How Convolutional Neural Networks Deal with Aliasing

Antônio H. Ribeiro, Thomas B. Schön

Uppsala University, Sweden

UPPSALA
UNIVERSITET

IEEE International Conference on Acoustics, Speech and Signal Processing
ICASSP, 2021

Aliasing

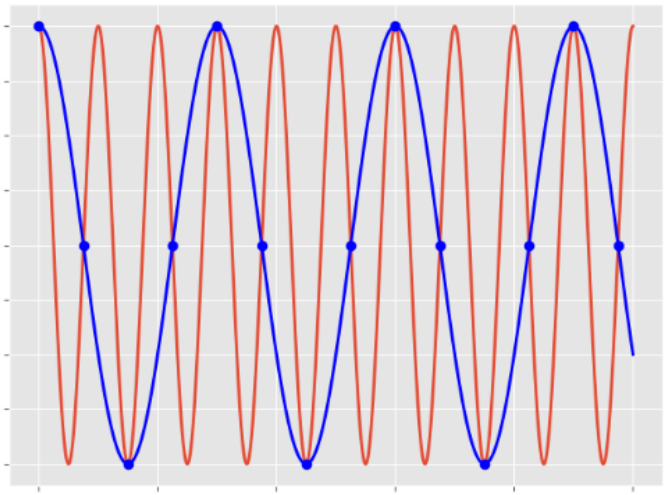

Figure: Aliasing illustration

Video: Perfectly synced rotor and camera frame rate.

Convolutional neural networks and downsampling

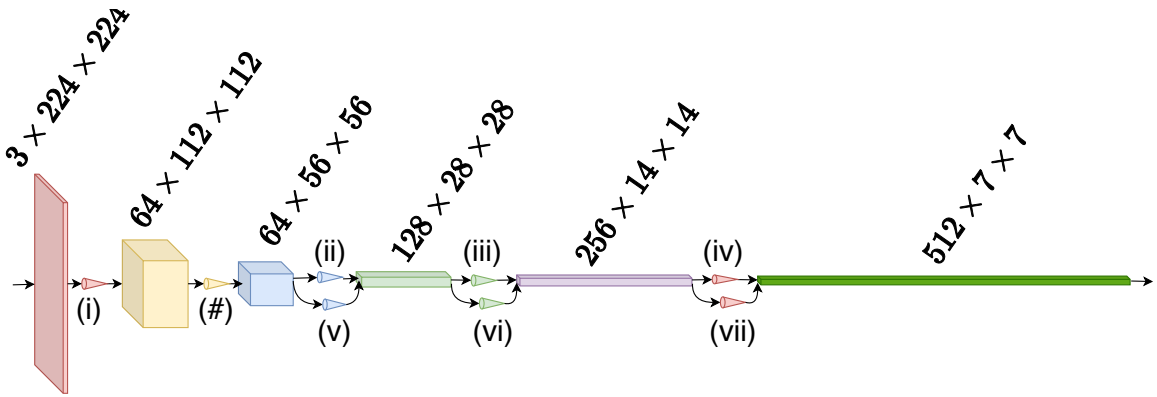

Figure: Dimensions of the intermediary tensors of ResNet34.

Previous work: anti-aliased convolutional neural networks

C. Vasconcelos, H. Larochelle, V. Dumoulin, *et al.*, "An Effective Anti-Aliasing Approach for Residual Networks," *arXiv:2011.10675*, Nov. 2020. arXiv: 2011.10675.

X. Zou, "Delving Deeper into Anti-aliasing in ConvNets," en, in *Proceedings of the 31st British Machine Vision Virtual Conference (BMVC)*, 2020.

R. Zhang, "Making Convolutional Networks Shift-Invariant Again," in *Proceedings of the 36th International Conference on Machine Learning (ICML)*, Jun. 2019. arXiv: 1904.11486.

How do convolutional neural networks manage to be successful without explicit anti-aliasing mechanisms?

How do convolutional neural networks manage to be successful without explicit anti-aliasing mechanisms?

1. Can it resolve between oscillations at its input?

How do convolutional neural networks manage to be successful without explicit anti-aliasing mechanisms?

1. Can it resolve between oscillations at its input?
2. Does it learn anti-aliasing filters?

Experiment description

Can it resolve between oscillations at its input?

Classification problem with **400** different frequencies (ω_1, ω_2).

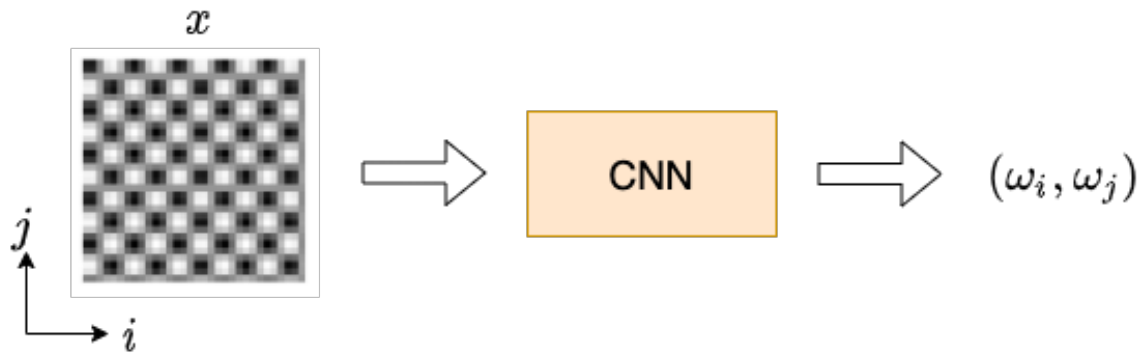

Experiment description

Can it resolve between oscillations at its input?

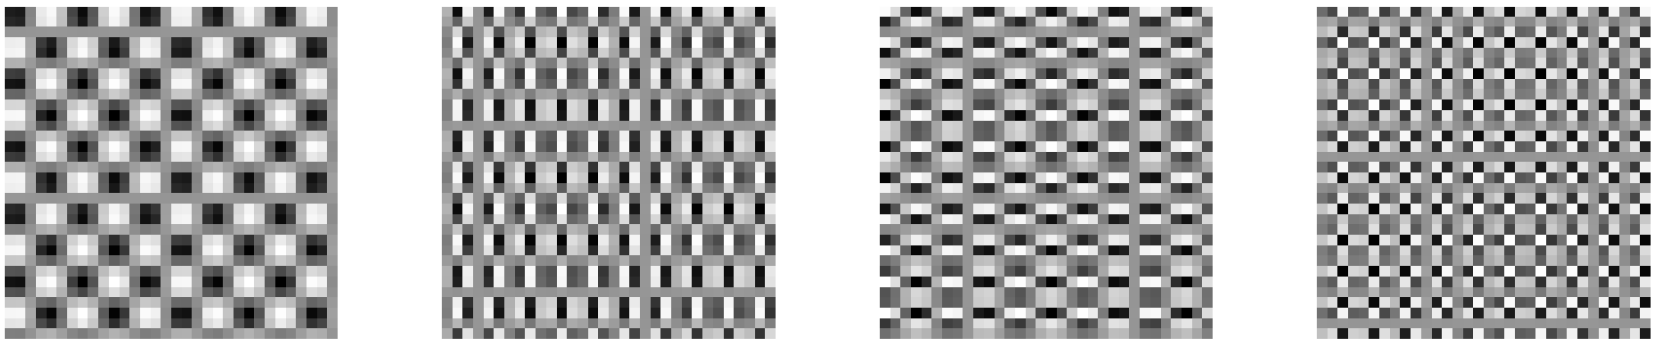

(a) (b) (c) (d)

Figure: **Four oscillatory patterns**. If the image is downsampled by a factor of two, the patterns become *indistinguishable due to aliasing*.

Results and the role of redundancy

Can it resolve between oscillations at its input?

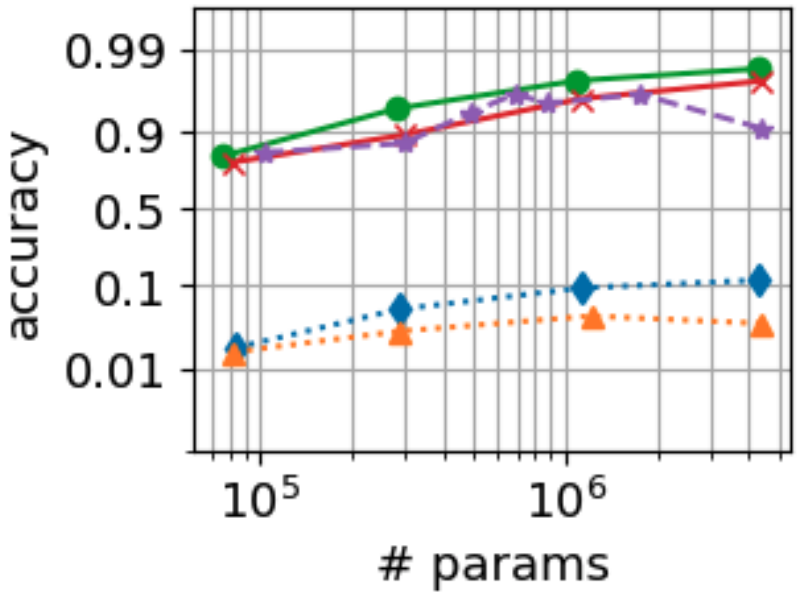

Fully connected neural network

- 1 hidden layer
- 2 hidden layers

ResNet

- fixed depth and and increasing # of channels.
- Fixed # of channels and increasing depth channel (constant proportion)
- Fixed # of channels and increasing depth channel (increasing proportion)

Quantifying aliasing

Does it learn anti-aliasing filters?

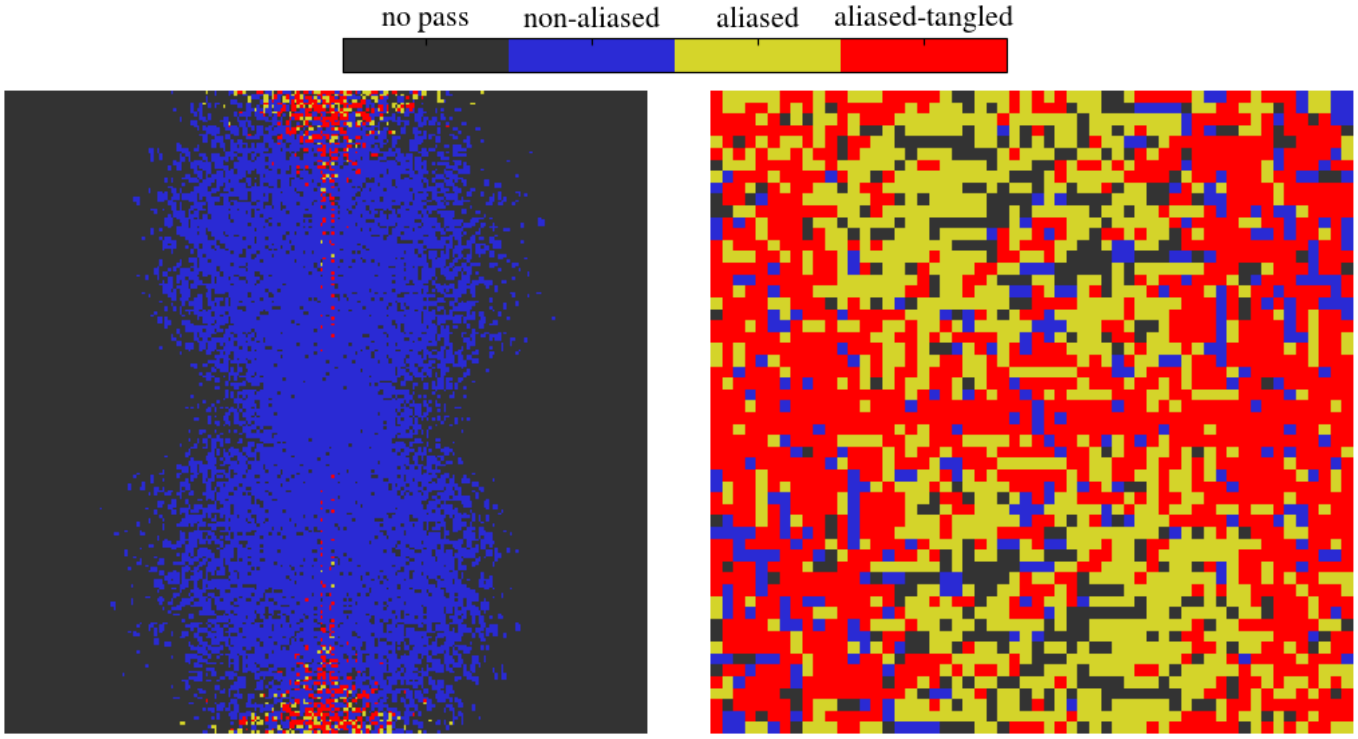

(a) 65%, 33%, 1%, 1% (b) 15%, 10%, 35%, 39%

Figure: **Aliasing in CNNs**. DFT points classified according to for the intermediate signals of the ResNet34 evaluated on an ImageNet test sample.

Does it learn anti-aliasing filters?

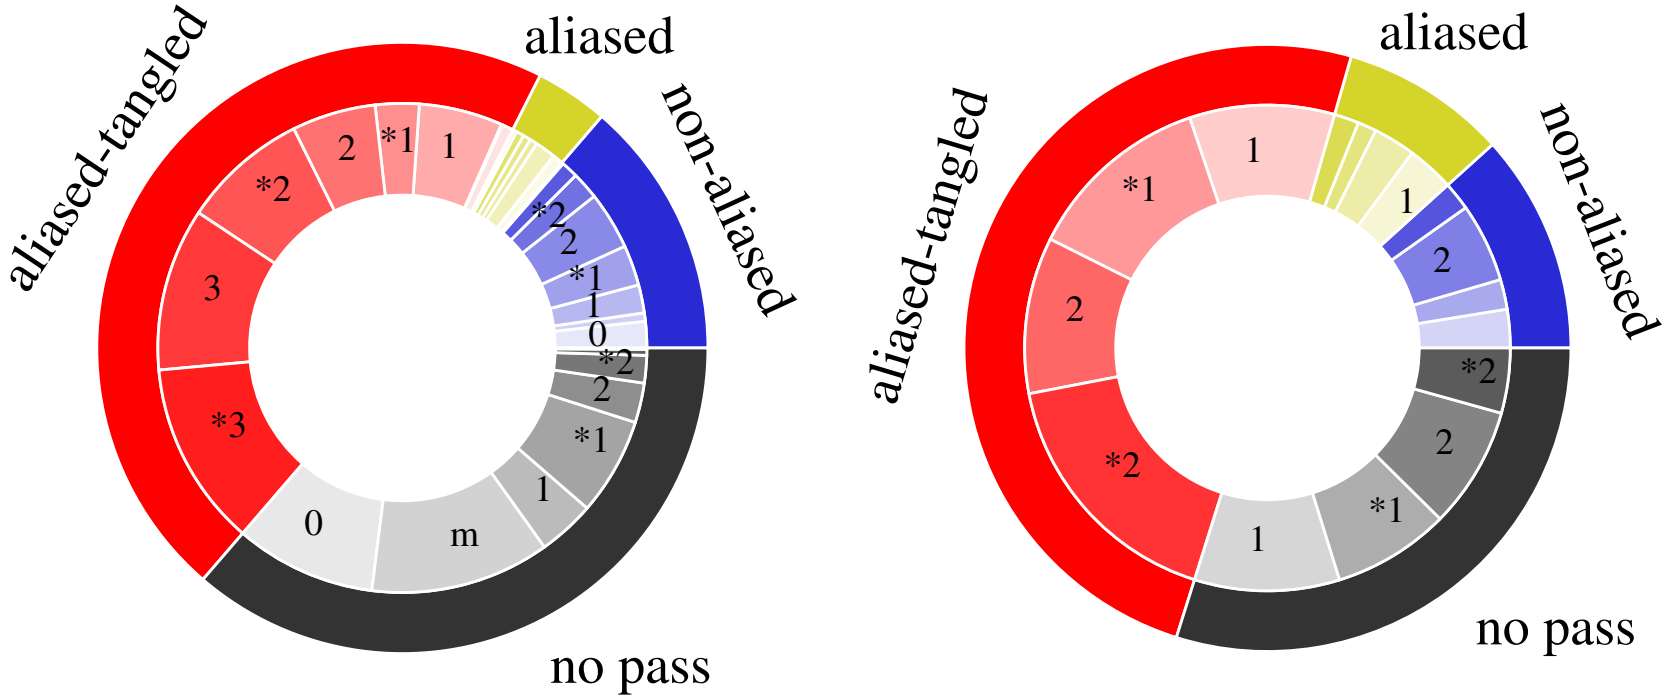

(a) ImageNet.

(b) Oscilation classif.

Figure: **Fraction of samples suffering aliasing**. Pie chart indicating the fraction of intermediate signals suffering aliasing.

Does it impact performance?

Does it learn anti-aliasing filters?

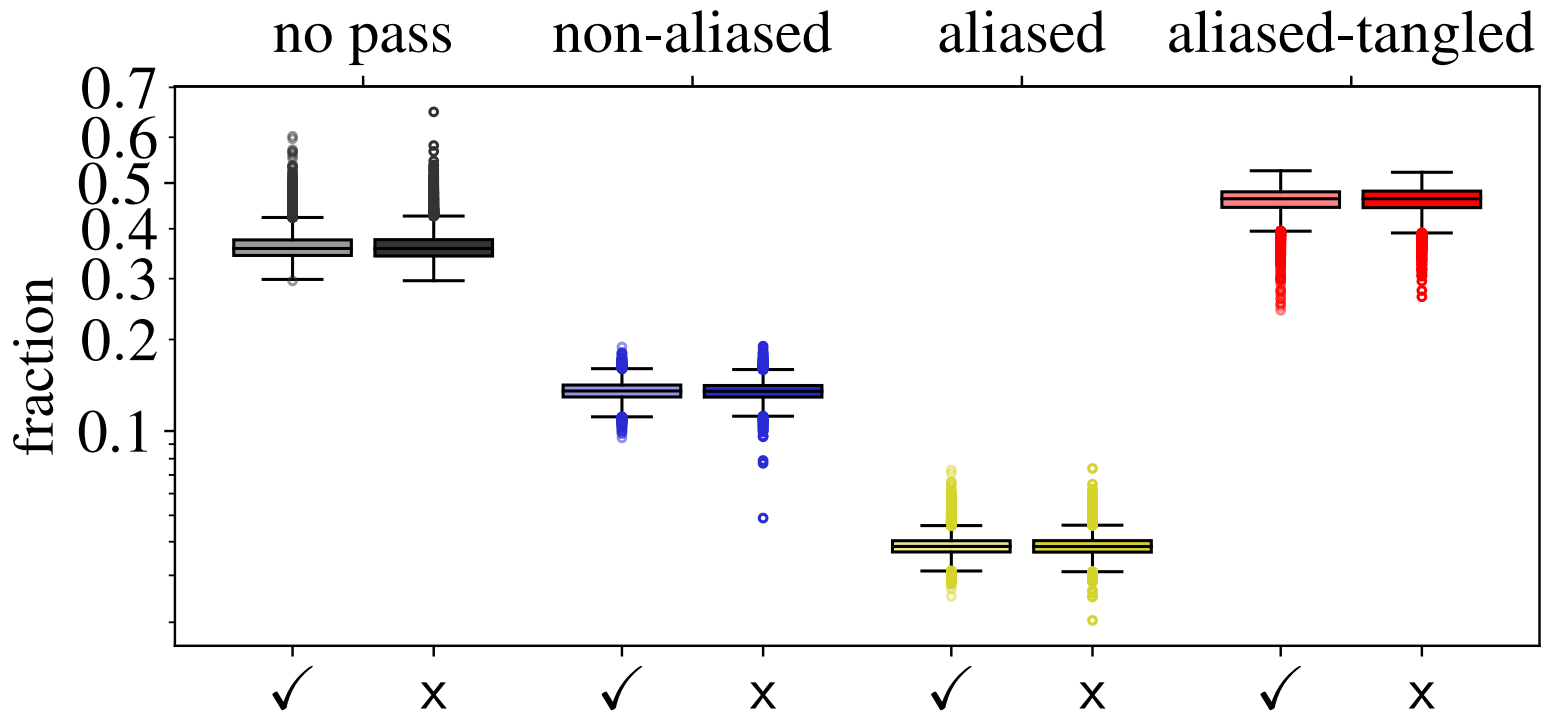

Figure: **Correct (x) vs incorrect (✓)** classified examples in ImageNet test set.

Conclusion

- Nyquist's sampling theorem gives a *sufficient* reconstruction criterion.

Conclusion

- Nyquist's sampling theorem gives a *sufficient* reconstruction criterion.
- It is possible to reconstruct signals sampled below the Nyquist rate (i.e. compressive sensing).

Conclusion

- Nyquist's sampling theorem gives a *sufficient* reconstruction criterion.
- It is possible to reconstruct signals sampled below the Nyquist rate (i.e. compressive sensing).
- In the case of CNNs, the possibility of reconstruction is simplified by the channel redundancy.

Thank you!

Contact info:


antonio.horta.ribeiro@it.uu.se
thomas.schon@it.uu.se


@ahortaribeiro


antonior92.github.io
user.it.uu.se/~thosc112


github.com/antonior92
github.com/thomasschon