# UNRAVELING NEURAL CELLULAR AUTOMATA FOR LIGHTWEIGHT IMAGE COMPRESSION

Anonymous authors

Paper under double-blind review

### Abstract

Neural Cellular Automata (NCA) are computational models inspired by cellular growth, capable of learning complex behaviors through local interactions. While NCAs have been applied to various tasks like image restoration and synthesis, their potential for image compression remains largely unexplored. This paper aims to unravel the capabilities of NCAs for lightweight image compression by introducing a Grid Neural Cellular Automata (GNCA) training strategy. Unlike traditional methods that depend on large deep learning models, NCAs offer a low-cost compact and highly parallelizable alternative with intrinsic robustness to noise. Through experiments on the COCO 2017 dataset, we compare the compression performance of NCAs against JPEG, JPEG-2000 and WebP, using the metrics PSNR, SSIM, and MSE and Compression Rate. Our results demonstrate that NCAs achieve competitive compression rates and image quality reconstruction, highlighting their potential as a lightweight solution for efficient image compression. The code will be available upon acceptance.

023 024 025

026

004

010 011

012

013

014

015

016

017

018

019

021

### 1 INTRODUCTION

Cellular Automata (CA) (Von Neumann, 2017) are computational models inspired by the process of cellular growth and differentiation, where these cells evolve over discrete time steps according to predefined rules and neighbouring cell states. Originating from the pioneering work of John von Neumann and Stephen Wolfram, with Conway's Game of Life (Adamatzky, 2010), Cellular Automata have found applications across various fields, including physics, biology, computer science, and social sciences (Hoekstra et al., 2010), due to their ability to capture and simulate a wide range of phenomena with simple yet powerful computational principles.

Building on the foundational concepts of traditional Cellular Automata, Neural Cellular Automata (NCA) introduces an advanced approach by integrating Deep Neural Network architectures to define their transition functions. Unlike CAs, the updated rules are not predefined but learned through training. Applications of NCA (Mordvintsev et al., 2020; Palm et al., 2022) have demonstrated the ability to reconstruct an image from a pixel seed by learning local neighbour interactions in a self-organized manner. Although many promising applications have been proposed, ranging from content generation (Sudhakaran et al., 2021) to self-organized control (Variengien et al., 2021), NCAs have not been explored in the context of data representation for image compression.

- State-of-the-art (SOTA) methods for image compression typically rely on large deep neural networks (Zhu et al., 2022; Sadeeq et al., 2021). While autoencoders have been successfully used for image compression (Liu et al., 2022), these models often come with significant storage demands due to their size, which is frequently overlooked when calculating the storage size for compressed representations. By contrast, NCAs present a lightweight alternative with several compelling properties, such as asynchronous operation with local interactions, high parallelization, and noise tolerance due to their self-organized reconstruction capabilities.
- This work proposes a novel approach to image compression using Neural Cellular Automata (NCAs). The key advantages of this approach include its parallelization capability, robustness to data corruption, and model compactness. Unlike traditional neural network-based solutions where the compressed representation is a latent vector, the proposed solution based on NCAs use the model weights themselves as the compressed data, reconstructing the image from a single pixel seed. The central research question we address is: : *How efficient are Neural Cellular Automata for Image*

*Compression*? To answer this, we utilized the NCA model proposed by Mordvintsev et al. (2020) and compared its performance against classical compression methods such as JPEG (Wallace, 1991), JPEG-2000 (Marcellin et al., 2000) and WebP (Si & Shen, 2016). We assessed the quality of compressed images using the metrics Peak Signal-to-Noise Ratio (PSNR) (Shirani, 2008), Structural Similarity Index Measure (SSIM) (Ding et al., 2020), Mean Squared Error (MSE) Thompson (1968) and Compress Rate (CR). The experiments were conducted on images from the COCO 2017 dataset (Lin et al., 2014).

- The main contributions of this paper are:
  - To the best of our knowledge, our work is the first to investigate the potential use of Neural Cellular Automata for image compression, identifying benefits and limitations compared to traditional methods;
  - We propose a grid training strategy to compress high-resolution images using NCAs, demonstrating high parallelization performance;
  - We investigate NCA parameters for image compression and highlight NCAs as a novel paradigm for compact image representation.

### 2 BACKGROUND

075 076

063

064

065

066 067

068

069 070

071

072 073 074

Neural net-based formulations of CAs can be traced back to the early work of Wulff & Hertz (1992).
Recent formulations of NCAs have shown their ability to learn complex desired behaviour, such as semantic segmentation (Sandler et al., 2020), common reinforcement learning tasks (Variengien et al., 2021), 3D locomotion (Najarro et al., 2022), Atari game playing (Najarro et al., 2022), and image synthesis (Palm et al., 2022).

Works by Mordvintsev et al. (2020) and Palm et al. (2022) show the capability of NCAs to reconstruct an image from a seed through learned neighbour interaction. Menta et al. (2024) uses NCA for resource-efficient image restoration, showing its potential for noise tolerance. In the context of image compression, the early work of Paul et al. (1999) proposed using Cellular Automata transformations to compress images. However, CAs have not been investigated as an efficient compression method since then.

Recent advancements in image compression have been driven by deep learning techniques (Mishra et al., 2022). Various compression algorithms based on convolutional neural networks (CNN) (Li et al., 2021), recurrent neural networks (RNN) (Medsker & Jain, 1999), autoencoders (Liu et al., 2022), and generative adversarial networks (GAN) (Gui et al., 2021) have shown significant improvements over traditional codecs like JPEG and WebP, particularly in terms of perceptual quality and rate-distortion performance (Sadeeq et al., 2021).

094 Using autoencoders and deep neural networks for image compression has some disadvantages de-095 spite their potential for high compression ratios and improved perceptual quality. These drawbacks 096 include the need for large amounts of training data, large models, and significant processing power 097 and memory resources for training. The SOTA MLIC++ (Jiang & Wang, 2023) uses a model with 098 83.5 million parameters, and the model size is not considered when calculating the compressed stor-099 age size, although these parameters are necessary to reconstruct the image. Large architectures like these demand extensive computational resources to train, encode and decode images (Mishra et al., 100 2022), making them unsuitable for lightweight devices. 101

This work proposes the use of NCAs to provide a lightweight solution with high parallelization capabilities, learning to compress and reconstruct images as self-organizing structures. Unlike autoencoders, which generate a latent compressed representation, we propose that NCAs use the model weights as encrypted, compressed representation form. This approach enables the image to be reconstructed from a pixel seed using the trained model, positioning the model itself as the compressed representation. To the best of our knowledge, this is the first study to explore NCAs for image com-

107 representation. To the best of our knowledge, this is the first study to explore NCAs for image compression.

### <sup>108</sup> 3 METHOD

## 110 3.1 PROBLEM DEFINITION111

Formally, image compression can be defined as a function  $C : \mathcal{I} \to C(\mathcal{I})$  that maps an original image I to a compressed representation C(I), where C(I) contains fewer bits than the original image while preserving important perceptual or structural features.

The goal of image compression is to minimize the distortion introduced during the compression process while significantly reducing data size. Image compression algorithms aim to minimize some distortion metric  $\mathcal{D}$  under a constraint on the size of C(I). This can be formulated as an optimization problem min<sub>C</sub>  $\mathcal{D}(\mathcal{I}, \mathcal{C}(\mathcal{I}))$  subject to a size constrain on  $\mathcal{C}(\mathcal{I})$ , where  $\mathcal{D}(\mathcal{I}, \mathcal{C}(\mathcal{I}))$  represents the distortion between the original and compressed images.

120 121

### 3.2 NEURAL CELLULAR AUTOMATA (NCA)

122 123 Cellular Automata (CA) are defined as a tuple  $M = (L, S, \eta, \phi)$ , where L is a d-dimensional lattice, 124 S is the set of possible states,  $\eta : L \to 2^L$  is the neighbourhood function and  $\phi : S^n \to S$  is the 125 transition function. Let  $s_i^t \in S$  indicate the state of cell I at time step t and  $\Omega_i^t = s_i^t : j \in \eta(i)$  the 126 set of neighbours of the cell I according to  $\eta$ . Then, the state update for cell i at each iteration t is 127 computed as  $s_i^{t+1} = \phi(s_i^t \cup \Omega_i^t)$ .

Neural Cellular Automata (NCAs) extend the CA framework by replacing the pre-defined transi-128 tion function  $\phi$  with a learnable function represented by a neural network  $\phi_{\theta}$  with parameters  $\theta$ , 129 preserving the essential feature of the locality. NCA can be trained to reconstruct images starting 130 from pixel seeds, which means to start the reconstruction from a blank image with a seed pixel and, 131 through learned neighbour rules, reconstruct the entire image (Mordvintsev et al., 2020), as shown 132 in Figure 1. The original NCA approach trains a single model for each image (Mordvintsev et al., 133 2020), which is a lightweight model by design. Although alternative methods have been proposed to 134 train one model for multiple images, they often require larger models or external support (Hernandez 135 et al., 2021), which we do not explore in this work.



Figure 1: Reconstruction of an image using a trained NCA, starting from pixel seed. The figure shows reconstruction through iterations.

The NCA, as defined by Mordvintsev et al. (2020), starts from a constant zero-filled initial state called seed and evolves this state over time according to its update rule. In this model, each cell is composed of a *C*-depth state. The update rule of this NCA consists of two parts: Perception and Stochastic Update. At the perception stage, each cell gathers information from its surrounding neighbours, forming the perception vector  $z_{ij} \in \mathbb{R}^{3C}$ , defined as:

149 150

141

142

143

$$z_{ij} = \operatorname{concat}(s_{ij}, K_x * s_{ij}, K_y * s_{ij}), \tag{1}$$

where  $K_x$  and  $K_y$  are predefined Sobel filters related to x and y axes, and \* stands for 2D-depthwise convolution operation using a  $3 \times 3$  kernel size.

The perceived information is passed through a Multilayer Perceptron (MLP) (Almeida, 2020) neural network to compute the new state  $s_{ij}^{t+1} = \text{MLP}(z_{ij}^t)$ . The NCA is trained using backpropagation through time (BPTT) to compute gradients of the loss concerning the neural network weights. The reconstruction loss is defined by the mean squared error between the final state  $S^t$  and target state  $S^*$ , as follows:

159

160

160  
161
$$\mathcal{L} = \frac{1}{T} \sum_{t=1}^{T} \mathcal{L}^t,$$
(2)

 $\mathcal{L}^{t} = \frac{1}{HW} \sum_{i,j} (S_{ij}^{t} - S^{*})^{2},$ (3)

163 164

162

165 166

167

184

185

where H and W are the height and width of a 2-dimensional lattice.

### 168 3.3 NCA FOR IMAGE COMPRESSION

170 We propose using NCA as a compression model, as they can reconstruct an image from a pixel 171 seed using a lightweight model. Different from Autoencoders, which extract a compressed latent representation of an image, in the proposed application of NCA, the compressed representation is 172 the model itself since it can reconstruct the image from a pixel seed. As proposed by Mordvintsev 173 et al. (2020), the NCA requires a different model for each input image to generate the compressed 174 representation. Although some NCA variations enable one model to build many images (Menta 175 et al., 2024), they require larger models and are not considered in this analysis. As the NCA proposed 176 by Mordvintsev et al. (2020) is a low-resource model, the size of the required weights to reconstruct 177 the image is lower than the input image size. Therefore, it is capable of generating a compressed 178 and efficient representation. 179

180 When defining an NCA architecture for image compression, the objective is to define a model with 181 the lowest number of parameters but the highest reconstruction quality. From a compression per-182 spective, the compression rate is based on the size of the model weights compared to the size of the 183 input image. The total number of parameters  $n_p$  used in this model is defined as follows:

$$n_p = h_s \cdot (4 \cdot n_c + 1), \tag{4}$$

where  $h_s$  is the hidden size (i.e. the number of neurons in the hidden layer) and  $n_c$  is the number of channels or states for each cell.

Reconstruction quality depends mainly on the parameters  $h_s$  and  $n_c$ , as will be shown in section 5.1. The main problem is that the number of steps necessary to reconstruct the image increases as the image size increases. Thus, backpropagation through time will be longer, implying more training time and a tendency for vanishing gradients. Also, larger images will require a larger number of parameters as the number of reconstruction patterns increases.

To address this problem, we propose the Grid Neural Cellular Automata (GNCA) training strategy for image compression. In this approach, we divide the input images into patches with size  $P_w \times P_h$ , where  $P_w$  and  $P_h$  are the width and height of each patch. For each patch, we train a different NCA to reconstruct that patch. The compressed representation is defined by  $R = \sum^n \text{size}(\theta_n)$ , where  $\theta_n$  are the parameters of the NCA model trained at patch n. Figure 2 shows the encoding stage to obtain the compressed representation. The decoding stage uses the trained model to rebuild each patch through inference. The merged patches reconstruct the entire image.

The GNCA approach addresses the scaling issue by using smaller models for each patch, which can be trained and reconstructed in parallel. This allows for efficient compression of large images while maintaining high reconstruction quality. With this, no matter the size of the original input image, a set of NCAs is able to learn to compress the representation of each image patch and then reconstruct it. The process of reconstructing an image starts by providing an empty image with a pixel seed to the trained model, as shown in Figure 1.

206 207

208 209

210

### 4 EXPERIMENTAL ENVIRONMENT

#### 4.1 DATASET

For our experiments, we used images from the COCO 2017 dataset (Lin et al., 2014) for the experiments. This dataset includes a diverse range of images of complex everyday scenes containing common objects in their natural context, with images of different sizes. For our experiments, we selected 30 random images from the dataset. We used a reduced number of images because each image requires a different model to train. We resized the original images to the sizes  $40 \times 40$ ,  $80 \times 80$  and  $120 \times 120$  pixels to our experiments.



Figure 2: Encoding and decoding stages from Grid Neural Cellular Automata (GNCA). In the encoding stage, the image is divided into patches, which are trained by different NCA models, represented as M. The model weights are the resulting compressed representation. In the decoding stage, we use inference of each model to reconstruct the image.

### 4.2 METRICS

To asses the quality of the compressed image, we used the metrics Peak Signal-to-Noise Ratio (PSNR), Mean Squared Error (MSE) and Structural Similarity Index Measure (SSIM), which are metrics used in literature for evaluating the quality of compressed images.

MSE is a measure of distance between the reconstructed image and the original image. The smaller the MSE, the more similar are the images. MSE is defined as follows: 

$$MSE = \frac{1}{MN} \sum_{i=1}^{M} \sum_{j=1}^{N} \left( I(i,j) - \hat{I}(i,j) \right)^2,$$
(5)

where I and  $\hat{I}$  are the original and reconstructed image, with dimensions  $M \times N$ . 

PSNR is calculated as the ratio of the maximum possible power of a signal to the power of corrupting noise that affects the fidelity of its representation. In the context of image compression, PSNR is often expressed in decibels (dB) and is calculated using the mean squared error (MSE) between the original and compressed images. Higher PSNR values indicate better image quality. The equation for Peak Signal-to-Noise Ratio (PSNR) is expressed as follows: 

$$PSNR = 10 \cdot \log_{10} \left( \frac{I_{Max}^2}{MSE} \right)$$
(6)

where  $I_{Max}$  is the maximum pixel value of the image. 

SSIM assesses the visual impact of three characteristics of an image: luminance, contrast, and structure. The metric compares local patterns of pixel intensities that have been normalized for luminance and contrast. This makes SSIM particularly effective at modelling perceived changes in structural information, which are more important to human visual perception than changes in luminance or contrast alone. The equation for SSIM is given by : 

$$SSIM(x,y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)}$$
(7)

where x and y are the uncompressed and compressed images, respectively,  $\mu_x$  and  $\mu_y$  are sample means,  $\sigma_x$  and  $\sigma_y$  the standard deviations and  $c_1$  and  $c_2$  the stabilization coefficients. This metric leverages a kernel to extract the structural information of images rather than focusing on single-pixel errors. This metric varies in the interval [0, 1], with higher values being better.

We also evaluate the compress rate (CR) for each model. For the GNCA, we compare the original image size with the number of parameters of the model, considering a quantisation of one byte per parameter. Viewing the model weights as the compressed representation, the compression rate is defined by  $CR = 1 - \frac{\text{Size}(\theta_I)}{\text{Size}(I)}$ , where *I* is the input image and  $\theta_I$  the model parameters to reconstruct  $\hat{I}$ . When using quantised models,  $\theta_I$  is defined by the number of parameters. According to Mordvintsev & Niklasson (2021), experimentation with quantisation of the NCA models to just one byte per parameter showed no evidence of degradation in quality.

280 281

297

306 307

308

4.3 IMPLEMENTATION

282 Our NCA model is based on the architecture proposed by Mordvintsev et al. (2020), and adapted 283 for Pytorch. All experiments were conducted using an NVIDIA GPU 3090 (16GB) and Intel Core 284 i5 processor. Each experiment consisted of training the model for 40,000 epochs, with a batch size 285 of 8, a learning rate of 2e-3, and 224 iterations per epoch. For every image, we trained a separate 286 model with learned parameters specific to that image. The patch size for all images was set to 287  $40 \times 40$  pixels. The original NCA supports both synchronous and asynchronous training, through 288 the parameter fire rate. For our experiments, we set the fire rate to 0 to train synchronously, as 289 this configuration leads to faster convergence. However, asynchronous training resulted in similar 290 performance but with increased training times.

We proposed three versions of GNCA, using low, medium and high amounts of parameters, to be competitive with the compressed image size of the compared methods. These versions vary the number of neurons in the hidden layer  $h_s$  and the number of channels  $n_c$ , producing the total number of parameters, as defined in Eq. 4. The parameters and corresponding model sizes are shown in Table 1.

Table 1: GNCA configuration used in the experiments.

Name	$n_c$	$h_s$	#Params
GNCA-small	14	20	1140
GNCA-medium	16	32	2080
GNCA-large	16	40	2600

5 Results

\_\_\_\_

### 5.1 ANALYSIS OF NCA

We evaluated how the parameters of the NCA model influence its ability to reconstruct images 310 and the total model size. Initially, we varied the hidden layer size and observed its effect on the 311 PSNR for image reconstruction after model training. Figures 3(a) and 3(b) show how PSNR and 312 the number of parameters increase as the size of the hidden layer increases, using a 40x40 pixel 313 image from the COCO dataset. Figures 3(c) and 3(d) illustrate how the number of channels affects 314 both the PSNR value and the number of parameters. By increasing the hidden layer and number of 315 channels, we observed that the model learns to reconstruct the image more closely to the original one. However, increasing model size to improve image quality reduces the compression rate as the 316 number of parameters increases. Thus, the tradeoff between image quality and compression rate 317 should be tuned by adjusting the number of model parameters. 318

We also analyzed how the reconstruction quality of NCA scales with image size. Figure 4 shows the
 PSNR metric with different image sizes when training NCAs while keeping the same parameters.
 To maintain high PSNR as image size increases, NCA needs more parameters. However, using a
 single model for high-dimensional images may be infeasible due to high computational resource
 requirements for large models and the potential gradient vanishing problem as the number of back-propagation iterations increases. Using the GNCA strategy allows us to maintain the same model



Figure 3: Analysis of (a) PSNR versus Size of Hidden Layer, (b) Number of Parameters versus Size of Hidden Layer, (c) PSNR versus Number of Channels and (d) Number of Parameters versus Number of Channels, when evaluating an NCA trained to reconstruct an image of size  $40 \times 40$  from COCO dataset. 



Figure 4: Analysis of PSRN metric versus the image size and the parameter hidden size. GNCA scales better as the image size increases compared with the standard NCA strategy.

size while preserving the PSNR value. The total compressed representation size is the sum of the parameters from the NCA models used for each patch.

We also evaluated the quality of reconstructed images using GNCA-small, GNCA-medium and GNCA-large. Figure 5 shows this comparison, for  $120 \times 120$  images from COCO dataset. We can notice there is a tradeoff between model size and image quality, but it show the capability of the NCA to reconstruct the image from a single pixel, being close to the original image, using a lightweight model. 

#### 5.2 **COMPARISSON WITH COMPRESSION METHODS**

We evaluated the compression capabilities of NCA compared to JPEG, JPEG-2000 and WebP, using maximum quality settings for each method. Table 2 presents the average results of each method for compressing  $40 \times 40$ ,  $80 \times 80$  and  $120 \times 120$  images. Results include average values of PSNR, MSE, SSIM and CR metrics. 

The results demonstrate that GNCA achieves competitive results in terms of compressed image quality over differ image resolutions. GNCA-medium showed a good balance between image quality and compression rate, for  $40 \times 40$  and  $80 \times 80$  resolution. While JPEG and WebP achieved higher



Figure 5: Visualization of reconstructed images using GNCA-small, GNCA-medium and GNCA-large, for  $120 \times 120$  image size.

Table 2: Average results of image compression for COCO dataset using GNCA, JPEG, JPEG200and WebP compression methods.

402	Image Size	Method	PSRN↑	MSE↓	SSIM↑	CR↑
403		JPEG	32.72±2.91	43.58±31.06	$0.97 {\pm} 0.01$	62.84±13.72
404	$40 \times 40$	JPEG-2000	$33.28 {\pm} 3.09$	$38.73 \pm 29.50$	$0.96 {\pm} 0.02$	$66.16 \pm 10.60$
405		WebP	$33.81 \pm 3.52$	$37.56 \pm 31.43$	$0.98{\pm}0.01$	$65.36{\pm}10.53$
406		GNCA-small	$28.06 \pm 3.14$	$129.65 \pm 92.61$	$0.90 {\pm} 0.05$	75.98±7.40
407		GNCA-medium	$32.58 {\pm} 3.14$	$46.20{\pm}34.81$	$0.96 {\pm} 0.02$	56.17±13.50
408		GNCA-large	$\textbf{34.66}{\pm}\textbf{3.16}$	$28.74{\pm}21.88$	$0.97{\pm}0.01$	$45.22{\pm}16.87$
409		JPEG	33.48±2.67	35.97±27.14	$0.94{\pm}0.95$	67.67±5.52
410	$80 \times 80$	JPEG-2000	$34.54{\pm}5.03$	$27.57 \pm 2.87$	$0.96{\pm}0.02$	$68.46{\pm}5.04$
411		WebP	$35.42 \pm 2.55$	$22.40{\pm}11.81$	$0.97 {\pm} 0.12$	$67.60 {\pm} 4.07$
412		GNCA-small	$30.34{\pm}2.72$	$71.64{\pm}40.64$	$0.89 {\pm} 0.04$	<b>68.85</b> ±5.24
413		GNCA-medium	$34.75 {\pm} 2.57$	$25.38{\pm}12.86$	$0.95 {\pm} 0.03$	$43.17 \pm 9.56$
414		GNCA-large	$\textbf{36.06}{\pm\textbf{2.37}}$	$18.24{\pm}8.60$	$0.96{\pm}0.03$	30.81±12.79
415		JPEG	33.81±2.74	$34.22{\pm}29.26$	0.98±0.01	69.55±3.87
416	$120 \times 120$	JPEG-2000	$36.23 {\pm} 2.66$	$18.30{\pm}10.63$	$0.97 {\pm} 0.01$	$65.95 {\pm} 4.02$
417		WebP	$34.87 \pm 3.23$	$29.29 \pm 29.50$	$0.98{\pm}0.01$	$66.64 {\pm} 2.50$
418		GNCA-small	$30.79 {\pm} 2.51$	$63.62 \pm 36.81$	$0.94{\pm}0.03$	$66.52 \pm 4.37$
419		GNCA-medium	$35.49 {\pm} 2.50$	$21.46 \pm 12.05$	$0.98{\pm}0.11$	$38.91 \pm 7.98$
420		GNCA-large	37.24±2.54	$14.50{\pm}8.49$	$0.98{\pm}0.01$	$23.64 {\pm} 9.98$

421 422

397

398 399

423 compression rates for the similar PSNR values, the GNCA models offer flexibility in enhancing
 424 image quality or compression rate (defined by the number of GNCA parameters) by tuning the
 425 hidden layer size, as shown in Figure 4.

In this study, we applied the Friedman test (Friedman, 1937) and the pairwise Mann-Whitney U
test (Sheskin, 2003) to evaluate whether there are statistically significant differences among the
performance of the different image compression methods evaluated. The Friedman test was initially
conducted to detect any overall differences across the methods. The results indicated a statistically
significant difference among the groups. To identify which specific methods differed from each
other, we performed post-hoc pairwise comparisons using the Mann-Whitney U test. The results of
these tests revealed significant differences between several pairs of methods, as shown in Table 3.

432 No significant differences were found between JPEG, JPEG-2000 and WEBP when compared to
 433 GNCA-medium int terms of PSNR metric. c.

434 435

Table 3: Comparison of Methods after using Pairwise Mann-Whitney U Test. The symbol  $\checkmark$  means there is a statistically significant difference (p-value < 0.05) between the methods for the evaluated metrics

120	metrics.							
439		Image Size	$40 \times$	40	$80 \times$	80	$120 \times$	120
440		Comparison	PSNR	CR	PSNR	CR	PSNR	CR
441		GNCA-small vs JPEG	$\checkmark$	$\checkmark$	$\checkmark$	Х	$\checkmark$	Х
442		GNCA-small vs JPEG-2000	$\checkmark$	Х	$\checkmark$	Х	$\checkmark$	Х
443		GNCA-small vs WebP	$\checkmark$	Х	$\checkmark$	Х	$\checkmark$	Х
444		GNCA-small vs GNCA-medium	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
445		GNCA-small vs GNCA-large	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
446		GNCA-medium vs JPEG	Х	Х	X	$\checkmark$	Х	$\checkmark$
447		GNCA-medium vs JPEG-2000	Х	Х	X	$\checkmark$	Х	$\checkmark$
448		GNCA-medium vs WebP	X	Х	X	$\checkmark$	Х	$\checkmark$
449		GNCA-medium vs GNCA-large	$\checkmark$	Х	X	Х	Х	$\checkmark$
450		GNCA-large vs JPEG	X	Х	X	$\checkmark$	$\checkmark$	$\checkmark$
151		GNCA-large vs JPEG-2000	Х	$\checkmark$	X	$\checkmark$	Х	$\checkmark$
451		GNCA-large vs WebP	X	$\checkmark$	X	$\checkmark$	Х	$\checkmark$
452								

In terms of encoding and decoding time, NCA requires model training to learn image compression for each image, making it suitable only for offline compression. The training time, using the setup described in section 4.3, is approximately 3.2 hours per patch, which can be trained in parallel. The inference time for reconstruction is 0.05 seconds.

6 DISCUSSION

459 460

458

453

461 Results demonstrate that NCAs can be used for data compression by treating the trained model 462 weights as the compressed representation. The architecture of the model plays a crucial role in determining the compression efficiency, especially in balancing the tradeoff between the number 463 of parameters and the quality of the reconstructed image. Our findings indicate that this approach 464 is more efficient for low-resolution images, but the compression rate tends to decrease as image 465 size increases, especially when compared to JPEG, JPEG-2000, and WebP. Despite this, the GNCA 466 method shows competitive performance against these classical methods, particularly in scenarios 467 that benefit from high parallelization. 468

Compared to SOTA neural network compression methods, the solution based on NCA proposes a
much lighter approach. Our largest model (GNCA-large) uses only 2,600 parameters, compared to
state-of-the-art deep learning-based compression methods, such as MLIC++ (Jiang & Wang, 2023),
which requires 83 million parameters. This lightweight characteristic makes GNCA an attractive
option for devices with limited computational resources, as it supports highly parallelized and asynchronous operations. However, due to the training time required for each image, NCAs are better
suited for offline compression rather than real-time applications.

Some new approaches have been proposed to train a single model for many images (Hernandez et al., 2021), but they require larger models. Otimizing NCA architecture in future work can lead to more efficient compression. In terms of reconstruction quality, NCAs can generate images with higher quality than JPEG and WebP, given larger models. Below, we discuss the benefits and drawbacks of using NCA for image compression.

- 481
- 482 6.1 ADVANTAGES

The main advantages of using NCAs for compression are:

485 *Noise Robustness*: NCA models are naturally robust to noise and are often used in denoising applications. This analysis has been done by the work of Menta et al. (2024), where NCAs successfully reconstructed images even when parts of the data were corrupted. The ability to reconstruct images
 from minimal information, such as a single seed pixel, further enhances this robustness

Parallelization: NCAs can easily be parallelized due to their localized operations. Each cell inter acts only with its immediate neighbors, allowing for asynchronous updates across different parts of the image. This makes NCAs particularly well-suited for large-scale or distributed computing environments where parallel processing is essential.

Low Computational Resources: Compared to deep learning-based methods, NCAs require fewer
 parameters, making them highly efficient in terms of memory and computational power. This efficiency enables NCAs to be deployed on devices with limited resources, such as mobile or embedded
 systems, which are often unable to support large and complex neural networks.

497 *Predictable Compression Size*: Another advantage of NCAs is the ability to predict the size of the
 498 compressed representation before the training or encoding process begins. Since the compressed
 499 size corresponds directly to the number of model parameters, it is possible to determine the storage
 500 requirements beforehand, offering a level of control in memory-constrained environments.

Furthermore, the compact representation through model weights can serve as an encrypted compression.

- 503
- 504 505

506

6.2 LIMITATIONS

507 Despite the advantages, there are some limitations to the current NCA-based approach: 508

Training Time: NCAs require a significant amount of training time for each individual image, which
 limits their applicability in real-time compression scenarios. This makes the current approach more
 suitable for offline compression rather than dynamic, real-time environments.

*Model-per-Image Limitation*: In this work, we used NCA based on Mordvintsev et al. (2020), which
trains one model per image. We chose this method because it provided higher-quality reconstruction
with a low number of parameters. However, strategies to train one model to many images (Hernandez et al., 2021) can be explored in future works.

- 516
- 517 518

### 7 CONCLUSION

519 520

This work investigates the potential of Neural Cellular Automata for image compression. NCAs
 present compelling properties compared to SOTA approaches, such as low computational cost, asyn chronous operation with local interactions, high parallelization, and robustness to noise. Our exper iments compared NCA-based compression to traditional methods such as JPEG, JPEG-2000, and
 WebP, showing that NCAs can provide competitive image quality with far fewer parameters. We
 also introduce a grid strategy called GNCA to train NCAs in higher-resolution images.

Experiments show that NCA can store a compressed image representation with a high reconstruction quality. In this approach, the model weights store the compressed representation, whose size is smaller than the input image. NCA demonstrates competitive results compared to JPEG, JPEG-2000 and WebP, and it is capable of compressing and reconstructing images while maintaining a high level of similarity with the original image.

The key benefits of the proposed method include its lightweight model architecture, high parallelization capability, and robustness to noise. These properties make NCAs well-suited for use in devices with limited computational resources or environments that require asynchronous processing.

Although the current approach has some limitations—such as the need for long training times and
 separate models for each image—there is potential for optimization and improvement. Future work
 may focus on reducing the model size for higher-resolution images and developing methods to train
 NCAs for multiple images without significantly increasing model complexity. With these advance ments, NCAs could emerge as a viable alternative to current state-of-the-art compression techniques,
 offering both efficiency and flexibility.

540	REFERENCES
541	

543

549

552

553

554

555

559

560

561

566

567

568

569

579

592

- 542 Andrew Adamatzky. *Game of life cellular automata*, volume 1. Springer, 2010.
- Luis B Almeida. Multilayer perceptrons. In *Handbook of Neural Computation*, pp. C1–2. CRC Press, 2020.
- Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence*, 44(5):2567–2581, 2020.
- Milton Friedman. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the american statistical association*, 32(200):675–701, 1937.
  - Jie Gui, Zhenan Sun, Yonggang Wen, Dacheng Tao, and Jieping Ye. A review on generative adversarial networks: Algorithms, theory, and applications. *IEEE transactions on knowledge and data engineering*, 35(4):3313–3332, 2021.
- Alejandro Hernandez, Armand Vilalta, and Francesc Moreno-Noguer. Neural cellular automata manifold. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10020–10028, 2021.
  - Alfons G Hoekstra, Jiri Kroc, and Peter MA Sloot. Simulating complex systems by cellular automata. Springer, 2010.
- Wei Jiang and Ronggang Wang. Mlic++: Linear complexity multi-reference entropy model ing for learned image compression. In *ICML 2023 Workshop Neural Compression: From In- formation Theory to Applications*, 2023. URL https://openreview.net/forum?id=
   hxIpcSoz2t.
  - Zewen Li, Fan Liu, Wenjie Yang, Shouheng Peng, and Jun Zhou. A survey of convolutional neural networks: analysis, applications, and prospects. *IEEE transactions on neural networks and learning systems*, 33(12):6999–7019, 2021.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr
   Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. *arXiv preprint arXiv:1405.0312*, 2014.
- 573
   574
   574
   575
   575
   576
   576
   576
   576
   576
   577
   576
   578
   578
   579
   579
   570
   570
   571
   572
   574
   574
   575
   576
   576
   576
   577
   578
   578
   579
   579
   570
   570
   570
   571
   572
   572
   574
   574
   575
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
   576
- 577 Michael W Marcellin, Michael J Gormish, Ali Bilgin, and Martin P Boliek. An overview of jpeg 578 2000. In *Proceedings DCC 2000. Data compression conference*, pp. 523–541. IEEE, 2000.
- Larry Medsker and Lakhmi C Jain. *Recurrent neural networks: design and applications*. CRC press, 1999.
- Andrea Menta, Alberto Archetti, and Matteo Matteucci. Latent neural cellular automata for resource-efficient image restoration. *arXiv preprint arXiv:2403.15525*, 2024.
- 585 Dipti Mishra, Satish Kumar Singh, and Rajat Kumar Singh. Deep architectures for image compression: a critical review. *Signal Processing*, 191:108346, 2022.
- Alexander Mordvintsev and Eyvind Niklasson.  $\mu$ nca: Texture generation with ultra-compact neural cellular automata. *arXiv preprint arXiv:2111.13545*, 2021.
- Alexander Mordvintsev, Ettore Randazzo, Eyvind Niklasson, and Michael Levin. Growing neural
   cellular automata. *Distill*, 5(2):e23, 2020.
- 593 Elias Najarro, Shyam Sudhakaran, Claire Glanois, and Sebastian Risi. Hypernca: Growing developmental networks with neural cellular automata. *arXiv preprint arXiv:2204.11674*, 2022.

594 595	Rasmus Berg Palm, Miguel González-Duque, Shyam Sudhakaran, and Sebastian Risi. Variational neural cellular automata. In 10th International Conference on Learning Representations, ICLR
596	2022, 2022.
597	
598	Kolin Paul, D Roy Choudnury, and P Pal Chaudnuri. Cellular automata based transform coding for
599	Calcutta India December 17 20, 1000 Proceedings 6, pp. 260, 273 Springer 1000
600	Calcula, India, December 17-20, 1999. Proceedings 0, pp. 209–275. Springer, 1999.
601	Haval T Sadeeq, Thamer H Hameed, Abdo S Abdi, and Ayman N Abdulfatah. Image compression
602 603	using neural networks: a review. International Journal of Online and Biomedical Engineering ( <i>iJOE</i> ), 17(14):135–153, 2021.
604	
605	Mark Sandler, Andrey Zhmoginov, Liangcheng Luo, Alexander Mordvintsev, Ettore Randazzo, et al. Image segmentation via cellular automata. <i>arXiv preprint arXiv:2008.04965</i> , 2020.
000	
607 608	David J Sheskin. <i>Handbook of parametric and nonparametric statistical procedures</i> . Chapman and hall/CRC, 2003.
609	Shahram Shirani. Data compression: The complete reference (by d. selemon; 2007)[heal; review]
610 611	<i>IEEE Signal Processing Magazine</i> , 25(2):147–149, 2008.
612	Zhanjun Si and Ke Shen. Research on the webp image format. In Advanced graphic communica-
613	tions, packaging technology and materials, pp. 271–277. Springer, 2016.
615	Shyam Sudhakaran, Djordje Grbic, Siyan Li, Adam Katona, Elias Najarro, Claire Glanois, and
610	Sebastian Risi. Growing 3d artefacts and functional machines with neural cellular automata. In
010	Artificial Life Conference Proceedings 33, volume 2021, pp. 108. MIT Press One Rogers Street,
017	Cambridge, MA 02142-1209, USA journals-info, 2021.
618	James D. Thompson, Some christians techniques for estimating the mean Journal of the American
619 620	Statistical Association, 63(321):113–122, 1968.
621	Alexandre Variengien Stefano Nichele, Tom Glover and Sidney Pontes-Filho. Towards self-
622 623	organized control: Using neural cellular automata to robustly control a cart-pole agent. arXiv preprint arXiv:2106.15240, 2021
624	proprim winter 2100.10210, 2021.
625 626	John Von Neumann. The general and logical theory of automata. In <i>Systems research for behavioral science</i> , pp. 97–107. Routledge, 2017.
627	Creasery K Walloss. The inex still nicture compression standard. Communications of the ACM 24
628 629	(4):30–44, 1991.
630	N Wulff and J A Hertz. Learning cellular automaton dynamics with neural networks. Advances in
631	Neural Information Processing Systems, 5, 1992.
632	Xiaosu Zhu, Jingkuan Song, Lianli Gao, Feng Zheng, and Heng Tao Shen. Unified multivariate
633	gaussian mixture for efficient neural image compression. In Proceedings of the IEEE/CVF Con-
634	ference on Computer Vision and Pattern Recognition, pp. 17612–17621, 2022.
635	
636	
637	
638	
639	
640	
641	
642	
643	
644	
645	
646	
647	