



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.

Cardiac Copilot: Automatic Probe Guidance for Echocardiography with World Model

Haojun Jiang^{1,2*}, Zhenguo Sun^{2*}, Ning Jia², Meng Li², Yu Sun², Shaqi Luo², Shiji Song¹, and Gao Huang^{1,2}✉

¹ Department of Automation, BNRist, Tsinghua University, Beijing, China

² Beijing Academy of Artificial Intelligence, Beijing, China

jhj20@mails.tsinghua.edu.cn, hitsunzhenguo@gmail.com,

gaohuang@tsinghua.edu.cn

Abstract. Echocardiography is the only technique capable of real-time imaging of the heart and is vital for diagnosing the majority of cardiac diseases. However, there is a severe shortage of experienced cardiac sonographers, due to the heart’s complex structure and significant operational challenges. To mitigate this situation, we present a Cardiac Copilot system capable of providing real-time probe movement guidance to assist less experienced sonographers in conducting freehand echocardiography. This system can enable non-experts, especially in primary departments and medically underserved areas, to perform cardiac ultrasound examinations, potentially improving global healthcare delivery. The core innovation lies in proposing a data-driven world model, named Cardiac Dreamer, for representing cardiac spatial structures. This world model can provide structure features of any cardiac planes around the current probe position in the latent space, serving as an precise navigation map for autonomous plane localization. We train our model with real-world ultrasound data and corresponding probe motion from 110 routine clinical scans with 151K sample pairs by three certified sonographers. Evaluations on three standard planes with 37K sample pairs demonstrate that the world model can reduce navigation errors by up to 33% and exhibit more stable performance.

Keywords: Echocardiography · World Model · Probe Guidance · Plane Navigation.

1 Introduction

Cardiovascular diseases, as the primary cause of death and health risks in humans, affected 422.7 million people globally and led to 17.9 million deaths, accounting for 31% of all global fatalities in 2015 [14, 17]. When screening for heart diseases, echocardiography is indispensable, as it is the only imaging technique

*These authors contributed equally to this work. This work was done while Haojun Jiang was an intern at Beijing Academy of Artificial Intelligence.

✉Corresponding author.

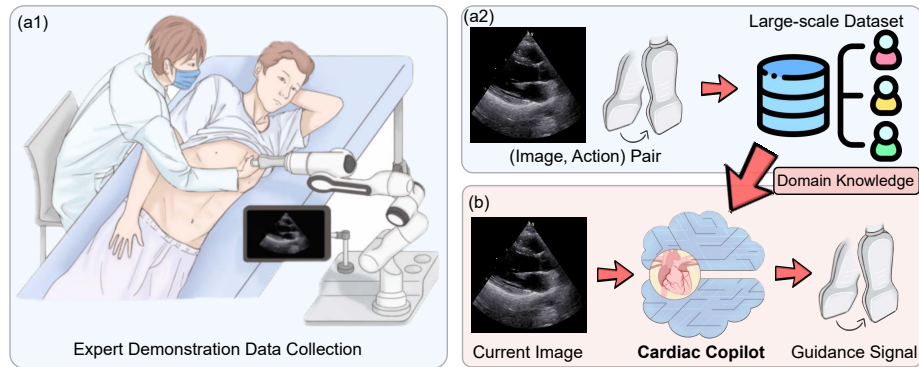


Fig. 1. Schematic diagram of cardiac copilot system. (a1-2) We collect a large-scale expert demonstration dataset with extensive ultrasound image and probe motion sample pairs using a robotic arm. The demonstration data contains critical medical knowledge required for performing echocardiography. **(b)** We encapsulate experts’ knowledge within deep neural networks to facilitate automatic probe guidance.

that allows real-time observation of the heart’s structure and function. However, cardiac ultrasound is challenging to operate due to the following reasons: (1) the heart’s complex structure; (2) limited observation windows due to rib obstruction; (3) the need for very precise adjustments; (4) particularly significant individual variations. Consequently, echocardiography require highly skilled sonographers who possess extensive anatomical knowledge and rich operational experience. According to the guideline of the American Society of Echocardiography [3, 4], a trainee with a medical background requires at least 12 months of clinical training before potentially being qualified for the position. The long training period combined with the rising incidence of cardiovascular diseases has resulted in a significant shortage of skilled cardiac sonographers, especially pronounced in primary departments and medical underserved areas.

With technological advancements [5, 8, 10–12, 19], AI-driven systems show immense potential to significantly improve the efficiency of medical ultrasound processes. For example, Droste *et al.* [1] proposed an US-GuideNet to provide only rotation movement guidance signal for acquiring standard planes in routine freehand obstetric ultrasound scanning. However, in routine ultrasound scanning, the first step often involves moving close to the target plane through translation. Thus, providing a six-dimension (6D) guidance signal is essential, which includes signals for both translation and rotation. Narang *et al.* [13] proposed an AI-driven commercial solution to guide nurses to acquire echocardiogram by suggesting approximate adjustment directions in one degree of freedom. They further conducted clinical trials and found that it could assist nurses in locating cardiac target planes that meet diagnostic requirements. Nonetheless, the specifics of the methodological implementation remain undisclosed, impeding progress in cardiac ultrasound domain. Later, Shida *et al.* [15, 16] investigated the poten-

tial of using manually predefined search strategies to acquire the Parasternal Long-Axis (PLAX) plane. The authors conducted a small-scale experiment on five healthy subjects, with the upper limit of the success rate being only 63.3%. This indicates that the rule-based guidance strategy lacks generalizability and practicality. Despite on-going efforts, there is still no publicly available research on cardiac probe navigation, which hampers further advancement in the field.

In this paper, we present an AI-driven cardiac probe guidance system, named Cardiac Copilot, capable of providing real-time probe guidance signals for locating target standard planes in routine freehand echocardiography. As illustrated in Fig. 1, the fundamental logic involves an artificial intelligence network that, based on the current ultrasound image, outputs a six-dimension action command to guide the operator towards the target standard plane. In clinical practice, highly skilled sonographers can locate the target section with fewer explorations, as they can imagine the potential image that might result from a specific adjustment based on the current image, thereby achieving accurate adjustment. Inspired by the above domain expertise, we propose a data-driven world model named Cardiac Dreamer to represent the spatial structure of the heart. Specifically, it can “imagine” the features of adjacent continuous planes in the latent space based on the current ultrasound image and relative positional relationships. Upon integrating Cardiac Dreamer into the system, it effectively gains a precise map in latent space, significantly enhancing its navigational proficiency. We collected expert operation data on acquiring standard planes from 125 routine clinical scans with 188K sample pairs by three certified sonographer, to drive the model’s learning of generalizable guidance capabilities. The data were collected by manually operating a probe attached to the end of a robotic arm, with precise probe 6D pose information obtained through the arm’s internal sensors. Assessments conducted on three standard planes using 37K pairs of samples reveal that the world model can decrease navigation errors by up to 33% and demonstrate enhanced stability in performance.

2 Method

In this section, we provide details regarding the methodological implementation about automatic probe guidance for freehand echocardiography. We first introduce the target-oriented guidance framework which enables the base capabilities of locating standard planes in Sec 2.1. Moreover, in Sec 2.2, we present the Cardiac Dreamer, which foresees the state that would result from the executing predicted action. By collaborating with Cardiac Dreamer, the base model is able to exhibit improved guidance performance.

2.1 Target-oriented Guidance Framework

In clinical practice, the ultimate goal of cardiac ultrasound is to locate each standard two-dimensional plane and observe the heart’s structure and function on these standard planes. Consequently, we aim to train a guidance network whose

goal is to directly predict the relative positional relationship between the current probe location and the target plane, without the need to consider the intermediate path, thereby making it target-oriented. Achieving this functionality is challenging as it requires the network to understand cardiac anatomy, interpret ultrasound images, and have spatial imagination skills to map 2D ultrasound images onto 3D cardiac structures. In fact, during routine examinations, sonographers’ decisions on each probe adjustment is the result of comprehensively applying these knowledge and skills. Thus, to acquire guidance capabilities, we adopt the data-driven imitation learning strategy [2, 9, 10] to learn extensive domain knowledge from the large-scale expert demonstration dataset.

Given an input ultrasound image at time t , representing the current state $\mathbf{S}_t \in \mathbb{R}^{H \times W}$, the operator needs a 6D guidance signal $\mathbf{a}_t \in \mathbb{R}^6$ to locate the target plane. This six-dimensional vector \mathbf{a}_t comprises the first three dimensions representing translational movements in the x, y, and z directions, while the last three dimensions indicate rotational movements in roll, pitch, and yaw. Let $\mathcal{D} = \{d_1, d_2, \dots, d_n\}$ be all the expert demonstration data collected from the routine examinations. Specifically, $d_i = \{(\mathbf{S}_t, \mathbf{p}_t) | \mathbf{S}_t \in \mathbb{R}^{H \times W}, \mathbf{p}_t \in \mathbb{R}^6, t \in \mathcal{T}\}$ is one demonstration of standard plane acquisition, where \mathbf{p}_t is 6D probe’s pose information and corresponding ultrasound image \mathbf{S}_t at time t . When $t = T$, \mathbf{S}_T is the standard section and p_T is the probe’s target pose. By comparing other poses p_t in the demonstration sequence with p_T , we obtain the \mathcal{D}' with each sequence d'_i as $\{(\mathbf{S}_t, \hat{\mathbf{a}}_t) | t \in \mathcal{T} \setminus T\}$ with groundtruth supervision $\hat{\mathbf{a}}_t$. Consequently, we train a policy network $\pi_\theta : \mathbf{S}_t \mapsto \mathbf{a}_t$ with parameters θ . As shown in Fig. 2, this network consists of a feature encoder $E(\cdot)$ and a guidance layer $G(\cdot)$:

$$\mathbf{a}_t = G(E(\mathbf{S}_t), \mathbf{q}_i), \quad (1)$$

where $\mathbf{q}_i, i \in [0, Q - 1]$ is a learnable query feature, indicating one of the target planes. Finally, the policy network π_θ is optimized with the following loss:

$$\mathcal{L}_{total} = \frac{1}{|\mathcal{D}'|} \sum_{d'_i \in \mathcal{D}'} \frac{1}{|d'_i|} \sum_{(\mathbf{S}_t, \hat{\mathbf{a}}_t) \in d'_i} \mathcal{L}_{\text{SmoothL1}}(\mathbf{a}_t, \hat{\mathbf{a}}_t), \quad (2)$$

where $|\mathcal{D}|$ and $|D_i|$ are the number of elements inside each set.

2.2 Foresee States by Cardiac Dreamer

This task can be likened to an autonomous navigation mission with the probe near the chest area, where it is usually necessary to construct an precise dense map to understand the relative positional relationships between any points. However, the aforementioned network only utilizes sparse relative positional information between arbitrary planes and the target plane. Upon closer examination of the demonstration data, we find a wealth of unused relative positional information between non-target planes. Assuming there are M intermediate planes in a demonstration sequence, the unused structural information accounts for




Fig. 2. Diagram of target-oriented guidance framework. *Left* is the policy network which provides the basic guidance signal for locating the target plane. *Right* is the Cardiac Dreamer which foresees states reached by executing actions output by the policy network and refines the actions based on these states. The details of the action combining operator is shown in Eq. 6.

$(M - 1)/M$ of the total data. Therefore, based on the extensive positional relationships, we can construct a world model [5–7] for heart, named Cardiac Dreamer, representing dense cardiac structures in latent space.

Given the processed demonstration dataset D'_i , we create the permutation dataset D'' with $d''_i = \{(\mathbf{S}_{t_1}, \mathbf{S}_{t_2}, \hat{\mathbf{a}}_{t_1 \rightarrow t_2}, \hat{\mathbf{a}}_{t_1}, \hat{\mathbf{a}}_{t_2}) | t_1 \neq t_2\}$, where $\hat{\mathbf{a}}_{t_1 \rightarrow t_2}$ is relative positional between S_{t_1} and S_{t_2} , and $\hat{\mathbf{a}}_{t_1}, \hat{\mathbf{a}}_{t_2}$ is relative positional between S_{t_1}, S_{t_2} and the target plane. Then, \mathbf{S}_{t_1} is processed by the feature encoder $E(\cdot)$ inside the policy network to yield \mathbf{f}_{t_1} :

$$\mathbf{f}_{t_1} = E(\mathbf{S}_{t_1}). \quad (3)$$

Cardiac Dreamer $I(\cdot)$ receives \mathbf{f}_{t_1} and $\hat{\mathbf{a}}_{t_1 \rightarrow t_2}$ as input and outputs the feature \mathbf{f}'_{t_2} of the corresponding plane:

$$\mathbf{f}'_{t_2} = I(\mathbf{f}_{t_1}, \hat{\mathbf{a}}_{t_1 \rightarrow t_2}). \quad (4)$$

We aim for Cardiac Dreamer to learn the spatial relationships between various planes, thereby gaining a comprehensive understanding of cardiac structures. Furthermore, the generated \mathbf{f}'_{t_2} is used to predict the guidance signal \mathbf{a}'_{t_2} for locating the target plane:

$$\mathbf{a}'_{t_2} = G(\mathbf{f}'_{t_2}, \mathbf{q}_i). \quad (5)$$

Then, the predicted guidance signal \mathbf{a}'_{t_2} is combined with $\hat{\mathbf{a}}_{t_1 \rightarrow t_2}$ to get \mathbf{a}'_{t_1} :

$$\mathbf{a}'_{t_1} = U^{-1}(U(\hat{\mathbf{a}}_{t_1 \rightarrow t_2})U(\mathbf{a}'_{t_2})), \quad (6)$$

where $U(\cdot)$ is the function that converts the vector-form action \mathbf{a} into the standard homogeneous transformation matrix $U(\mathbf{a}) \in \mathbb{R}^{4 \times 4}$. Thus, the total loss for joint optimization of the policy network with a Cardiac Dreamer is:

$$\mathcal{L}_{total} = \mathcal{L}_{\text{SmoothL1}}(\mathbf{a}'_{t_1}, \hat{\mathbf{a}}_{t_1}) + \mathcal{L}_{\text{SmoothL1}}(\mathbf{a}'_{t_2}, \hat{\mathbf{a}}_{t_2}). \quad (7)$$

In experiments, we observe that Cardiac Dreamer can improve the guidance performance and provide more stable guidance signal.

3 Experiments

3.1 Setup

In the experimental setup, three standard echocardiographic planes, *i.e.*, the Parasternal Long-Axis (PLAX), the Parasternal Short-Axis Aortic Valve (PSAX-AV), and the Parasternal Short-Axis Mitral Valve (PSAX-MV), were chosen for their clinical importance, offering critical insights into cardiac anatomy and functionality. The ultrasound image data was acquired from health adult male volunteers using a Franka Panda robot arm with an ultrasound probe (M5S probe, GE Vivid E7 machine) manipulated by professional echocardiographers, meanwhile the pose information of the probe was also recorded synchronously by utilizing Robot Operation System (ROS). Overall, approximately 188K cardiac ultrasound images paired with probe motion were collected from 125 routine clinical scans, forming a diverse dataset for training the proposed Cardiac Copilot models. The whole data collecting process was approved and supervised by The Tsinghua University Science and Technology Ethics Committee. Data processing involved aligning relative pose coordinates of each image with target planes which were identified by professionals, ensuring a unified coordinate system for each plane. The dataset was split into 110 scans (about 151K images) for training and 15 scans (about 37K images) for testing. Note that the test set included individuals not seen during training to evaluate the algorithm’s generalization.

3.2 Implementation Details

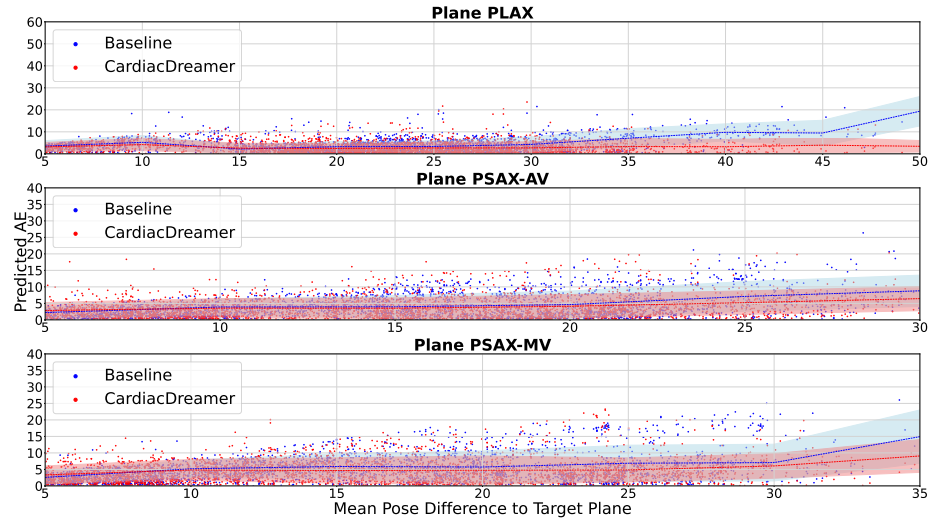
As shown in Fig. 2, the model architecture includes a ResNet-34 for feature extraction, a guidance layer with three fully connected layers (the output dimensions are 2048, 512, and 6), and a Cardiac Dreamer network based on the original transformer architecture proposed in [18]. Additionally, one-hot encoding for plane selection is introduced, allowing the model to adapt relative pose predictions based on the chosen target plane. Model training, conducted with Distributed Data Parallel (DDP) training on four NVIDIA RTX 3090 GPUs, spanned 300 epochs. The training process employed the CosineAnnealingLR scheduler with 0.0001 learning rate, AdamW optimizer, and Automatic Mixed Precision (AMP). For inference, we tested on an Nvidia RTX 3090, and the inference time for a single image is 59.4ms which meets practical requirements.

3.3 Evaluation Metrics

In order to assess the accuracy and robustness of our proposed Cardiac Dreamer compared to the policy network which comprises merely feature encoder and guidance layer as baseline, we designed the evaluation metrics deliberately. The primary metric for evaluation is the Mean Absolute Error (MAE) between the predicted and ground truth probe poses. This quantitative measure assesses the accuracy of our model in imitating the spatial movement of the echocardiography probe. Simultaneously, the stability of the model was evaluated by measuring

Table 1. Comparisons between the cardiac dreamer and the baseline. We report MAE results representing probe guidance errors across three standard planes.

Plane	Model	Translation (mm)			Rotation (degree)		
		x	y	z	rx	ry	rz
PLAX	Baseline	7.85	5.97	3.64	6.77	4.92	10.46
	Cardiac Dreamer	6.10 (-22.3%)	4.85 (-18.8%)	2.64 (-27.5%)	5.76 (-14.9%)	4.46 (-9.3%)	9.57 (-8.5%)
PSAX-AV	Baseline	7.70	6.01	3.30	5.17	7.00	13.26
	Cardiac Dreamer	7.48 (-2.9%)	4.83 (-19.6%)	2.67 (-19.1%)	4.20 (-18.8%)	5.65 (-19.3%)	9.77 (-26.3%)
PSAX-MV	Baseline	10.11	8.21	4.45	5.50	8.57	13.20
	Cardiac Dreamer	8.53 (-15.6%)	6.93 (-15.6%)	3.62 (-18.7%)	4.57 (-16.9%)	5.73 (-33.1%)	10.21 (-22.7%)

**Fig. 3. Detailed analysis.** For optimal viewing, please consider enlarging and viewing in color. The change of predicted absolute error of different samples with respect to the mean pose difference to target plane on three standard planes. The 'mean pose difference' refers to calculating the mean absolute error between the pose of the current plane and that of the target plane. The dashed line represents the mean value, and the shading represents the standard deviation.

the standard deviation of the absolute error to the target plane for all poses from the evaluation scans.

3.4 Results and Analysis

Accuracy. The Table. 1 presents the MAE results for probe guidance prediction across three standard echocardiographic planes: PLAX, PSAX-AV, and PSAX-MV. The results indicate that the Cardiac Dreamer model consistently outperforms the baseline model across all translation and rotation components. The improvements are most pronounced in the z translation component for PLAX, rz rotation component for PSAX-AV, and ry rotation component for PSAX-MV,

with 27.5%, 26.3%, and 33.1% reduction in MAE respectively. These findings suggest that Cardiac Dreamer’s ability to understand and model the spatial relationships and anatomical structures involved in echocardiographic imaging contributes to its enhanced guidance accuracy.

Stability. The relationship between all probe poses and their Absolute Error (AE) to target planes is shown in Fig. 3. The figure shows, the Cardiac Dreamer model achieves significantly lower AE compared to the baseline model across different mean pose differences to the target plane. Moreover, the shaded regions for Cardiac Dreamer are narrower, indicating a smaller standard deviation and more consistent performance with different poses. The inclusion of standard deviation in the plots provides valuable insights into the reliability and robustness of the models’ predictions. A smaller standard deviation, as observed for Cardiac Dreamer, suggests that the model’s predictions are more tightly clustered around the mean, indicating greater precision and consistency across different scans. Conversely, a larger standard deviation, as seen for the baseline model, implies greater variability in the predictions, potentially indicating sensitivity to specific probe poses or scans. This variability could lead to less reliable performance in practical applications.

Analysis. Overall, the results demonstrate that the proposed Cardiac Dreamer model not only achieves superior general performance in terms of navigation guidance accuracy but also exhibits greater consistency and robustness, as evidenced by the lower mean absolute error and the smaller standard deviations across all three echocardiographic planes evaluated. This combination of improved accuracy and reliability makes Cardiac Dreamer a more suitable choice for real-world scenarios in intelligent probe guidance scanning.

4 Conclusion

In this work, we present a Cardiac Copilot system designed to offer real-time guidance on probe movements, aiding novice sonographers in performing free-hand echocardiography. We innovatively introduce a data-driven world model named Cardiac Dreamer, which represents cardiac structures in feature space and serves as a latent map for automatic probe navigation. Evaluations on the large-scale dataset show that Cardiac Dreamer consistently reduces the guidance errors across six-dimensions and exhibits more stable performance. AI-driven probe guidance system shows the potential to reduce sonographer skill requirements, thereby increasing the capacity for cardiac ultrasound examinations in primary hospitals and areas with scarce medical resources. In the future, the probe guidance system can also be used to direct a robotic arm for fully autonomous cardiac ultrasound scanning, serving as a core module in autonomous robots. This is a promising future direction that can significantly reduce repetitive manual operations, greatly enhance the supply capability of cardiac ultrasound, and improve the standardization of the examination process.

Acknowledgement. This work was supported in part by the National Key R&D Program of China (2021ZD0140407), the NSFC (62321005) and the Deng Feng Fund.

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

References

1. Droste, R., Drukker, L., Papageorghiou, A.T., Noble, J.A.: Automatic probe movement guidance for freehand obstetric ultrasound. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part III 23. pp. 583–592. Springer (2020)
2. Duan, Y., Andrychowicz, M., Stadie, B., Jonathan Ho, O., Schneider, J., Sutskever, I., Abbeel, P., Zaremba, W.: One-shot imitation learning. *Advances in neural information processing systems* **30** (2017)
3. Ehler, D., Carney, D.K., Dempsey, A.L., Rigling, R., Kraft, C., Witt, S.A., Kimball, T.R., Sisk, E.J., Geiser, E.A., Gresser, C.D., et al.: Guidelines for cardiac sonographer education: recommendations of the american society of echocardiography sonographer training and education committee. *Journal of the American society of echocardiography* **14**(1), 77–84 (2001)
4. Gardner, C.J., Brown, S., Hagen-Ansert, S., Harrigan, P., Kisslo, J., Kisslo, K., Kwan, O.L., Menapace, F., Otto, C., Pandian, N., et al.: Guidelines for cardiac sonographer education: report of the american society of echocardiography sonographer education and training committee. *Journal of the American Society of Echocardiography* **5**(6), 635–639 (1992)
5. Ha, D., Schmidhuber, J.: World models. arXiv preprint arXiv:1803.10122 (2018)
6. Hafner, D., Lillicrap, T., Ba, J., Norouzi, M.: Dream to control: Learning behaviors by latent imagination. arXiv preprint arXiv:1912.01603 (2019)
7. Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., Davidson, J.: Learning latent dynamics for planning from pixels. In: International conference on machine learning. pp. 2555–2565. PMLR (2019)
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. pp. 770–778 (2016)
9. Ho, J., Ermon, S.: Generative adversarial imitation learning. *Advances in neural information processing systems* **29** (2016)
10. Hussein, A., Gaber, M.M., Elyan, E., Jayne, C.: Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)* **50**(2), 1–35 (2017)
11. Jiang, H., Lin, Y., Han, D., Song, S., Huang, G.: Pseudo-q: Generating pseudo language queries for visual grounding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15513–15523 (2022)
12. Jiang, H., Zhang, J., Huang, R., Ge, C., Ni, Z., Lu, J., Zhou, J., Song, S., Huang, G.: Cross-modal adapter for text-video retrieval. arXiv preprint arXiv:2211.09623 (2022)
13. Narang, A., Bae, R., Hong, H., Thomas, Y., Surette, S., Cadieu, C., Chaudhry, A., Martin, R.P., McCarthy, P.M., Rubenson, D.S., et al.: Utility of a deep-learning

- algorithm to guide novices to acquire echocardiograms for limited diagnostic use. *JAMA cardiology* **6**(6), 624–632 (2021)
14. Roth, G.A., Johnson, C., Abajobir, A., Abd-Allah, F., Abera, S.F., Abyu, G., Ahmed, M., Aksut, B., Alam, T., Alam, K., et al.: Global, regional, and national burden of cardiovascular diseases for 10 causes, 1990 to 2015. *Journal of the American college of cardiology* **70**(1), 1–25 (2017)
 15. Shida, Y., Kumagai, S., Tsumura, R., Iwata, H.: Automated image acquisition of parasternal long-axis view with robotic echocardiography. *IEEE Robotics and Automation Letters* (2023)
 16. Shida, Y., Sugawara, M., Tsumura, R., Chiba, H., Uejima, T., Iwata, H.: Diagnostic posture control system for seated-style echocardiography robot. *International Journal of Computer Assisted Radiology and Surgery* **18**(5), 887–897 (2023)
 17. Song, P., Fang, Z., Wang, H., Cai, Y., Rahimi, K., Zhu, Y., Fowkes, F.G.R., Fowkes, F.J., Rudan, I.: Global and regional prevalence, burden, and risk factors for carotid atherosclerosis: a systematic review, meta-analysis, and modelling study. *The Lancet Global Health* **8**(5), e721–e729 (2020)
 18. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. p. 6000–6010. Curran Associates Inc. (2017)
 19. Yang, L., Jiang, H., Cai, R., Wang, Y., Song, S., Huang, G., Tian, Q.: Condensenet v2: Sparse feature reactivation for deep networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3569–3578 (2021)