FISHing in Uncertainty: Synthetic Contrastive Learning for Genetic Aberration Detection - Supplementary Material

Simon Gutwein[1,2,3][0009−0004−8406−0736], Martin Kampel[2][0000−0002−5217−2854], Sabine Taschner-Mandl[1][0000−0002−1439−5301], and Roxane Licandro[3][0000−0001−9066−4473]

[1] St. Anna Children's Cancer Research Institute, Vienna, Austria
[2] TU Wien, Institute of Visual Computing and Human-Centered Technology, Computer Vision Lab, Vienna, Austria
[3] Medical University of Vienna, Department of Biomedical Imaging and Image-guided Therapy, Computational Imaging Research, ELIA Group, Vienna, Austria

1 Supplementary Material

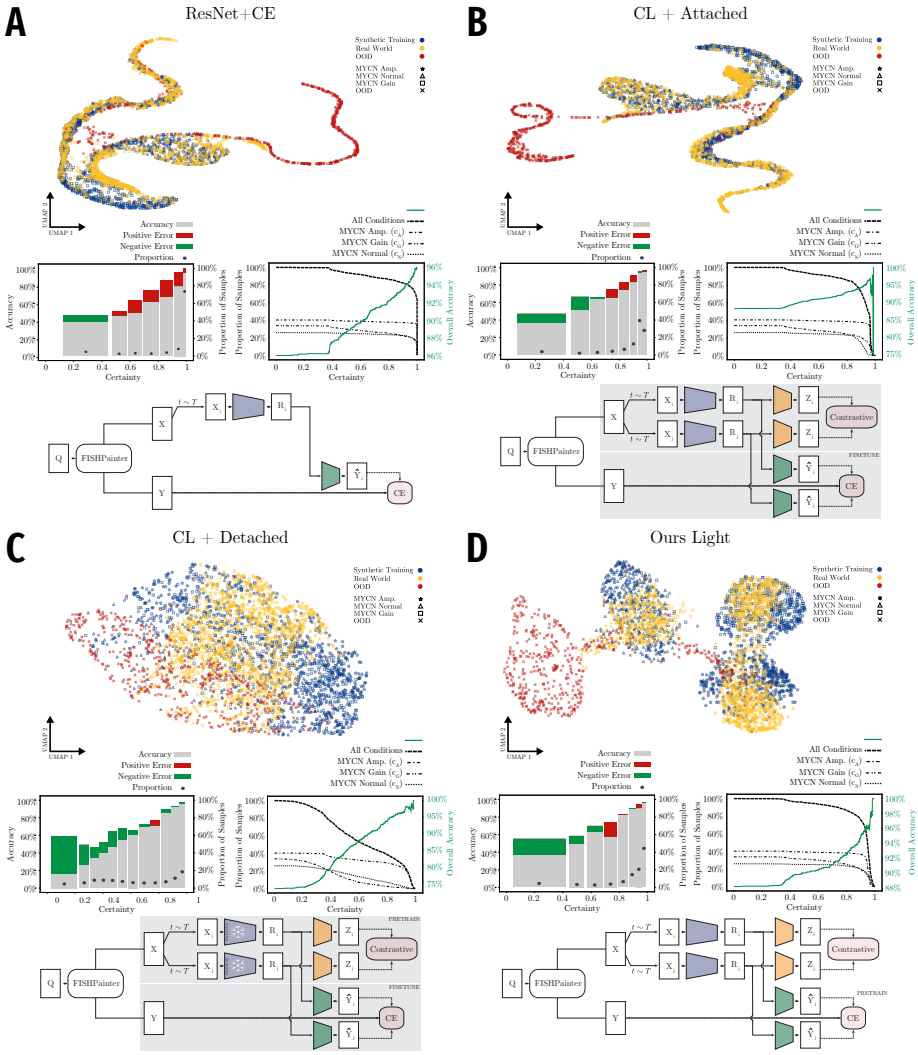

Fig. S1: The experimental results are shown in subplots **A-D** for Resnet+CE, CL+Attached, CL+Detached, and Ours Light. Each subplot displays latent representations of the synthetic training dataset (blue), real-world test set (yellow), and OOD dataset (red). Below, the left plot shows ECE (green for underconfidence, red for overconfidence), and the right plot shows accuracy and class distribution across certainty thresholds. A data flow diagram below these plots illustrates the training methods for each model. The snowflake in C indicates frozen weights during fine-tuning.

A

#Patches	Class	R. Spots	G. Spots	G. Clusters	G.Cluster Size	Signal Size
10000	MYCN Normal	2	2	0	0	1-2
10000	MYCN Gain	2	3-7	0	0	1-2
5000	MYCN Amp. Cluster	2	2-8	1-2	6-12	1-2
5000	MYCN Amp. Signals	2	8-20	0-1	8-16	1-2

B

	Affine		Gradient		Intensity		Noise		Blurr		Flip	
	Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value
Ablation	degrees_min	1	low	0	low RGB	[.75, .75, .75]	mean	[-0.1, 0.1]	std	[10^{-6}, 1]	hori. p	0.5
	degrees_max	360	high	0.5	high RGB	[1.5, 1.5, 1.5]	std	[10^{-6}, 0.07]			vert. p	0.5
	scale_min	0.5										
	scale_max	2.0										
Ours Light	not used		low	0	low RGB	[.75, .75, .75]	mean	[-0.1, 0.1]	std	[10^{-6}, 1]	hori. p	0.5
			high	0.5	high RGB	[1.5, 1.5, 1.5]	std	[10^{-6}, 0.07]			vert. p	0.5
Ours Heavy	degrees_min	1	low	0	low RGB	[.50, .50, .50]	mean	[-0.15, 0.15]	std	[10^{-6}, 1]	hori. p	0.5
	degrees_max	360	high	0.5	high RGB	[2.0, 2.0, 2.0]	std	[10^{-6}, 0.12]			vert. p	0.5
	scale_min	0.5										
	scale_max	2.0										
Baselines	degrees_min	1	low	0	low RGB	[.50, .50, .50]	mean	[-0.15, 0.15]	std	[10^{-6}, 1]	hori. p	0.5
	degrees_max	360	high	0.5	high RGB	[2.0, 2.0, 2.0]	std	[10^{-6}, 0.12]			vert. p	0.5
	scale_min	0.5										
	scale_max	2.0										

Examples

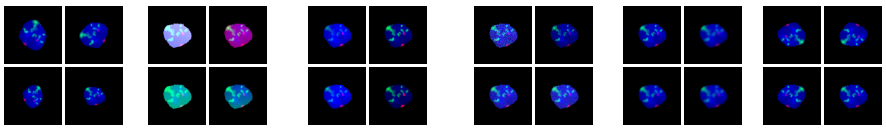

Fig. S2: **A**: Configurations Q of synthetic dataset: Classes are *MYCN Normal*, *MYCN Gain*, and *MYCN Amplified*, with details on spot counts, cluster presence, and signal sizes to simulate diverse FISH image scenarios. *MYCN Amplified* is divided into "Cluster" and "Signals". R.: red, G.: green. **B**: Specification of augmentations for the ablation experiment and the two implementations of our approach, *Ours Heavy* and *Ours Light*, along with examples showing the effect of the augmentations on the same cell.

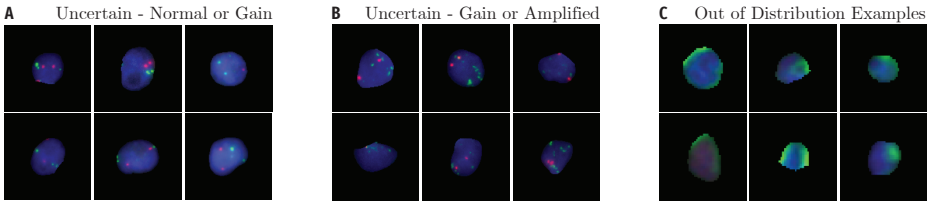

Fig. S3: **A-B**: High Uncertainty Cases for *Ours Heavy*. **A**: Images challenging for both humans and Ours Heavy to classify as *MYCN Normal* or *MYCN Gain*. **B**: Examples with ambiguous *MYCN* copy numbers, hard to classify as *MYCN Gain* or *MYCN Amplified*. **C**: OOD images embedded close to ID samples.