Efficient Out-of-Distribution Detection in Digital Pathology

Jasper Linmans, Jeroen van der Laak, Geert Litjens

MIDL
Montréal 2020

Radboudumc

Out-of-Distribution (OOD) Detection

- CNNs fail silently

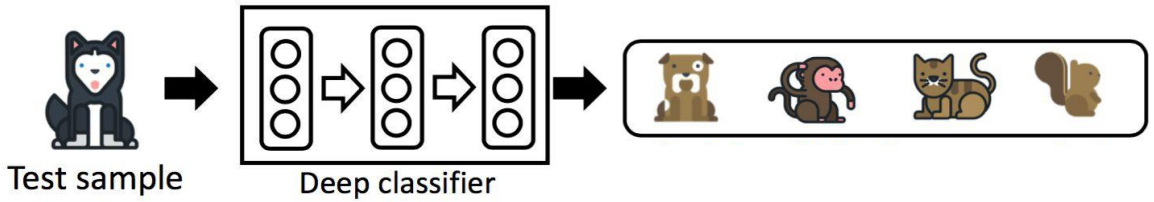

Out-of-Distribution (OOD) Detection

- CNNs fail silently

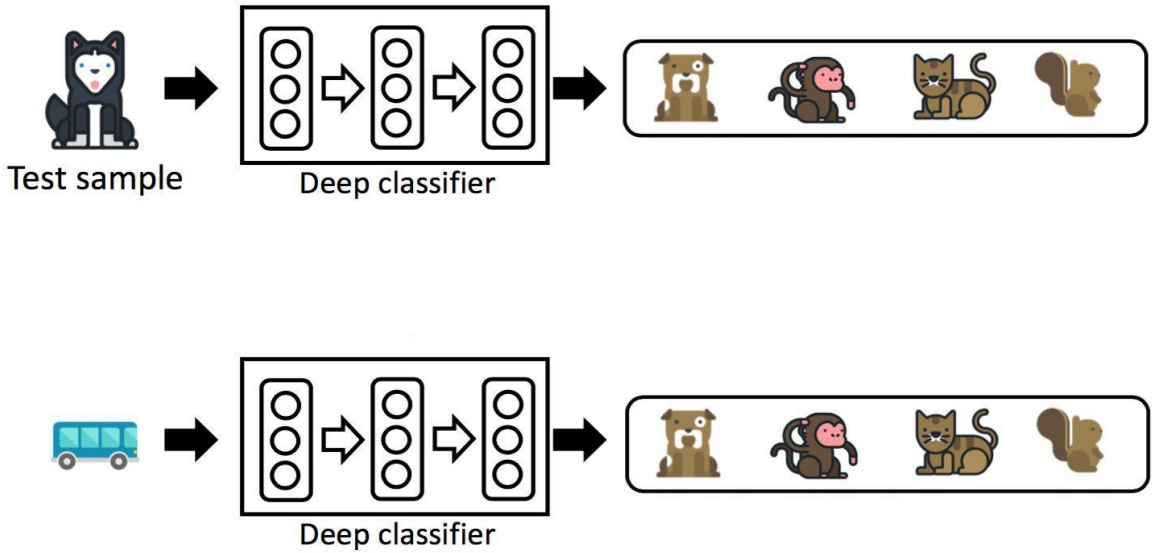

Out-of-Distribution (OOD) Detection

- CNNs fail silently
- Goal: fail loudly on OOD data

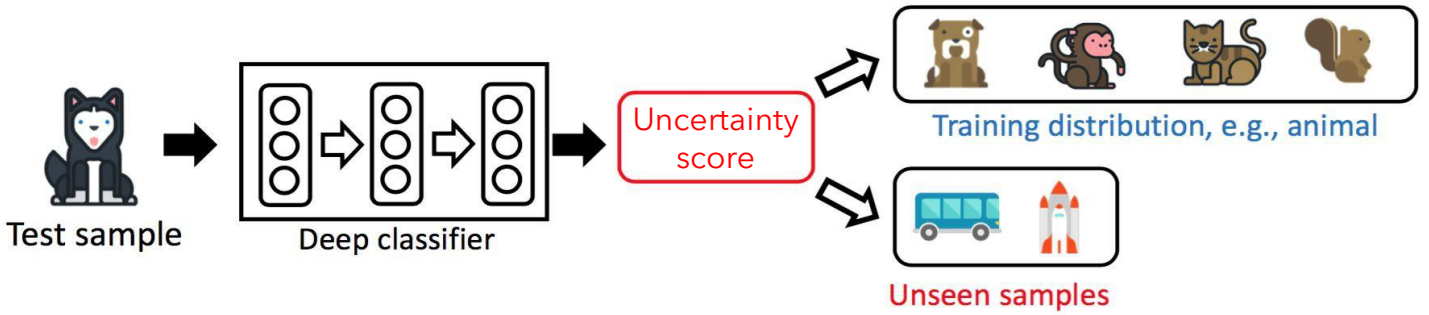

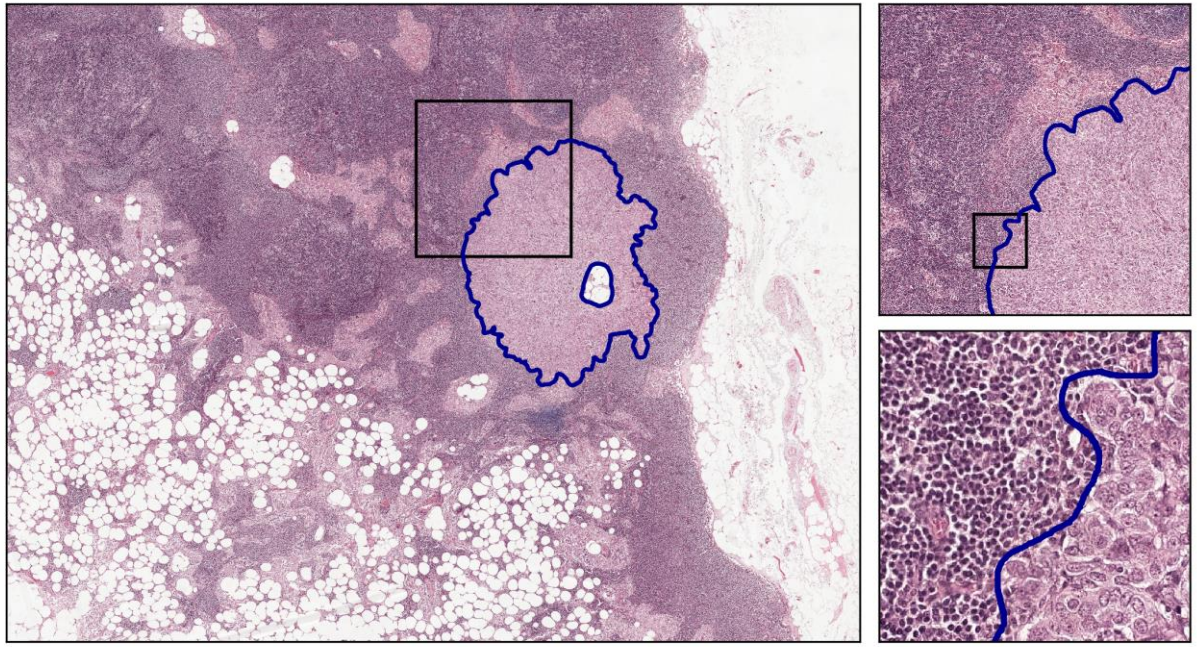

Training data

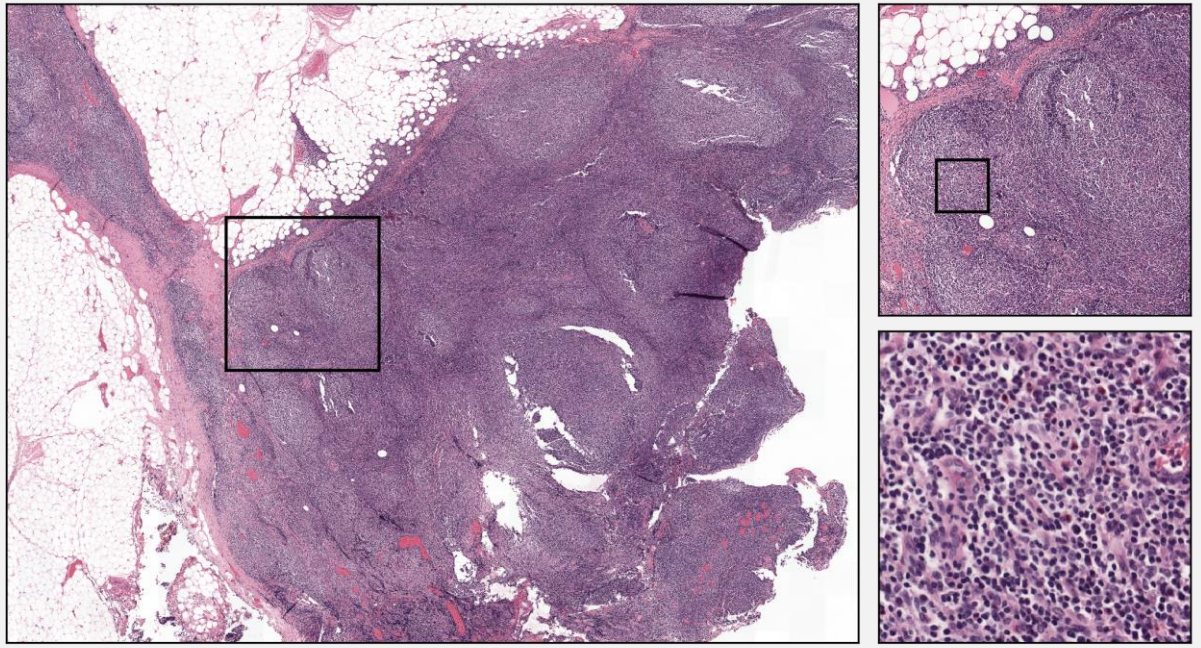

Out-of-Distribution data

Detecting OOD

Most popular approaches measure entropy using:

- Deep Ensembles
- Mc-Dropout

We propose to use **Multi-Head CNNs**

- **Computationally efficient**: require only a single feed forward pass
- **Memory efficient**: entire model is trained at once

Multi-Head CNNs – Distributing Gradients

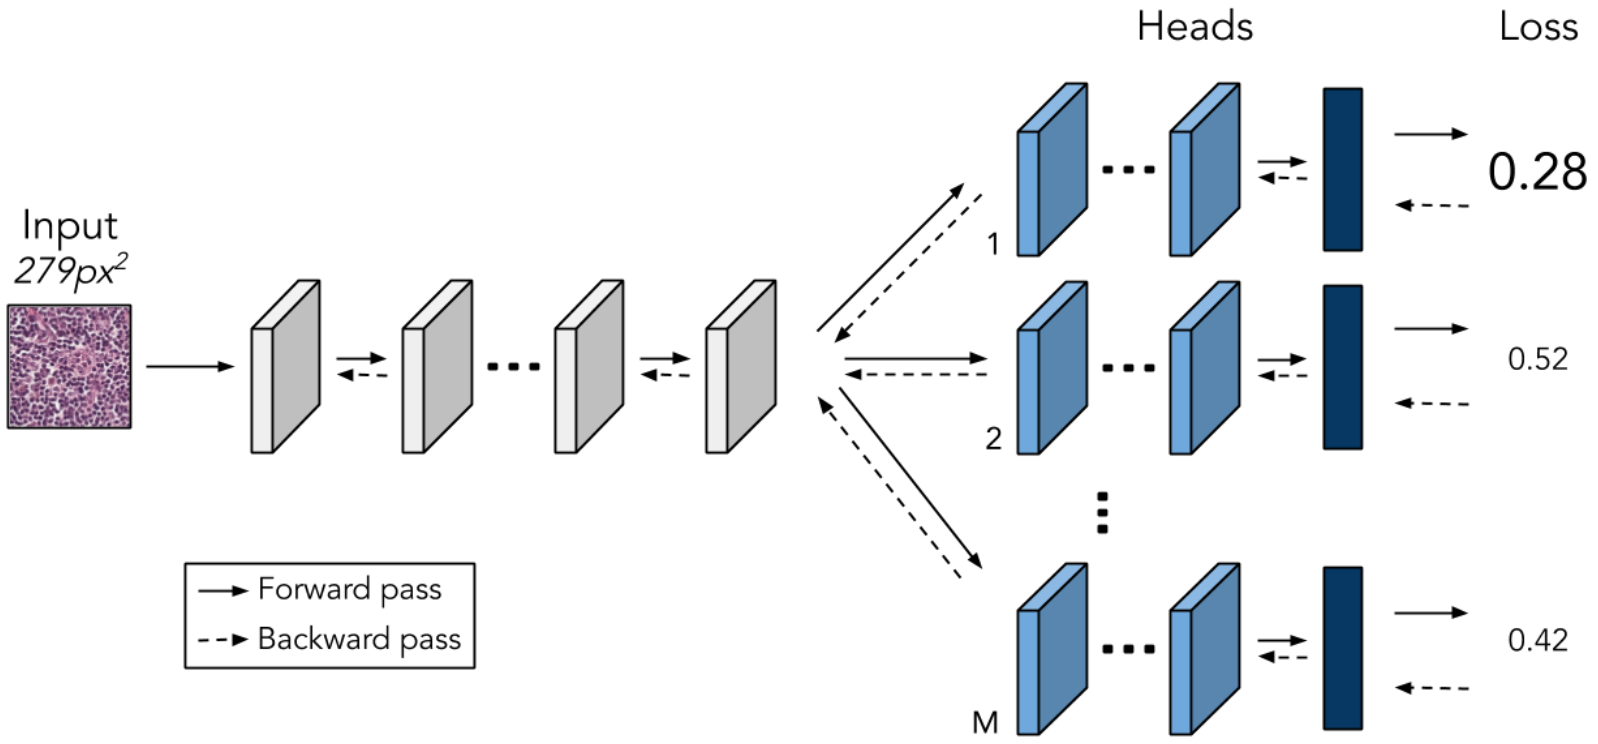

Multi-Head CNNs – Distributing Gradients

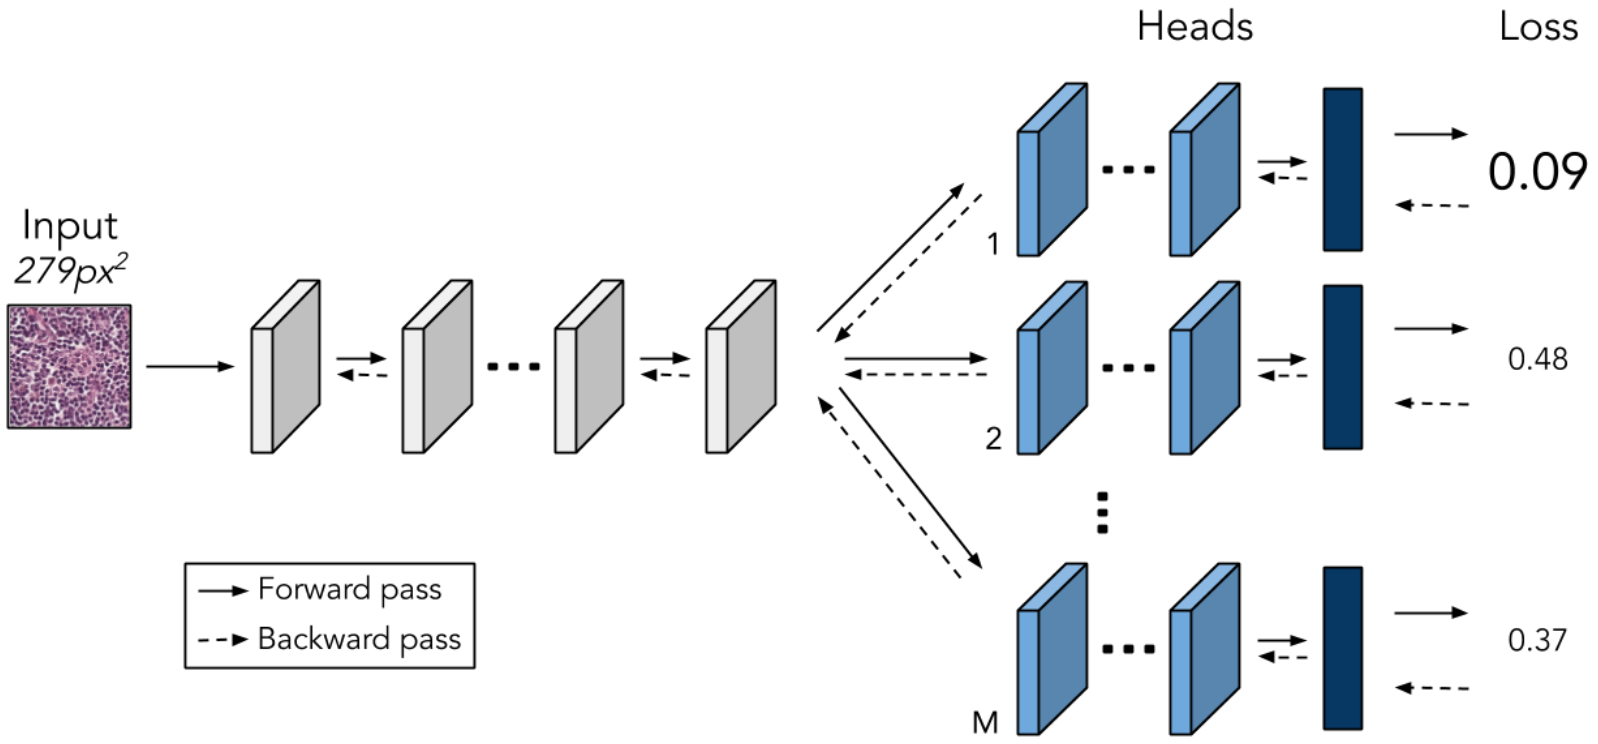

Multi-Head CNNs – Distributing Gradients

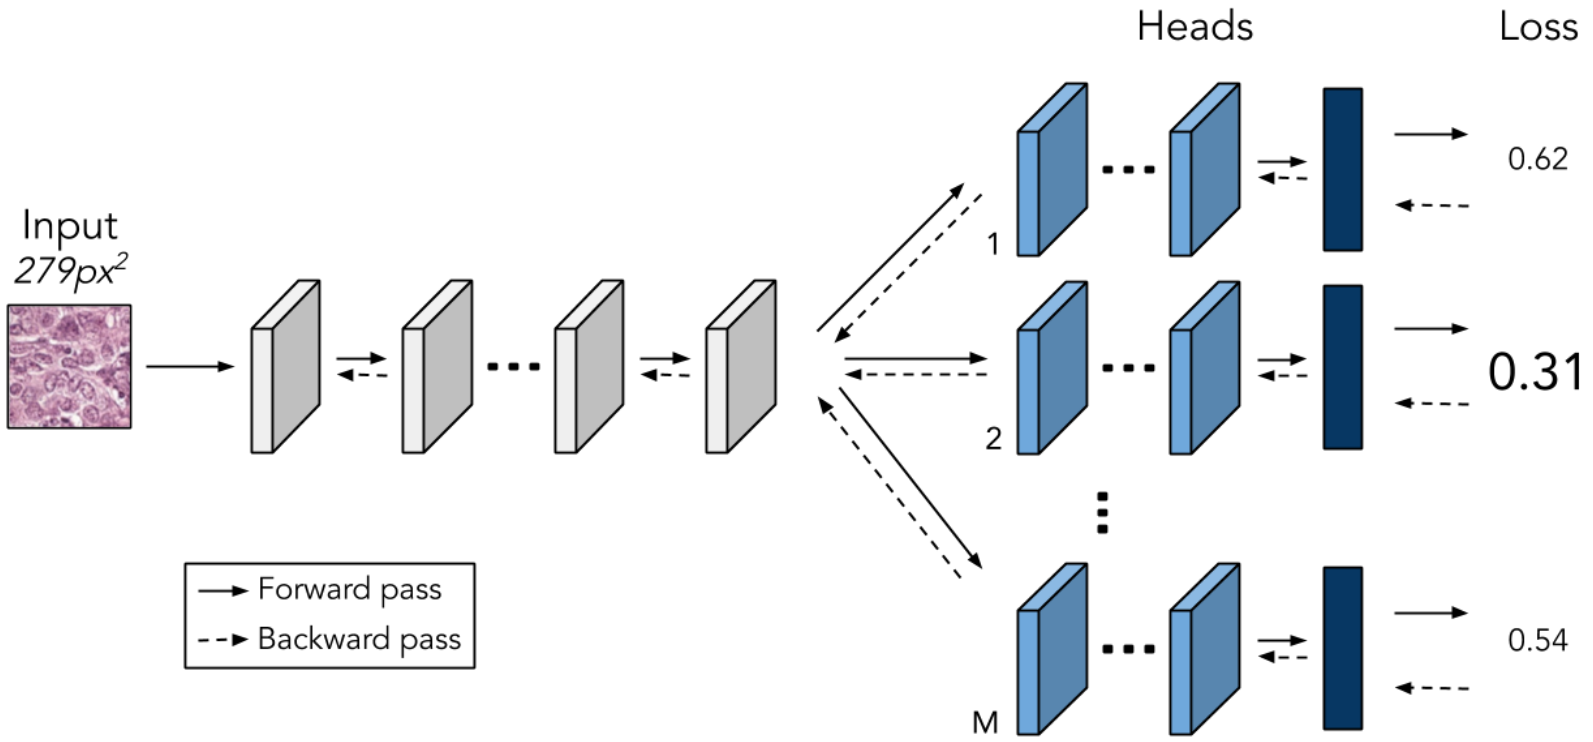

Multi-Head CNNs – Distributing Gradients

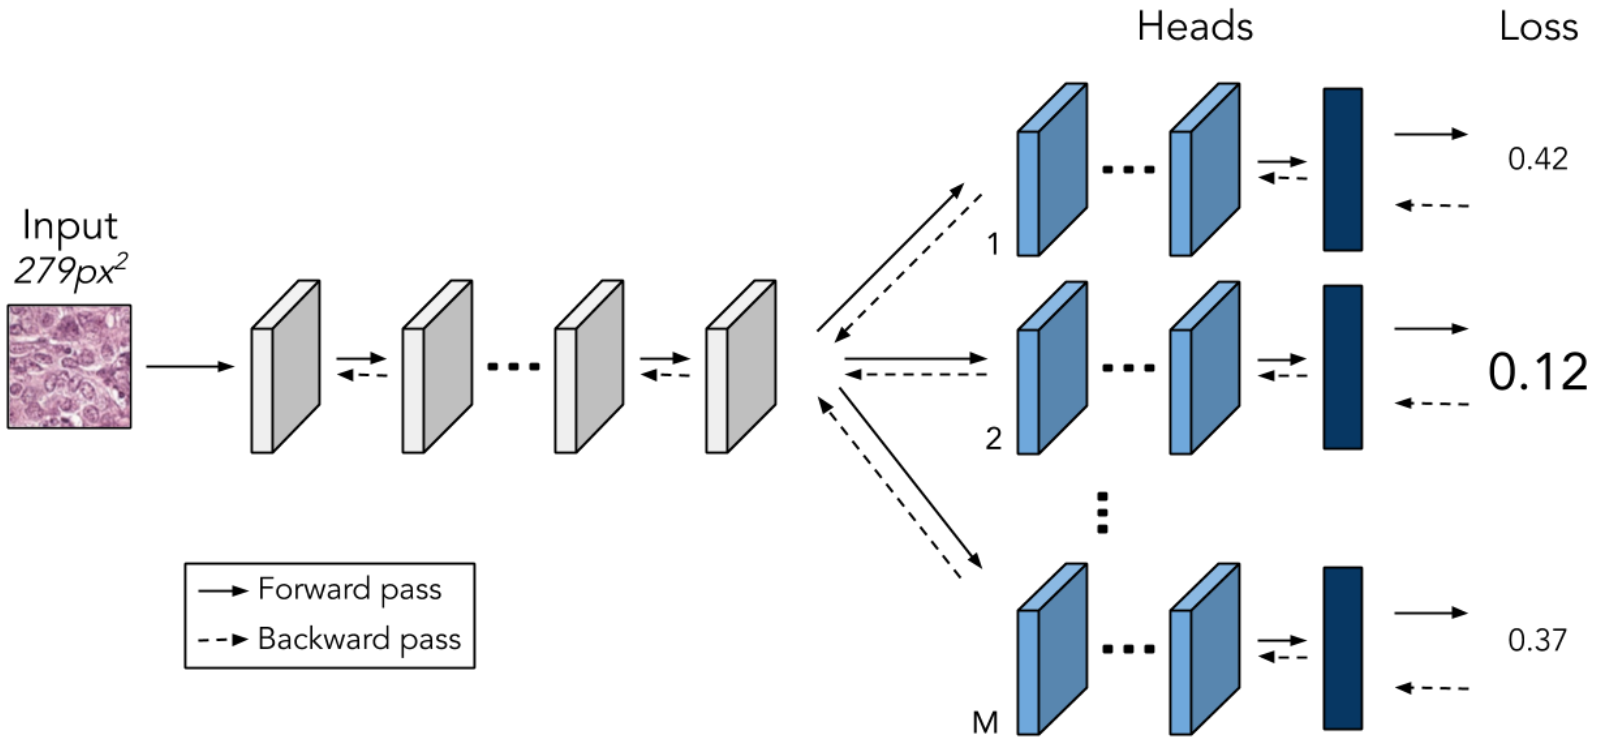

Results

Input & ground-truth

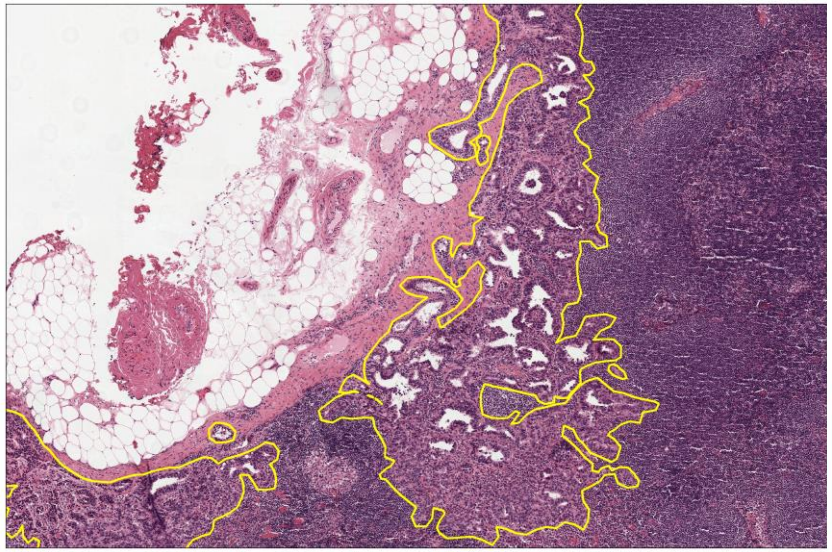

Prediction

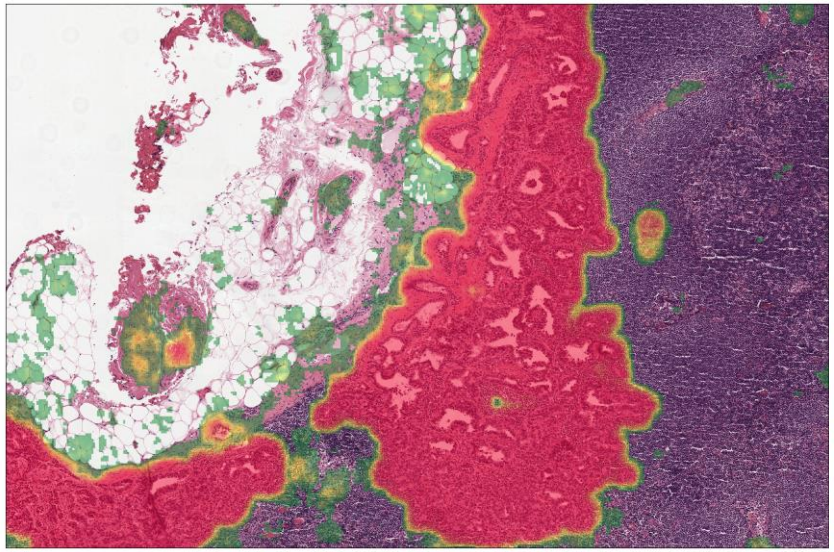

Uncertainty

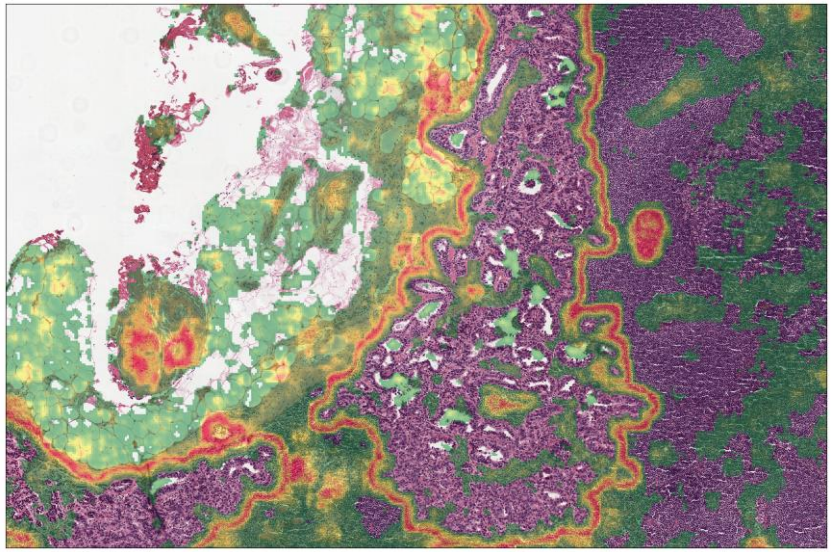

Radboudumc

Results

Input & ground-truth

Prediction

Uncertainty

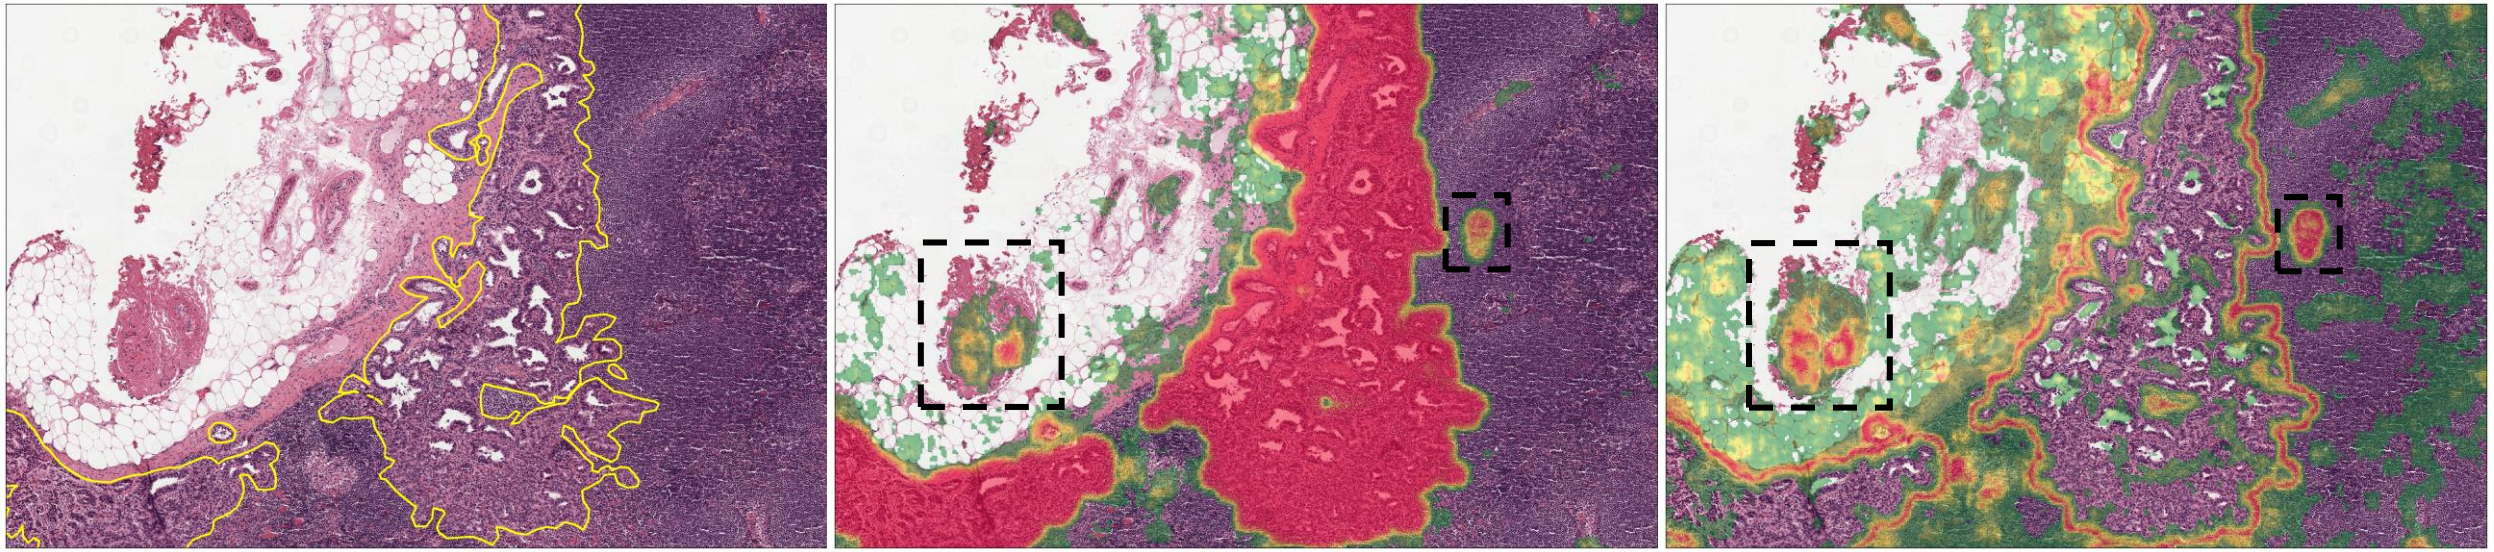

Radboudumc

Results

Model	FPR @ 95TPR ↓	AUROC
Baseline	45.2 (25.1, 65.4)	84.2 (77.5, 91.3)
MC-Dropout	48.3 (26.9, 68.2)	88.3 (81.5, 94.1)
Ensemble (10)	43.4 (24.0, 62.5)	86.8 (79.9, 92.9)
M-heads (10)	28.9 (12.0, 46.2)	91.7 (86.3, 96.5)

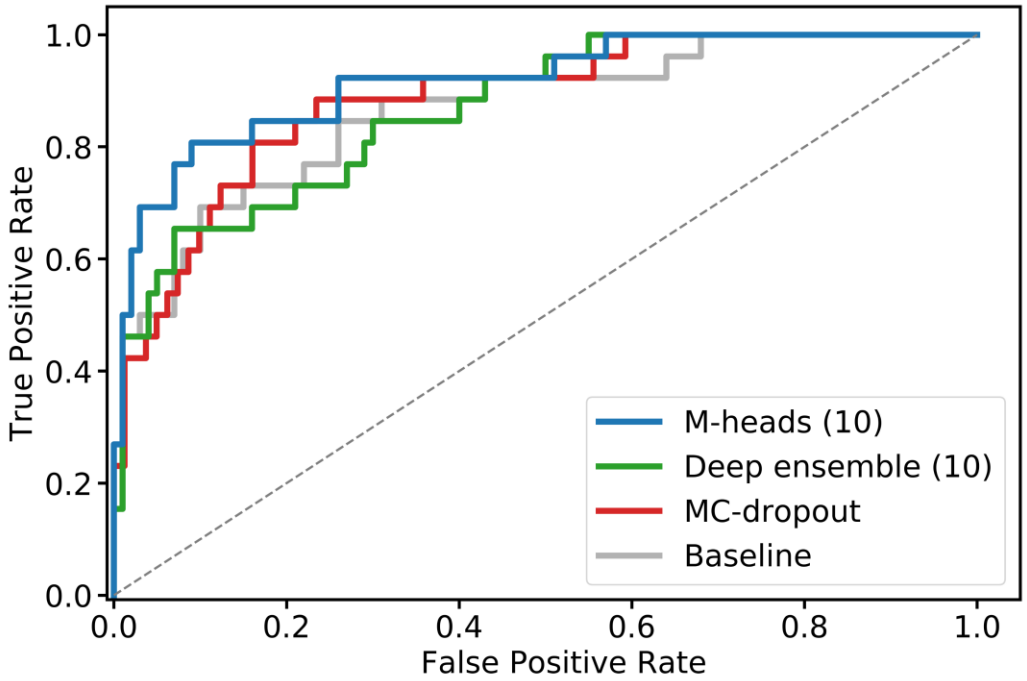

Radboudumc

Specialisation

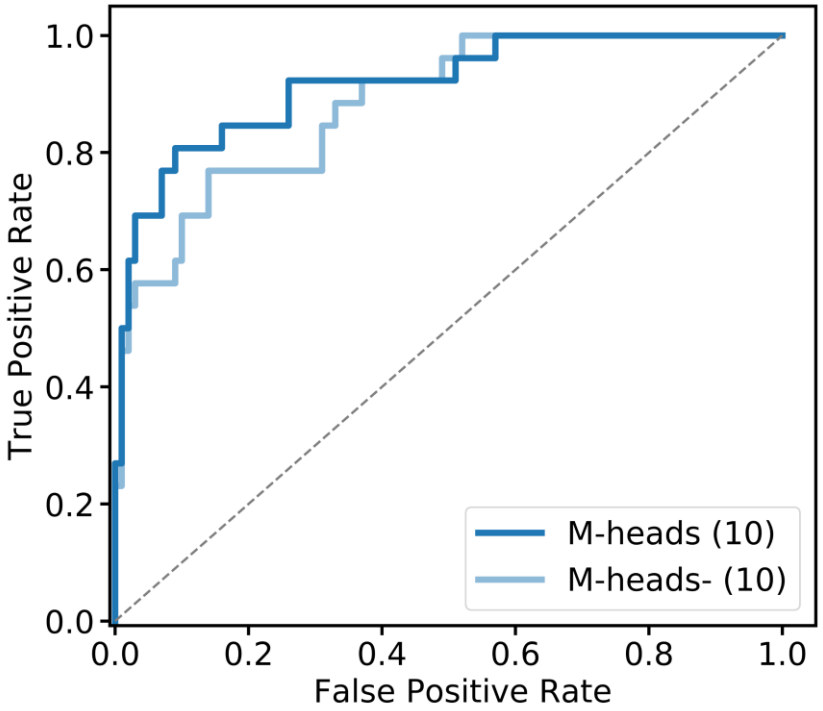

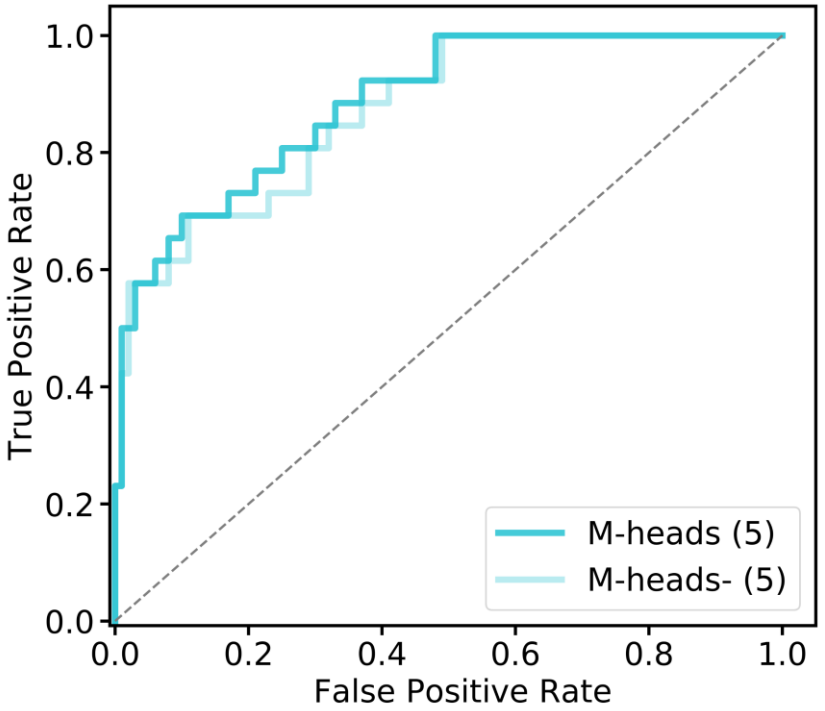

Takeaway Messages

- We can **fail loudly** on OOD data
- M-heads can outperform de current SOTA: deep ensembles
- Head specialisation improves OOD detection

Thanks for Watching

Efficient Out-of-Distribution Detection in Digital Pathology
Using Multi-Head Convolutional Neural Networks

Jasper Linmans, Jeroen van der Laak, Geert Litjens

Code available at: https://github.com/JasperLinmans/m-heads

Radboudumc