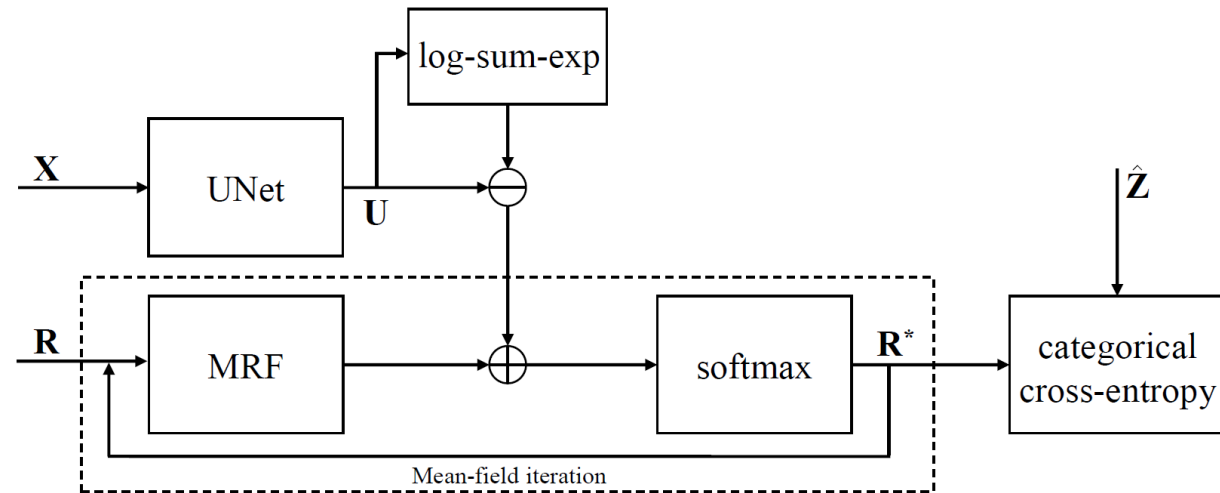


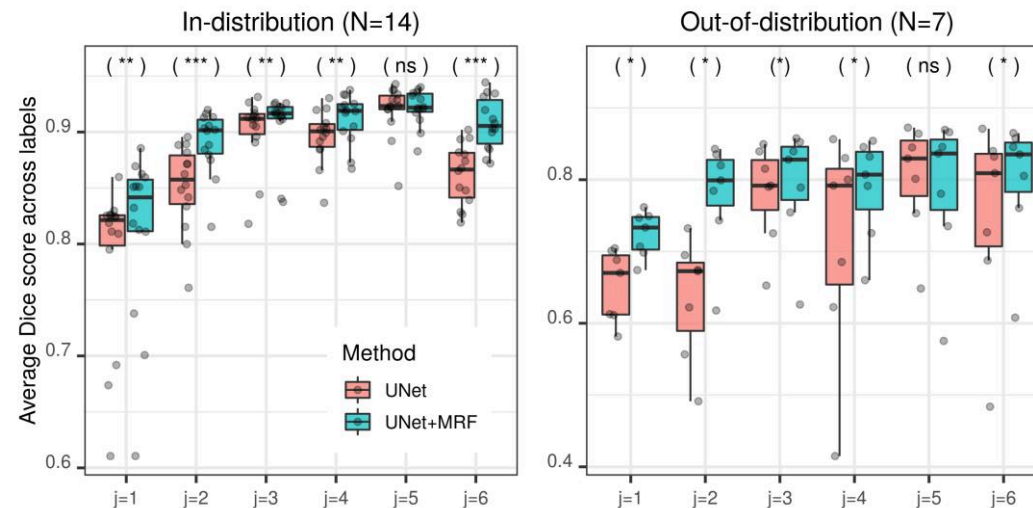
Introducing MRF priors into CNNs

For improving out-of-distribution generalisability and decreasing model size

Methods. An image (\mathbf{X}) is passed forward, through a UNet, whose logit outputs are then fused with the current estimate of the responsibility map (\mathbf{R}). The responsibility map is updated in an iterative fashion (\mathbf{R}^*). For training, the categorical cross entropy between the reference segmentation (\mathbf{Z}) and the responsibilities is computed.



Experiments. Average Dice scores for segmenting the in- and out-of-distribution test images, using both the UNet and the MRF-UNet. For both networks, we vary the number of convolutional filters as: $(2^j, 2^{(j+1)}, 2^{(j+2)}, 2^{(j+3)}, 2^{(j+4)})$. The asterisks above the boxes indicate statistical significance of paired Wilcoxon tests.



Results. The MRF-UNet outperforms the baseline UNet for almost all parameter configurations. Furthermore, the results imply that the MRF-UNet model allows for using fewer UNet parameters, with retained Dice scores.

Paper 107

An MRF-UNet Product of Experts
for Image Segmentation

Mikael Brudfors, Yaël Balbastre, John Ashburner
Geraint Rees, Parashkev Nachev, Sebastien Ourselin
M. Jorge Cardoso

AMIGO
Artificial Medical Intelligence Group

KING'S
College
LONDON

MIDL
Lübeck 2021