Train Once, Generate Anywhere: Discretization Agnostic Neural Cellular Automata using SPH Method

Hyunsoo Kim
Korea Advanced Institute of Science and Technology (KAIST)
Daejeon, Republic of Korea
khskhs@kaist.ac.kr

Jinah Park
Korea Advanced Institute of Science and Technology (KAIST)
Daejeon, Republic of Korea
jinahpark@kaist.ac.kr

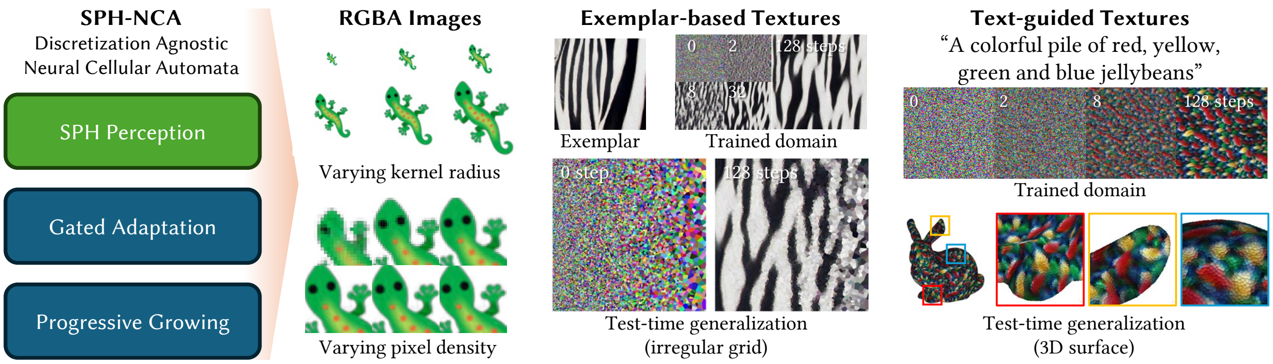

Figure 1: While trained on a regular grid of fixed resolution, our method, SPH-NCA, can generalize in diverse samplings.

CCS Concepts

• **Computing methodologies** → **Self-organization**; **Appearance and texture representations**; • **Mathematics of computing** → *Numerical differentiation.*

Keywords

Neural Cellular Automata, Smoothed Particle Hydrodynamics, Texture synthesis, Discretization agnostic

ACM Reference Format:
Hyunsoo Kim and Jinah Park. 2025. Train Once, Generate Anywhere: Discretization Agnostic Neural Cellular Automata using SPH Method. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters (SIGGRAPH Posters '25), August 10–14, 2025, Vancouver, BC, Canada.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3721250.3742998

1 Introduction

Image and texture synthesis is crucial for visually rich content, yet most techniques are limited to regularly sampled data. Neural Cellular Automata (NCA) introduce a self-organizing approach to neural image and texture generation, extending classical cellular automata by integrating trainable neural networks to govern local cell state transitions, making them lightweight and robust [Mordvintsev et al. 2020]. While recent studies have explored various perception schemes for diverse data structures (e.g., GraphNCA [Grattarola et al. 2021] and MeshNCA [Pajouheshgar et al. 2024]), these methods often overfit to the connectivity of the training data, making it difficult to generalize to topologies and geometries outside the trained domain.

To overcome these limitations, we propose SPH-NCA[1], a novel framework that combines the self-organizing NCA architecture with the mesh-free spatial discretization of Smoothed Particle Hydrodynamics (SPH) methods for perception. In contrast to other works, our SPH perception can estimate the cell's local gradient on arbitrary spatial discretizations without relying on explicit connectivity, circumventing the need for training data augmentation and achieving discretization-agnostic generation. Alongside SPH perception, we also present gated adaptation and progressive growing training scheme to stabilize the training of SPH-NCA models. We demonstrate SPH-NCA's test-time generalization across diverse surface samplings, from regular 2D grids to complex irregular 3D surfaces, showcasing its ability to train once on regular data and generalize to arbitrary geometries.

2 Methods

This section details the SPH-NCA method, focusing on our key contributions: SPH perception, gated adaptation, and the progressive growing training scheme. Figure 2 provides an overview of the architecture and its main components, while additional details can be found in the supplemental material.

SPH Perception. We replace Sobel filter-based gradient estimation with a diferentiable SPH method with the kernel radius h: $\nabla S_i = \texttt{SPHGrad}(\{\mathbf{x}_i, S_i\}, h)$. The resulting $3C$-dimension perception vector Z_i is concatenated from C-dimensional S_i and $2C$-dimensional 2D

ACM ISBN 979-8-4007-1549-5/25/08
https://doi.org/10.1145/3721250.3742998

[1] The code is available at https://github.com/hyunsoo0o0o0/SPH-NCA

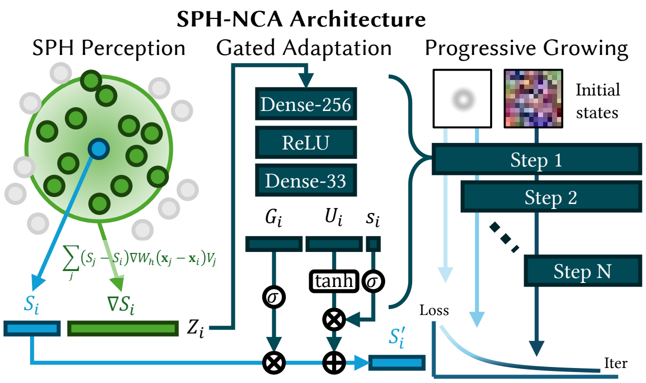

Figure 2: Architecture of SPH-NCA. The perception stage utilized the SPH method to compute the gradient from irregular samples. The adaptation stage uses a gating mechanism to update particle states. During the training of SPH-NCA models, progressively increasing the number of steps stabilizes the training.

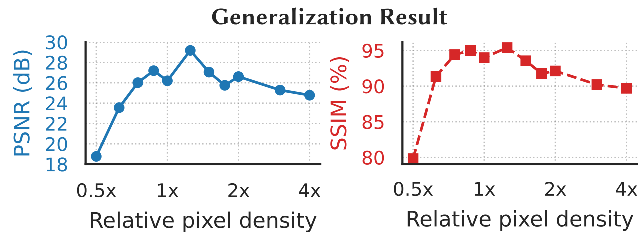

Figure 3: SPH-NCA can preserve the image quality while inferring outside of the trained pixel density.

projection of ∇S_i. On 3D surfaces, ∇S_i is projected onto the local tangent planes.

Gated Adaptation. SPH-NCA replaces the typical simple additive update with a novel gating mechanism to prevent state explosion and saturation during training. The perception vector Z_i is fed through a two-layer MLP, yielding update components (G_i, U_i, s_i). The next cell state S'_i is then computed via $S'_i = \sigma(G_i) \odot S_i + \sigma(s_i)\tanh(U_i)$, where G_i controls the retention of the previous state, U_i determines the change, and s_i controls its magnitude.

Progressive Growing. We introduce a progressive growing strategy for stable SPH-NCA training. Initial training begins with a single NCA step, incrementally increasing the step count by one every k iterations until a maximum of N steps is reached. Subsequently, the step size is uniformly sampled within $[N-\alpha, N+\alpha]$. For image synthesis tasks, we employ $(k, N, \alpha) = (10, 40, 8)$, ensuring robust convergence.

3 Experiments and Results

RGBA Image Synthesis and Scale Control. Using the Noto emoji dataset and an initial single-blob state, we demonstrate SPH-NCA's ability to control feature scale during inference by varying the SPH kernel radius h. This is achieved by scaling the gradient term in Z_i as $(S_i, \frac{h}{h_0}\nabla S_i)$. Figure 1 shows the generation result while changing kernel radius and pixel density. The feature scale is directly governed by the SPH kernel radius (h), which controls the perception range and thus information propagation speed. Our SPH formulation makes the model inherently agnostic to pixel density, decoupling feature scale from the underlying grid resolution. This allows for high-quality generation across a wide range of densities from a single trained model, a claim validated by the quantitative analysis in Figure 3. Performance remains stable except at extremely sparse discretizations where SPH gradient estimation is compromised by an insufficient particle neighborhood.

Exemplar-based Texture Synthesis and Generalization over Grid Regularity. We employ the dataset and optimal transport style loss from [Mordvintsev and Niklasson 2021]. We show that SPH-NCA models trained solely on a regular grid can generalize to irregular grids created by adding random jitters and removing pixels randomly. The result shown in Figure 1 demonstrates the robustness of SPH-NCA on a grid with varying regularity.

Text-guided Texture Synthesis and 2D-to-3D Transfer. We use OpenCLIP [Cherti et al. 2023] for a multi-scale image-text alignment loss to perform text-guided texture synthesis. Figure 1 shows the text-guided textures and its generation on a 3D surface. SPH-NCA can synthesize even in high-curvature regions on the Stanford Bunny, demonstrating its robust 2D-to-3D transfer capabilities.

4 Conclusion

We introduce SPH-NCA, a discretization-agnostic neural cellular automata framework built upon SPH perception, gated adaptation, and progressive growing. We demonstrate that SPH-NCA can generalize across diverse surface types–regardless of resolution, grid regularity, or geometry–and can truly *train once, generate anywhere.*

Our work highlights the potential of SPH methods for spatial data learning within NCA architectures. We aim to explore the application of differentiable SPH methods in other spatial data processing for generative modeling on diverse geometries.

Acknowledgments

This research was partly supported by Institute for Information & communications Technology Promotion(IITP) grant funded by the Korea government(MSIT) (No.00223446, Development of object-oriented synthetic data generation and evaluation methods) and the KAIST Venture Research Program for Graduate & Ph.D students.

References

Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. 2023. Reproducible scaling laws for contrastive language-image learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2818–2829.

Daniele Grattarola, Lorenzo Livi, and Cesare Alippi. 2021. Learning graph cellular automata. *Advances in Neural Information Processing Systems* 34 (2021), 20983–20994.

Alexander Mordvintsev and Eyvind Niklasson. 2021. μ NCA: Texture Generation with Ultra-Compact Neural Cellular Automata. *arXiv preprint arXiv:2111.13545* (2021).

Alexander Mordvintsev, Ettore Randazzo, Eyvind Niklasson, and Michael Levin. 2020. Growing neural cellular automata. *Distill* 5, 2 (2020), e23.

Ehsan Pajouheshgar, Yitao Xu, Alexander Mordvintsev, Eyvind Niklasson, Tong Zhang, and Sabine Süsstrunk. 2024. Mesh neural cellular automata. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–16.