. Plappert, C. Mandery, and T. Asfour, "The kit motion-language dataset," *Big data*, 2016.
- [51] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, "NTU RGB+D: A large scale dataset for 3D human activity analysis," in *CVPR*, 2016.
- [52] W. Li, Z. Zhang, and Z. Liu, "Action recognition based on a bag of 3d points," in *CVPR Workshops*, 2010.
- [53] J. Romero, D. Tzionas, and M. J. Black, "Embodied hands: Modeling and capturing hands and bodies together," in *SIGGRAPH Asia*, 2017.
- [54] S. Muralidhar Jayanthi, D. Pruthi, and G. Neubig, "Neuspell: A neural spelling correction toolkit," in *EMNLP*, 2020.
- [55] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," in *EMNLP*, 2014.
- [56] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *EMNLP*, 2014.
- [57] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter," *NeurIPS workshop*, 2019.
- [58] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. Osman, D. Tzionas, and M. J. Black, "Expressive body capture: 3D hands, face, and body from a single image," in *CVPR*, 2019.
- [59] N. Vo, L. Jiang, C. Sun, K. Murphy, L.-J. Li, L. Fei-Fei, and J. Hays, "Composing text and image for image retrieval-an empirical odyssey," in *CVPR*, 2019.
- [60] H. Joo, N. Neverova, and A. Vedaldi, "Exemplar fine-tuning for 3D human model fitting towards in-the-wild 3D human pose estimation," in *3DV*, 2020.
- [61] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *ICLR*, 2014.
- [62] O. Rybkin, K. Daniilidis, and S. Levine, "Simple and effective vae training with calibrated decoders," in *ICML*, 2021.
- [63] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, "On the continuity of rotation representations in neural networks," in *CVPR*, 2019.
- [64] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black, "Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image," in *ECCV*, 2016.
- [65] K. Papineni, S. Roukos, T. Ward, and W. jing Zhu, "Bleu: a method for automatic evaluation of machine translation," in *ACL*, 2002.
- [66] C.-Y. Lin, "Rouge: A package for automatic evaluation of summaries," in *Text summarization branches out*, 2004.
- [67] S. Banerjee and A. Lavie, "METEOR: An automatic metric for MT evaluation with improved correlation with human judgments," in *ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, 2005.
- [68] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.



Vision applications 2016), outstanding reviewer awards at ECCV 2012 and CVPR 2014, 2021, and Google and Amazon Faculty Research Awards in 2017 and 2019, respectively. He has (co)authored over 150 publications in refereed journals and conferences (including 10 IEEE Transactions on PAMI, 5 Intl. Journal of Computer Vision, 28 CVPR, 11 ECCV and 9 ICCV).



2015), a Research Scientist with the LEAR/THOTH team at Inria Grenoble Rhône-Alpes (2015-2018) and since 2019 he is a Senior Research Scientist at NAVER LABS Europe where he leads the computer vision group. His research interests include computer vision and machine learning, with a special focus on understanding people from visual data.