Hardware-Mappable Cellular Neural Networks for Distributed Wavefront Detection in Next-Generation Cardiac Implants

Zhuolin Yang, Lei Zhang, Kedar Aras, Igor R. Efimov, and Gina C. Adam*

Artificial intelligence algorithms are being adopted to analyze medical data, promising faster interpretation to support doctors' diagnostics. The next frontier is to bring these powerful algorithms to implantable medical devices. Herein, a closed-loop solution is proposed, where a cellular neural network is used to detect abnormal wavefronts and wavebrakes in cardiac signals recorded in human tissue is trained to achieve >96% accuracy, >92% precision, >99% specificity, and >93% sensitivity, when floating point precision weights are assumed. Unfortunately, the current hardware technologies for floating point precision are too bulky or energy intensive for compact standalone applications in medical implants. Emerging device technologies, such as memristors, can provide the compact and energy-efficient hardware fabric to support these efforts and can be reliably embedded with existing sensor and actuator platforms in implantable devices. A distributed design that considers the hardware limitations in terms of overhead and limited bit precision is also discussed. The proposed distributed solution can be easily adapted to other medical technologies that require compact and efficient computing, like wearable devices and lab-on-chip platforms.

1. Introduction

Machine learning (ML) algorithms are being adopted to analyze medical data in specialties like radiology, oncology, and cardiology, promising faster interpretation with accuracy close to doctors' diagnostics.[1] The next frontier in computing

Z. Yang, L. Zhang, G. C. Adam
Department of Electrical and Computer Engineering
The George Washington University
Washington, DC 20052, USA
E-mail: ginaadam@email.gwu.edu

K. Aras, I. R. Efimov
Department of Biomedical Engineering
The George Washington University
Washington, DC 20052, USA

The ORCID identification number(s) for the author(s) of this article can be found under https://doi.org/10.1002/aisy.202200032.

technology is to bring these powerful algorithms to implantable medical devices, which requires automation of real-time life-saving therapeutic decisions without the physician's presence. An example is the need for improved medical solutions for life-saving cardiac defibrillation therapies, that can detect bioelectric anomalies (e.g., cardiac arrhythmias) and act on this data locally for real-time therapy delivered within tens of seconds or minutes since the onset of life-threatening ventricular fibrillation (VF). The statistics put this challenging technological need in perspective: ventricular arrhythmias such as VF are responsible for over 700 000 sudden cardiac deaths a year in the USA and Europe.[2] VF is a common, life-threatening arrhythmia characterized by chaotic asynchronous electrical activity of the cardiac muscle, which results in death within 10 minutes.

Individual differences in physiological mechanisms, anatomic and genetic determinants, and etiologies of various arrhythmias impact the course of treatment. Ablation therapy, while promising, remains a work in progress. Therefore, on average, defibrillation therapy delivered by implantable cardioverter defibrillators (ICDs) remains the most effective treatment as antiarrhythmic drugs have limited efficacy and can be associated with adverse side effects. Implants have to be biocompatible, organ conformal, and small enough to minimize the tissue damage and be capable of independent autonomous operation without external intervention. Low power is an essential characteristic to avoid the heat damage to the tissue and prolong the lifetime of the embedded battery for many years without recharging.[3] Currently, most volume of the ICD has been occupied by batteries, which has limited the volume reduction and the computing capacity. ICD has local computing based on a microprocessor to detect and differentiate arrhythmia to offer different treatments, but the resolution provided by ICD is really low typically limited to only one or a couple of sensors; as such, the ability to detect arrhythmia wavefronts is non-existent. The data can be read wirelessly by the physician during periodic checkups. Increasing the sensing resolution is desired but the local computing capacity has to also be increased which is difficult due to power constraints. Wireless data transmission for processing of data outside of the body is not a viable solution either, as real-time data transfer between

**ADVANCED
SCIENCE NEWS**
www.advancedsciencenews.com

**ADVANCED
INTELLIGENT
SYSTEMS**
Open Access

www.advintellsyst.com

the implant and the external world requires a significant amount of power, even increasing the volume of the implant while introducing delays and additional security risks. That is why these systems have focused mostly on local real-time signal processing.

However, due to the low resolution for sensing and therapy, high-energy biphasic shocks are needed to effectively terminate life-threatening high-frequency arrhythmias. These high-energy shocks can lead to myocardial damage and associated comorbidities and it is a painful and traumatic experience for patients, especially when delivered inappropriately when arrhythmia is not present due to poor sensing. On the other hand, multipulse therapy (MPT) utilizes well-timed trains of low-energy electric pulses. Experiments on animal and human heart tissue showed that when appropriately timed, MPT significantly decreases the high-energy defibrillation threshold by almost an order of magnitude. Moreover, recent first-in-human clinical trial demonstrated safety and efficacy of MPT in patients with atrial fibrillation,[4] which is not possible with current high-energy ICDs due to pain and discomfort caused by high-energy shocks. However, as it is administered also through transvenous leads, the issue of low resolution remains.

To study the mechanisms of arrhythmias and develop suitable MPT for clinical use, high-definition electrically or optically mapped electrocardiograms (ECG) data must be used, which requires a large number of sensors to map the cardiac tissue surface. High-definition ventricular arrhythmia sensing integrated with electrotherapy is an emerging concept enabled by organ-conformal electronics platforms. Prototype organ-conformal electronic platforms have been developed with noncontact sensors and actuators and tested in vivo[5] but have limited resolution. Increasing the density of sensors and actuators is underway,[6] promising a personalized electrotherapy solution to terminate life-threatening tachycardias with two orders of magnitude less energy than a typical shock.[7] Such platforms could be used to predict fibrillatory wavefronts and enable their prevention using high-definition sensing and ultralow-energy electrotherapy that does not cause pain and discomfort.

The high definition is a critical requirement as multiple rotors can be simultaneously present in the myocardium[8] during an arrhythmia event and generate the seemingly chaotic pattern on the electrocardiogram that is the hallmark of atrial and ventricular fibrillation. The ventricular fibrillation rotors can be identified based on individual wavefronts, and wavebreaks are represented by phase singularities. The wavefront is defined as isolines of the phase that terminate either at boundaries or at singular points with the phase field (phase singularities[9]). Although the exact data resolution needed to extract these chaotic wavefronts is still under investigation, we estimated that >10 000 sensors, sampled at ≈500 Hz with 12-bit digitization, can produce an accurate map for the entire human heart. Such a system would produce >60 MB s^{-1} of data which must be processed in milliseconds, an insurmountable task for serial computation, especially on microprocessors of miniature implantable devices with limited energy resources. Real-time smart and energy-efficient computation is needed to process the data and trigger the local activation of actuators. To our knowledge, no organ conformal electronics platform has embedded computing for local data interpretation and millisecond decision-making,

as needed for real-time life-saving therapy such as arrhythmia electrotherapy.

In this work, we propose the use of distributed computing neural network algorithms which are hardware mappable, to provide high classification sensitivity, specificity, accuracy, and precision in determining the challenging spatiotemporal dynamics of cardiac electrical signals. Artificial neural networks can process a large amount of data in a parallel fashion and "learn" its patterns. As their name suggests, artificial neural networks are inspired by biological brain and can provide intelligent computing solutions. Deep learning techniques, such as convolutional neural networks, have been demonstrated to perform with >93% accuracy for the classification of ECG heartbeats.[10–12] These complex networks can be used for classification of heartbeat by heartbeat of data obtained from bedside ECG recording equipment, but they have yet to be applied in current low-resolution ICDs that shock the entire heart due to computational complexities and limited microprocessor capabilities.[13] However, for new types of high-definition organ-conformal platforms, they are impractical to physically realize due to their complexity for a large number of recording channels and also unsuitably centralized for the spatiotemporal tracking of wavefronts and wavebreaks as needed for precise therapy by distributed electric field. To our knowledge, no neural network algorithm has been proposed for the identification of wavefronts.

This work describes a distributed computing algorithm based on cellular neural networks that is readily mappable to memristor-based hardware circuitry and could enable a closed-loop solution
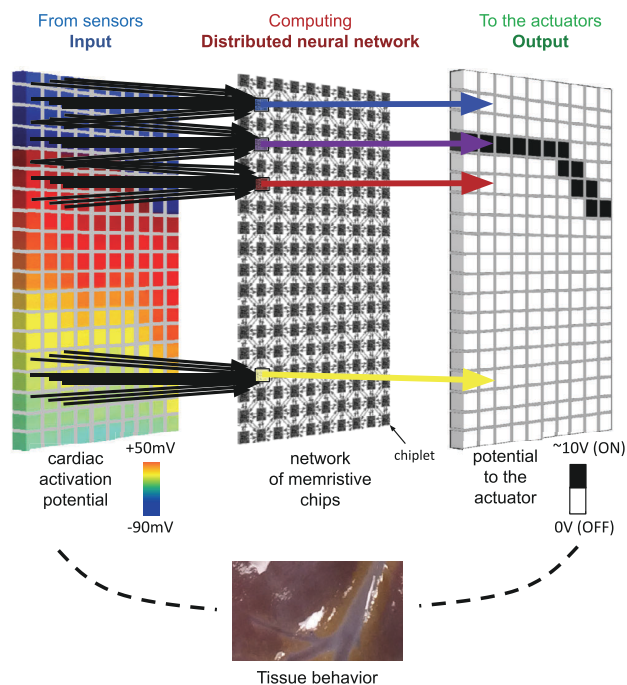
Figure 1. Distributed computing for electrical wavefront determination: Proposed technology using integrated network of sensors, computing chiplets distributed in a cellular neural network architecture, and actuators that will allow high-definition mapping, interpretation, and therapeutic response in a closed-loop fashion.

**ADVANCED
SCIENCE NEWS**
www.advancedsciencenews.com

**ADVANCED
INTELLIGENT
SYSTEMS**
Open Access
www.advintellsyst.com

that includes spatially distributed sensing, data processing, and any required actuation for therapy (**Figure 1**). The cellular neural network maps well to a spatiotemporally distributed architecture and would enable a high-speed high-data-throughput computing solution. Any other type of neural network, for example, a multi-layer perceptron or a convolutional neural network, would require hardware implementation in a single chip which would have to be connected to a multitude of sensors and actuators, with density limitations due to the interconnects. This proposed cellular neural network architecture was chosen as most suitable because it takes advantage of its natural tiled organization to easily map it to a distributed network of identical computing chiplets, as shown in Figure 1. We consider a chiplet to be a small integrated circuit (IC) of submillimeter dimensions that contains a well-defined subset of functionality and is designed to be combined with other chiplets in the organ-conformal platform. Each chiplet implements a cell unit of the cellular neural network, processing only local sensor information from itself and its neighbors and providing an output only to its local actuator.

The size, area, and power constraints are particularly important for this application. Emerging computing technologies, like memristor crossbars,[14] have significant potential in the More-than-Moore era, promising orders of magnitude better energy efficiency and compact implementation[15,16] of use in novel computing systems for implantable devices. A memristor commonly uses metal/insulator/metal sandwich structures, which include two layers of electrodes and an intermediate layer of memristive functional material, which is called the insulator.[17,18] Memristor devices can be fabricated as small as 2 nm, and[19] their resistance transition characteristics are closely associated with their electrodes and the switching materials. The device needs "forming" to create filamentary path(s) in the insulator and then reversibly set and reset to program the device to a desired conductance state between low (OFF) and high (ON) states. Thanks to its ionic transport, the programmed state is retained without static power dissipation. Memristor devices can be integrated with complementary metal–oxide–semiconductor (CMOS) control circuitry as dense matrices (crossbars) of artificial synapses to implement vector matrix multiplication using Ohm's law,[14,20–22] which is a fundamental operation in neural networks. This behavior enables a natural solution for the implementation of templates for the proposed cellular neural network computing, to be integrated directly with sensors and actuators. This approach allows for flexibility, requiring the design of only one chiplet and its tape-out in as many samples as needed for the size of the network at hand. The proposed solution can be used to develop the next-generation implantable devices that can provide low-energy therapy, thanks to high-resolution sensing, local computing, and precise actuation.

The remainder of the article is organized as follows. Section 2 describes the methodological details, the data obtained from human cardiac tissue, as well as the algorithm and the performance metrics used. Section 3 introduces the evaluation of the proposed methodology on the dataset, considering the optimization of hyperparameters such as the learning rate, weight initialization, binarization, as well as the impact of noise and quantization in the input and templates on the inference results. Section 4 concludes with a discussion of the algorithmic results and their potential mapping to a memristor-based hardware implementation.

2. Experimental Section

2.1. Data Gathering and Preprocessing

This study utilized representative data obtained from a deidentified donor human heart from the Washington Regional Transplant Community (Church Falls, VA). The study was approved by the Institutional Review Board at the George Washington University.
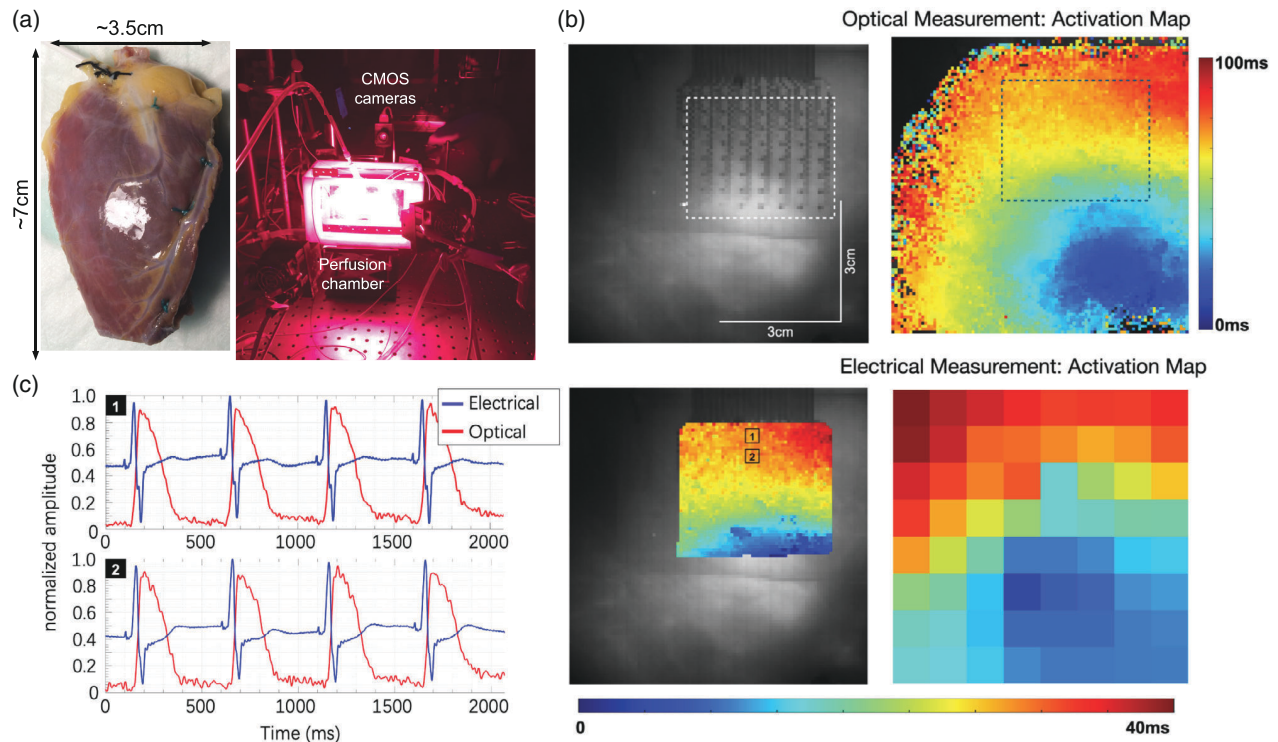
The experimental apparatus and procedures are explained in detail in the study by Aras et al.[23] Briefly, the ventricular tissue was prepared as a wedge with average dimension of 7 cm × 3.5 cm (**Figure 2**a). The tissue was then mounted in a temperature-controlled, pressure-controlled, and an oxygenated optical mapping setup (Figure 2a). Optical action potentials were mapped from ≈7 cm × 7 cm field of view using a MiCAM05 (SciMedia, CA) CMOS camera (100 × 100 pixels) and sampled at 1 KHz sampling rate.

The dataset consisted of 1000 optical mapping images of the epicardium tissue recorded at 1 kHz sampling rate with a size of 100 × 50 pixels. 800 images were used for training and 200 for testing. The dataset included complete recordings of several fibrillation events, enabling the analysis of various wavefront patterns during fibrillation as part of this work. Optical recordings were used because they provide higher-resolution mapping than flexible electrode arrays. However, these results were directly applicable to electrically recorded data, as shown in Figure 2b,c.[24] Studies into the resolution required to extract any possible chaotic rotors in human tissue are still under investigation and higher-resolution setups are being developed.

Analysis in the phase domain was typically done for such studies, as the wavefront propagation and the singularities could be easily detected in the phase domain. The time domain optical raw data recorded by the cameras was preprocessed to transform it into the phase domain with a scale between $-\pi$ and π through the Hilbert transform.[25] The Hilbert transform is an efficient signal analysis method for nonstationary time series, especially in determining the instantaneous frequency of time-varying signals, such as ventricular arrhythmias. Detection of these subtle frequency changes and potentially recognizing the initiation and/or termination of VT/VF is very important in understanding the mechanisms of arrhythmia. Given a real-time function $x(t)$, its Hilbert transform was defined as[26]

$$\widehat{x}(t) = H[x(t)] = x(t) * \frac{1}{\pi t} = \frac{1}{\pi} \int_{-\infty}^{+\infty} \frac{x(\tau)}{t - \tau} \mathrm{d}\tau \tag{1}$$

Figure 3a shows a raw optical signal and Figure 3b shows its phase-domain equivalent that was further preprocessed before looking at the wavefront. More details are presented in prior work.[23] A wavefront was located at the edge of phase $\varnothing(t) = \pi$ (red) and phase $\varnothing(t) = -\pi$ (blue) on the blue side. The wavefronts were labeled manually because the noise and the undesirable artifacts of the pacing electrode might affect the precision of the labels and affect the training results afterwards. For each data sample, a corresponding phase map 3b and its wavefront mapping 3c served as input and desired output, respectively, for the neural network core.

**ADVANCED
SCIENCE NEWS**
www.advancedsciencenews.com

**ADVANCED
INTELLIGENT
SYSTEMS**
Open Access
www.advintellsyst.com

Figure 2. Data gathering. a) Human left ventricular tissue wedge and experimental setup. b) Simultaneous optical and electrical cardiac mapping. c) Corresponding representative electrical and optical signals. Figure 2b,c are reproduced with permission.[24] Copyright 2022, American Heart Association.
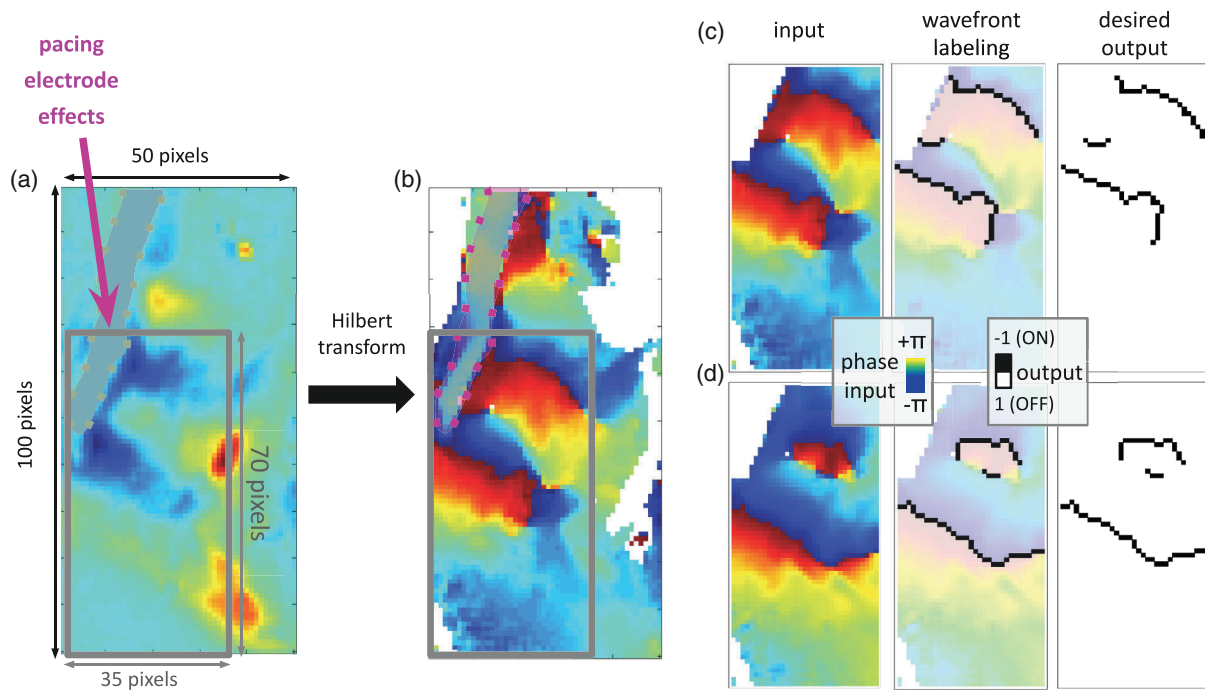
Figure 3. Data preprocessing. a) Example of raw optical phase map (100 × 50 pixels) recorded during VF in the human heart preparation showing the influence of the pacing electrode on the obtained signal. b) Example of Hilbert-transformed optical phase map. A subset (70 × 35 pixels) was selected to avoid network confusion due to pacing electrode effects. c) Example of input data used for training and its labeling. d) Example of input data used for testing and its labeling.

**ADVANCED
SCIENCE NEWS**
www.advancedsciencenews.com

**ADVANCED
INTELLIGENT
SYSTEMS**
Open Access

www.advintellsyst.com

Due to the unavoidable interrupts during the hour-long experiments and the underlying condition of the available human heart tissue, noise was an inevitable occurrence in the dataset. Noise is regarded as the irregular small section of pixels rapidly changing in the range from π to $-\pi$, as well as the value of pixel remaining constant throughout the measurement. The pacing electrode could also introduce significant artifacts due to its large size, needed to provide mechanical robustness during insertion into the rather stiff human cardiac muscle tissue. To avoid these unwanted effects, the data was cropped to 70×35 pixels and the pixels containing the pacing electrode were removed, as shown in 3b vs. 3c.

2.2. Cellular Neural Networks

Given the tight requirements for high speed and low-power hardware, the cellular neural network is a promising topology for distributed computing, based on a fixed number of interconnected processing units called "cells." Each unit, for example, unit ij at row i and column j, could be implemented by a computing chiplet, processing only local information from itself and its neighbors, with small size and energy requirements. The inputs $u_{ij}(t)$ at time t were fed into the network and outputs $y_{ij}(t)$ were obtained. The output of a processing cell ij was determined by the state of the cell $x_{ij}(t)$ according to Equation (2).

$$y_{ij}(t) = (|x_{ij}(t) + 1| - |x_{ij}(t) - 1|) \tag{2}$$

The state of cell ij at time t was calculated using the differential equation (3) taking into consideration all the cells in the neighborhood of size $M \times N$. This work included only the nearest neighbors (neighborhood size $= 3 \times 3$) to keep the results mappable to a potential compact hardware implementation. However, the neighborhood could include further away neighbors, for example, a neighborhood of size 7×7 included one central cell and 48 neighbors.

$$\frac{dx_{ij}(t)}{dt} = -x_{ij}(t) + \sum_{\substack{1 \le i \le M \\ 1 \le j \le N}} a_{mn} y_{mn}(t) + \sum_{\substack{1 \le i \le M \\ 1 \le j \le N}} b_{mn} u_{mn}(t) + I \tag{3}$$

In Equation (3), the inputs u_{mn} and outputs y_{mn} of its cell and neighboring cells were weighted via the matrix elements a_{mn} and b_{mn} of two matrices A and B of size M and N. The matrix A linked the outputs y_{mn} to the state x via its elements a_{mn}, while template B similarly linked the inputs u_{mn} to the state x, respectively. These matrices were called templates and were used repeatedly for each cell. Training the network means determining the values of templates A and B and of bias I.

Several algorithms were used for training these networks, including, random weights change,[27] Kalman filters,[28] genetic algorithms,[29] and backpropagation.[30] The random weight change[27] is a hardware-friendly algorithm for on-chip training on a wide range of tasks, but it involves large number of training epochs to obtain accurate templates. Kalman filters have been used to obtain accurate output from the inaccurate input information, minimizing the mean of squared error by estimating the

inner states of any dynamic process.[30] Genetic algorithms have been shown to train the network with desirable accuracy and robustness, but the evaluation of the fitness functions is computationally very expensive.[29]

We have defined a training algorithm based on backpropagation and batch updates robust to template nonidealities. Following initialization, the network will calculate the corresponding error for each image in the batch. The templates A, B, and bias I will be updated after each batch calculation. The process was repeated for all images in the training dataset to minimize the error between the obtained wavefront map output and the desired output. The network took several epochs to converge and several performance metrics, as shown in the next section, could be used to track the convergence.

For the case of the typical adapted stochastic gradient descent backpropagation training algorithm, the error was calculated based on

$$e_{ij}[k] = \frac{1}{2}(d_{ij} - y_{ij}^*[k]) \tag{4}$$

where $y_{ij}^*[k]$ is the output calculated by the algorithm at iteration k and d_{ij} is the desired cell output according to the image label. The templates A, B, and bias I are updated based on

$$a_{mn}[k+1] = a_{mn}[k] + \eta \Delta a_{mn}[k] \tag{5}$$

$$b_{mn}[k+1] = b_{mn}[k] + \eta \Delta b_{mn}[k] \tag{6}$$

$$I[k+1] = I[k] + \eta \Delta I[k] \tag{7}$$

with the updates of Δ weights

$$\Delta a_{mn}[k] = \frac{1}{MN} \sum_{\substack{1 \le i \le M \\ 1 \le j \le N}} e_{ij}[k] y_{i+m-2,j+n-2}^*[k] \tag{8}$$

$$\Delta b_{mn}[k] = \frac{1}{MN} \sum_{\substack{1 \le i \le M \\ 1 \le j \le N}} e_{ij}[k] u_{i+m-2,j+n-2}[k] \tag{9}$$

$$\Delta I[k] = \frac{1}{MN} \sum_{\substack{1 \le i \le M \\ 1 \le j \le N}} e_{ij}[k] \tag{10}$$

where m and n are the row and column indices, respectively, of the templates A and B. η is the learning rate, typically a small number always >0, that defines the range of weight updates in each iteration. As seen in Equation (8), the update $\Delta a_{mn}[k]$ for the feedback template A was calculated via the weighted sum of the error and the desired output for each cell. A similar update $\Delta b_{mn}[k]$ was calculated for control template B based on the error and the respective input. The bias I was also updated accordingly based on the average error of each cell to increase the performance of the network.

To improve the wall-clock time, we used batch training as defined by

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.

fastcase®
Smarter legal research.