

Pose-Conditioned Joint Angle Limits for 3D Human Pose Reconstruction

Ijaz Akhter, Michael J. Black

Max Planck Institute for Intelligent Systems, Tübingen, Germany



Figure 1: **Joint-limit dataset.** We captured a new dataset for learning pose-dependent joint angle limits. This includes an extensive variety of stretching poses. A few sample images are shown here. We use this dataset to learn pose-conditioned joint-angle limits. The dataset and the learned joint-angle model will be made publicly available.

Accurate modeling of priors over 3D human pose is fundamental to many problems in computer vision. Most previous priors are either not general enough for the diverse nature of human poses or not restrictive enough to avoid invalid 3D poses. We propose a physically-motivated prior that only allows anthropometrically valid poses and restricts the ones that are invalid.

One can use joint-angle limits to evaluate whether two connected bones are valid or not. However, it is established in biomechanics that there are dependencies in joint-angle limits between certain pair of bones [4, 6]. For example how much one can flex one’s arm depends on whether it is in front of, or behind, the back. Medical textbooks only provide joint-angle limits in a few positions [2, 8] and the complete configuration of pose-dependent joint-angle limits for the full body is unknown.

We found that existing mocap datasets (like the CMU dataset) are insufficient to learn true joint angle limits, in particular limits that are pose dependent. Therefore we captured a new dataset of human motions that includes an extensive variety of stretching poses performed by trained athletes and gymnasts (see Fig. 1). We learn *pose-dependent joint angle limits* from this data and propose a novel prior based on these limits.

The proposed prior can be used for problems where estimating 3D human pose is ambiguous. Our pose parametrization is particularly simple and general in that the 3D pose of the kinematic skeleton is defined by the two endpoints of each bone in Cartesian coordinates. Constraining a 3D pose to remain valid during optimization simply requires the addition of our penalty term in the objective function. We also show that our prior can be combined with a sparse representation of poses, selected from an overcomplete dictionary, to define a general yet accurate parametrization of human pose.

We use our prior to estimate 3D human pose from 2D joint locations. Given a single view in Fig. 2, the 3D pose is ambiguous [9] and there exist several plausible 3D poses all resulting in the same 2D observations. Thus no generic prior information about static body pose is sufficient to guarantee a single correct 3D pose. Here we seek the most probable, valid, human pose.

We show that a critical step for 3D pose estimation given 2D point locations is the estimation of camera parameters. Given the diversity of human poses, incorrect camera parameters can lead to an incorrect pose estimate. To solve this problem we propose a grouping of body parts, called the “extended-torso,” consisting of the torso, head, and upper-legs. Exploiting the fact that the pose variations for the extended-torso are fewer than for the full-body, we estimate its 3D pose and the corresponding camera parameters more easily. The estimated camera parameters are then used for full-body




Figure 2: We use our joint-angle-limit prior for 3D pose estimation given 2D joint locations in an image. The proposed prior helps in reducing the space of possible solutions to only valid 3D human poses. Our prior can be also used for many other problems where estimating 3D human pose is ambiguous.

pose estimation. The proposed multi-step solution gives substantially improved results over previous methods.

We evaluate 2D pose estimation from 2D for a wide range of poses and camera views using activities from the CMU motion capture dataset. These are more complex and varied than the data used by previous methods [3, 7] and we show that previous methods have trouble in this case. We also report superior results on manual annotations and automatic part-based detections [5] on the Leeds sports pose dataset. The data used for evaluation and all software is publicly available for other researchers to compare with our results [1]

- [1] <http://poseprior.is.tue.mpg.de/>.
- [2] United States. National Aeronautics and Space Administration. *NASA-STD-3000: Man-systems integration standards*. Number v. 3 in NASA-STD. National Aeronautics and Space Administration, 1995. URL <http://msis.jsc.nasa.gov/sections/section03.htm>.
- [3] Xiaochuan Fan, Kang Zheng, Youjie Zhou, and Song Wang. Pose locality constrained representation for 3d human pose reconstruction. In *Computer Vision—ECCV 2014*, pages 174–188. Springer, 2014.
- [4] H Hatze. A three-dimensional multivariate model of passive human joint torques and articular boundaries. *Clinical Biomechanics*, 12(2): 128–135, 1997.
- [5] Martin Kiefel and Peter Gehler. Human pose estimation with fields of parts. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, volume 8693 of *Lecture Notes in Computer Science*, pages 331–346. Springer International Publishing, September 2014.
- [6] T Kodek and M Munich. Identifying shoulder and elbow passive moments and muscle contributions. In *IEEE Int. Conf. on Intelligent Robots and Systems*, volume 2, pages 1391–1396, 2002.
- [7] V. Ramakrishna, T. Kanade, and Y. Sheikh. Reconstructing 3D human pose from 2D image landmarks. *European Conference on Computer Vision*, pages 573–586, 2012.
- [8] M. Schünke, E. Schulte, and U. Schumacher. *Prometheus: Allgemeine Anatomie und Bewegungssystem : LernAtlas der Anatomie*. Prometheus LernAtlas der Anatomie. Thieme, 2005. ISBN 9783131395214.
- [9] C. Sminchisescu and B. Triggs. Building roadmaps of local minima of visual models. In *European Conference on Computer Vision*, volume 1, pages 566–582, Copenhagen, 2002.