University of Amsterdam

AMLAB Amsterdam Machine Learning Lab

Group Equivariant Deep Learning

Lecture 1 - Regular group convolutions

Lecture 1.4 - SE(2) Equivariant NN Example | With histopathology images

Visual example for roto-translation equivariance (SE(2))

Erik Bekkers, Amsterdam Machine Learning Lab, University of Amsterdam

This mini-course serves as a module with the UvA Master AI course Deep Learning 2 https://uvadl2c.github.io/

Architecture for rotation invariant mitotic cell detection

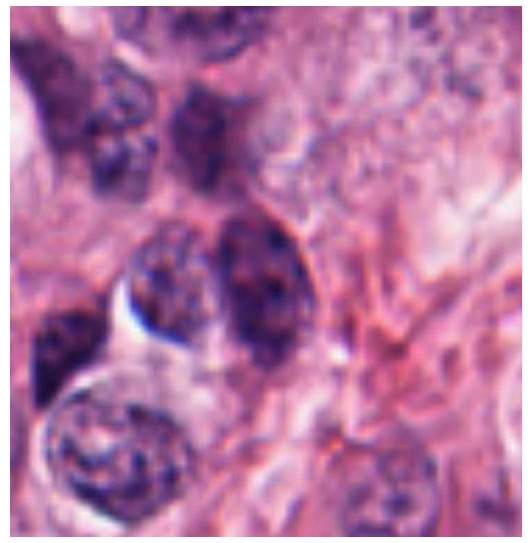

Architecture for rotation invariant mitotic cell detection

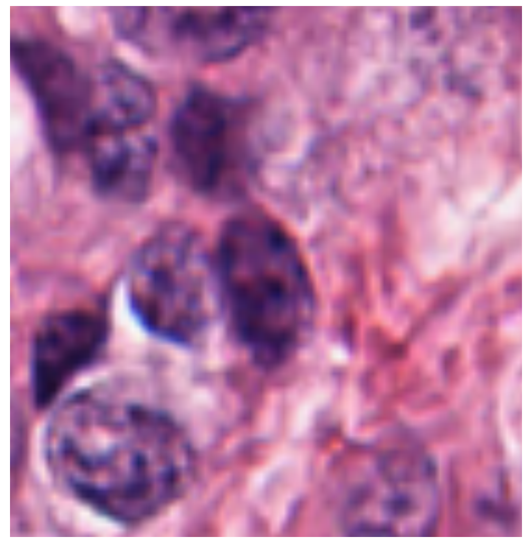

“normal” (0)
vs
“mitotic” (1)

Architecture for rotation invariant mitotic cell detection

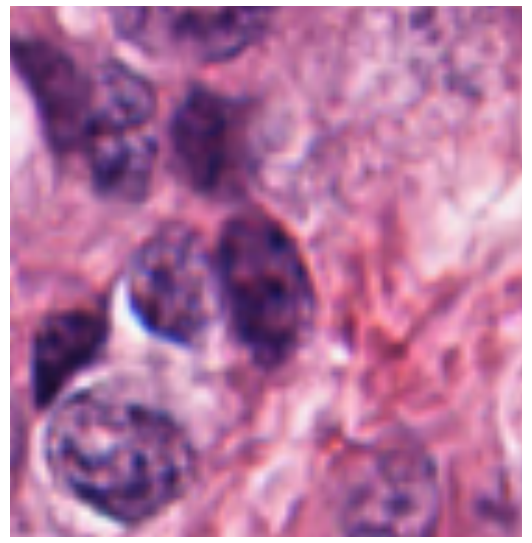

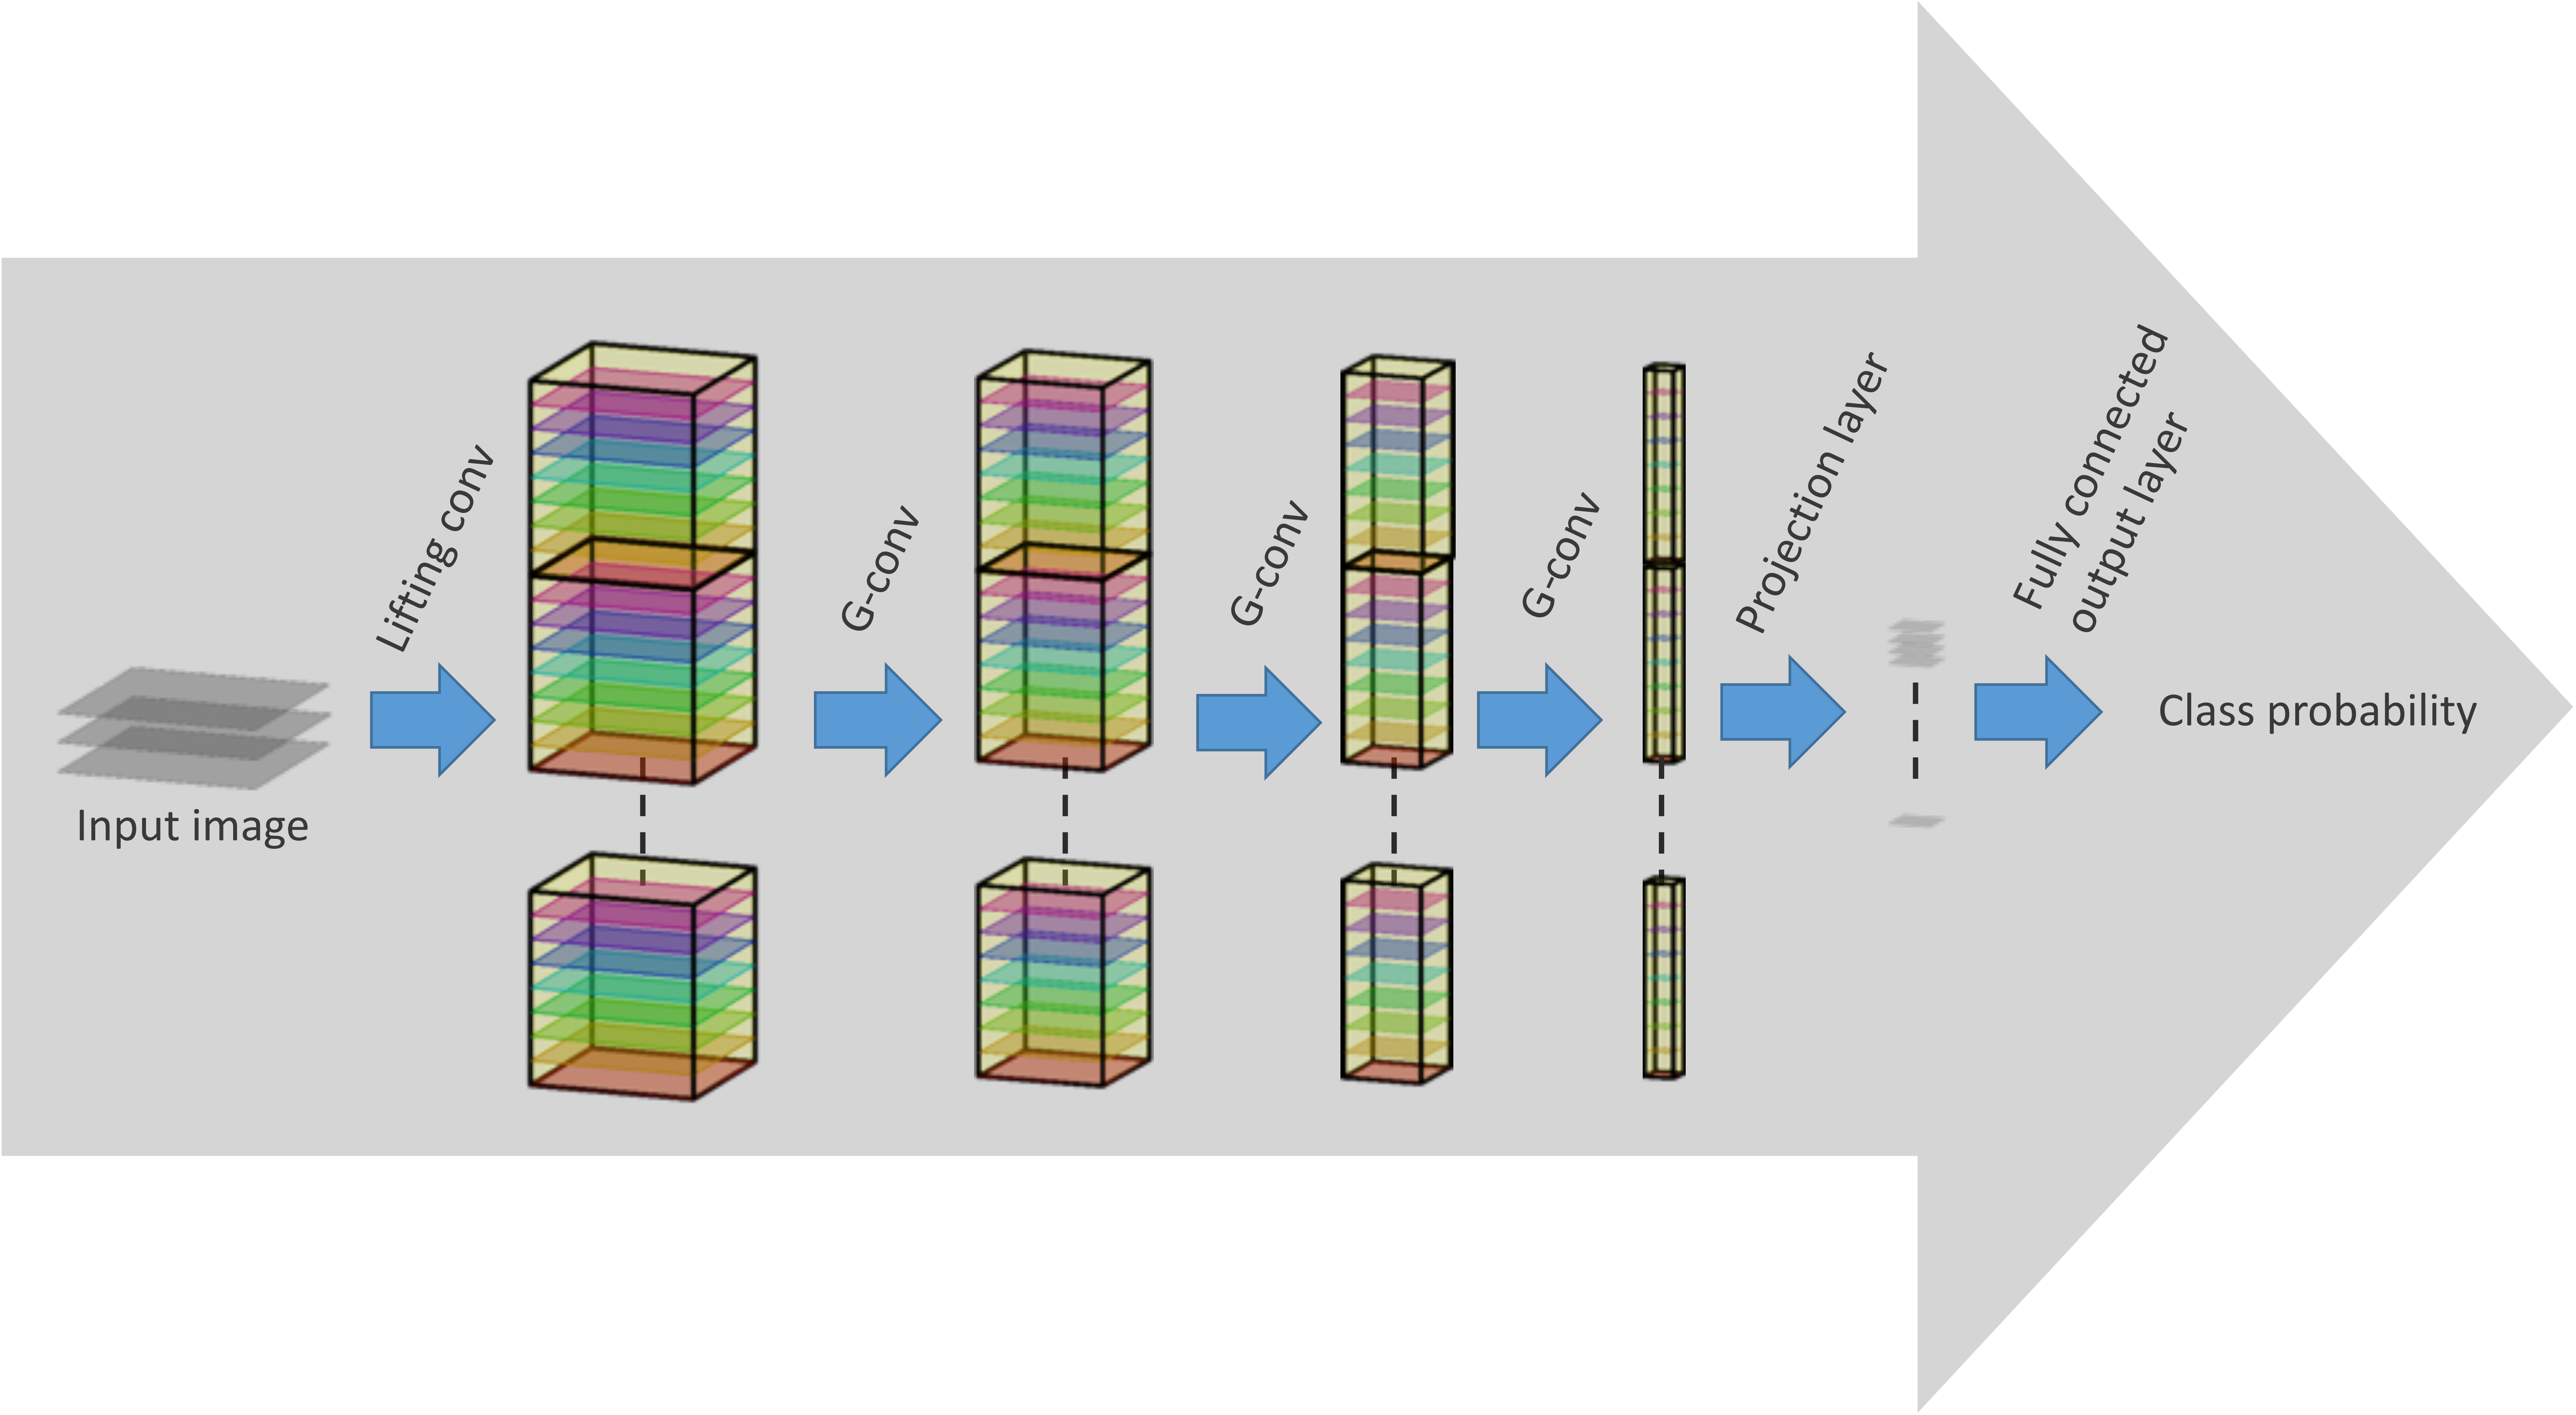

"normal" (0)
vs
"mitotic" (1)

Architecture for rotation invariant mitotic cell detection

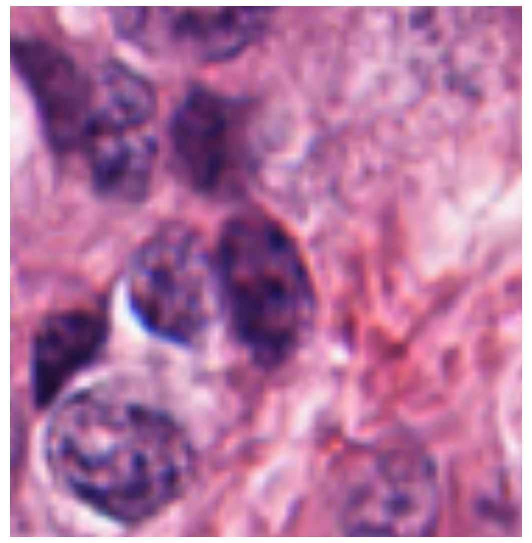

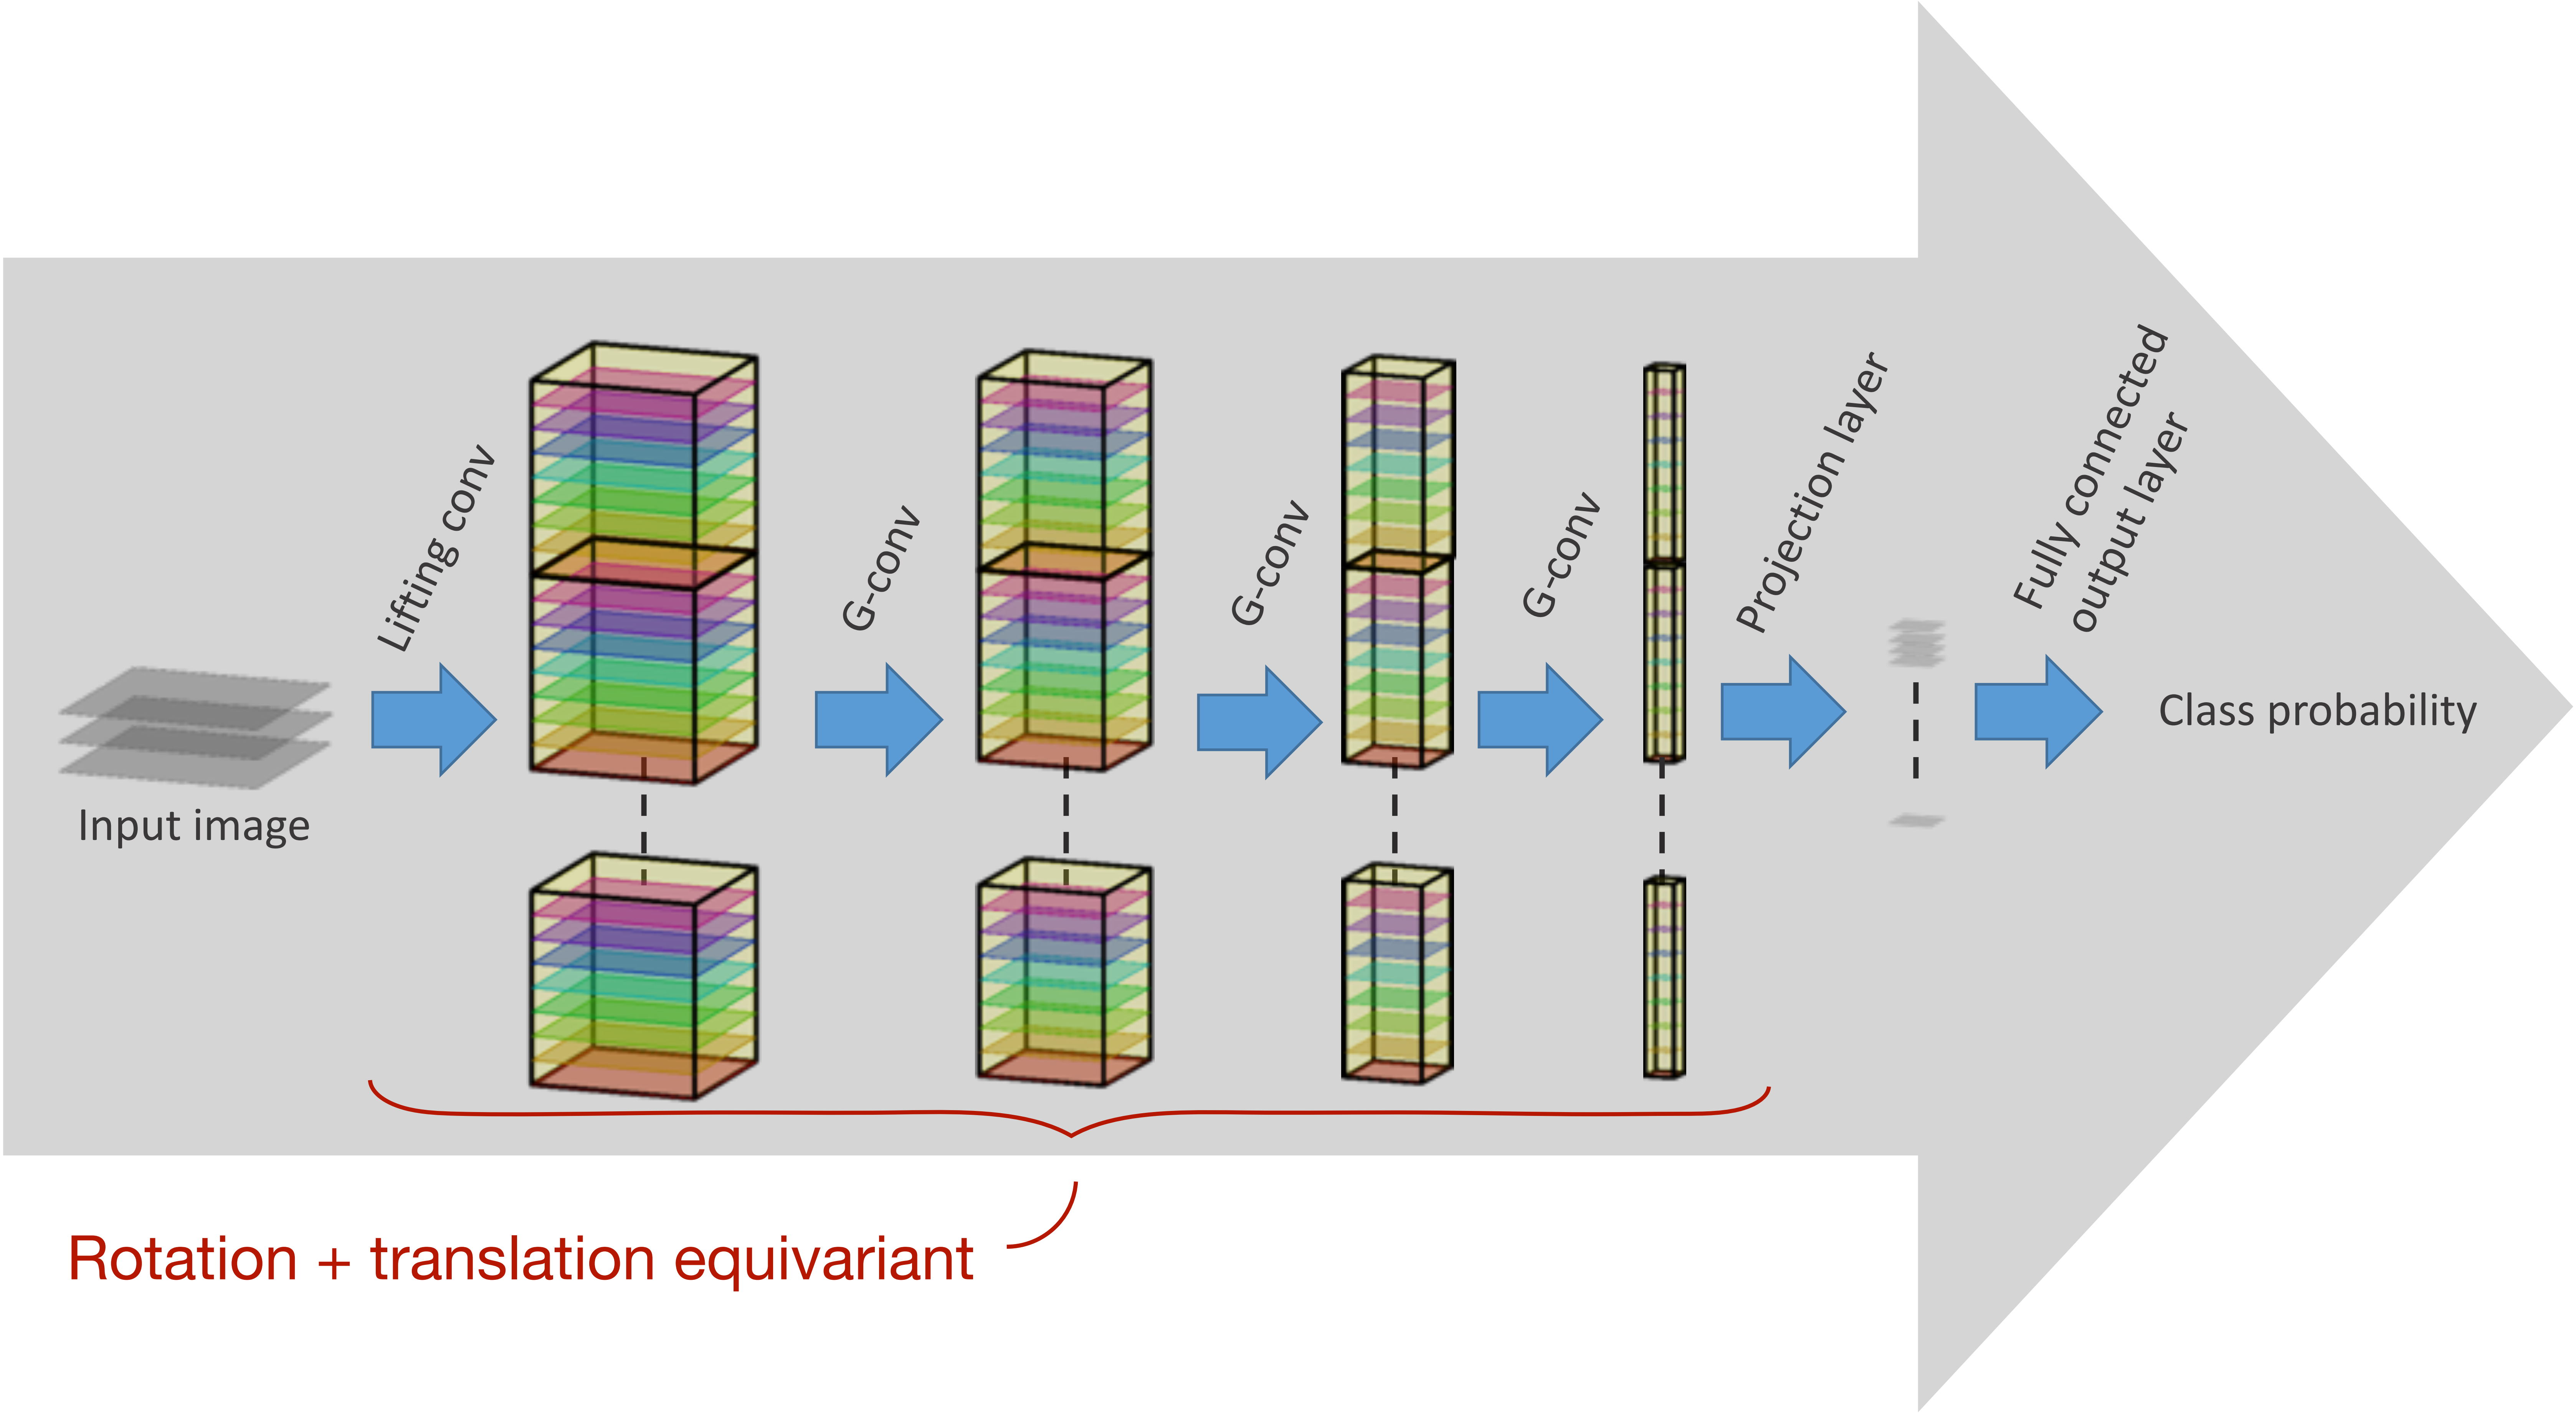

"normal" (0)
vs
"mitotic" (1)

Architecture for rotation invariant mitotic cell detection

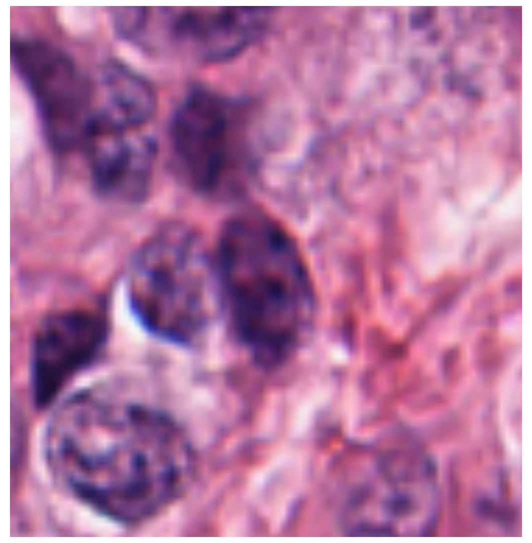

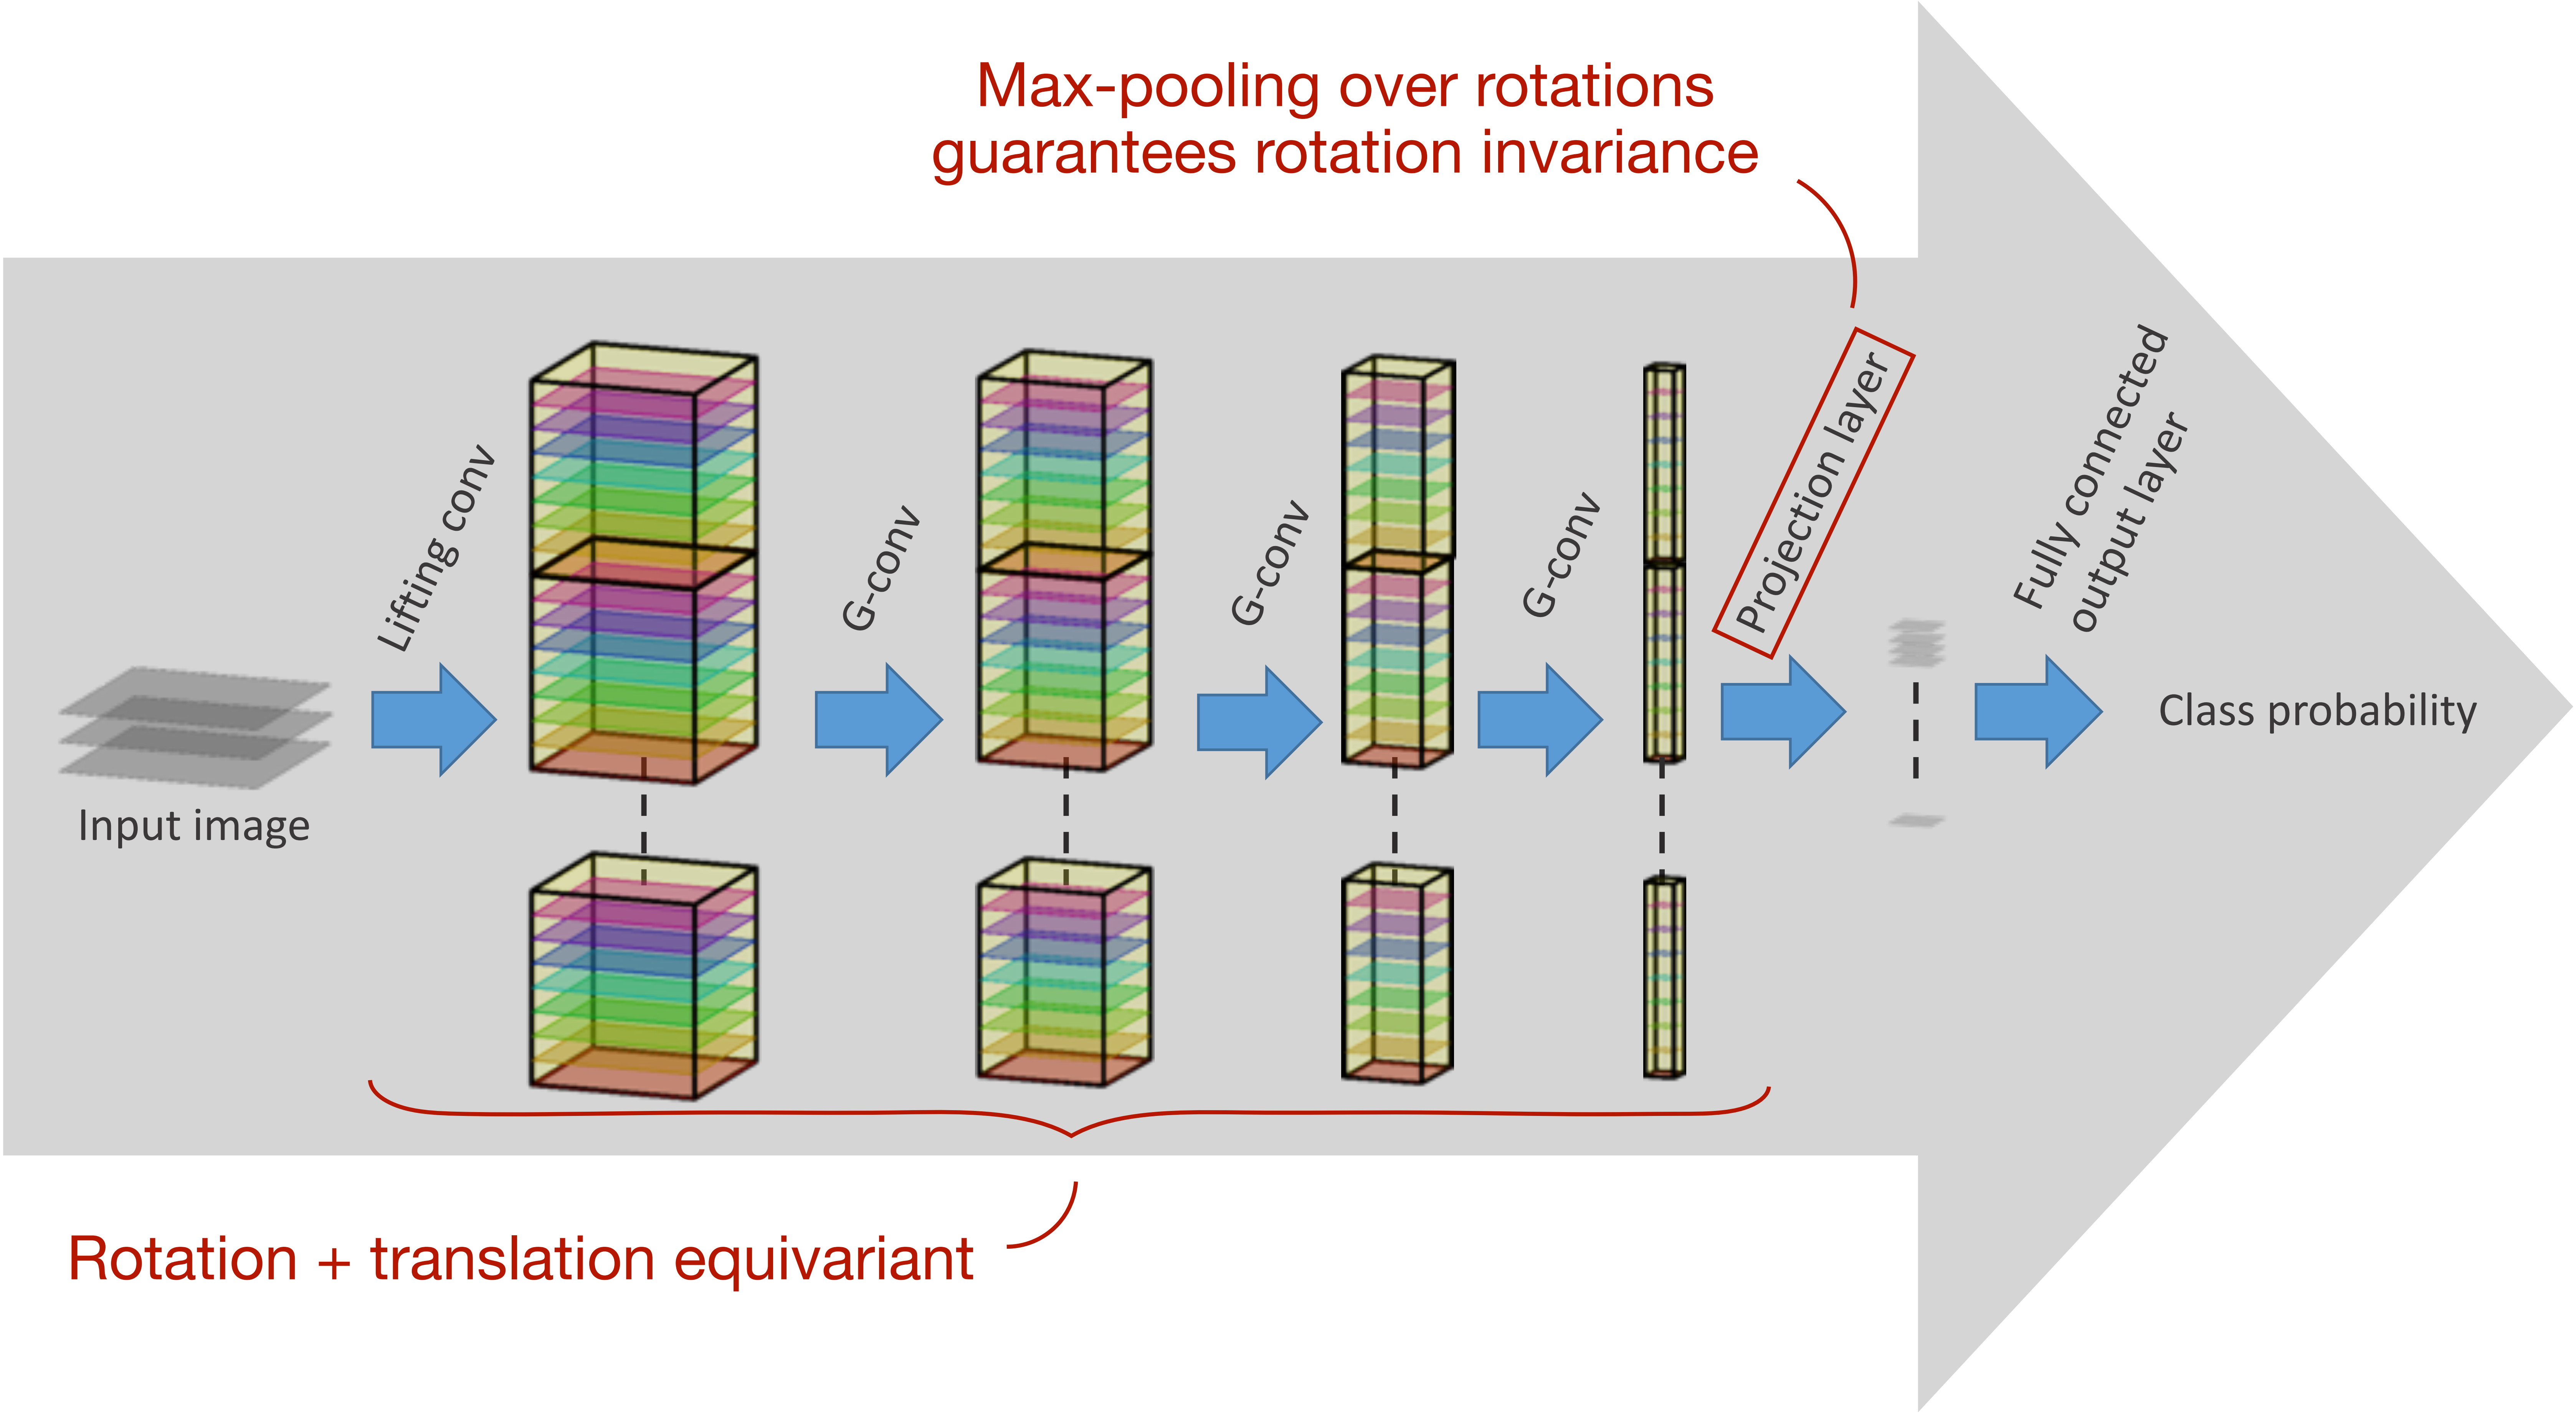

“normal” (0)
vs
“mitotic” (1)

Bekkers & Lafarge et al. MICCAI 2018

Architecture for rotation invariant mitotic cell detection

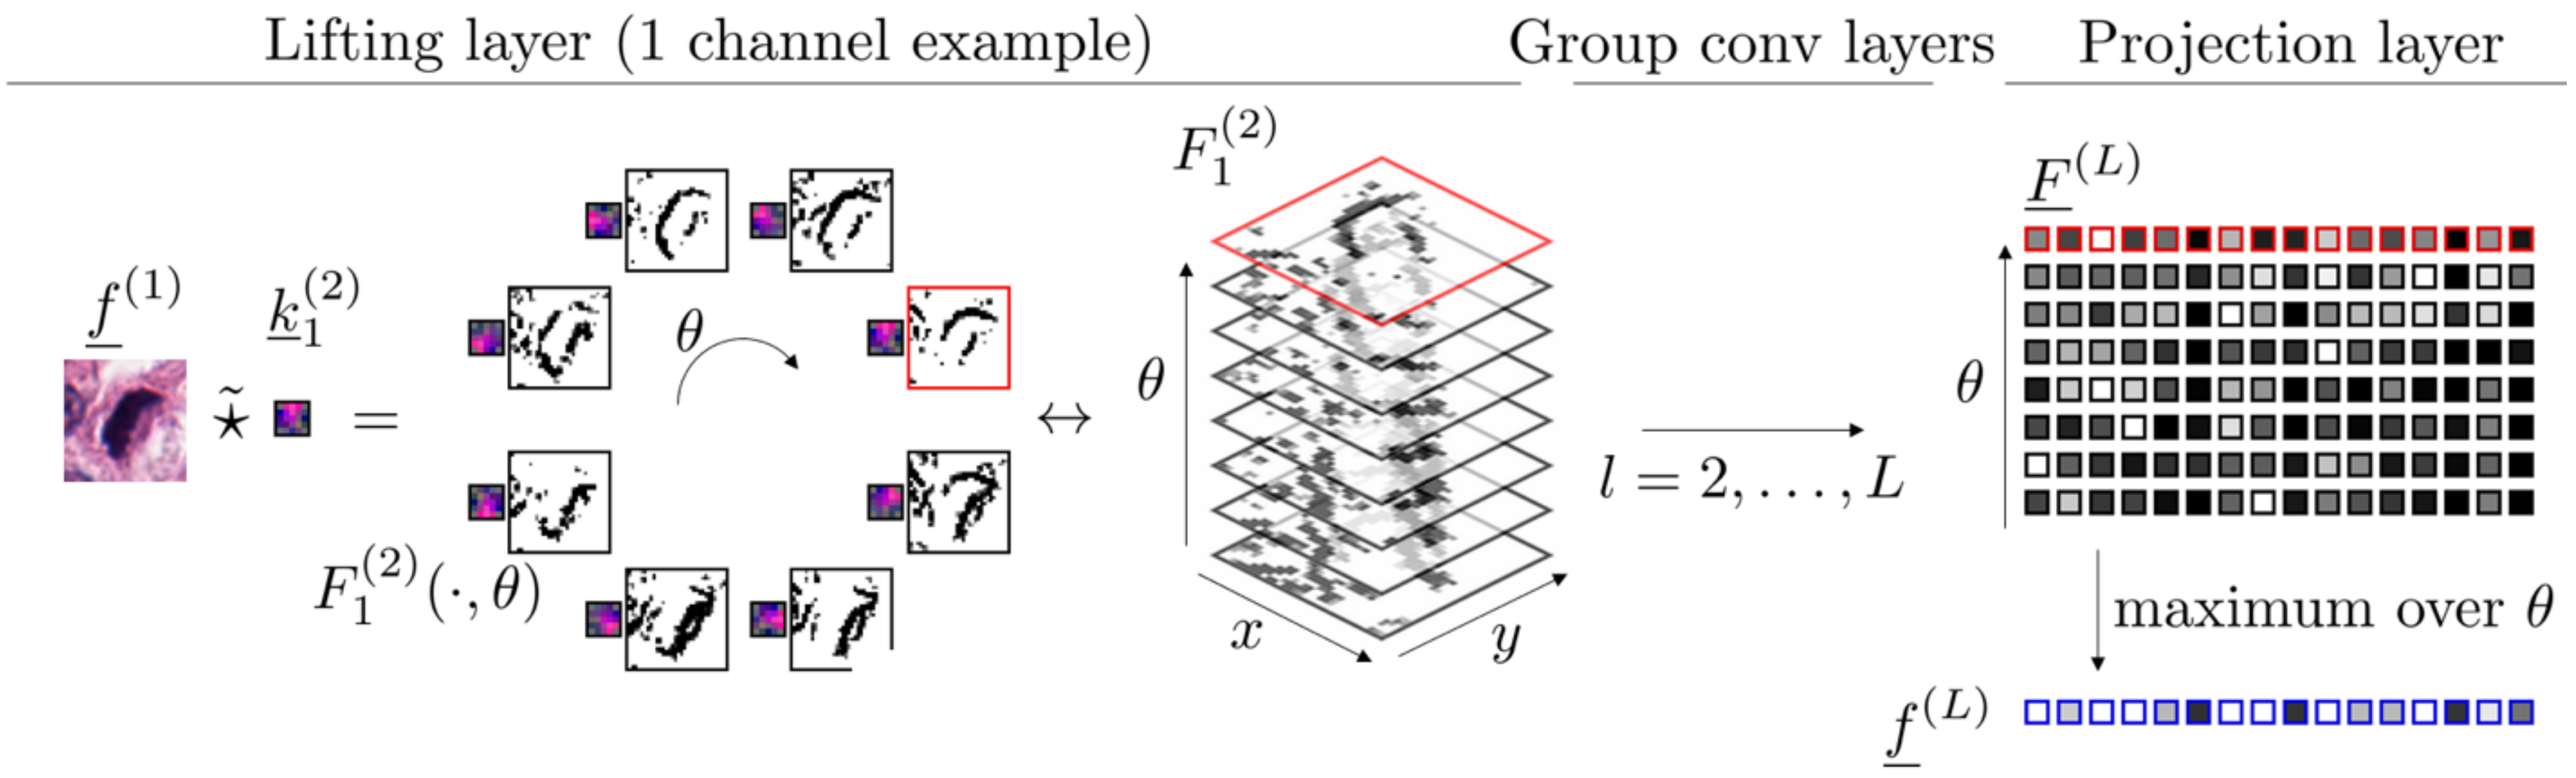

Bekkers & Lafarge et al. MICCAI 2018

Architecture for rotation invariant mitotic cell detection

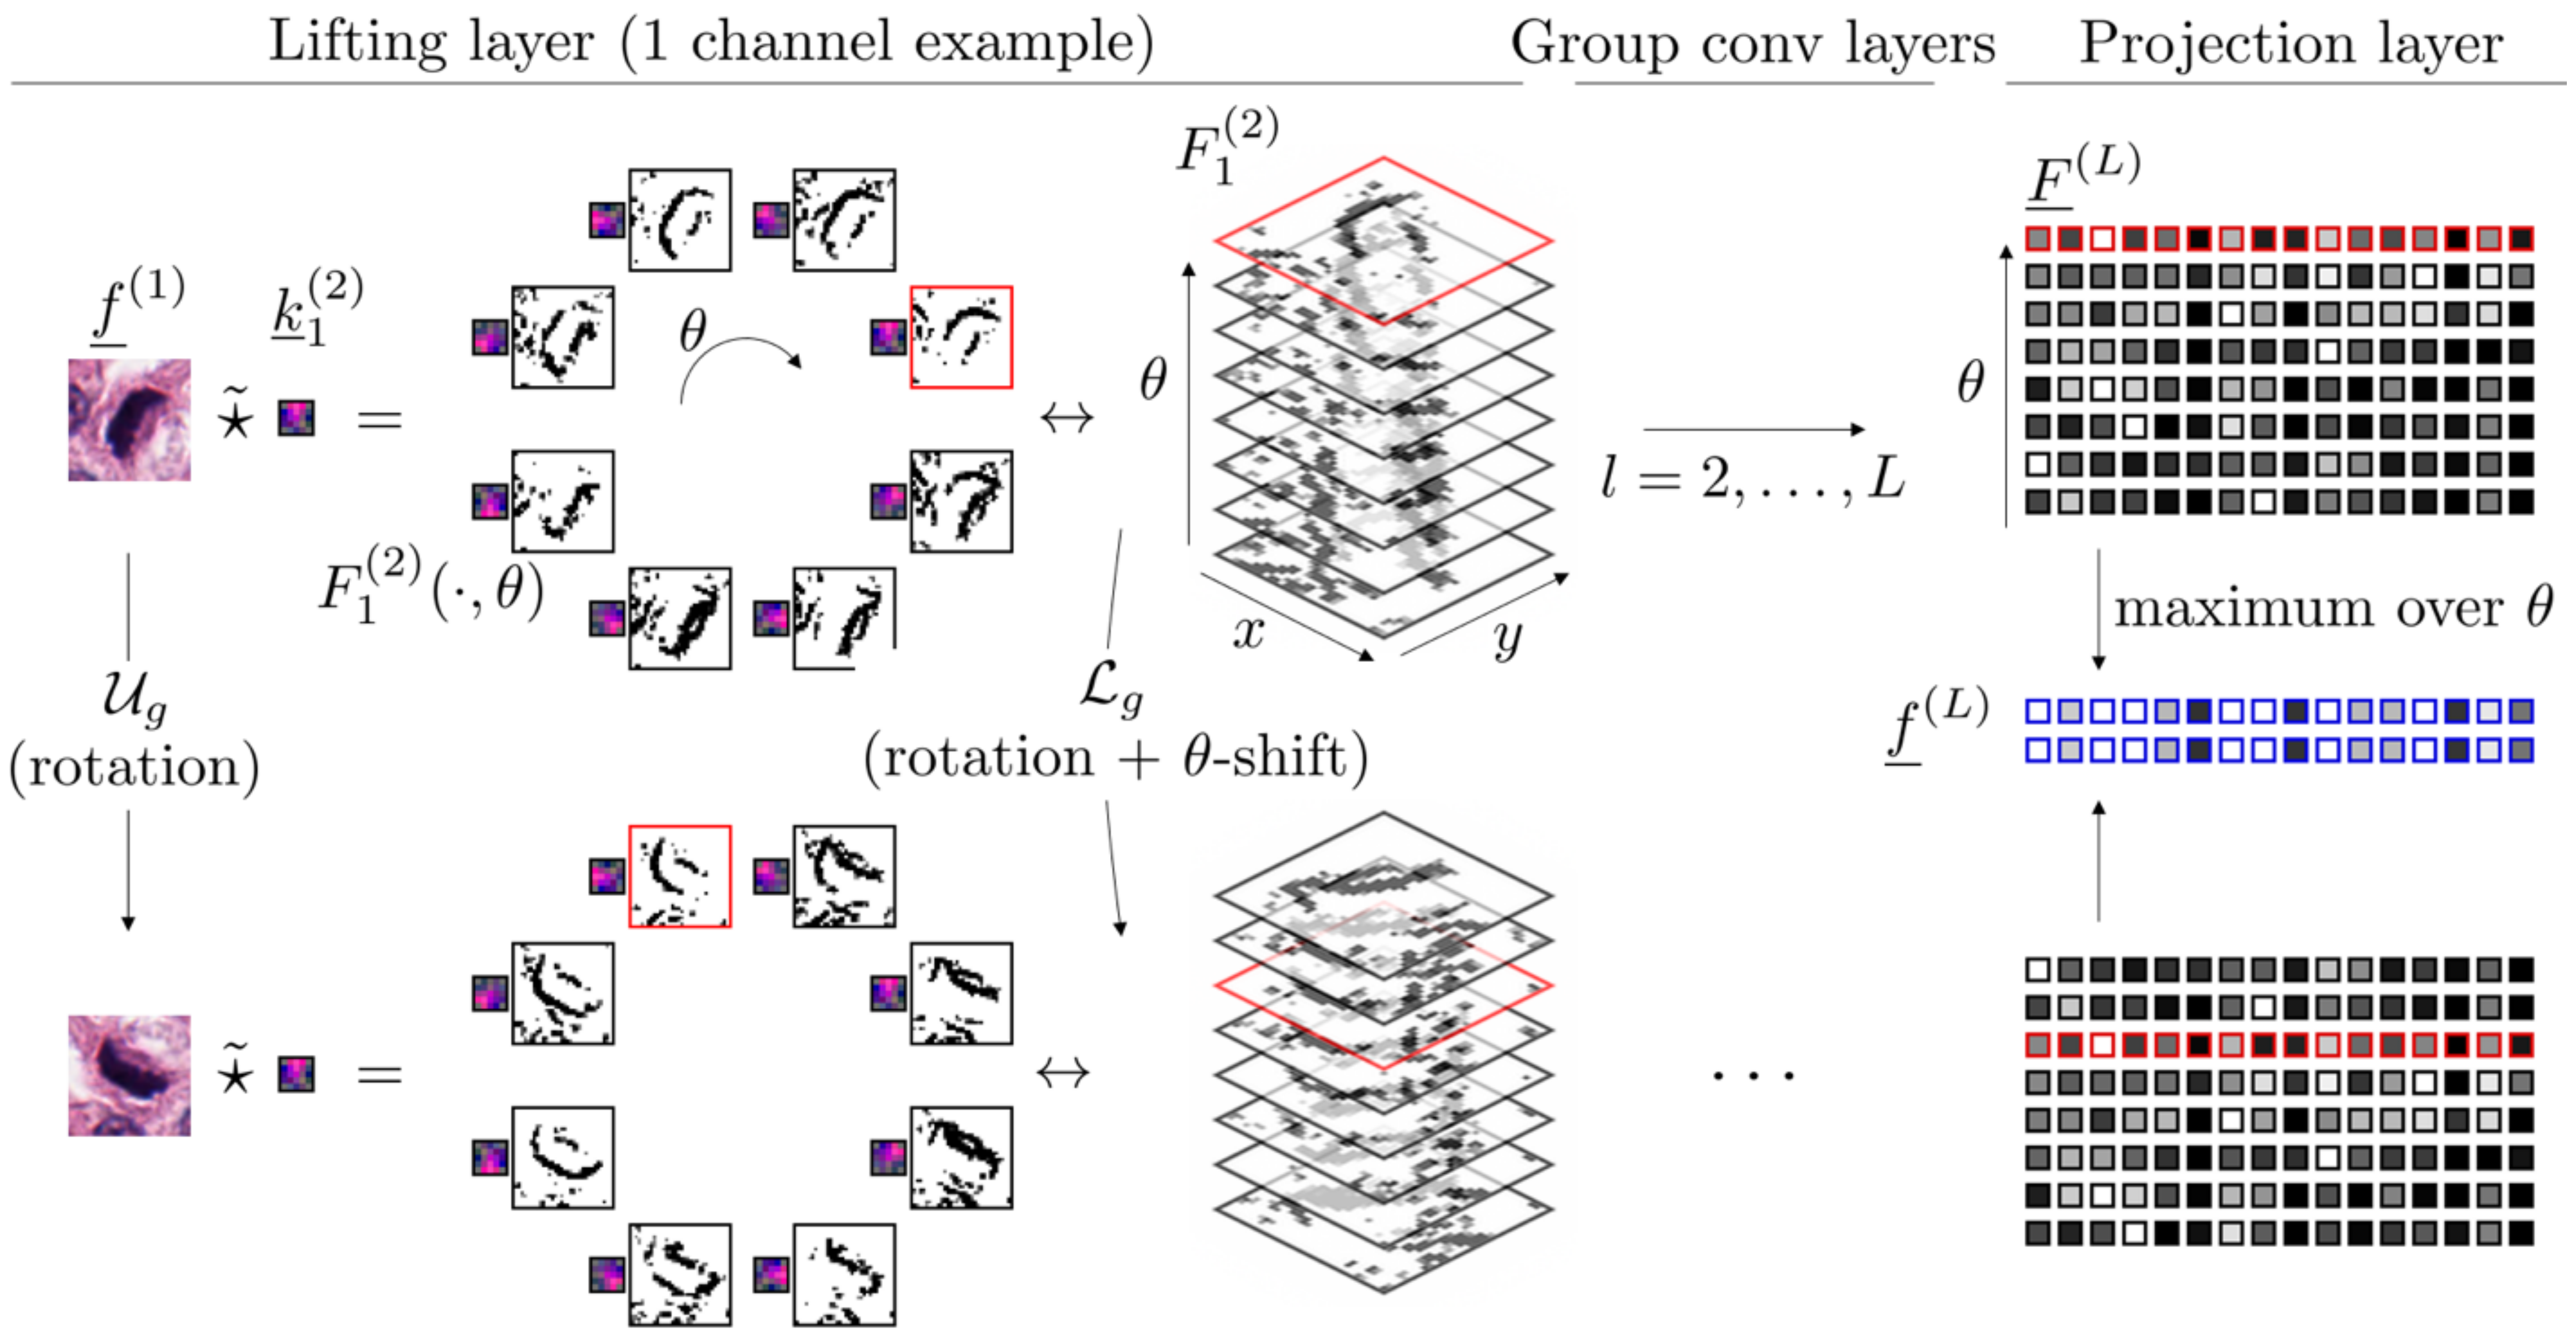

Architecture for rotation invariant mitotic cell detection

Bekkers & Lafarge et al. MICCAI 2018

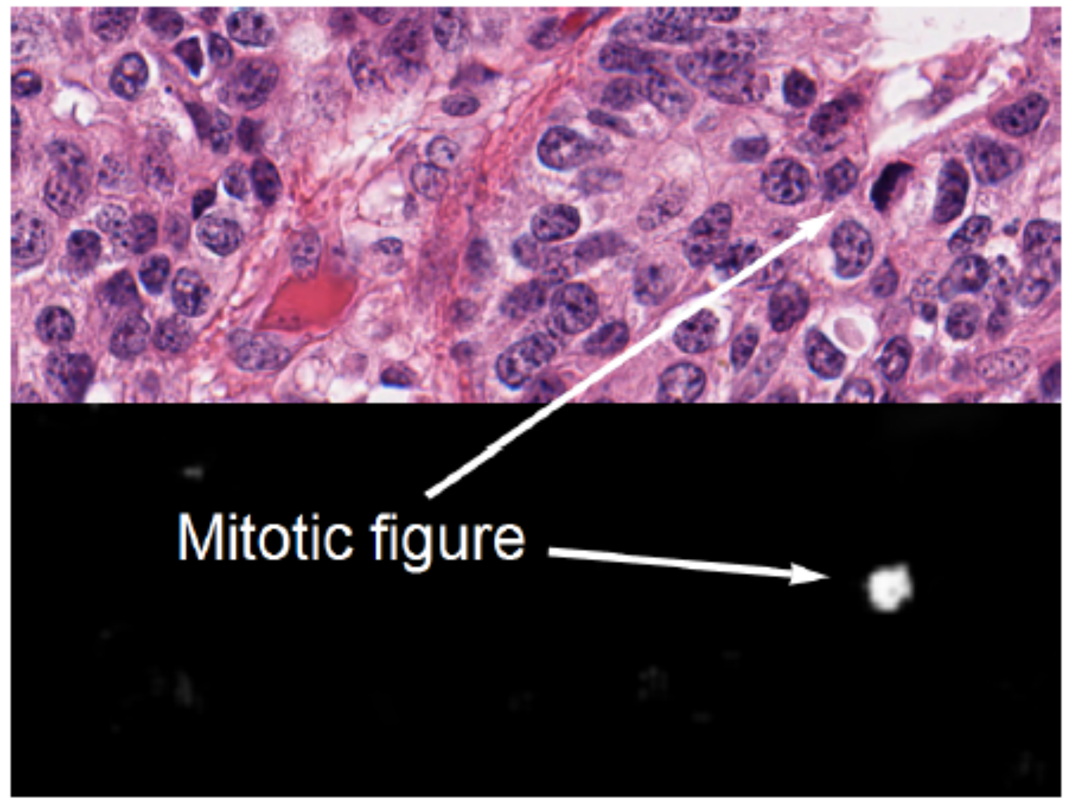

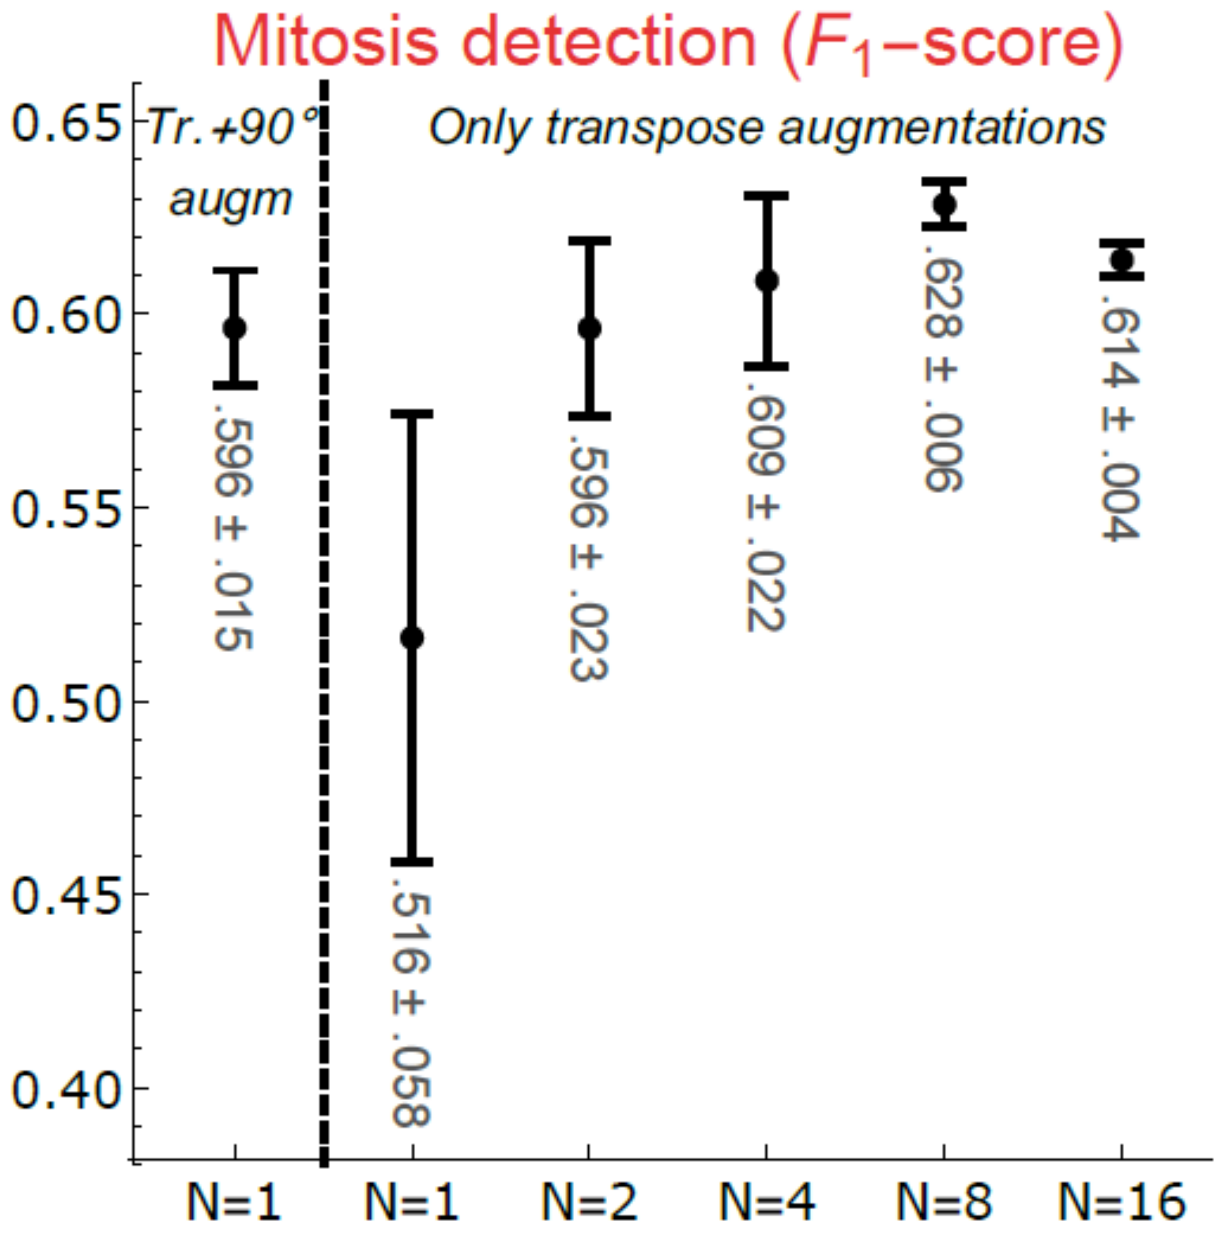

G-CNNs without data-augmentation
outperform
CNNs with data-augmentation

Architecture for rotation invariant mitotic cell detection

Bekkers & Lafarge et al. MICCAI 2018

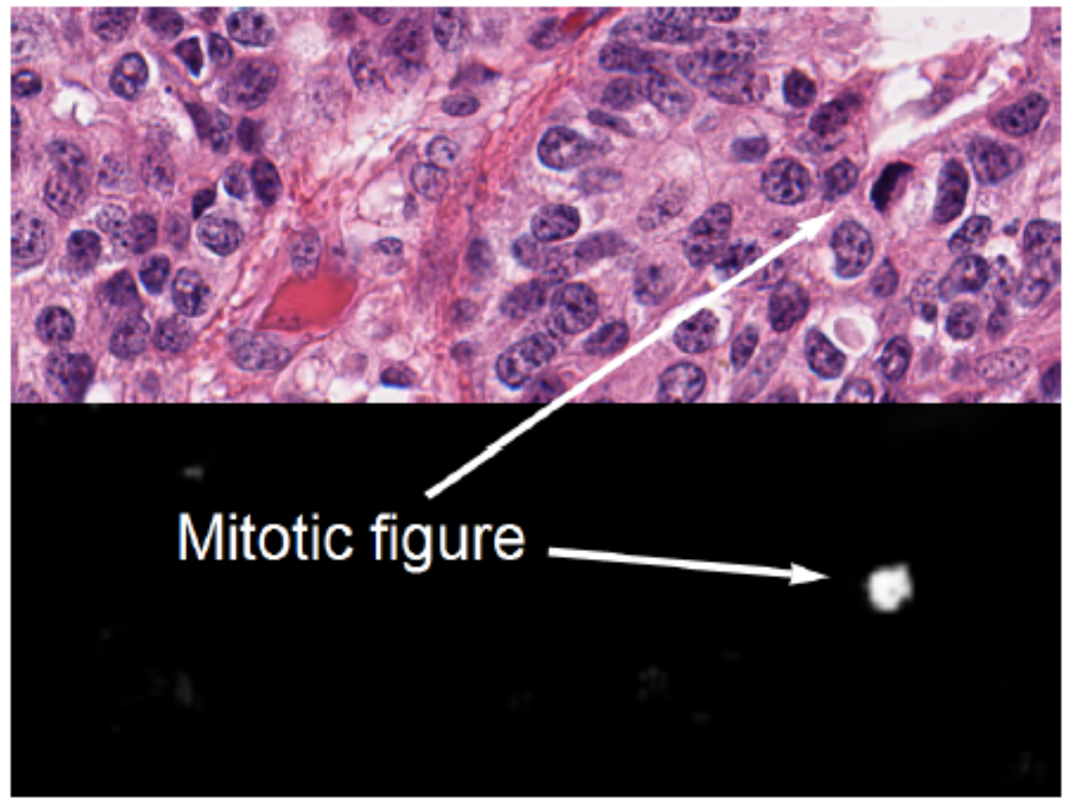

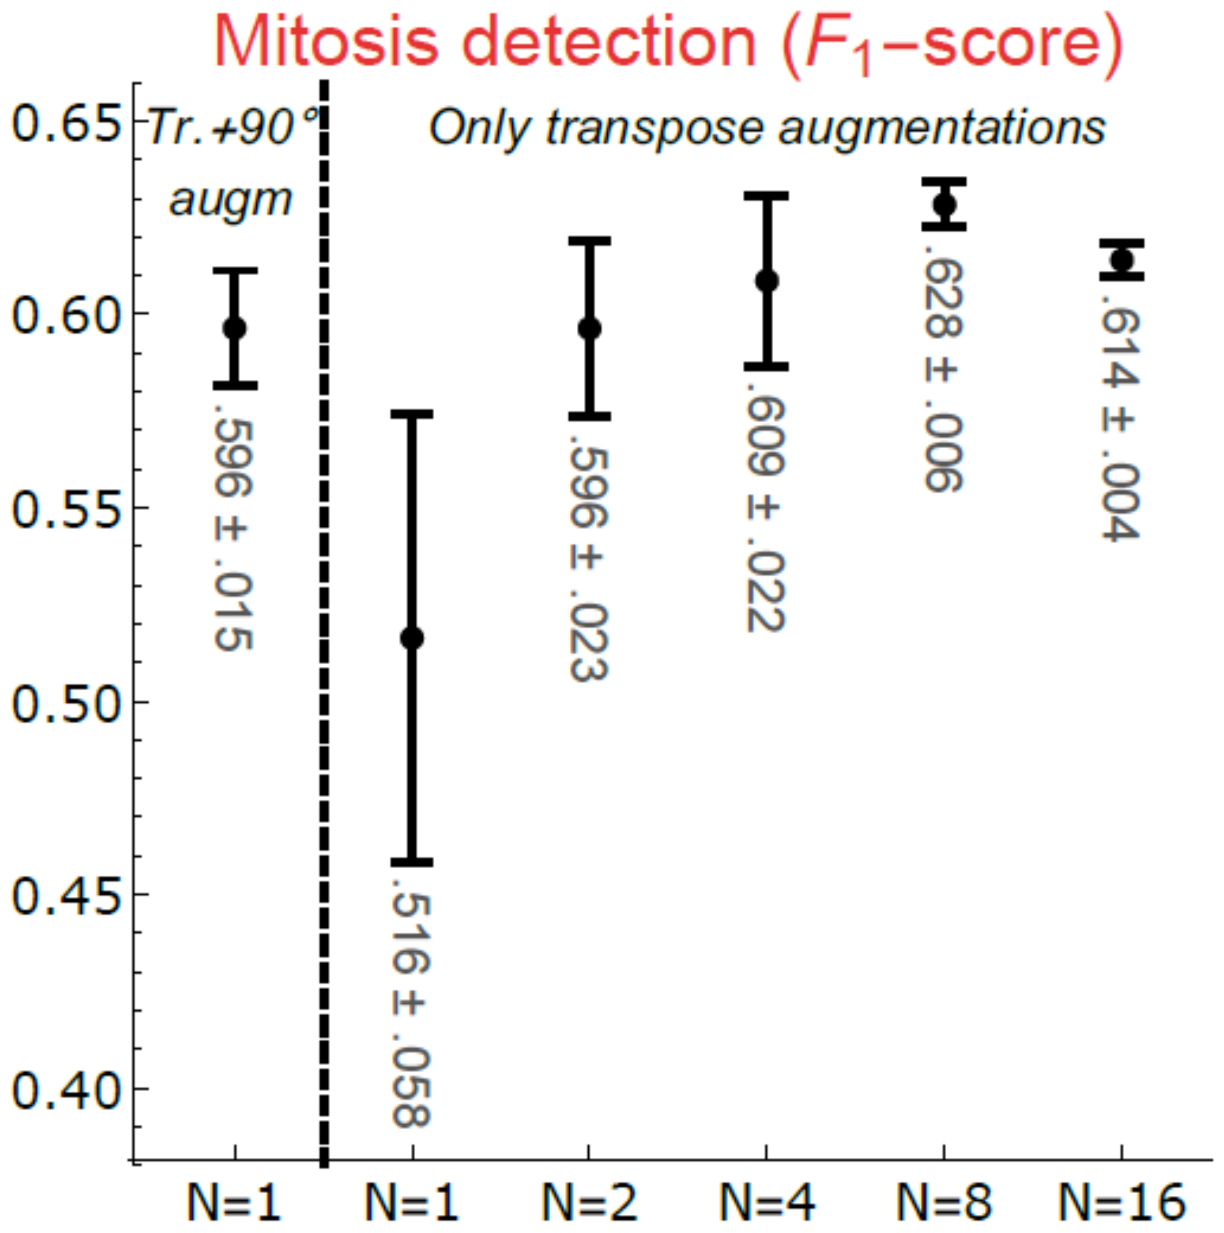

Lafarge et al. MedIA 2020

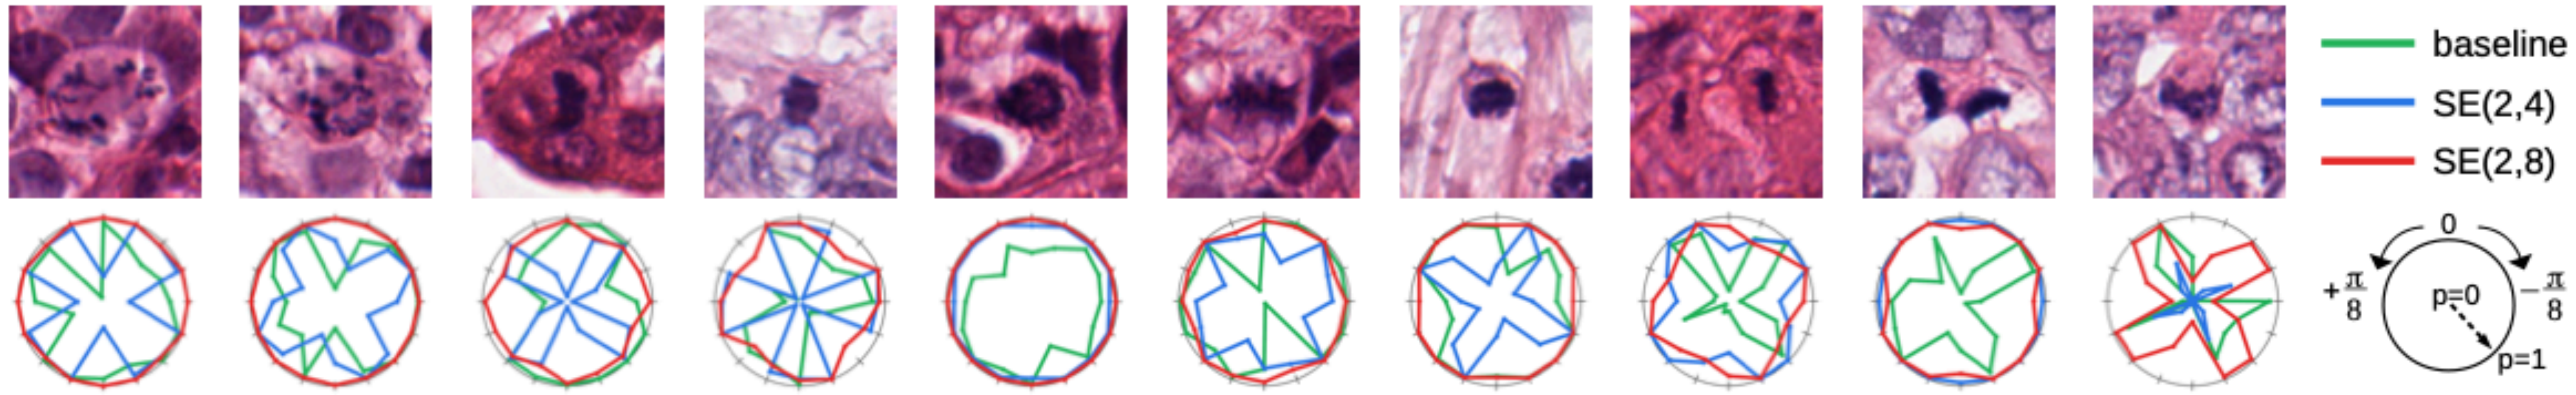

G-CNNs guarantee geometric stability. They are ***robust to input distortions****, regular CNNs aren't…*

G-CNNs without data-augmentation *outperform CNNs with data-augmentation*

Architecture for rotation invariant mitotic cell detection

Bekkers & Lafarge et al. MICCAI 2018

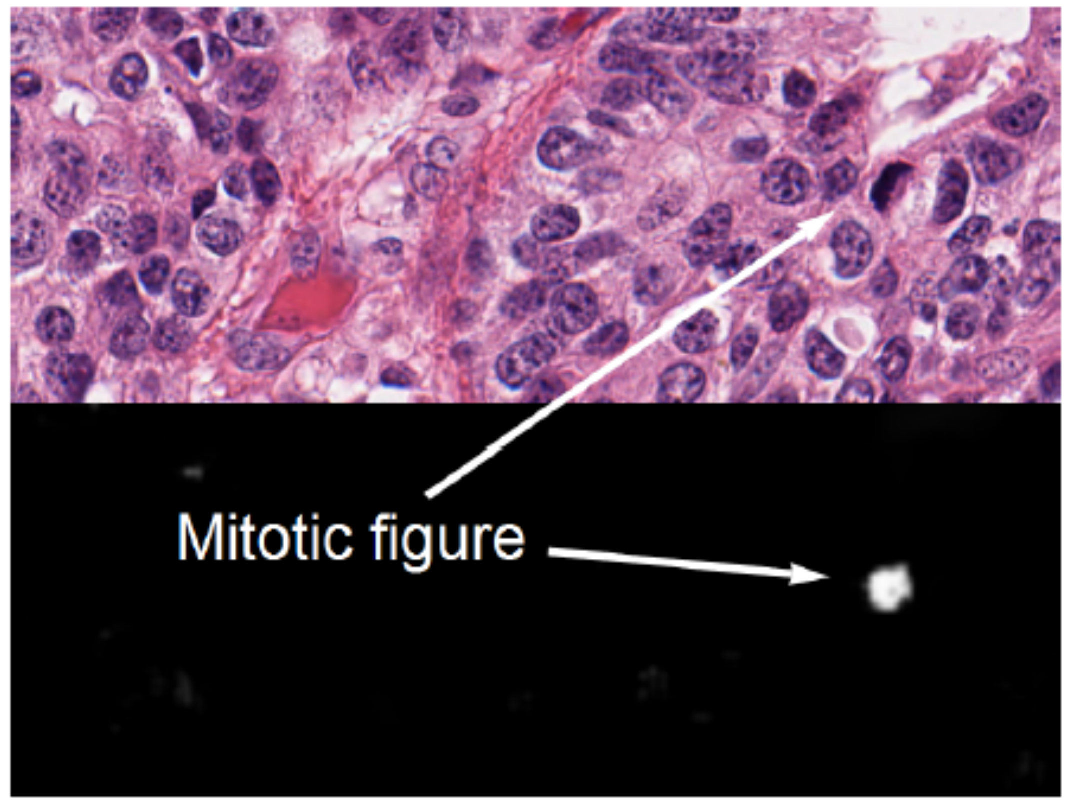

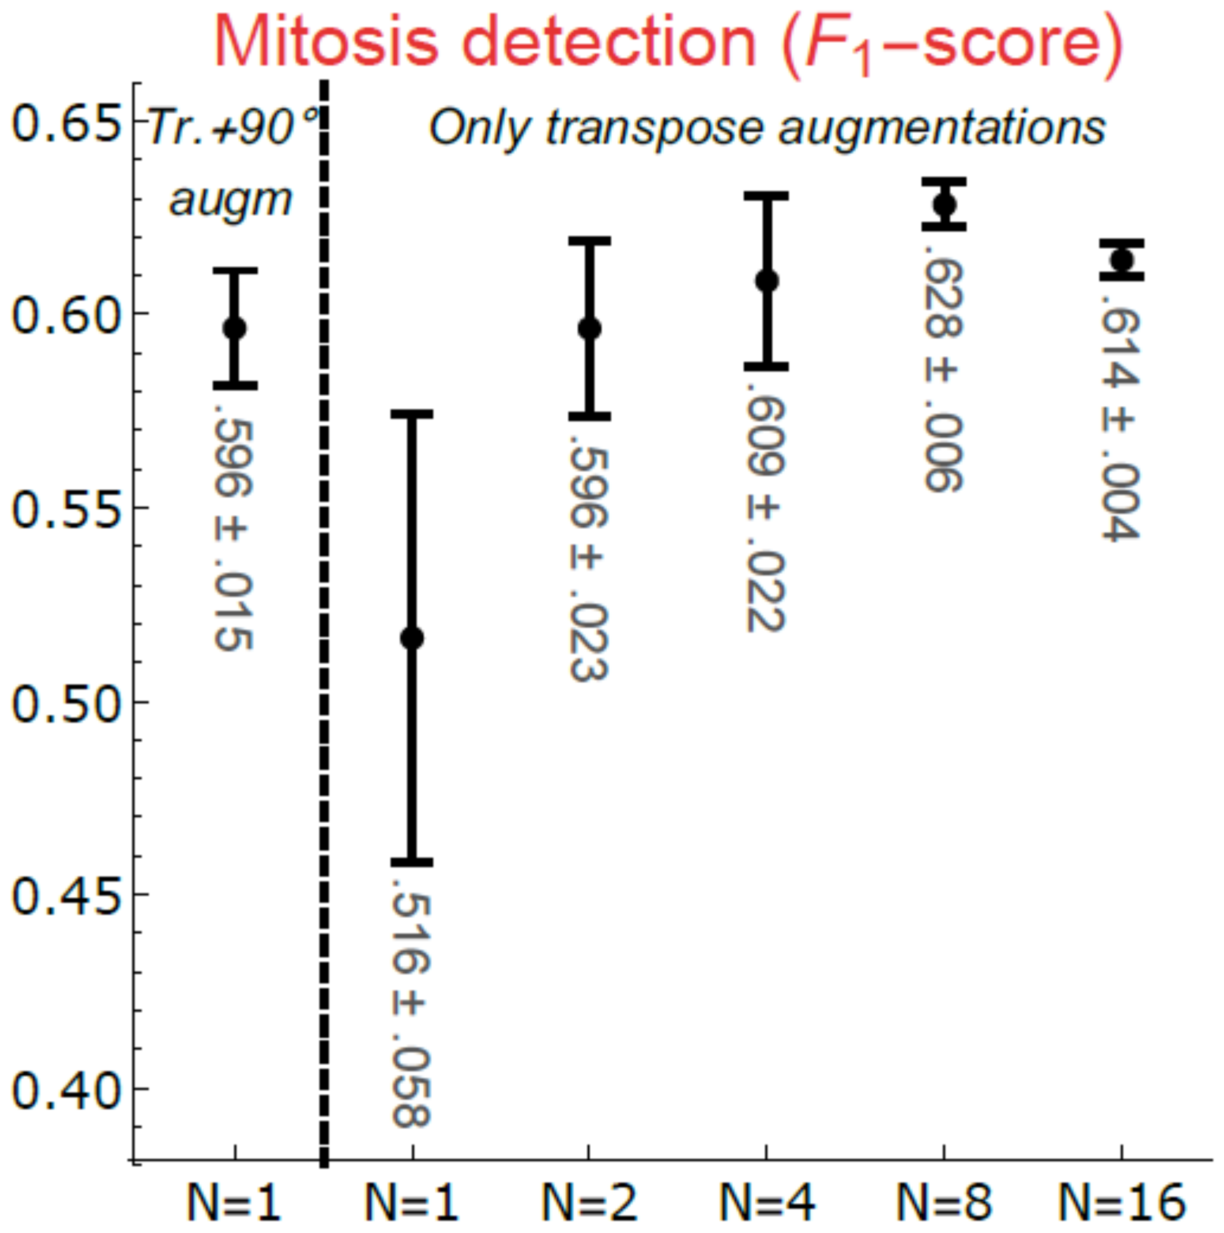

Lafarge et al. MedIA 2020

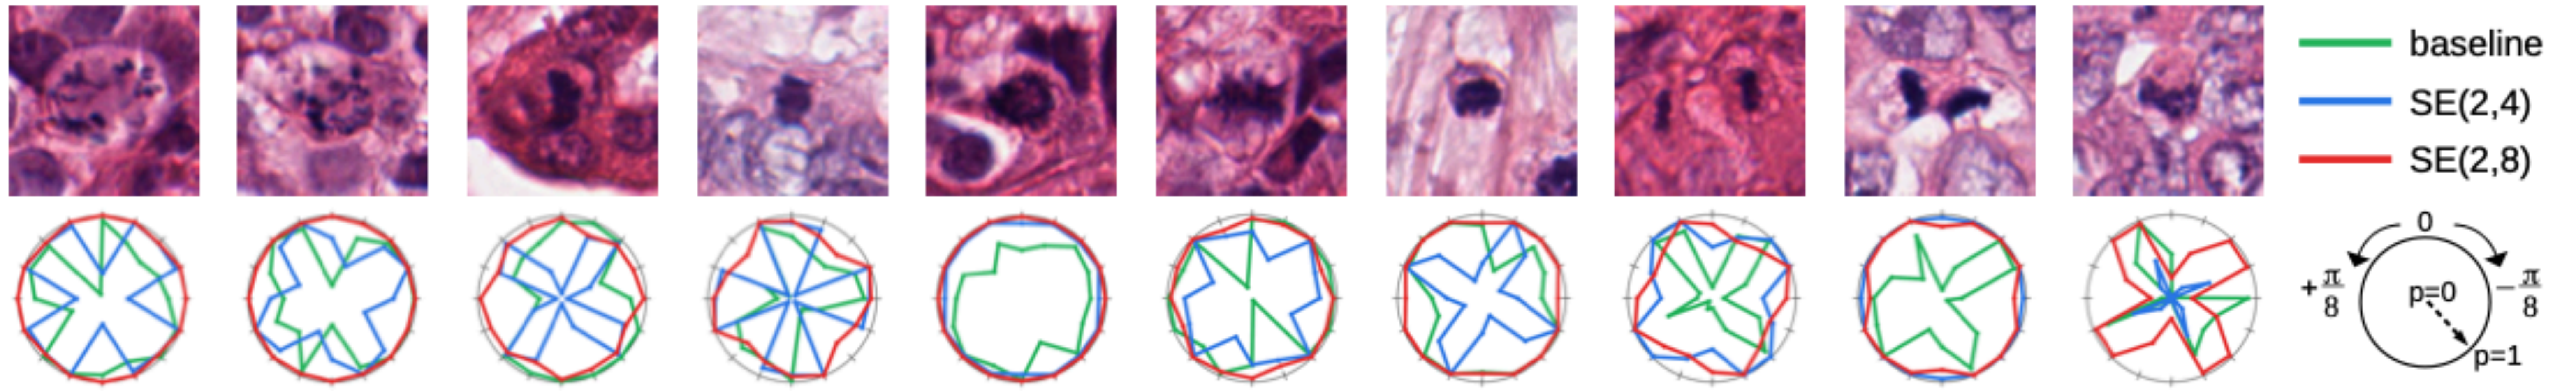

G-CNNs guarantee geometric stability. They are ***robust to input distortions****, regular CNNs aren't…*

G-CNNs are more ***sample efficient****!*
G-CNNs (25% data) > CNNs (100% data)

G-CNNs without data-augmentation
outperform
CNNs with data-augmentation

Lafarge et al. ArXiv/MedIA 2020

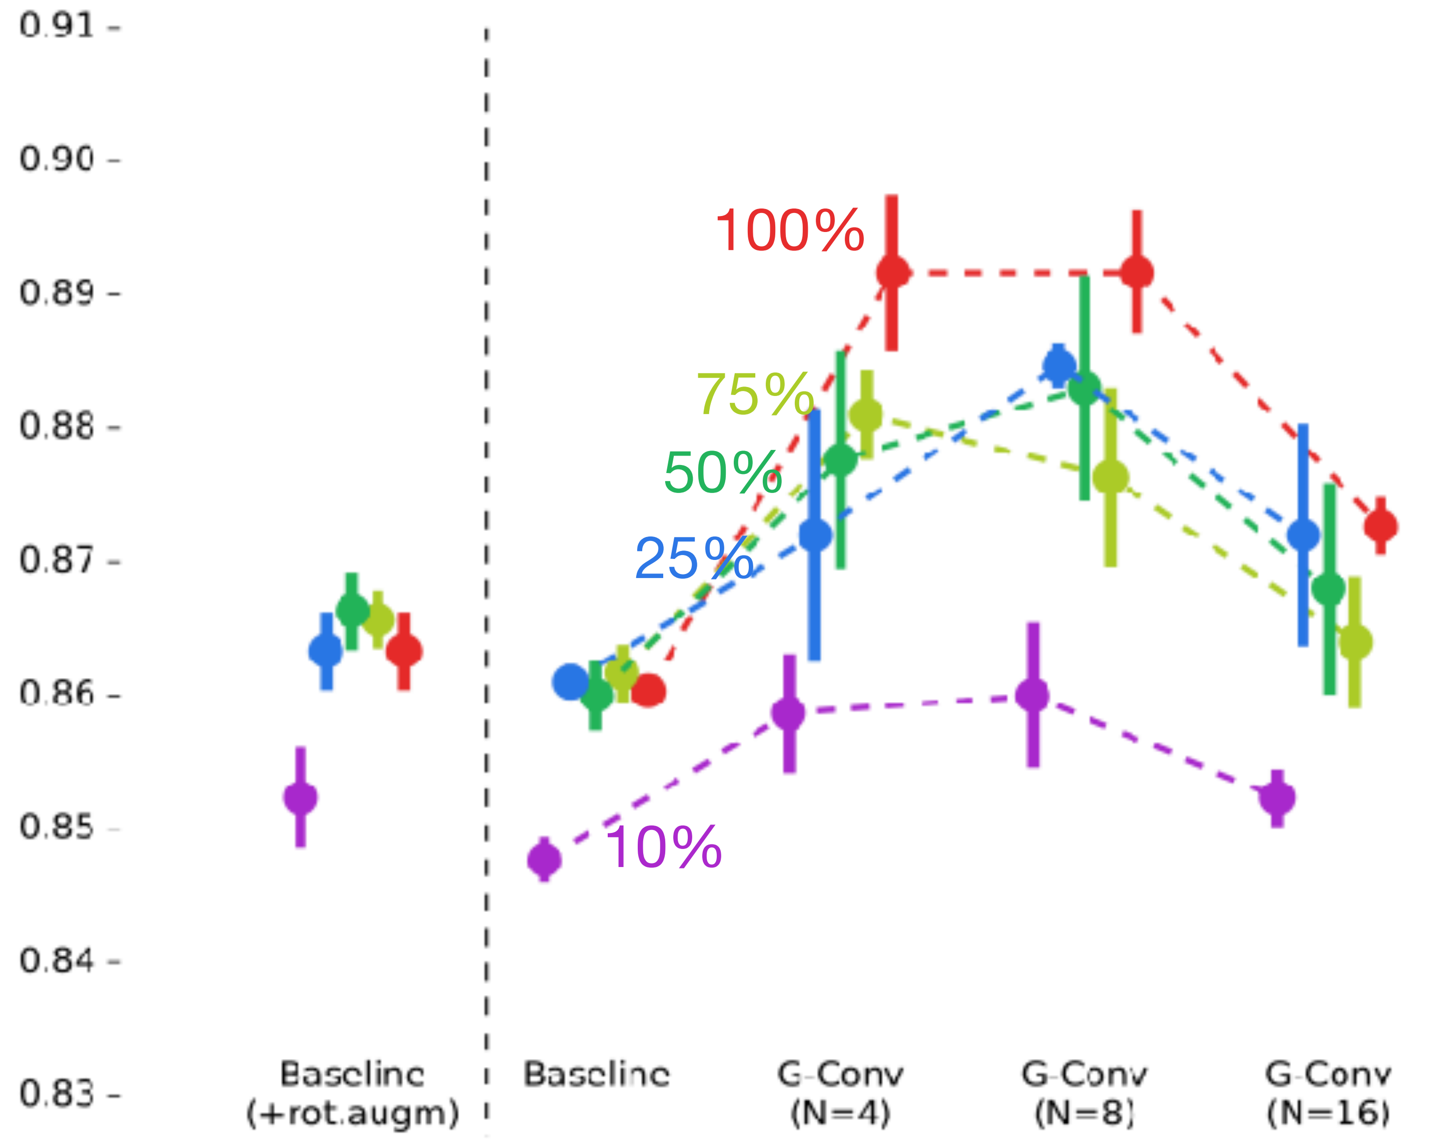

Figure 7: Mean and Standard Deviation plots summarizing the accuracy of the tumor classification models. Mean ± standard deviation is indicated. Color identifies the different data regime (red: 100%; lime: 75%; green: 50%; blue: 25%; purple: 10%).

Experiments in medical image analysis

Bekkers & Lafarge et al. MICCAI 2018

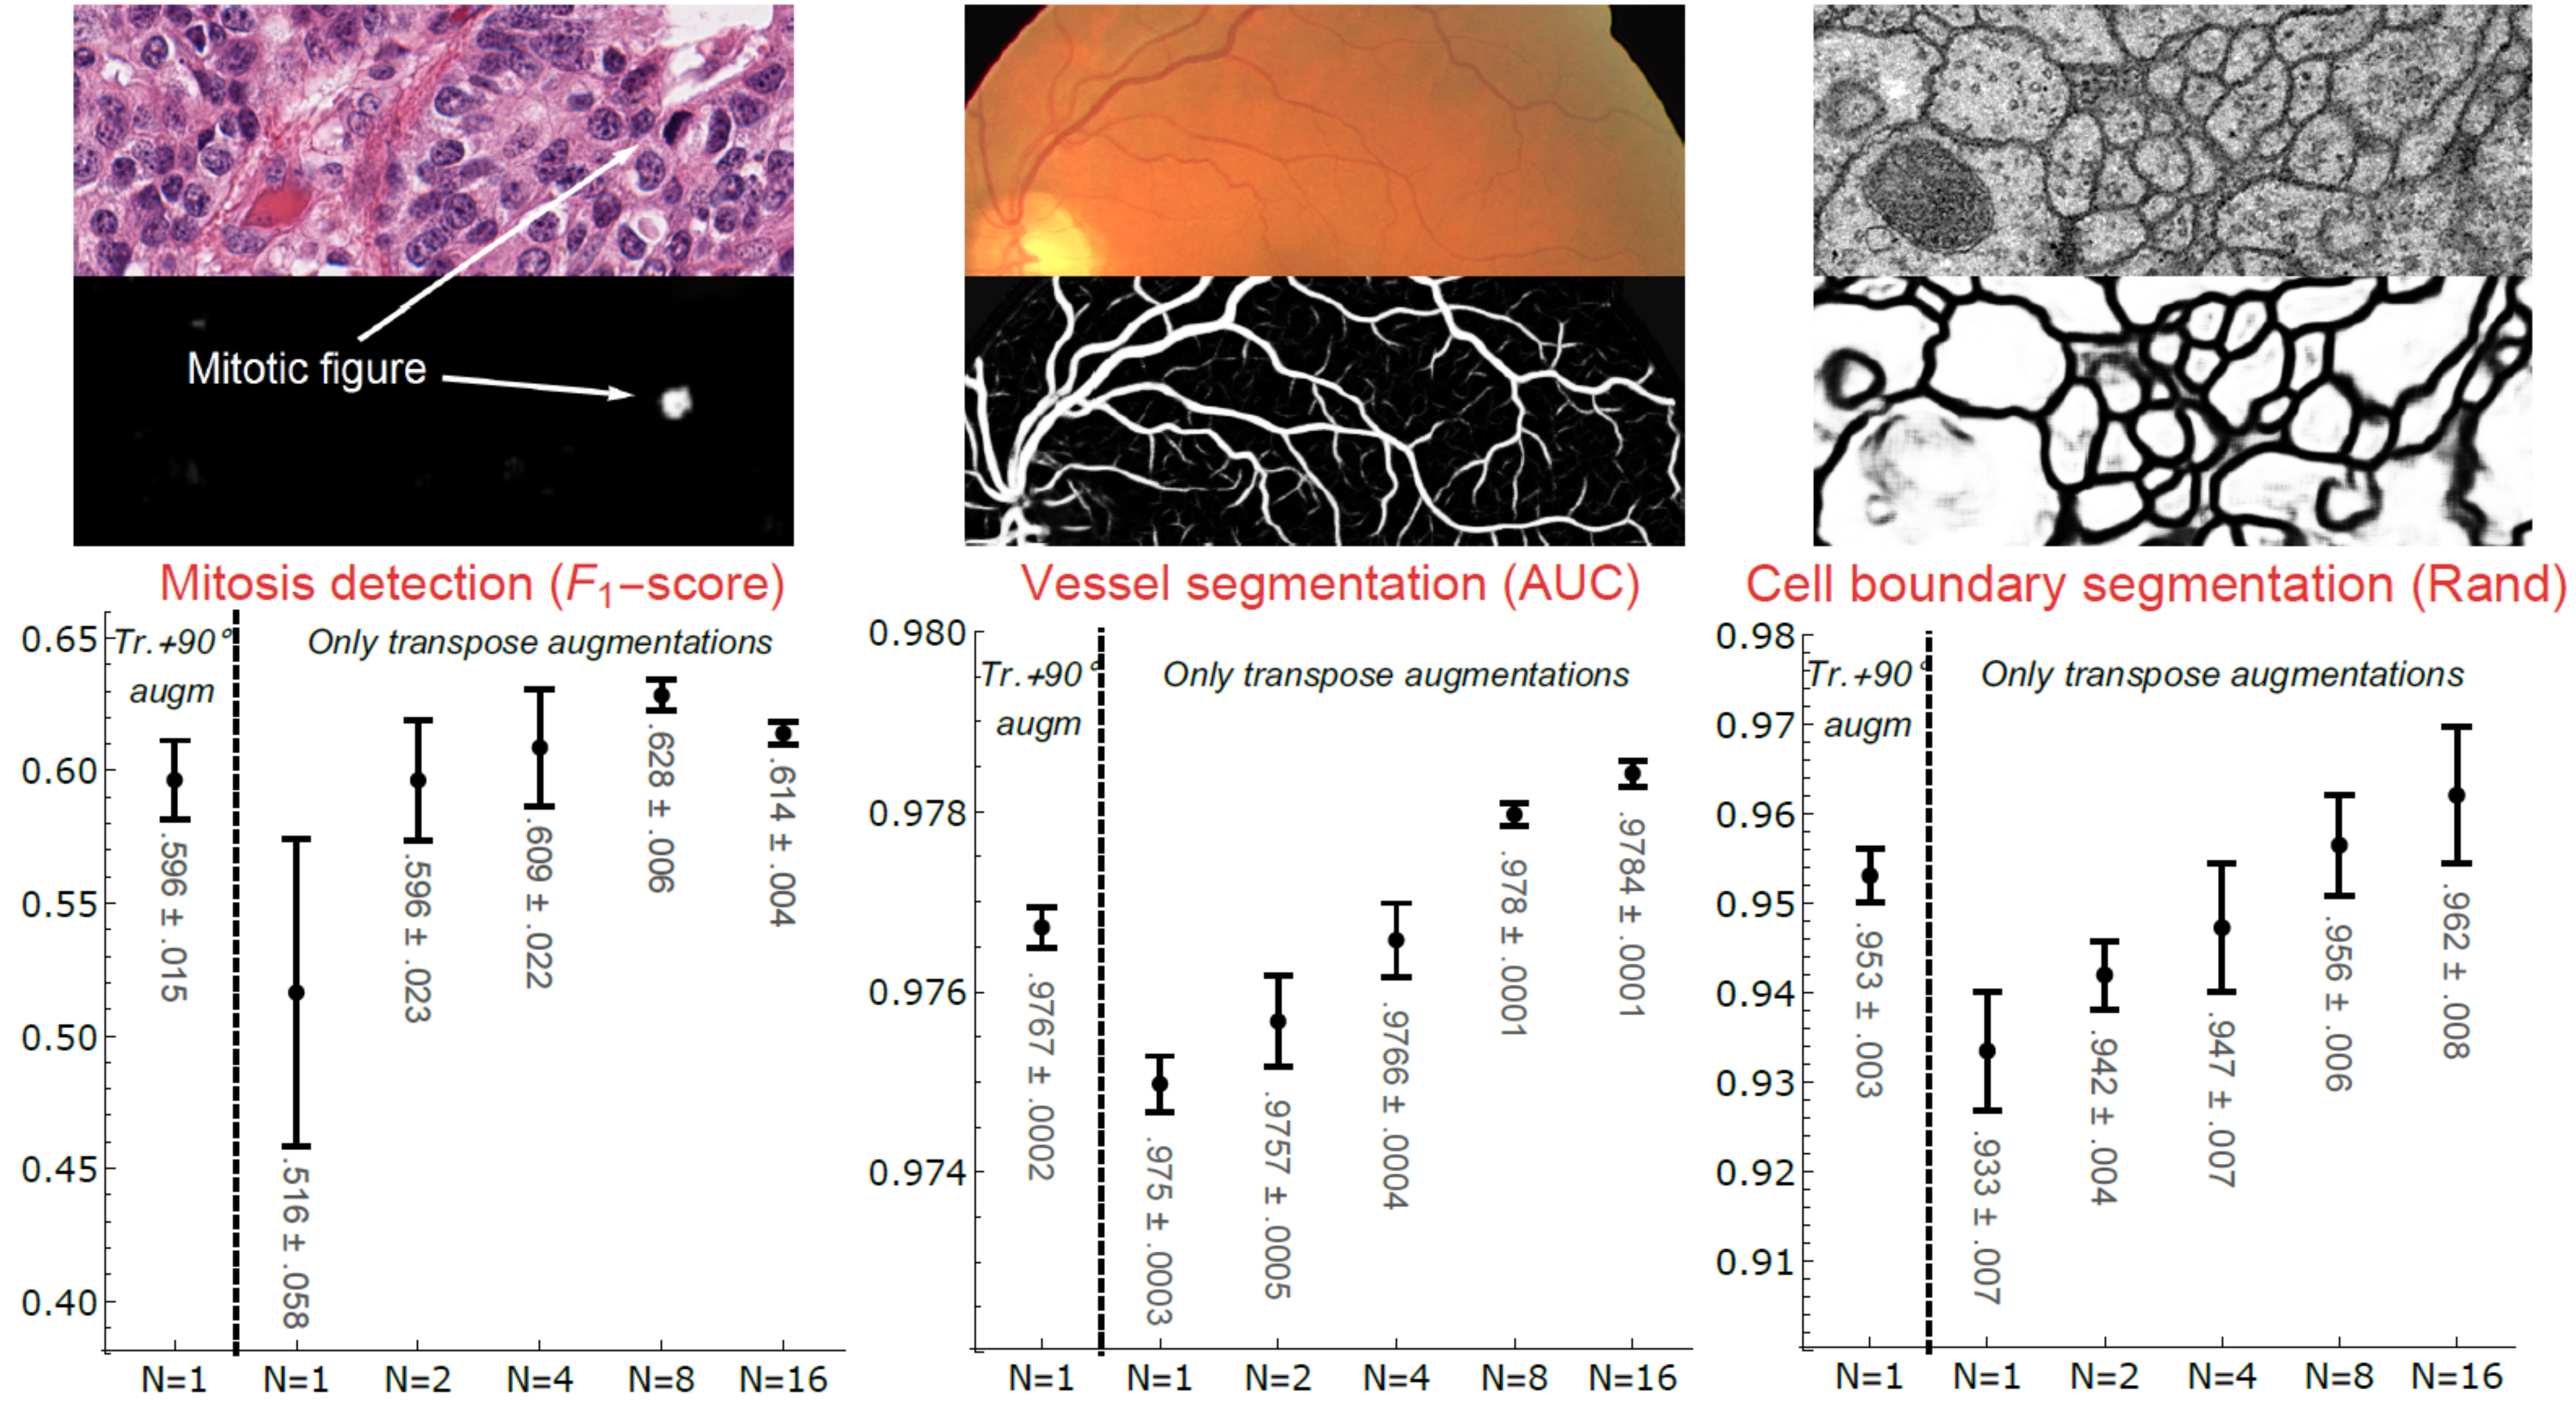

From rotation to scale equivariant CNNs

Bekkers ICLR 2020

Translation + *scale equivariant* G-CNNs

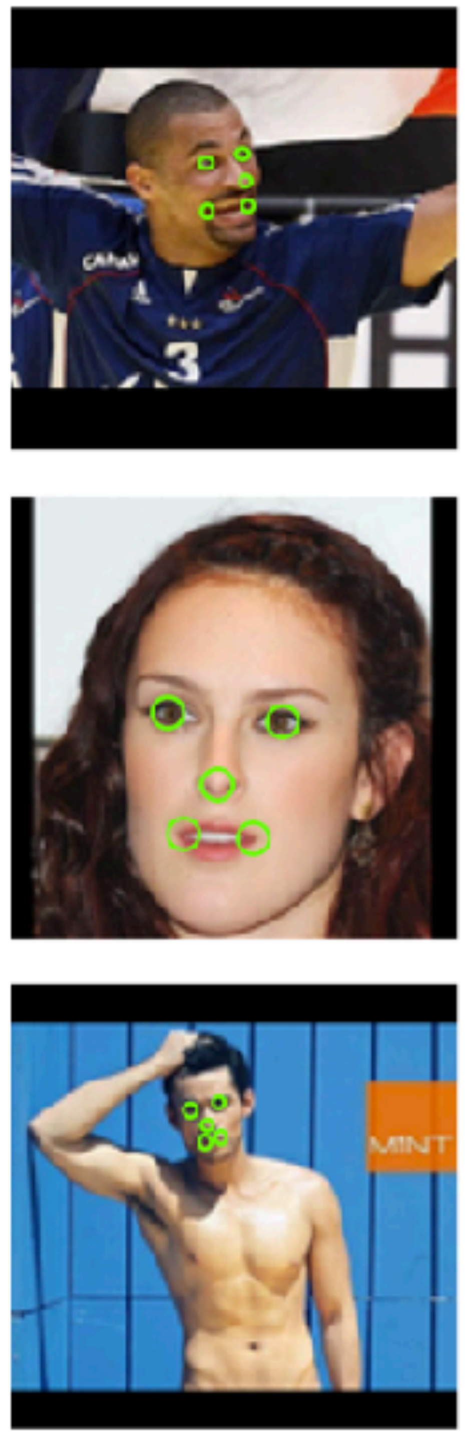

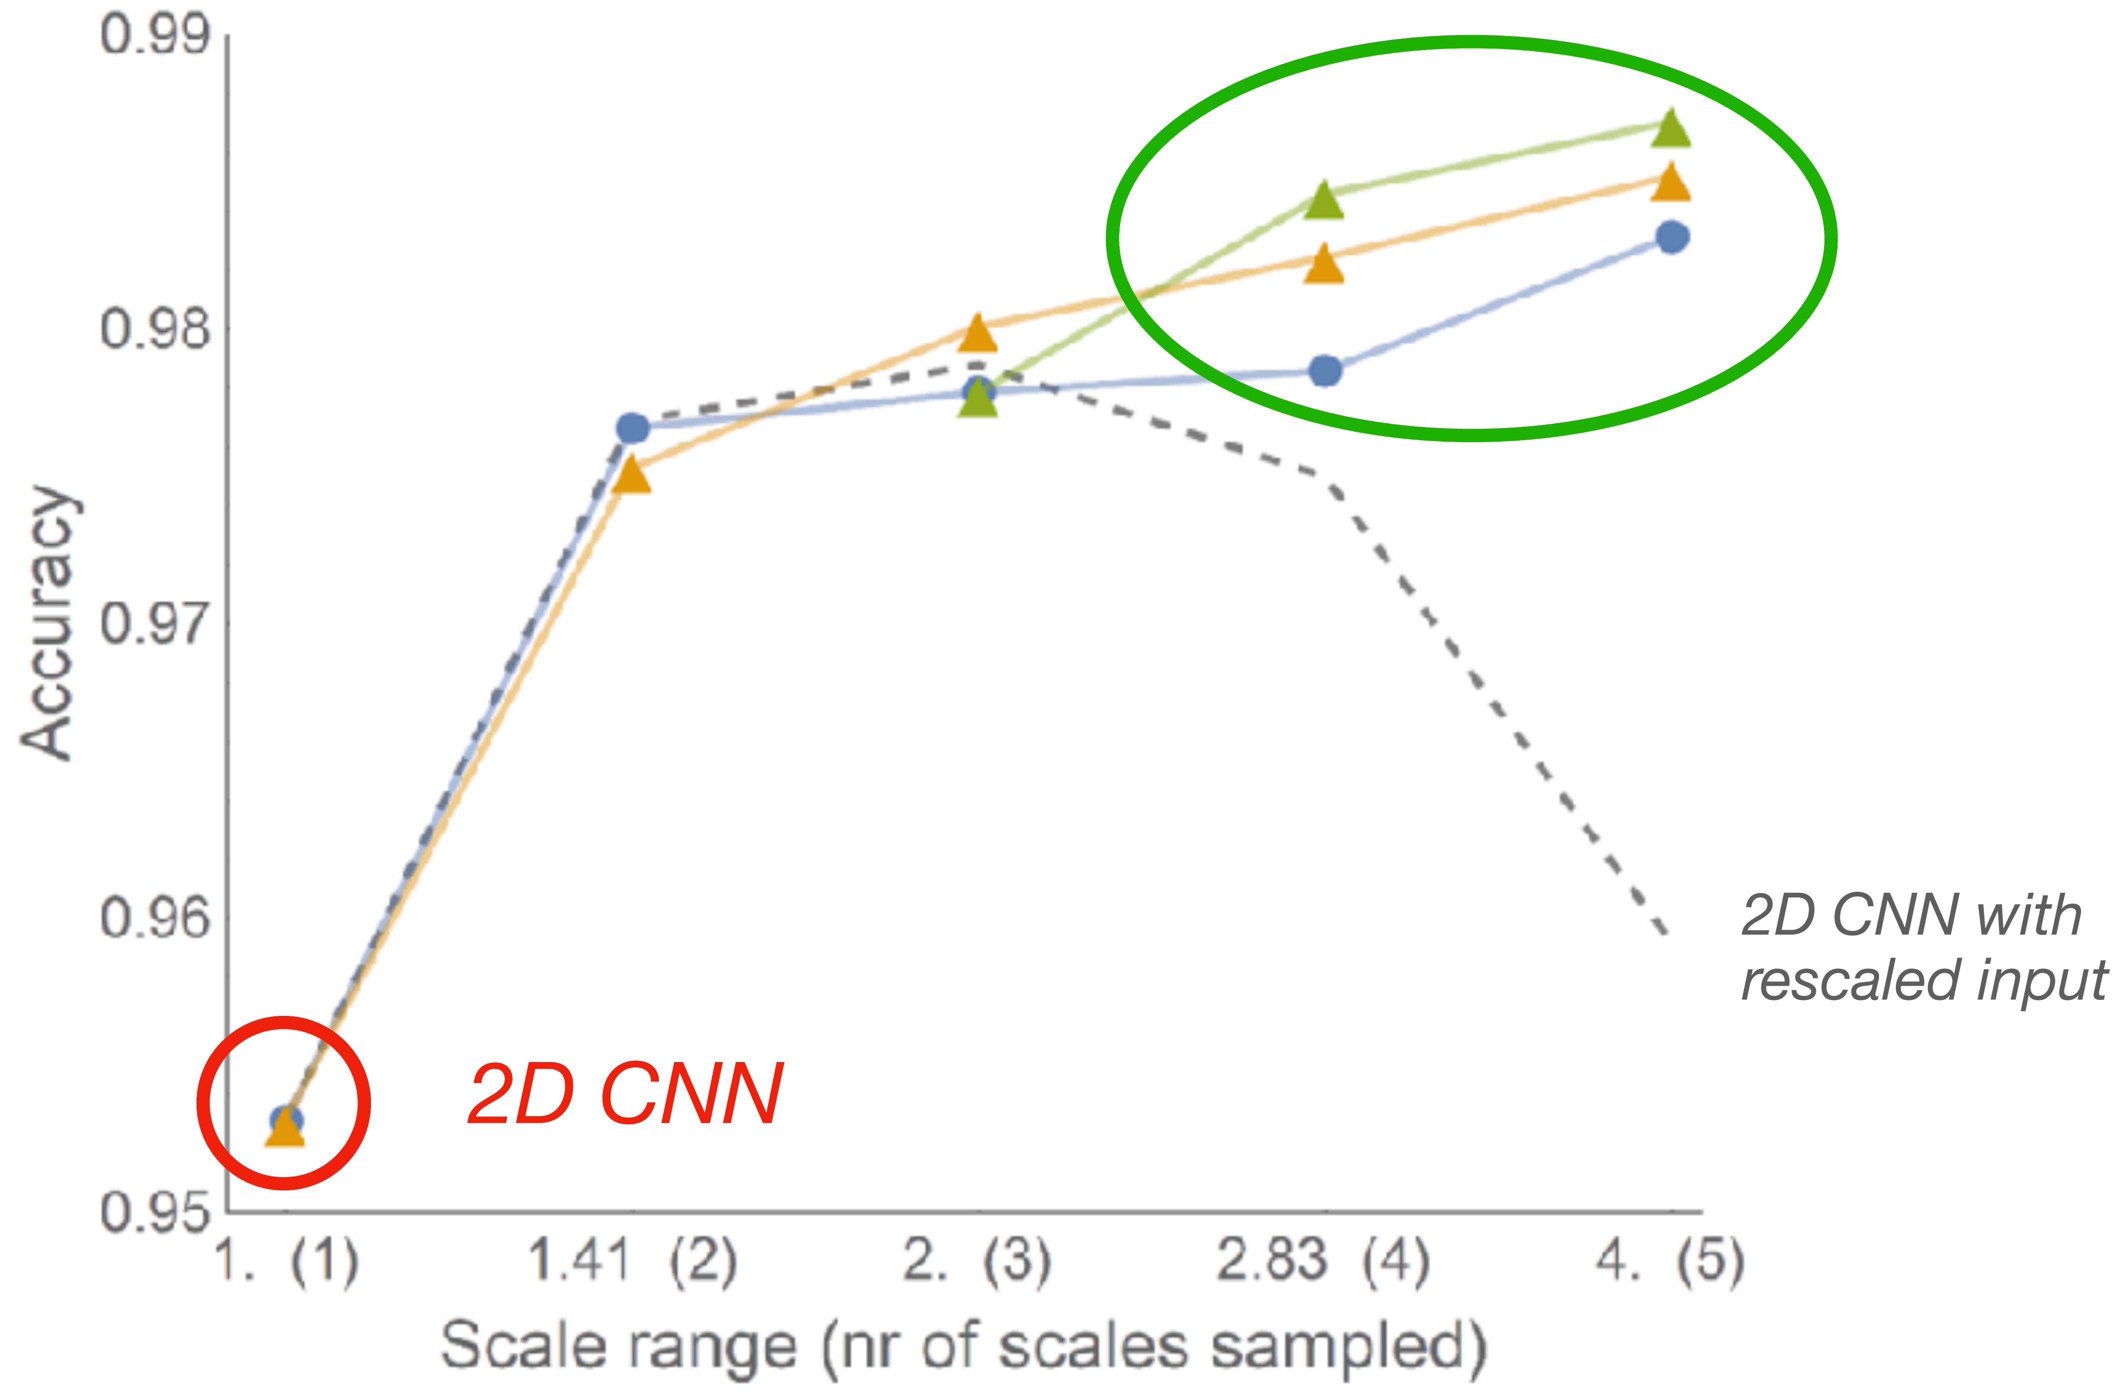

From rotation to scale equivariant CNNs

Romero, Bekkers, Tomczak, Hoogeboom
Wavelet Networks: Scale Equivariant Learning From Raw Waveforms - arXiv:2006.05259

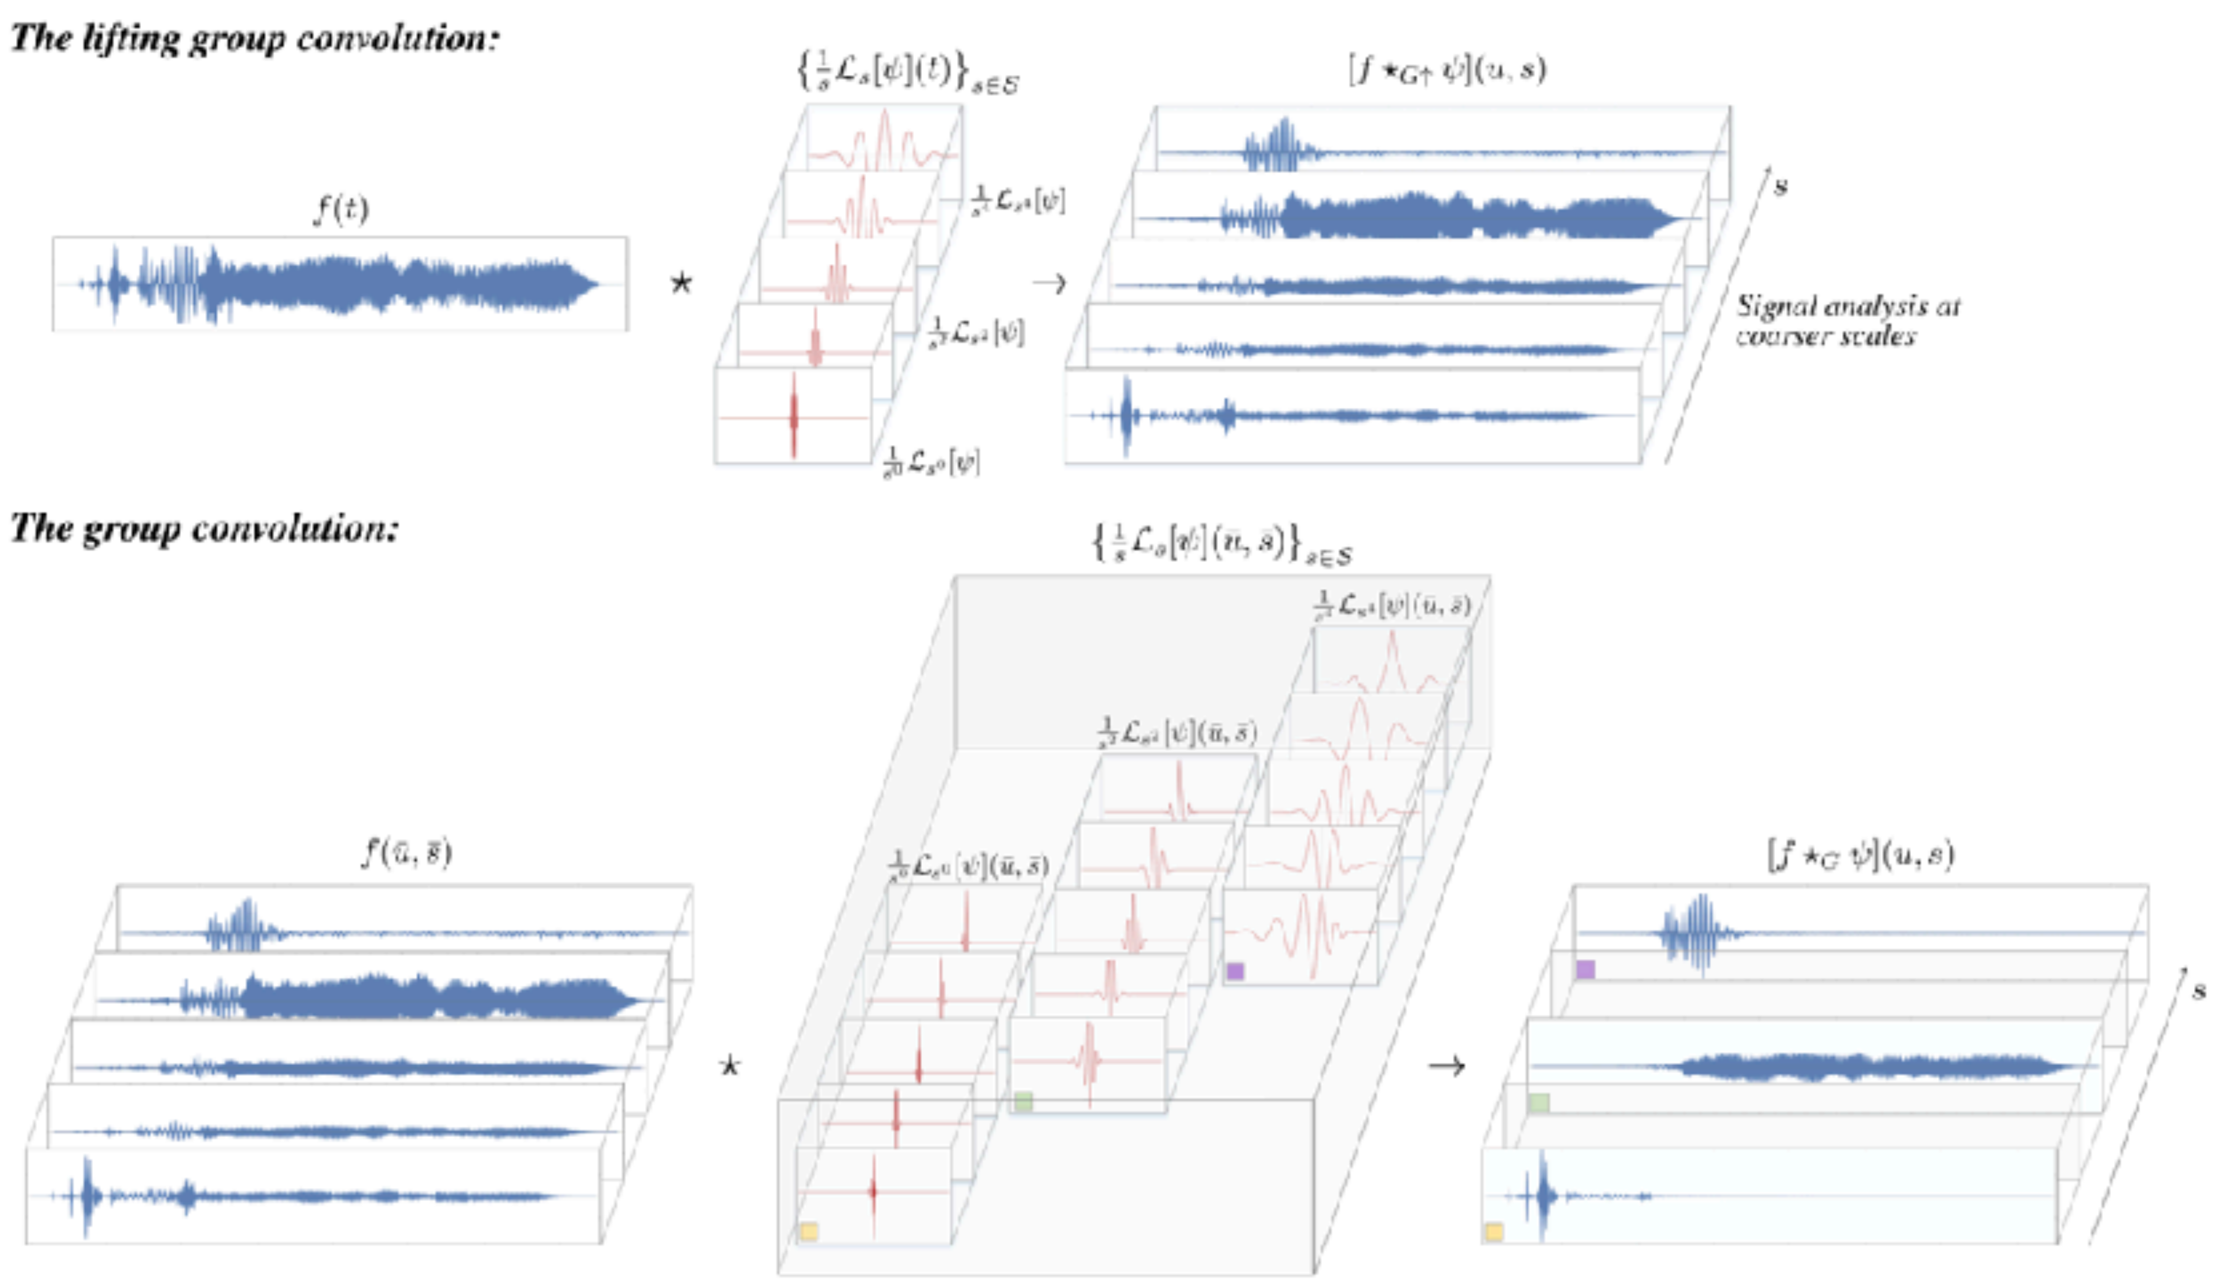

G-CNNs rule!

- The right inductive bias: guaranteed equivariance
  (no loss of information)
- Performance gains that can't be obtained by data-augmentation alone
  (both local and global equivariance/invariance)
- Increased sample efficiency
  (increased weight sharing, no geometric augmentation necessary)

G-CNNs rule!

- The right inductive bias: guaranteed equivariance
  (no loss of information)
- Performance gains that can’t be obtained by data-augmentation alone
  (both local and global equivariance/invariance)
- Increased sample efficiency
  (increased weight sharing, no geometric augmentation necessary)