

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

Kaiwen Wan¹, Boming Wang¹, Fuping Wu², Haiyu Gong¹, and Xiaohai 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




Fig. 1. Illustration of 2D CgMtQ network architecture. The outputted vector represents the utility of the action taken by the agent. The magnitude of this utility determines the Q value. $h(w)$ represents width (height) of cropped patches and n represents the time lag term. J represents the number of targets.




Fig. 2. The visualization of predictions with different methods and settings on handsets. The dots (\bullet) represent Ground Truth (GD) landmarks. The dots (\bullet) represent predictions of landmarks which are close to GD while the crosses (\times) represent abnormal predictions. Noted that every image contains predictions for all 17 landmarks, but some abnormal predictions may overlapped.

Algorithm 1: Training stage of CgCRL

Data: Sequential Training Set $\mathcal{T}_k = \{(I_{k,i}, Y_{k,i}) | i \in \{1, \dots, N_k\}\}$, number of subtasks K , maximum number of episodes N , budget T , greedy parameter ϵ .

Result: Optimal network parameters ω

- 1 Initialize CgMtQ with random weights ω , the lag weight $\omega^- = \omega$, replay memory \mathbf{D} , prompt library Γ ;
- 2 When progressing to k^{th} subtask, update Γ with \mathcal{T}_k . If $k > 1$, augment \mathcal{T}_k with Γ using CMRM;
- 3 **for** $episode = 1, N$ **do**
- 4 Select a random sample pair from \mathcal{T}_k : a random image $I_{k,i}$ with j -th target landmark y_j ; a prompt Γ_j^p from Γ_j ;
- 5 Initialize the starting position $\hat{y}^{[0]}$ randomly, and the state S accordingly;
- 6 **for** $t = 1, T$ **do**
- 7 If agent is not terminated, performs $A^{[t]}$ and move to $\hat{y}^{[t+1]}$ according to CgMtQ with probability $1 - \epsilon$, otherwise selects a random action;
- 8 Get $\hat{y}^{[t+1]}$, $R^{[t]}$, $S^{[t+1]}$ and store transition $(S^{[t]}, A^{[t]}, R^{[t]}, S^{[t+1]})$ in \mathbf{D} ;
- 9 Sample a random batch of transitions from \mathbf{D} , perform a gradient descent step with Eq. (1) for CgMtQ parameters ω ;
- 10 Terminate agent if it finds the target, oscillates or reaches maximum steps, break when all agents are terminated;
- 11 **end**
- 12 For every c_1 steps, reset $\omega^- = \omega$ and update ϵ ;
- 13 **end**

Algorithm 2: Test stage of CgMtQ

Data: Test image set $\{I_i | i=1, \dots, N'\}$, number of test images N' , index of all target landmarks \mathcal{J} , prompt library Γ , budget of CgMtQ T , maximum number of iterations M' , number of prompts for j -th target N_j .

Result: Prediction for j -th landmark $\{y_{ij}^{pred} | i=1, \dots, N'\}$

- 1 Select image I_i from $\{I_i | i=1, \dots, N'\}$;
- 2 **for** $p = 1, N_j$ **do**
- 3 Initialise $y^{[0]}$ with random landmarks, select Γ_j^m from Γ ;
- 4 **for** $t = 1, T$ **do**
- 5 If agent is not terminated, perform $A^{[t]}$ according to CgMtQ, move from $y^{[t]}$ to $y^{[t+1]}$;
- 6 Terminate agent if it oscillates or reaches max steps, update \hat{y}^p with $y^{[t+1]}$, crop Γ_j^p ;
- 7 **end**
- 8 **end**
- 9 get y_{ij}^{pred} with Eq. (2) in the main manuscript.
