DARK: Dynamic Graphs based Angle-aware Registration of Knee Ultrasound Point Clouds

Injune Hwang, Stephen Mellon, and S. Jack Tu

Nuffield Department of Orthopaedics, Rheumatology and Musculoskeletal Sciences, University of Oxford, Oxfordshire, UK
injune.hwang@ndorms.ox.ac.uk

Abstract. Real world medical imaging tasks often deviate sharply from the clean, controlled conditions assumed in standard computer vision benchmarks. This gap is particularly evident when registering point clouds of bony anatomy derived via 3D freehand ultrasound. Noise from the transducer tracking methods used, along with image segmentation errors, leads to spatially irregular data, resulting in minimal or no point-wise correspondence between scans. Consequently, many registration methods that perform well on open datasets, such as ModelNet40, struggle to generalise effectively in this medical domain. In this work, we introduce **DARK** (Dynamic Graphs based Angle-aware Registration of Knee Ultrasound Point Clouds), a registration framework purpose-built for this highly challenging setting. DARK integrates dynamic graph-based denoising with a quaternion-based Multi-Layer Perceptron (MLP) head, trained using a geodesic loss defined on the SO(3) rotation group. This design enables robust alignment of sparse, noisy and clinically acquired 3D ultrasound point clouds. Critically, unlike many prior methods that rely on simulated data or strong anatomical priors, DARK is trained and evaluated entirely on 3D freehand ultrasound data. Tested on 32 difficult cases involving knee flexion at varying angles with no point-wise overlap, DARK achieves a mean geodesic loss of 33.8°, substantially outperforming both classical and learning-based baselines. This research highlights the value of applying geometric ideas to medical registration tasks, particularly for challenging modalities like ultrasound.

Keywords: Knee Ultrasound Point Clouds · Geometric Registration · Dynamic Graphs.

1 Introduction

Patellofemoral joint (PFJ) disorders, originating from complex biomechanics, are a major source of pain and instability, affecting up to 20% of the population and up to 30% of patients with total knee arthroplasty (TKA) [6, 13, 15, 25]. Patient outcomes are often poor, mainly because current technologies cannot adequately assess dynamic 3D joint movement. Conventional imaging such as CT and MRI is limited by radiation, cost, and metal artefacts, while skin-marker-based gait analysis is inaccurate for mobile bones like the patella [11, 17, 18, 21]. Single

and biplanar fluoroscopic imaging has become popular for PFJ motion analysis; however, these methods are limited by the patella being obscured by the femoral component in patients with TKA [1, 2]. To address this gap, the Computer Aided Tracking and Motion Analysis with Ultrasound System (CAT&MAUS) is in development [10, 14]. CAT&MAUS leverages 3D freehand ultrasound (US), automatic bone segmentation [9] and point cloud registration to output joint kinematics.

To estimate PFJ motion, we register two ultrasound-derived point clouds of distal femur and the patella acquired at different positions during knee flexion. This process can be hindered by significant noise inherent to ultrasound and soft tissue artefacts in ultrasound-derived point clouds, which, in turn, impairs the accuracy of the registration. Therefore, the first stage of our DARK pipeline is to de-noise the point clouds using our previously developed DG-PPU method. [8].

As in [8,10], each point in our 3D US point cloud represents a point of contact of the ultrasound probe computed via sensor fusion. Even when scanning the same knee, contact points vary between acquisitions. Thus, we frame registration as **aligning point clouds with shared underlying geometry, focusing on pose rather than point-wise correspondence**.

Iterative Close Points (ICP) and Globally Optimal ICP (GO-ICP) have shown strong performance in rigid alignment tasks, particularly when there is strict point-wise correspondence [3, 27]. However, these Mean Squared Error (MSE)-based approaches can underperform in our setting, as our goal is to estimate joint motion by registering point clouds extracted from two independent scans and there is little to no point-wise correspondence, even if both sweeps are performed in the same area. Although not based on an MSE framework, Coherent Point Drift (CPD) based methods are also not suited for our task, as they assume dense point-wise correspondences and smooth, Gaussian-distributed noise [16], which we know do not hold for US scans. This motivates the approach of using a feature-based model to tackle our task. Following the strategy in DG-PPU, we leverage learned geometric features from embedding networks such as PointNet++ and Dynamic Graph CNN (DGCNN) [19, 24] to guide registration. The complete DARK pipeline is described in the following sections.

2 DARK Architecture

Fig.1 illustrates the full DARK pipeline. We begin by applying DG-PPU [8], our denoising module built on a dynamic-graphs framework inspired by DGCNN. DG-PPU is trained to capture the intrinsic geometry of knee anatomy by constructing a $k-$regular graph, connecting each point to its k nearest neighbours according to an estimated geodesic distance along the anatomical manifold.

We empirically fine tuned $k = 20$, likewise to [8], to achieve optimal denoising performance on our data set. The geodesic distances and related geometric quantities (e.g., curvature) are estimated within the DGCNN module via Multi-Layer Perceptron (MLP) that takes pointwise position differences as its input. In essence, the network learns a composition of linear maps and nonlinear

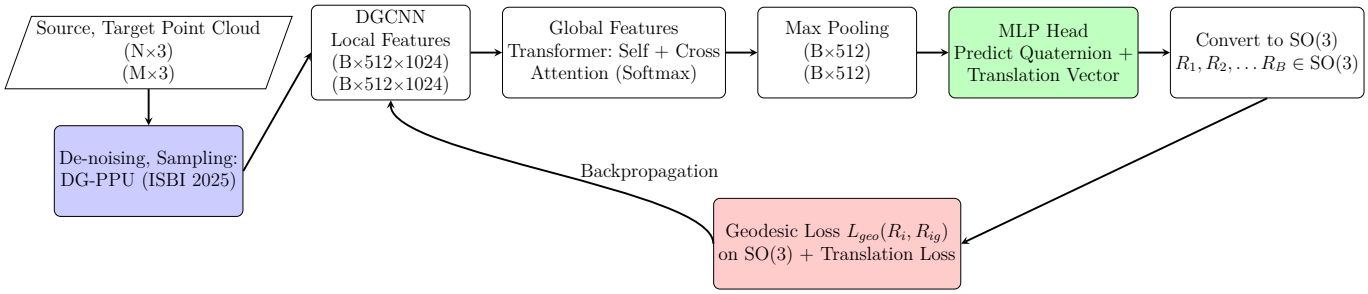

Fig. 1: Overview of the DARK Pipeline. *Dimensions (e.g. ($B \times 512 \times 1024$)) describe the output at each stage. The MLP predicts a quaternion and a translation vector; the quaternion is normalised and converted to a rotation matrix in SO(3) via Quat2Mat. The geodesic loss is computed between predicted ($R_i; 1 \leq i \leq B$) and ground-truth rotations ($R_{ig}; 1 \leq i \leq B$) and is backpropagated to train the weights in the DGCNN, Transformer, MLP stages.*

activation functions whose derivatives approximate the data-driven geometry encoded by the dynamic graphs (see [24] for details). This dynamic graphs approach allows us to remove points that are close in Euclidean space but disconnected anatomically, effectively eliminating false positives. The resulting denoised point clouds contain fewer points than the raw input.

Next, we perform Monte Carlo sampling, following [8], to generate fixed-size batches ($B = 32$, each with 1024 points). These are fed into the feature extraction stage of the pipeline. This strategy enables the model to learn rotation matrices and translation vectors between source and target point clouds of differing sizes, a common scenario in clinical scans taken at varying knee flexion angles.

2.1 Feature Extraction

We extract both local and global features from the denoised, sampled point clouds using DGCNN and Transformer-based modules. While DGCNN is also used in the DG-PPU stage, its roles are distinct: the DGCNN step in DG-PPU is used to learn the anatomical manifold, whereas the second DGCNN re-extracts local features from the cleaned input to support accurate alignment. For global feature extraction, we employ a Transformer architecture. Instead of hard correspondences, we use attention-based soft correspondences, similar in spirit to RPM-Net [28] but without iterative refinement or Singular Value Decomposition (SVD) heads. We apply multi-head self-attention, allowing each point to attend to all others via softmax-weighted scores, and incorporate cross-attention between source and target clouds to model global contextual relationships for alignment.

After feature extraction, we apply global max-pooling across all point-wise features, reducing the per-point feature map (B × 512 × 1024) to a global em-

bedding vector (B $\times$ 512), where 512 is the dimension of features in this embedding. This global representation captures the aggregated contextual information across each point cloud (with 1024 points) in each batch ($B = 32$) and serves as the input to the MLP head. The MLP predicts a quaternion (converted to a rotation matrix) and a translation vector, enabling pose regression between source and target point clouds.

2.2 Pose Regression

As noted earlier, DARK employs an MLP head instead of an SVD-based head. This enables end-to-end differentiability across the DGCNN, Transformer, and MLP modules, facilitating smooth gradient flow throughout training. In contrast, SVD-based heads (like in [23, 28]) impose hard constraints, limiting the flexibility of back propagation. Our MLP head directly predicts quaternions, which are normalised to ensure valid rotations (i.e., orthogonal matrices with determinant 1) and converted to 3×3 rotation matrices via a `Quat2Mat` operation. This approach offers a principled and numerically stable alternative to projecting arbitrary matrices onto SO(3), which often introduces discontinuities.

To evaluate rotation accuracy, we use a geodesic loss on the rotation group SO(3), as in [7]. This loss measures the minimal angle required to align two rotation matrices, or in other words, the shortest angular (geodesic) distance in SO(3). This geodesic distance can be computed via the trace identity:

$$L_{geo}(R_1, R_2) = arccos(\frac{Tr(R_1^T R_2) - 1}{2})$$

This manifold-aware metric offers a more faithful measure of rotation error than Euclidean alternatives like cosine similarity or $||RS - T||$ (where R is the rotation matrix, S, T the source and target), which averages point-wise MSE across the point cloud. Using geodesic loss also avoids issues inherent to Euler angles, such as axis-wise ambiguity and gimbal lock, which make them less reliable for measuring rotation error. We measure Euler angle discrepancies to provide an intuitive sense of alignment, but they are not the loss we train with. Given the absence of point-wise correspondence, translation cannot be perfectly constrained. Thus, our total loss is defined as:

$$L_{geo} + 10^{-7} MSE$$

This formulation prioritises geodesic loss while still accounting for translation.

3 Experiments and Results

3.1 Data Acquisition and Experimental Methods

Our training data (identical to [8]) consists of 2000 point clouds derived from ultrasound scans acquired at four knee positions (P0−P3), where P0 represents full flexion and P3 represents 90-degree extension. The bone surface was segmented

from ultrasound images and the segmented pixels were then transformed into 3D as points using the method described in [9,10]. With the trajectory captured from the tracking method, the transformed points are finally accumulated into point clouds according to the sweeps.

To mitigate bias from the limited number of scans, we applied Monte Carlo sampling for data augmentation. As reported in [8], DG-PPU achieves 98.2% precision in removing false positives from point clouds at positions P1−P3, collected using fewer sweeps but with the same acquisition method and joint configurations as the corresponding scans in the training data set. For evaluation of DARK, we sampled 1500 test point clouds using the same Monte Carlo strategy.

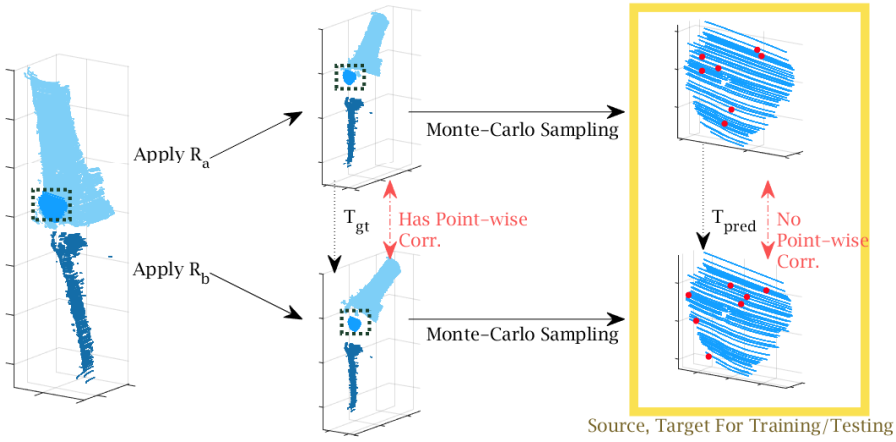

Fig. 2: Our training/testing method. *T_{gt} is the ground truth matrix and T_{pred} the matrix prediction. We boxed the sampled source and target clouds and represented a batch with randomly dotted points (actual size: 1024 points) to emphasise the absence of point-wise correspondence.*

We train our registration model using filtered point clouds produced after the DG-PPU step. During data collection, point clouds were classified into femur, patella, and tibia categories, with corrections applied via DG-PPU.

To deal with non-overlapping inputs, we designed a custom training and testing framework (Fig.2). We would apply R_A, R_B (resp.) to a point cloud in the testing data set to create a point cloud pair from which we independently Monte-Carlo sample 1024 points to create the source and target. This produces source-target pairs with little to no point-wise correspondence but a known alignment matrix (i.e., the ground truth). We would repeat this 32 times with different matrices each time. Training used the Adam optimiser [12] with a learning rate of 0.001. Test data included rotations with angular displacement $|\theta| \in [0°, 135°]$ and translations in the range $[-1, 1]$.

3.2 Results

Table 1 presents the differences in translation, rotation, and geodesic loss, evaluated over a batch of 32 independently Monte Carlo-sampled point cloud pairs. Each source-target pair is generated using a distinct ground truth transformation prior to sampling. We, first, debugged DCP [23], with minor changes in architecture, to apply to our input and observed that the geodesic loss does not converge at all. Hence, we compared the performance of DARK to DCP-variants with significant changes in architecture: DARK trained on cosine similarity loss rather than geodesic loss (referred to as DARK−Geo), DCP trained on geodesic loss with original SVD head (referred to as DCP + Geo). Results of ICP were included as a sanity check.

Table 1. Registration Results *Translation loss t_x, t_y, t_z are the MSE differences in each axis of $||RS - T||$. Rotation loss R_x, R_y, R_z are the differences in Euler angles to T_{gt} per axis. Best results are highlighted in bold.*

Method	R_x	R_y	R_z	t_x	t_y	t_z	L_{geo}
ICP	96.2°	51.8°	89.4°	**0.23**	0.25	**0.22**	150.9°
DCP+Geo	65.7°	41.5°	72.6°	1.00	0.99	0.99	85.4°
DARK−Geo	87.3°	19.8°	60.4°	0.31	0.27	0.26	48.4°
DARK	**30.4°**	**9.17°**	**39.5°**	0.31	**0.23**	0.24	**33.8°**

In terms of runtime, DARK achieves real-time inference speed when tested on an Apple M2 system, further supporting its usage in clinical settings.

4 Evaluation

Comparisons to other Registration methods MSE-based registration methods like ICP tend to minimise translation error while failing to recover accurate rotations, resulting in geodesic losses worse than a random guess. In contrast, DCP + Geo, which leverages learned geometric features, improves rotation and geodesic accuracy but struggles with translation. These results highlight the fundamental challenge of learning to predict rotation and translation without point-wise correspondence.

Most registration methods, classical or learning-based, assume at least partial point-wise correspondence. Even feature-based approaches such as DCP [23] and RPM-Net [28] rely on reasonably aligned inputs to succeed. However, in our clinical setting, we must register 3D point clouds extracted from ultrasound scans acquired at different knee flexion angles, which introduces three major challenges: **(1) no point-wise correspondence, (2) inherently significant ultrasound noise, (3) complexity of PFJ motion**. Many existing registration methods fail outright when applied to our data, often producing NaN values due to invalid

outputs outside the SO(3) rotation group, making geodesic loss computation impossible. Adapting DCP, without any architectural changes like we made for DARK, required significant debugging to achieve convergence. Classical methods like ICP appear entirely broken under these conditions, producing geodesic errors worse than chance.

We selected DCP as our main baseline, given that PointNetLK [5] underperforms DCP in high-variance testing conditions, as shown in experiments on ModelNet40 [26]. While RPM-Net has demonstrated superior performance in scenarios with moderate noise and partial correspondence [28], it is sensitive to initialisation, making it unsuitable (without further adjustments) for our highly sparse, non-overlapping ultrasound point clouds. Also, RPM-Net's architecture is not easily adaptable to our context. Unlike DCP, which uses a DGCNN backbone that aligns naturally with our pipeline, RPM-Net relies on iterative soft-correspondence refinement. Given these potential mismatches, we opted not to include RPM-Net in this initial study. However, incorporating elements from RPM-Net remains an interesting direction perhaps for refinement after coarse alignment using DARK.

SVD vs. MLP The comparison between DCP + Geo and DARK − Geo (ablation study) reveals that much of DARK's performance gains stem from adopting the quaternion-based MLP head, with the geodesic loss also contributing but not working well alongside a SVD head. In contrast, the original DCP paper [23] reported that DCP with a SVD head outperforms DCP with a MLP head on ModelNet40, where there exists point-wise correspondence. This is because SVD directly minimises alignment error and avoids over-fitting to a data set, not having learnable parameters. In contrast, our US-derived point clouds are sparse, noisy and lack any point-wise correspondence, where assumptions of the SVD head break down. Hence, a learnable MLP head becomes advantageous.

Results of DARK Despite the difficulty of the task, DARK achieves a geodesic loss of 33.8°. It's important to clarify that this does not imply each anatomical axis is off by 33.8°; rather, this value represents the total angular displacement between rotation matrices on SO(3). DARK outperforms other baselines significantly in terms of geodesic loss and rotation angle discrepancies. Regarding translation, even perfectly aligned point clouds may yield non-zero $||RS - T||$ due to the absence of point-wise correspondence. Thus, translation loss alone is not a reliable metric for accurate registration. The ground truth translation loss may be larger/smaller than what we achieve for ICP. Still, DARK's predictions are clearly more reasonable than those of other DCP-variants, and only marginally off ICP's results. We visualise DARK's performance in Fig. 3, where ground truth and predicted rotations are applied to a CAD mesh model of the knee, allowing comparison of alignment results in real anatomical scenarios.

We emphasise that if registration were performed on clean geometries, such as those from the CAD models in Fig.3, methods like DARK, DCP + Geo, and ICP would all yield results much closer to the ground truth. Fig.3. is intended to illustrate how registration might appear in a clinical setting; our testing data is inherently sparse and irregular, making it difficult to visualise. Samples 3, 7, and

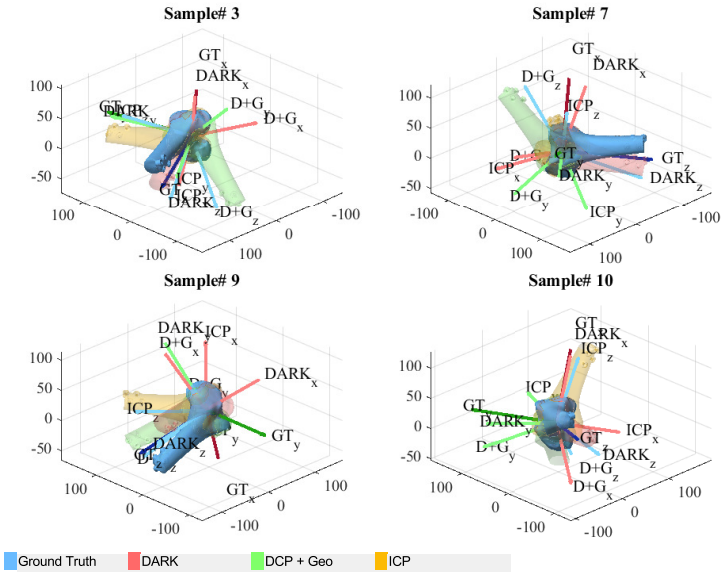

Fig. 3: Registration results applied to a CAD model. *Sample 3,7,10 demonstrate that DARK (pink) works as a good coarse-alignment to the ground truth (blue). Sample 9 is when DARK had the largest geodesic loss across the test data set.*

10 demonstrate that DARK yields significantly improved registration outcomes compared to baseline methods.

In sample 9, while DARK's alignment with the ground truth is not sufficient for clinical deployment in this instance, it still surpasses traditional methods under similar failure conditions (Sample 3, 7, 10). This is crucial for real-world applications, where ground truth transformations are unavailable and no corrective feedback can be used. The fact that DARK maintains a consistently lower error range, even in failure cases, supports its robustness and reliability.

5 Conclusion and Further Discussions

In this work, we introduced DARK, a novel feature-based registration framework tailored for aligning sparse and noisy 3D point clouds derived from knee ultrasound (US) scans. DARK achieves a mean geodesic loss of 33.8° across 32 challenging test cases—an impressive result given the complexity of patellofemoral joint (PFJ) motion, the inherent noise and sparsity of US data, and the lack of point-wise correspondence.

To contextualise this performance, we benchmarked DARK against pre-existing registration methods including ICP and DCP + Geo, both of which are commonly used in rigid alignment tasks. Notably, even DCP-style models, designed to handle partial overlap, struggled under our setting, often failing to con-

verge. DARK significantly outperformed these baselines. For reference, prior work (e.g., [7]) reports $10° - 20°$ geodesic errors in more structured settings, often with stronger supervision and less noisy modalities. We believe that DARK achieving 33.8° in ultrasound-based PFJ motion tracking without point-wise correspondence rivals these accomplishments.

The successful performance of the DARK pipeline marks a substantial advancement in 3D freehand ultrasound registration. By outperforming existing methods on our test data, it holds significant promise for integration into the CAT&MAUS pipeline [10]. A robust and reliable registration method of this nature is the key to unlocking the potential of 3D freehand ultrasound for the dynamic assessment of complex joints like the PFJ.

Future work will concentrate on enhancing the convergence of DARK towards clinically acceptable performance. One direction involves incorporating post-alignment modules, such as DCP + ICP in [23], to boost registration accuracy, potentially using local geometric features. As in [8], a practical limitation we've had whilst training the model was the scarcity of high-quality, annotated US data. While our Monte Carlo sampling helps address this, larger real-world datasets would likely enhance anatomical precision and generalisation. Inspired by [22], we also plan to investigate incorporating a group equivariant framework, adapting principles from [4, 20]. Specifically, our aim is to develop a feature extraction pipeline that respects the symmetry properties of SO(3), which DARK does not enforce, as this is an important avenue for improving model robustness.

Acknowledgments and Disclosure of Interests This study has been delivered through the National Institute for Health and Care Research (NIHR) Oxford Biomedical Research Centre (BRC). The views expressed are those of the author(s) and not necessarily those of the NIHR. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Asano, T., Akagi, M., Nakamura, T.: The functional flexion-extension axis of the knee corresponds to the surgical epicondylar axis: in vivo analysis using a biplanar image-matching technique. Journal of Arthroplasty **20**(8), 1060–1067 (2005). https://doi.org/10.1016/j.arth.2004.08.005
2. Banks, S., Markovich, G., Hodge, W.: The mechanics of knee replacements during gait. in vivo fluoroscopic analysis of two designs. American Journal of Knee Surgery **10**(4), 261–267 (1997)
3. Besl, P., McKay, N.: A method for registration of 3-d shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence **14**(2) (1992). https://doi.org/10.1109/34.121791
4. Fuchs, F.B., Worrall, D.E., Fischer, V., Welling, M.: Se(3)-transformers: 3d roto-translation equivariant attention networks. CoRR **abs/2006.10503** (2020), https://arxiv.org/abs/2006.10503
5. Goforth, H., Aoki, Y., Rangaprasad, A.S., Lucey, S.: Pointnetlk: Robust efficient point cloud registration using pointnet. In: Proceedings of (CVPR) Computer Vi-

sion and Pattern Recognition. pp. 7156 – 7165 (2019). https://doi.org/10.1109/CVPR.2019.00733
6. Harris, M., Edwards, S., Rio, E., Cook, J., Cencini, S., Hannington, M.C., et al.: Nearly 40% of adolescent athletes report anterior knee pain regardless of maturation status, age, sex or sport played. Physical Therapy in Sport **51**, 29–35 (2021). https://doi.org/10.1016/j.ptsp.2021.06.005
7. Hou, B., Miolane, N., Khanal, B., Lee, M.C.H., Alansary, A., McDonagh, S.G., Hajnal, J.V., Rueckert, D., Glocker, B., Kainz, B.: Computing CNN loss and gradients for pose estimation with riemannian geometry. CoRR **abs/1805.01026** (2018), http://arxiv.org/abs/1805.01026
8. Hwang, I., Sarvanan, K., Coralli, C.V., Tu, S.J., Mellon, S.J.: Dg-ppu: Dynamical graphs based post-processing of point clouds extracted from knee ultrasounds. 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI) pp. 1–5 (2025). https://doi.org/10.1109/ISBI60581.2025.10980873
9. Jia, R., Mellon, S.J., Hansjee, S., Monk, A.P., Murray, D.W., Noble, J.A.: Automatic bone segmentation in ultrasound images using local phase features and dynamic programming. In: 2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI). pp. 1005–1008 (2016). https://doi.org/10.1109/ISBI.2016.7493435
10. Jia, R., Monk, P., Murray, D., Noble, J.A., Mellon, S.: CAT& MAUS: A novel system for true dynamic motion measurement of underlying bony structures with compensation for soft tissue movement. Journal of Biomechanics **62**, 156–164 (2017). https://doi.org/10.1016/j.jbiomech.2017.04.015
11. Khoury, V., Cardinal, Bureau, N.J.: Musculoskeletal sonography: A dynamic tool for usual and unusual disorders. American Journal of Roentgenology **188**(1), W63–W73 (2007). https://doi.org/10.2214/AJR.06.0579
12. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014). https://doi.org/10.48550/arXiv.1412.6980
13. Laubach, M., Hellmann, J., Dirrichs, T., Gatz, M., Quack, V., Tingart, M., Betsch, M.: Anterior knee pain after total knee arthroplasty: A multifactorial analysis. Journal of Orthopaedic Surgery **28**(2), 2309499020918947 (2020). https://doi.org/10.1177/2309499020918947
14. Monk, A.P., Chen, M., Mellon, S., Gibbons, C.L.M.H., Beard, D.J., Gill, H.S., Murray, D.W.: Measurement of in-vivo patella kinematics using motion analysis and ultrasound (maus). In: 2013 IEEE International Symposium on Medical Measurements and Applications (MeMeA). pp. 257 – 260. IEEE (2013). https://doi.org/10.1109/MeMeA.2013.6549747
15. Myer, G.D., Ford, K.R., Stasi, S.L.D., Foss, K.D.B., Micheli, L.J., Hewett, T.E.: High knee abduction moments are common risk factors for patellofemoral pain (PFP) and anterior cruciate ligament (ACL) injury in girls: Is PFP itself a predictor for subsequent ACL injury? British Journal of Sports Medicine **49**(2), 118–122 (2015). https://doi.org/10.1136/bjsports-2013-092536
16. Myronenko, A., Song, X.: Point set registration: Coherent point drift. IEEE Transactions on Pattern Analysis and Machine Intelligence **32**(12), 2262–2275 (2010). https://doi.org/10.1109/TPAMI.2010.46
17. Nazarian, L.N.: The top 10 reasons musculoskeletal sonography is an important complementary or alternative technique to mri. American Journal of Roentgenology **190**(6), 1621 – 1626 (2008). https://doi.org/10.2214/AJR.07.3385
18. Neustadter, J., Raikin, S.M., Nazarian, L.N.: Dynamic sonographic evaluation of peroneal tendon subluxation. American Journal of Roentgenology **183**(4), 985–988 (2004). https://doi.org/10.2214/ajr.183.4.1830985

19. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: Advances in Neural Information Processing Systems. pp. 5099–5108. NeurIPS (2017). https://doi.org/10.48550/arXiv.1706.02413
20. Satorras, V.G., Hoogeboom, E., Welling, M.: E(n) equivariant graph neural networks. CoRR **abs/2102.09844** (2021), https://arxiv.org/abs/2102.09844
21. Smet, A.A.D., Winter, T.C., Best, T.M., Bernhardt, D.T.: Dynamic sonography with valgus stress to assess elbow ulnar collateral ligament injury in baseball pitchers. Skeletal Radiology **31**(11), 671–676 (2002). https://doi.org/10.1007/s00256-002-0558-0
22. Wang, H., Liu, Y., Hu, Q., Wang, B., Chen, J., Dong, Z., Guo, Y., Wang, W., Yang, B.: Roreg: Pairwise point cloud registration with oriented descriptors and local rotations. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(8), 10376–10393 (2023). https://doi.org/10.1109/TPAMI.2023.3244951
23. Wang, Y., Solomon, J.M.: Deep closest point: Learning representations for point cloud registration. International Conference on Computer Vision (2019). https://doi.org/10.1109/iccv.2019.00362
24. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph CNN for learning on point clouds. CoRR **abs/1801.07829** (2018), http://arxiv.org/abs/1801.07829
25. Witvrouw, E., Lysens, R., Bellemans, J., Cambier, D., Vanderstraeten, G.: Intrinsic risk factors for the development of anterior knee pain in an athletic population: A two-year prospective study. The American Journal of Sports Medicine **28**(4), 480–489 (2000). https://doi.org/10.1177/03635465000280040701
26. Wu, Z., Song, S., Khosla, A., et al.: 3d shapenets: A deep representation for volumetric shapes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1912–1920 (2015). https://doi.org/10.1109/CVPR.2015.7298801
27. Yang, J., Li, H., Jia, Y.: Go-icp: Solving 3d registration efficiently and globally optimally. IEEE International Conference on Computer Vision (2013). https://doi.org/10.1109/ICCV.2013.184
28. Yew, Z.Y., Lee, G.H.: Rpm-net: Robust point matching using learned features. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11824–11833 (2020). https://doi.org/10.1109/CVPR42600.2020.01184