
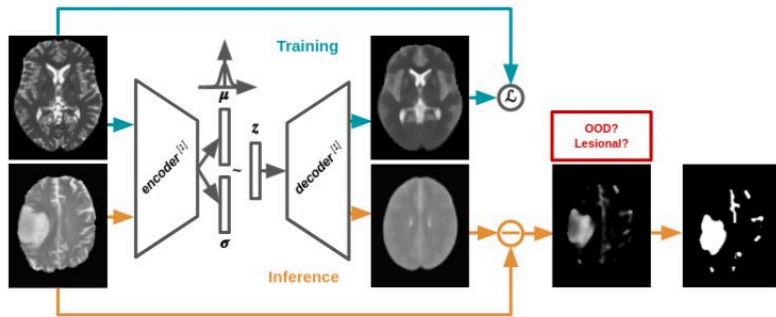



The Out-of-Distribution Blindspot of Unsupervised Lesion Detection

Matthäus Heer*, Janis Postels**, Xiaoran Chen**, Ender Konukoglu**, Shadi Albarqouni***



Training Distribution

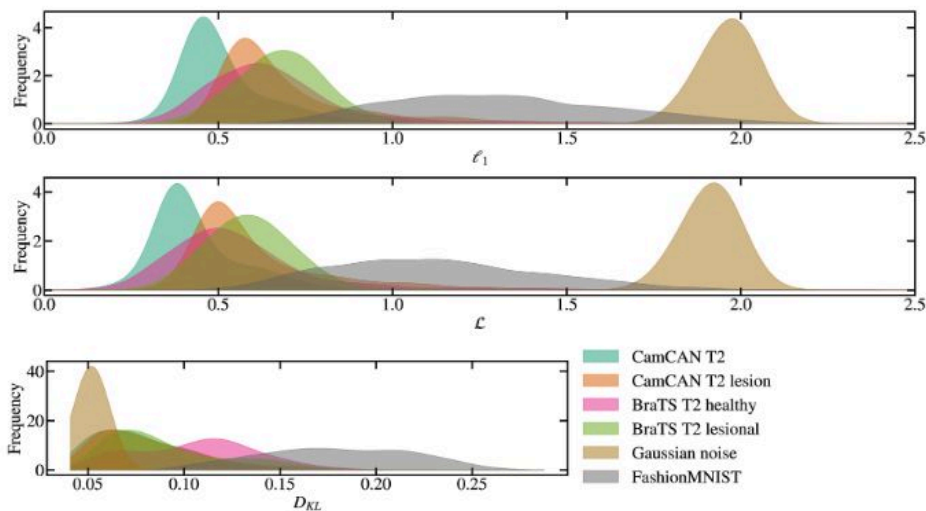


Healthy & In-distribution

Lesional & OOD

Healthy & OOD

Possible Inference Samples



$$\log p(\mathbf{x}) \geq \mathcal{L} = \underbrace{E_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})]}_{\text{Reconstruction Error}} - \underbrace{D_{KL}[q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]}_{\text{Prior Loss}}$$