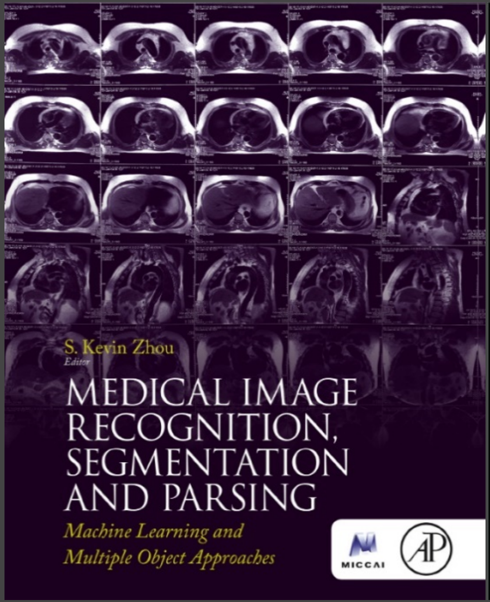

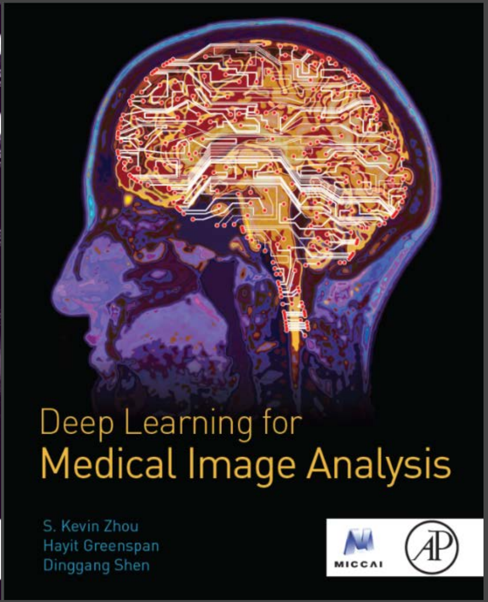

PARAMETRIC DETECTION & REGISTRATION USING DEEP REINFORCEMENT LEARNING

Dr. S. Kevin Zhou, 2018/09/16

Disclaimer

SIEMENS Healthineers

中國科学院
CHINESE ACADEMY OF SCIENCES

Institute of Computing Technology

Princeton

Beijing

~14 years

3 months

Talk outline

- DRL formulation for detection and registration
- Example 1: Landmark detection
- Example 2: 2D/3D image registration
- Example 3: Supervised action classification

1. F. Ghesu et al., Towards intelligent robust detection of anatomical structures in incomplete volumetric data, Medical Image Analysis 2018.
2. R. Liao et al., An Artificial Agent for Robust Image Registration, AAAI 2017.
3. Xu et al., Supervised Action Classifier: Approaching Landmark Detection as Image Partitioning, MICCAI 2017.

Walking from a random place to Albert Einstein Home virtually → Landmark detection

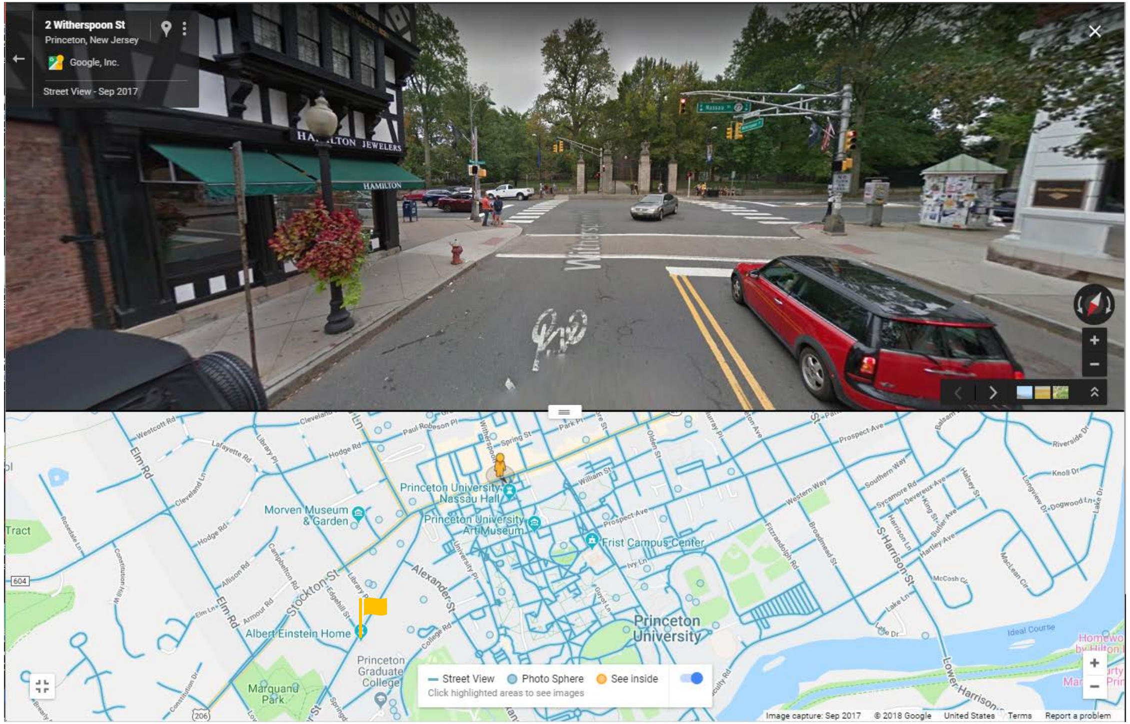

Albert Einstein Home

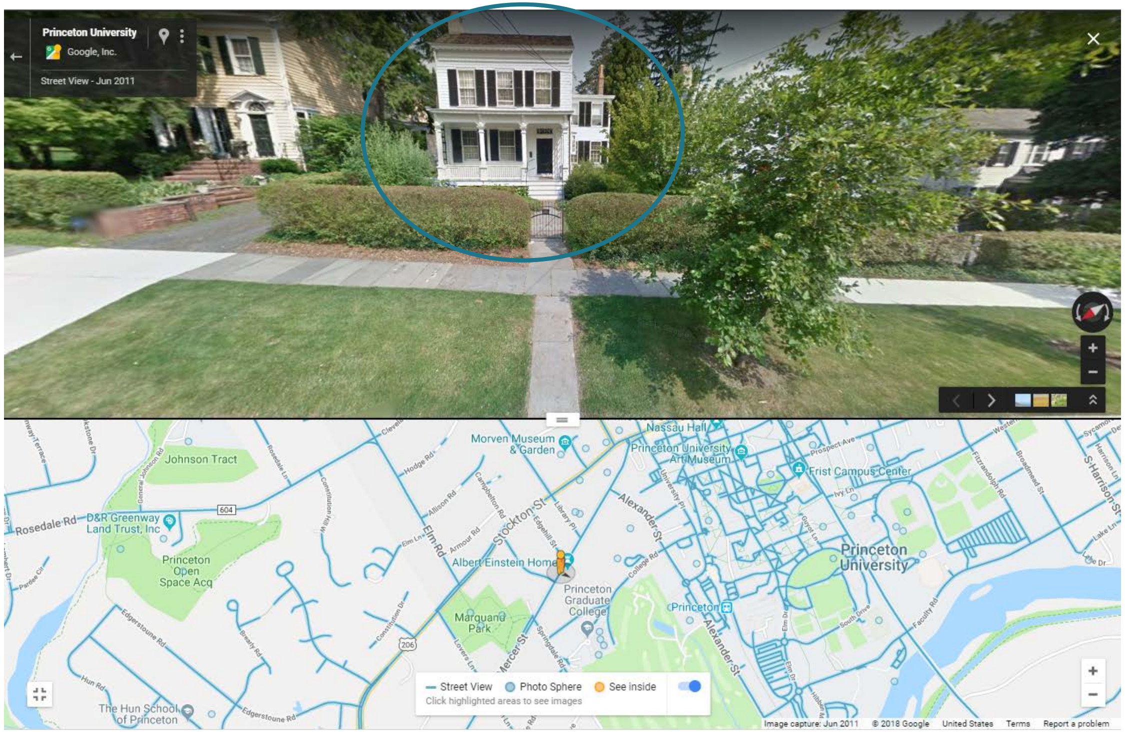

Parametric detection and registration

The parameters $\theta = [\theta_1, \theta_2, ..., \theta_n]$

- 2D Landmark detection $\theta = [x,y]$
- 3D Landmark detection $\theta = [x,y,z]$
- Rigid 2D object detection $\theta = [x,y, \alpha, s]$
- Rigid 3D object detection $\theta = [x,y,z, \alpha,\beta,\gamma, s]$
- Rigid 2D/3D registration $\theta = [x,y,z, \alpha,\beta,\gamma]$
- Rigid 3D/3D registration $\theta = [x,y,z, \alpha,\beta,\gamma, s]$

DRL formulation

Action, state, reward

Action $\underline{a}$

- Move each parameter by $\pm\delta\theta_i$ while keeping the other parameters the same
- A: action space, $|A|=2n$

State $\underline{s}$

- The observations with all actions taken so far
- $\theta_t = \theta_{t-1} + \underline{a}_t = \theta_0 + \Sigma_i \underline{a}_i$
- $<I, \theta_t>$, $I[\theta_t]$: image (or image patch) 'centered' at θ_t

Reward $\underline{r}$

- Rewards when the target is hit or closer.
- $\underline{r}(\underline{s}_{t+1}, \underline{s}_t, \underline{a}_t) = |\theta_t - \theta_0|^2 - |\theta_{t+1} - \theta_0|^2$

DRL formulation

Q-learning

- Learn Q(<u>s</u>,<u>a</u>) function, <u>s</u>: state, <u>a</u>: action

Value iteration:

$$Q^{new}(s_t, a_t) \leftarrow (1-\alpha) \cdot \underbrace{Q(s_t, a_t)}_{\text{old value}} + \underbrace{\alpha}_{\text{learning rate}} \cdot \overbrace{\left(\underbrace{r_t}_{\text{reward}} + \underbrace{\gamma}_{\text{discount factor}} \cdot \underbrace{\max_a Q(s_{t+1}, a)}_{\text{estimate of optimal future value}} \right)}^{\text{learned value}}$$

- Approximate Q(<u>s</u>,<u>a</u>) with a deep neural network

Talk outline

- DRL formulation for detection and registration
- Example 1: Landmark detection
- Example 2: 2D/3D image registration
- Example 3: Supervised action classification

- F. Ghesu et al., Robust Multi-Scale Anatomical Landmark Detection in Incomplete 3D-CT Data, MICCAI 2017.
- F. Ghesu et al., Towards intelligent robust detection of anatomical structures in incomplete volumetric data, Medical Image Analysis 2018.

Agents

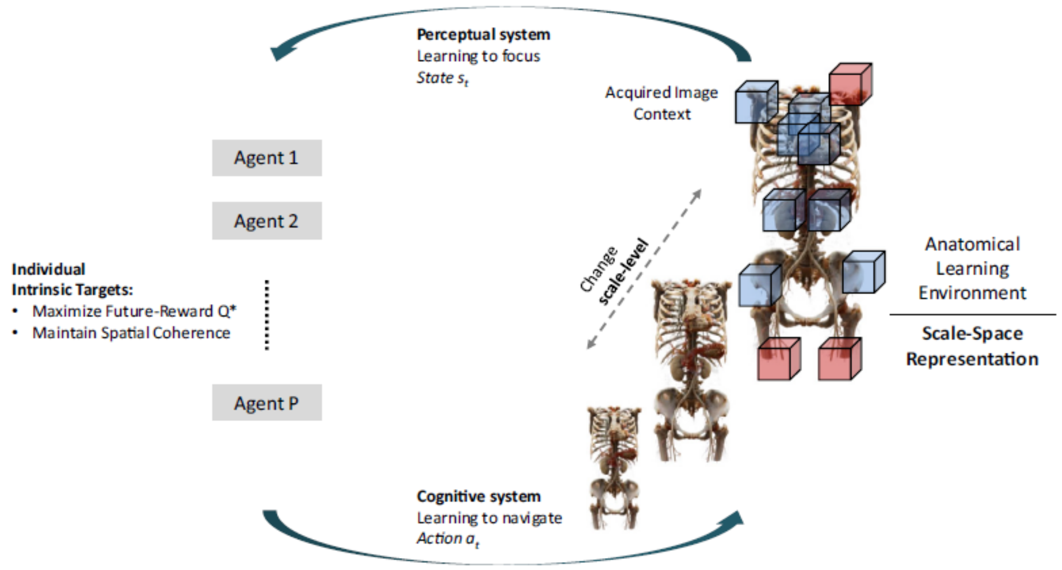

Perceptual system
Learning to focus
State s_t
Acquired Image
Context
Agent 1
Agent 2
Individual
Intrinsic Targets:
• Maximize Future-Reward Q*
• Maintain Spatial Coherence
Agent P
Change
scale-level
Anatomical
Learning
Environment
Scale-Space
Representation
Cognitive system
Learning to navigate
Action a_t

Search from coarse to fine

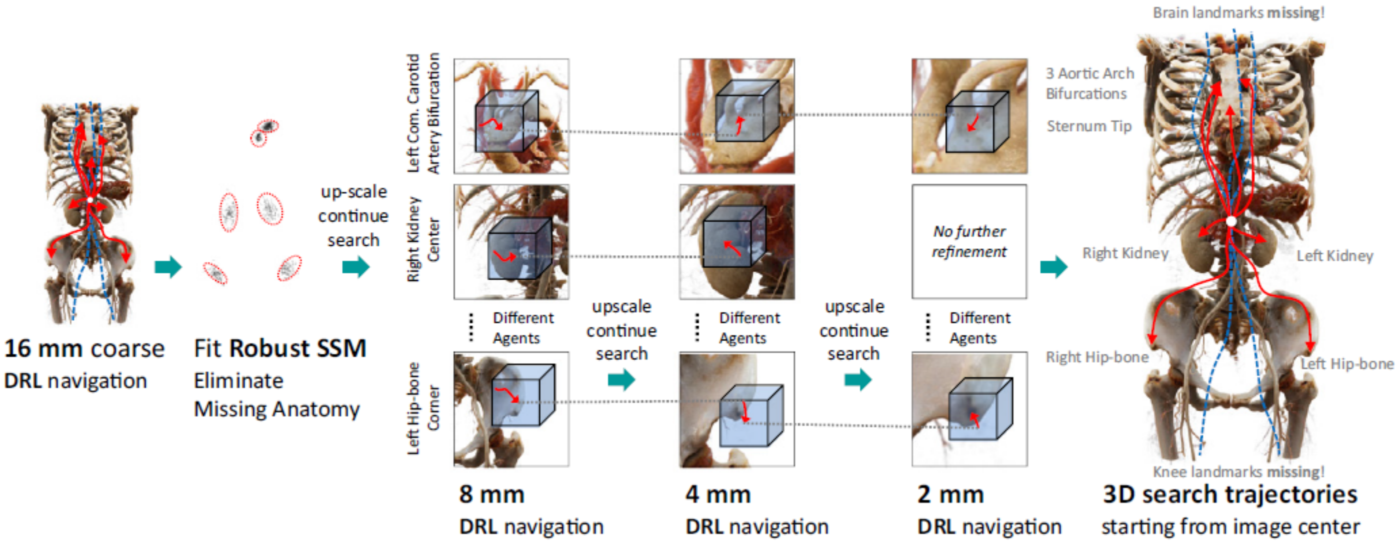

Outlier removal

- Robust statistical shape model

Each landmark follows a multi-normal distribution

$$\boldsymbol{p}_i \sim \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$$

Robust fitting via M-estimator sample consensus

$$\hat{S} \leftarrow \underset{S \in I_3(\tilde{\boldsymbol{P}})}{\arg\min} \sum_{i=0}^{|\tilde{\boldsymbol{P}}|} \min\left[\frac{1}{Z_i}\left(\phi(\tilde{\boldsymbol{p}}_i) - \boldsymbol{\mu}_i\right)^{\top} \boldsymbol{\Sigma}_i^{-1}\left(\phi(\tilde{\boldsymbol{p}}_i) - \boldsymbol{\mu}_i\right), 1\right]$$

all triples $I_3(\tilde{\boldsymbol{P}})$

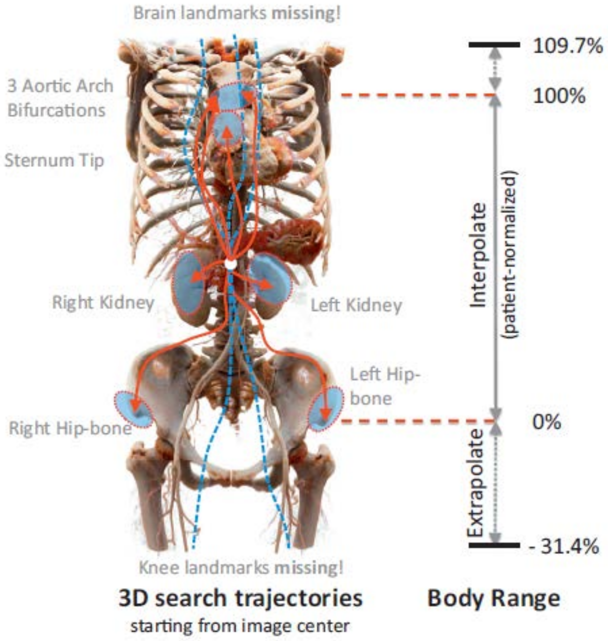

Dealing with missing landmark

- Crop from a volume with a known landmark to
  - Create a new ‘incomplete’ volume with this landmark outside of the volume
  - Record the landmark in the ‘incomplete’ volume though it is outside
- Put this into the training pool

Experiment

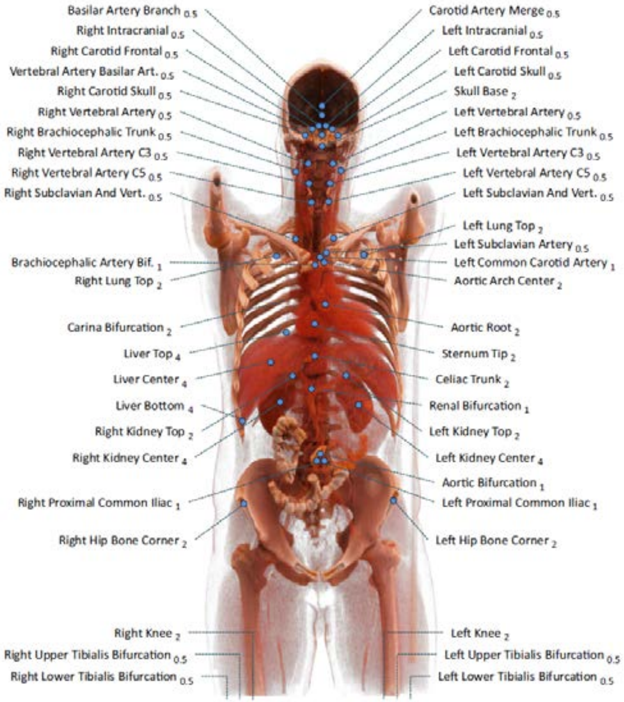

- 5043 volumes from 2000+ patients
- 49 landmarks: bones, organs, vessel bifurcations, muscles, etc.
- Evaluation excludes the landmarks very close to the boarder (<3cm)

Results

Landmark		FPR	FNR	Mean	STD	Med.
R. Kidney Center	MSDL	1.5%	9.4%	6.36	3.06	5.80
	noSSM	4.5%	1.4%	7.57	5.07	6.63
	Ours	**0%**	**0%**	6.98	3.83	6.63
L. Kidney Center	MSDL	1.5%	13.9%	6.17	3.32	5.64
	noSSM	3.0%	1.6%	7.15	4.37	6.36
	Ours	**0%**	**0%**	6.83	3.52	6.32
R. Hip Bone Corner	MSDL	1.2%	0.4%	3.66	1.83	3.44
	noSSM	0%	0.8%	2.87	3.01	2.52
	Ours	**0%**	**0%**	2.63	1.53	2.49
L. Hip Bone Corner	MSDL	1.1%	1.2%	4.92	2.09	4.70
	noSSM	0.5%	1.6%	3.82	3.02	2.83
	Ours	**0%**	**0%**	3.61	2.08	2.83
Left Common Car. Artery Bifurcation	MSDL	1.0%	10.8%	4.78	3.30	4.17
	noSSM	1.5%	3.1%	4.22	4.56	2.88
	Ours	**0%**	**0%**	4.02	3.33	2.86
Brachiocephalic Art. Bifurcation	MSDL	1.0%	11.3%	5.05	3.02	4.54
	noSSM	2.0%	2.2%	4.35	3.61	3.46
	Ours	**0%**	**0%**	4.26	2.97	3.46
Left Subclavian Art. Bifurcation	MSDL	1.1%	7.2%	5.25	3.51	4.62
	noSSM	3.0%	2.2%	4.41	4.57	3.23
	Ours	**0%**	**0%**	4.23	3.37	3.21
Carina Bifurcation	MSDL	1.0%	4.9%	5.10	2.82	4.53
	noSSM	2.0%	0.8%	4.09	2.17	3.78
	Ours	**0%**	**0%**	4.07	2.16	3.77

No FP or FN with the aid of SSM

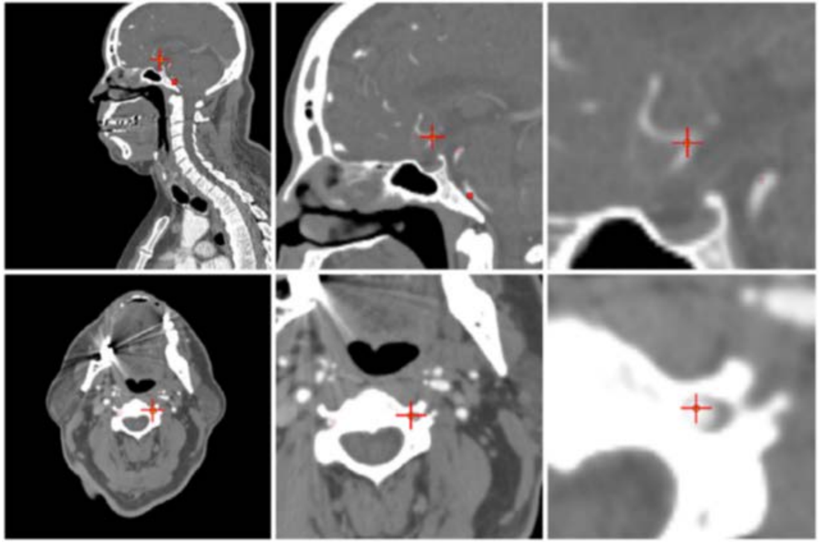

52ms per landmark (Intel 8-core)
28ms per landmark (Nvidia Pascal)

Talk outline

- DRL formulation for detection and registration
- Example 1: Landmark detection
- Example 2: 2D/3D image registration
- Example 3: Supervised action classification

- R. Liao et al., An Artificial Agent for Robust Image Registration, AAAI 2017.

2D/3D registration

- The goal of 2D/3D registration is to find the 3D pose (6 DoFs) of a 3D volume to match with 2D X-ray image(s).

Preoperative CT

Intraoperative X-Ray

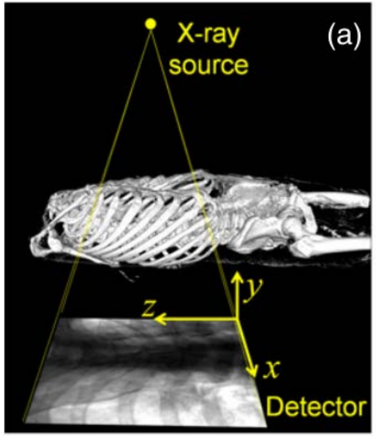

3D Position

Overlay Preoperative Labels

3D pose of CT/CBCT

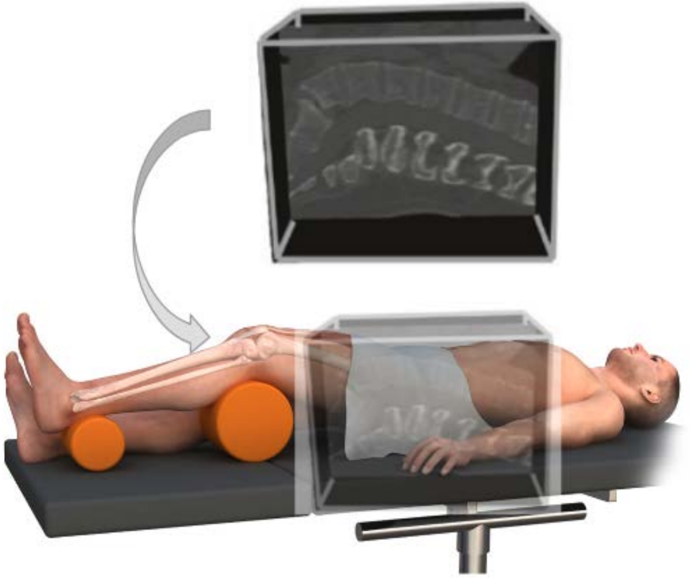

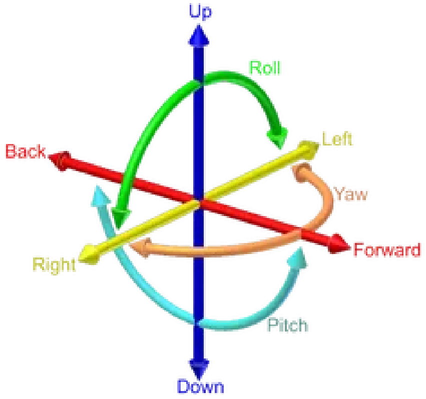

$$T(t_x,t_y,t_z,\theta_x,\theta_y,\theta_z)=\begin{bmatrix}1&0&0&t_x\\0&c\theta_x&-s\theta_x&t_y\\0&s\theta_x&c\theta_x&t_z\\0&0&0&1\end{bmatrix}$$

$$\times\begin{bmatrix}c\theta_y&0&s\theta_y&0\\0&0&0&0\\-s\theta_y&0&c\theta_y&0\\0&0&0&1\end{bmatrix}\times\begin{bmatrix}c\theta_z&-s\theta_z&0&0\\s\theta_z&c\theta_z&0&0\\0&0&1&0\\0&0&0&1\end{bmatrix}$$

Digitally reconstructed radiography (DRR)

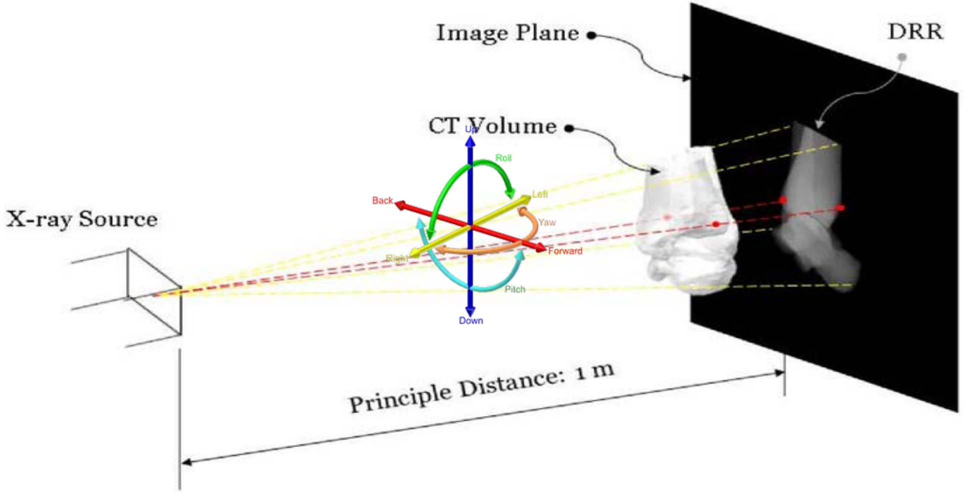

X-Ray projection

DRL based method

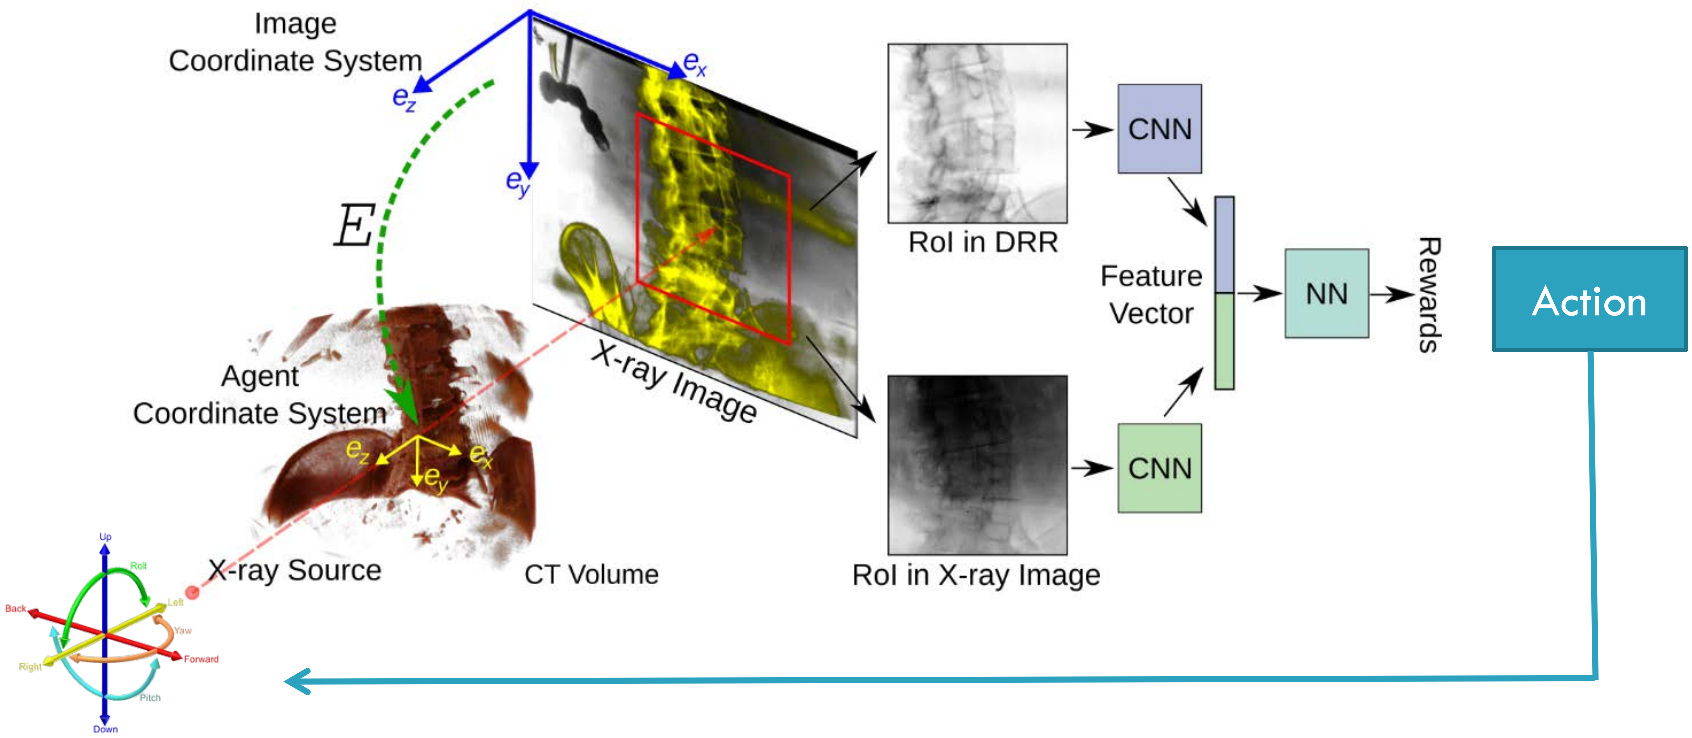

Experiments

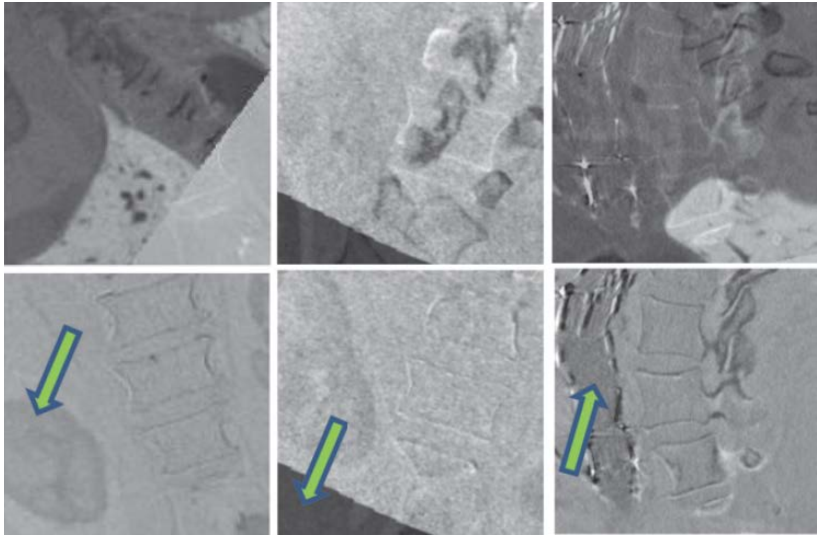

Spine

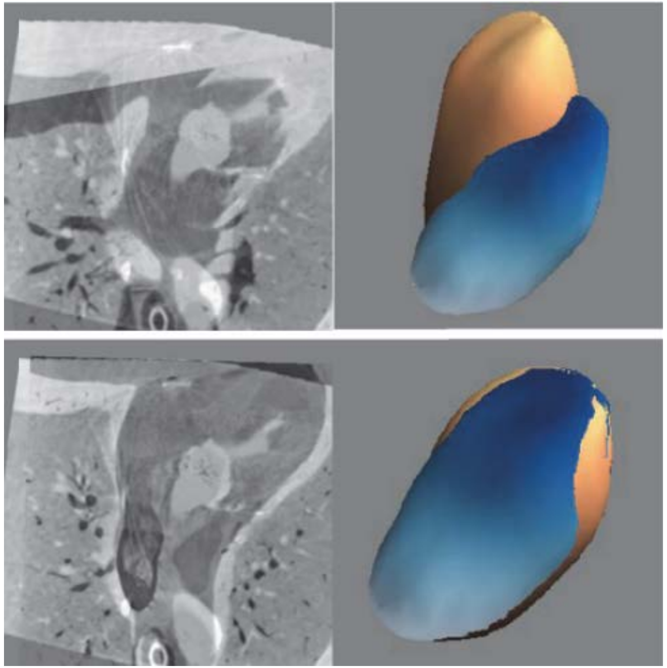

Heart

Quantitative results

	Spine (E1) (TRE mm)				Heart (E2) (MME mm)			
Methods	**Success**	**10th**	**50th**	**90th**	**Success**	**10th**	**50th**	**90th**
Ground Truth	N/A	0.8	0.9	1.2	N/A	2.1	4.0	5.9
Initial Position	N/A	35.5	73.9	116.2	N/A	9.2	22.8	30.5
ITK(Ibanez et al. 2005)	12%	1.9	77.3	130.4	14%	14.9	34.9	47.6
Quasi-global(Miao et al. 2013)	20%	1.6	60.9	136.2	14%	16.2	35.9	58.7
Semantic registration(Neumann et al. 2015)	24%	3.0	34.9	71.0	72%	7.6	15.3	30.6
Proposed method	92%	1.7	2.5	3.8	100%	3.2	4.8	6.9
Human registration	70%	0.8	1.6	15.8	96%	4.0	6.2	13.4

Talk outline

- DRL formulation for detection and registration
- Example 1: Landmark detection
- Example 2: 2D/3D image registration
- Example 3: Supervised action classification

• Xu et al., Supervised Action Classifier: Approaching Landmark Detection as Image Partitioning, MICCAI 2017.

Landmark representation:

spatially local vs <u>spatially global</u>

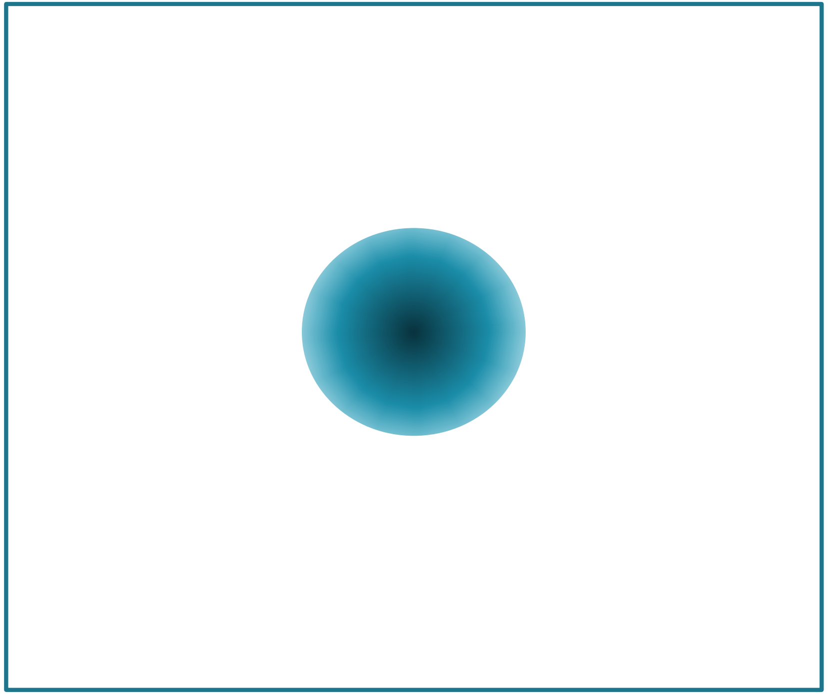

Heat map

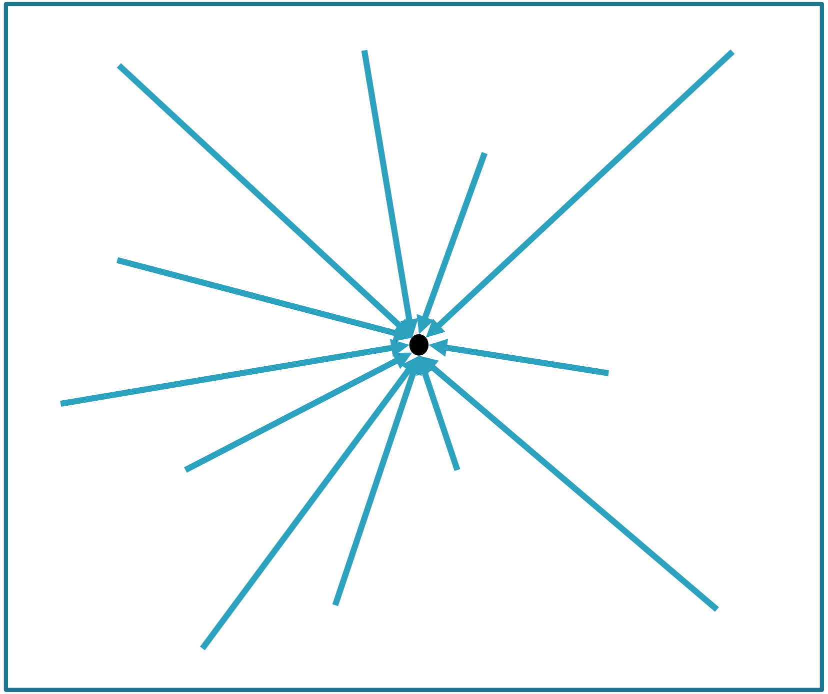

Relative offset vector

- Zhou et al., Shape regression machine, Medical Image Analysis 2010

Landmark representation:

spatially global vs spatially distributed

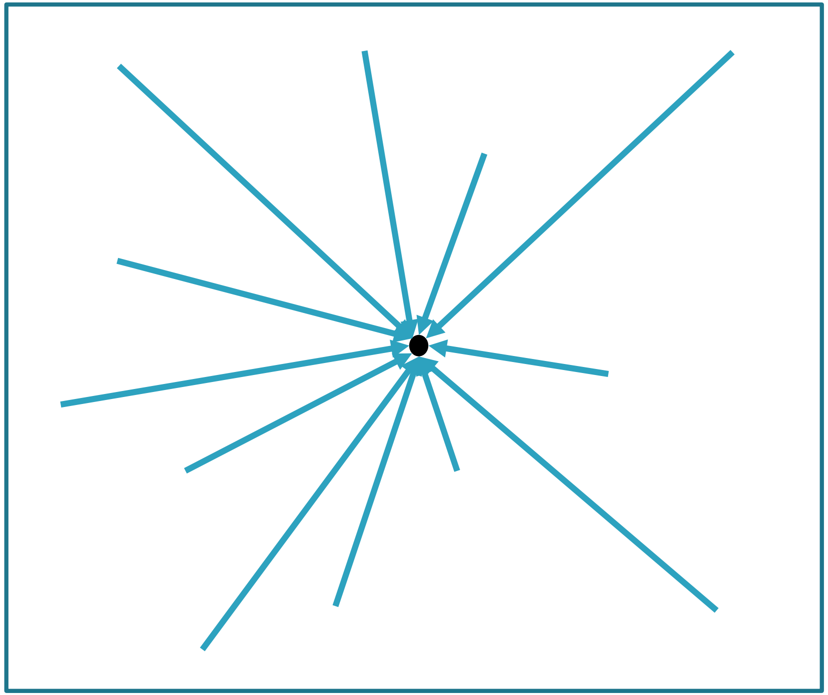

(continuous) Relative offset vector

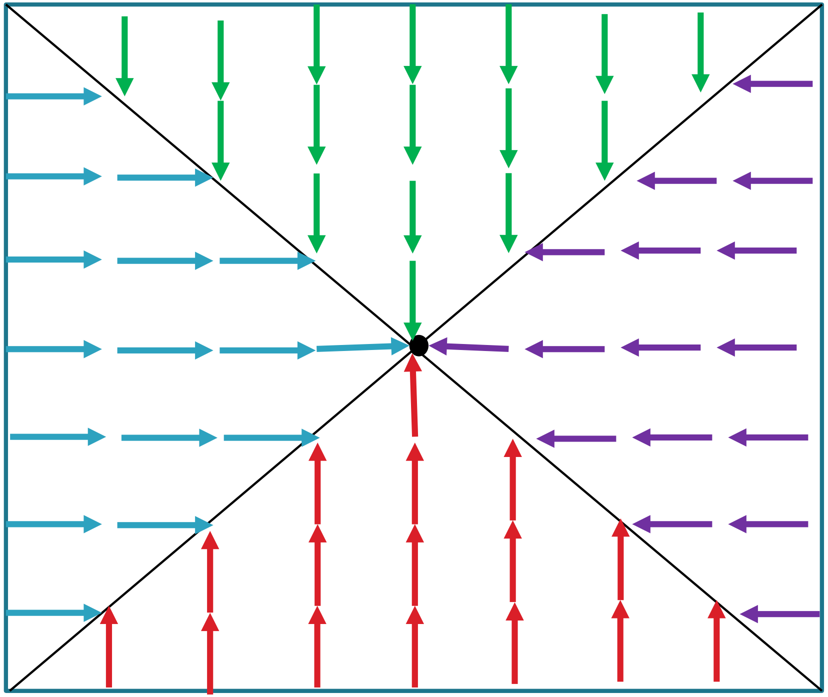

Discrete action map

• Xu et al., Supervised Action Classifier: Approaching Landmark Detection as Image Partitioning, MICCAI 2017.

Discrete action map

UP $(d_x^{(0)} = 0, d_y^{(0)} = -1)$

RIGHT $(d_x^{(1)} = 1, d_y^{(1)} = 0)$

$\|d_x^{(a)}\|^2 + \|d_y^{(a)}\|^2 = 1$

DOWN $(d_x^{(2)} = 0, d_y^{(2)} = 1)$

LEFT $(d_x^{(3)} = -1, d_y^{(3)} = 0)$

Supervised path : $\hat{a} = \operatorname{argmin}_a \sqrt{\left(x - x_t + d_x^{(a)}\right)^2 + \left(y - y_t + d_y^{(a)}\right)^2}$

Solutions : 2 lines $y = x + (\hat{y} - \hat{x})$ $y = -x + (\hat{x} + \hat{y})$

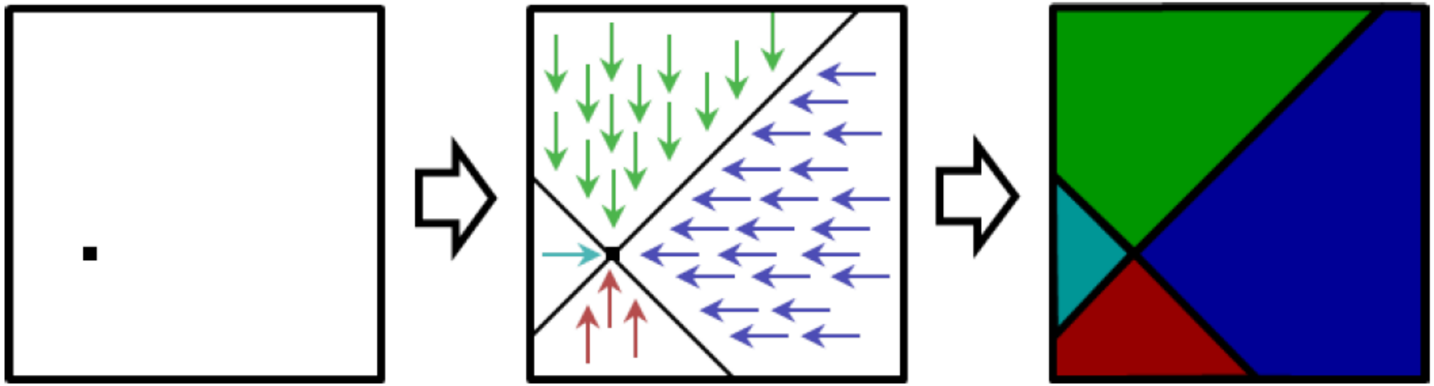

Action aggregation

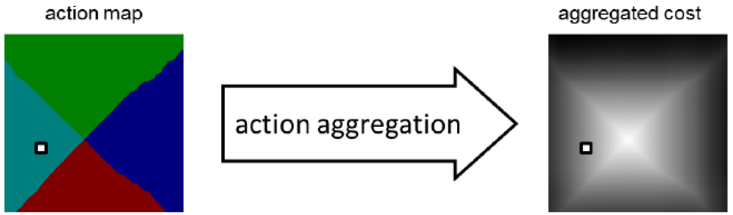

$$(x', y') = \arg\max_{(x,y)} C(x,y) = \arg\max_{(x,y)} \sum_a C_a(x,y)$$

$$C_a(x,y) = \begin{cases} d_x^{(a)}\{\sum_{i=x}^{\infty}\delta(A(i,y)==a) - \sum_{i=-\infty}^{x}\delta(A(i,y)==a)\} & \text{if } \|d_x^{(a)}\|=1,\\ d_y^{(a)}\{\sum_{i=y}^{\infty}\delta(A(x,j)==a)-\sum_{i=-\infty}^{y}\delta(A(x,j)==a)\} & \text{if } \|d_y^{(a)}\|=1.\end{cases}$$

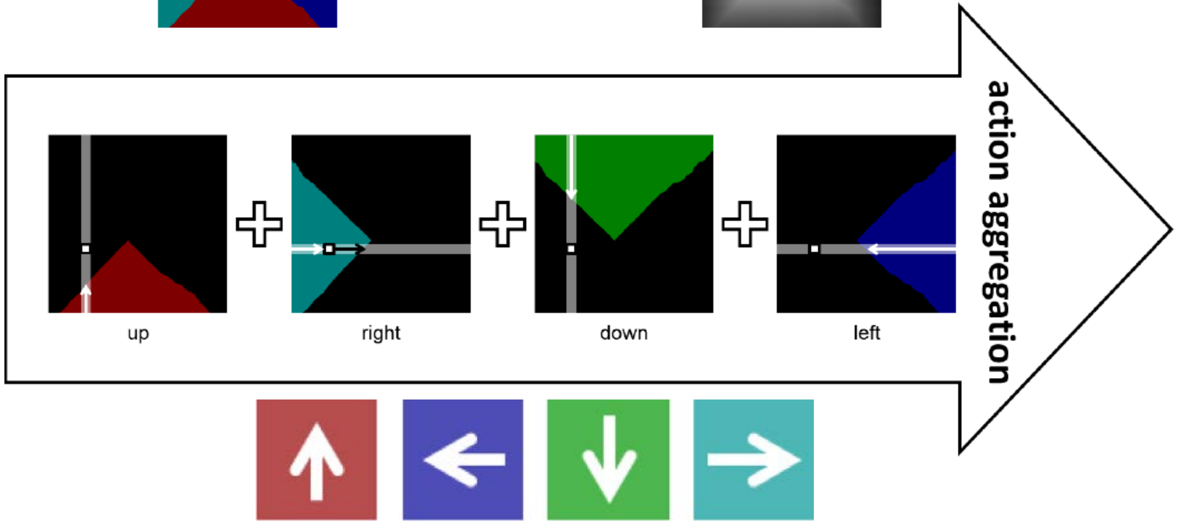

Method recap

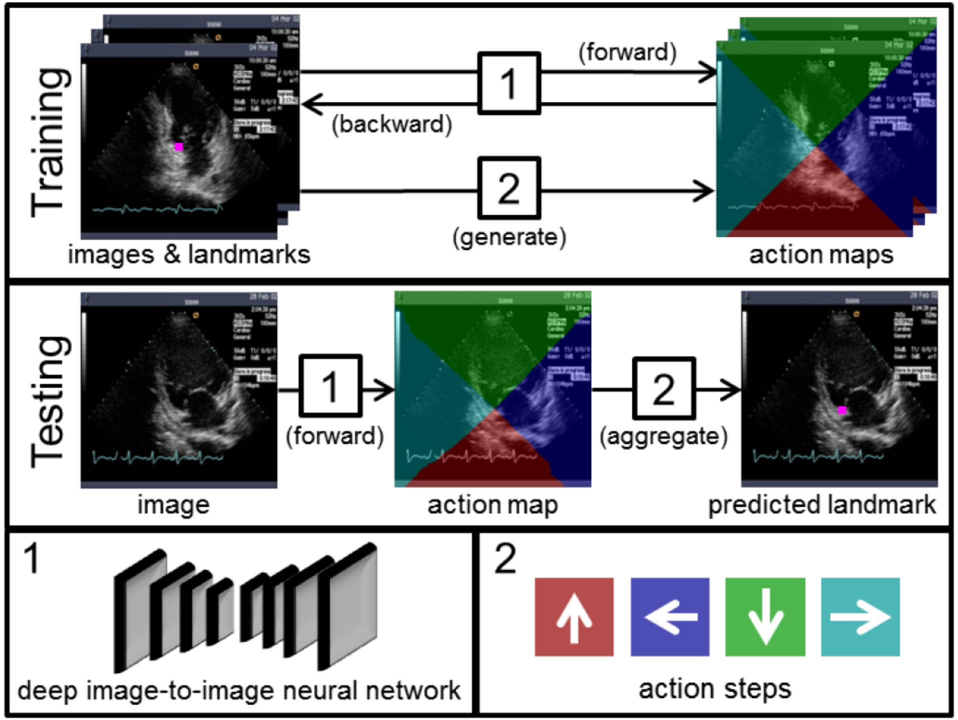

Experimental results

		PBT		DRL		I2I		SAC	
		lmk1	lmk2	lmk1	lmk2	lmk1	lmk2	lmk1	lmk2
CA	mean	10.45	13.85	7.69	10.02	6.73	9.02	**6.31**	**8.01**
	50%	5.74	8.11	5.43	7.63	5.00	6.40	**4.35**	**5.88**
	80%	11.11	16.18	9.33	13.73	8.54	11.40	**7.54**	**10.83**
OB	mean	59.23	130.66	29.99	32.45	30.07	21.97	**14.94**	**16.76**
	50%	35.31	139.49	11.69	13.17	5.39	6.08	**4.85**	**5.91**
	80%	109.84	193.64	43.98	45.76	13.34	15.54	**11.76**	**13.67**

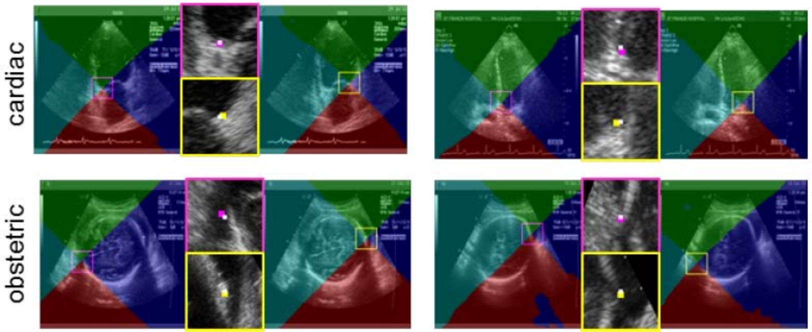

1353 cardiac A4C and A4C images
1643 OB images

Summary

- DRL is a powerful framework for image-based parameter inference
- The trick is to define actions, states, and reward.
- Success has been achieved in various medical imaging problems.
- Still difficult to learn a high-dimensional Q- function