

SIMSPINE: A Biomechanics-Aware Simulation Framework for 3D Spine Motion Annotation and Benchmarking

Muhammad Saif Ullah Khan Didier Stricker
German Research Center for Artificial Intelligence (DFKI)

<https://saifkhichi96.github.io/research/simspine/>

Modeling spinal motion is fundamental to understanding human biomechanics, yet remains underexplored in computer vision due to the spine’s complex multi-joint kinematics and the lack of large-scale 3D annotations. We present a biomechanics-aware keypoint simulation framework that augments existing human pose datasets with anatomically consistent 3D spinal keypoints derived from musculoskeletal modeling. Using this framework, we create the first open dataset, named SIMSPINE, which provides sparse vertebral-level 3D spinal annotations for natural full-body motions in indoor multi-camera capture without external restraints. With 2.14 million frames, this enables data-driven learning of vertebral kinematics from subtle posture variations and bridges the gap between musculoskeletal simulation and computer vision. In addition, we release pretrained baselines covering fine-tuned 2D detectors, monocular 3D pose lifting models, and multi-view reconstruction pipelines, establishing a unified benchmark for biomechanically valid spine motion estimation. Specifically, our 2D spine baselines improve the state-of-the-art from 0.63 to 0.80 AUC in controlled environments, and from 0.91 to 0.93 AP for in-the-wild spine tracking. Together, the simulation framework and SIMSPINE dataset advance research in vision-based biomechanics, motion analysis, and digital human modeling by enabling reproducible, anatomically grounded 3D spine estimation under natural conditions.

1. Introduction

The vertebral column, together with the pelvic girdle, forms the biomechanical core of the human skeleton—bearing axial loads, enabling locomotion, and protecting the spinal cord. Comprised of over two dozen articulating vertebrae and intervertebral joints, each with multiple degrees of freedom, the spine exhibits highly nonlinear, interdependent motion patterns [71]. Despite decades of work across anatomy, biomechanics, neuromechanics, and rehabilitation science, precise intervertebral kinematics remain debated [18, 30, 51], with poor agreement on ranges of motion across *in vivo* studies [34, 61, 73], cadaveric experi-




Figure 1. **SIMSPINE annotations.** Neutral pose of the simulated spine model (left), divided into three anatomical regions—cervical (pink), thoracic (blue), and lumbar (green)—with 15 annotated landmarks: 9 along the vertebral column, 2 on the skull, 2 at the clavicle joints, and 2 on the shoulder blades (*last 4 not shown*). Cervical and thoracic motion is limited to transition junctions (indicated by anatomical axis markers), while intermediate vertebrae remain rigidly coupled. The lumbar segment (L1–L5) is fully articulated, with intervertebral rotations simulated for all degrees of freedom (bottom right). Overall spinal curvature is characterized by thoracic kyphosis (θ_k) and lumbar lordosis (θ_l) angles. Motion was generated using [Beaucage-Gauvreau et al.](#)’s musculoskeletal model [5] as a function of full-body movement in Human3.6M, with 5 training and 2 validation subjects performing 15 actions.

ments [52, 65], and computational models. Among these, *in vivo* imaging (fluoroscopy, MRI, X-ray) and motion capture are most relevant for vision-based human motion research.

Traditional motion capture excels at tracking large-scale limb motion for applications such as action recognition and

human–computer interaction, but it misses subtle movements—vertebral rotations, postural sway, compensatory pelvic tilts—that influence spinal stability, load distribution, and injury risk. This limits use in sports injury prevention, ergonomics, rehabilitation, and clinical contexts.

Recent RGB approaches include *marker-assisted* tracking with perforated kinesiology tape [19] and a real-world dataset/model for 2D spinal keypoint tracking [31]. The former is difficult to deploy outside controlled settings; the latter, while scalable, remains 2D with opaque annotations and limited biomechanical verification. Thus, achieving biomechanically accurate, clinically relevant, and *unconstrained* 3D estimation of healthy and pathological spine motion remains open, where “unconstrained” means not restricted to stationary subjects, close-up unclothed views, or fixed camera angles. Following [31], we argue that the *first step* towards solving this problem is to curate a comprehensive 3D spine motion dataset that can serve both as a baseline and a stepping stone towards finding a robust solution.

In this work, we use **biomechanics-aware keypoint simulation** to annotate an existing large-scale 3D pose dataset [27] with vertebra-level keypoints, enabling learning of spinal micro-movements from subtle whole-body posture shifts. We release: (1) SIMSPINE, a comprehensive **dataset** containing anatomically valid 3D spinal annotations for unconstrained motions, (2) a **simulation pipeline** for generating biomechanically consistent spinal motion using musculoskeletal models, and (3) a set of **pretrained baselines** including fine-tuned 2D and 3D pose estimation models. Together, these resources provide the first open benchmark for vision-based spinal motion analysis.

Scope and limitations. Annotations are simulation-derived rather than measured *in vivo*. The cervical and thoracic regions are modeled as rigid segments with motion permitted only at transition junctions (Fig. 1); the lumbar spine (L1–L5) is fully articulated. Motions are sourced from Human3.6M indoor activities, so appearance diversity and clinical pathologies are out of scope, and the resource is intended for research—not diagnostic—use. The biomechanically constrained 3D spine motion also supports improved realism in downstream animation and avatar synthesis tasks.

Contributions:

- Biomechanics-aware simulation framework that generates anatomically valid 3D spine motion data, augmenting existing human pose datasets with precise spinal labels.
- Pretrained spine motion baselines, including 6 fine-tuned 2D detectors for in-the-wild full-body pose estimation with state-of-the-art spine tracking, 2 monocular 3D baselines for root-relative 3D, and one multi-view 3D method for high-precision 3D in absolute metric coordinate space.
- Public release of the full simulation pipeline, pretrained models, and dataset—the first open 3D resource for spine motion estimation—supporting reliability, reproducibil-

ity, and benchmarking in spine-aware pose estimation.

2. Related Work

The Human Spine. The human spine (Fig. 1) consists of 33 vertebrae forming an S-shaped curve in the sagittal plane and appearing straight coronally in neutral posture [18]. Of these, 24 are mobile (7 cervical, 12 thoracic, 5 lumbar). Each functional motion segment permits three rotations (flexion/extension, lateral bending, axial rotation) with small, constrained translations [51]. Reported ranges of motion (ROM) vary widely across cadaveric, *in vivo*, and computational studies [18, 52, 65], highlighting the lack of a single “ground-truth” kinematic profile. *Cervical (neck)*. Most mobile region: C0–C1 mainly flexion–extension, C1–C2 primarily axial rotation, and C3–C7 distributed coupled motion [41, 51]. Dual-fluoroscopy with model-based tracking (MBT) measures sub-degree motions but reveals task- and level-specific coupling [3, 39, 43]. *Thoracic (upper back)*. Constrained by the rib cage and coronal facets, it allows moderate lateral bending and greater axial rotation in upper levels, decreasing toward the thoracolumbar junction [18, 38]. Upright biplanar radiography shows posture-dependent coupling not visible in supine imaging [15, 26]. *Lumbar (lower back)*. Large flexion/extension but limited axial rotation due to sagittal facet orientation and tall discs; rotation and lateral bending generally decrease caudally [18, 51]. Imaging studies report small per-segment rotations and load-dependent coupling [39, 49].

Clinical Spine Tracking. Clinical assessment of *static shape and function* relies on radiographs using metrics such as Cobb’s angle and regional curvature [11]. Low-dose EOS enables upright 3D reconstructions from simultaneous AP/lateral views with improved reproducibility [15, 17, 26, 48]. CT/MRI segmentation supports vertebra-level analysis and planning but typically in static postures [8, 37, 64]. For *dynamic motion*, time-varying vertebral motion is measured with *biplane fluoroscopy* plus MBT, achieving sub-millimeter accuracy when validated against RSA [25, 29, 39]. Dynamic MRI offers non-ionizing alternatives for posture-dependent deformation, but frame rate and artifacts limit kinematic fidelity [35, 43, 49]. Indirect RGB-based gait biomarkers for scoliosis screening [22, 32, 57, 77] scale to large cohorts but lack vertebra-level precision.

Medical Imaging-Based Spine Analysis. Radiographs remain standard for global alignment and deformity metrics [11]. EOS extends these to upright 3D reconstructions [15, 26]. CT provides detailed bone morphology with strong performance on segmentation benchmarks [8, 37, 64]. MRI captures soft-tissue and some dynamic changes but with lower geometric fidelity for bone [35, 49]. Dual-fluoroscopy with MBT yields the most accurate *in*

Table 1. **Comparison of existing spine datasets.** Summary of public datasets relevant to spinal imaging, scoliosis analysis, and motion estimation. Our proposed dataset uniquely provides vertebra-level 3D kinematics from RGB videos, bridging clinical imaging, biomechanics, and computer vision domains. *Availability:* ✓ Public, † On request, ✗ Private.

Dataset	Year	Input(s)	Label(s)	# Samples	Tasks	Annotation Notes	Avail.
Kim et al. [33]	2019	Radiographs	Lumbar spine positions (sagittal, 2D)	797 images	Lumbar vertebrae segmentation	2x expert (radiologists)	✗
Hornig et al. [24]	2019	Radiographs	Vertebral segmentation masks (coronal, 2D)	595 images	Vertebral segmentation	Multi-expert (clinical)	†
VinDr-SpineXR [58]	2021	Radiographs	ROI boxes for 13 abnormalities	10k images / 5k cases	Lesion detection	Expert annotations	✓
VerSe [64]	2021	CT Scans	Vertebra centroid / segmentation masks (sagittal, 2D)	374 scans / 355 patients	Vertebral segmentation	Hybrid (voxel)	✓
Scoliosis1K [77]	2024	Silhouette Images	Demographics, Scoliosis Diagnosis, 2D Body Pose	1.5k videos (447k frames) from 1k walking subjects	Scoliosis detection	Evaluation by medical professionals	✓
USTC&SYSU-Scoliosis [79]	2025	Radiographs, RGB	Regional Cobb’s angles	1898 images / 1067 patients aged 10-18	Scoliosis detection	Multi-expert agreement	†
SpineTrack [31]	2025	Images (RGB) Videos (RGB)	9 spine keypoints (various viewpoints, 2D) 3D spine position, vertebral rotations	33k images 1.56M train, 0.58M test images / 7 people, 15 actions	Keypoint Estimation Keypoint Estimation, Rotation Regression	Hybrid, non-expert, 2D, outdoor, inconsistent Real+sim, biomech. plausible, indoor	✓ ✓
SIMSPINE (Ours)	2025	Videos (RGB)	3D spine position, vertebral rotations	1.56M train, 0.58M test images / 7 people, 15 actions	Keypoint Estimation, Rotation Regression	Real+sim, biomech. plausible, indoor	✓

vivo kinematics, though costly and dose-intensive; synthetic CT approaches seek to mitigate this [36, 39].

RGB-Based Spine Motion Tracking. RGB systems use marker-assisted tracking for structured back motion capture [19], but require exposed skin and controlled views. Markerless models estimate 2D spinal keypoints in natural scenes [31], facing projection ambiguity and limited biomechanical validity. No RGB-only method yet reconstructs biomechanically consistent 3D vertebral motion under unconstrained conditions, motivating simulation-driven supervision with kinematic priors. To provide further context, Tab. 1 compares key spine datasets across imaging, clinical, and vision domains. Our dataset uniquely combines real RGB sequences with anatomically constrained 3D motion.

OpenSim Simulation. OpenSim [12] is a standard platform for musculoskeletal modeling, inverse kinematics/dynamics, and forward simulation. Widely used baselines [59] and SimTK repositories provide spine-focused models such as: full-body and lumbar-detailed variants [5, 10], thoracolumbar and rib-cage models with articulated T1–L5 [6], personalizable and pediatric spines [2, 63], and cervical/neck and impact-oriented models [7, 50]. Toolchains like Pose2Sim [53] connect RGB-derived trajectories with OpenSim for simulation-ready kinematics. AdBIOmechanics [70] provides datasets and a tool to automatically add kinematics and dynamics to joint trajectories. These provide components which our framework links with vision for large-scale image annotation.

3D Pose Estimation. 2D-to-3D lifting predicts absolute or root-relative 3D coordinates from monocular keypoints, most often trained on Human3.6M [27]. Foundational models range from simple MLP baselines [46] to temporal convolutional networks with semi-supervision [54]. Recent transformer-based approaches integrate motion priors for spatial–temporal consistency, such as MotionBERT [78],

MotionAGFormer [47], EvoPose with structure priors [76], and PriorFormer for real-time geometric lifting [1]. While these achieve scalable full-body reconstruction, they overlook anatomical validity and intervertebral coherence.

3. Methodology

Biomechanics-aware keypoint simulation, illustrated in Fig. 2, augments Human3.6M with sparse 3D positions of vertebral bodies, per-vertebra rotational kinematics, and subject-scaled biomechanical models. We follow the standard Human3.6M splits (train: S1, S5, S6, S7, S8; test: S9, S11), preserving Human3.6M time stamps.

3.1. Biomechanics-aware keypoint simulation

(1) Multi-view spinal detection and triangulation. For each synchronized frame t and view v , a pretrained detector [31] predicts 2D spinal keypoints $\hat{\mathbf{u}}_{v,t} \in \mathbb{R}^{2 \times K_s}$ (with $K_s = 9$). Using calibrated intrinsics/extrinsics $\{\mathbf{K}_v, \mathbf{R}_v, \mathbf{t}_v\}$ from Human3.6M, we recover pseudo-3D points $\tilde{\mathbf{X}}_t \in \mathbb{R}^{3 \times K_s}$ by robust triangulation:

$$\tilde{\mathbf{X}}_t = \arg \min_{\mathbf{X}} \sum_{v \in \mathcal{V}} \rho \left(\left\| \Pi(\mathbf{K}_v[\mathbf{R}_v|\mathbf{t}_v]\mathbf{X}) - \hat{\mathbf{u}}_{v,t} \right\|_2^2 \right),$$

where $\Pi(\cdot)$ denotes perspective projection and ρ is a robust penalty (Huber). We prune outliers by view-consistency and reprojection-error thresholds, and apply zero-phase low-pass filtering to suppress frame-to-frame jitter.

(2) Merging with ground-truth body markers. Let $\mathbf{Y}_t \in \mathbb{R}^{3 \times K_h}$ be the Human3.6M 3D markers¹ (K_h joints) in the camera/world frame. We align the pseudo-3D spinal set $\tilde{\mathbf{X}}_t$ and $\mathbf{Y}_t$ into a common OpenSim marker set $\mathcal{M}$ by (i)

¹ Pelvis, Spine, Neck, HeadTop, and Nose from Human3.6M are not used because of overlap with our spinal points and, for the latter two points, inconsistent labels [42].




Figure 2. **Biomechanics-aware keypoint simulation pipeline.** From synchronized multi-view RGB, a 2D detector [31] predicts spinal landmarks that are robustly triangulated using calibrated cameras to obtain pseudo-3D spinal keypoints. These pseudo labels are temporally aligned and merged with known Human3.6M [27] 3D markers (GT 3D Pose). OpenSim inverse kinematics (IK) [12] fits a subject-scaled full-body model [5, 53, 59] to the merged trajectories. We attach virtual markers to vertebral bodies and, using the IK joint angles and subject-specific anthropometrics, generate anatomically consistent spine keypoints via forward kinematics (FK). The pipeline also outputs biomechanical parameters (e.g., per-vertebra rotations).

selecting semantic correspondences, (ii) temporal synchronization at Human3.6M frame times, and (iii) filling missing views by short-horizon interpolation. The merged markers $\mathbf{Z}_t = \{\mathbf{Y}_t, \tilde{\mathbf{X}}_t\} \in \mathbb{R}^{3 \times K}$ serve as IK targets.

(3) Subject scaling and inverse kinematics. We adopt a full-body OpenSim model based on Rajagopal et al. [59] with lumbar spine details adapted from Beaucage-Gauvreau et al. [5]; we use Pose2Sim [53] for data I/O and utilities. Scaling uses subject height/mass estimates derived from TRC anthropometrics (with conservative outlier trimming). IK solves, per time t , the weighted least-squares problem

$$\mathbf{q}_t^* = \arg \min_{\mathbf{q}_t} \sum_{m \in \mathcal{M}} w_m \|\mathbf{z}_{m,t} - \hat{\mathbf{z}}_m(\mathbf{q}_t)\|_2^2 + \lambda \|\mathbf{D}\mathbf{q}\|_2^2,$$

where $\hat{\mathbf{z}}_m(\mathbf{q}_t)$ are model marker positions from FK of joint state $\mathbf{q}_t$, w_m are per-marker confidences (higher for Human3.6M markers, lower for pseudo spinal points), and $\mathbf{D}$ penalizes joint-velocity/acceleration for temporal smoothness. The model includes a *fully articulated lumbar spine* with three rotational DOFs at intervertebral joints from T12–L1 through L5–S1; a single 3-DOF joint at the cervicothoracic junction provides an aggregate neck DOF. Thoracic and cervical bodies beyond this aggregate are treated as rigid segments with neutral baseline curvature. This choice approximates thoracic rib-cage constraints while keeping the model identifiable with RGB-derived inputs.

(4) Virtual vertebral markers and forward kinematics. We attach virtual markers to vertebral bodies (centroidal locations) and compute their 3D trajectories from the IK solution $\{\mathbf{q}_t^*\}$ via FK. These markers define the 3D spinal keypoints distributed along the column (sacral base to lower cervical). In parallel, we export per-vertebra Euler rotations

about anatomical axes (flexion/extension, lateral bending, axial rotation) as biomechanical parameters.

(5) Quality control and curation. We (i) reject frames with implausible curvature, (ii) clamp rare angle discontinuities from gimbal wrap, and (iii) apply temporal smoothing and interpolation to fill small gaps and ensure motion continuity. Subject-specific scaled OpenSim models, marker positions, and joint angles, time synchronized with RGB frames from Human3.6M, are generated. **Markers:** 37 total, with 12 limb points from Human3.6M, 15 high-precision new points that directly model the spine, and 10 pseudo-labels on feet and face. **Kinematic axes:** 62, including 56 Euler angles. Further details are in Supplementary Sec. B.

3.2. Biomechanical validity of simulated data

We use statistical analysis to understand the simulated spine kinematics, and compare with biomechanics literature to determine the dataset’s suitability for downstream tasks.

Thoracolumbar spine curvature. The violin plots in Fig. 3 (left) show distributions of the Lumbar Lordotic Angle (LLA) and Thoracic Kyphotic Angle (TKA), computed using Cobb’s method [11] between L1–S1 and T3–T12 endplates, respectively. Outliers were conservatively removed. The observed distributions align with *in vivo* studies and reveal activity-specific trends: seated motions (e.g., *Sitting*, *SittingDown*) exhibit reduced lumbar lordosis compared to standing actions (e.g., *Walking*, *Greeting*) due to gravitational effects on intervertebral spacing [4, 9, 66, 69], while actions involving arm elevation (e.g., *Photo*, *Posing*, *Greeting*) show reduced thoracic kyphosis from posterior shoulder displacement [60]. Across all actions, mean lordosis (33–39°) and kyphosis (29–37°) fall within normative adult




Figure 3. **Thoracolumbar spine in SIMSPINE.** *Left:* Distributions of thoracolumbar curvature across actions, defined by the Lumbar Lordotic Angle (LLA, θ_l) in the lower back and Thoracic Kyphotic Angle (TKA, θ_k) in the upper back. LLA and TKA average within 1 SD at 32–41° and 28–39°, respectively, indicating greater curvature in the lower spine but higher variability in the upper. Values fall within reported biomechanical ranges [16, 40], confirming that SIMSPINE produces anatomically plausible curvatures and captures expected action-specific postural trends [4, 9, 60, 66, 69]. *Right:* Per-vertebra range of motion (ROM) on y-axis for the three lumbar rotational DOFs. Our simulated data (solid) follows similar trends as reported by White and Panjabi (1978) [72] (dashed).




Figure 4. **Cervical spine in SIMSPINE.** The distributions (per action) remain centered near neutral with task-dependent spread, reflecting that our model uses a single 3-DOF aggregate neck joint while keeping the thoracic/cervical bodies rigid beyond the cervicothoracic junction. This is within the neck ROM reported in [14] with approximately half coverage, which indicates the presence of only small head movements in the dataset.

ranges [16, 40]. These bounded, unimodal, and action-sensitive distributions reflect realistic postural variability in healthy adults, underscoring the biomechanical fidelity of the simulated data. Details of subject and action-specific ranges is provided in Supplementary Tab. B2.

Lumbar spine ROM trends. The lumbar ROM curves in Fig. 3, (right) display known qualitative gradients: flexion/extension increases from the thoracolumbar junction toward the caudal levels with a peak near L4–L5, then reduces at L5–S1; lateral bending peaks mid-lumbar; axial rota-

tion is modest overall and highest in mid-lumbar segments. These monotone or near-monotone trends match widely reported patterns [72] for upright motion and are difficult to reproduce with unconstrained pose-only priors; their presence suggests the IK solution respects both morphology and motion constraints. Simulated values follow similar trends as the reference data. It should be noted that exact ROM values for lumbar vertebrae are debated [62].

Cervical proxy. Despite modeling the cervical spine as a single aggregate 3-DOF joint, action-conditioned neck ROM distributions in Fig. 4 stay within physiologically credible envelopes and broaden for tasks with head motion. For example, the *Phone* action has the largest lateral bending range indicative of a sideways head tilt typical during phone conversations. This supports usefulness for coarse head–neck kinematics and for supervising 2D-to-3D lifting in the absence of vertebra-resolved cervical detail.

These graphs jointly indicate that our simulation yields *anatomically plausible* and *action-sensitive* kinematics. This supports the dataset’s suitability as supervision for learning vertebral motion from natural movements and as a benchmark for biomechanics-aware spine motion models.

4. Baselines and Experiments

We benchmark three tasks enabled by SIMSPINE: (1) 2D pose estimation from RGB, (2) multiview 3D pose reconstruction, and (3) monocular 3D pose lifting. We formally define each task, describe the baselines, training methods, evaluation protocols, and present ablation studies. The aim

is to provide deployment-ready, robust pretrained models, and quantitative results which future spine research can use as a reference. Known limitations are also discussed.

Notation. Let B be batch size, T frames per clip ($T=1$ for image-based; $T>1$ for video-based), K vertebral key-points, and $C_{\text{img}}=3$ image channels. For each sequence, $\mathbf{I} \in \mathbb{R}^{B \times T \times C_{\text{img}} \times H \times W}$ are RGB frame(s), $\mathbf{U} \in \mathbb{R}^{B \times T \times K \times 2}$ are 2D keypoints, and $\mathbf{Y} \in \mathbb{R}^{B \times T \times K \times 3}$ are 3D coordinates in a global camera or anatomical frame.

4.1. 2D Pose Estimation

Definition. This task aims to learn a function

$$f_{2D} : \mathbb{R}^{B \times T \times C_{\text{img}} \times H \times W} \rightarrow \mathbb{R}^{B \times T \times K \times 2},$$

predicting $\hat{\mathbf{U}}$ from images with supervision $\mathbf{U}$ and visibility weights $\mathbf{V} \in [0, 1]^{B \times T \times K}$.

Baselines. We fine-tune three variants of SpinePose [31] and compare them with representative architectures from both CNN and transformer families: HRNet-w32 [67], RTMPose-m [28], and ViTPose-b [74]. HRNet and ViTPose are heatmap-based methods while RTMPose is a coordinate classification method.

Training. All models are initialized with pretrained weights and fine-tuned on a balanced combination of outdoor (SpineTrack) and indoor (SIMSPINE) images, with equal samples per batch. Only 2% of SIMSPINE train split are used to avoid overfitting while still transferring simulation-derived knowledge. Each model is fine-tuned for 10 epochs with a three-stage curriculum where data augmentation transitions from hard to easy (see Supplementary Sec. A). Optimizers and losses follow SpinePose [31].

Table 2. **Baselines for 2D Spine Pose Estimation.** Performance of various CNN and transformer architectures on the SpineTrack and SIMSPINE benchmarks.

Method	Dataset		SpineTrack				Ours
	Pretrain	Finetune	AP ^B	AR ^B	AP ^S	AR ^S	
SpinePose-s	SpineTrack	-	0.792	0.821	0.896	0.908	0.611
SpinePose-m			0.840	0.864	0.914	0.926	0.633
SpinePose-l			0.854	0.877	0.910	0.922	0.633
SpinePose-s-ft	SpineTrack	SpineTrack + Ours	0.788	0.815	0.920	0.929	0.790
SpinePose-m-ft			0.821	0.846	0.928	0.937	0.798
SpinePose-l-ft			0.840	0.862	0.917	0.927	0.803
HRNet-w32	COCO	SpineTrack + Ours	0.776	0.806	0.905	0.918	0.769
RTMPose-m			0.832	0.858	0.925	0.935	0.794
ViTPose-b			0.835	0.866	0.921	0.933	0.794

Evaluation. On SpineTrack we report AP/AR using COCO-style OKS with body (B) and spine (S) subsets following [31]. On SIMSPINE we report AUC of PCK over thresholds $\tau \in [0, 0.5]$, with distances normalized by the shorter side of the person bounding box. Higher is better for all metrics. Comparisons with the original SpinePose [31] quantify the effect of consistent labeling in SIMSPINE.

4.2. Multiview 3D Reconstruction

Definition. Given synchronized detections $\{\hat{\mathbf{U}}^{(m)}\}_{m=1}^M$ and calibrated projections $\{P^{(m)}\}_{m=1}^M$, triangulate

$$\hat{\mathbf{Y}} = f_{\text{tri}}(\{\hat{\mathbf{U}}^{(m)}, P^{(m)}\}_{m=1}^M),$$

by minimizing multi-view reprojection error.

Baselines. We use a weighted linear least-squares triangulation with confidence-based outlier rejection [20, 21]. Two 2D detectors are compared: the zero-shot *SpinePose-m* pretrained on SpineTrack [31], and its fine-tuned variant *SpinePose-m-ft* adapted on SIMSPINE.

Table 3. **Multiview 3D Spine Reconstruction.** MPJPE (mm) across actions and spinal regions using linear triangulation. S_C : Cervical, S_T : Thoracic, S_L : Lumbar, S : Full spine, B : Body, and All: Complete skeleton.

Action	S_C	S_T	S_L	S	B	All
Mean (GT 2D)	12.36	8.80	5.09	9.00	7.07	7.85
Directions	33.69	39.86	40.16	37.88	27.63	31.79
Discussion	36.35	40.55	41.44	39.39	28.33	32.81
Eating	38.47	43.04	35.40	39.48	27.44	32.32
Greeting	30.66	35.19	33.73	33.29	25.57	28.70
Phoning	42.41	44.64	39.14	42.43	27.23	33.40
Photo	36.40	42.13	42.27	40.25	30.35	34.37
Posing	33.48	37.46	40.14	36.85	28.02	31.60
Purchases	34.81	35.67	39.63	36.44	26.64	30.61
Sitting	40.94	43.16	31.21	39.23	25.42	31.02
SittingDown	34.85	35.95	23.63	32.30	23.31	26.95
Smoking	38.51	44.83	35.66	40.28	27.26	32.54
Waiting	35.17	37.40	33.22	35.54	24.78	29.14
WalkDog	35.53	41.44	40.83	39.31	29.49	33.47
WalkTogether	34.05	43.88	40.40	39.68	29.28	33.50
Walking	33.02	43.80	41.53	39.60	29.06	33.33
Mean (Finetuned)	36.50	41.07	37.13	38.50	27.27	31.82
Mean (Zero-Shot)	54.39	69.90	48.13	58.92	42.74	49.30

Evaluation. For triangulated poses in global world coordinates, we report Mean Per-Joint Position Error (MPJPE):

$$\text{MPJPE} = \frac{1}{BTK} \sum_{b,t,j} \|\mathbf{Y}_{b,t,j}^{\text{pred}} - \mathbf{Y}_{b,t,j}\|_2.$$

Results are summarized in Tab. 3. To further isolate geometric shape from global alignment, Tab. 4 reports Procrustes-aligned MPJPE (P-MPJPE).

Table 4. **Triangulation Baseline: P-MPJPE (mm).** Same setup as Table 3, evaluated after similarity alignment.

Action	S_C	S_T	S_L	S	B	All
Mean (GT 2D)	0.33	0.43	0.06	0.67	1.81	1.79
Mean (Finetuned)	19.68	20.67	12.21	26.67	24.21	29.53
Mean (Zero-Shot)	26.88	23.40	9.92	39.01	42.25	46.48

Sub-millimeter values when using GT 2D confirm geometric consistency within floating-point precision. However, detector-based reconstructions remain in the 20–40 mm range because alignment corrects translation and rotation but not inter-view noise.

Table 5. **Simple Baselines for Monocular 3D Spine Pose Lifting.** Evaluation of Martinez et al.’s lifting model trained on spine-only (15 joints) and full-body (37 joints) keypoint sets. Reported metric: Procrustes-aligned MPJPE (P-MPJPE, mm) per action. Training on full-body joints improves spine localization accuracy. Evaluation is on 15 spine joints only.

Train Set	2D	Direct	Discuss	Eating	Greet	Phone	Photo	Pose	Purchase	Sit	SitD	Smoke	Wait	WalkD	WalkT	Walk	Avg
Spine Only	Det.	16.58	17.92	17.79	18.32	21.81	19.74	15.69	17.70	19.28	22.32	18.43	18.82	18.91	16.24	15.96	18.58
Full-Body	Det.	<u>15.55</u>	<u>16.40</u>	<u>14.45</u>	<u>15.98</u>	<u>19.84</u>	<u>17.66</u>	<u>14.25</u>	<u>15.68</u>	<u>15.97</u>	<u>19.23</u>	<u>16.38</u>	<u>17.17</u>	<u>17.41</u>	<u>12.24</u>	<u>12.12</u>	<u>16.28</u>
Spine Only	GT	15.66	17.01	16.05	17.43	20.64	18.34	14.87	16.40	17.91	22.10	16.92	17.78	17.96	15.39	15.28	17.52
Full-Body	GT	<u>11.56</u>	<u>13.05</u>	<u>11.36</u>	<u>13.61</u>	<u>18.06</u>	<u>14.61</u>	<u>10.60</u>	<u>12.39</u>	<u>13.64</u>	<u>17.36</u>	<u>13.25</u>	<u>14.47</u>	<u>14.51</u>	<u>9.15</u>	<u>9.98</u>	<u>13.48</u>

Table 6. **Monocular 3D Spine Pose Lifting (Component-Wise).** Same experimental setup as Table 5, but results are aggregated across actions and broken down by spine components. Columns report P-MPJPE (mm) for S_C (cervical), S_T (thoracic), S_L (lumbar), and S (full spine). Rows compare training on spine-only (15 joints) vs. full-body (37 joints); 2D inputs are either detected (Det.) or ground truth (GT).

Train Set	2D	P-MPJPE				MPJPE			
		S_C	S_T	S_L	S	S_C	S_T	S_L	S
Spine Only	Det.	12.98	14.29	3.32	18.58	96.88	75.96	18.83	67.70
Full-Body	Det.	<u>12.81</u>	<u>10.12</u>	<u>3.51</u>	<u>16.28</u>	<u>100.50</u>	<u>72.61</u>	<u>18.16</u>	<u>67.39</u>
Spine Only	GT	12.03	13.45	3.32	17.52	77.93	58.70	15.54	53.60
Full-Body	GT	<u>10.94</u>	<u>7.98</u>	<u>2.03</u>	<u>13.48</u>	<u>39.24</u>	<u>26.97</u>	<u>7.76</u>	<u>25.94</u>

This experiment quantifies the geometric upper bound achievable conditioned on 2D detection accuracy alone.

4.3. Monocular 3D Lifting

Definition. Learn

$$f_{\text{lift}} : \mathbb{R}^{B \times T \times K \times 2} \rightarrow \mathbb{R}^{B \times T \times (K-1) \times 3},$$

mapping 2D keypoints to root-centered 3D. Let $\mathcal{R}$ be the root index set (here $\mathcal{R}=\{0\}$). For frame t , $\bar{\mathbf{y}}_{b,t} = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \mathbf{Y}_{b,t,r}$; and $\tilde{\mathbf{Y}}_{b,t,j,:} = \mathbf{Y}_{b,t,j,:} - \bar{\mathbf{y}}_{b,t}$; the network minimizes $\|\hat{\mathbf{Y}} - \tilde{\mathbf{Y}}\|_2^2$.

Baselines. We adopt SimpleBaseline3D [46] as the reference architecture due to its simplicity and interpretability. It encodes per-joint features through linear layers with residual connections, operating on either spine-only or full-body keypoints.

Training. 3D targets are root-centered and standardized per joint. The root joint is excluded during training and reinserted at inference. Frames are sampled at 1 Hz to preserve motion diversity while maintaining manageable sequence length. Optimization uses AdamW [44] with step scheduling. As in [46], a standard MSE loss is used for training.

Evaluation. We report root-relative Mean Per-Joint Position Error (MPJPE) and Procrustes-aligned MPJPE (P-MPJPE) on decoded 3D predictions. For P-MPJPE, each predicted frame is first aligned to the ground truth using the Procrustes similarity alignment ($s, \mathbf{R}, \mathbf{t}$). All metrics are reported in millimeters. Tab. 5 summarizes the P-MPJPE of

spine keypoints across all dataset actions, and Tab. 6 compares both P-MPJPE and root-relative MPJPE with further breakdown of three key spinal regions. Exact keypoints evaluated in each region are in Supplementary Sec. B.

4.4. Ablation Studies

We report here controlled studies isolating the effect of data mixing, sampling, and training choices.




Figure 5. **Ablation Study: Mixup Composition.** We examine how the fraction of SIMSPINE used in training influences indoor (AUC) and outdoor (AP) performance. Increasing the SIMSPINE fraction improves indoor performance up to 10%, while outdoor gains saturate by 2–5%. Per-batch (PB) mixup maintains the best balance between indoor and outdoor metrics, whereas per-epoch (PE) alternation favors one domain at the expense of the other. Sampling only 2% of SIMSPINE achieves near-saturated results on both datasets, indicating diminishing returns from larger fractions.

Mixup Composition and Dataset Fraction. Figure 5 analyzes how the fraction of SIMSPINE contributes to 2D fine-tuning performance under different data mixing strategies. This ablation is motivated by the need for stable and accurate 2D keypoint detectors, which form the foundation for reliable 3D triangulation and lifting. Existing off-the-shelf 2D detectors trained on manually annotated spine datasets often exhibit label noise and jitter due to inconsistent supervision, making it essential to study how synthetic and real data can be best combined. Increasing the proportion of SIMSPINE from 0% to 10% steadily improves indoor AUC (from 0.61 to 0.79), with minimal degradation on outdoor AP. Beyond 2–5%, however, outdoor AP plateaus around 0.84, suggesting that only a small subset of high-quality indoor samples is sufficient to enhance generalization. Comparing mixing strategies, per-batch (PB) mixup consistently outperforms per-epoch (PE) alternation, yielding the most stable trade-off across domains. PB works better because

the model and optimizer jointly observe samples from both datasets in each iteration, allowing AdamW to maintain smoother gradient statistics and consistent momentum estimates across domains. These results justify using 2% of SIMSPINE—about 31k indoor images, roughly matching the 33k outdoor samples in SpineTrack—with per-batch mixing as the default configuration, achieving balanced exposure across domains and efficient use of synthetic data while preserving real-world performance.

4.5. Discussion

Across our three experiments—2D detection, multiview triangulation, and monocular lifting—we establish reference baselines for spine-aware pose estimation. In 2D, fine-tuning across SpineTrack and SIMSPINE consistently improves *spine* metrics (indoor AUC from 0.61 to 0.80; SpineTrack AP^S from 0.91 to 0.93), with a small trade-off on body AP^B relative to the strongest SpinePose pre-trained model. The multiview triangulation baseline attains **31.8 mm** MPJPE and **29.5 mm** P-MPJPE, and our oracle using GT 2D reaches sub-millimeter P-MPJPE, confirming geometric consistency. For monocular lifting, the full-body variant outperforms the spine-only variant (detected 2D: 18.6 mm $\rightarrow$ 16.3 mm P-MPJPE; GT 2D: 17.5 mm $\rightarrow$ 13.5 mm), indicating that global context aids vertebral localization. Together, these results provide a practical baseline suite that links strong 2D cues to geometrically faithful 3D reconstructions and highlights where temporal or biomechanical priors add the most value.

Limitations. Our simulation framework is a kinematics-only, simulation-derived resource and carries several limitations. Anatomically, we articulate five intervertebral lumbar joints and a single 3-DOF cervicothoracic joint, while treating the remaining thoracic and cervical segments as rigid. This choice, made for numerical stability and identifiability from RGB inputs, neglects rib-cage coupling and soft-tissue effects that materially constrain and distribute thoracic motion [38]. Intervertebral translations are not modeled; although small in healthy spines, nonzero translations have been measured in vivo with stereoradiography [55, 56]. Pelvis–lumbar coupling is also simplified, which can underrepresent lumbopelvic rhythm during trunk motion [68]. Subjects are implicitly healthy and scaled by height/mass; age-, sex-, and pathology-dependent variation in sagittal alignment is not modeled, so the dataset encodes nominal healthy-motion priors rather than the diversity seen clinically [45, 75]. Because all motions are sourced from Human3.6M, the visual domain is limited to indoor captures with fixed multi-view cameras and a restricted action set [27]. This may limit generalization to fully unconstrained, in-the-wild scenarios. Our OpenSim step solves inverse kinematics only; muscle actuation, ground-reaction forces, and load equilibria are not enforced,

so trajectories are geometrically plausible but not physically validated [23]. Extending to dynamics-consistent optimal control would allow force-consistent motion generation [13]. Our empirical validation (Sec. 3.2) target geometric plausibility (curvature envelopes, ROM profiles) rather than absolute accuracy against in vivo ground truth such as dual-fluoroscopy/biplane tracking or standing biplanar reconstructions, which provide the most precise in vivo vertebral kinematics and alignment [3, 48]. Accordingly, biomechanics-aware keypoint simulation framework should be regarded as a scalable proxy for method development and benchmarking rather than clinical measurement, while SIMSPINE serves best as a large-scale pretraining resource for spine pose estimation models later fine-tuned on smaller, biomechanically validated datasets.

Future work. Angular annotations in SIMSPINE, not evaluated in this work because of additional design choices (representation, normalization, interpretation) orthogonal to this initial benchmark, should be benchmarked. In addition, we see three priorities: (i) expand anatomical fidelity by adding rib-cage articulation and small intervertebral translations, (ii) couple IK with dynamics (inverse dynamics or optimal control) to ensure force-consistent motion [13], and (iii) broaden the visual domain with in-the-wild sequences and pathology-specific cohorts. Longer term, combining upright clinical imaging (EOS/fluoroscopy) with RGB could anchor subject-specific spine priors and reduce ambiguity in vertebra-level motion.

5. Conclusion

We introduced biomechanics-aware keypoint simulation, a pipeline that fuses calibrated RGB, subject-scaled musculoskeletal models, and virtual vertebral markers to produce anatomically constrained 3D spine motion from standard multi-view footage. The resulting SIMSPINE provides 15 keypoints driving spine with per-segment rotations over 2.14M frames and a set of pretrained baselines spanning 2D detection, multiview triangulation, and monocular lifting. The evidence—near-oracle P-MPJPE under GT 2D, consistent thoracolumbar curvature envelopes, and action-sensitive ROM profiles—shows that simulation-driven annotation can extend existing datasets with biomechanically meaningful structure. While simplified and domain-limited, the framework offers a practical bridge between computer vision and musculoskeletal modeling, enabling models that reason about posture and vertebral kinematics, not only joint geometry.

Availability and licensing. Code, models, and SIMSPINE annotations for the spine markers and kinematic parameters will be released for research use only. Due to licensing, full-body keypoints are reproducible by running our pipeline on the licensed Human3.6M data.

Acknowledgement

This work was co-funded by the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101135724 (LUMINOUS) and Grant Agreement No 101092889 (SHARESPACE).

References

- [1] Mohamed Adjel and Vincent Bonnet. Priorformer: A transformer for real-time monocular 3d human pose estimation with versatile geometric priors. *arXiv preprint arXiv:2508.18238*, 2025. 3
- [2] Dennis Anderson, Hossein Mokhtarzadeh, Brett Allaire, Katelyn Burkhart, and M Bouxsein. Subject-specific spine models for 250 individuals from the framingham heart study. *Harvard Dataverse*. doi: doi/10.7910/DVN/SJ5MVM, 2020. 3
- [3] William J Anderst, Emma Baillargeon, William F Donaldson III, Joon Y Lee, and James D Kang. Validation of a noninvasive technique to precisely measure in vivo three-dimensional cervical spine movement. *Spine*, 36(6):E393–E400, 2011. 2, 8
- [4] Jun Seok Bae, Jee-Soo Jang, Sang-Ho Lee, and Jin Uk Kim. A comparison study on the change in lumbar lordosis when standing, sitting on a chair, and sitting on the floor in normal individuals. *Journal of Korean Neurosurgical Society*, 51(1): 20–23, 2012. 4, 5
- [5] Erica Beaucage-Gauvreau, William SP Robertson, Scott CE Brandon, Robert Fraser, Brian JC Freeman, Ryan B Graham, Dominic Thewlis, and Claire F Jones. Validation of an open-sim full-body model with detailed lumbar spine for estimating lower lumbar spine loads during symmetric and asymmetric lifting tasks. *Computer methods in biomechanics and biomedical engineering*, 22(5):451–464, 2019. 1, 3, 4
- [6] Alexander G Bruno, Mary L Bouxsein, and Dennis E Anderson. Development and validation of a musculoskeletal model of the fully articulated thoracolumbar spine and rib cage. *Journal of biomechanical engineering*, 137(8):081003, 2015. 3
- [7] Dario Cazzola, Timothy P Holsgrove, Ezio Preatoni, Harinderjit S Gill, and Grant Trewartha. Cervical spine injuries: a whole-body musculoskeletal model for the analysis of spinal loading. *PloS one*, 12(1):e0169329, 2017. 3
- [8] Pengfei Cheng, Yusheng Yang, Huiqiang Yu, and Yongyi He. Automatic vertebrae localization and segmentation in ct with a two-stage dense-u-net. *Scientific Reports*, 11(1): 22156, 2021. 2
- [9] Il Youp Cho, Si Young Park, Jong Hoon Park, Tae Kwon Kim, Tae Wan Jung, and Hyun Min Lee. The effect of standing and different sitting positions on lumbar lordosis: radiographic study of 30 healthy volunteers. *Asian spine journal*, 9(5):762, 2015. 4, 5
- [10] Miguel Christophy, Nur Adila Faruk Senan, Jeffrey C Lotz, and Oliver M O'Reilly. A musculoskeletal model for the lumbar spine. *Biomechanics and modeling in mechanobiology*, 11(1):19–34, 2012. 3
- [11] JR Cobb. Outline for the study of scoliosis. *Instructional course lecture*, 1948. 2, 4
- [12] Scott L Delp, Frank C Anderson, Allison S Arnold, Peter Loan, Ayman Habib, Chand T John, Eran Guendelman, and Darryl G Thelen. Opensim: open-source software to create and analyze dynamic simulations of movement. *IEEE transactions on biomedical engineering*, 54(11):1940–1950, 2007. 3, 4
- [13] Christopher L Dembia, Nicholas A Bianco, Antoine Falisse, Jennifer L Hicks, and Scott L Delp. Opensim moco: Musculoskeletal optimal control. *PLOS Computational Biology*, 16(12):e1008493, 2020. 8
- [14] Arockia Selvakumar Arockia Doss, Pavan Kalyan Lingampally, Gia Minh Thao Nguyen, and Daniel Schilberg. A comprehensive review of wearable assistive robotic devices used for head and neck rehabilitation. *Results in Engineering*, 19: 101306, 2023. 5
- [15] Jean Dubousset, Georges Charpak, Irene Dorion, Wafa Skalli, Francois Lavaste, Jacques Deguise, Gabriel Kalifa, and Solene Ferey. A new 2d and 3d imaging approach to musculoskeletal physiology and pathology with low-dose radiation and the standing position: the eos system. *Bulletin de l'Academie nationale de medecine*, 189(2):287–97, 2005. 2
- [16] Gerald T Fon, Michael J Pitt, and A Cole Thies Jr. Thoracic kyphosis: range in normal subjects. *American Journal of Roentgenology*, 134(5):979–983, 1980. 5
- [17] Laurent Gajny, Shahin Ebrahimi, Claudio Vergari, Elsa Angelini, and Wafa Skalli. Quasi-automatic 3d reconstruction of the full spine from low-dose biplanar x-rays based on statistical inferences and image analysis. *European Spine Journal*, 28(4):658–664, 2019. 2
- [18] Fabio Galbusera and Hans-Joachim Wilke. *Biomechanics of the spine: Basic concepts, spinal disorders and treatments*. Academic Press, 2018. 1, 2
- [19] Hendrik Hachmann and Bodo Rosenhahn. Human spine motion capture using perforated kinesiology tape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5149–5157, 2023. 2, 3
- [20] Richard Hartley. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 6
- [21] Richard I Hartley and Peter Sturm. Triangulation. *Computer vision and image understanding*, 68(2):146–157, 1997. 6
- [22] Zhong He, Neng Lu, Yi Chen, Elvis Chun-Sing Chui, Zhen Liu, Xiaodong Qin, Jie Li, Shengru Wang, Junlin Yang, Zhiwei Wang, et al. Conditional generative adversarial network-assisted system for radiation-free evaluation of scoliosis using a single smartphone photograph: a model development and validation study. *Eclinicalmedicine*, 75, 2024. 2
- [23] Jennifer L Hicks, Thomas K Uchida, Ajay Seth, Apoorva Rajagopal, and Scott L Delp. Is my model good enough? best practices for verification and validation of musculoskeletal models and simulations of movement. *Journal of biomechanical engineering*, 137(2):020905, 2015. 8
- [24] Ming-Huwi Horng, Chan-Pang Kuok, Min-Jun Fu, Chii-Jen Lin, and Yung-Nien Sun. Cobb angle measurement of spine from x-ray images using convolutional neural network. *Computational and mathematical methods in medicine*, 2019(1): 6357171, 2019. 3

- [25] Ali Humadi, Sulaf Dawood, Klas Halldin, and Brian Freeman. Rsa in spine: a review. *Global Spine Journal*, 7(8): 811–820, 2017. 2
- [26] Ludovic Humbert, Jacques A De Guise, Benjamin Aubert, Benoît Godbout, and Wafa Skalli. 3d reconstruction of the spine from biplanar x-rays using parametric models based on transversal and longitudinal inferences. *Medical engineering & physics*, 31(6):681–687, 2009. 2
- [27] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, 2014. 2, 3, 4, 8
- [28] Tao Jiang, Peng Lu, Li Zhang, Ningsheng Ma, Rui Han, Chengqi Lyu, Yining Li, and Kai Chen. RtmPose: Real-time multi-person pose estimation based on mmpose, 2023. 6
- [29] Craig C Kage, Mohsen Akbari-Shandiz, Mary H Foltz, Rebekah L Lawrence, Taycia L Brandon, Nathaniel E Helwig, and Arin M Ellingson. Validation of an automated shape-matching algorithm for biplane radiographic spine osteokinematics and radiostereometric analysis error quantification. *PLoS One*, 15(2):e0228594, 2020. 2
- [30] Ibrahim A Kapandji. The physiology of the joints, volume i, upper limb. *American Journal of Physical Medicine & Rehabilitation*, 50(2):96, 1971. 1
- [31] Muhammad Saif Ullah Khan, Stephan Krauß, and Didier Stricker. Towards unconstrained 2d pose estimation of the human spine. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 6171–6180, 2025. 2, 3, 4, 6
- [32] Dong Suk Kim, Shi Hwan Park, Tae Sik Goh, Seung Min Son, and Jung Sub Lee. A meta-analysis of gait in adolescent idiopathic scoliosis. *Journal of Clinical Neuroscience*, 81: 196–200, 2020. 2
- [33] Kang Cheol Kim, Hyun Cheol Cho, Tae Jun Jang, Jong Mun Choi, and Jin Keun Seo. Automatic detection and segmentation of lumbar vertebrae from x-ray images for compression fracture evaluation. *Computer Methods and Programs in Biomedicine*, 200:105833, 2021. 3
- [34] Michal Kozanek, Shaobai Wang, Peter G Passias, Qun Xia, Gang Li, Christopher M Bono, Kirkham B Wood, and Guoan Li. Range of motion and orientation of the lumbar facet joints in vivo. *Spine*, 34(19):E689–E696, 2009. 1
- [35] Kornelia Kulig, Robert F Landel, and Christopher M Powers. Assessment of lumbar spine kinematics using dynamic mri: a proposed mechanism of sagittal plane motion induced by manual posterior-to-anterior mobilization. *Journal of Orthopaedic & Sports Physical Therapy*, 34(2):57–64, 2004. 2
- [36] Seth J Kussow, Jared L Zitnay, Penny R Atkins, and Andrew E Anderson. Accuracy and reliability of synthetic computed tomography for model-based tracking of biplane videoradiography data: Sj kussow et al. *Annals of Biomedical Engineering*, pages 1–14, 2025. 3
- [37] Nikolas Lessmann, Bram Van Ginneken, Pim A De Jong, and Ivana Išgum. Iterative fully convolutional neural networks for automatic vertebra segmentation and identification. *Medical image analysis*, 53:142–155, 2019. 2
- [38] Christian Liebsch and Hans-Joachim Wilke. Basic biomechanics of the thoracic spine and rib cage. In *Biomechanics of the Spine*, pages 35–50. Elsevier, 2018. 2, 8
- [39] Cheng-Chung Lin, Tung-Wu Lu, Ting-Ming Wang, Chao-Yu Hsu, Shih-Jung Hsu, and Ting-Fang Shih. In vivo three-dimensional intervertebral kinematics of the subaxial cervical spine during seated axial rotation and lateral bending via a fluoroscopy-to-ct registration approach. *Journal of biomechanics*, 47(13):3310–3317, 2014. 2, 3
- [40] RM Lin, I-Ming Jou, and Chin-Yin Yu. Lumbar lordosis: normal adults. *Journal of the Formosan Medical Association= Taiwan yi zhi*, 91(3):329–333, 1992. 5
- [41] Sara Lindenmann, Christos Tsagkaris, Mazda Farshad, and Jonas Widmer. Kinematics of the cervical spine under healthy and degenerative conditions: a systematic review. *Annals of Biomedical Engineering*, 50(12):1705–1733, 2022. 2
- [42] Filipa Lino, Carlos Santiago, and Manuel Marques. Benchmarking 3d human pose estimation models under occlusions. *arXiv preprint arXiv:2504.10350*, 2025. 3
- [43] Ao Liu, Nan-Hai Qiu, Xue-Ren Zhong, Xiang Fang, Jun-Jian Liao, Zhi-Peng Zhang, Pei-Feng Zheng, Yong-Yu Hu, Kong-He Hu, Ying-Hui Xiong, et al. Dynamic evaluation of the cervical spine by kinematic mri in patients with cervical spinal cord injury without fracture and dislocation. *Journal of Orthopaedic Surgery and Research*, 18(1):249, 2023. 2
- [44] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 7
- [45] Oliver Ludwig, Carlo Dindorf, Jens Kelm, Steven Simon, Felix Nimmrichter, and Michael Fröhlich. Reference values for sagittal clinical posture assessment in people aged 10 to 69 years. *International Journal of Environmental Research and Public Health*, 20(5):4131, 2023. 8
- [46] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 2640–2649, 2017. 3, 7
- [47] Soroush Mehraban, Vida Adeli, and Babak Taati. Motionagformer: Enhancing 3d human pose estimation with a transformer-gcnformer network. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 6920–6930, 2024. 3
- [48] Elias Melhem, Ayman Assi, Rami El Rachkidi, and Ismat Ghanem. Eos® biplanar x-ray imaging: concept, developments, benefits, and limitations. *Journal of children's orthopaedics*, 10(1):1–14, 2016. 2, 8
- [49] Giulia Michelini, Antonella Corridore, Silvia Torlone, Federico Bruno, Claudia Marsecano, Raffaella Capasso, Ferdinando Caranci, Antonio Barile, Carlo Masciocchi, and Alessandra Splendiani. Dynamic mri in the evaluation of the spine: state of the art. *Acta Bio Medica: Atenei Parmensis*, 89(Suppl 1):89, 2018. 2
- [50] Jonathan D Mortensen, Anita N Vasavada, and Andrew S Merryweather. The inclusion of hyoid muscles improve moment generating capacity and dynamic simulations in musculoskeletal models of the head and neck. *PLoS one*, 13(6): e0199912, 2018. 3

- [51] Donald A Neumann. Kinesiology of the musculoskeletal system: foundations for rehabilitation. *Mosby Elsevier*, 2010. 1, 2
- [52] Itaru Oda, Kuniyoshi Abumi, Bryan W Cunningham, Kiyoshi Kaneda, and Paul C McAfee. An in vitro human cadaveric study investigating the biomechanical properties of the thoracic spine. *Spine*, 27(3):E64–E70, 2002. 1, 2
- [53] David Pagnon, Mathieu Domalain, and Lionel Reveret. Pose2sim: an end-to-end workflow for 3d markerless sports kinematics—part 2: accuracy. *Sensors*, 22(7):2712, 2022. 3, 4
- [54] Dario Pavllo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7753–7762, 2019. 3
- [55] Mark Pearcy, IAN Portek, and Janis Shepherd. Three-dimensional x-ray analysis of normal movement in the lumbar spine. *Spine*, 9(3):294–297, 1984. 8
- [56] Mark J Pearcy. Stereo radiography of lumbar spine motion. *Acta Orthopaedica Scandinavica*, 56(sup212):1–45, 1985. 8
- [57] Zizhao Peng, Zihan Wang, Mengying Sun, Zheng Lv, Yan Wang, Ping Li, and Fengwei An. Graph convolutional networks for 3d skeleton-based scoliosis screening using gait sequences: Z. peng et al. *The Visual Computer*, pages 1–13, 2025. 2
- [58] Hieu Huy Pham, H Nguyen Trung, and Ha Quy Nguyen. Vindr-spinexr: A large annotated medical image dataset for spinal lesions detection and classification from radiographs. *PhysioNet*, pages RRID–SCR.007345, 2021. 3
- [59] Apoorva Rajagopal, Christopher L Dembia, Matthew S Demers, Denny D Delp, Jennifer L Hicks, and Scott L Delp. Full-body musculoskeletal model for muscle-driven simulation of human gait. *IEEE transactions on biomedical engineering*, 63(10):2068–2079, 2016. 3, 4
- [60] Alexandra Rören, Robin Ogiez, Laurent Gajny, Aurore Blasco, Fadila Mihoubi Bouvier, Antoine Feydy, François Rannou, Marie-Martine Lefèvre-Colau, and Agnès Roby-Brami. Arm elevation involves changes in the whole spine: an exploratory study using eos imaging. *BMC Musculoskeletal Disorders*, 25(1):993, 2024. 4, 5
- [61] Adam Rozumalski, Michael H Schwartz, Roy Werve, Andrew Swanson, Daryll C Dykes, and Tom Novacheck. The in vivo three-dimensional motion of the human lumbar spine during gait. *Gait & posture*, 28(3):378–384, 2008. 1
- [62] Janis Savlovskis. Range of the motion (rom) of the cervical, thoracic and lumbar spine in the traditional anatomical planes. *Anatomy Standard*, 2020. 5
- [63] Stefan Schmid, Katelyn A Burkhart, Brett T Allaire, Daniel Grindle, and Dennis E Anderson. Musculoskeletal full-body models including a detailed thoracolumbar spine for children and adolescents aged 6–18 years. *Journal of biomechanics*, 102:109305, 2020. 3
- [64] Anjany Sekuboyina, Malek E Hussein, Amirhossein Bayat, Maximilian Löffler, Hans Liebl, Hongwei Li, Giles Tetteh, Jan Kukačka, Christian Payer, Darko Štern, et al. Verse: a vertebrae labelling and segmentation benchmark for multi-detector ct images. *Medical image analysis*, 73:102166, 2021. 2, 3
- [65] Jeremy D Shaw, Daniel L Shaw, Daniel R Cooperman, Jason D Eubanks, Ling Li, and David H Kim. Characterization of lumbar spinous process morphology: a cadaveric study of 2,955 human lumbar vertebrae. *The Spine Journal*, 15(7):1645–1652, 2015. 1, 2
- [66] Tomasz Sipko, Katarzyna Barczyk-Pawełec, Mirela Piksa, and Joanna Mencil. Impact of standing and sitting postures on spinal curvature and muscle mechanical properties in young women: A photogrammetric and myotonpro analysis. *Medical Science Monitor: International Medical Journal of Experimental and Clinical Research*, 30:e944930–1, 2024. 4, 5
- [67] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5693–5703, 2019. 6
- [68] A Tafazzol, N Arjmand, Aboulfazl Shirazi-Adl, and M Parianpour. Lumbopelvic rhythm during forward and backward sagittal trunk rotations: combined in vivo measurement with inertial tracking device and biomechanical modeling. *Clinical Biomechanics*, 29(1):7–13, 2014. 8
- [69] Christos Tsagkaris, Jonas Widmer, Florian Wanivenhaus, Andrea Redaelli, Claudio Lamartina, and Mazda Farshad. The sitting vs standing spine. *North American Spine Society Journal (NASSJ)*, 9:100108, 2022. 4, 5
- [70] Keenon Werling, Nicholas A Bianco, Michael Raitor, Jon Stingel, Jennifer L Hicks, Steven H Collins, Scott L Delp, and C Karen Liu. Addbiomechanics: Automating model scaling, inverse kinematics, and inverse dynamics from human motion data through sequential optimization. *Plos one*, 18(11):e0295152, 2023. 3
- [71] AUGUSTUS A WHITE III. Clinical biomechanics of cervical spine implants. *Spine*, 14(10):1040–1045, 1989. 1
- [72] AUGUSTUS A WHITE III and Manohar M Panjabi. The basic kinematics of the human spine: a review of past and current knowledge. *Spine*, 3(1):12–20, 1978. 5
- [73] Qun Xia, Shaobai Wang, Peter G Passias, Michal Kozanek, Gang Li, Brian E Grottkau, Kirkham B Wood, and Guoan Li. In vivo range of motion of the lumbar spinous processes. *European Spine Journal*, 18(9):1355–1362, 2009. 1
- [74] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. ViTPose: Simple vision transformer baselines for human pose estimation. In *Advances in Neural Information Processing Systems*, 2022. 6
- [75] Mattia Zappalá, Stephen Lightbourne, and Nicola R Heneghan. The relationship between thoracic kyphosis and age, and normative values across age groups: a systematic review of healthy adults. *Journal of orthopaedic surgery and research*, 16(1):447, 2021. 8
- [76] Yaqi Zhang, Yan Lu, Bin Liu, Zhiwei Zhao, Qi Chu, and Nenghai Yu. Evopose: A recursive transformer for 3d human pose estimation with kinematic structure priors. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. 3

- [77] Zirui Zhou, Junhao Liang, Zizhao Peng, Chao Fan, Fengwei An, and Shiqi Yu. Gait patterns as biomarkers: A video-based approach for classifying scoliosis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 284–294. Springer, 2024. [2](#), [3](#)
- [78] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. Motionbert: A unified perspective on learning human motion representations. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15085–15099, 2023. [3](#)
- [79] Xiaojia Zhu, Rui Chen, Zhiwen Shao, Ming Zhang, Yuhu Dai, Wenzhi Zhang, and Chuandong Lang. Mgscoliosis: Multi-grained scoliosis detection with joint ordinal regression from natural image. *Alexandria Engineering Journal*, 111:329–340, 2025. [3](#)