



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.

EchoTracker: Advancing Myocardial Point Tracking in Echocardiography

Md Abulkalam Azad^{1,2}[0009-0006-7177-4961], Artem Chernyshov¹, John Nyberg¹, Ingrid Tveten^{1,3}, Lasse Lovstakken¹, Håvard Dalen^{1,2}, Bjørnar Grenne^{1,2}, and Andreas Østvik^{1,2,3}[0000-0003-3895-2683]

¹ Norwegian University of Science and Technology, Trondheim, Norway

² Clinic of Cardiology, St. Olavs Hospital, Trondheim, Norway

³ SINTEF Digital, Trondheim, Norway

{md.a.azad, andreas.ostvik}@ntnu.no

Abstract. Tissue tracking in echocardiography is challenging due to the complex cardiac motion and the inherent nature of ultrasound acquisitions. Although optical flow methods are considered state-of-the-art (SOTA), they struggle with long-range tracking, noise occlusions, and drift throughout the cardiac cycle. Recently, novel learning-based point tracking techniques have been introduced to tackle some of these issues. In this paper, we build upon these techniques and introduce EchoTracker, a two-fold coarse-to-fine model that facilitates the tracking of queried points on a tissue surface across ultrasound image sequences. The architecture contains a preliminary coarse initialization of the trajectories, followed by reinforcement iterations based on fine-grained appearance changes. It is efficient, light, and can run on mid-range GPUs. Experiments demonstrate that the model outperforms SOTA methods, with an average position accuracy of 67% and a median trajectory error of 2.86 pixels. Furthermore, we show a relative improvement of 25% when using our model to calculate the global longitudinal strain (GLS) in a clinical test-retest dataset compared to other methods. This implies that learning-based point tracking can potentially improve performance and yield a higher diagnostic and prognostic value for clinical measurements than current techniques. Our source code is available at: <https://github.com/riponazad/echotracker/>.

Keywords: Deep learning · Point-tracking · Motion estimation · Ultrasound imaging · Strain measurements.

1 Introduction

Myocardial imaging in echocardiography uses ultrasound (US) image analysis to assess and quantify the morphology and function of the cardiac muscle. These methods can be used to uncover reduced pumping efficiency, detect muscle irregularities, and diagnose various heart conditions, facilitating early identification of cardiac dysfunction. Myocardial strain, a measure of deformation, has shown

high sensitivity with superior diagnostic and prognostic value compared to common anatomical measurements, such as ejection fraction [6]. Motion estimation is vital for precise strain, but is hampered by variability in image acquisition, measurements, and inherent limitations of US. Currently, motion estimation and strain imaging rely on speckle tracking using block- and feature-matching approaches [18]. Recent advances in learning-based techniques, such as optical flow-based architectures like FlowNet 2.0 [9] and PWC-Net [17], have inspired researchers to use and adapt them for US imaging and enhanced strain calculations [5, 11, 12]. However, optical flow estimates dense displacement fields between consecutive frames without considering long-range temporal context. This limitation makes tracking susceptible to noise, out-of-plane motion, and decorrelation of speckle patterns. Therefore, it hinders optimal tracking across multiple frames and causes drift within the cardiac cycle.

In this study, we propose EchoTracker, a novel method for tissue tracking in echocardiography. It is designed based on the latest point tracking approaches for general applications [2, 3, 7, 10, 21]. To the best of our knowledge, this is the first work that uses such an approach in the field of medical imaging. It addresses the limitations of dense optical flow by leveraging the temporal context of longer sequences. Herein, we build upon this and design our model architecture tailored for US data, enabling efficient learning with enhanced performance and a lightweight structure. We assess the tracking performance of our model compared to related approaches. Finally, we utilize our model for GLS calculations and compare the clinical performance with other SOTA solutions.

2 Tracking Any Point (TAP)

Tracking of arbitrary points in video sequences is a new research area in deep learning, evolving to mitigate the limitations of optical flow-based tracking. It also possesses the capability to tackle deformations when queried points are selected on the surface of non-rigid objects. Doersch et al. were the first to formalize the TAP problem, provided a benchmark comprising real and synthetic data, and proposed a simple baseline, TAP-Net, for evaluation [2]. The TAP algorithm takes a video and query points from any frame t as input and predicts locations (x_t, y_t) and binary occlusions (o_t) as output for each point in all other frames. Doersch et al. also mentioned that the output location (x_t, y_t) is meaningless when occluded ($o_t = 1$). However, this poses a contradiction to our problem, as we still require the location to compute GLS even if the point is out of the plane. Our modifications to the formal definition of TAP align with Persistent Independent Particles (PIPs) [7]. It aimed to track particles as pixels across long-range video sequences, inspired by the classic Particle Video approach [15].

However, PIPs has notable drawbacks, as it operates on long videos using temporal sliding, which can lead to drift when points are occluded for more than a single window. Expanding the temporal window makes the model slow and unsuitable for parallel computation. To this end, Doersch et al. later introduced TAPIR [3], leveraging both TAP-Net and PIPs. Also in parallel, Zheng et al.

proposed PIPs++ by modifying PIPs with expanded temporal receptive fields and a multi-template update mechanism [21].

Although these approaches differ, a common factor is that they all track points independently, not sharing information between trajectories. This limitation may hinder the tracking of deformable objects like myocardial tissue and lead to drift for points outside the US probe view. This issue has been addressed by CoTracker [10] and OmniMotion [19]. CoTracker iteratively refines trajectories using a transformer architecture in a sliding window manner after a naive initialization. Consequently, it shares the same disadvantage as PIPs for long-range tracking and also exponentially increases the time complexity in case of longer temporal windows. On the other hand, OmniMotion provides a test-time optimization approach based on the canonical 3D volume of the input video.

3 Methods

The overall goal of our learning-based model is to track tissue points through the cardiac cycle while dealing with complex motion, deformation, and noise. As depicted in Fig. 1, the input comprises a sequence of US images $U = \{u_s \in \mathbb{R}^{H \times W}\}$ for $s = 0, 1, \dots, S$, and a set of query points on the first frame $p_0 = \{(x_0^n, y_0^n)\}$ for $n = 0, 1, \dots, N$. Here, H is the height of the image, W is the width, while x and y are the horizontal and vertical pixel locations of a given point. The proposed solution outputs the locations of the queried points in all other frames, $P = \{p_s \in (x_s^n, y_s^n)\}$. Hence, the problem can be summarized as $\text{EchoTracker}(u_s, p_0) = p_s \in (x_s^n, y_s^n)$.




Fig. 1. An illustration of tracking queried points (highlighted in red) from the first frame throughout one heart cycle.

3.1 EchoTracker

Our proposed model, named EchoTracker, includes two stages as shown in Fig. 2, *initialization* and *iterative reinforcement*. The approach follows a two-fold coarse-to-fine strategy inspired by TAPIR [3]. In the initial stage, trajectories are initialized based on the coarse resolution of the feature maps using a coarse network. Subsequently, in the second stage, the trajectories are iteratively refined

using fine-grained feature maps by a fine network, thus constituting a two-fold coarse-to-fine approach. This technique not only speeds up computation but also prevents the loss of important information due to downsampling. Although the networks in both stages estimate trajectories independently, they exploit point locations from the first frame to maintain spatial correlation and estimate coherent trajectories. Additionally, frame flow (i.e. $u_s - u_{s-1}$), representing the difference between consecutive frames, is naively passed to the model to make it aware of global appearance changes. The model can run on US sequences of any length and with any number of query points, depending on available memory.

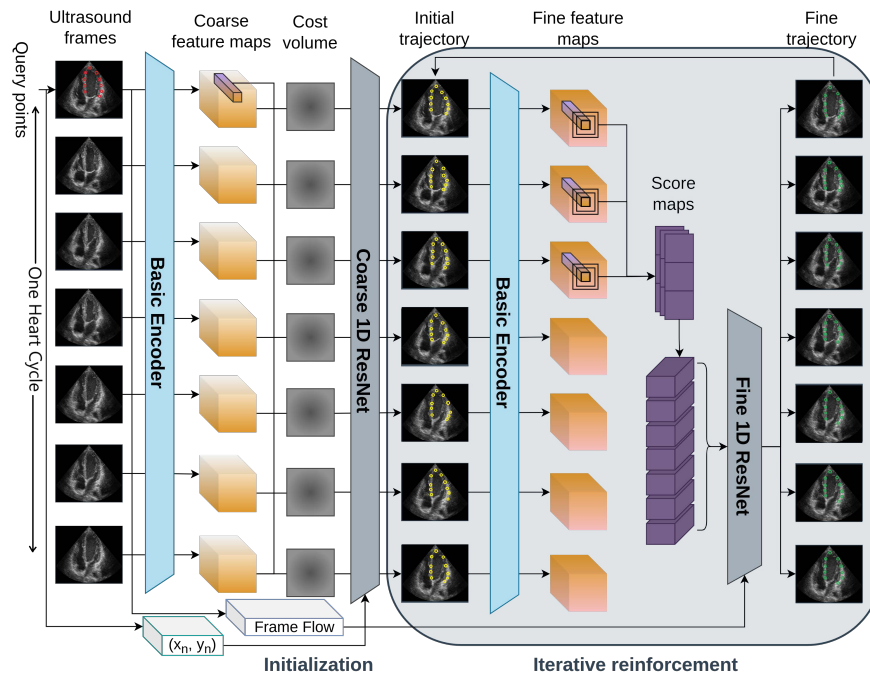


Fig. 2. EchoTracker is a two-fold coarse-to-fine model. Initially, it estimates coarse trajectories (yellow points) based on the cost volume for the given query points (red). It then imposes iterative reinforcement to obtain the fine trajectories (green points).

Initialization. The input contains S ultrasound images and N number of query points, as highlighted by red dots in Fig. 2. We utilize a pruned 2D residual convolutional network (basic encoder) [8] to generate coarse feature maps $F_s \in \mathbb{R}^{d \times \frac{H}{k} \times \frac{W}{k}}$ for each frame with $k = 8$ and $d = 64$. The pruning is motivated by a reduction in computational costs and the limited variability of data representation (e.g. grayscale, cyclic, velocity) present in echocardiog-

raphy. Given a query point in the first frame p_0^n , we extract a feature vector $f_{p_0^n} = \text{sample}(F_0, p_0^n)$, which captures the appearance of the point, through a bilinear sampling of the feature maps F_0 at that specific point location. Following that, we compute the cost volume $C_s^n = f_{p_0^n} \cdot \text{pyramid}(F_s)$ by correlating $f_{p_0^n}$ with features across the entire video, using multi-scale feature pyramids with $L = 4$ levels and kernel size $r = 3$. Finally, the cost volume is fed to a coarse 1D ResNet, similar to PIPs++ [21], which convolves across time to estimate the trajectory P_s^n for $s = 0, 1, \dots, S$. This is highlighted in yellow in Fig. 2. We choose a 1D ConvNet here rather than a 2D to prioritize temporal information, assuming that bilinear sampling already perceives the required spatial information present in the US images.

Iterative Reinforcement. Taking inspiration from PIPs [7] and more recent point tracking methods, we also refine the initial coarse trajectories through an iterative reinforcement process illustrated in Fig. 2. We hypothesize that our strong initialization can substantially decrease the number of iterations necessary to converge to a refined trajectory, and thus keep it at $I = 4$ for both training and evaluation. After initialization, we use the same basic encoder with a reduced downsampling factor $k = 2$ to produce fine feature maps $F_s \in \mathbb{R}^{d \times \frac{H}{2} \times \frac{W}{2}}$. Subsequently, unlike the initialization, we compute the cost volume $C_s^n = f_{p_s^n} \cdot \text{multicrop-pyramid}(F_s)$ for the query point by correlating the feature vector $f_{p_s^n}$ from the current frame s with multi-scale crops of pyramid features on the fine feature maps surrounding the point location. Inspired by the feature extraction within a fixed temporal span of the current frame [1] and multi-template update [21], we track changes in the point appearance by acquiring additional cost volumes (i.e., C_{s-2}^n, C_{s-4}^n) at a fixed temporal span from the current frame and always for the first C_0^n . Afterwards, we concatenate these cost volumes and pass them through a linear layer to obtain the score maps for the current frame. Similarly, we obtain score maps for all frames and concatenate them before passing to the next layer to generate updates Δp_s^n for the trajectory. This layer contains a fine 1D ResNet, which is a deeper and more weighted version of the coarse ResNet utilized in the initialization. Finally, updates are applied to the points to obtain the refined trajectory $p_{s,i}^n = p_{s,i-1}^n + \Delta p_{s,i-1}^n$, and the iterative reinforcement ensures the most fine-grained estimated trajectory. We supervise the model by considering all iterations in an end-to-end manner using the weighted L_1 distance between estimated and ground-truth trajectories as the loss function similar to PIPs++ [21].

4 Experiments and Results

4.1 Datasets

Learning-based methods are usually trained using synthetic data and tested on real data annotated by humans [5]. Unfortunately, the availability of high-quality synthetic point tracking data for echocardiography that can adequately

replicate real-world data remains scarce. In this work, we chose to rely on human-annotated data, similar to real-world data annotations in TAP benchmarks [2]. Our trajectories are generated in a semi-supervised fashion using a traditional tracking algorithm and undergo tuning for optimal tracking by clinical experts with over 20 years of experience (ASE-level III [20]). In addition, experts perform quality assurance by rejecting points that do not follow the tissue properly, and we discard those from training and evaluation. We retrieve four splits, as illustrated in Table 1, focusing on tracking the left ventricle myocardium from three acoustic views, namely the apical four-chamber, two-chamber and long-axis. All data were collected using GE Vivid E95 scanners, with phased array probes. All data originated from patient studies approved by the Regional Committee for Medical Research Ethics and conducted according to the Helsinki Declaration. Written informed consent was obtained from all patients. DS-A is a test-retest dataset, meaning that the same patient is being scanned twice in immediate succession by two different operators. Therefore, we expect the patient to be in the same physical condition for both exams, serving as a reference for the reproducibility of the method. This dataset is independent and used exclusively for testing. On the other hand, DS-B, DS-C, and DS-D are subsets of the same dataset but analyzed by three different observers through random selection. We train the model using these datasets and evaluate its performance on the DS-A dataset.

Table 1. Ultrasound point tracking datasets and selected characteristics. The number of points and frames, as well as the height and width of the images, are given as average with range (min, max) in parenthesis.

Dataset	Patients	Videos	Points	Frames	Height	Width
DS-A	40	210	73 (57-91)	93 (57-150)	593 (583-640)	571 (496-650)
DS-B	913	2731	81 (49-115)	85 (44-151)	592 (583-640)	630 (453-844)
DS-C	615	1837	80 (49-111)	84 (44-150)	592 (583-640)	631 (496-778)
DS-D	643	1922	66 (45-111)	84 (46-151)	592 (421-640)	634 (392-856)

4.2 Implementation details

Similar to Zheng et al. [21], we use a learning rate of $5 \cdot 10^{-4}$, a one-cycle scheduler [16], and the AdamW optimizer. Images are resized to 256×256 for both training and evaluation by default to limit the GPU memory footprint. Initially, we train our model on DS-B for 22 epochs using a sequence length of $S = 36$. We fine-tune the model on DS-C for 50 epochs with $S = 68$, and then on DS-D for a single epoch with $S = 68$. Throughout the training process, we consistently use all available points with a batch size of 1. The training time on a single GPU, completing 50 epochs on the DS-B dataset typically takes over one week. We implemented our framework in PyTorch and used an NVIDIA GeForce RTX 3090 (24GB) GPU for both training and evaluation of models.

4.3 Evaluation

Evaluation Metrics. Following SOTA point tracking literature, we report average position accuracy ($< \delta_{avg}^x$) as proposed in TAP-Vid [2] for the technical performance. Position accuracy ($< \delta^x$), measures the proportion (%) of all the query points that fall within the threshold distance of pixels (e.g., $x = 1, 2$) from their ground truth across the entire video and $< \delta_{avg}^x$ averages over five thresholds: 1, 2, 4, 8, and 16 pixels. Given the sensitivity of tracking in US, we also report $< \delta^x$ for individual threshold and median trajectory error (MTE) [21], which calculates the distance in pixels between the estimated and ground truth trajectories. To show the efficiency of our model, we also include the average inference time (AIT) per video measured in seconds. Furthermore, to evaluate our model in a clinical setting, we calculate the peak GLS, which is defined as the relative change of the longitudinal ventricular length from end-diastole to minimum ventricular length.

Technical Performance. For reference, we conduct an evaluation of current SOTA models directly on our DS-A dataset without fine-tuning. These results are summarized in Table 2 which shows that PIPs++ and CoTracker stand out. Thus, we use these two models as baselines for our subsequent experiments.

Table 2. Performance comparison of the state-of-the-art methods on the DS-A dataset without fine-tuning.

Method	$< \delta^1 \uparrow$	$< \delta^2 \uparrow$	$< \delta^4 \uparrow$	$< \delta^8 \uparrow$	$< \delta^{16} \uparrow$	$< \delta_{avg}^x \uparrow$	MTE \downarrow
PIPs	5	12	28	54	81	36	10.22
TAP-Net	5	17	41	70	89	44	6.50
TAPIR	7	21	47	78	95	50	5.63
PIPs++	8	22	47	78	96	50	5.31
CoTracker	13	27	52	79	95	53	5.28

δ : Position accuracy (%), MTE: Median trajectory error (pixel)

The performance of EchoTracker compared with PIPs++ [21] and CoTracker [10] fine-tuned with the same training datasets is displayed in Table 3. CoTracker shows a slight decline in performance after fine-tuning, likely attributed to sub-optimal implementation details or data processing, necessitating in-depth investigation in future studies. Our model demonstrates superior performance across all metrics, surpassing other methods by a significant margin. Moreover, it shows faster inference times compared to these alternatives.

To investigate if our two-stage architecture can reduce the number of iterations in the reinforcement stage, we solely train the initialization part on a limited part of the DS-B dataset. This initialization part yields $< \delta_{avg}^x = 48\%$, surpassing the performance of baseline PIPs and TAP-Net. Thus, it can be expected that the iterative reinforcement stage would require less effort to refine and smooth the initial trajectory. Our empirical investigations found that models trained on short sequences struggle when tracking longer videos. Therefore,

Table 3. Technical performance of EchoTracker on the DS-A test-retest dataset compared to state-of-the-art methods.

Method	$< \delta^1 \uparrow$	$< \delta^2 \uparrow$	$< \delta^4 \uparrow$	$< \delta^8 \uparrow$	$< \delta^{16} \uparrow$	$< \delta_{avg}^x \uparrow$	MTE \downarrow	AIT \downarrow
PIPs++	15	36	70	94	100	63	3.28	0.42
CoTracker	8	22	47	78	96	50	5.45	1.67
EchoTracker (ours)	19	43	76	96	100	67	2.86	0.24

δ : Position accuracy (%), MTE: Median trajectory error (pixel), AIT: Average inference time (s)

a straightforward improvement is to train on longer sequences. One limitation of this work is tracking points undergoing large motions, such as those on the cardiac valves. This was anticipated since the training dataset does not include such representations. Additionally, the spatial information is limited by bilinear sampling and the 1D ResNet used in the model’s initialization. Furthermore, we experiment with replacing the fine ResNet with a transformer, following the approach in CoTracker [10]. Surprisingly, this modification led to a drop in performance. The reason may be attributed to the fact that transformers typically require pretraining on large-scale datasets to outperform ConvNets [4]. Exploring this aspect further could be an intriguing avenue for future research as we witness CoTracker outperform all the other methods without fine-tuning.

Clinical Performance. In Table 4, we present the GLS results compared to the reference measurements and in a test-retest situation. We also list results from solutions developed by others, albeit calculated on their private datasets. Our method, as well as fine-tuned PIPs++, performs favourably compared to other published work. However, the methods are tested on different private datasets, so a direct comparison was not possible.

Table 4. Clinical results for GLS calculations compared to reference measurements and in a test-retest scenario.

Method	Reference			Test-retest		
	μ	$\sigma \downarrow$	$MAD \downarrow$	μ	$\sigma \downarrow$	$MAD \downarrow$
c-PWC-Net-60A [5]	1.85	2.73	N/A			
us2ai [11]	0.68	2.52	2.0			
EchoPWCNet [13, 14]	-1.4	1.9	1.8	0.0	1.9	1.6
PIPs++	-1.21	1.95	1.76	0.11	1.62	1.28
CoTracker	-0.82	2.40	1.98	-0.11	2.47	1.96
EchoTracker (ours)	-0.13	1.78	1.36	-0.13	1.55	1.21

μ : Bias (%), σ : Standard deviation (%), MAD : Mean absolute deviation (%)

5 Conclusion

We have introduced modern general-purpose point tracking approaches in echocardiography. Through a comprehensive analysis of several SOTA methods, we de-

sign a two-fold coarse-to-fine architecture and propose a novel model for tracking myocardial tissue. Our assessment demonstrates that this new approach not only outperforms SOTA solutions but also enhances the measurements of GLS in a test-retest scenario. Thus, learning-based point tracking holds the potential to elevate both the diagnostic and prognostic utility of myocardial function measurements, marking a notable step forward in the field of echocardiography.

Acknowledgments. This work was funded by the Liaison Committee between the Central Norway Regional Health Authority and the Norwegian University of Science and Technology (2023-34187, 2024-37009), the Centre for Innovative Ultrasound Solutions (CIUS) and the Precision Health Center for optimized cardiac care (ProCardio) - both Norwegian Research Council centers for research-based innovation, St. Olavs University Hospital, and SINTEF Digital.

Disclosure of Interests. A. Chernyshov, J. Nyberg, H. Dalen, L. Lovstakken, B. Grenne and A. Østvik hold positions at ProCardio, where GE HealthCare is a consortium partner. H. Dalen, L. Lovstakken, A. Østvik and B. Grenne hold positions at CIUS, where GE HealthCare also is a consortium partner. L. Lovstakken also acts as a part-time consultant for GE HealthCare.

References

1. Azad, M.A., Mohammed, A., Waszak, M., Elvesæter, B., Ludvigsen, M.: Multi-label video classification for underwater ship inspection. In: OCEANS 2023 - Limerick. pp. 1–10 (2023)
2. Doersch, C., Gupta, A., Markeeva, L., Recasens, A., Smaira, L., Aytar, Y., Carreira, J., Zisserman, A., Yang, Y.: Tap-vid: A benchmark for tracking any point in a video. *Advances in Neural Information Processing Systems (NeurIPS)* **35**, 13610–13626 (2022)
3. Doersch, C., Yang, Y., Vecerik, M., Gokay, D., Gupta, A., Aytar, Y., Carreira, J., Zisserman, A.: Tapir: Tracking any point with per-frame initialization and temporal refinement. *ICCV* (2023)
4. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations (ICLR)* (2021)
5. Evain, E., Sun, Y., Faraz, K., Garcia, D., Saloux, E., Gerber, B.L., De Craene, M., Bernard, O.: Motion estimation by deep learning in 2d echocardiography: synthetic dataset and validation. *IEEE transactions on medical imaging* **41**(8), 1911–1924 (2022)
6. Farsalinos, K.E., Daraban, A.M., Ünlü, S., Thomas, J.D., Badano, L.P., Voigt, J.U.: Head-to-head comparison of global longitudinal strain measurements among nine different vendors: the eacvi/ase inter-vendor comparison study. *Journal of the American Society of Echocardiography* **28**(10), 1171–1181 (2015)
7. Harley, A.W., Fang, Z., Fragkiadaki, K.: Particle video revisited: Tracking through occlusions using point trajectories. In: *European Conference on Computer Vision (ECCV)*. pp. 59–75. Springer (2022)

8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR). pp. 770–778 (2016)
9. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: FlowNet 2.0: Evolution of optical flow estimation with deep networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR). pp. 2462–2470 (2017)
10. Karaev, N., Rocco, I., Graham, B., Neverova, N., Vedaldi, A., Ruppel, C.: Cotracker: It is better to track together. arXiv:2307.07635 (2023)
11. Myhre, P.L., Hung, C.L., Frost, M.J., Jiang, Z., Ouwerkerk, W., Teramoto, K., Svedlund, S., Saraste, A., Hage, C., Tan, R.S., et al.: External validation of a deep learning algorithm for automated echocardiographic strain measurements. *European Heart Journal-Digital Health* **5**(1), 60–68 (2024)
12. Østvik, A., Salte, I.M., Smistad, E., Nguyen, T.M., Melichova, D., Brunvand, H., Haugaa, K., Edvardsen, T., Grenne, B., Lovstakken, L.: Myocardial function imaging in echocardiography using deep learning. *IEEE Transactions on Medical Imaging* **40**(5), 1340–1351 (2021)
13. Salte, I.M., Østvik, A., Olaisen, S.H., Karlsen, S., Dahlslett, T., Smistad, E., Eriksen-Volnes, T.K., Brunvand, H., Haugaa, K.H., Edvardsen, T., et al.: Deep learning for improved precision and reproducibility of left ventricular strain in echocardiography: A test-retest study. *Journal of the American Society of Echocardiography* (2023)
14. Salte, I.M., Østvik, A., Smistad, E., Melichova, D., Nguyen, T.M., Karlsen, S., Brunvand, H., Haugaa, K.H., Edvardsen, T., Lovstakken, L., et al.: Artificial intelligence for automatic measurement of left ventricular strain in echocardiography. *Cardiovascular Imaging* **14**(10), 1918–1928 (2021)
15. Sand, P., Teller, S.: Particle video: Long-range motion estimation using point trajectories. *International journal of computer vision* **80**, 72–91 (2008)
16. Smith, L.N., Topin, N.: Super-convergence: Very fast training of neural networks using large learning rates. In: *Artificial intelligence and machine learning for multi-domain operations applications*. vol. 11006, pp. 369–386. SPIE (2019)
17. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR). pp. 8934–8943 (2018)
18. Voigt, J.U., Pedrizzetti, G., Lysyansky, P., Marwick, T.H., Houle, H., Baumann, R., Pedri, S., Ito, Y., Abe, Y., Metz, S., et al.: Definitions for a common standard for 2d speckle tracking echocardiography: consensus document of the eacvi/ase/industry task force to standardize deformation imaging. *European Heart Journal-Cardiovascular Imaging* **16**(1), 1–11 (2015)
19. Wang, Q., Chang, Y.Y., Cai, R., Li, Z., Hariharan, B., Holynski, A., Snavely, N.: Tracking everything everywhere all at once. In: *International Conference on Computer Vision (ICCV)* (2023)
20. Weigers, S., Ryan, T., Arrighi, J., Brown, S., Canaday, B., Damp, J., Diaz-Gomez, J., Figueredo, V., Garcia, M., Gillam, L., et al.: Acc/aha/ase advanced training statement on echocardiography (revision of the 2003 acc/aha clinic competence statement on echocardiography). *Journal of the American College of Cardiology* **74**, 377–402 (2019)
21. Zheng, Y., Harley, A.W., Shen, B., Wetzstein, G., Guibas, L.J.: Pointodyssey: A large-scale synthetic dataset for long-term point tracking. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 19855–19865 (2023)