Accelerated Intravascular Ultrasound Imaging using Deep Reinforcement Learning

Tristan S.W. Stevens[1], Nishith Chennakeshava[1], Frederik J. de Bruijn[2], Martin Pekar[2] and Ruud J.G. van Sloun[1,2]

4195

OVERVIEW

- ✓ **Clinical challenge:** Intravascular Ultrasound Imaging (IVUS) has a limited number of physical channels which imposes a *trade-off* between image quality and frame-rate.
- ✓ **Current approach:** Hand-crafted solutions such as sub-aperture or sparse array designs are used to decrease number of measurements needed for reconstruction. More advanced solutions are focused on MRI and not applicable to IVUS [1].
- ✓ **Hypothesis:** An adaptive learned algorithm can be deployed to leverage the *trade-off* more efficiently.
- ✓ **Goal:** Optimally extract information from the acoustic scene by adaptive sampling of ultrasound measurements.
- ✓ **Method:** A Deep Reinforcement Learning agent (AiVUS) is trained to adaptively subsample all measurements needed for a single IVUS reconstruction.

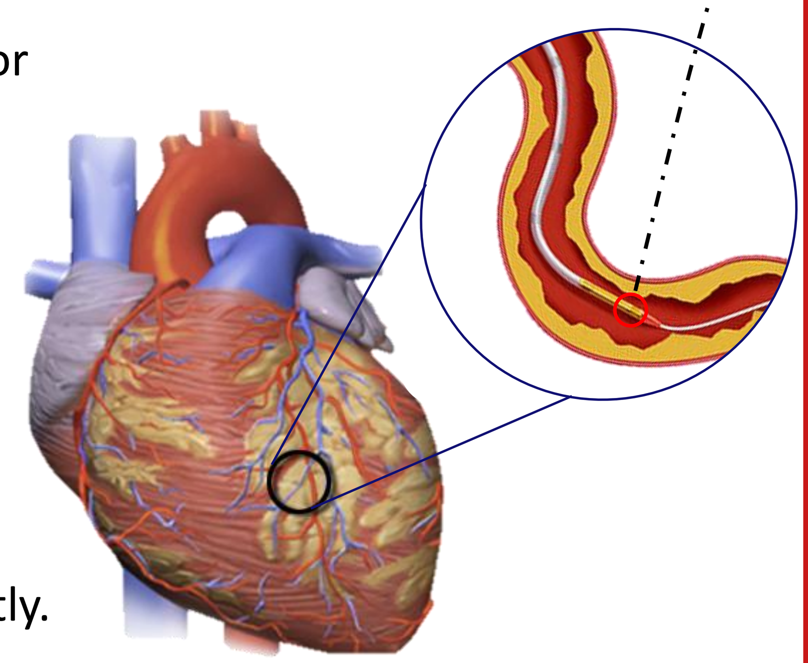

THE GOOD STUFF

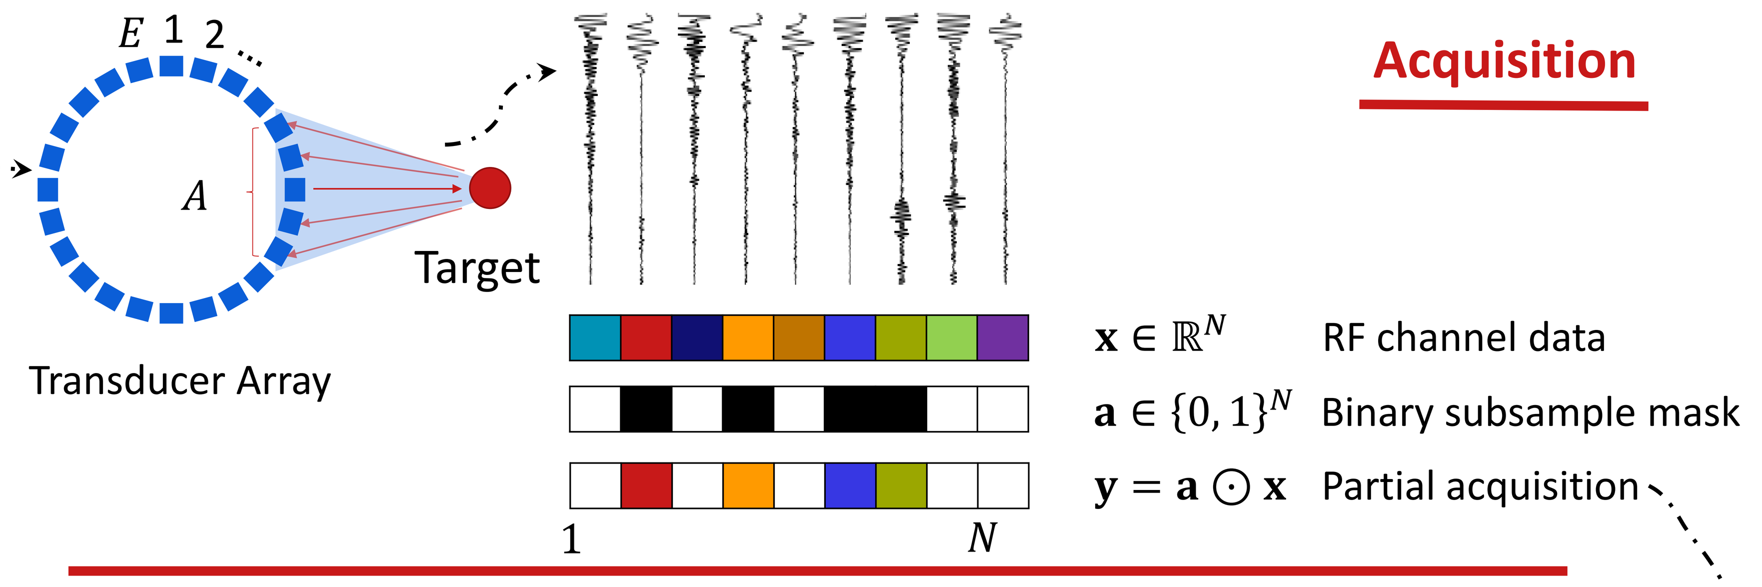

Exploit knowledge of the current frame to provide an optimal acquisition sequence for the next reconstruction

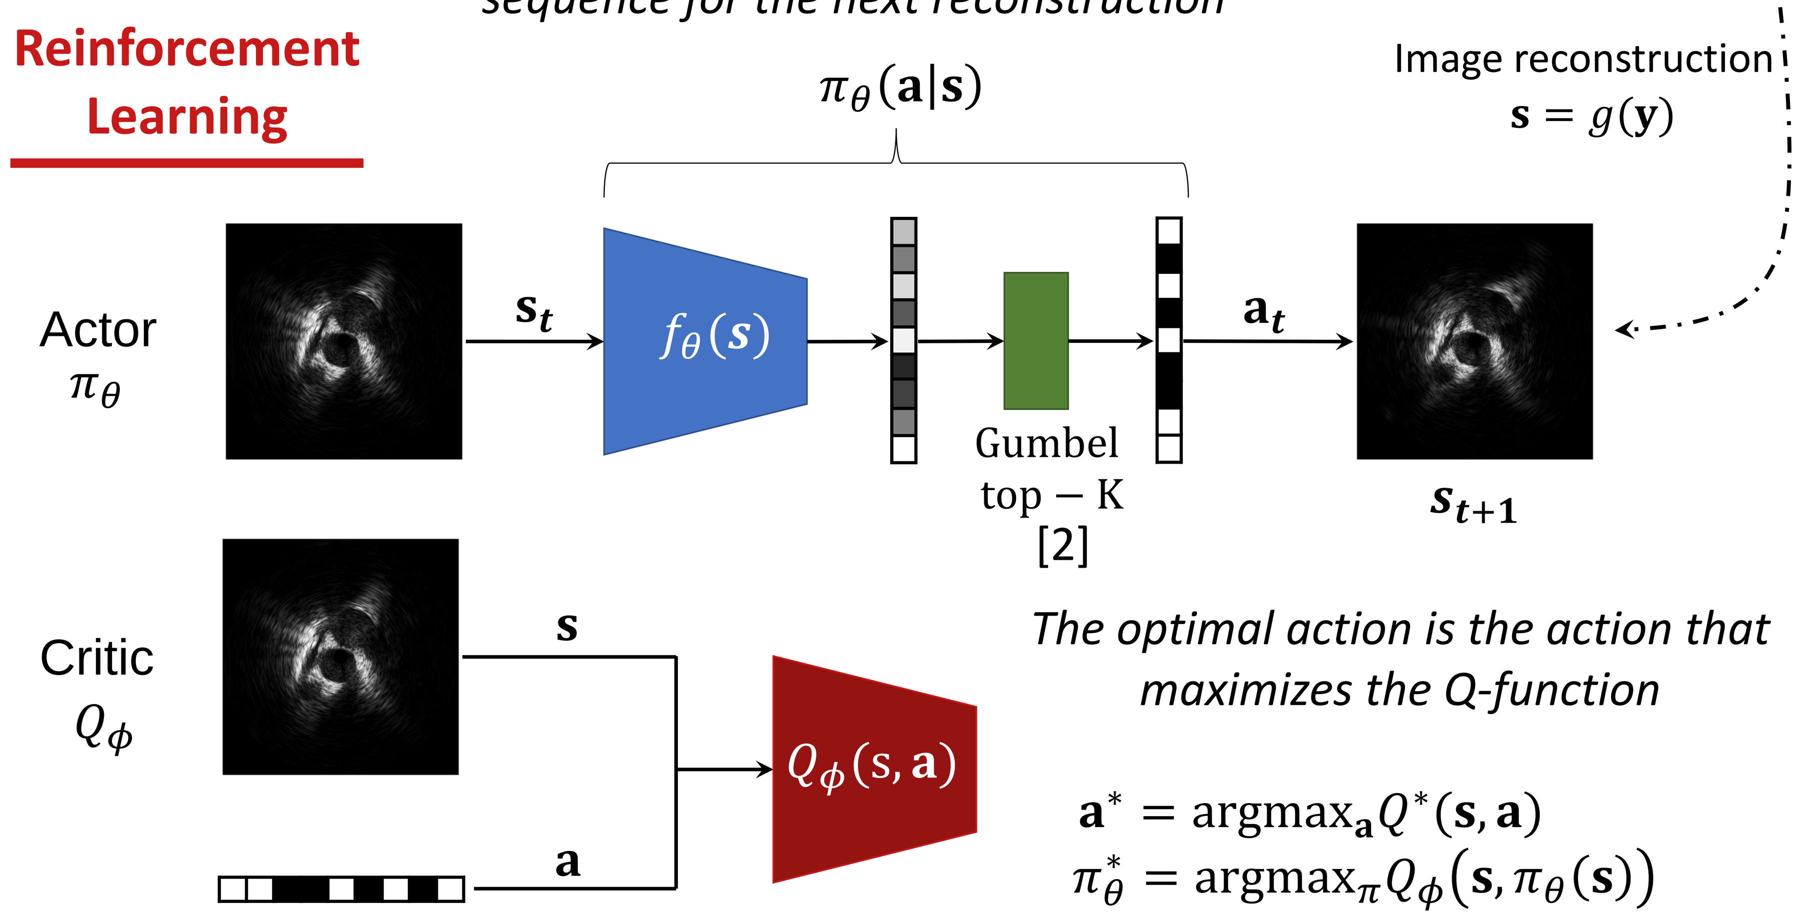

The optimal action is the action that maximizes the Q-function

$$\mathbf{a}^* = \text{argmax}_\mathbf{a} Q^*(\mathbf{s}, \mathbf{a})$$
$$\pi^*_\theta = \text{argmax}_\pi Q_\phi(\mathbf{s}, \pi_\theta(\mathbf{s}))$$

DATA

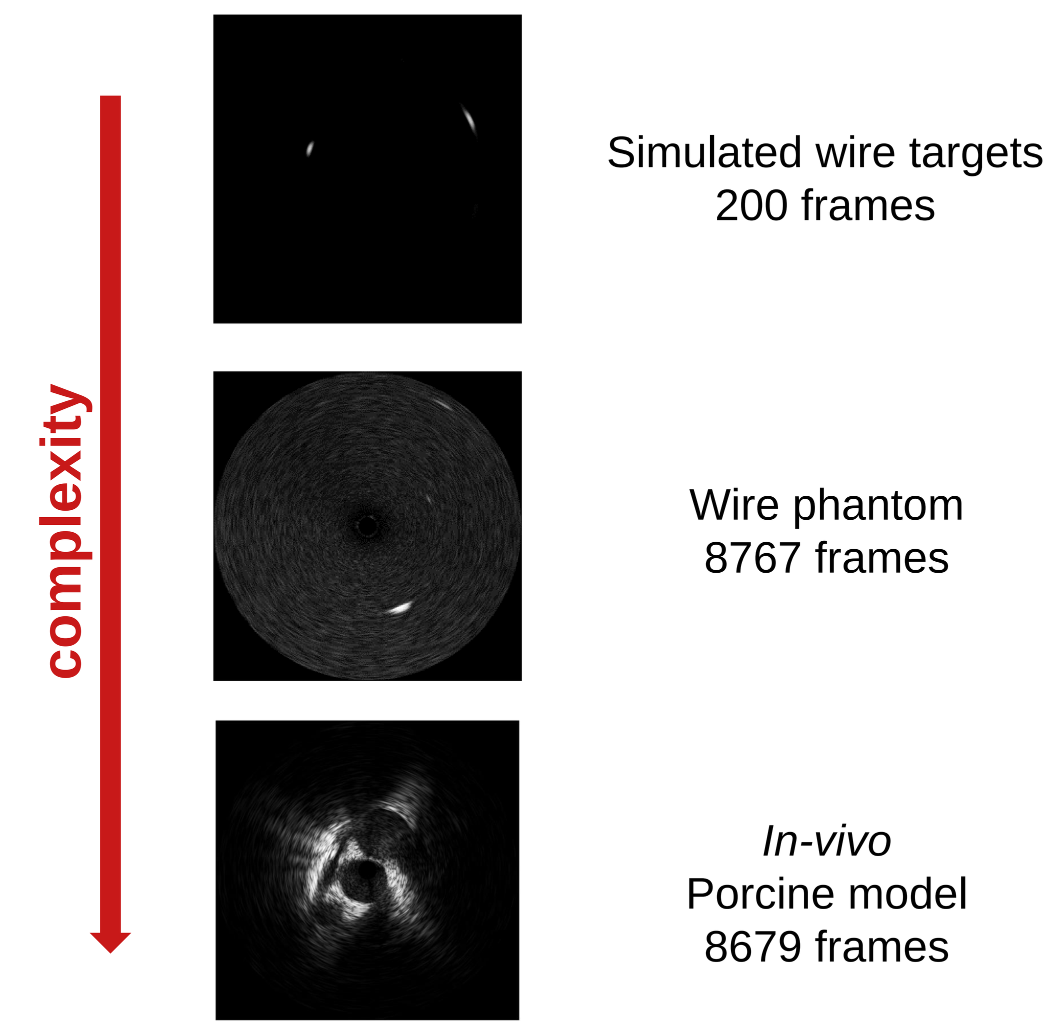

Conclusions

AiVUS can navigate in controlled IVUS environments with high dimensional state and action spaces. AiVUS outperforms a random agent using a learned acquisition strategy.

RESULTS

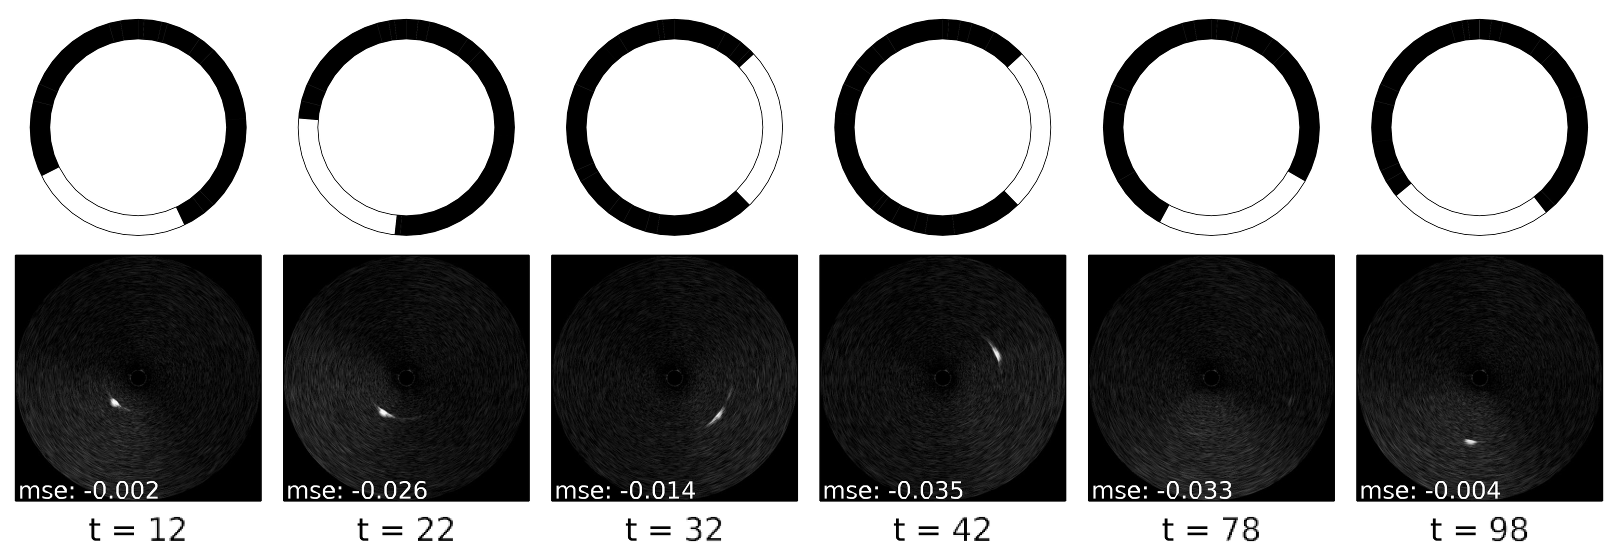

Six successive wire phantom frames constructed using AiVUS. The agent's action is displayed in the top row, where the circle represents the elements in the transducer array.

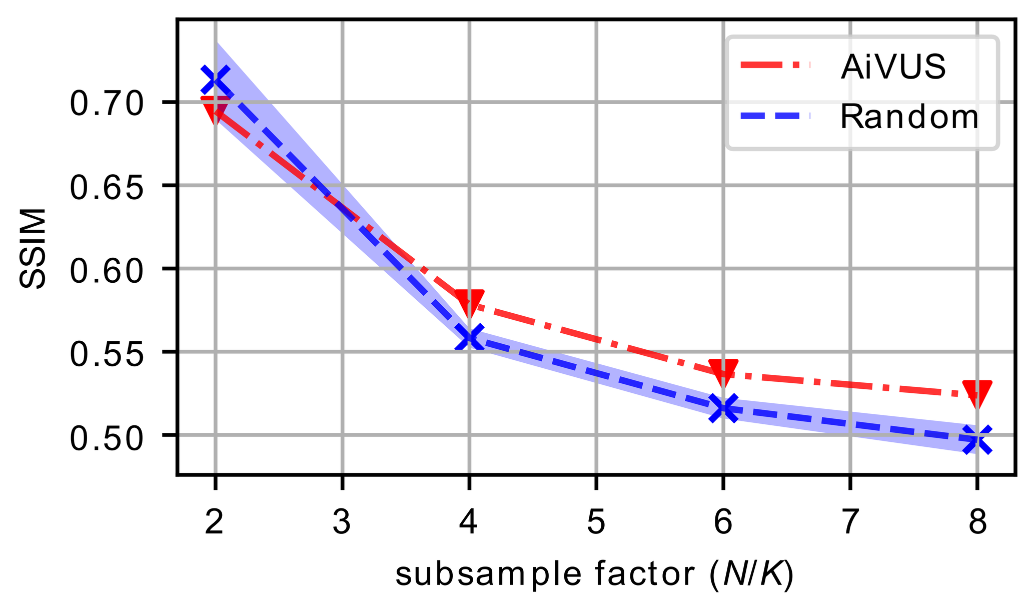

SSIM performance on the *in-vivo* test data for both learned (AiVUS) and random sampling strategies.

Dataset	Sim.		Phantom		In-vivo	
Agent	I	II	I	II	I	II
MSE ↓	3.25	1.42	0.067	0.047	0.078	0.070
MAE ↓	0.069	0.034	0.169	0.133	0.211	0.200
PSNR ↑	44.33	49.95	61.59	64.48	59.43	59.90
SSIM ↑	0.996	0.998	0.308	0.447	0.552	0.578

Quantitative results comparing a random agent (I) with a trained agent (II) (AiVUS).

Future work

Extend actions beyond subsampling; i.e., frequency, pulse duration, etc.

References

[1] Bakker, Tim, Herke van Hoof, and Max Welling. "Experimental design for MRI by greedy policy search." Advances in Neural Information Processing Systems 33 (2020)

[2] Huijben, Iris AM, et al. "A Review of the Gumbel-max Trick and its Extensions for Discrete Stochasticity in Machine Learning." IEEE Transactions on Pattern Analysis and Machine Intelligence (2022).

Paper

[1]Eindhoven University of Technology, [2]Philips Research. Eindhoven, The Netherlands.

t.s.w.stevens@tue.nl