

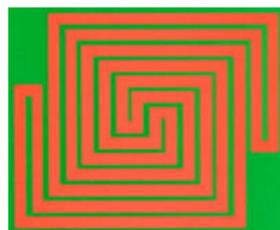
Do Neural Networks for Segmentation Understand Insiderness?

Kimberly Villalobos*, Vilim Štih*, Amineh Ahmadinejad*, Shobhita Sundaram, Jamell Dozier, Andrew Francl, Frederico Azevedo, Tomotake Sasaki+, Xavier Boix+

* and +: equal contributions



1. What pixels are inside a closed curve?

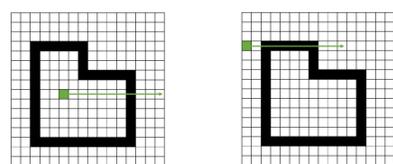


- Challenging long-range dependencies
- In isolation to avoid confounding factors
- Helps to debug image segmentation networks

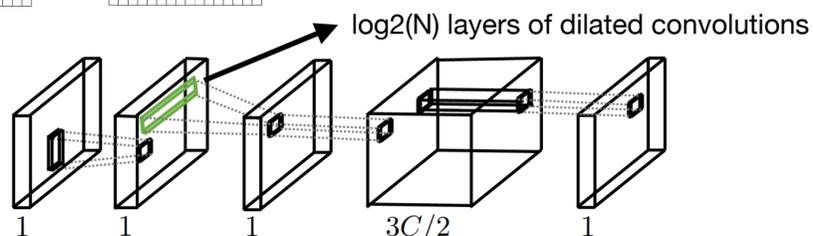
2. Can DNNs for segmentation solve insiderness? YES

We analytically demonstrate that segmentation network architectures have sufficient complexity to exactly solve the insiderness for any given curve with few neurons.

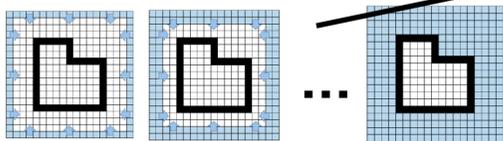
Dilated convolutions



Ray-intersection method can be implemented with at least $\log_2(N)$ layers and $3N/2$ kernels (3×3), in which N = image width or height.



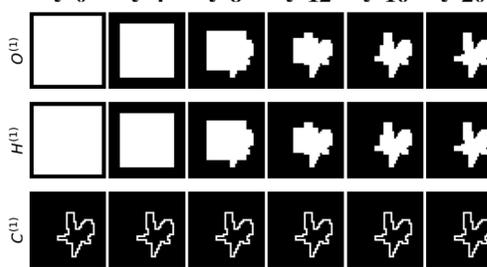
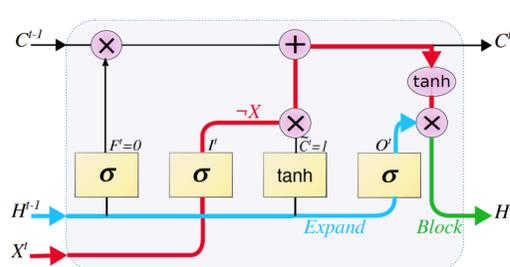
Recurrent Networks



Breaks the long-range dependencies into local operations

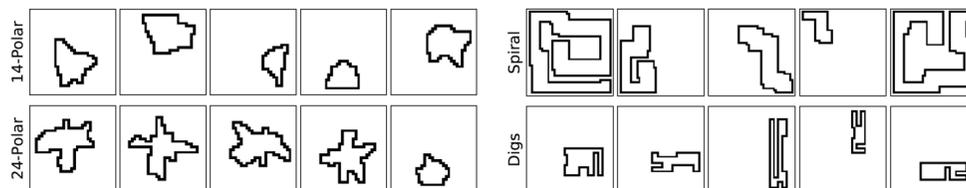
The colouring routine can be implemented with at least one convLSTM with 1 kernel (3×3).

t=0 t=4 t=8 t=12 t=16 t=20

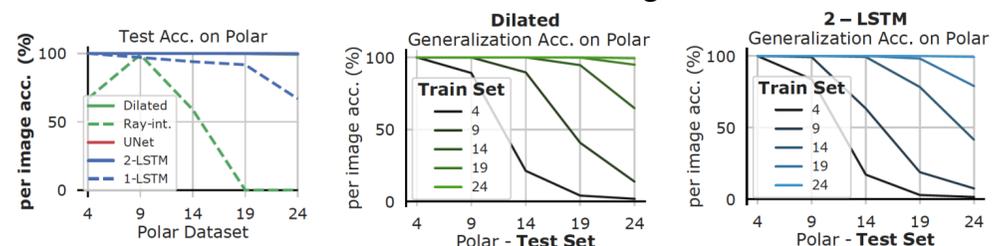


3. Can DNNs for segmentation learn insiderness? NO

Dataset to train the networks:

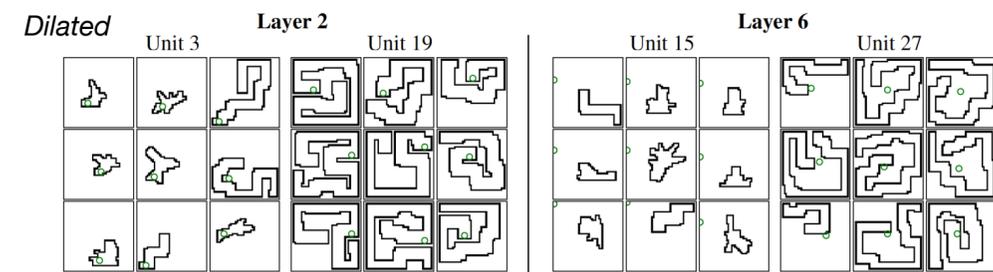


Cross-dataset evaluation shows lack of generalization:

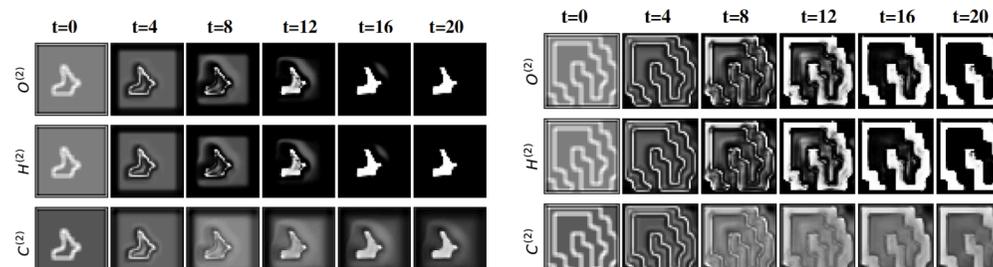


Test Set	Dilated			CNN			UNet			2 - LSTM			2 - LSTM, w/oinit.		
	Polar	Spiral	Both	Polar	Spiral	Both	Polar	Spiral	Both	Polar	Spiral	Both	Polar	Spiral	Both
Polar	99	4.6	94	99	3.1	97	99	5.3	98	99	14	99	99	56	99
Spiral	18	99	92	24	99	98	0.56	99	98	30	99	99	31	99	99
Digs	16	0.88	8.7	18	0.69	11	32	1.3	10	23	3.7	30	58	14	28

Neural representations do not capture general solutions:

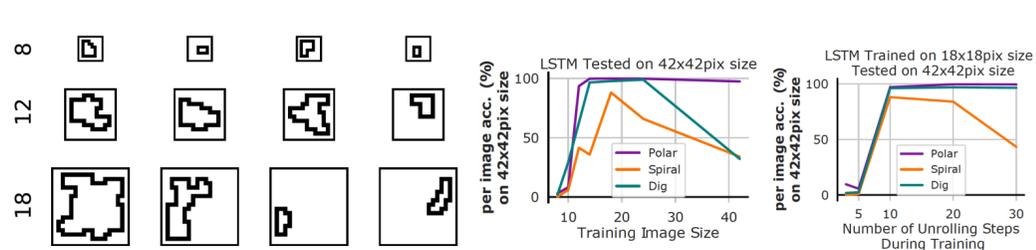


2-LSTM



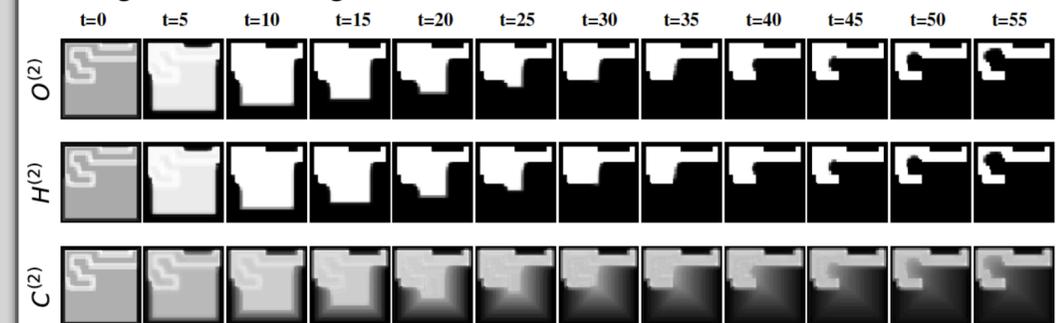
BUT 4. Learning the Coloring routine with small images

Controlling the number of unrolling steps during training via small images helps.



The network generalises on all size images when training on 18x18 pix images and 10 unrolling steps.

Emergent colouring routine:



Conclusions

- Networks for segmentation are unable to learn general solution for insiderness even though they have sufficient complexity to do so.
- Recurrent networks trained on small images with controlled number of unrolling steps are effective to learn general solutions.

Acknowledgements

We are grateful to Tomaso Poggio for his insightful advice and warm encouragement. We also thank Shimon Ullman and Pawan Sinha for helpful feedback and discussions. This work is supported by the Center for Brains, Minds and Machines (funded by NSF STC award CCF-1231216), Fujitsu Laboratories Ltd. (Contract No. 40008401) and the MIT-Sensetime Alliance on Artificial Intelligence.