Rotation-Scale Equivariant Steerable Filters

Author name(s) withheld EMAIL(S) WITHHELD
Address withheld

Editors: Under Review for MIDL 2023

Abstract

This is a great paper and it has a concise abstract.
Keywords: List of keywords, comma separated.

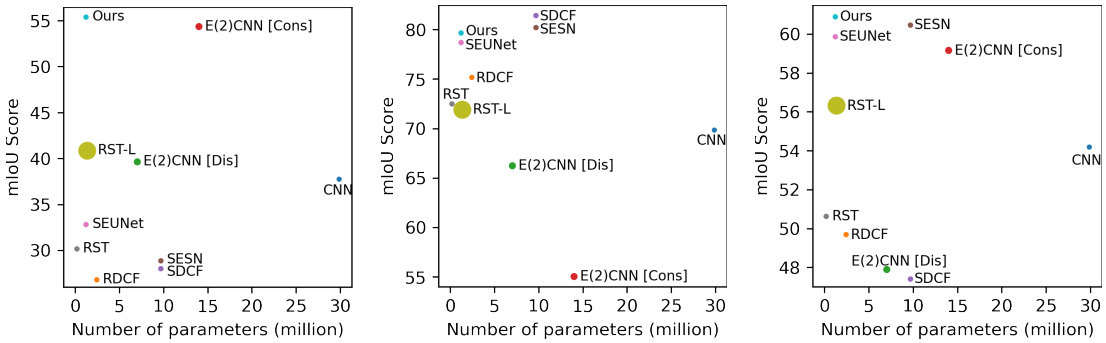

(a) Texture (b) CRAG (c) GlaS

Figure 1: **Model Size vs. Segmentation Performance vs. GPU Requirement.** Our Rotation-Scale Equivariant Steerable Filter outperforms other filters. In particular, RSESF achieves state-of-the-art mIoU but is much smaller and GPU efficient. Details are in Table 2 and 4.

1. Introduction

This is where the content of your paper goes. Some random notes[1]:

- You should use LaTeX(Lamport, 1986).
- JMLR/PMLR uses natbib for references. For simplicity, here, `\cite` defaults to parenthetical citations, i.e. `\citep`. You can of course also use `\citet` for textual citations.
- Eprints such as arXiv papers can of course be cited (Hinton et al., 2015), Hinton et al. (2015). We recomend using a `@misc` bibtex entry for these as shown in the sample bibliography.
- You should follow the guidelines provided by the conference.

1. Random footnote are discouraged

© 2023 Anonymous.

- Read through the JMLR template documentation for specific LaTeXusage questions.
- Note that the JMLR template provides many handy functionalities such as `\figureref` to refer to a figure, e.g. Figure **??**, `\tableref` to refer to a table, e.g. Table 3 and `\equationref` to refer to an equation, e.g. Equation (1).

Acknowledgments

Acknowledgments withheld.

References

Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. Eprint arXiv:1503.02531, 2015.

Leslie Lamport. *Latex: A Document Preparation System.* Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1986. ISBN 0-201-15790-X.

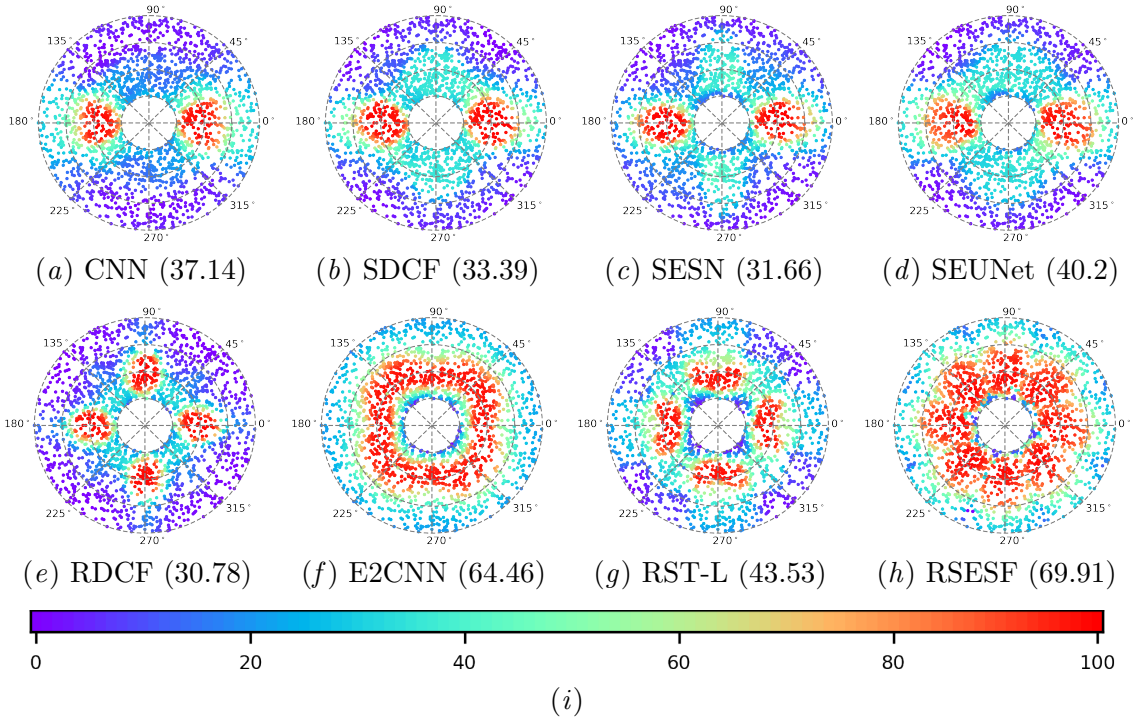

Figure 2: Polar plots show model predictions as a function of the orientation and scale of the input.

Table 1: Segmentation performance for in distribution (ID) and out-of-distribution settings.

Filter Types	Year	**In Distribution**			**Out-of-Distribution**		
		GlaS	**CRAG**	**Texture**	**GlaS**	**CRAG**	**Texture**
CNN	2015	67.09	87.83	99.11	54.19	69.86	37.78
RDCF	2019	65.48	87.69	99.05	49.69	75.19	26.82
E(2)CNN [Dis]	2019	**68.31**	**89.73**	99.57	47.91	66.26	39.67
E(2)CNN [Con]	2019	65.90	82.87	**99.71**	59.16	55.07	54.37
SDCF	2022	66.87	88.63	99.16	47.41	**81.45**	28.03
SESN	2020	67.73	88.51	98.83	60.47	80.21	28.87
SEUNet	2022	67.15	87.63	98.55	59.87	78.73	32.82
RST	2022	65.88	83.41	95.92	50.63	72.52	30.19
RST-L	2022	67.82	87.94	99.25	56.33	71.90	40.87
ours	2022	66.53	88.16	98.00	**60.89**	79.69	**55.42**

Table 2: Segmentation performance for in distribution (ID) and out-of-distribution settings.

Channel Squeeze at the Last Layer	**In Distribution**			**Out-of-Distribution**		
	GlaS	**CRAG**	**Texture**	**GlaS**	**CRAG**	**Texture**
Pooling over R	**68.02**	**88.79**	75.58	60.81	**80.49**	30.89
Selecting R^* over C	66.53	88.16	**98.00**	**60.89**	79.69	**55.42**
Selecting R^c & for each c	64.24	86.18	78.87	51.98	67.92	32.04

Table 3: Model comparison in terms of number of trainable parameters (Million), amount of GPU (GB) required during training, and the training time (seconds) per epoch (averaged over 70 epochs).

	CNN	RDCF	E2CNN [Dis]	E2CNN [Con]	SDCF	SESN	SEUNet	RST	RST-L	Ours
Params	28.95	2.42	7.03	13.98	9.66	9.66	1.22	0.15	1.36	1.22
GPU	4.68	5.07	7.65	7.65	5.09	5.06	5.05	5.06	15.12	5.05
Speed	93	383	182	45	427	406	372	365	895	372

Appendix A. Proof of Theorem 1

This is a boring technical proof of

$$\cos^2\theta + \sin^2\theta \equiv 1. \tag{1}$$

Appendix B. Proof of Theorem 2

This is a complete version of a proof sketched in the main text.