

# ZooBuilder: 2D and 3D Pose Estimation for Quadrupeds Using Synthetic Data


A. S. Fangbemi<sup>1</sup>, Y. F. Lu<sup>1</sup>, M. Y. Xu<sup>2</sup>, X. W. Luo<sup>1</sup>, A. Rolland<sup>1</sup>, C. Raissi<sup>3,4</sup>

<sup>1</sup>Ubisoft China AI & Data Lab

<sup>2</sup>Sichuan University

<sup>3</sup>Ubisoft, Singapore

<sup>4</sup>INRIA Nancy Grand Est, France



**Figure 1:** Process from data generation to 3D pose estimation: (a) Original FBX animation file, (b) Camera setup in MotionBuilder, (c) sample training data, (d) 2D pose inference, (e) 3D pose inference.

## Abstract

This work introduces a novel strategy for generating synthetic training data for 2D and 3D pose estimation of animals using keyframe animations. With the objective to automate the process of creating animations for wildlife, we train several 2D and 3D pose estimation models with synthetic data, and put in place an end-to-end pipeline called **ZooBuilder**. The pipeline takes as input a video of an animal in the wild, and generates the corresponding 2D and 3D coordinates for each joint of the animal's skeleton. With this approach, we produce motion capture data that can be used to create animations for wildlife.


## CCS Concepts

• *Computing methodologies* → *Animation*;

## 1 Introduction

Thanks to the relative abundance of training datasets for humans, researchers have developed reliable deep learning approaches for 2D and 3D pose estimation [CHS\*18]. These datasets are usually created in a controlled environment such as a motion capture room where actors perform a variety of actions. In contrast to humans, bringing wild animals in a mocap room can be a challenging task which limits the availability of training datasets and research on the topic. In this paper, we address the difficult work of generating animal datasets for quadrupeds pose estimation by creating synthetic data from existing keyframe animations. Given an FBX animation file of a cougar, we create a set of cameras around it to build a virtual mocap room (Figure 1). From each camera view and for each frame of the animation, we extract the 2D and 3D coordinates of the skeleton, an image, as well as the parameters of the camera. Using this data, we retrain 2D and a 3D pose estimation models originally developed for humans. We include these models in an end-to-end pipeline called “ZooBuilder”, which takes as input a video of a cougar and outputs the corresponding 3D animation.

## 2 Approach



**Figure 2:** Samples of the dataset used for training 2D pose estimation models.

**Training data generation.** We import FBX files containing keyframe animations of a cougar into Autodesk MotionBuilder and

create a set of twelve cameras. For each camera and frame, we generate: (1) the 3D world coordinates for 37 joints, (2) a rendered image, (3) the 3D world coordinates of the camera, its projection matrix and the 3D world coordinates of its focal point. We further generate the 3D view coordinates of the joints from the camera perspective by multiplying their 3D world coordinates by the view matrix. The 3D view coordinates are then used to compute the 2D coordinates of the joints on the projection plane, which are finally converted to 2D image coordinates. [DP\*10].

In addition, we post process images generated from MotionBuilder to add random backgrounds in order to increase their realism [ZKL\*16]. The final training dataset consists of (1) cougar images with background, (2) 2D image coordinates (Figure 2), and (3) 3D view coordinates of each joint for each image.

**ZooBuilder pipeline.** ZooBuilder is an end-to-end pipeline that generates FBX animation files from 2D videos of animals. It involves three core steps leveraging machine learning models (ML) for (1) object detection (YOLOv3 [RF18]), (2) 2D pose estimation (OpenPose [CHS\*18]), and (3) 3D pose estimation (Pose\_3D [RIHL18]). We adapt the 2D and 3D pose estimation models to our synthetic cougar data and retrain them.

### 3 Challenges

**Irrealistic training data.** The main challenge to overcome with our approach resides in the difference of distribution between the synthetic training data and the real test data collected from the wild. To ensure that models trained on synthetic data generalize well on realistic videos, we explore two main data transformation techniques, changing both training and test images into (1) gray-scale images or (2) using a style transfer technique to standardize the data (Figure 3). Although training and testing our 2D pose estimation models on gray-scale images results in a significant improvement in the prediction accuracy, the results with style transfer under perform predictions done without style transfer.




**Figure 3:** Transformation to standardize train data (above) and test data (below) with style transfer (c, d) or gray-scale (e, f).

**Training data volume.** ML models are usually trained with a large quantity of images. Using transfer learning strategies, we leverage features learned by Resnet50 on the ImageNet dataset, and further

train our 2D pose estimation model with our synthetic data composed of about 170 thousand images. To increase the diversity of images, we introduce different data augmentation techniques during the training phase such as image and joints rotation, scaling, flipping, Gaussian noise, color jittering, random brightness and contrast.

**Machine learning models tuning.** Apart from the challenges mentioned above, retraining human based models for a cougar skeleton and dataset requires to invest efforts in fine-tuning the models. We explore different hyper-parameters configurations for the learning rate, optimizer, activation functions. We adopt a layer-specific learning rate approach coupled with a learning scheduler with the SGD optimizer to train a model that generalizes relatively well.



**Figure 4:** Examples of 2D (above) and 3D (below) results on test data from the wild

### 4 Conclusion

Quadrupeds pose estimation is a challenging task due to the lack of training data. In this work, we proposed a strategy to fill this gap by generating a synthetic dataset based on existing keyframe animations. Using the data generated, we successfully trained 2D and 3D pose estimation models, and integrated them into an end-to-end pipeline that infers 3D animations from videos of real cougars (Figure 4). The pipeline still shows limitations when exposed to complex videos including multiple subjects, major occlusions, low contrast, etc. providing opportunities for further improvements on both the dataset generation and ML models development. This work could benefit from the contributions of both academic and industry peers who focus on quadrupeds pose estimation and animations.

### References

- [CHS\*18] CAO Z., HIDALGO G., SIMON T., WEI S.-E., SHEIKH Y.: Openpose: realtime multi-person 2d pose estimation using part affinity fields. *arXiv preprint arXiv:1812.08008* (2018). 1, 2
- [DP\*10] DUNN F., PARBERRY I., ET AL.: *3D math primer for graphics and game development*. Jones & Bartlett Publishers, 2010. 2
- [RF18] REDMON J., FARHADI A.: Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767* (2018). 2
- [RIHL18] RAYAT IMTIAZ HOSSAIN M., LITTLE J. J.: Exploiting temporal information for 3d human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)* (2018), pp. 68–84. 2
- [ZKL\*16] ZHOU B., KHOSLA A., LAPEDRIZA A., TORRALBA A., OLIVA A.: Places: An image database for deep scene understanding. *arXiv preprint arXiv:1610.02055* (2016). 2