

Measuring Prediction Uncertainty in Neural Cellular Automata

Ario Sadafi^{1,2,3,4}, Michael Deutges¹, Nassir Navab^{3,4}, and
Carsten Marr^{1,2,4,5,6,7}

¹ Computational Health Center, Helmholtz Munich, Neuherberg, Germany

² Helmholtz AI, Helmholtz Munich, Neuherberg, Germany

³ Computer Aided Medical Procedures, Technical University of Munich, Munich, Germany

⁴ Munich Center for Machine Learning, Munich, Germany

⁵ Department of Medicine III, Ludwig-Maximilian-University Hospital, Munich, Germany

⁶ Department of Physics, University of Munich, Munich, Germany

⁷ German Cancer Consortium (DKTK), partner site Munich, Germany

Abstract. Neural cellular automata (NCA) provide a lightweight alternative to encoder–decoder segmentation networks. However, it can be difficult to decide when a prediction should be trusted. Here, we study uncertainty estimation for NCA-based medical image segmentation without modifying the underlying architecture or retraining the model. Our approach is motivated by viewing the NCA as a dynamical system where convergent attractors correspond to confident predictions.

Concretely, we propose *resilience*, a simple measure that leverages the intrinsic iterative structure of NCAs by probing the stability of the final prediction under small perturbations of the automaton state. Predictions that return to the same solution are deemed confident, while those that change substantially are flagged as uncertain.

We evaluate uncertainty by its ability to predict segmentation quality using selective prediction metrics ($\Delta\text{Dice@90}$ and AURC) and ranking metrics (AUROC and AUPRC). Across multiple medical segmentation benchmarks, *resilience* identifies failure cases more reliably than baselines, improving trust and safety in NCA-based models.

Keywords: Neural Cellular Automata · Uncertainty Quantification · Resilience.

1 Introduction

Medical image segmentation is a foundational step in many clinical and biomedical workflows, from quantifying structures in microscopy to delineating lesions and organs in radiology and endoscopy [16]. Yet even strong segmentation models can fail silently on unusual anatomy, poor image quality, or domain shifts across

Correspondence: {carsten.marr, ario.sadafi}@helmholtz-munich.de

devices and sites. In practice, this makes reliability as important as accuracy. A usable system should not only output a mask, but also provide a meaningful signal of when that mask can be trusted, reviewed, or rejected [8].

Uncertainty estimation has become the standard route to reliability, enabling selective prediction, quality control, and human-in-the-loop triage [6]. Most existing approaches, however, are developed for feed-forward architectures that produce predictions in a single pass, where uncertainty is tied to predictive probabilities, ensembling, or test-time augmentation. While effective in many settings, these techniques do not explicitly account for models whose predictions emerge from an iterative process and may therefore miss uncertainty cues that arise from the dynamics of inference itself.

Neural Cellular Automata (NCA) offer a lightweight and conceptually appealing alternative for segmentation [15]. Rather than relying on an encoder-decoder backbone, an NCA maintains a spatial state and repeatedly applies a shared local update rule until a segmentation is formed [5,10,12,13]. This dynamical view changes the nature of uncertainty. For an iterative model, an apparently sharp final mask can still be unreliable if the NCA dynamics have not converged to a stable prediction, i.e. if it oscillates between competing solutions, or if small disturbances steer the system toward a different outcome. These behaviors are described in terms of stability and attractors—properties that are not captured by a single forward-pass confidence score. Furthermore, NCA models have been applied beyond segmentation to tasks such as image classification and generative modeling [4,15,17,18]. This broader applicability motivates uncertainty measures that operate directly on the underlying dynamical system rather than on a task-specific architectures.

This work proposes a simple, NCA-native uncertainty measure that we call *resilience*. The central idea is to probe whether the final segmentation behaves like a stable solution of the learned dynamics. Starting from the final NCA state, a small perturbation is injected into the latent state and the automaton is allowed to evolve for a few additional steps. If it recovers to the same segmentation, the prediction is considered confident; if the prediction changes substantially, the output is considered uncertain. Unlike Monte Carlo dropout [7], which samples variability through stochastic subnetworks and requires architectural modification to include dropout layers, *resilience* tests the stability of the learned dynamical system and can be applied without retraining.

To place *resilience* in context, we evaluate several complementary uncertainty signals as baselines under a common protocol that assigns one scalar uncertainty score per image and measures how effectively uncertainty ranks predictions. Evaluation is performed across diverse medical segmentation benchmarks spanning microscopy, histology, dermoscopy, and endoscopy. Uncertainty is quantified using selective prediction criteria and failure-detection metrics enabling a consistent comparison across datasets with different baseline accuracies.

The main contributions are: (1) We provide, to our knowledge, the first systematic study of uncertainty estimation for NCA-based medical segmentation, covering four imaging modalities and explicit distribution shifts under a uni-

fied evaluation protocol. (2) We introduce *resilience*, a training-free perturb-and-recover uncertainty measure that directly probes the stability of the NCA’s learned dynamics, requiring neither architectural changes nor retraining. We test our approach on five benchmarks and show that *resilience* yields a strong and consistent signal for uncertainty-aware deployment of NCA-based models.

Our implementation, including all uncertainty baselines and evaluation scripts is publicly available at <https://github.com/marrlab/resilience>.

2 Methods

2.1 Neural Cellular Automata for Segmentation

We model image segmentation using Neural Cellular Automata (NCA), which defines a discrete dynamical system over a latent state [15, 11]. Given an RGB input image $x \in \mathbb{R}^{H \times W \times 3}$, the NCA maintains a state tensor

$$S_t \in \mathbb{R}^{H \times W \times C}, \quad t = 0, 1, \dots, T,$$

where C denotes the number of state channels. The initial state S_0 is constructed from the input image by concatenation with zero-initialized latent channels.

The states evolve according to a learned update rule $S_{t+1} = F_\theta(S_t)$, where F_θ is a local neural network applied identically at every spatial location and time step. After T iterations, a segmentation prediction is obtained from a designated output channel $m = o(S_T)$, where we denote $o(\cdot)$ as readout operator, and $m \in \{0, 1\}^{H \times W}$ as the predicted binary mask. Crucially, the NCA defines an explicit trajectory $\{S_t\}$ whose convergence behavior is governed by the learned dynamics.

2.2 Resilience Uncertainty via Perturb-and-Recover

We exploit the dynamical nature of the NCA to estimate prediction uncertainty by probing the stability of the final state under perturbations (Figure 1).

For a given test image, we first run the NCA for T steps to obtain the final state S_T and its associated prediction m . We then perturb the final state by adding Gaussian noise,

$$\tilde{S}_T = S_T + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 I),$$

where $\sigma > 0$ controls the perturbation magnitude and I denotes the identity over all state dimensions. Thus, a separate Gaussian noise value is sampled for every element of the state tensor, i.e., for every pixel location and every state channel; the same noise value is not shared across channels. Importantly, the input image remains fixed; only the latent NCA state is perturbed.

Starting from the perturbed state $\tilde{S}_T$, we apply an additional T' NCA update steps,

$$S_{T+T'} = F_\theta^{T'}(\tilde{S}_T),$$




Fig. 1. Uncertainty quantification via perturb-and-recover in Neural Cellular Automata (NCA). After iterating the NCA for T steps on an input image, a final state S_T and segmentation mask m are obtained. To probe the stability of the learned dynamics, Gaussian noise ϵ is injected and the system is allowed to relax for an additional T' steps. The resulting prediction m' is compared to the original mask using the Intersection-over-Union (IoU). The *resilience* method $u_{\text{res}} = 1 - \text{IoU}(m, m')$ measures sensitivity to state perturbations: low values indicate convergence to a stable attractor, while high values indicate unstable (and uncertain) segmentations.

and compute a segmentation prediction after recovery as $m' = o(S_{T+T'})$.

We quantify the degree of recovery by measuring the Intersection over Union $\text{IoU}(m, m') = \frac{|m \cap m'|}{|m \cup m'|}$ between the original and relaxed predictions. We define the uncertainty measure *resilience* as

$$u_{\text{res}} = 1 - \text{IoU}(m, m').$$

In all experiments, we used a perturbation magnitude of $\sigma = 0.02$ and a relaxation length of $T' = 12$. *Resilience* was computed using one perturb-and-recover rollout per image.

Low values of u_{res} indicate that the NCA dynamics recover to an equivalent segmentation after perturbation, suggesting a stable attractor (see Figure 1). Conversely, high values indicate sensitivity to small state perturbations and correspond to unstable predictions.

Our proposed uncertainty measure, *resilience*, can be interpreted as a local stability probe of the learned dynamical system around the final state S_T . When $\tilde{S}_T$ lies within the basin of attraction of a fixed point, perturbations return back to S_T under further NCA iterations and $m' \approx m$. In contrast, for challenging samples, the dynamics may result in divergent predictions and therefore increased uncertainty.

3 Experiments and Results

3.1 Datasets

We evaluate *resilience* on five medical image segmentation benchmarks. All tasks are formulated as binary segmentation.

- ClinicDB [1]: ClinicDB is a colonoscopy benchmark consisting of endoscopic frames with corresponding pixel-level polyp masks. We use a deterministic split to obtain 428 training, 91 validation, and 93 test images. Masks are binary polyp silhouettes.
- DSB 2018 [2]: We utilize the 2018 Data Science Bowl benchmark for nuclei segmentation. We generate a deterministic split by sorting the 671 available case folders and partitioning them 80/10/10, resulting in 537 training, 67 validation, and 67 test images. Instance masks are unified into a single binary foreground mask per image.
- ISIC 2017 [3]: This dataset for skin lesion segmentation is used with its official splits: 2,000 training, 150 validation, and 600 test images.
- Kvasir-SEG [9]: Consisting of 1,000 colonoscopy images, we define this task as a polyp-vs-background segmentation. We apply a 70/15/15 split resulting in 700 train, 150 validation, and 150 test images.
- NuInsSeg [14]: This benchmark spans multiple histology tissue domains for nuclei segmentation. We split the dataset into 465 training, 99 validation, and 101 test images, ensuring no slide-level overlap across splits.

To assess uncertainty under distribution shifts and harsher scenarios, we augmented the test set for ClinicDB and Kvasir-SEG datasets by applying geometric and photometric perturbations to the original test samples

3.2 Training

The same training settings across all datasets are used without dataset-specific tuning. While this is not optimal for maximizing segmentation performance on each benchmark, it provides a consistent and sufficiently challenging regime in which prediction errors, and therefore uncertainty, remain observable and comparable across datasets. We train an NCA segmentation model with $C=64$ state channels, fire rate 0.5, hidden size 128, and three RGB input channels. The fire rate denotes the probability that a cell updates its state at each NCA step. All models are optimized with Adam (learning rate 10^{-3} , weight decay 10^{-2}) for 50 epochs using a batch size of 2. Training uses a rollout-length schedule: at each iteration we sample the number of NCA update steps $T \sim \mathcal{U}(32, 64)$ and compute the loss on the final prediction after T steps. We minimize standard softmax cross-entropy over the output logits. No additional class weighting or auxiliary losses are used.

The segmentation performance of the NCA model across all benchmarks is summarized in Table 1. We report Pixel Accuracy, Mean IoU, and Dice Coefficient to provide a baseline for the model’s efficacy prior to uncertainty evaluation.

While the NCA achieves high performance on datasets like DSB2018, the lower mIoU scores on Kvasir-SEG and ClinicDB reflect the inherent difficulty of these tasks, providing a diverse range of error regimes for evaluating uncertainty methods.

Table 1. Segmentation performance across five medical benchmarks. Mean and standard deviation is reported for 3 runs.

Dataset	Pixel Acc.	Mean IoU	Dice
ClinicDB	0.91±0.01	0.58±0.01	0.74±0.01
DSB 2018	0.97±0.01	0.89±0.01	0.94±0.01
ISIC 2017	0.86±0.01	0.67±0.01	0.80±0.01
Kvasir-SEG	0.83±0.01	0.43±0.01	0.60±0.01
NuInsSeg	0.89±0.01	0.62±0.04	0.77±0.03

3.3 Baselines for Uncertainty Quantification

The proposed *resilience* method is compared against uncertainty signals derived from NCA outputs and trajectories. Each method assigns a scalar uncertainty score u per test image which is used to rank predictions from most to least confident for selective prediction and failure detection (See Table 2). For map-based methods, a pixel-wise uncertainty map $U(x)$ is computed and aggregated over a thin boundary band $B(m)$ around the predicted mask m (dilation minus erosion). Ranking uses the boundary mean $u = \frac{1}{|B|} \sum_{x \in B(m)} U(x)$ while the 95th percentile (p95) within $B(m)$ is used only as an additional reporting statistic and not for ranking.

Single-pass entropy ("single"). A single rollout to T yields softmax probabilities; uncertainty is the per-pixel entropy map $U(x) = -\sum_c p_c(x) \log p_c(x)$.

Stopping-time disagreement ("stoptime"). Masks from intermediate rollouts with $T_k \sim \mathcal{U}(T_{\min}, T_{\max})$ are compared to the final mask; $U(x)$ is the fraction of steps where $m^{(T_k)}(x) \neq m^{(T_{\max})}(x)$.

Late-stage drift ("stability"). Over the last W steps, $U(x)$ is the mean absolute change in foreground probability, $U(x) = \frac{1}{W} \sum |p_t(x) - p_{t-1}(x)|$.

Binary flicker ("flicker"). Foreground probabilities are thresholded ($\tau=0.5$) and $U(x)$ is the flip frequency over a window of W steps.

Test-time augmentation ("TTA"). Deterministic rotations/flips are applied to the input, predictions are inverse-warped and averaged; uncertainty is taken as the entropy of the averaged probability map.

3.4 Uncertainty Measures

We evaluate uncertainty using selective prediction and failure detection. Each method produces an image-level uncertainty score u ; images are ranked from most confident (low u) to least confident (high u). For selective prediction, we

Table 2. Quantitative evaluation of uncertainty measures across datasets. Selective prediction is reported via $\Delta\text{Dice@90}$ (percentage-point Dice gain when retaining the 90% most confident images) and AURC, and failure detection via AUROC and AUPRC. Values are mean $\pm$ std over 3 runs.

Dataset	Method	$\Delta\text{Dice@90} \uparrow$	AURC $\downarrow$	AUROC $\uparrow$	AUPRC $\uparrow$
ClinicDB	single	3.76%	0.798 $\pm$ 0.033	0.348 $\pm$ 0.057	0.199 $\pm$ 0.044
	stability	-5.92%	0.835 $\pm$ 0.020	0.437 $\pm$ 0.078	0.272 $\pm$ 0.034
	stoptime	-7.17%	0.821 $\pm$ 0.010	0.477 $\pm$ 0.037	0.279 $\pm$ 0.026
	flicker	-1.17%	0.738 $\pm$ 0.021	0.517 $\pm$ 0.091	0.245 $\pm$ 0.040
	TTA	-0.08%	0.779 $\pm$ 0.027	0.392 $\pm$ 0.069	0.248 $\pm$ 0.071
	resilience	8.13%	0.719$\pm$0.012	0.641$\pm$0.014	0.443$\pm$0.056
DSB 2018	single	3.57%	0.111 $\pm$ 0.004	0.887 $\pm$ 0.037	0.703 $\pm$ 0.095
	stability	4.18%	0.106 $\pm$ 0.002	0.948 $\pm$ 0.029	0.819 $\pm$ 0.101
	stoptime	4.40%	0.114 $\pm$ 0.002	0.921 $\pm$ 0.029	0.794 $\pm$ 0.057
	flicker	4.00%	0.109 $\pm$ 0.001	0.937 $\pm$ 0.016	0.824 $\pm$ 0.037
	TTA	4.04%	0.111 $\pm$ 0.004	0.873 $\pm$ 0.048	0.686 $\pm$ 0.116
	resilience	4.06%	0.103$\pm$0.003	0.972$\pm$0.009	0.914$\pm$0.034
ISIC 2017	single	1.27%	0.335 $\pm$ 0.016	0.534 $\pm$ 0.008	0.230 $\pm$ 0.011
	stability	1.22%	0.280 $\pm$ 0.023	0.623 $\pm$ 0.025	0.271 $\pm$ 0.021
	stoptime	1.59%	0.293 $\pm$ 0.014	0.592 $\pm$ 0.013	0.262 $\pm$ 0.014
	flicker	3.39%	0.249 $\pm$ 0.010	0.708 $\pm$ 0.017	0.336 $\pm$ 0.002
	TTA	4.47%	0.228 $\pm$ 0.008	0.758 $\pm$ 0.015	0.462 $\pm$ 0.026
	resilience	7.61%	0.176$\pm$0.003	0.910$\pm$0.007	0.755$\pm$0.004
Kvasir-SEG	single	-1.73%	0.886 $\pm$ 0.074	0.476 $\pm$ 0.155	0.229 $\pm$ 0.043
	stability	7.46%	0.873 $\pm$ 0.015	0.513 $\pm$ 0.203	0.276 $\pm$ 0.097
	stoptime	5.45%	0.817 $\pm$ 0.025	0.519 $\pm$ 0.091	0.334 $\pm$ 0.094
	flicker	0.01%	0.873 $\pm$ 0.058	0.538 $\pm$ 0.135	0.325 $\pm$ 0.116
	TTA	2.08%	0.794$\pm$0.059	0.616 $\pm$ 0.159	0.344 $\pm$ 0.102
	resilience	9.26%	0.816 $\pm$ 0.043	0.718$\pm$0.105	0.500$\pm$0.176
NuInsSeg	single	0.55%	0.502 $\pm$ 0.072	0.436 $\pm$ 0.102	0.193 $\pm$ 0.033
	stability	-0.81%	0.546 $\pm$ 0.082	0.492 $\pm$ 0.280	0.369 $\pm$ 0.255
	stoptime	-2.23%	0.533 $\pm$ 0.069	0.465 $\pm$ 0.189	0.289 $\pm$ 0.112
	flicker	1.09%	0.391$\pm$0.022	0.857$\pm$0.062	0.587 $\pm$ 0.123
	TTA	1.21%	0.472 $\pm$ 0.052	0.503 $\pm$ 0.075	0.227 $\pm$ 0.051
	resilience	7.63%	0.394 $\pm$ 0.006	0.797 $\pm$ 0.150	0.688$\pm$0.208

report $\Delta\text{Dice@90}$, defined as $100 \times (\text{Dice@90} - \text{Dice})$, i.e., the percentage-point change in Dice when keeping only the 90% most confident images compared to using all images, and AURC, the area under the risk-coverage curve with risk = $1 - \text{Dice}$ on the retained set (lower is better). For failure detection, we treat u as a score for identifying low-quality segmentations and report AUROC and AUPRC, which measure how well uncertainty separates poor predictions from good ones. AUPRC is most informative under class imbalance.




Fig. 2. Our *resilience* method for uncertainty estimation highlights confident and failed segmentations. For each dataset, one low- and one high-uncertainty example are shown; low uncertainty corresponds to more accurate masks, while high uncertainty captures failures.

3.5 Results and Discussion

Table 2 reports selective prediction and failure-detection performance across datasets. Overall, *resilience* provides the most consistent uncertainty signal, yielding strong $\Delta\text{Dice}@90$ gains and competitive failure detection. Method preferences are dataset-dependent in a few cases: *stoptime* attains the best $\Delta\text{Dice}@90$ on DSB 2018, *TTA* gives the lowest AURC on Kvasir-SEG, and *flicker* performs best on NuInsSeg in AURC/AUROC, suggesting complementary uncertainty cues across datasets.

Figure 2 shows typical low and high uncertainty examples. Low uncertainty cases exhibit clearer structure and more stable masks, whereas high uncertainty cases tend to involve ambiguous boundaries, low contrast, artifacts, or atypical appearance. On challenging datasets (e.g., Kvasir-SEG), even the most confident predictions may be imperfect, reflecting the difficulty of the problem rather than a failure of the uncertainty estimate.

The overall pattern indicates that probing the stability of NCA dynamics is a reliable proxy for segmentation trustworthiness: *resilience* is a strong default across datasets and metrics, particularly when selective prediction at high coverage is the target. At the same time, the strongest competing baselines vary by dataset, indicating that different uncertainty families capture different failure modes. For example, *stoptime* reflects sensitivity to rollout length, *TTA* reflects appearance sensitivity under input transformations, and *flicker* captures temporal oscillations in the predicted mask. These observations motivate the conclusion that *resilience* can be used as a robust general-purpose uncertainty signal for NCA models, while hybrid or dataset-adaptive combinations may be beneficial when prioritizing a specific operating point or metric, such as minimizing AURC versus maximizing $\Delta\text{Dice}@90$.

4 Conclusion

We introduced *resilience*, a highly efficient training-free uncertainty estimation method tailored to NCA that exploits the model’s iterative dynamics by testing whether a final prediction returns to the same solution after small state perturbations. Across five diverse biomedical segmentation benchmarks, *resilience* produced a consistently strong reliability signal: it improved selective prediction (higher $\Delta\text{Dice@90}$, lower AURC) and achieved competitive failure detection (AUROC / AUPRC), while remaining lightweight and easy to integrate with existing NCA pipelines. Qualitative results show that high uncertainty concentrates on inherently challenging cases. Furthermore, *resilience* is readout-agnostic and only probes the stability of the underlying NCA dynamics, making it directly applicable to other NCA tasks beyond segmentation. Overall, these findings highlight that stability-driven uncertainty can serve as a simple, model-agnostic reliability layer for NCA, helping bridge the gap between label-efficient methods and clinical practicality by making errors more predictable, enabling selective use of outputs, and focusing expert attention where it matters most.

Acknowledgments. C.M. acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant Agreement No. 866411 & 101113551 & 101213822) and support from the High-tech Agenda Bayern.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bernal, J., Sánchez, F.J., Fernández-Esparrach, G., Gil, D., Rodríguez, C., Vilar-íño, F.: Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Computerized medical imaging and graphics* **43**, 99–111 (2015)
2. Caicedo, J.C., Goodman, A., Karhohs, K.W., Cimini, B.A., Ackerman, J., Haghighi, M., Heng, C., Becker, T., Doan, M., McQuin, C., et al.: Nucleus segmentation across imaging experiments: the 2018 data science bowl. *Nature methods* **16**(12), 1247–1253 (2019)
3. Codella, N.C., Gutman, D., Celebi, M.E., Helba, B., Marchetti, M.A., Dusza, S.W., Kalloo, A., Liopyris, K., Mishra, N., Kittler, H., et al.: Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In: 2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018). pp. 168–172. IEEE (2018)
4. Deutges, M., Sadafi, A., Navab, N., Marr, C.: Neural cellular automata for lightweight, robust and explainable classification of white blood cell images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 693–702. Springer (2024)

5. Deutges, M., Yang, C., Salehi, R., Navab, N., Marr, C., Sadafi, A.: Neural cellular automata for weakly supervised segmentation of white blood cells. In: International Workshop on Efficient Medical Artificial Intelligence. pp. 289–298. Springer (2025)
6. Dolezal, J.M., Srisuwananukorn, A., Karpeyev, D., Ramesh, S., Kochanny, S., Cody, B., Mansfield, A.S., Rakshit, S., Bansal, R., Bois, M.C., et al.: Uncertainty-informed deep learning models enable high-confidence predictions for digital histopathology. *Nature communications* **13**(1), 6572 (2022)
7. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: international conference on machine learning. pp. 1050–1059. PMLR (2016)
8. González, C., Gotkowski, K., Fuchs, M., Bucher, A., Dadras, A., Fischbach, R., Kaltenborn, I.J., Mukhopadhyay, A.: Distance-based detection of out-of-distribution silent failures for covid-19 lung lesion segmentation. *Medical image analysis* **82**, 102596 (2022)
9. Jha, D., Smedsrud, P.H., Riegler, M.A., Halvorsen, P., De Lange, T., Johansen, D., Johansen, H.D.: Kvasir-seg: A segmented polyp dataset. In: International conference on multimedia modeling. pp. 451–462. Springer (2019)
10. Kalkhof, J., González, C., Mukhopadhyay, A.: Med-nca: Robust and lightweight segmentation with neural cellular automata. In: International Conference on Information Processing in Medical Imaging. pp. 705–716. Springer (2023)
11. Kalkhof, J., Ihm, N., Koehler, T., Gregori, B., Mukhopadhyay, A.: Med-nca: Bio-inspired medical image segmentation. *Medical Image Analysis* **103**, 103601 (2025)
12. Kalkhof, J., Mukhopadhyay, A.: M3d-nca: Robust 3d segmentation with built-in quality control. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 169–178. Springer (2023)
13. Lemke, N., Kalkhof, J., Babendererde, N., Mukhopadhyay, A.: Octreenca: Single-pass 184 mp segmentation on consumer hardware. *arXiv preprint arXiv:2508.06993* (2025)
14. Mahbod, A., Polak, C., Feldmann, K., Khan, R., Gelles, K., Dorffner, G., Woitek, R., Hatamikia, S., Ellinger, I.: Nuinsseg: a fully annotated dataset for nuclei instance segmentation in h&e-stained histological images. *Scientific Data* **11**(1), 295 (2024)
15. Mordvintsev, A., Randazzo, E., Niklasson, E., Levin, M.: Growing neural cellular automata. *Distill* **5**(2), e23 (2020)
16. Wang, R., Lei, T., Cui, R., Zhang, B., Meng, H., Nandi, A.K.: Medical image segmentation using deep learning: A survey. *IET image processing* **16**(5), 1243–1267 (2022)
17. Yang, C., Deutges, M., Liu, J., Li, H., Navab, N., Marr, C., Sadafi, A.: Attention pooling enhances nca-based classification of microscopy images. In: International Workshop on Machine Learning in Medical Imaging. pp. 583–593. Springer (2025)
18. Yang, C., Deutges, M., Navab, N., Sadafi, A., Marr, C.: Hierarchical neural cellular automata for lightweight microscopy image classification. In: International Conference on Information Processing in Medical Imaging. pp. 19–32. Springer (2025)