

NCA-Morph: Medical Image Registration with Neural Cellular Automata

Amin Ranem
amin.ranem@gris.informatik.tu-darmstadt.de
John Kalkhof
Anirban Mukhopadhyay

Technical University of Darmstadt,
Karolinenpl. 5,
64289 Darmstadt, Germany

Abstract

Medical image registration is a critical process that aligns various patient scans, facilitating tasks like diagnosis, surgical planning, and tracking. Traditional optimization-based methods are slow, prompting the use of Deep Learning (DL) techniques, such as VoxelMorph and Transformer-based strategies, for faster results. However, these DL methods often impose significant resource demands. In response to these challenges, we present **NCA-Morph**, an innovative approach that seamlessly blends DL with a bio-inspired communication and networking approach, enabled by Neural Cellular Automata (NCAs). NCA-Morph not only harnesses the power of DL for efficient image registration but also builds a network of local communications between cells and respective voxels over time, mimicking the interaction observed in living systems. In our extensive experiments, we subject NCA-Morph to evaluations across three distinct 3D registration tasks, encompassing Brain, Prostate and Hippocampus images from both healthy and diseased patients. The results showcase NCA-Morph's ability to achieve state-of-the-art performance. Notably, NCA-Morph distinguishes itself as a lightweight architecture with significantly fewer parameters; 60% and 99.7% less than VoxelMorph and TransMorph. This characteristic positions NCA-Morph as an ideal solution for resource-constrained medical applications, such as primary care settings and operating rooms.

1 Introduction

In healthcare, the process of image registration is a core technical problem to combine different images of one (or more) patients [16, 46] into a common spatio-temporal coordinate system [0]. Successful registration enables a wide range of applications from diagnosis to surgical planning and tracking [12, 18, 30, 43]. The registration process infers the deformation of a medical image by using a *displacement vector field* [36] or *deformation field* [33]. Traditional – non-learning – registration is a time-consuming process that requires optimizing the deformation field based on a set of parameters [0, 3, 42]. The introduction of Deep Learning (DL) methods like VoxelMorph- and Transformer-based strategies [0, 3, 37] allow learning the deformation field end-to-end without requiring a ground truth warp field. Though DL makes inference faster, the computational requirement is ever increasing. In scenarios like cluttered operating rooms or primary care disease progression analysis, where

deploying large computers is often impractical, the need for low-resource inference becomes crucial for democratizing medical Artificial Intelligence [14].




Figure 1: (a) Traditional optimization techniques like SyN focus on *local displacements of voxels* while optimizing a set of parameters during inference [6, 27]. (b) Traditional DL techniques leverage global information while estimating a *one-shot deformation field during inference* based on displacement rules learned during training [9, 10, 21]. (c) NCA-Morph learns *local displacement rules* in a DL-fashion while *regularizing the deformation field during inference* based on local interactions and learned rules.

In this paper, we introduce a bio-inspired paradigm for image registration as an emergent phenomenon, drawing inspiration from the dynamics of cellular communication and collective intelligence, Figure 1. Building on this, we explore the recent advances in Neural Cellular Automata (NCA) for lightweight and adaptive medical applications. NCA is an emerging field with promise for various applications where a common local rule is uniformly applied to all cells in an image or 3D volume. The particular appeal of NCAs lies in their ability to communicate solely with their nearby neighbors. A complex goal can be reached by repeating this same rule multiple times. The one-cell architecture makes NCAs lightweight with only around 13k parameters. To date, most publications in this area have been limited to toy examples, such as growing an image from a single cell [39], self-classifying robots [24], and segmentation [39] on image sizes of 64x64. This is partially due to the substantial video random access memory (VRAM) requirements during training and the more difficult convergence in general. However, recent developments like M3D-NCA [20] have expanded the potential use cases.

Given the success of NCAs in segmenting medical images [19], the question naturally arises as to whether they can be applied to another important area of medical image analysis: *registration*. However, this problem poses a greater challenge since the goal of image registration is generally to match two 3D volumes. Our NCA-based approach stands apart, as it does not involve variational minimization, elastic deformations or convolutional networks and spatial transformations like traditional and DL approaches do. It offers a novel, efficient pathway for achieving high-precision registration by leveraging the self-organizing properties of NCAs.

NCAs are initially conceived as a *solitary active cell* that can be seen as mimicking a single voxel taking the initiative to build a network of communication between its neighboring cells. These connections facilitate communication and collaboration among cells to collectively influence their transformations, i.e. a 3D deformation field for optimal registration. During the training of the NCA, the initial network of connections between the cells evolves

and matures into a resilient and complex communication network among the NCA cells. Remarkably, despite the growing complexity of the network, each individual cell adheres to the same rule, mirroring the uniformity of action within a developing organism.

We demonstrate how this bio-inspired approach is (1) **parameter efficient**, (2) holds the potential to significantly **enhance the precision and efficiency of medical image registration** while (3) **being light-weight** and able to run on a Raspberry Pi.

NCA-Morph establishes the basis for utilizing NCA architectures in the context of medical image registration. The contributions of this work are two-fold: (1) NCA-Morph successfully introduces the *NCA*s in *3D medical image registration* by learning local displacement rules. (2) We perform an extensive *analysis of the NCA architecture and its components* on three different organs (Brains, Prostate and Hippocampus) in the form of inter-subject setups to get a proper understanding of their impact on medical image registration. The results show that our method achieves accuracy comparable to the state-of-the-art. With the help of *NCA-Morph*, registration can be completed in a few seconds using a GPU while being a lightweight architecture with 60% and 99.7% less parameters compared to VoxelMorph and TransMorph respectively.

2 Related Work

Traditional medical image registration: There are several traditional 3D medical image registration techniques based on basic – non-learning – assumptions like the concept of *attraction* using a deformable model like Maxwell’s demons [1]. Another strategy uses large deformation metric mappings via geodesic flows [2], which is based on variational minimization on vector fields using Euler-Lagrange equations. Topology-preserving diffeomorphic transformations have demonstrated success in diverse studies related to computational anatomy. Approaches like elastic matching [3] or symmetric diffeomorphic registration [4] are based on local elastic deformations in combination with methods to maximize the cross-correlation for diffeomorphic cases in the space of topology preserving maps. Different work uses a Levenberg–Marquardt strategy and a constant Eulerian velocity framework [5] to perform registration. One of the well-known traditional registration strategies is standard symmetric normalization, SyN and SyNCC using Normalized Cross Correlation [6]. However, traditional methods like SyNCC are known for their slow processing speed.

(Deep-)Learning-based medical image registration: Recent research focuses on DL-based techniques for medical image registration. For instance, there are methods that are trained end-to-end using segmentations as their ground truth or optical flows like warp fields [7, 8, 9, 10]. Being dependent on segmentations or ground truth warp fields for training is a crucial bottleneck as to why another group of methods has emerged that focuses on unsupervised medical image registration.

The prime example of unsupervised medical image registration is the VoxelMorph framework [11]. VoxelMorph and several widely adopted methods [8, 12, 13, 14] rely on Convolutional Neural Network (CNN) structures along with a spatial transformation function [15]. This function is employed to deform images by applying an estimated deformation field denoted as ϕ .

Since the publication of VoxelMorph, advanced and more complex methods based on this framework were published. Notably, given the current success of Transformers in the field of medical imaging, recent publications have explored image registration using Transformer- or

Attention-based networks [9, 10, 11, 26, 57, 60]. Additionally, techniques like CycleMorph incorporate cycle loss and local cross-correlation during their training process [21].

In contrast, our NCA-Morph introduces a bio-inspired method for image registration, following the adaptability observed in living organisms. NCA-Morph enhances the precision and adaptability of medical image alignment by guiding the gradual evolution of the cells in the NCA network. This approach sets itself apart by achieving superior performance with notable parameter efficiency, offering a lightweight solution for precise medical image registration.

3 Methodology

To our knowledge, no work to date addresses the problem of 3D registration using NCAs.

Let $\mathcal{O} \subset \Omega$ be a dataset \mathcal{O} , while $\Omega \subset \mathbb{R}^3$ is a 3D spatial domain. \mathcal{O} consists of m pairs $\left\{ \left(I_{img}^k, I_{seg}^k \right) \right\}_{k \leq |\mathcal{O}|}$, where I_{img} is a 3D scan and I_{seg} the corresponding segmentation mask. For medical image registration, patient pairs are used, whereas one image pair k is referenced as **fixed (F)** and the other one as **moving (M)** – $\left\{ \left(F_{img}^k, F_{seg}^k \right) \mid \left(M_{img}^k, M_{seg}^k \right) \right\}$. Registration is a Transformation \mathcal{T} which maps an image pair from the *moving*, i.e. *source* domain Ω_M into the *target* domain Ω_F of the fixed image pair using a parameteric representation ϕ of the deformation field: $\mathcal{T}_\phi \left(M_{img}^k, M_{seg}^k \right) \cong \phi \circ \left[\left(M_{img}^k, M_{seg}^k \right) \in \Omega_M \right] = \left(M_{\phi(img)}^k, M_{\phi(seg)}^k \right) \in \Omega_F$.

Neural Cellular Automata for registration: Different NCA variations are introduced throughout this work based on their settings using the NCA-Morph_{kernel}^{steps} notation. When designing *NCA-Morph* for registration, we formulate a localized update rule represented as $g_\theta(F, M) = u$. This rule acts to incrementally adjust the cell state of each voxel, progressively aligning it with a deformation field ϕ . Through a series of iterations we gradually approach the desired deformation field by applying this rule locally to every voxel within a 3D image.

We define ϕ as the accumulated outcome of applying the update function g over a series of time steps. This can be expressed as $\phi = g(x_{t=n}) \odot g(x_{t=n-1}) \odot \dots \odot g(x_{t=0})$, where \odot represents a simple concatenation, t are the time steps and x represents the combined 3D volume of the moving image and the fixed image: $x_0 := F_{img}^k \odot M_{img}^k$. The objective is to identify anatomically similar regions in these images.

In Figure 2, we give an overview of our NCA-Morph architecture and the underlying inference setup. We initiate the model with the input images M_{img}^k and F_{img}^k . At each time step t , we execute the function g , which is parameterized by θ . This function guides updates for each cell in the input images, gradually refining the alignment, i.e. *propagation of the deformation field over time* following our bio-inspired communication approach among cells. We train these parameters by applying ϕ alongside a spatial transformation function, resulting in $M_{\phi(img)}^k$. This transformation enables the computation of the similarity loss $\mathcal{L}_{sim} \left(F_{img}^k, M_{\phi(img)}^k \right)$ between the *fixed* and *moved* images, facilitating the alignment process.

Additionally, we incorporate a smoothing technique $\mathcal{L}_{smooth}(\phi)$ for the deformation field, as proposed in VoxelMorph, enforcing a spatially smooth deformation. Since we possess segmentation masks, we include a segmentation loss term $\mathcal{L}_{seg} \left(F_{seg}^k, M_{\phi(seg)}^k \right)$, measuring the alignment between the fixed and moved segmentation masks.

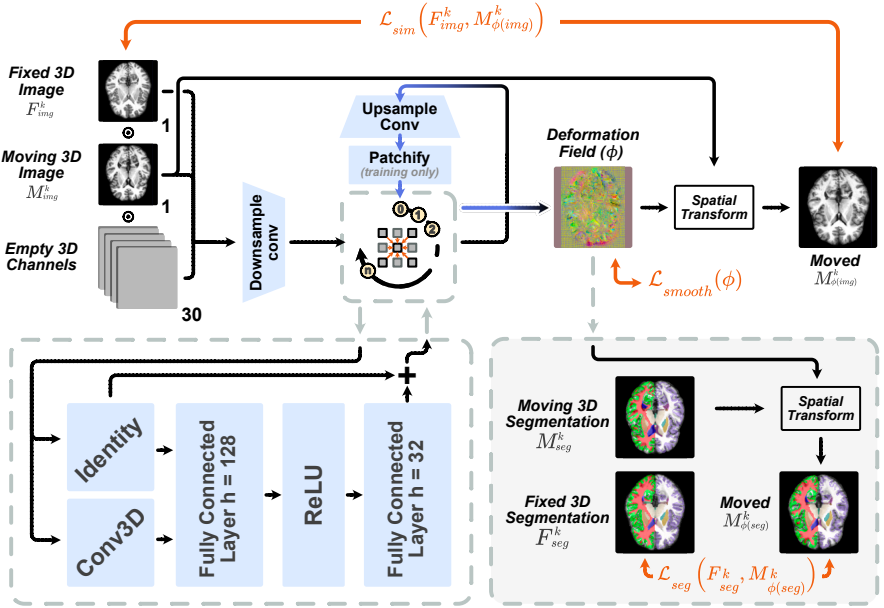


Figure 2: *NCA-Morph architecture for registration*: The input is a concatenation \odot of the fixed and moving image accompanied by 30 empty channels. Upon reducing the image size by a quarter, the initial NCA forecasts the overall flow of cells. Subsequently, this rough flow is combined with the detailed images, and a second NCA completes the deformation. While the second NCA solely observes patches during training, it operates at full resolution during inference. The resulting deformation field can be applied to warp both the moving image and its segmentation.

Model architecture: M3D-NCA [14], initially designed for medical segmentation, presents limitations when applied to registration tasks. This is because registration requires the alignment of two volumes, which demands a different architecture to handle the increased complexity. Our introduced NCA-Morph uses an n-level approach, where NCAs are running on different scales of the image, starting from a coarse to a fine worldview, increasing the global information. This global information is then upscaled to the next image scale and can thus be combined with high-resolution image information. NCA-Morph solves the high VRAM requirements by using patches on the high-resolution layers of the model. Due to the one-cell architecture, NCAs can be trained on patches and inferred on the full-scale image.

The update rule on each level g is represented by an NCA, which uses one 3D convolution with a kernel size of $k = 3$ as its input, describing a simplification of all neighborhood information into a vector v_n . We concatenate this information (\odot) with the cell's current state c , which leaves us with a state vector of $v = v_n \odot c$. In the next step, v is updated by two fully connected hidden layers, where the first one maps v onto a vector of size $h = 128$, the Rectified Linear Unit (ReLU) activation function is applied and the second layer maps it back to the input channel size i_c , resulting in a vector v_i . We then update c by adding v_i thus making $c_{t+1} = c_t + v_i$.

4 Datasets

Open Access Series of Imaging Studies (OASIS): The OASIS-1 dataset [52] is a cross-sectional collection of 416 subjects, aged from 18 to 96. For each subject, up to four individual T1-weighted MRI scans were obtained during a single session. The dataset was obtained within the *Learn2Reg* Challenge [45]. For this work, we randomly sampled 364 inter-patient image pairs from the training distribution and another 50 inter-patient pairs from the testing distribution.

Prostate task: Our prostate registration focuses on two distinct datasets: the I2CVB [24] and PROMISE12 dataset [27]. These datasets are used in their original form as part of the *Multi-site Dataset for Prostate MRI Segmentation Challenge* [28, 29, 30]. They consist of T2-weighted MRIs. For our registration preparation, we randomly selected 240 image pairs between different patients for training and an additional 60 pairs for evaluation.

Hippocampus task: The registration task for the Hippocampus involves two datasets: the *Harmonized Hippocampal Protocol* dataset [6] and the *Dryad* dataset [23]. The Hippocampus task comprises T1-weighted MRIs of elderly healthy individuals and patients with Alzheimer’s disease. We follow the same registration preparation as for prostate (240-60).

5 Results

We conduct a comprehensive analysis on computation vs. performance of our NCA-Morph variations in comparison to traditional methods like SyN [2] and state-of-the-art registration techniques, such as VoxelMorph [4], TransMorph [10], ViTVNet [9], and NICE-Trans [52].

5.1 Computation vs. Performance

To get a proper comparison over the computational burden of our proposed NCA-Morph variations and state-of-the-art (SOTA) methods, the Dice performance distribution across the number of model parameters is shown in Figure 3 and Table 1.




Figure 3: NCA-Morph compared to VoxelMorph, TransMorph, ViTVNet and NICE-Trans in terms of the number of parameters and Dice performance based on the OASIS registration task. NCA-Morph uses 60% and 99.7% fewer parameters than VoxelMorph and TransMorph.

Figure 3 shows that the NCA-Morph has a significant advantage in terms of the number of parameters. Our best NCA-Morph achieves a superior Dice performance than the SOTA VoxelMorph while using 60% less parameters.

Table 1 gives a proper overview of allocated resources during training and inference. Jacobian determinants are used as a measure for deformation fields as negative determinants indicate that the transformation causes local volume compression or folding. The runtime indicates the number of seconds to train a network for a single epoch along with the model size in megabyte (MB). The number of parameters describes the amount of trainable model parameters, whereas the inference time to register a single 3D volume on a GPU is indicated as "GPU sec". Although our NCA-Morph takes 3 seconds more for a single registration, it has a significantly lower number of parameters compared to VoxelMorph and TransMorph, while achieving still similar performances.

| Anatomy | Method | SSIM \uparrow | $ J_\theta \leq 0 \downarrow$ | Dice \uparrow [%] | Runtime \downarrow [sec] | Model size \downarrow [MB] | Nr. parameters \downarrow | GPU sec \downarrow |
|--------------------------------------|--------------------------------------|---------------------|--------------------------------|----------------------|----------------------------|------------------------------|-----------------------------|----------------------|
| Brain (OASIS) | SyN | 74.69 (1.26) | --- | 68.58 (1.56) | --- | --- | --- | 11.38 \times |
| | SyNCC | 78.93 (1.33) | --- | 72.00 (1.53) | --- | --- | --- | 448 \times |
| | VoxelMorph | 82.84 (1.36) | 254217 (20553) | 90.82 (1.03) | 722 | 1.30 | 327,331 | 7.79 |
| | TransMorph _{bin} | 81.20 (1.53) | 377100 (32541) | 90.49 (1.29) | 361 | 4.60 | 244,527 | 6.62 |
| | TransMorph _{default} | 86.21 (1.26) | 283837 (20883) | 93.80 (0.82) | 445 | 190.70 | 46,771,251 | 10.89 |
| | ViTVNet | 84.76 (1.38) | 310839 (25593) | 93.25 (0.88) | 551 | 126.3 | 31,560,079 | 8.98 |
| | NICE-Trans | 88.31 (1.29) | 305220 (29010) | 94.95 (0.62) | 654 | 24.90 | 5,706,508 | 6.59 |
| | NCA _{3\times3} | 81.00 (1.70) | 291620 (21842) | 90.99 (1.20) | 560 | 0.52 | 129,504 | 11.51 |
| | NCA _{5\times3} | 65.45 (2.36) | 323441 (25429) | 90.67 (1.22) | 567 | 0.52 | 130,641 | 9.90 |
| | NCA _{7\times7} | 82.70 (1.66) | 257241 (24826) | 92.14 (1.09) | 630 | 0.52 | 129,504 | 16.79 |
| NCA _{7\times7} | 82.05 (1.57) | 336983 (23095) | 91.29 (1.18) | 636 | 0.53 | 130,641 | 12.92 | |
| Prostate | SyN | 95.50 (1.09) | --- | 49.32 (21.02) | --- | --- | --- | 7.22 \times |
| | SyNCC | 95.59 (1.15) | --- | 51.88 (22.13) | --- | --- | --- | 420 \times |
| | VoxelMorph [⊗] | 95.24 (1.31) | 0.00 (0.00) | 53.39 (16.11) | 282 | 1.30 | 327,331 | 7.44 |
| | TransMorph _{bin} | 95.24 (1.31) | 0.00 (0.00) | 53.38 (16.16) | 294 | 4.60 | 244,527 | 6.49 |
| | TransMorph _{default} | 95.23 (1.22) | 96881 (94986) | 51.78 (15.82) | 367 | 190.70 | 46,771,251 | 6.68 |
| | ViTVNet | 94.98 (1.04) | 111511 (72308) | 27.61 (19.76) | 417 | 126.30 | 31,560,079 | 6.34 |
| | NICE-Trans [⊙] | 96.06 (0.91) | 2343699 (26870) | 0.00 (0.00) | 520 | 24.90 | 5,706,508 | 7.44 |
| | NCA _{3\times3} | 95.19 (1.31) | 53195 (7134) | 53.17 (15.80) | 428 | 0.52 | 129,504 | 9.43 |
| | NCA _{5\times3} | 95.12 (1.13) | 107357 (19479) | 35.53 (18.44) | 430 | 0.53 | 130,641 | 8.09 |
| | NCA _{7\times7} | 94.95 (1.30) | 333012 (29795) | 46.18 (15.54) | 477 | 0.52 | 129,504 | 13.30 |
| NCA _{7\times7} | 95.13 (1.10) | 83298 (14387) | 43.70 (15.46) | 479 | 0.53 | 130,641 | 10.25 | |
| Hippocampus | SyN | 83.49 (4.19) | --- | 35.26 (24.61) | --- | --- | --- | 0.35 \times |
| | SyNCC | 83.24 (4.91) | --- | 32.03 (27.12) | --- | --- | --- | 14.21 \times |
| | VoxelMorph [⊗] | 80.76 (3.46) | 6.36 (48.9) | 32.4 (22.75) | 33 | 1.30 | 327,331 | 0.68 |
| | TransMorph _{bin} | 79.67 (4.47) | 49379 (8187) | 33.64 (25.97) | 62 | 4.60 | 244,527 | 0.56 |
| | TransMorph _{default} | 84.70 (4.52) | 58104 (9120) | 48.28 (33.41) | 69 | 190.70 | 46,771,251 | 1.08 |
| | ViTVNet | 82.23 (3.83) | 29941 (7983) | 33.64 (26.71) | 66 | 126.30 | 31,560,079 | 0.93 |
| | NICE-Trans [⊙] | 67.82 (6.08) | 104669 (5135) | 4.80 (5.09) | 96 | 24.90 | 5,706,508 | 1.33 |
| | NCA _{3\times3} | 82.08 (3.92) | 29393 (8424) | 36.05 (27.88) | 45 | 0.52 | 129,504 | 0.53 |
| | NCA _{5\times3} | 80.77 (4.25) | 44966 (11952) | 35.46 (27.34) | 42 | 0.53 | 130,641 | 0.57 |
| | NCA _{7\times7} | 82.85 (3.81) | 28982 (8357) | 37.77 (28.54) | 57 | 0.52 | 129,504 | 0.46 |
| NCA _{7\times7} | 80.80 (4.34) | 51756 (14247) | 33.18 (27.39) | 55 | 0.53 | 130,641 | 0.46 | |

Table 1: *Computation vs. Performance*: Structural Similarity Index Measure (SSIM), negative Jacobian and Dice score for every trained DL model with runtime per epoch, model size, number of parameters and inference time on a GPU. For traditional methods, inference time on CPU is provided and marked with \times . NCA* generates the deformation field with an additional CNN layer. Due to non-convergence: \otimes is trained for 150 epochs, \odot for 100 epochs and \ominus performance dropped to 0% after few epochs of training.

The outcomes and resource allocations for our most successful NCA-Morph ablation – 7x7 kernel, 10 steps, 32 channels, 128 hidden size and directly extracting the flow from the NCA-Morph – in addition to VoxelMorph, two TransMorph ablations, the ViTVNet and NICE-Trans distinctly illustrate that NCA-Morph models use significantly fewer parameters than all other models. Compared to the default TransMorph model, our *NCA-Morph uses 99.7% less parameters while achieving nearly identical Dice performance.*

5.2 NCA-Morph ablation studies

We conducted a series of ablation experiments to investigate the impact of various hyperparameters on the performance of our NCA-Morph. Specifically, we explored the effects of altering the kernel size, the number of steps, channel size, and hidden size. Our objective was to understand which of these alterations lead to substantial changes in model performance to correctly determine a well suited setup for registration, Table 2.

| Method | SSIM \uparrow | $ J_\phi \leq 0 \downarrow$ | Dice \uparrow [%] |
|---|---------------------|------------------------------|---------------------|
| <i>Kernel size ablation (10 steps, 16 channels, 64 hidden size)</i> | | | |
| $NCA_{3 \times 3}^{10}$ | 78.82 (1.44) | 456804 (22491) | 88.99 (1.44) |
| $NCA_{5 \times 5}^{10}$ | 78.83 (1.48) | 460264 (22363) | 88.95 (1.41) |
| $NCA_{7 \times 7}^{10}$ | 78.93 (1.53) | 445223 (23424) | 88.91 (1.43) |
| $NCA_{9 \times 9}^{10}$ | 78.88 (1.49) | 444250 (22278) | 88.96 (1.43) |
| <i>Step size ablation (3×3 kernel, 16 channels, 64 hidden size)</i> | | | |
| $NCA_{3 \times 3}^5$ | 79.14(1.49) | 360049 (27372) | 89.13 (1.27) |
| $NCA_{3 \times 3}^{10}$ | 78.82 (1.44) | 456804 (22491) | 88.99 (1.44) |
| $NCA_{3 \times 3}^{30}$ | 78.47 (1.46) | 445687 (20268) | 88.96 (1.42) |
| $NCA_{3 \times 3}^{50}$ | 78.45 (1.47) | 436240 (20486) | 88.59 (1.53) |
| <i>Channel and hidden size ablation (3×3 kernel, 5 steps)</i> | | | |
| $NCA_{3 \times 3}^5(16, 64)$ | 79.14(1.49) | 360049 (27372) | 89.13 (1.27) |
| $NCA_{3 \times 3}^5(16, 128)$ | 78.64 (1.41) | 439126 (24152) | 88.22 (1.40) |
| $NCA_{3 \times 3}^5(32, 64)$ | 80.68 (1.64) | 315900 (30167) | 90.57 (1.21) |
| $NCA_{3 \times 3}^5(32, 128)$ | 81.00 (1.70) | 291620 (21842) | 90.99 (1.20) |

Table 2: Quantitative comparison of kernel size, number of steps, channel size and hidden dimensions based on SSIM, negative Jacobian determinants and Dice.

The results of our ablation studies shows that hyperparameter variations did not yield a significant effect on the model’s overall performance. The modifications made to the NCA-Morph kernel size, channel size, and hidden size seemed to have an impact that could be regarded as random, given the absence of a significant difference. It is important to mention, that by increasing the number of steps, the perceptive range is increased leading to worse registration performance as shown in Table 2. With this insight in mind, we proceeded to perform our experiments with the NCA-Morph using 3×3 kernel with 5 steps and a 7×7 kernel with 10 steps, both having 32 channels and a hidden size of 128.

5.3 NCA-Morph produces stable flows

To assess the variance level of our NCA-Morph, we conducted ten consecutive inference registrations on three random images from the OASIS registration task.

We computed the standard deviation of the ten resulting flow fields, offering a measure of variability in the spatial transformations generated during the registration process. Additionally, we assessed the standard deviation of the corresponding segmentation masks, shedding light on the uncertainty associated with our model’s performance. Figure 4 clearly indicates the reliability and consistency of our NCA-Morph in handling registration tasks.




Figure 4: Stability assessment over 10 predictions on different test samples from the OASIS registration task. Left: variance for the flow, right: variance for the corresponding segmentation masks – using $NCA-Morph_{7 \times 7}^{10}$.

5.4 NCA-Morph runs on a Raspberry Pi

Incorporating DL models into resource-limited platforms like the Raspberry Pi (RPI) has emerged as an important endeavor. We showcase the unique benefits of our lightweight NCA-Morph architecture by deploying it on a Raspberry Pi 4, Model B (2 GB RAM), and illustrate its potential for edge computing applications. We conducted experiments to evaluate the performance and efficiency of our architecture in this constrained environment.




Figure 5: Inference times of our $NCA-Morph_{7 \times 7}^{10}$ architecture across different input sizes on a Raspberry Pi.

NCA-Morph’s efficiency in resource utilization sets it apart as a particularly well-suited DL model for edge computing. Engineered with a focus on minimal memory and computational requirements, it seamlessly adapts to the RPi’s constrained hardware of 2GB. This efficiency ensures not only smooth operation but also the potential for additional software components, making NCA-Morph a valuable addition to the Raspberry Pi ecosystem. With its focus on resource efficiency, real-time capabilities, and versatility, our NCA-Morph is positioned to contribute to advancements in various domains, while ensuring accessibility to a wide range of users.

6 Conclusion

We introduce *NCA-Morph*, a robust medical image registration technique that uses an NCA architecture for registration instead of U-Net-based strategies. We evaluate our approach across three different registration tasks and show that it achieves state-of-the-art performance while being lightweight and using a significantly lower amount of parameters. Based on our bio-inspired approach, NCA-Morph successfully builds a complex network between its cells

while learning local displacement rules in a DL fashion. Future work should explore the NCA-related limitations of VRAM and training times to further improve the applicability of our method.

7 Reproducibility

The used dataset is publicly available under the corresponding citation. The code with all implementations is made publicly available under <https://github.com/MECLabTUDA/NCA-Morph>.

8 Acknowledgements

This work is partially supported by Norwegian Research Council project number 322600 Capsnetwork.

References

- [1] John Ashburner. A fast diffeomorphic image registration algorithm. *Neuroimage*, 38(1):95–113, 2007.
- [2] Brian B Avants, Charles L Epstein, Murray Grossman, and James C Gee. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis*, 12(1):26–41, 2008.
- [3] Ruzena Bajcsy and Stane Kovačič. Multiresolution elastic matching. *Computer vision, graphics, and image processing*, 46(1):1–21, 1989.
- [4] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019.
- [5] M Faisal Beg, Michael I Miller, Alain Trouvé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision*, 61:139–157, 2005.
- [6] Marina Boccardi, Martina Bocchetta, Félix C Morency, D Louis Collins, Masami Nishikawa, Rossana Ganzola, Michel J Grothe, Dominik Wolf, Alberto Redolfi, Michela Pievani, et al. Training labels for hippocampal segmentation based on the eadc-adni harmonized hippocampal protocol. *Alzheimer’s & Dementia*, 11(2):175–183, 2015.
- [7] Lisa Gottesfeld Brown. A survey of image registration techniques. *ACM computing surveys (CSUR)*, 24(4):325–376, 1992.
- [8] Xiaohuan Cao, Jianhua Yang, Jun Zhang, Dong Nie, Minjeong Kim, Qian Wang, and Dinggang Shen. Deformable image registration based on similarity-steered cnn regression. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part I 20*, pages 300–308. Springer, 2017.

- [9] Junyu Chen, Yufan He, Eric C Frey, Ye Li, and Yong Du. Vit-v-net: Vision transformer for unsupervised volumetric medical image registration. *arXiv preprint arXiv:2104.06468*, 2021.
- [10] Junyu Chen, Eric C Frey, Yufan He, William P Segars, Ye Li, and Yong Du. Transmorph: Transformer for unsupervised medical image registration. *Medical image analysis*, 82:102615, 2022.
- [11] Junyu Chen, Yihao Liu, Yufan He, and Yong Du. Deformable cross-attention transformer for medical image registration. *arXiv preprint arXiv:2303.06179*, 2023.
- [12] Bob D De Vos, Floris F Berendsen, Max A Viergever, Marius Staring, and Ivana Išgum. End-to-end unsupervised deformable image registration with a convolutional neural network. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings 3*, pages 204–212. Springer, 2017.
- [13] G Eggers, J Mühling, and R Marmulla. Image-to-patient registration techniques in head surgery. *International journal of oral and maxillofacial surgery*, 35(12):1081–1095, 2006.
- [14] Moritz Fuchs, Camila Gonzalez, Yannik Frisch, Paul Hahn, Philipp Matthies, Maximilian Gruening, Daniel Pinto Dos Santos, Thomas Dratsch, Moon Kim, Felix Nensa, et al. Closing the loop for ai-ready radiology. In *RöFo-Fortschritte auf dem Gebiet der Röntgenstrahlen und der bildgebenden Verfahren*. Georg Thieme Verlag KG, 2023.
- [15] Alessa Hering, Lasse Hansen, Tony CW Mok, Albert CS Chung, Hanna Siebert, Stephanie Häger, Annkristin Lange, Sven Kuckertz, Stefan Heldmann, Wei Shao, et al. Learn2reg: comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning. *IEEE Transactions on Medical Imaging*, 2022.
- [16] Derek LG Hill, Philipp G Batchelor, Mark Holden, and David J Hawkes. Medical image registration. *Physics in medicine & biology*, 46(3):R1, 2001.
- [17] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *Advances in neural information processing systems*, 28, 2015.
- [18] Songbai Ji, Xiaoyao Fan, David W Roberts, Alex Hartov, and Keith D Paulsen. Cortical surface shift estimation using stereovision and optical flow motion tracking via projection image registration. *Medical image analysis*, 18(7):1169–1183, 2014.
- [19] John Kalkhof and Anirban Mukhopadhyay. M3d-nca: Robust 3d segmentation with built-in quality control. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 169–178. Springer, 2023.
- [20] John Kalkhof, Camila González, and Anirban Mukhopadhyay. Med-nca: Robust and lightweight segmentation with neural cellular automata. *arXiv preprint arXiv:2302.03473*, 2023.
- [21] Boah Kim, Dong Hwan Kim, Seong Ho Park, Jieun Kim, June-Goo Lee, and Jong Chul Ye. Cyclemorph: cycle consistent unsupervised deformable image registration. *Medical image analysis*, 71:102036, 2021.

- [22] Julian Krebs, Tommaso Mansi, Hervé Delingette, Li Zhang, Florin C Ghesu, Shun Miao, Andreas K Maier, Nicholas Ayache, Rui Liao, and Ali Kamen. Robust non-rigid registration through agent-based action learning. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part I 20*, pages 344–352. Springer, 2017.
- [23] Jessie Kulaga-Yoskovitz, Boris C Bernhardt, Seok-Jun Hong, Tommaso Mansi, Kevin E Liang, Andre JW Van Der Kouwe, Jonathan Smallwood, Andrea Bernasconi, and Neda Bernasconi. Multi-contrast submillimetric 3 tesla hippocampal subfield segmentation protocol and dataset. *Scientific Data*, 2(1):1–9, 2015.
- [24] Guillaume Lemaître, Robert Martí, Jordi Freixenet, Joan C Vilanova, Paul M Walker, and Fabrice Meriaudeau. Computer-aided detection and diagnosis for prostate cancer based on mono and multi-parametric mri: a review. *Computers in biology and medicine*, 60:8–31, 2015.
- [25] Hongming Li and Yong Fan. Non-rigid image registration using fully convolutional networks with deep self-supervision. *arXiv preprint arXiv:1709.00799*, 2017.
- [26] Yong-xin Li, Hui Tang, Wei Wang, Xiu-feng Zhang, and Hang Qu. Dual attention network for unsupervised medical image registration based on voxelmorph. *Scientific Reports*, 12(1):16250, 2022.
- [27] Geert Litjens, Robert Toth, Wendy Van De Ven, Caroline Hoeks, Sjoerd Kerkstra, Bram van Ginneken, Graham Vincent, Gwenael Guillard, Neil Birbeck, Jindang Zhang, et al. Evaluation of prostate segmentation algorithms for mri: the promise12 challenge. *Medical image analysis*, 18(2):359–373, 2014.
- [28] Quande Liu. A Multi-site Dataset for Prostate MRI Segmentation. <https://liuquande.github.io/SAML/>.
- [29] Quande Liu, Qi Dou, and Pheng Ann Heng. Shape-aware meta-learning for generalizing prostate mri segmentation to unseen domains. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2020.
- [30] Quande Liu, Qi Dou, Lequan Yu, and Pheng Ann Heng. Ms-net: Multi-site network for improving prostate segmentation with heterogeneous mri data. *IEEE Transactions on Medical Imaging*, 2020.
- [31] JB Antoine Maintz and Max A Viergever. A survey of medical image registration. *Medical image analysis*, 2(1):1–36, 1998.
- [32] Daniel S Marcus, Tracy H Wang, Jamie Parker, John G Csernansky, John C Morris, and Randy L Buckner. Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *Journal of cognitive neuroscience*, 19(9):1498–1507, 2007.
- [33] Erik HW Meijering, Karel J Zuiderveld, and Max A Viergever. Image registration for digital subtraction angiography. *International Journal of Computer Vision*, 31:227–246, 1999.

- [34] Mingyuan Meng, Lei Bi, Michael Fulham, Dagan Feng, and Jinman Kim. Non-iterative coarse-to-fine transformer networks for joint affine and deformable image registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 750–760. Springer, 2023.
- [35] Alexander Mordvintsev, Ettore Randazzo, Eyvind Niklasson, and Michael Levin. Growing neural cellular automata. *Distill*, 5(2):e23, 2020.
- [36] Hans-Hellmut Nagel and Wilfried Enkelmann. An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 565–593, 1986.
- [37] Javier Pérez de Frutos, André Pedersen, Egidijus Pelanis, David Bouget, Shanmugapriya Survarachakan, Thomas Langø, Ole-Jakob Elle, and Frank Lindseth. Learning deep abdominal ct registration through adaptive loss weighting and synthetic data generation. *Plos one*, 18(2):e0282110, 2023.
- [38] Marc-Michel Rohé, Manasi Datar, Tobias Heimann, Maxime Sermesant, and Xavier Pennec. Svf-net: learning deformable image registration using shape matching. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part I 20*, pages 266–274. Springer, 2017.
- [39] Mark Sandler, Andrey Zhmoginov, Liangcheng Luo, Alexander Mordvintsev, Ettore Randazzo, et al. Image segmentation via cellular automata. *arXiv preprint arXiv:2008.04965*, 2020.
- [40] Jiacheng Shi, Yuting He, Youyong Kong, Jean-Louis Coatrieux, Huazhong Shu, Guanyu Yang, and Shuo Li. Xmorpher: Full transformer for deformable medical image registration via cross attention. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VI*, pages 217–226. Springer, 2022.
- [41] Hessam Sokooti, Bob De Vos, Floris Berendsen, Boudewijn PF Lelieveldt, Ivana Išgum, and Marius Staring. Nonrigid image registration using multi-scale 3d convolutional neural networks. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part I 20*, pages 232–239. Springer, 2017.
- [42] J-P Thirion. Image matching as a diffusion process: an analogy with maxwell’s demons. *Medical image analysis*, 2(3):243–260, 1998.
- [43] A Uneri, Y Otake, AS Wang, G Kleinszig, S Vogt, A Jay Khanna, and JH Siewerdsen. 3d–2d registration for surgical guidance: effect of projection view angles on registration accuracy. *Physics in Medicine & Biology*, 59(2):271, 2013.
- [44] Kathryn Walker, Rasmus Berg Palm, Rodrigo Moreno, Andres Faina, Kasper Stoy, and Sebastian Risi. Physical neural cellular automata for 2d shape classification. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12667–12673. IEEE, 2022.

- [45] Xiao Yang, Roland Kwitt, Martin Styner, and Marc Niethammer. Quicksilver: Fast predictive image registration—a deep learning approach. *NeuroImage*, 158:378–396, 2017.
- [46] Barbara Zitova and Jan Flusser. Image registration methods: a survey. *Image and vision computing*, 21(11):977–1000, 2003.