DG-PPU: DYNAMICAL GRAPHS BASED POST-PROCESSING OF POINT CLOUDS EXTRACTED FROM KNEE ULTRASOUNDS

Injune Hwang [*] *Karthik Saravanan* [*] *Caterina Vanelli Coralli* [*] *S. Jack Tu* [†] *Stephen J. Mellon*[†]

[*]University of Oxford
[†]NDORMs, University of Oxford

ABSTRACT

Patellofemoral joint (PFJ) pain affects one in four people, with one in five experiencing chronic knee pain despite treatment. Incorrect patellar tracking after arthroplasty may contribute to poor outcomes and ongoing pain. Traditional imaging methods such as CT and MRI have limitations when it comes to visualising PFJ motion. Our goal is to improve the visualisation of patellar tracking and PFJ motion by utilising 3D registration of point clouds obtained from freehand ultrasound scans taken at various flexion angles. Soft tissues are often misidentified as bone during segmentation, leading to noisy 3D point clouds that hinder accurate registration of the bony joint anatomy. Utilising machine learning to analyse the intrinsic geometry of the knee may help eliminate these false positives, as the geometry of the knee remains consistent during PFJ motion. Our dynamical graphs-based post-processing of ultrasound (DG-PPU) algorithm effectively generates smoother point clouds that accurately represent the bony knee anatomy at various joint flexion angles. Point clouds were converted back to 2D and visually evaluated against the original ultrasound images. DG-PPU outperformed manual data cleaning performed by author CVC, achieving a precision of 98.2% in deleting false positives and noise across three different angles of joint flexion. DG-PPU is the first algorithm specifically developed to directly clean 3D point clouds generated from ultrasound scans, bypassing traditional 2D cleaning methods. Hence, it facilitates the development of a novel assessment system for patellar mal-tracking, which currently lacks a viable solution.

***Index Terms*—** Ultrasound, 3D Point Clouds, Deep Learning, Graphs, Post-processing

1. INTRODUCTION

Up to 30% of patients experience *non-specific anterior knee pain* after TKA [1]. The exact mechanisms are unclear, but abnormal patellofemoral instability with the implant is a suggested cause. While previous studies used static imaging methods like CT and MRI, factors such as limited field view, radiation exposure from CT, and MRI costs have increased the use of ultrasound for dynamic MSK imaging [2–5].

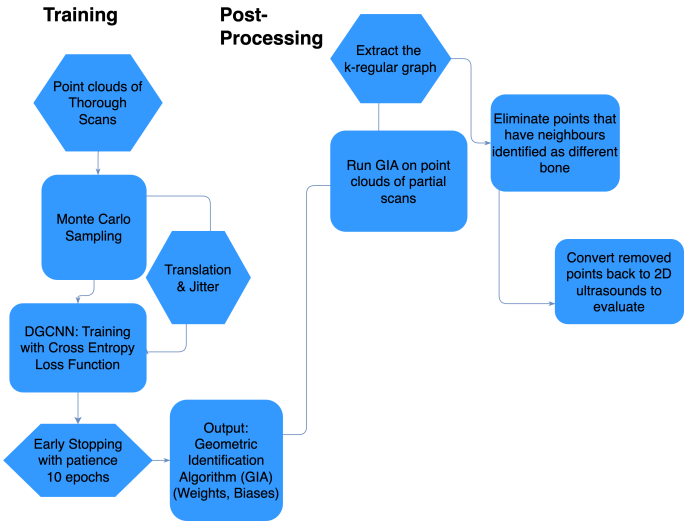

Fig. 1: DG-PPU

Computer-Aided Tracking and Motion Analysis with Ultrasound System (CATMAUS) [6, 7] was developed to visualise patellofemoral interactions. Recent work highlights the capability to reconstruct 3D point clouds from 2D freehand ultrasound scans using motion tracking. This was accomplished by combining semantic segmentation and point cloud registration with the integration of various tracking methods [8].

Ultrasound data is inherently noisy, and point clouds extracted from 2D ultrasound images often contain false positives caused by soft tissue. In the past, manual cleaning and de-noising were necessary to eliminate these erroneous points before reconstructing a 3D mesh of the joint. Manual data cleaning is impractical for real-time visualisation of PFJ motion in clinical settings. Additionally, noise in the point clouds interferes with the Iterative Closest Points (ICP) algorithm [9], complicating the rigid transformations necessary for effective point cloud registration [8]. Accurate registration is essential for visualising 3D PFJ motion.

Detailed ultrasound scans are a necessity when reconstructing a realistic point cloud representation of the PFJ. To collect such data, it requires very thorough, time-consuming

scanning of the joint that would not be feasible within the time frame of a standard clinic appointment. We use the term *thorough scans* to describe the freehand ultrasound scans we achieve in the lab and *partial scans* for a less detailed scan one may expect to be able to use in a clinical setting. We aim to achieve accurate visualisation of PFJ motion via registration of point clouds extracted from partial scans across different angles of the knee joint.

With the long term goal of obtaining point cloud registration, this project focuses on de-noising and filtering the 3D point clouds extracted from the ultrasound images by CATMAUS. For this purpose, we've designed our novel algorithm: Dynamical Graph-based Post-processing for Ultrasounds (DG-PPU). We elaborate more on the methodology below.

2. DG-PPU

2.1. Dynamical Graphs Approach

Dynamical Graph Convolutional Neural Networks (DGCNN) [10] allow one to achieve classification, segmentation, and surface reconstruction tasks on point cloud data. These are tasks that a standard CNN would struggle with due to the unordered and irregular structure of point clouds [11]. The key idea of this approach is that the graph is updated per layer as we learn more features from the input data. Precisely, DGCNN uses a graph that connects a point to its k-nearest neighbours. As each layer learns more about the geometry, the graph would update the k-nearest neighbours accordingly. By the final epoch of training, the graph is theoretically expected to represent the geometry of the original data depicted by the point clouds (e.g., geometry of the knee represented as 3D point clouds). This is supported by test results achieved by the DGCNN on the ModelNet40 dataset [12].

The primary idea of DG-PPU is to use the dynamical graphs structure that is produced by the DGCNN to post-process our point clouds. The dynamical graphs at the final stage of training would let us understand the distance between two points along the knee joint.

Mislabelling soft tissue as bone, and including them in our point clouds, may occur due to noise accumulated from the ultrasound data itself, or due to close proximity of the ultrasound probe to neighbouring soft tissue. For the former case, the false positive would stand out amongst a cluster of identically labelled points. For the latter, in terms of Euclidean distance, it may not be obvious. However, learning the intrinsic geometry of the knee bones from 3D point clouds allows us to distinguish the tip of two adjacent bones. False positives between two bones cause the estimation of both bones to touch at boundaries—their k-nearest neighbours would therefore not be classified as the same type of bone. This motivated the point elimination criterion of our post-processing algorithm DG-PPU which we describe in Fig. 1. We describe the algorithm in more detail below.

2.2. Monte Carlo Sampling and Batches

Ultrasound scans and the point clouds we extract from them present two challenges. First, the number of points we extract from a thorough scan tend to be very large. Applying the k-regular graph approach to such would be computationally expensive. Second, creating a 3D point cloud reconstruction from a thorough ultrasound scan is time consuming. This motivates the usage of probabilistic models to generate batches of point clouds. We designed the algorithm to generate 500 point clouds each containing 1024 points per ultrasound scan. Batches of 1024 points make the algorithm very computationally efficient when drawing k-regular graphs. Our algorithm samples points with replacement, as we would like to represent the density accurately. As we sample 500 times (hence, 512,000 points with duplicates), even with replacement, we cover most of our point clouds when sampling.

Since our data set has a large imbalance in data points per bone, we implemented translation and jittering [13, 14] to augment the minority class, doubling the number of patellar points in our dataset and minimising the likelihood that they are mislabelled by our DGCNN. We set $k = 20$ for the k-regular graphs as a result of the following calculation. When $k = 20$, it is of 0.9695 probability that all 500 samples have more than 20 patellar points. As patellar points are the minority class, we used them as the criterion. We could set $k > 20$, which would allow the graphs to capture the geometry in more detail. However, this would reduce the likelihood that all samples contain more than k patellar points.

2.3. Training and Post-Processing

After generating 500 point clouds of 1024 points per scan, we implemented DGCNN using the Pytorch Geometric python package [15] with Adam optimizer (learning rate 0.001) [16], and early stopping callbacks monitoring validation loss values with 10 epochs of patience during training [17, 18]. Model performance metrics, including accuracy, validation loss, precision, recall, f1-score, intersection over union (IoU), and interactive 3D segmentation objects, were logged for each epoch during training via the wandb python package [19].

As discussed above, we trained using 20 regular graphs. Each batch creates a feature point that contains geometric information updated using the graph structure per layer and epoch. The algorithm would gather the information from all these feature points to fix the weights and biases of a geometric identification algorithm. The geometric identification algorithm captures the geometric information of a point cloud we utilise to draw a 20 regular graph for a general point cloud. To avoid over-fitting, we trained using scans across four dif-

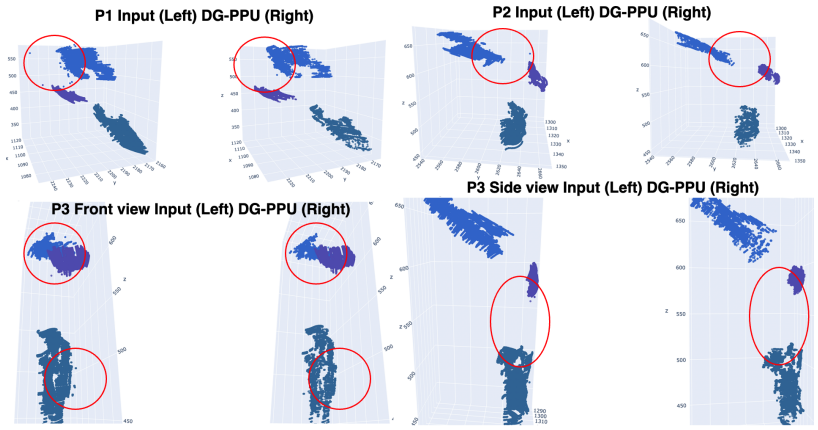

Fig. 2: Post-processing results of DG-PPU on point clouds extracted from partial scans P1, P2, P3—*circled in red are the key areas where DG-PPU outperforms the manual cleaning/de-noising done by our third author, CVC.*

ferent angles of knee flexion, implementing early stopping callbacks as described earlier.

DG-PPU operates according to the geometric identification algorithm above. Our algorithm searches the neighbours of each point in a batch and deletes points that have at least one neighbouring point identified as a different bone to itself. This would eliminate the false positives where soft tissue was mislabelled as bone, improving the accuracy with which we are able to visualise PFJ motion.

3. EXPERIMENTS AND RESULTS

3.1. Data sets and Experimental Setup

Our training data-set consists of 2,000 point clouds generated from thorough ultrasound scans of 4 different positions (i.e., 4 different angles of the knee joint). We denote the 4 different positions as P0, P1, P2, P3 for future reference where P0 is full flexion and P3 full extension. At full flexion, the patellar bone is distinctive; yet, for other positions, the outline of the patella becomes less evident. We tested our post-processing algorithm on 1,500 point clouds generated from partial scans ranging from P1 to P3. Partial scans are sparse scans collected at the same positions as their corresponding thorough scans. We tested on partial scans to assess whether the post-processing algorithm is robust to a lack of detail and consequently has potential clinical use.

3.2. Training and Post-Processing Results

We attach our Github Repository that provides both the training and post-processing algorithm of DG-PPU. Our post-processing algorithm, when applied to point clouds of partial scans, was successful in getting rid of the false positives and noise of our original 3D point cloud data. We can see this in Fig. 2 where we compare side by side the input data and DG-PPU outcome for positions: P1, P2, P3. At P1, we can see that DG-PPU effectively gets rid of points floating between the cluster of patellar and femoral points. Across all positions, we can also see that DG-PPU is effective in making the point clouds smoother. Note that point clouds of partial scans have not been previously seen by our model. This demonstrates that DG-PPU works for non-complete point clouds and has potential use in clinical settings where thorough scans are not feasible due to time constraints.

We would like to remark that DG-PPU provides an automated post-processing procedure that trains on 3D point clouds extracted from freehand ultrasound scanning and removes false positives in real-time; this has not been done before. DG-PPU has also got rid of false positives for data sets that it has not trained on. Yet, due to the currently limited amount of data available, we are uncertain that this will translate across other knee ultrasound images. Further testing with more scans is required to determine whether the post-processing algorithm, without additional training, can be used in clinical settings.

4. EVALUATION

DG-PPU produced smoother point clouds that align more with the 3D anatomy of the knee joint. Yet, the following question remains: *is the algorithm actually deleting false positives?* As we do not have a ground truth for 3D points, it is challenging to determine what the points we've delete using DG-PPU represent. Medical specialists, however, can look at the original 2D ultrasound images and assess whether the part of an image that DG-PPU has deleted contradicts our knowledge of bone anatomy. Hence, we did an inverse operation to convert the 3D points back to their corresponding ultrasound frames.

After comparing the frames, we found many examples (Fig. 3) where DG-PPU removed false positives that manual cleaning carried out by CVC had missed. From the in-

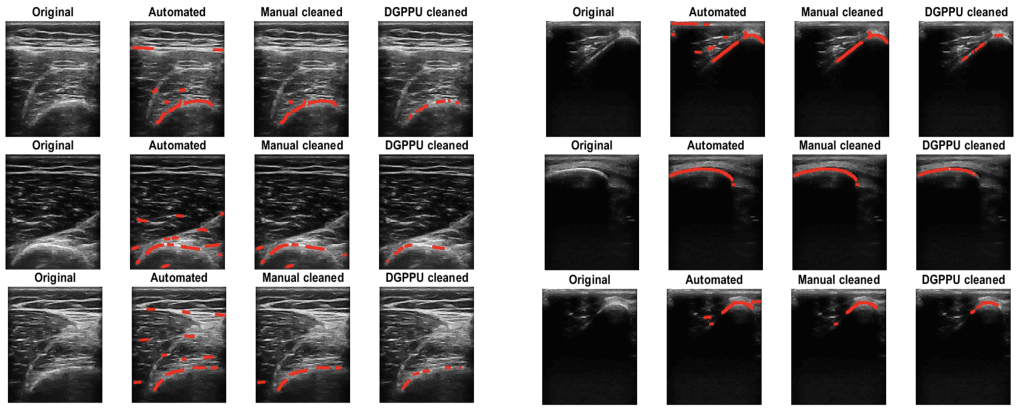

Fig. 3: DG-PPU compared to medical specialist (manual cleaned) when inverted back to ultrasound scans—*lines dotted in red are the parts of the freehand ultrasound scan included in each 3D point cloud. The images labelled* ***Automated*** *are the input data and the ones labelled as* ***Manual Cleaned*** *are data our third author generated for a different project. These images were taken across P1, P2, and P3, representing femoral, patellar, and tibial scans.*

verted ultrasound scans, we also calculated DG-PPU's precision. Precision was measured by excluding frames where DG-PPU either deleted a whole line labelled as bone or left only a single point representing the bone, making it impossible to reconstruct from the output of that frame. Across P1, P2, P3, DG-PPU had precision of 98.4% (183 out of 186), 94% (194 out of 200) and 99.2% (247 out of 249) respectively. Mean precision was 98.2% overall.

It is worthy to highlight that although evaluation was done on 2D ultrasound scans, DG-PPU's primary goal is to post-process the 3D point clouds we extract from the ultrasound scans and aid the CATMAUS point cloud registration pipeline. Evaluating 2D ultrasound scans allow us to point out some mistakes an algorithm can make, which justifies our definition for precision. Segmentation per pixel of 2D ultrasound scans has been studied before using classical segmentation methods like U-net [20]. We would like to emphasise that such methods are designed for segmentation per pixel in 2D ultrasound scans, not for 3D point cloud registration as in the CATMAUS pipeline.

5. CONCLUSION AND FURTHER DISCUSSION

We've established that the dynamical graphs approach allows us to create a novel post-processing algorithm that eliminates false positives and noise of 3D point clouds generated from freehand ultrasound scans. DG-PPU produced smooth point clouds that would be more suitable inputs for point cloud registration algorithms and by evaluating the 2D ultrasound frames, we've established that the points the algorithm deleted did not contradict our knowledge of knee bone anatomy. DG-PPU also runs in real-time including both training and post-processing and is robust across different angles of knee joint flexion. In addition, DG-PPU successfully filters partial scans that the algorithm has not trained on, which justifies that DG-PPU has potential clinical usage. However, the limited amount of 3D ultrasound point cloud data available restricted our ability to do a wider evaluation of DG-PPU. This lack of data makes it challenging to determine if DG-PPU can consistently produce smooth post-processing results as shown in Fig. 2. Testing whether this novel algorithm has immediate clinical application in the assessment of patients would be an interesting avenue to explore.

DG-PPU currently struggles with noisy datasets that contain clusters of false positives larger than k in size. We aim to resolve this issue by implementing graph-based augmentations such as adjusting edge weights or adding random edges to k-regular graphs. This may be proven useful to ensure that DG-PPU does not over-fit due to the lack of currently available data. Introducing higher-order message passing algorithms [21] which utilise simplicial complices rather than graphs may resolve this issue fundamentally.

Since dynamical graphs allow us to learn the intrinsic geometry that is invariant to knee flexion, we believe that we could integrate DG-PPU within the CATMAUS pipeline [7] to achieve 3D point cloud registration. We aim to accomplish this by replacing ICP within the current CATMAUS pipeline with the Deep Closest Point (DCP) algorithm [22] in future work. DCP makes use of the DGCNN to determine the rigid transformations from one point cloud to another., integrating naturally with DG-PPU. Although DG-PPU has corrected a lot of false positives, it may still contain noise for which DCP is robust to compared to its classical counterpart. This further supports our future trajectory of utilising DCP instead of ICP for the point cloud registration of our DG-PPU outputs. If DG-PPU and DCP can successfully achieve point cloud registration in future work, this would enable the development of a novel patellar tracking assessment system with ultrasound, which currently does not occur in clinical practice.

6. COMPLIANCE WITH ETHICAL STANDARDS AND ACKNOWLEDGEMENTS

This study was conducted in accordance with ethical guidelines and approval from the relevant Institutional Review Board (IRB). The authors have no relevant financial or non-financial interests to disclose.

7. REFERENCES

[1] M. Laubach, J. TR Hellmann, et al. Anterior knee pain after total knee arthroplasty: A multifactorial analysis. *Journal of Orthopaedic Surgery*, 28(2):2309499020918947, 2020. PMID: 32338135.

[2] L. N. Nazarian. The top 10 reasons musculoskeletal sonography is an important complementary or alternative technique to mri. *American Journal of Roentgenology*, 190(6):1621–1626, 2008.

[3] V. Khoury, É. Cardinal, and N. J. Bureau. Musculoskeletal sonography: A dynamic tool for usual and unusual disorders. *American Journal of Roentgenology*, 188(1):W63–W73, 2007.

[4] J. Neustadter, S. M. Raikin, and L. N. Nazarian. Dynamic sonographic evaluation of peroneal tendon subluxation. *American Journal of Roentgenology*, 183(4):985–988, 2004.

[5] A. A. De Smet, T. C. Winter, T. M. Best, and D. T. Bernhardt. Dynamic sonography with valgus stress to assess elbow ulnar collateral ligament injury in baseball pitchers. *Skeletal Radiology*, 31(11):671–676, 2002.

[6] A. P. Monk, M. Chen, S. Mellon, et al. Measurement of in-vivo patella kinematics using motion analysis and ultrasound (maus). In *2013 IEEE International Symposium on Medical Measurements and Applications (MeMeA)*, page 257–260. IEEE, 2013.

[7] R. Jia, P. Monk, D. Murray, J. A. Noble, and S. Mellon. 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.

[8] R. Buchanan, S. J. Tu, M. Camurri, S. J. Mellon, and M. Fallon. 3d freehand ultrasound using visual inertial and deep inertial odometry for measuring patellar tracking. *MEMEA*, 10:1–6, 2024.

[9] P.J. Besl and N.D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2), 1992.

[10] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon. Dynamic graph CNN for learning on point clouds. *CoRR*, abs/1801.07829, 2018.

[11] S. A. Bello, S.Yu, C. Wang, et al. Review: Deep learning on 3d point clouds. *Remote Sensing*, 12(11), 2020.

[12] Z. Wu, S. Song, A. Khosla, et al. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015.

[13] H. Li, L. Zhang, et al. Pointaugment: An auto-augmentation framework for point cloud classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6378–6387, 2020.

[14] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems*, pages 5099–5108. NeurIPS, 2017.

[15] M. Fey and J. E. Lenssen. Fast graph representation learning with pytorch geometric. *CoRR*, abs/1903.02428, 2019.

[16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[17] Towards AI Team. Pause for performance: The guide to using early stopping in ml and dl model training. *Towards AI*, 2024.

[18] Various authors. Keeping deep learning models in check: A history-based approach to mitigate overfitting. *arXiv preprint arXiv:2401.10359*, 2024.

[19] L. Biewald. Experiment tracking with weights and biases, 2020. Software available from wandb.com.

[20] M. Dunnhofer et al. Siam-u-net: encoder-decoder siamese network for knee cartilage tracking in ultrasound images. *Medical Image Analysis*, 60(1), 2019.

[21] Mustafa Hajij, Ghada Zamzmi, Theodore Papamarkou, Nina Miolane, Aldo Guzmán-Sáenz, Karthikeyan Natesan Ramamurthy, Tolga Birdal, Tamal K. Dey, Soham Mukherjee, Shreyas N. Samaga, Neal Livesay, Robin Walters, Paul Rosen, and Michael T. Schaub. Topological deep learning: Going beyond graph data. *arXiv preprint arXiv:2206.00606*, 2023.

[22] Y. Wang and J. M. Solomon. Deep closest point: Learning representations for point cloud registration. *International Conference on Computer Vision*, 2019.