This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

SpineSAC: Hierarchical Multi-Agent Soft Actor-Critic for Patient-Agnostic 2D–3D Spine Registration in Image-Guided Spine Surgery

Darshan K. N.[1,2], Sriram L.[2*], Sarath P.[2*], Vivek Maik[2], Manojkumar Lakshmanan[2], and Mohanasankar Sivaprakasam[1†]

[1] Dept. of Electrical Engineering, Indian Institute of Technology Madras, India
[2] Healthcare Technology Innovation Centre (HTIC), IIT Madras, India
mohan@ee.iitm.ac.in

Abstract. Intraoperative 2D-3D spine registration is a critical step in Image-Guided Spinal Surgery, requiring accurate alignment of preoperative CT volumes with intraoperative C-Arm X-ray images under stringent real-time constraints. In this paper, we propose **SpineSAC**, a Hierarchical Multi-agent Soft Actor-Critic Reinforcement Learning framework for autonomous six degree of freedom vertebral pose-based 2D-3D registration. SpineSAC integrates a Principal Component Analysis (PCA)-based Canonical Coordinate System (CCS) to establish a common anatomical reference frame across different CT and C-Arm coordinate frames and a novel Rotational Feature Encoder (RFE) that encodes Antero-Posterior and Lateral DRR and X-ray rotational misalignment into a compact 16-dimensional latent observation for the RL agent. SpineSAC achieves a Total Registration Error (TRE) of 1.45 ± 0.5 mm on phantoms and 2.31 mm on cadaver data, with a $9\times$ inference speedup over the computational approach, demonstrating its suitability for real-time intraoperative deployment. The Code is available at: SpineSAC.

Keywords: Image-Guided Spinal Surgery · 2D-3D registration · Rotational Feature Encoder · Canonical Coordinate System · Reinforcement Learning.

1 Introduction

Image-Guided Spinal Surgery (IGSS) relies on real-time imaging to improve surgical precision, particularly in Minimally Invasive Spine Surgery [14], where the surgeon has limited direct view of the anatomy [15]. A fundamental component in IGSS procedure is 2D-3D registration, which aligns a preoperative CT volume with an intraoperative C-Arm X-ray to perform the surgical planning within the CT coordinate frame [13,11]. The primary objective is to estimate the rigid six degree of freedom (6-DoF) transformation matrix that aligns the CT volume

*Equal contribution
†Corresponding author

with the C-Arm images. This is achieved by generating Digitally Reconstructed Radiographs (DRRs) from the CT and maximizing intensity-based similarity metrics such as Normalized Cross Correlation (NCC), Local NCC (LNCC), or Mutual Information against the intraoperative images using iterative optimization methods [5,22]. However, this iterative optimization until the maximum similarity score is achieved is computationally expensive across all 6-DoF combinations and always achieves multiple false convergences at different DoF. To overcome these limitations, learning-based approaches based on a deep iterative registration framework using a Point-to-Plane Correspondence pose solver avoids repeated DRR generation and mitigates local optima. However, it relies heavily on accurate DRR-to-X-ray correspondence estimation, requires large annotated training datasets, and its performance degrades under varying C-arm image quality and imaging artifacts [8].A weakly supervised 2D–3D registration framework based on differentiable rendering and ROI-based segmentation removes the need for explicit correspondence estimation, but its accuracy depends heavily on segmentation quality [2]. More recently, a Proximal Policy Optimization (PPO)-based Reinforcement Learning (RL) framework performs registration through iterative pose actions, reducing the reliance on ROI segmentation quality and explicit optimization procedures. However, it is patient-specific which requires a new model to be trained for each patient's CT data, thereby reducing its effectiveness for unseen anatomies [23]. To address the above mentioned drawbacks, we propose SpineSAC, a novel patient-agnostic RL framework for 2D-3D spine registration, with the following contributions:

- **PCA-Based Canonical Pose Representation.** We incorporate an existing PCA-based CCS to standardize 6-DoF pose representations across CT and C-Arm coordinate systems, enabling consistent RL action semantics and improved policy generalization.
- **Rotational Feature Encoder.** We propose a novel RFE that extracts a compact 16-dimensional latent embedding (z) from fixed AP, LP X-rays and their moving DRR pairs. The z captures rotational misalignment between the X-ray and DRR projections, enabling the SAC agent to select appropriate rotational actions along the r_x, r_y, and r_z axes during policy execution.
- **Hierarchical Multi-Agent SAC.** We design a coarse-to-fine hierarchical SAC framework with three resolution-specific agents performing progressive 6 DOF pose refinement, stabilized through curriculum learning, replay filtering, and knowledge transfer.

2 Methodology

In 2D-3D registration, the X-ray images get pre-processed to improve consistency with the corresponding DRRs [16]. This reduces the X-ray and DRR attribute gap and provides more reliable observations for RL-based registration. The registration is then initialized using user-defined anatomical landmark correspondences on the Antero-Posterior (AP) and Lateral (LP) X-ray images

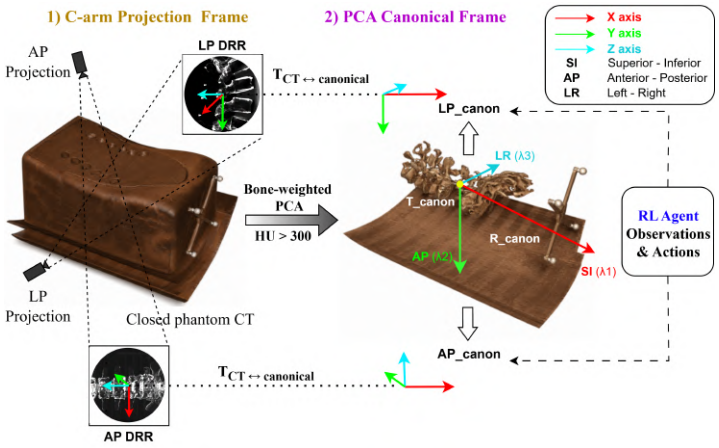


Fig. 1: PCA based canonical coordinate system.

and the corresponding CT volume, following the 6-DoF initial pose estimation method [17]. The resulting rigid transformation $T = [R|t]$ is used together with the AP and LP projection matrices to generate DRRs using an existing DRRGenerator framework given in [10], which employs VTK-based ray casting.

2.1 PCA-Based Canonical Pose Representation

Variations in C-Arm geometry and CT coordinate systems make RL policy learning challenging. To address this, we incorporate the PCA-based CCS proposed in [20,9] to represent all 6-DoF vertebral poses in a unified anatomical reference frame. PCA is applied to the spatial coordinates of bone voxels extracted from the CT volume using Hounsfield Unit (HU) thresholding. This process allows us to fix the superior-inferior (SI), anterior-posterior (AP), and left-right (LR) anatomical axes. Together with the centroid of the bone voxels, these eigenvectors define a rigid transformation $T_{\text{CT}\to\text{canonical}}$. It remaps all 6-DoF vertebral poses from the original CT frame into the CCS, which is then used throughout our SpineSAC RL environment as mentioned in Fig 1. By transforming all poses into the CCS, every translational and rotational component carries an identical anatomical meaning across all datasets, enabling direct generalization on patients CT data.

2.2 Rotational Feature Encoder

We propose a novel RFE, depicted in Fig 2, that adopts an equivariant Siamese backbone of [6], where all four images, fixed AP, LP X-ray ($\mathbf{F}^{ap}_{\text{fix}}$, $\mathbf{F}^{lp}_{\text{fix}}$), moved AP, LP DRR ($\mathbf{F}^{ap}_{\text{mov}}$, $\mathbf{F}^{lp}_{\text{mov}}$), are independently passed for feature extraction. Regular Group Equivariant Convolution (R2Conv) backbone [21] makes the extracted features robust to variations through equivariant in-plane rotations.

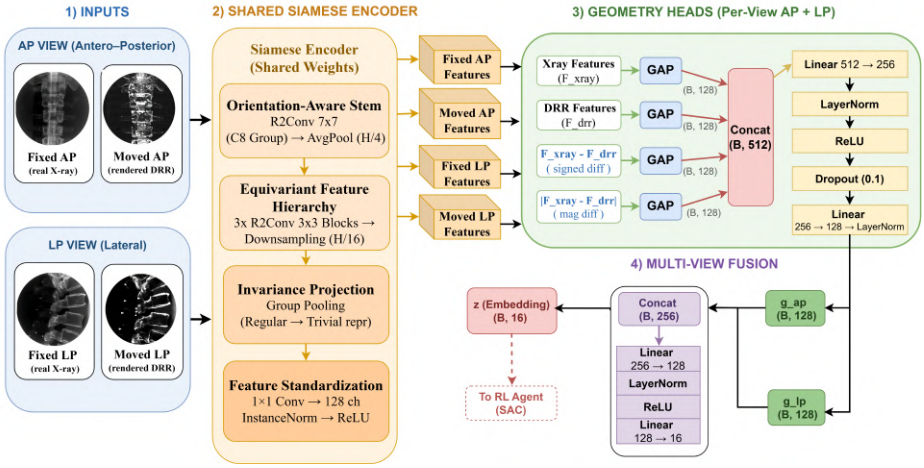


Fig. 2: Rotational feature encoder to extract a compact 16-D z from fixed AP, LP X-rays and their moving DRR pairs.

Group Pooling layer and InstanceNorm convert the equivariant features into rotation-invariant feature maps $\mathbf{F}^{ap}_{\text{fix}}$, $\mathbf{F}^{ap}_{\text{mov}}$, $\mathbf{F}^{lp}_{\text{fix}}$, and $\mathbf{F}^{lp}_{\text{mov}} \in \mathbb{R}^{B\times 128\times \text{Height}\times \text{Width}}$, where B is the batch size. For each view, the feature maps are then collapsed through Global Average Pooling (GAP) into 128-dimensional vectors. These lower-dimensional vectors are then projected via a Multi-Layer Perceptron (MLP), representing the rotational misalignment between the fixed X-ray and moved DRR for the AP and LP views, respectively. The final 16-D embedding $\mathbf{z} \in \mathbb{R}^{16}$ serves as the rotational state observation for the RL agent; trained across 3 resolution levels (0.25×, 0.5×, 1.0×) with degrees of rotation between ±20°, ±15°, and ±5° using:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{geo}} + 0.2\,\mathcal{L}_{\ell_1} + 0.1\,\mathcal{L}_{\text{emb}} \tag{1}$$

where $\mathcal{L}_{\text{geo}}$ is a geodesic rotation loss measuring angular distance on the $SO(3)$ manifold [18], $\mathcal{L}_{\ell_1}$ is a smooth ℓ_1 (Huber) loss for outlier-robust regression [4], and $\mathcal{L}_{\text{emb}}$ is an InfoNCE contrastive loss [1] that shapes $\mathbf{z}$ into a discriminative space by pulling similar rotational states together and pushing dissimilar ones apart. Once trained, the frozen RFE provides a stable $\mathbf{z}$, preventing embedding drift and accelerating policy convergence across all resolution levels.

2.3 Hierarchical Multi-Agent SAC with Training Stabilizers

Problem Formulation: As illustrated in Fig. 3, we formulate the 2D-3D registration as a Markov Decision Process, where the goal is to estimate the rigid 6-DoF transformation aligning a preoperative CT volume with intraoperative AP, LP C-Arm images. At each time step t, the agent observes state s_t and outputs a continuous action $a_t = [\Delta r_x, \Delta r_y, \Delta r_z, \Delta t_x, \Delta t_y, \Delta t_z]$, representing incremental pose updates in the PCA-based CCS. The updated pose renders

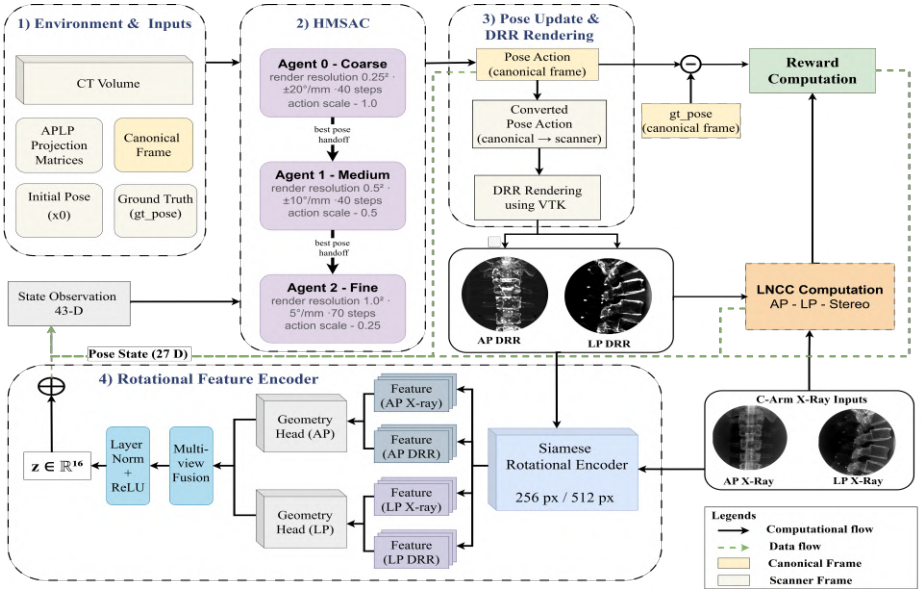


Fig. 3: Hierarchical multi-agent SAC framework.

a new DRR via a VTK-based renderer, which is compared against the X-rays to compute the next state and reward, with the policy optimized to maximize expected cumulative discounted reward.

Hierarchical Registration Framework: We propose a coarse-to-fine Hierarchical Multi-Agent Soft Actor-Critic (HMSAC) framework with three sequential agents at increasing image resolutions. Agent 0 performs coarse registration at 0.25× resolution with a $\pm 20°/\text{mm}$ capture range, Agent 1 refines at 0.5× within $\pm 10°/\text{mm}$, and Agent 2 conducts fine registration at full 1.0× resolution within $\pm 5°/\text{mm}$. Each agent initializes from the best pose estimated by its predecessor and predicts incremental 6-DoF corrections, with progressively reduced action scales to prevent overshooting during fine refinement.

State Representation and Reward: Each agent observes a 43-dimensional state vector comprising the normalized 6-DoF pose in the CCS, multi-scale stereo LNCC similarity scores, optimization progress, visual features from phase correlation and edge alignment, and the 16-dimensional frozen rotational feature embedding z produced by the RFE. Stereo LNCC is computed by concatenating the AP, LP X-ray and DRR pairs and evaluating LNCC within the vertebral ROI, preventing a good AP alignment from masking a poor LP alignment.

$$r_t = \lambda_P \cdot \Delta d_{gt} + \lambda_L \cdot \Delta S_{\text{stereo}} + r_{\text{terminal}} - r_{\text{penalty}}, \tag{2}$$

The reward equation (2) uses $\lambda_P = 10$ and $\lambda_L = 20$, where Δd_{gt} and ΔS_{stereo} measure the reduction in pose error and improvement in stereo LNCC, respectively. The terminal reward r_{terminal} encourages rapid convergence, while r_{penalty} discourages invalid or unstable pose updates.

3 Results and Discussion

Implementation and Datasets: Our in-house spine dataset comprises 59 AP, LP X-ray scenarios with corresponding projection matrices and CT volumes, collected from three sawbone phantoms (one closed & two open phantom). Images were acquired using different C-Arm systems varying in X-ray intensity, orientation, and source-to-patient distance. The dataset was split into 40, 10, and 9 samples for training, validation, and testing, respectively. The 40/10/9 train/validation/test split was performed at the scenario level, with all three phantoms represented in each partition (Open Phantom: 33/8/7 ;Closed Phantom: 7/2/2), ensuring exposure to diverse C-arm imaging conditions during training and evaluation. To confirm clinical reliability, SpineSAC was evaluated on a held-out cadaver dataset, yielding a TRE of 2.31 mm. The model is trained for 100 epochs on a single NVIDIA L40S GPU (48 GB VRAM), taking 13-15 hr. Training uses SAC [7] with $\alpha = 1 \times 10^{-4}$, batch size $B = 512$, $\gamma = 0.99$, $\tau = 0.005$, and four gradient updates per environment step; entropy targets were set to -1.5, -3.0, and -6.0 for A_0, A_1, and A_2 to progressively reduce exploration, following a curriculum allocating 20%, 30%, 50% of epochs to A_0, $A_0\&A_1$, and joint optimization of $A_0\&A_1\&A_2$ with sequential warm-starts, while each agent maintained an independent replay buffer of 60,000-80,000 transitions.

Table 1: Intra-architecture comparison and Ablation TRE (mm, mean ± std).

Method	**Phantom TRE ± std (↓)**		**LNCC (↑)**			**Global NCC (↑)**		**Gradient NCC (↑)**	
			AP	**LP**	**Stereo**	**AP**	**LP**	**AP**	**LP**
Inter-Architecture Comparison									
TD3	Closed	5.76 ± 0.8	0.286	0.294	0.290	0.520	0.550	0.430	0.470
	Open	5.84 ± 0.8	0.291	0.320	0.306	0.510	0.540	0.420	0.460
DQN	Closed	7.1 ± 0.6	0.245	0.253	0.249	0.450	0.480	0.380	0.410
	Open	7.3 ± 0.6	0.238	0.249	0.244	0.430	0.460	0.360	0.390
PPO	Closed	6.6 ± 0.4	0.274	0.297	0.265	0.513	0.500	0.410	0.450
	Open	6.9 ± 0.4	0.293	0.294	0.289	0.514	0.507	0.400	0.440
SpineSAC	Closed	$\mathbf{1.4 \pm 0.5}$	**0.582**	**0.571**	**0.577**	**0.881**	**0.961**	**0.648**	**0.784**
	Open	$\mathbf{1.5 \pm 0.5}$	**0.574**	**0.553**	**0.563**	**0.921**	**0.921**	**0.756**	**0.770**
Intra-Ablation Study									
Base SAC	Closed	4.9 ± 0.5	0.356	0.392	0.374	0.610	0.650	0.520	0.560
	Open	5.2 ± 0.4	0.321	0.358	0.340	0.590	0.620	0.500	0.530
HMSAC	Closed	2.74 ± 0.4	0.489	0.510	0.500	0.770	0.790	0.560	0.700
	Open	2.89 ± 0.4	0.461	0.487	0.474	0.750	0.770	0.590	0.640
HMSAC + CCS	Closed	1.97 ± 0.5	0.486	0.469	0.477	0.780	0.805	0.600	0.690
	Open	2.12 ± 0.5	0.451	0.455	0.453	0.750	0.780	0.580	0.680
HMSAC + RFE	Closed	$\mathbf{1.68 \pm 0.5}$	**0.524**	**0.546**	**0.535**	**0.830**	**0.850**	**0.630**	**0.710**
	Open	$\mathbf{1.73 \pm 0.5}$	**0.516**	**0.537**	**0.527**	**0.810**	**0.840**	**0.650**	**0.720**

Comparison and Ablation Analysis: Table 1 reports results across two evaluations. In the inter-architecture comparison, Twin Delayed Deep Deterministic Policy Gradient (TD3) [3], Proximal Policy Optimization (PPO) [19], and Deep Q-Network (DQN) [12] produce weak results, with TRE of 5.76 ± 0.8 mm,

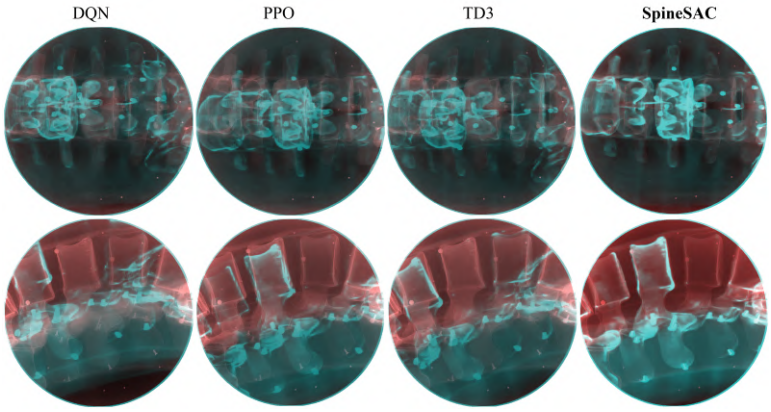


Fig. 4: Blend view comparison of SpineSAC and other methods.

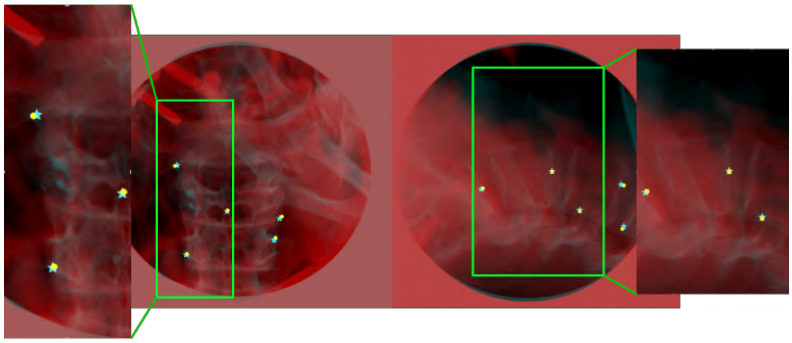


Fig. 5: Blended DRR–X-ray overlay with predicted landmarks on cadaver dataset.

6.6 ± 0.4 mm, and 7.1 ± 0.6 mm respectively and low similarity scores, confirming that deterministic, on-policy clipped, and value-based discrete methods all struggle with the continuous 6-DoF registration task. The proposed SpineSAC achieves a TRE of 1.4 ± 0.5 mm with substantially higher similarity scores across all metrics, demonstrating improved RL performance for this task. The ablation study validates each proposed component. The base SAC model yields a TRE of 4.9 ± 0.5 mm. Adding the hierarchical multi-agent framework (HMSAC) reduces this to 2.74 ± 0.4 mm. Incorporating the Canonical Coordinate System (HMSAC + CCS) further improves accuracy to 1.97 ± 0.5 mm by providing a consistent anatomical reference frame. Adding the Rotational Feature Encoder (HMSAC + RFE) achieves the best result of 1.68 ± 0.5 mm, confirming that each component contributes meaningfully to registration accuracy. The views of the X-ray-DRR blend in Fig. 4 demonstrate the superior anatomical alignment achieved by SpineSAC compared to DQN, TD3 and PPO.

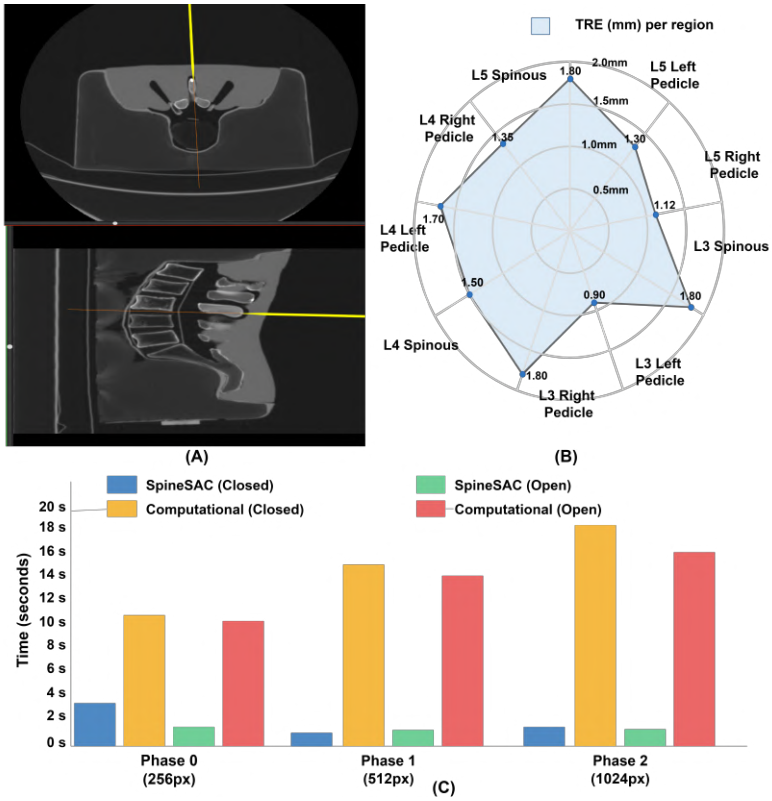


Fig. 6: (A) Intraoperative tracking view with registered CT volume in axial and sagittal views. (B) TRE radar plot for L3, L4, and L5 vertebrae. (C) Computational and SpineSAC time comparison across three resolution levels.

Fig. 5 presents the clinical evaluation on a cadaver dataset showing the blended DRR, X-ray overlay for AP (left) and LP (right) views. Star markers (★) indicate landmarks projected from the fixed X-ray and dot markers (•) indicate landmarks projected from the predicted pose, where the spatial offset between corresponding pairs reflects the residual TRE, which ranges from 2.0–2.5 mm on the cadaver dataset.

Fig. 6(A) shows the intraoperative tracking view with the registration result obtained using SpineSAC, while Fig. 6(B) shows the per-region TRE radar plot for L3, L4, and L5 vertebrae across the right pedicle, left pedicle, and spinous process, where TRE values remain below 1.8 mm across all anatomical landmarks. Fig. 6(C) compares the per-phase inference time across three resolution levels (0.25×, 0.5×, 1.0×) between the computational approach [22] and SpineSAC, demonstrating that our SpineSAC achieves registration in 5 s compared to 45 s for the computational method, a 9× speedup.

4 Conclusion

We present SpineSAC, a hierarchical multi-agent SAC framework for intraoperative 2D-3D spine registration integrating a PCA-based CCS and RFE, achieving a TRE of 1.45 mm on phantoms and 2.31 mm on cadaver data with a 9× inference speedup over the computational baseline. Ablation studies confirm CCS and RFE together reduce TRE from 4.9 mm to 1.4 mm. Future work targets TRE below 1.5 mm on cadaver and below 1.0 mm on phantoms through enhanced RL exploration across broader clinical datasets.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: Proceedings of the International Conference on Machine Learning (ICML). pp. 1597–1607. PMLR (2020)
2. Cui, Y., Meng, M.Q.H., Min, Z.: Weakly-Supervised 2D/3D Image Registration via Differentiable X-ray Rendering and ROI Segmentation, p. 648–658. Springer Nature Switzerland (September 2025). https://doi.org/10.1007/978-3-032-04984-1_62
3. Fujimoto, S., Hoof, H., Meger, D.: Addressing function approximation error in actor-critic methods. In: International conference on machine learning. pp. 1587–1596. Pmlr (2018)
4. Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 1440–1448 (2015). https://doi.org/10.1109/ICCV.2015.169
5. Gopalakrishnan, V., Dey, N., Golland, P.: Intraoperative 2d/3d image registration via differentiable x-ray rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11662–11672 (06 2024). https://doi.org/10.1109/CVPR52733.2024.01108
6. Gupta, D.K., Arya, D., Gavves, E.: Rotation equivariant siamese networks for tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12362–12371 (June 2021). https://doi.org/10.1109/CVPR46437.2021.01218
7. Haarnoja, T., Zhou, A., Abbeel, P., Levine, S.: Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In: Proceedings of the 35th International Conference on Machine Learning (ICML). pp. 1861–1870. PMLR (2018)
8. Jaganathan, S., Wang, J., Borsdorf, A., Shetty, K., Maier, A.: Deep iterative 2d/3d registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 383–392. Springer (2021)
9. Lang, T.F., Keyak, M.J., et al.: A hierarchical 3d segmentation method and the definition of vertebral body coordinate systems for qct of the lumbar spine. Medical Image Analysis **10**(4), 510–520 (2006). https://doi.org/10.1016/j.media.2006.05.005

10. Levine, L., Levine, M.: Drrgenerator: A three-dimensional slicer extension for the rapid and easy development of digitally reconstructed radiographs. Journal of Clinical Imaging Science **10**, 69 (oct 2020). https://doi.org/https://doi.org/10.25259/jcis_105_2020
11. Maik, V., Purayath, A., R, D., Lakshmanan, M., Sivaprakasam, M.: Gui-based pedicle screw planning on fluoroscopic images utilizing vertebral segmentation. In: 2024 IEEE International Symposium on Medical Measurements and Applications (MeMeA). pp. 1–6 (2024). https://doi.org/10.1109/MeMeA60663.2024.10596889
12. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A., Veness, J., Bellemare, M., Graves, A., Riedmiller, M., Fidjeland, A., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D.: Human-level control through deep reinforcement learning. Nature **518**, 529–33 (02 2015). https://doi.org/10.1038/nature14236
13. Neeraja, R., Devadharshiniinst, S., Venkateswaran, N., Maik, V., Purayath, A., Lakshmanan, M., Sivaprakasam, M.: Spinestyle: Conceptualizing style transfer for image-guided spine surgery on radiographs. In: International Workshop on Machine Learning in Medical Imaging. pp. 372–381. Springer (2024)
14. Purayath, A., Maik, V., Chakkaravarthy, A., Lakshmanan, M., Sivaprakasam, M.: A novel spatio2-frequency blob detection algorithm for enhancing precision in image guided surgery. In: 2024 IEEE International Symposium on Medical Measurements and Applications (MeMeA). pp. 1–6 (2024). https://doi.org/10.1109/MeMeA60663.2024.10596832
15. R, D., B, D., S, P.H., Chakkaravarthy, A., Purayath, A., Maik, V., Lakshmanan, M., Sivaprakasam, M.: Enhancing precision in image-guided spine surgery through the prediction of occluded fiducials utilizing resnet architecture. In: ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 1–5 (2025). https://doi.org/10.1109/ICASSP49660.2025.10888585
16. Ren, Y., Wu, S., Wang, M., Cen, Z.: Study on construction of a medical x-ray direct digital radiography system and hybrid preprocessing methods. Computational and Mathematical Methods in Medicine **2014**, 495729 (2014). https://doi.org/10.1155/2014/495729
17. Roth, T., Carrillo, F., Wieczorek, M., Ceschi, G., Esfandiari, H., Sutter, R., Vlachopoulos, L., Wein, W., Fucentese, S.F., Fürnstahl, P.: Three-dimensional preoperative planning in the weight-bearing state: Validation and clinical evaluation. Insights into Imaging **12**(1), 44 (2021). https://doi.org/10.1186/s13244-021-00994-8
18. Salehi, S.S.M., Khan, S., Erdogmus, D., Gholipour, A.: Real-time deep pose estimation with geodesic loss for image-to-template rigid registration. IEEE Transactions on Medical Imaging **38**(2), 470–481 (2019). https://doi.org/10.1109/TMI.2018.2866442
19. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (07 2017). https://doi.org/10.48550/arXiv.1707.06347
20. Vrtovec, T., Likar, B., Pernuš, F.: Spine-based coordinate system. In: Conference Proceedings of the IEEE Engineering in Medicine and Biology Society. vol. 2005, pp. 5120–5123. IEEE (2005). https://doi.org/10.1109/IEMBS.2005.1615629
21. Weiler, M., Cesa, G.: General e(2)-equivariant steerable cnns. In: Advances in Neural Information Processing Systems (NeurIPS). vol. 32 (2019)

22. Yang, D., Shi, H., Zeng, B., Chen, X.: 2d/3d registration based on biplanar x-ray and ct images for surgical navigation. Computer Methods and Programs in Biomedicine **257**, 108444 (10 2024). https://doi.org/10.1016/j.cmpb.2024.108444
23. Zhang, H., He, L., Cheng, Z., Kadkhodamohammadi, A., Stoyanov, D., Davidson, B.R., Mazomenos, E.B., Clarkson, M.J.: Warm-started reinforcement learning for iterative 3d/2d liver registration. International Journal of Computer Assisted Radiology and Surgery pp. 1–9 (2026)