



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.

Auxiliary Input in Training: Incorporating Catheter Features into Deep Learning Models for ECG-Free Dynamic Coronary Roadmapping

Yikang Liu¹[0000-0003-1069-1215], Lin Zhao¹, Eric Z. Chen¹, Xiao Chen¹, Terrence Chen¹, and Shanhui Sun¹

United Imaging Intelligence, Boston, MA, USA
shanhui.sun@uii-ai.com

Abstract. Dynamic coronary roadmapping is a technology that overlays the vessel maps (the "roadmap") extracted from an offline image sequence of X-ray angiography onto a live stream of X-ray fluoroscopy in real-time. It aims to offer navigational guidance for interventional surgeries without the need for repeated contrast agent injections, thereby reducing the risks associated with radiation exposure and kidney failure. The precision of the roadmaps is contingent upon the accurate alignment of angiographic and fluoroscopic images based on their cardiac phases, as well as precise catheter tip tracking. The former ensures the selection of a roadmap that closely matches the vessel shape in the current frame, while the latter uses catheter tips as reference points to adjust for translational motion between the roadmap and the present vessel tree. Training deep learning models for both tasks is challenging and underexplored. However, incorporating catheter features into the models could offer substantial benefits, given humans heavily rely on catheters to complete the tasks. To this end, we introduce a simple but effective method, auxiliary input in training (AIT), and demonstrate that it enhances model performance across both tasks, outperforming baseline methods in knowledge incorporation and transfer learning.

Keywords: Dynamic Coronary Roadmapping · Cardiac Phase Detection · Catheter Tip Tracking · Knowledge Incorporation

1 Introduction

X-ray angiographic image sequences are frequently used in interventional cardiology to assist with the navigation of devices in coronary arteries during procedures such as angioplasty or stent placement. However, it exposes patients to certain risks, including radiation exposure and potential kidney failure due to the use of contrast agents [16]. Dynamic coronary roadmapping (Fig. 1) is a technology aimed at minimizing the doses of radiation and contrast agent required. It works by superimposing a live contrast-free X-ray fluoroscopic image of the patient with a detailed coronary artery map (the "roadmap"), which is

obtained in advance through X-ray angiography. The roadmap is updated real-time accounting for the movement of the heart and the patient’s breathing. This approach effectively reduces the necessity for repeated angiography [16].




Fig. 1. The Workflow of Dynamic Coronary Roadmapping.

Dynamic coronary roadmapping can be implemented with three modules (Fig. 1): vessel segmentation, cardiac phase matching, and catheter tip detection and tracking. The vessel segmentation module extracts vessel masks from angiographic videos. As coronary vessels deform in cycles due to heartbeats, the cardiac phase matching module ensures the similarity of vessel shapes in the roadmap and live image by selecting the angiographic video frame that best matches the cardiac phase of the live fluoroscopic image. The catheter tip detection and tracking module locates the tip of the guiding catheter, which remains stationary in the vessel and serves as a reference for heart and breathing motion, in both the angiographic and fluoroscopic videos. Integrating these components, the system overlays the vessel mask from the selected angiographic frame onto the live image after compensating for translational motion. This overlay serves as the dynamic navigational roadmap.

This paper focuses on developing deep learning models for the cardiac phase matching and catheter tip tracking functions. Humans naturally leverage information from the catheter body to assist with both cardiac phase matching and catheter tip tracking, due to their challenging natures. In cardiac phase matching, the invisibility of vessels in fluoroscopy images makes it difficult to identify the cardiac phase. However, in both angiographic and fluoroscopic image sequences, catheter moves and deforms periodically with heartbeat. A person can compare catheter position and shape across two videos and match their cardiac phases. In catheter tip tracking, the challenges arise from the tip is frequently obscured by contrast agents or other devices. In this case, an individual can approximate their location by examining the shapes of the catheter in adjacent frames where the tips are visible. However, leveraging catheter information proves challenging for deep learning models. As we will demonstrate in this paper, models struggle to

converge or learn shortcut rules [11, 8] that fail to generalize in more challenging scenarios (e.g. the presence of thick contrast agents).

Training deep learning models to leverage catheter body features can be seen as an application of domain knowledge incorporation [21, 5] or transfer learning [24]. Typical approaches for medical images under this umbrella include fine-tuning a pretrained network [17, 19], multi-task learning [23, 1], custom neural network architecture [9, 20], teacher-student models [10, 12], generative style transfer [23], and constraining image features with attention maps [13, 4].

In the context of catheter tip tracking with deep learning, prior research has enhanced model capabilities with a dedicated branch for predicting the catheter mask [6, 14] and custom architectures with emphasis on catheter motion [6]. In contrast, incorporating catheter knowledge into cardiac phase matching models is less explored. Prior research has either used ECG signals to match cardiac phases [14] or analyzed the catheter curvature to assess the similarity between two frames’ cardiac phases [16]. These methods face limitations, as ECG signals may not always be accessible in clinical settings, and manually engineered features tend to be less reliable than machine learning approaches, especially when the catheter’s shape is complicated by foreshortening or overlaps with other objects. Ciusdel et al. [3] developed a deep learning model to identify cardiac phase, but its application is limited to angiographic images with visible vessels.

In this paper, we introduce Auxiliary Input in Training (AIT), a simple yet effective method that leverages catheter masks as auxiliary signals to incorporate catheter features into deep learning models for cardiac phase matching and catheter tip tracking, thereby enhancing representation quality and accelerating convergence. By appending the catheter mask as an additional input channel and gradually ablate it to a zero matrix during training, AIT facilitates the integration of catheter information. This approach has enabled us to create, to our knowledge, the first end-to-end deep learning framework capable of accurately matching coronary X-ray frames by cardiac phases. Moreover, we demonstrate that AIT also improves catheter tip tracking models, outperforming baseline methods in knowledge incorporation or transfer learning despite its simplicity.

2 Methodology

In this section, we first introduce the concept of AIT, then formally define the problems of cardiac phase matching and catheter tip tracking, and present our models and loss functions. More details are in the supplementary material.

AIT Suppose we want to train a deep learning model $f : x \mapsto y$ on a dataset $D = \{x, y\}$, but the model is hard to train due to the unsmoothness of the loss landscape or a complicated relationship between x and y . As a consequence, directly training f on D may take very long to converge or converge to minima that fail to generalize to the test dataset (e.g. shortcut learning [11, 8]). The key idea of AIT is to introduce an auxiliary input z to guide the training process so that the training converges faster and better representations are learned. More

specifically, we initially train $f(x; z)$ on a dataset $D_z = \{x, y, z\}$, and gradually ablate z throughout the training process so that the model’s reliance from the auxiliary information z is transferred back to the primary input x , which is the only variable needed for inference. The auxiliary input z , which though can be inferred from x , has a more obvious relationship to y (with respect to the network’s architecture or inductive bias) or is an indispensable step to infer y with x based on prior knowledge (e.g. humans rely on catheter shape consistency to identify cardiac phases and track tips).

In the applications of cardiac phase matching and catheter tip tracking, z is a binary catheter mask, which is concatenated with the input image x along the channel dimension. The ablation of z is done by adding Gaussian noise and concurrently decreasing the signal magnitude:

$$\tilde{z} = (1 - \alpha)((1 - \alpha)z + \alpha\mathcal{N}(0, 1))$$

, where \tilde{z} is the ablated z and α is a parameter that adjusts the intensity of the ablation. Throughout the training process, α is progressively increased from 0 to 1. At the point where $\alpha = 1$, z is entirely ablated into a matrix of zeros, thus becoming unnecessary for inference. In our default setting, we increment α by 0.1 at each step. After achieving network convergence at a given ablation level without any signs of overfitting, we escalate to the subsequent ablation level.

Cardiac Phase Matching Given a sequence of recorded cardiac angiographic images $\{I_i^A\}$ and a live fluoroscopic image stream $\{I_i^F\}$, the cardiac phase matching function finds the image in $\{I_i^A\}$ that best matches the cardiac phase of the current image I_i^F in real-time. We achieve this by using a CNN encoder to extract features from each image and a temporal neural network to infer temporal relations between image features, for which we experimented with both ConvLSTM and Transformer backbones to show the effectiveness of AIT on different architectures. The model outputs a feature vector v_i for each image and the cosine similarity between two feature vectors measures how close the cardiac phases of the corresponding images are.

The CNN-ConvLSTM model (hereafter denoted as CNN-C) comprises a series of alternating UNetResBlock [2] and ConvLSTM [18] layers, followed by a global max pooling and a fully connected layer to transform a 4D image tensor into a 1D feature vector. $\{I_i^A\}$ and $\{I_i^F\}$ are concatenated along the temporal dimension and sequentially fed into the model (Fig. S1).

The CNN-Transformer model (hereafter denoted as CNN-T) comprises a ResNet encoder and stacked attention layers. The outputs from the ResNet are flattened into 1D vectors before being passed to the attention layers. The attention layers run with self-attention for $\{I_i^A\}$. For real-time inference of $\{I_i^F\}$, the features extracted from the current fluoroscopic image are used as the query vector, while features from previous frames are used as key and value vectors.

Both models were trained with a triplet loss

$$\mathcal{L} = \max(-S(v^F, v_p^A) + S(v^F, v_n^A) + \epsilon, 0)$$

S denotes the cosine similarity. v^F is the feature vector of a fluoroscopic image I_i^F . v_p^A is the feature vector of an angiographic image I_i^A with the same cardiac phase, whereas v_n^A is the feature vector of one with a different cardiac phase. ϵ represents a positive margin, which we set to 0.8 in all experiments.

Catheter Tip Tracking Catheter tip tracking involves determining the coordinates (x, y) of a catheter tip within an image, based on one or multiple previous images and their corresponding tip coordinates. Similar to cardiac phase matching, we used both CNN-C and CNN-T models to show the effectiveness of AIT on different architectures.

The CNN-C model is a UNet with ConvLSTM layers in the skip connections. The input is a sequence of 3-channel tensors, with each channel containing the reference image (the image where tip location is known), the reference tip heatmap, and the current image to inference. The tensors are sequentially inputted into the network, which then outputs a tip heatmap for each frame.

The CNN-T model is similar to STARK-S[22]. It takes a template obtained by cropping the reference image around the tip and a search image, and passes them through a ResNet encoder. The encoder’s outputs are flattened, concatenated, and then forwarded to a transformer encoder. Subsequently, a trainable target query and the transformer encoder’s output are sent to a transformer decoder. The resulting output is further processed by a CNN head for heatmap regression. Similar to [6], three templates are used to adapt to variations in the tip’s appearance, which includes the initial template and two from the latest tracked tips, selected if their probabilities in the heatmaps exceeding a certain threshold (0.5). In AIT, catheter masks are concatenated with both the templates and search images.

Both models were trained with the L1 loss between the predicted heatmaps and the labels.

3 Experiments

Baseline Methods In addition to vanilla supervised learning, we also compared AIT with three other methods on both the cardiac phase matching and the catheter tip tracking tasks (hereafter denoted as CPM and CTT, respectively).

The first method (denoted as FT) fine-tunes a model trained for catheter segmentation. In the CPM task, the segmentation network builds upon the original CNN-C or CNN-T architecture (Section 2) and appends a CNN decoder after the last ConvLSTM/Attention layer for mask regression. In the CTT task, the segmentation models share the same structure as the tracking models.

The second method employs a multi-task learning (MTL) approach, wherein the models feature two branches simultaneously trained to predict both catheter masks and task-specific outputs. In the CPM task, the catheter segmentation branch uses the same CNN decoder structure in the FT method. In the CTT task, the catheter segmentation branch parallels the UNet decoder (in the CNN-C) or the heatmap regression head (in the CNN-T), having the same structures.

The third method uses a teacher-student (T-S) model, where the teacher network, pre-trained for catheter segmentation, guides the student (target) model with the maximum mean discrepancy (MMD) loss on the student’s and teacher’s features. The teacher models share the same architectures with the segmentation networks in FT. For fair comparison with MTL, the MMD loss is applied to the features before the segmentation branch, where networks start to use separate features for segmentation and target tasks.

In all experiments, we set the learning rates to 10^{-5} and used the Adam optimizer, with betas configured to 0.9 and 0.999.

Ablation Studies We investigated how AIT was affected by the percentage of data that trained with auxiliary inputs. This question is important since acquiring extra labels can be expensive. We run AIT on both tasks with partial inclusions (20%, 40%, 60%, and 80%) of catheter masks, where zero matrices were used as placeholders for the missing catheter masks.

Datasets and Evaluation Metrics All experiments were conducted using in-house data, obtained with institutional committee approval. The datasets for CPM and CTT contain 2483 pairs of angiographic and fluoroscopic videos (174228 frames in total) and 4098 videos (255432 frames) respectively, with frame rates equal to 7.5, 15, or 30 fps and image sizes range from 492×492 to 624×624 after normalizing pixel spacing to isotropic 0.2 mm. Cardiac phases, catheter tips, and catheter masks were manually labeled. The datasets were divided into training, validation, and testing sets with a 7:2:1 ratio.

The performance of models on the CPM task is assessed with matching accuracy, defined as the temporal distance between the predicted frame and the nearest ground-truth frame (multiple frames may exhibit the same cardiac phase due to the periodic nature of heartbeats). For the distance metric, we employed two units of measurement: frame counts and the percentage of a cardiac cycle. For example, if the distance is 2 frames within a cardiac cycle spanning 12 frames, the corresponding percentage is calculated as $1/6$ or 16.7%. To evaluate the performance on the CTT task, we deemed a tracking attempt successful if the distance between prediction and ground-truth did not exceed 2 mm. This threshold is consistent with the outside diameter of a guiding catheter [15] and meets the conventional requirements of roadmap accuracy [7, 16]. Using this criterion, we calculated the precision (P) and recall (R) for tracking. We also calculated the distance mean and standard deviation for true positives (TPs) and all cases.

4 Results and Discussion

Both CNN-C and CNN-T backbones have inference times (on an Nvidia V100 GPU) under 25 ms and 40 ms for the CPM and CTT tasks, respectively, satisfying the requirements for clinical application (15 fps).

Table 1. Performance on the Cardiac Phase Matching Task. α is the ablation strength on auxiliary inputs. **Bold** font indicates the best method (paired t-tests, $p < 0.05$) excluding intermediate AIT results ($\alpha \neq 1$).

Methods	CNN-ConvLSTM		CNN-Transformer	
	dist(frame) \downarrow	dist(%) \downarrow	dist(frame) \downarrow	dist(%) \downarrow
Vanilla	3.26 \pm 1.87	25.00 \pm 14.42	3.25 \pm 1.88	25.00 \pm 14.45
FT	3.19 \pm 1.85	24.01 \pm 13.81	3.21 \pm 1.84	24.49 \pm 14.13
MTL	3.25 \pm 1.87	24.99 \pm 14.43	3.30 \pm 1.91	24.46 \pm 14.10
T-S	2.35 \pm 1.35	17.42 \pm 10.08	2.44 \pm 1.41	19.05 \pm 10.98
AIT ($\alpha = 0$)	0.92 \pm 0.54	7.26 \pm 4.22	0.87 \pm 0.51	6.85 \pm 4.00
AIT ($\alpha = 0.5$)	0.96 \pm 0.57	7.69 \pm 4.45	0.97 \pm 0.59	7.75 \pm 4.59
AIT ($\alpha = 0.8$)	0.96 \pm 0.56	7.34 \pm 4.34	0.95 \pm 0.54	7.32 \pm 4.31
AIT (final)	0.89\pm0.51	7.01\pm4.08	0.85\pm0.51	6.72\pm3.98

The performance of all the methods on the CPM task, along with AIT performance at different ablation strengths, is shown in Table 1. We trained CNN-C and CNN-T models using AIT and other methods for 300 and 800 epochs, respectively. It was observed that both networks failed to converge under the vanilla, FT, and MTL strategies, as evidenced by a distance metric around 25%, indicating that the networks were essentially making random guesses. The T-S method produced predictions above random chance, yet its performance was significantly inferior to that of AIT (paired t-test, $p < 0.05$) and did not meet clinical standards. To better understand the underlying causes, we visualized the features after the second ConvLSTM block in the CNN-C backbone (Fig. 2), which was done by concatenating the magnitude of the first three principal components as RGB channels. It can be observed that AIT($\alpha = 1$) was able to learn strong features related to cardiac phase, located at catheter, wire, heart contour (circled in Fig. 2), whereas other methods learned much weaker features. These observations indicate that AIT is able to facilitate model convergence by incorporating catheter features.

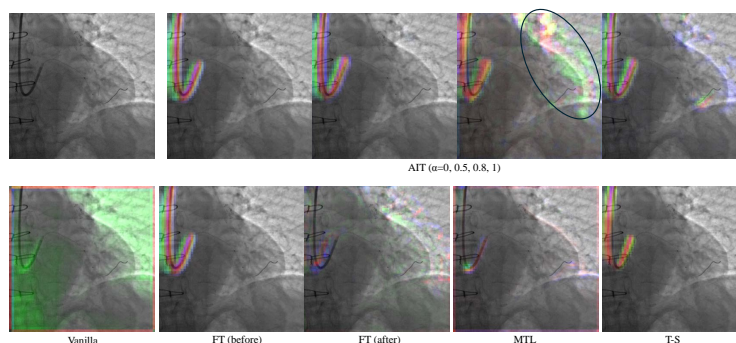


Fig. 2. Visualization of Features Learned by Different Methods in the CPM task. First row (left to right): original image and AIT ($\alpha=0, 0.5, 0.8, 1$). Second row (left to right): vanilla supervised learning, FT (before and after fine-tuning), MTL, and T-S.

Additionally, it is observed that AIT’s performance dip initially with the start of auxiliary input ablation but improved near the end of the ablation schedule. Together with the feature maps in Fig. 2, it suggests that the model initially relied on catheter features, which became weaker due to mask ablation. Subsequently, the model adapted by leveraging alternative features to offset the weakened catheter features, resulting in more robust predictions than those based solely on catheter information.

AIT achieved sub-frame average accuracy with both networks, indicating its potential to replace ECG-based phase matching methods. However, to assess its practical clinical use compared to the ECG-based method, the accuracy needs further evaluation by measuring the distance between the overlaid vessels and the interventional devices (e.g. guide wires) within the vessels, as demonstrated in [14]. This practical assessment will be included in other future work.

Table 2. Performance on the Catheter Tip Tracking Task. α is the ablation strength on auxiliary inputs. **Bold** font indicates the best method (paired t-tests, $p < 0.05$) excluding intermediate AIT results ($\alpha \neq 1$).

Methods	CNN-ConvLSTM				CNN-Transformer			
	P(%) \uparrow	R(%) \uparrow	dist(TP) \downarrow	dist(all) \downarrow	P(%) \uparrow	R(%) \uparrow	dist(TP) \downarrow	dist(all) \downarrow
Vanilla	91.3	94.5	0.90 \pm 0.63	1.67 \pm 2.88	92.5	95.2	0.92 \pm 0.62	1.45 \pm 2.24
FT	92.6	95.9	0.89 \pm 0.66	1.41 \pm 2.16	93.0	96.1	0.91 \pm 0.57	1.33 \pm 1.81
MTL	90.5	94.0	0.91 \pm 0.79	1.75 \pm 2.98	91.6	94.1	0.93 \pm 0.65	1.56 \pm 2.41
T-S	94.7	97.4	0.89 \pm 0.52	1.26 \pm 1.87	95.5	98.0	0.91 \pm 0.57	1.13 \pm 1.36
AIT ($\alpha = 0$)	99.7	99.8	0.87 \pm 0.49	0.89 \pm 0.62	99.7	99.8	0.90 \pm 0.51	0.91 \pm 0.61
AIT ($\alpha = 0.5$)	98.2	99.0	0.88 \pm 0.53	0.96 \pm 0.88	97.8	98.7	0.91 \pm 0.54	1.04 \pm 1.08
AIT ($\alpha = 0.8$)	97.7	98.6	0.88 \pm 0.51	1.00 \pm 1.04	97.2	98.5	0.91 \pm 0.53	1.08 \pm 1.22
AIT (final)	96.9	97.6	0.88\pm0.50	1.10\pm1.48	97.1	97.8	0.90\pm0.54	1.08\pm1.23

In the CTT task, AIT achieved the highest performance using the CNN-C backbone and ranked as either the best or second-best using the CNN-T backbone (Table 2). Generally, MTL negatively impacted performance (consistent with [6]), while other approaches contributed positively. AIT performance decreased with the progression of auxiliary input ablation, due to ease of inferring catheter tip positions from clean catheter masks. However, it still significantly outperformed the vanilla supervised learning method (paired t-tests, $p < 0.05$), indicating that catheter features were effectively incorporated into the models even in the absence of catheter masks as inputs, thereby improving performance. Furthermore, it should be noted that vanilla CNN-T resembles ConTrack without multitask, flow, or multi-templates (refer to Table 2 in [6]). AIT improved the average tracking distance by 25.5% to 1.08 mm, compared to ConTrack’s improvement by 24.9% to 1.63 mm (Table 2 in [6]). Although these models were trained and tested on different datasets, the results suggest that AIT’s approach of integrating catheter body information might achieve performance levels similar to those of specifically designed neural networks.

Finally, we demonstrate in the ablation study that even a small percentage of auxiliary inputs can yield significant benefits, especially in the CPM task, as shown in Fig. 3.




Fig. 3. Ablation Study

5 Conclusion

This study introduces a straightforward yet effective approach, Auxiliary Input in Training (AIT), for incorporating prior knowledge into deep learning models. We applied this method to train models for cardiac phase matching and catheter tip tracking—two demanding tasks in dynamic coronary roadmapping—and showcased its efficacy in enhancing performance across both tasks. Despite its simplicity, AIT’s superior performance stands out in comparison to other techniques, underscoring its value in complex medical imaging tasks.

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

References

1. Bakalo, R., Ben-Ari, R., Goldberger, J.: Classification and detection in mammograms with weak supervision via dual branch deep neural net. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). pp. 1905–1909. IEEE (2019)
2. Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., et al.: Monai: An open-source framework for deep learning in healthcare. arXiv preprint arXiv:2211.02701 (2022)
3. Ciusdel, C., Turcea, A., Puiu, A., Itu, L., Calmac, L., Weiss, E., Margineanu, C., Badila, E., Berger, M., Redel, T., et al.: Deep neural networks for ecg-free cardiac phase and end-diastolic frame detection on coronary angiographies. *Computerized Medical Imaging and Graphics* **84**, 101749 (2020)
4. Cui, H., Xu, Y., Li, W., Wang, L., Duh, H.: Collaborative learning of cross-channel clinical attention for radiotherapy-related esophageal fistula prediction from ct. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part I* 23. pp. 212–220. Springer (2020)
5. Dash, T., Chitlangia, S., Ahuja, A., Srinivasan, A.: A review of some techniques for inclusion of domain-knowledge into deep neural networks. *Scientific Reports* **12**(1), 1040 (2022)
6. Demoustier, M., Zhang, Y., Narasimha Murthy, V., Ghesu, F.C., Comaniciu, D.: Contrack: Contextual transformer for device tracking in x-ray. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 679–688. Springer (2023)

7. Faranesh, A.Z., Kellman, P., Ratnayaka, K., Lederman, R.J.: Integration of cardiac and respiratory motion into mri roadmaps fused with x-ray. *Medical physics* **40**(3), 032302 (2013)
8. Geirhos, R., Jacobsen, J.H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., Wichmann, F.A.: Shortcut learning in deep neural networks. *Nature Machine Intelligence* **2**(11), 665–673 (2020)
9. Guan, Q., Huang, Y., Zhong, Z., Zheng, Z., Zheng, L., Yang, Y.: Diagnose like a radiologist: Attention guided convolutional neural network for thorax disease classification. *arXiv preprint arXiv:1801.09927* (2018)
10. Han, X., Wang, J., Zhou, W., Chang, C., Ying, S., Shi, J.: Deep doubly supervised transfer network for diagnosis of breast cancer with imbalanced ultrasound imaging modalities. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VI* 23. pp. 141–149. Springer (2020)
11. Hermann, K.L., Mobahi, H., Fel, T., Mozer, M.C.: On the foundations of shortcut learning. *arXiv preprint arXiv:2310.16228* (2023)
12. Hu, C., Li, X., Liu, D., Chen, X., Wang, J., Liu, X.: Teacher-student architecture for knowledge learning: A survey. *arXiv preprint arXiv:2210.17332* (2022)
13. Li, L., Xu, M., Wang, X., Jiang, L., Liu, H.: Attention based glaucoma detection: A large-scale database and cnn model. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10571–10580 (2019)
14. Ma, H., Smal, I., Daemen, J., van Walsum, T.: Dynamic coronary roadmapping via catheter tip tracking in x-ray fluoroscopy with deep learning based bayesian filtering. *Medical image analysis* **61**, 101634 (2020)
15. Ojha, V., Raju, S.N., Deshpande, A., Ganga, K.P., Kumar, S.: Catheters in vascular interventional radiology: an illustrated review. *Diagnostic and interventional radiology (Ankara, Turkey)* **29**(1), 138–145 (2023)
16. Piayda, K., Kleinebrecht, L., Afzal, S., Bullens, R., Ter Horst, I., Polzin, A., Veulemans, V., Dannenberg, L., Wimmer, A.C., Jung, C., et al.: Dynamic coronary roadmapping during percutaneous coronary intervention: a feasibility study. *European journal of medical research* **23**, 1–7 (2018)
17. Samala, R.K., Chan, H.P., Hadjiiski, L., Helvie, M.A., Richter, C.D., Cha, K.H.: Breast cancer diagnosis in digital breast tomosynthesis: Effects of training sample size on multi-stage transfer learning using deep neural nets. *IEEE Transactions on Medical Imaging* **38**(3), 686–696 (2019). <https://doi.org/10.1109/TMI.2018.2870343>
18. Shi, X., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.K., Woo, W.c.: Convolutional lstm network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems* **28** (2015)
19. Shin, H.C., Roth, H.R., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D., Summers, R.M.: Deep convolutional neural networks for computer-aided detection: Cnn architectures, dataset characteristics and transfer learning. *IEEE transactions on medical imaging* **35**(5), 1285–1298 (2016)
20. Wang, X., Peng, Y., Lu, L., Lu, Z., Summers, R.M.: Tienet: Text-image embedding network for common thorax disease classification and reporting in chest x-rays. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 9049–9058 (2018)
21. Xie, X., Niu, J., Liu, X., Chen, Z., Tang, S., Yu, S.: A survey on incorporating domain knowledge into deep learning for medical image analysis. *Medical Image Analysis* **69**, 101985 (2021)

22. Yan, B., Peng, H., Fu, J., Wang, D., Lu, H.: Learning spatio-temporal transformer for visual tracking. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10448–10457 (2021)
23. Zhao, J., Li, D., Kassam, Z., Howey, J., Chong, J., Chen, B., Li, S.: Tripartitegan: Synthesizing liver contrast-enhanced mri to improve tumor detection. *Medical image analysis* **63**, 101667 (2020)
24. Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., He, Q.: A comprehensive survey on transfer learning. *Proceedings of the IEEE* **109**(1), 43–76 (2020)