

Context-guided Continual Reinforcement Learning for Landmark Detection with Incomplete Data

Kaiwen Wan¹, Bomin Wang¹, Fuping Wu², Haiyu Gong¹, and Xiahai Zhuang¹

✉

¹ School of Data Science, Fudan University, Shanghai, 200433, China

² Nuffield Department of Population Health, University of Oxford, Oxford, UK
{16210180100,zxh}@fudan.edu.cn, {bmwang21,hygong23}@m.fudan.edu.cn,
Fuping.Wu@ndph.ox.ac.uk

Abstract. Existing landmark detection methods are primarily designed for centralized learning scenarios where all training data and labels are complete and available throughout the entire training phase. In real-world scenarios, training data may be collected sequentially, covering only part of the region of interest or providing incomplete landmark labels. In this work, we propose a novel continual reinforcement learning framework to tackle this complex situation in landmark detection. To handle the increasing number of landmark targets during training, we introduce a Q-learning network that takes both observations and prompts as input. The prompts are stored in a buffer and utilized to guide the prediction for each landmark, enabling our method to adapt to the intricacies of the data collection process. We validate our approach on two datasets: the RSNA-PBA dataset, representing scenarios with complete images and incomplete labels, and the WB-DXA dataset, representing situations where both images and labels are incomplete. The results demonstrate the effectiveness of the proposed method in landmark detection tasks with complex data structures. The source code will be available from <https://github.com/kevinwolcano/CgCRL>.

Keywords: Landmark Detection · Reinforcement Learning · Continual Learning · Incomplete Data.

1 Introduction

Landmark detection is a foundational task in medical image analysis, providing crucial information for depicting morphological features, pre-localizing regions of interest (ROIs), and facilitating downstream tasks like segmentation and registration [11,12]. With the evolution of deep learning, automatic landmark detection algorithms have progressed rapidly, diminishing the reliance on time-consuming and labor-intensive manual annotations by experts.

However, most existing landmark detection algorithms are built on the assumptions of complete and centralized training data. Under such settings, all




Fig. 1. Illustration of scenarios involving sequential incomplete data in our work. Here, I^c (L^c) and I^{inc} (L^{inc}) respectively refer to complete images (labels) and incomplete images (labels); the dots ($\bullet$) and the crosses ($\times$) represent known landmarks and missing ones, respectively.

training data are available throughout the training phase, and both images and the labels are complete. In clinical scenarios, these assumptions are often not met due to incomplete and sequential data collection. Specifically, in real-world clinical practice, incomplete data may arise due to researchers focusing on different subjects, resulting in varying ROIs in images. Additionally, the high cost of labeling may deter annotators from tagging every label. Both result in non-standard datasets. Many deep learning models require fixed sizes for input and output images, making them unsuitable for handling such incomplete data. For instance, Deep Multi-agent Q-learning (DeepMaQ) addresses multi-target landmark detection in incomplete images using multiple agents [2]. However, as the number of landmarks increases, the network parameters of DeepMaQ also increase. Moreover, DeepMaQ must be pre-given the number of landmarks, making it ineffective for coping with the emergence of new targets. Sequential data collection can lead to the problem of catastrophic forgetting in deep learning models, where knowledge learned from old data may gradually be forgotten when the model adapts to new data. We illustrate the problem considered in Fig. 1. At present, there has been relevant work in response to these two challenges [3,4]. However, currently, there are no works that take into account the existence of both simultaneously.

In this work, we propose a Context-guided Continual Reinforcement Learning model (CgCRL) for landmark detection with sequential incomplete training data. The model comprises a Context-guided Multi-target Q-learning (CgMtQ) network and a Context Memory Replay Mechanism (CMRM). The CgMtQ network utilizes a single agent to explore trajectories leading to target landmarks. Besides the agent’s observations, a context associated with the target landmark is input into the network as a prompt to guide the agent to perform actions. Simultaneously, with the assistance of CMRM, we maintain a prompt library containing representative patches for target regions during training. The collaborative functioning of these two components enhances the robustness of the landmark detection model and mitigates the issue of catastrophic forgetting.

We summarize our contributions as follows: (1) We undertake the pioneering investigation into landmark detection using sequential and incomplete training data. (2) We introduce a novel Continual Reinforcement Learning (CRL) framework tailored for landmark detection with sequential incomplete data. Our framework maintains a consistent network size even when faced with a varying and expanding number of target landmarks and effectively tackles the challenge of catastrophic forgetting. (3) We demonstrate the effectiveness of our proposed method through experiments conducted on two datasets.

2 Methodology

Let $\mathcal{T} = \{\mathcal{T}_k | k=1, \dots, K\}$ be K sequential training datasets, where $\mathcal{T}_k = \{(I_{k,i}, Y_{k,i}) | i \in \{1, \dots, N_k\}\}$ with $I_{k,i}$ and $Y_{k,i} = \{y_{k,i}^{[j]} \in \mathbb{R}^2 | j \in \mathcal{J}_k\}$ being an image and its corresponding landmark locations (coordinates), and $\mathcal{J}_k$ the index set. Note that image $I_{k,i}$ could be incomplete, covering a partial ROIs and varying by task. Accordingly, the set of target indices $\mathcal{J}_k$ for the k -th task could also be a subset of the complete index set $\mathcal{J}$, i.e., $\mathcal{J}_k \subset \mathcal{J}$. For simplification, we refer to a fully visible image as I^c and a partially visible one as I^{inc} . In the context of continual learning, datasets $\{\mathcal{T}_k\}$ are collected sequentially, and previously collected data may not be completely retrievable in later stages of training due to the storage limit. The goal here is to develop a model for detecting landmarks, denoted by Q , capable of identifying all potential landmarks in a complete image.

To tackle the complex data structure during training, we develop a context-guided continual reinforcement learning framework. As illustrated in Fig. 2, we propose a Context-Guided Multi-Target Q-Network (CgMtQ), which takes both observation patches and instructional prompts as input. This approach can accommodate the increasing target indexes without changing the network structure (refer to Section 2.1 for details). Moreover, as elaborated in Section 2.2, a Context Memory Replay Mechanism (CMRM) is introduced to maintain a prompt library. As training progresses, CMRM actively updates, storing representative patches for all emerged target indices within a limited memory space. These patches serve as prompts to predict corresponding targets, as well as facilitate memory replay of previous tasks, thereby mitigating knowledge forgetting.

2.1 Context-guided Multi-target Q-learning

CgMtQ utilizes a singular agent to identify all landmarks, differentiating each landmark through the use of distinct, informative prompts. As shown in Fig. 2, throughout the training phase, for every targeted landmark, a specific prompt is selected from the prompt library. This prompt, alongside an observation patch, is then input into CgMtQ to undergo reinforcement learning. This methodology allows for the accommodation of new targets without changing the network structure. We detail the basic elements of Reinforcement Learning in the following.




Fig. 2. Illustration of the Context-guided Continual Reinforcement Learning (CgCRL) consisting of CgMtQ and CMRM.

State: The current state S is defined as $S = \{obs, \Gamma_j^p\}$, where obs is the current observation of the agent and Γ_j^p is the p^{th} prompt for j^{th} target selected from the prompt library Γ . The observation can be obtained by cropping an image around the current location of the agent with a given length. Details about the prompt library Γ will be given in Sec 2.2.

Action: The current action A is defined as the agent’s movement in different directions in the image. In 2D images, it moves in four directions: up, down, left, and right by a fixed step length, i.e., $A \in \{M_{\text{left}}, M_{\text{right}}, M_{\text{up}}, M_{\text{down}}\}$. We also adopt a multi-scale strategy, that is, starting with a larger step length and reducing the step length once the agent converges to the target point.

Reward: The reward R after acting A is defined as $R = D(y, y^{[t-1]}) - D(y, y^{[t]})$, where y is the target landmark location, $y^{[t-1]}$ and $y^{[t]}$ are the agent’s locations before and after taking action at t -th step, respectively. For the distance measurement D , we employ the Euclidean distance. Hence, R is positive if the agent’s distance to the target point becomes smaller after acting.

CgMtQ learns the Q-value function. Details of the network structure can be found in the Supplementary Material. Based on the Bellman equation [5], the loss function for training CgMtQ can be formulated as follows,

$$\mathcal{L}_Q(\omega) = \sum_{i=1}^N \mathbb{E}_{S,A,R,S'} \left(R - Q(S, A; \omega) + \gamma \max_{A'} Q(S', A'; \omega) \right)^2, \quad (1)$$

where S' and A' are the new state after taking the action A and the new action taken in the new state S' , respectively; $Q(\cdot, \cdot; \omega)$ represents the Q-value function and γ is the discount factor.

2.2 Context Memory Replay Mechanism

The prompt library maintained throughout training is designed to enable a single agent to adapt to an increasing number of targets, and is also utilized for memory replay in the context of continual learning.

Contextual Prompt The prompt library Γ is built by collecting representative patches for the target regions from the training data. Specifically, we have $\Gamma = \{\Gamma_j | j \in \mathcal{J}\}$ with each Γ_j a sub-library of prompts for j^{th} target, denoted by $\Gamma_j = \{\Gamma_j^p | p \in \{1, \dots, \mathcal{N}_j\}\}$. For any given index j , we extract patches with a given size centered on the relevant landmark from images that include it as prompts. When new images and labels emerge, new prompts will be added to the prompt library. Due to the memory limit, when the storage is full and new prompts are obtained, we discard the redundant old prompts based on their distances to the corresponding clustering centers to accommodate the new ones. During both training and testing, the appropriate prompt Γ_j^p for targeting the j^{th} landmark is retrieved from Γ_j and fed into the model to guide the prediction process. The prompt selection method and other details will be outlined in Section 2.3.

Replay Strategy The prompt library, Γ , contains representative exemplars for previously emerged targets and thus can be utilized to avoid knowledge-forgetting with memory replay strategy [13]. Specifically, during the training phase for a given task $\mathcal{T}_k$, if annotations for the j^{th} landmark are absent, *i.e.*, $j \notin \mathcal{J}_k$, it is possible that the landmark may either be present within or absent from the image I_k . These situations are addressed differently: (1) If the image I_k contains the landmark, the model, CgMtQ, generates predictions for the landmark’s location using each exemplar Γ_j^p within Γ_j for $p \in \{1, \dots, \mathcal{N}_j\}$, upon a pseudo-prompt patch, denoted $\Gamma_j^{\hat{p}}$, is extracted from I_k . We then calculate the prediction confidence $\mathbb{P}$ as follows:

$$\mathbb{P} = \max_p \mathcal{S}(\Gamma_j^{\hat{p}}, \Gamma_j^p), \quad (2)$$

where $\mathcal{S}$ represents a measure of image similarity, with Peak Signal-to-Noise Ratio (PSNR) utilized in this work. If $\mathbb{P}$ is larger than a predefined threshold, the corresponding prediction is considered reliable and added into Γ_j . (2) If the j^{th} landmark locates outside of I_k , a random prompt Γ_j^i is selected from Γ_j and embedded into I_k at a random coordinate y' to generate an augmented image I_k^{new} . Then, the new image-label pair $\{I_k^{new}, y'\}$ will be added to the current training set, as shown in Fig. 2. This approach not only enriches the training data with diverse examples but also fortifies the model’s ability to recall and apply previously acquired knowledge to new tasks.

2.3 Implementation Details

Training Stage During the training process with CgMtQ, the initial step involves choosing an appropriate prompt for every target landmark. A suitable

representative prompt for the target’s current image context is important for effective learning of the knowledge. Hence, similar to Eq. (2), the prompts in the prompt library will be compared with the cropped area centered at the target location, and the one with the highest similarity will be selected. Then, the agent will start with a random position and act under the guidance of the network until it reaches the terminal state. Subsequently, the agent resets its state and begins to search for the next target landmark. In addition, we employ several techniques to expedite network convergence, including the ϵ -greedy algorithm, parameter replay, frozen networks, and multi-scale strategy. Please refer to Algorithm 1 showing the pseudocode in the Supplementary Material for more details.

Testing Stage During the testing phase, for each landmark index j , CgMtQ generates predictions using all prompts from Γ_j , resulting in multiple potential landmark predictions. Similar to the patch selection strategy in Eq. (2), we select the prediction exhibiting the highest similarity as the final output. Please refer to the Supplementary Material for more details of the testing algorithm.

3 Experiment

The proposed CgCRL was tested across various scenarios using two distinct datasets: the RSNA Pediatric Bone Age Challenge (2017) dataset (referred to as RSNA-PBA) [15] and the whole-body dual-energy X-ray absorptiometry dataset (referred to as WB-DXA) [8,14], as depicted in Fig.1.

3.1 Dataset

(1) **RSNA-PBA** comprises 150 X-ray images of children’s hands, each down-sampled to a resolution of 1319×1664 pixels. Expert annotations highlighted 17 keypoints at the joints in each image, as illustrated in Fig. 1 (a). To create a sequentially labeled dataset that mimics incomplete labeling, we divided the dataset into 6 subsets, each with 20 images for training and the remaining 30 images for validation. Each of the 6 training subsets was assigned labels for a distinct set of 3, 3, 3, 3, 3 and 2 landmarks, covering all 17 target points. This setup resulted in six distinct sub-tasks characterized by partial labels. The evaluation of the methodologies was conducted on these sub-tasks, utilizing datasets with complete images (I^c) and incomplete labels (L^{inc}).

(2) **WB-DXA** consists of 99 DXA original images, each annotated with 40 key contour points as illustrated in Fig. 1 (b). This dataset was split into four parts: three training subsets each containing 20 images, and a fourth testing subset with the remaining 39 images. Our evaluation focused on scenarios featuring both incomplete images (I^{inc}) and incomplete labels (L^{inc}). To create conditions representing partial image availability, the images in the three training subsets underwent vertical cropping to exclude a predetermined portion, set at a missing proportion (mp) of 40%. Specifically, for the first subtask, we preserved the upper 40% of each image; for the second, the central 40%; and for the third,

Table 1. Comparative results of CgMtQ with different settings on RSNA-PBA dataset. Noted that ResNet and DensNet were trained with complete labels. ADE Values are reported here in *mm*.

Training: $I^c + L^{inc}$ (sequential)				Target: $I^c \rightarrow L^c$				
Training		Test						
		I^{inc}						I^c
	$\mathcal{T}_1$	$\mathcal{T}_2$	$\mathcal{T}_3$	$\mathcal{T}_4$	$\mathcal{T}_5$	$\mathcal{T}_6$		
CgMtQ+CMRM (CgCRL)	$\mathcal{T}_1$	2.16±4.57	N/A	N/A	N/A	N/A	N/A	N/A
	$\mathcal{T}_2$	1.97± 3.24	1.24±1.94	N/A	N/A	N/A	N/A	N/A
	$\mathcal{T}_3$	2.16± 3.65	1.60±2.38	0.95±0.717	N/A	N/A	N/A	N/A
	$\mathcal{T}_4$	2.53±4.62	0.92±0.964	1.82±3.09	1.91±4.41	N/A	N/A	N/A
	$\mathcal{T}_5$	1.43±0.881	2.67±5.74	1.51±2.34	4.93±8.74	2.11±4.97	N/A	N/A
	$\mathcal{T}_6$	3.64±7.05	2.52±4.43	1.22±1.25	1.41±1.78	2.25±3.47	1.87±1.18	2.17±2.11
CgMtQ	$\mathcal{T}_6$	88.8±7.96	103±10.9	6 113±12.6	117±12.8	71.7±8.36	1.29±0.79	87.2±8.28
CgMtQ+EWC	$\mathcal{T}_6$	91.7±9.09	98.9±7.45	110 ±9.14	113±11.7	72.7±8.29	1.39±1.36	86.0±6.87
CgMtQ+GPM	$\mathcal{T}_6$	87.5±7.73	108±7.64	108±8.05	112±11.4	71.3±8.03	1.26±0.41	84.9±6.46
Training: $I^c + L^{inc}$ (centralized)				Target: $I^c \rightarrow L^c$				
CgMtQ	2.42±4.31	1.76±2.99	1.45±2.30	1.78±4.43	1.40±1.35	1.82±1.88	1.77±1.68	
DeepMaQ	3.68±7.31	0.73±0.34	0.69±0.22	2.15±5.89	0.965±0.37	1.28±0.69	1.60±1.62	
ResNet(L^c)	6.03±3.78	5.24±2.66	5.37±2.46	6.47±2.99	6.34±2.96	6.25±3.03	5.93±2.08	
DensNet(L^c)	5.49±3.62	4.14±2.13	3.80±2.06	5.55±3.33	4.99±2.92	4.68±2.95	4.78±2.05	

the lower 40%. Note that in the testing phase of each sub-task within the continual learning process, the testing images were similarly cropped to match the respective training sub-task.

3.2 Experiment Results

The RSNA-PBA dataset was used to test our approach under conditions where images were fully available but labels were partially available in a sequential manner. Conversely, the WB-DXA dataset was chosen to assess performance in a more challenging scenario where both images and labels were partially available. To benchmark our method, experiments were also performed with datasets compiled in a centralized fashion, meaning all training data for the various sub-tasks were accessible at the same time.

For comparative analysis, we selected our proposed method, CgMtQ, alongside DeepMaQ [2], a ResNet-based keypoint detection technique for evaluation. Furthermore, to establish baseline performances, DeepMaQ was trained with fully complete images and labels utilizing both ResNet [10] and DenseNet architectures, with their outcomes serving as reference points for comparison. To measure prediction accuracy, Average Distance Error (ADE) was employed.

Results on RSNA-PBA dataset Table 1 summarizes reveals that the proposed CgCRL method delivers impressive outcomes on the hand dataset within a context of sequential incomplete data, which obtained an average error of 2.17 mm across all landmarks in the test dataset. This performance was on par with the centralized training outcomes of CgMtQ (1.77 mm) and DeepMaQ (1.60

Table 2. Comparative results of CgCRL with different settings on WB-DXA dataset, ADE values are reported here in *mm*. Noted CgMtQ and DeepMaQ need an artificial initial position when testing in I^c .

Training: $I^{inc} + L^{inc}$ (sequential)		Target: $I^c \rightarrow L^c$			
Training		Test			
		I^{inc}			I^c
		$\mathcal{T}_1$	$\mathcal{T}_2$	$\mathcal{T}_3$	
CgMtQ+CMRM (CgCRL)	$\mathcal{T}_1$	0.722±0.166	N/A	N/A	N/A
	$\mathcal{T}_2$	1.71±1.26	0.500±0.182	N/A	N/A
	$\mathcal{T}_3$	2.52±2.13	1.76±0.948	1.24±0.37	2.32±2.02
CgMtQ	$\mathcal{T}_3$	288±23.9	191±12.2	1.16±0.548	169±15.9
CgMtQ+EWC	$\mathcal{T}_3$	154±44.2	84.9±20.6	1.20±0.401	75.8±24.5
CgMtQ+ GPM	$\mathcal{T}_3$	252±21.3	176±19.4	1.19±0.35	153±13.7
Training: $I^{inc} + L^{inc}$ (centralized)		Target: $I^c \rightarrow L^c$			
CgMtQ		0.836±0.277	0.683±0.445	0.594±0.232	0.753±0.422
DeepMaQ		0.632 ±0.378	0.732±0.392	0.549±0.232	0.679±0.328
ResNet(I^c, L^c)		N/A	N/A	N/A	2.13±0.823
DensNet(I^c, L^c)		N/A	N/A	N/A	1.83±0.912

mm), and it outperformed the centralized training results of both ResNet and DenseNet, which were trained with fully complete labels.

When integrating GPM and EWC, CgMtQ almost fails to keep the knowledge learning from $\mathcal{T}_1$ to $\mathcal{T}_5$, which causes large ADE and most of the predictions to be abnormal predictions (predictions far from the true landmark). Details can be seen in Supplementary Material. This demonstrated the ineffectiveness of both EWC and GPM in preventing knowledge forgetting in our setting. The sequential setting results for CgMtQ exhibited significant forgetting, which demonstrated the effectiveness of CMRM in defying forgetting. Notably, for task $\mathcal{T}_4$, CgCRL recorded a lower error after the training of all tasks compared to post $\mathcal{T}_4$ completion. This suggests that CgCRL not only addresses catastrophic forgetting effectively but also leverages new data to enhance the optimization of contextual cues further.

Results on the WB-DXA dataset Table 2 shows that CgCRL achieved satisfactory results in the full-body dataset. In the context of continual learning with incomplete data, CgMtQ achieved an ADE of 2.32 cm on complete test images. Unlike the hand dataset, the performance on old subtasks progressively deteriorated during the training process. For example, for subtask $\mathcal{T}_1$, the initial ADE of CgMtQ was 0.722 cm. When the data of all sub-tasks were available, CgMtQ achieved an ADE of 0.753 cm, slightly larger than DeepMaQ’s 0.679 cm. For multi-object detection problems involving incomplete images, DeepMaQ was a better solution than CgMtQ. The reason could be that the context from incomplete images may differ from that of complete images due to cropping, which could disrupt guidance.

In contrast, ResNet achieved a prediction of 2.13 cm using complete data. Additionally, the ADE value of CgMtQ was 0.19 cm lower than that in the case

of learning with incomplete data, indicating that incomplete acquisition poses a more challenging training scenario for CgMtQ. When integrating GPM and EWC, CgMtQ also fails to keep memory learnt from tasks before while CMRM successes. This suggests that CgCRL can also deal with catastrophic forgetting when both images and labels are incomplete.

4 Conclusion

In this study, we introduce a novel reinforcement learning model for landmark detection in scenarios involving sequential and incomplete data. The proposed method enables the learning of an increasing number of landmarks during training by utilizing representative prompts selected from the prompt library, which is maintained during training. We validate the proposed method on two datasets, including the RSNA-PBA dataset for the scenario with complete images and incomplete labels, and the WB-DXA dataset for the scenario where both images and labels are incomplete. The results demonstrate the effectiveness of CgCRL in landmark detection with complex data structures.

Acknowledgments. This work was funded by the National Natural Science Foundation of China (grant No. 62372115, 61971142 and 62111530195).

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

References

1. Perez-Rua, J., Zhu, X., Hospedales, T. & Xiang, T. Incremental few-shot object detection. *Proceedings Of The IEEE/CVF Conference On Computer Vision And Pattern Recognition*. pp. 13846-13855 (2020)
2. Wan, K., Li, L., Jia, D., Gao, S., Qian, W., Wu, Y., Lin, H., Mu, X., Gao, X., Wang, S. & Others Multi-target landmark detection with incomplete images via reinforcement learning and shape prior embedding. *Medical Image Analysis*. pp. 102875 (2023)
3. Kang, S., Jeon, K., Kang, S. & Lee, S. 3D Cephalometric landmark detection by multiple stage deep reinforcement learning. *Scientific Reports*. **11**, 17509 (2021,9,1)
4. Shieh, J., Haq, Q., Haq, M., Karam, S., Chondro, P., Gao, D. & Ruan, S. Continual Learning Strategy in One-Stage Object Detection Framework Based on Experience Replay for Autonomous Driving Vehicle. *Sensors*. **20** (2020)
5. Alansary, A., Oktay, O., Li, Y., Le Folgoc, L., Hou, B., Vaillant, G., Kamnitsas, K., Vlontzos, A., Glocker, B., Kainz, B. & Others Evaluating reinforcement learning agents for anatomical landmark detection. *Medical Image Analysis*. **53** pp. 156-164 (2019)
6. Escobar, M., González, C., Torres, F., Daza, L., Triana, G. & Arbeláez, P. Hand pose estimation for pediatric bone age assessment. *Medical Image Computing And Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part VI 22*. pp. 531-539 (2019)

7. Ham, G. & Oh, K. Learning Spatial Configuration Feature for Landmark Localization in Hand X-rays. *Electronics*. **12** (2023),
8. Gao, X., Hofman, A., Hu, Y., Lin, H., Zhu, C. Jeekel, J., Xue-Jin, Ji-Wang, Gao, J., Yin, Y. & Zhao, N. The Shanghai Changfeng Study: a community-based prospective cohort study of chronic diseases among middle-aged and elderly: objectives and design. *European Journal Of Epidemiology*. **25** pp. 885-893 (2010)
9. Shepherd, J., Ng, B., Fan, B., Schwartz, A., Cawthon, P., Cummings, S., Kritchevsky, S., Nevitt, M., Santanasto, A. & Cootes, T. Modeling the shape and composition of the human body using dual energy X-ray absorptiometry images. *PLOS ONE*. **12** pp. e0175857 (2017,4)
10. Terada, T., Chen, Y. & Kimura, R. 3D Facial Landmark Detection Using Deep Convolutional Neural Networks. *2018 14th International Conference On Natural Computation, Fuzzy Systems And Knowledge Discovery (ICNC-FSKD)*. pp. 390-393 (2018)
11. Zhan, Y., Peng, Z., Hermosillo, G. & Zhou, X. Chapter 3 - Robust Multi-Landmark Detection Based on Information Theoretic Scheduling. *Medical Image Recognition, Segmentation And Parsing*. pp. 45-70 (2016)
12. Seghers, D., Slagmolen, P., Lambelin, Y., Hermans, J., Loeckx, D., Maes, F. & Suetens, P. Landmark based liver segmentation using local shape and local intensity models. *Proc. Workshop Of The 10th Int. Conf. On MICCAI, Workshop On 3D Segmentation In The Clinic: A Grand Challenge*. pp. 135-142 (2007)
13. Rios, A. & Itti, L. Closed-Loop Memory GAN for Continual Learning. *International Joint Conference On Artificial Intelligence*. (2018)
14. Zeng, H., Ge, J., Xu, W., Ma, H., Chen, L., Xia, M., Pan, B., Lin, H., Wang, S. & Gao, X. Twelve Loci Associated With Bone Density in Middle-aged and Elderly Chinese: The Shanghai Changfeng Study. *The Journal Of Clinical Endocrinology & Metabolism*. (2022,10), dgac597
15. Halabi, S., Prevedello, L., Kalpathy-Cramer, J., Mamonov, A., Bilbily, A., Cicero, M., Pan, I., Pereira, L., Sousa, R., Abdala, N., Kitamura, F., Thodberg, H., Chen, L., Shih, G., Andriole, K., Kohli, M., Erickson, B. & Flanders, A. The RSNA Pediatric Bone Age Machine Learning Challenge.. *Radiology*. **290** **2** pp. 498-503 (2019)