This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Supplementary Material

Table 3: **Comparing OOD scores:** Full version of Table 1 with more OOD scores and their tail-specific version. $P(tail)$ achieves the best performance.

Score	FPR	AUC
P(tail)	$\mathbf{25.03}^{\pm 2.16}$	$\mathbf{95.01}^{\pm 0.48}$
MSP [12] head	$30.98^{\pm 2.36}$	$93.85^{\pm 0.58}$
MSP [12] tail	$25.33^{\pm 1.91}$	$94.98^{\pm 0.47}$
Energy [21] head	$39.17^{\pm 5.30}$	$91.79^{\pm 0.51}$
Energy [21] tail	$89.44^{\pm 6.42}$	$63.90^{\pm 2.54}$
MaxLogit [11] head	$39.02^{\pm 5.27}$	$91.79^{\pm 0.50}$
MaxLogit [11] tail	$89.37^{\pm 6.47}$	$63.77^{\pm 2.49}$
Maha dist. [19] to head	$53.41^{\pm 3.17}$	$82.58^{\pm 1.52}$
Maha dist. [19] to tail	$98.46^{\pm 0.63}$	$23.42^{\pm 1.55}$
Maha dist. [19] to tail - to head	$34.72^{\pm 5.11}$	$89.23^{\pm 1.28}$
relative Maha distance [26]	$32.45^{\pm 4.06}$	$91.88^{\pm 0.75}$

Table 4: **LoRA training details.**

rank	4
learning rate	1e-04
train steps	15000
lr scheduler	cosine
resolution	512
augmentation	center-crop and random-flip

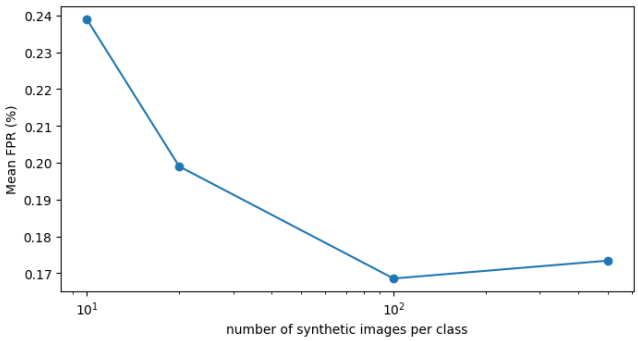


Fig. 3: Ablation on the number of synthetic images per tail class: Adding more than 100 samples leads to a slight increase in FPR, but still outperforms baselines. We hypothesize that the relative importance of the natural tail samples decreases when too many synthetic images are used.

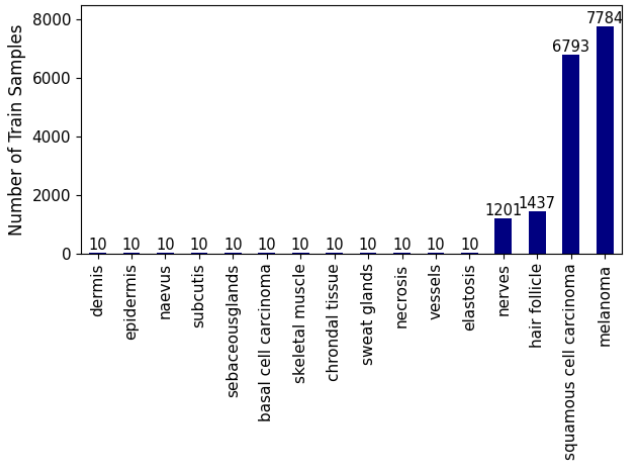


Fig. 4: Class distribution of the train dataset.

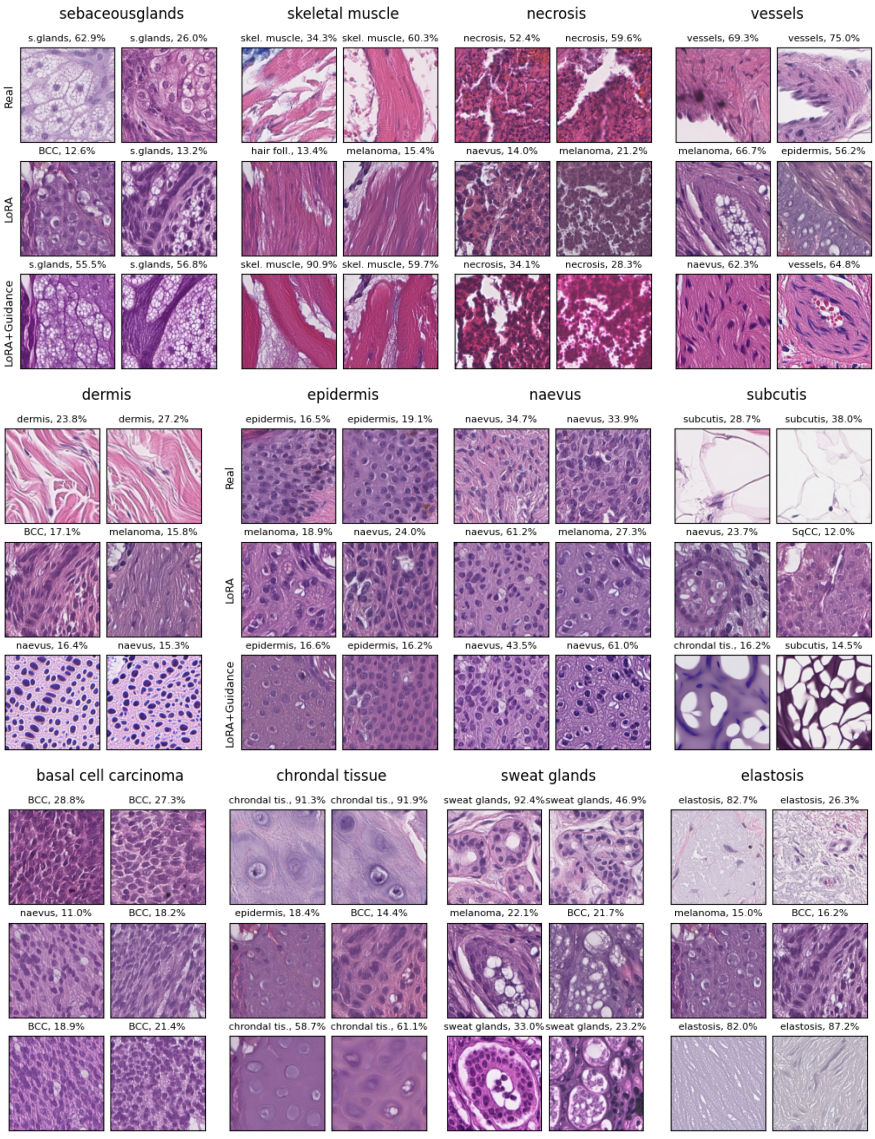


Fig. 5: **Samples from all classes:** We show samples from the train dataset (first row), samples generated with LoRA (second row), and generated with `LoGex` (third row). For each sample, we report the corresponding prediction of a classifier trained on the original dataset from [18] (achieving a tail accuracy of 96.1%). With `LoGex` the synthetic images are more often classified as the desired class.