MIRAGE: Multi-scale Lesion-Informed Representation with Auxiliary Guidance for MRI Contrast Enhancement

Andrea Borghesi, Xin Wang, Jonas Teuwen, George Yiasemis

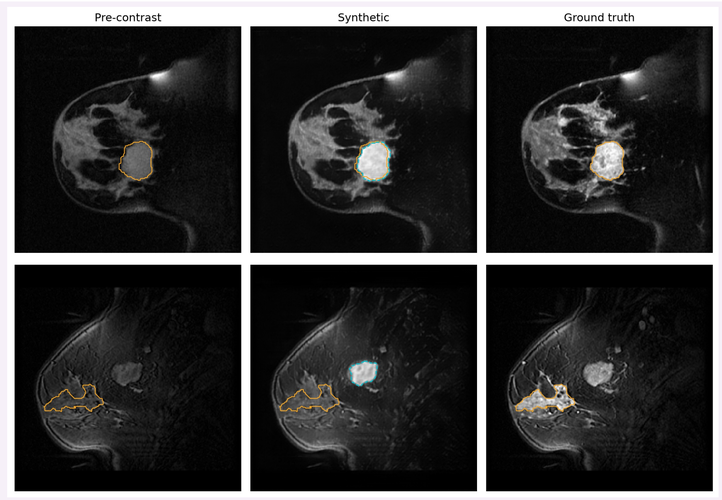
Netherlands Cancer Institute, Amsterdam

Motivation

To develop a Deep Learning method to synthesize Contrast-Enhacement in Breast MRI starting from a non-contrast MRI slice.

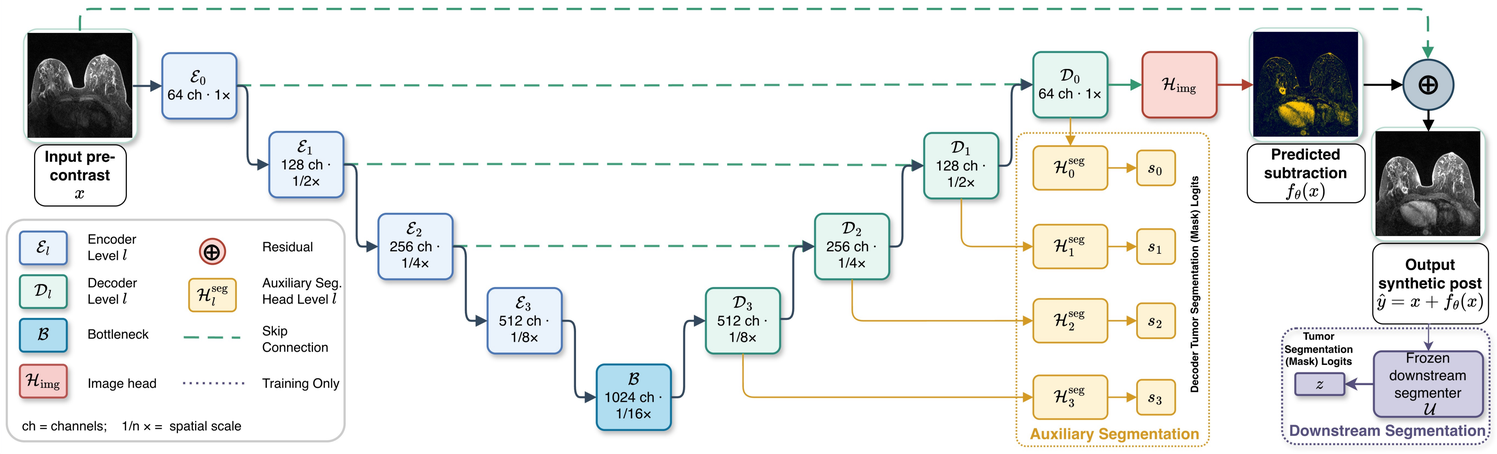
Introduction

- **DCE-MRI** is the most sensitive breast imaging exam: lesions are revealed by their uptake of a gadolinium contrast agent
- Contrast has a cost: IV access, longer exams, exclusion of some patients, gadolinium in the environment
- **Virtual contrast enhancement** predicts the post-contrast image from the pre-contrast scan alone, without injection



Method

We propose MIRAGE, a residual 2D U-Net trained with image-, lesion- and task-aware losses.
Tumour masks supervise training only; at inference the network sees the pre-contrast slice alone.
In this context the residual predicted is the contrast, preventing any anatomy hallucination.



$$\mathcal{L} = \mathcal{L}_1 + 0.3\,\mathcal{L}_{\text{LPIPS}} + 0.1\,\mathcal{L}_{\text{seg}} + 0.2\,\mathcal{L}_{\text{dseg}} + 0.05\,\mathcal{L}_{\text{under}}$$

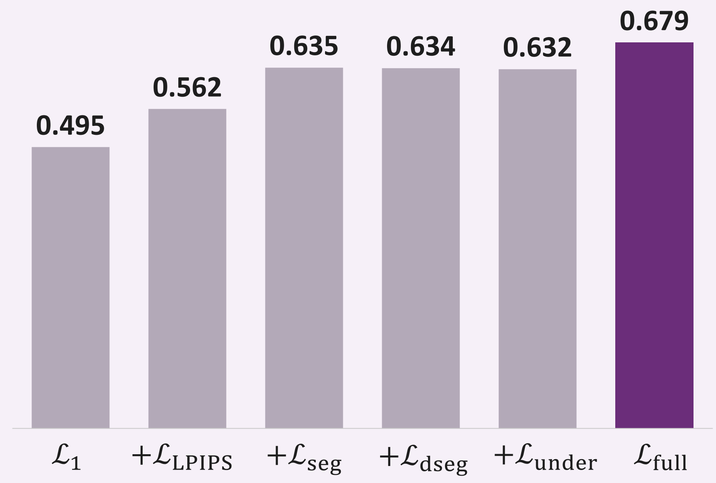
$\mathcal{L}_1$	**masked pixel reconstruction**
$\mathcal{L}_{\text{LPIPS}}$	**perceptual appearance**
$\mathcal{L}_{\text{under}}$	**penalises missed lesion enhancement**
$\mathcal{L}_{\text{seg}}$	**tumour masks at 4 decoder scales**
$\mathcal{L}_{\text{dseg}}$	**frozen nnU-Net segments** $\hat{y}$

Results

301 held-out validation cases · **1,506** exams in MAMA-SYNTH · **4** cohorts, 1.5 T and 3 T · **8** metrics in four families

Model	MSE ↓	LPIPS ↓	FRD ↓	SSIMtum ↑	AUROC-C ↑	AUROC-R ↑	Dice ↑	HD95 ↓	Avg rank ↓
MIRAGE (ours)	**0.607**	0.095	**6.56**	**0.540**	0.711	**0.718**	**0.679**	**34.8**	**1.50**
LBM	0.612	**0.079**	8.12	0.529	0.792	0.657	0.493	126.6	2.50
pix2pix	0.628	0.169	11.75	0.497	0.755	0.677	0.509	106.2	2.88
DDPM	0.795	0.122	9.07	0.389	**0.816**	0.632	0.495	122.4	3.13

Loss Ablation



Tumour Dice, leave-one-in from L1 baseline

Discussion

Task-aware synthesis:

a frozen downstream model steers outputs to clinically relevant content

Realism ≠ utility:

sharper, more contrast-like images localise tumours worse

Losses overlap:

any single lesion-aware term recovers most of the Dice, but only the full objective gets to 0.68 / HD95 35

g.yiasemis@nki.nl

aiforoncology.nl

Code & weights

Pre-print