

RESEARCH ARTICLE

Ultrasound Artifacts Imaging Recognition by Hebbian-Trained Hopfield Networks

Hatim Oulad Arifi ^{1*}, Xiao-Zhi Gao ²

¹School of Computing , University of Eastern Finland, Joensuu, Finland.

²School of Computing, University of Eastern Finland, Kuopio, Finland.

*Corresponding author E-mail: hatimo@uef.fi

Article Info.	Abstract
<i>Article history:</i> Received: 05/11/2025 Accepted: 23/01/2026 Published: 01/04/2026	Ultrasound imaging is a cornerstone of non-invasive diagnostics due to its real-time capability, safety, and low cost. However, image quality is often degraded by lingering vibrations of piezoelectric transducers after excitation, leading to artifacts. Hopfield Neural Networks (HNNs) have emerged as powerful tools for associative memory and pattern recognition, but their performance can be limited by dynamical instabilities, convergence to spurious states, and interference among overlapping patterns. This study investigates the principles and practical significance of HNNs in memory-based computational modeling. We examine limitations such as false minimum, memory capacity constraints, and instability, which can compromise accurate associative recall. Hebbian learning is applied to construct weight matrices that optimize memory storage, enhance fidelity of stored patterns, and reduce spurious attractors. Results show that properly tuned Hebbian-trained HNNs improve pattern stability and retrieval accuracy, especially in high-noise or memory-dense scenarios. enhancing structural consistency in both imaging and pattern analysis. Integrating Hebbian-trained HNNs with morphological filtering demonstrates superior performance in pattern recognition, retrieval reliability, and robustness under noisy or complex conditions. These findings highlight the potential of combining biologically inspired neural networks with image-processing techniques to enhance ultrasound image quality and computational modeling efficiency.

Keywords: Ultrasound imaging artifacts ; Diagnostic accuracy ; Hopfield Neural Networks (HNNs) ; Associative memory ;Hebbian learning.

2026 Center of Science.

Introduction

Sound exists across a wide range of frequencies, but human ears can detect only a portion of it known as audible sound typically ranging from 20 Hz to 20,000 Hz [3]. Sounds below this range are called infrasound, and though we cannot hear them, they can travel long distances and are often produced by natural phenomena like earthquakes and volcanic eruptions [4]. On the other end of the spectrum, ultrasound refers to sounds above 20,000 Hz, commonly used in medical imaging and industrial applications [5]. Together, infrasound, audible sound, and ultrasound represent the full spectrum of acoustic waves, each with unique characteristics and practical uses. In ultrasound imaging systems, piezoelectric crystals serve as both transmitters and receivers of acoustic waves [6]. When an acoustic wave propagates through a piezoelectric medium layered with materials of different acoustic impedance, multiple reflections and transmissions occur, producing constructive and destructive interference patterns [7]. This phenomenon is analogous to optical interference in thin films and is governed by phase shifts induced by differences in path lengths and propagation speeds [8]. The goal of this study is to investigate Hopfield Neural Networks (HNNs) to detect ring-down artifacts, which appear as repetitive high-intensity vertical streaks that may interfere with diagnostic interpretation [9].

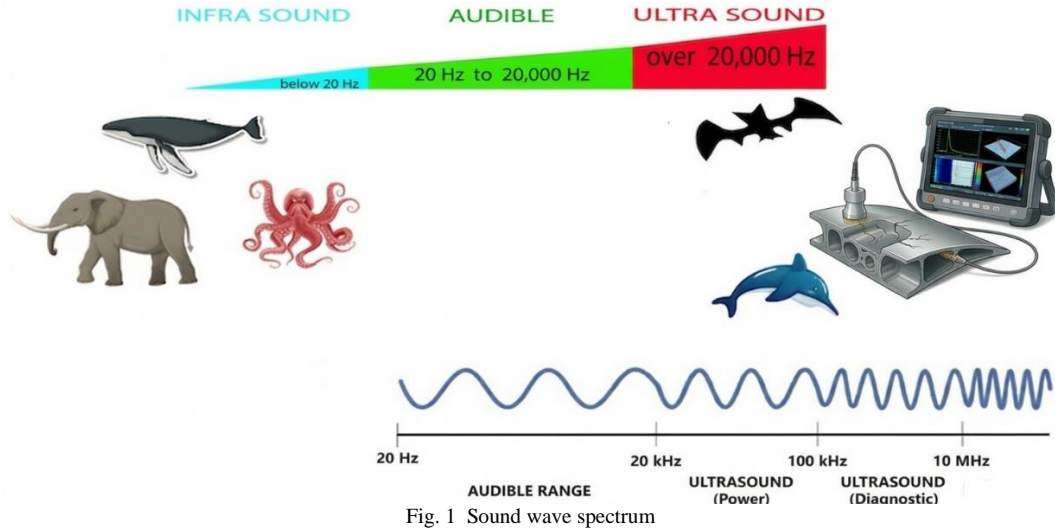



Fig. 1 Sound wave spectrum

A raster ultrasound image is represented by a matrix. Because grayscale images only contain bright information, they require less memory than color images [10]. Since differences between very similar shades of gray are often barely noticeable to the human eye, pixel intensity histograms are often used to represent brightness distributions in medical images [11]. In grayscale ultrasound imaging, pixel intensities range from 0 (black) to 255 (white). Anechoic areas (intensities near 0-10) appear black, indicating absence of echoes, as seen in fluid-filled structures (Hoskins2010). Hypoechoic regions (30-80) appear dark gray and produce fewer echoes than surrounding tissues, typical of soft tissues or solid organs. Isoechoic tissues (90-140) appear medium gray, resembling adjacent tissues and making boundaries harder to distinguish. Hyperechoic structures (180-255) appear light gray to white, reflecting strong echoes characteristic of bone, fat, calcifications, or fibrous tissues [12] [5]. A raster ultrasound image in this study is represented by a matrix of pixel intensities with a shape of 225 rows and 400 columns.



$$M_{i,j} = \begin{bmatrix} 33 & 36 & 37 & \dots & 17 & 18 & 13 \\ 18 & 18 & 21 & \dots & 2 & 2 & 2 \\ 8 & 9 & 14 & \dots & 36 & 41 & 43 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 10 & 12 & 15 & \dots & 22 & 26 & 23 \\ 16 & 16 & 16 & \dots & 20 & 20 & 18 \\ 21 & 24 & 26 & \dots & 29 & 21 & 21 \end{bmatrix}$$




Fig. 2 Ultrasound Image Matrix Pixel Classification

1. Methodology and Material

The histogram acts as a compact descriptor of the global intensity distribution of the ultrasound image, and it provides a balanced and informative representation of the pixel intensity distribution [13]. While these ranges are approximate and dependent on imaging settings such as gain and dynamic range, they naturally provide discrete intensity states that can be mapped to the state space of a Hopfield neural network. In this context, the discrete Hopfield model serves as a classifier where pixel intensities correspond to neuronal activations, and classification

emerges through convergence toward stable states of the network [14] [15]. The connections are resistive, and the connection strength over them is represented as w_{ij} [14]. Since negative resistors do not exist physically, excitatory connections use positive outputs, while inhibitory connections use inverted outputs [15]. Connections are excitatory if the output of a processing element is the same as its input, and inhibitory if the inputs differ from the output of the processing element [17]. A connection between the processing elements i and j is associated with a connection strength w_{ij} . This weight is positive if units i and j are both on, whereas a negative weight represents the situation of unit i being on and j being off. The weights are symmetrical, i.e., $w_{ij}=w_{ji}$ [1].

We obtained the Histogram of pixel intensities Matrix:

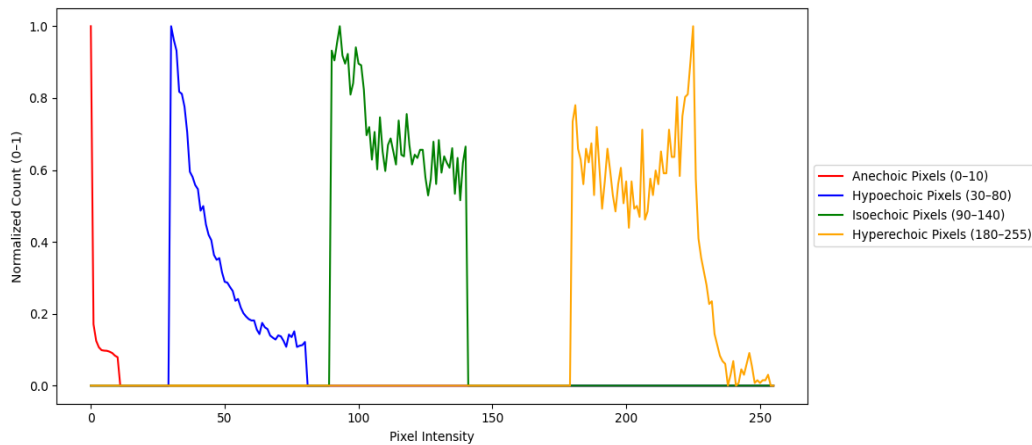
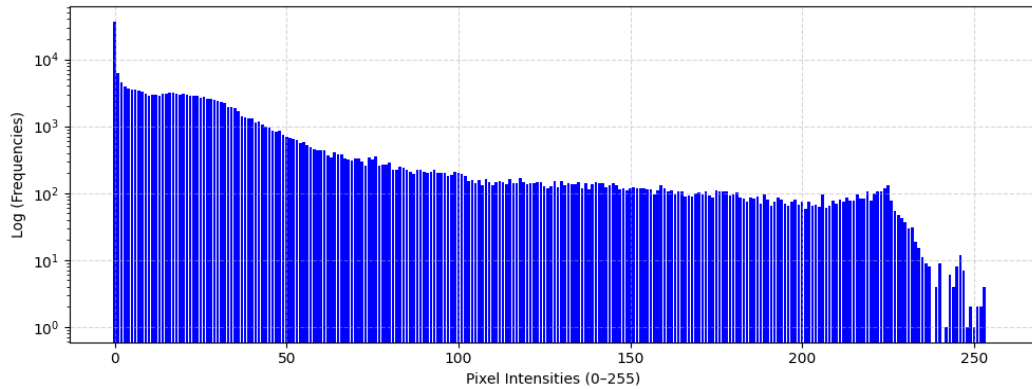
$$H = \begin{bmatrix} 36697 & 6280 & 4576 & 3920 & 3627 & 3577 & 3559 & 3469 & 3331 & 3068 \\ 2906 & 2937 & 2945 & 2854 & 3102 & 3088 & 3224 & 3142 & 3115 & 2956 \\ 3098 & 3004 & 2894 & 2894 & 2847 & 2675 & 2709 & 2611 & 2570 & 2486 \\ 2376 & 2286 & 2215 & 1943 & 1929 & 1844 & 1676 & 1414 & 1380 & 1326 \\ 1299 & 1157 & 1187 & 1066 & 997 & 962 & 864 & 832 & 843 & 750 \\ 687 & 681 & 652 & 626 & 561 & 573 & 517 & 479 & 457 & 440 \\ 431 & 431 & 371 & 341 & 415 & 386 & 374 & 331 & 317 & 305 \\ 332 & 326 & 296 & 257 & 338 & 321 & 359 & 256 & 264 & 268 \\ 289 & 223 & 226 & 247 & 243 & 223 & 209 & 195 & 220 & 221 \\ 206 & 200 & 211 & 221 & 203 & 198 & 204 & 179 & 186 & 208 \\ 198 & 197 & 182 & 154 & 159 & 139 & 156 & 133 & 165 & 144 \\ 132 & 148 & 152 & 144 & 136 & 163 & 142 & 141 & 167 & 148 \\ 136 & 142 & 140 & 145 & 145 & 128 & 117 & 127 & 150 & 124 \\ 151 & 131 & 141 & 137 & 134 & 146 & 118 & 140 & 114 & 137 \\ 147 & 141 & 143 & 121 & 133 & 140 & 133 & 115 & 117 & 109 \\ 119 & 124 & 117 & 120 & 119 & 116 & 113 & 97 & 111 & 131 \\ 117 & 105 & 112 & 97 & 108 & 107 & 88 & 91 & 90 & 99 \\ 104 & 95 & 105 & 93 & 87 & 111 & 108 & 106 & 106 & 92 \\ 97 & 103 & 87 & 83 & 74 & 87 & 82 & 89 & 70 & 95 \\ 80 & 65 & 75 & 87 & 79 & 70 & 64 & 74 & 80 & 67 \\ 75 & 58 & 75 & 65 & 66 & 62 & 94 & 61 & 64 & 76 \\ 70 & 79 & 74 & 86 & 78 & 78 & 94 & 84 & 84 & 106 \\ 77 & 99 & 106 & 107 & 119 & 132 & 76 & 54 & 47 & 42 \\ 37 & 30 & 31 & 19 & 15 & 11 & 9 & 8 & 0 & 4 \\ 9 & 0 & 1 & 6 & 4 & 8 & 12 & 7 & 1 & 2 \\ 1 & 2 & 2 & 4 & 0 & 0 & & & & \end{bmatrix}$$


Fig. 3 Histogram Classification

1.1 Edge detection

Edge detection is a fundamental step in ultrasound image analysis because many imaging artifacts are primarily characterized by abrupt spatial intensity variations rather than absolute pixel values. These variations are mathematically described by the image gradient, which represents the rate of change of image intensity in horizontal and vertical directions. For a grayscale image $I(x, y)$, the gradient is defined as

$$\nabla I = \left(\frac{\partial I}{\partial x}, \frac{\partial I}{\partial y} \right) \quad [10] (1)$$

and the edge strength is commonly measured by the gradient magnitude

$$|\nabla I| = \sqrt{\left(\frac{\partial I}{\partial x}\right)^2 + \left(\frac{\partial I}{\partial y}\right)^2} \quad [10] (2)$$

The Prewitt operator applies uniform weights, making it sensitive to horizontal and vertical transitions, which is useful for highlighting repetitive, regularly spaced structures such as reverberation and ring-down artifacts. The Prewitt filter approximates the horizontal and vertical derivatives using the following kernels for Horizontal derivative G_x And, for Vertical derivative G_y :

$$G_x = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}, G_y = \begin{bmatrix} -1 & -1 & -1 \\ 0 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix} \quad [10](3)$$

The Sobel operator introduces a smoothing effect through larger central weights, improving robustness to speckle noise commonly present in ultrasound images and enhancing continuous, high-contrast edges associated with mirror artifacts. The Sobel filter uses weighted kernels to improve noise robustness for Horizontal derivative G_x And, for Vertical derivative G_y :

$$G_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}, G_y = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix} \quad [10](4)$$

The gradient magnitude is then computed as:

$$|\nabla I| = \sqrt{G_x^2 + G_y^2} \quad [10](5)$$

In contrast, the Canny edge detector is a multi-stage algorithm that includes Gaussian smoothing, precise gradient computation, non-maximum suppression, and hysteresis thresholding, allowing it to detect thin and well-localized edges while suppressing noise. This makes Canny particularly effective for identifying strong, well-defined artifact boundaries. Together, these filters provide complementary information: Prewitt captures basic directional changes, Sobel offers noise-robust gradient estimation, and Canny delivers optimal edge localization. Their combined use is therefore highly important for reliable detection and analysis of ultrasound artifacts such as mirroring, ring-down, and reverberation, which are inherently edge-driven phenomena.



Fig. 4 Ultrasound Artifacts detection by filters operations

1.2 Patterns storing

After edge enhancement, a Discrete Hopfield network that operates with binary or bipolar states, where each unit represents a bin of the histogram, the network can serve as a content-addressable memory capable of recognizing and stabilizing toward characteristic histogram patterns [14] [15]. The network takes two-valued inputs: binary (0, 1) or bipolar (+1, -1), which is the standard representation in associative memory models [16] [17]. We load and convert ultrasound images into grayscale, then randomly extract small patches (e.g., 10×10 pixels) representing localized image regions. Each patch is binarized into vectors of +1 and -1 based on pixel intensity to make it suitable for Hopfield network processing.

$$M_{i,j} \rightarrow x \in \{-1, +1\}^N \quad [13](6)$$

is a 2D grayscale image patch extracted from a larger image. To be used by the Hopfield Network, it is first flattened and binarized into a 1D vector: Where N is the total number of pixels in the patch, and x is the input state vector representing the patch.




Fig. 5 Spinning Pixels Storing patterns

2. Dataset description

Nowadays, servers tend to store fewer purely static files and rely more on dynamic data, where templates are populated by variables at runtime to generate meaningful content. In a similar way, an ultrasound image can be viewed as consisting of a static component that represents the underlying anatomical structure, while the dynamic components correspond to imaging artifacts such as reverberation, ring-down, and mirror artifacts. These artifacts are not fixed features of the tissue but arise from variable interactions between the ultrasound waves, the imaging system, and the acoustic environment. Just as dynamic server-side data modifies a static template, ultrasound artifacts dynamically alter the base image, making their detection and modeling essential for accurate interpretation and analysis. We use Flask a Micro Framework written by Armin Ronacher in 2010 to train our model that will classify the dataset ultrasound images of Shenzhen People's Hospital [2].

Flask uses Web Server Gateway Interface WSGI that run python applications. And handle processing requests from the web server and create HTML or other markup formats that are returned to the user via an HTTP (Hypertext Transfer Protocol) request using Jinja2 template engine that can inherit from another Template. App routing is used to map a specific URL to the associated function that is intended to handle a request, enabling structured navigation and logical flow within a web application. The `url_for()` function dynamically builds a URL for a given function, ensuring flexibility, maintainability, and consistency even if route definitions change. URL stands for Uniform Resource Locator, which uniquely identifies the address of a resource on the web. In applications developed for medical imaging analysis—such as systems designed to study ultrasound imaging artifacts including mirror artifacts, ring-down effects, reverberations, and acoustic shadowing—proper routing allows different URLs to correspond to distinct processing modules, visualization pages, or diagnostic functions. The trained neural network model was saved to disk at the path `models/ultrasound_artifact_model.pth`, which stores the learned parameters and architecture state after completion of the training process. For storing a set of binary patterns $S^{(p)}$, $p = 1, 2, \dots, P$, the Hopfield network functions as an associative memory system that converges to stable states representing these stored patterns [18].

$$S^{(p)} = (S_1^{(p)}, S_2^{(p)}, \dots, S_N^{(p)}) \quad [11] (7)$$

the weight matrix W is given as:

$$w_{ij} = \sum_{p=1}^P [2 S_i^{(p)} - 1][2 S_j^{(p)} - 1] \text{ for } i \neq j \quad [11] (8)$$

Ultrasound Artifact Classifier




Fig. 6 Model Interface of Ultrasound Artifact classifier

The use of bipolar inputs makes the analysis easier. The network has symmetrical weights with no self-connections,

$$w_{ij}=w_{ji} \ ; \ w_{ii}=0 \quad [26] \ (9)$$

Hopfield network is guaranteed to converge to local minimum of its energy landscape, which corresponds to stored patterns or classification states. An energy function E_f of a discrete Hopfield network is characterized as a bound, nonincreasing function of the system state [14] [15]:

$$E_f = -\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n w_{ij} y_i y_j - \sum_{i=1}^n \theta_i y_i \quad [28] \ (10)$$

If the network is stable, the energy function decreases whenever the state of any node changes. If node i has changed its state from $y_i^{(k)}$ to $y_i^{(k+1)}$, i.e., the output has changed from +1 to -1 or from -1 to +1, the energy change E_f is given by:

$$\begin{aligned} \Delta E_f &= E_f(y_i^{(k)}) - E_f(y_i^{(k+1)}) \\ &= -(\sum_{j \neq i} w_{ij} y_j - \theta_i)(y_i^{(k)} - y_i^{(k+1)}) \\ &= -(net_i - \theta_i) \Delta y_i \end{aligned} \quad [27] \ (11)$$

The Hopfield network is trained using Hebbian learning, where each patch contributes to a collective memory matrix that captures the internal structure of known artifact patterns like ring-down or object Mirroring. The testing algorithm for the discrete Hopfield network is as follows:

Step 0: Initialize the weights to store patterns, i.e., weights obtained from training algorithm using Hebb rule.

Step 1: When the activations of the net are not converged, then perform Steps 2–8.

Step 2: Make the initial activations of the net equal to the external input vector X .

$$y_i = x_i \text{ for } i = 1 \text{ to } n \quad [4][5][20][28] \quad (12)$$

Step 3: Perform Steps 5–7 for each unit y_i . (Here, the units are updated in random order).

Step 4: Calculate the net input of the network:

$$y_{in_i} = x_i + \sum_j w_{ij} y_j \quad [4][5][20][28] \quad (13)$$

Step 5: Apply the activations over the net input to calculate the output:

$$\begin{cases} 1 & \text{if } y_{in_i} > \theta_i \\ y_i & \text{if } y_{in_i} = \theta_i \\ 0 & \text{if } y_{in_i} < \theta_i \end{cases} \quad [4][5][20][28]$$

where θ_i is the threshold and is normally taken as zero.

Step 6: Now feedback (transmit) the obtained output y_i to all other units. Thus, the activation vectors are updated.

Step 7: Finally, test the network for convergence.

A discrete Hopfield network can be modified to a continuous model, in which time is a continuous variable [19]. Continuous Hopfield networks can be applied to associative memory problems or optimization tasks, such as modeling ultrasound artifacts [16]. The operation of a Hopfield network is governed by an energy function $E_f(y)$, which always decreases whenever the state of a single node change. For a neuron y_i , the state update depends on whether its net input relative to threshold θ_i is sufficient to cause a change in activation. Specifically, where $\Delta y_i = y_i^{(k+1)} - y_i^{(k)}$, the change in energy ΔE_f exploits the facts that only one unit updates at a time, that $y_j^{(k+1)} = y_j^{(k)}$, for $j \neq i$, and that the weight matrix is symmetric, $w_{ij} = w_{ji}$ with zeros on the diagonal.

Two cases arise: if y_i is positive, it will change to zero when

$$x_i + \sum_j w_{ij} y_j < \theta_i \quad [12] \quad (14)$$

resulting in a negative Δy_i and hence $\Delta E_f < 0$. Conversely, if y_i is zero, it will change to positive when

$$x_i + \sum_j w_{ij} y_j > \theta_i \quad [12] \quad (15)$$

leading to a positive Δy_i , again ensuring $\Delta E_f < 0$. Thus, Δy_i is positive only when the net input is positive and negative when the net input is negative, guaranteeing that the energy cannot increase. Because the energy is bound below, the network must converge to a stable equilibrium where no further updates occur [16] [17]. From a theoretical perspective, this convergence property can be proven using the Lyapunov stability theorem. The energy function $E_f(y)$ associated with the Hopfield network is a Lyapunov function since:

$E_f(y)$ is continuous with respect to all components y_i , for $i = 1, \dots, n$.

$\frac{dE_f(y)}{dt} < 0$, indicating that the energy strictly decreases over time, ensuring asymptotic

The storage capacity of a Hopfield network sets a practical limit on how many distinct patterns (in this case, echogenicity intensity distributions) can be reliably stored and recalled [16] [17]. This capacity is approximately

$$C \approx 0.15n \approx \frac{n}{2 \log_2 n} \quad [12] \quad (16)$$

where n is the number of neurons.

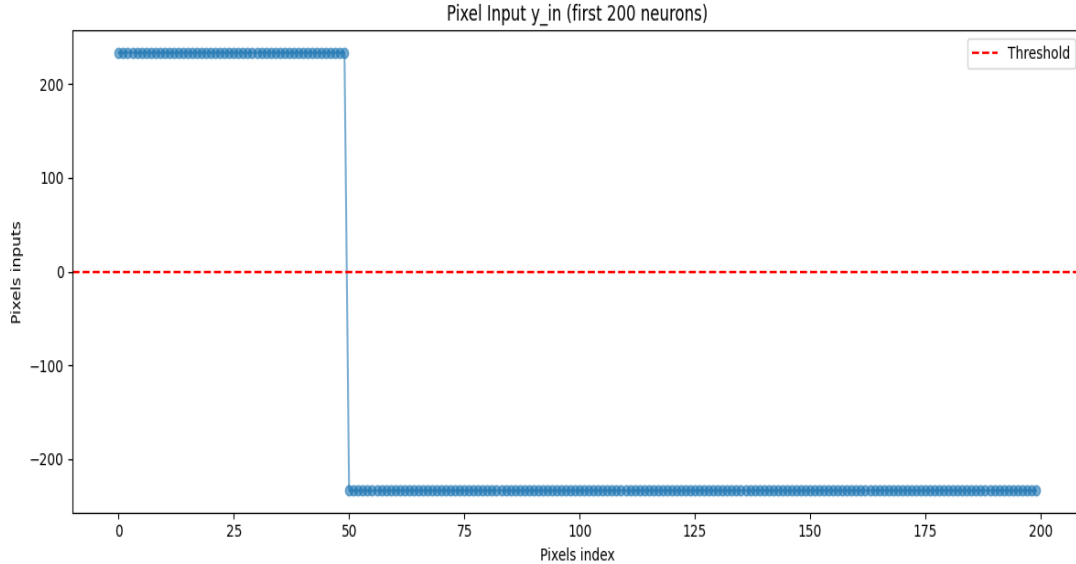


Fig. 7 Echogenicity Hopfield Network Storage Capacity

3. Results

3.1 Patterns Recognition by Hamming distance and energy

Hebbian learning is used to construct the network's weight matrix W from a set of known artifact patterns from a set of known artifact patterns $\{x^{(1)}, x^{(2)}, \dots, x^{(p)}\}$, where each $x^{(p)}$ is a binary training vector representing a reshaped image patch. The weight matrix is defined as:

$$W = \sum_{p=1}^p x^{(p)} (x^{(p)})^T - P.I \quad [15] \quad (17)$$

Where I , is the identity matrix used to zero out the diagonal elements and prevent self-connections. The result is a symmetric matrix $W \in \mathbb{R}^{N \times N}$ encoding the memory of the network. Each patch is reshaped into a binary vector and presented to the network. The state of the network is iteratively updated to minimize the energy function:

$$E = -\frac{1}{2} \sum_{ij} W_{ij} s_i s_j \quad [15] \quad (18)$$

This energy-minimization framework is particularly suitable for modeling reverberation artifacts, which arise when ultrasound waves repeatedly reflect between strong interfaces such as a needle and soft tissue or the transducer and bone, producing multiple, equally spaced false echoes that fade with depth. Continuous Hopfield networks with graded outputs leverage Lyapunov-stable energy functions to converge toward stable states corresponding to true tissue structures, thereby suppressing reverberation artifacts while preserving meaningful anatomical information.




Fig. 8 Reverberation artifact in ultrasound imaging.

where S_i is the state of neuron i and W_{ij} is the weight between neurons i and j . the patch is marked as containing a Mirror artifact. Mirror artifacts occur when the ultrasound beam encounters highly reflective surfaces, such as the diaphragm or pleura [9] [19]. Reflections cause a duplicated image of the structure on the opposite side of the reflective boundary, creating a false anatomical structure. Mirror artifacts can be conceptually related to temporal associative memories in neural networks, where sequences of patterns are recalled cyclically [20]. Similarly, acoustic impedance mismatches generate repeated echoes.

We define the Acoustic impedance as:

$$Z = \rho c \quad [25][28] \quad (19)$$

where ρ is the medium density and c is the speed of sound (Kremkau2015). The intensity reflection coefficient is

$$R_I = \left(\frac{z_2 - z_1}{z_2 + z_1} \right)^2 \quad [25][28] \quad (20)$$

In air gaps, multiple reflections generate a geometric series of diminishing echoes:

$$RTI_0, R^2TI_0, RTI_0, R^3TI_0, \dots$$

with cumulative intensity

$$I_a = \frac{r^2}{1-R^2} I_0 \quad [25][28] \quad (21)$$

Both temporal Bidirectional Associative Memory (BAM) networks mirror artifacts rely on recursive processes: BAM produces cyclic recall of stored patterns [20], while ultrasound recursive reflections produce duplicated echoes. During testing, the neurons update until convergence to a local minimum of E , which ideally corresponds to one of the stored patterns. If the final converged pattern matches a stored artifact template, the region is classified as containing a Mirror artifact. We perform image segmentation using clustering, we start by defining a distance function that measures how different two pixels are; for grayscale images, this is commonly the absolute or squared difference in pixel intensity. We find the local minimum in a Hopfield Neural Network (HNN), we use k , the number of clusters stored memories in a Hopfield Neural Network (HNN) where each cluster centroid serves as a prototype representing a group of similar pixels, while in a Hopfield network, stored patterns represent memory states that the network can recall through associative dynamics:

$$\arg \min_{i=1, \dots, k} d_{\text{Hamming}}(P, C_i) = \arg \min_{i=1, \dots, k} |P - C_i| \quad [25][28] \quad (22)$$

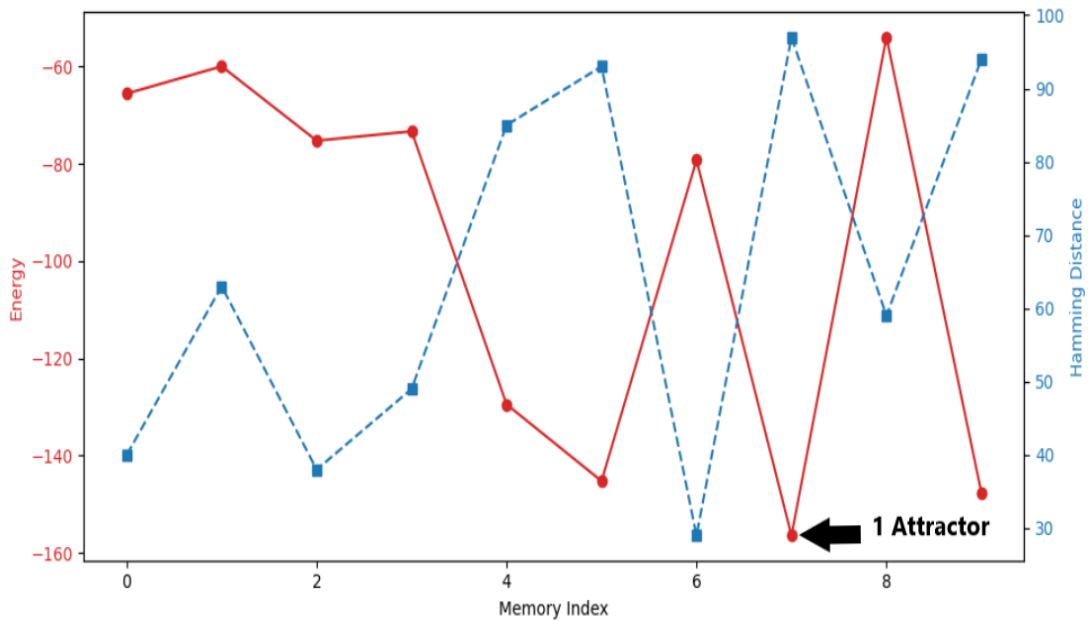




Fig. 9 Mirror artifact pattern recognition by Hamming distance and energy

A noisy test patch (simulating corrupted data) is then passed through the network, which iteratively updates the state to minimize an energy function essentially pulling the noisy patch toward a stored pattern (attractor). The network's energy convergence curve indicates how strongly the test input matches known patterns: rapid or deep convergence suggests similarity to a stored artifact. Ring-down artifact appears as a bright, continuous vertical streak extending posteriorly from gas-containing fluid structures [5] [9]. Iterative auto associative networks, such as Hopfield networks, can be trained to recognize and suppress such artifacts by converging toward stored patterns that represent true tissue structures.

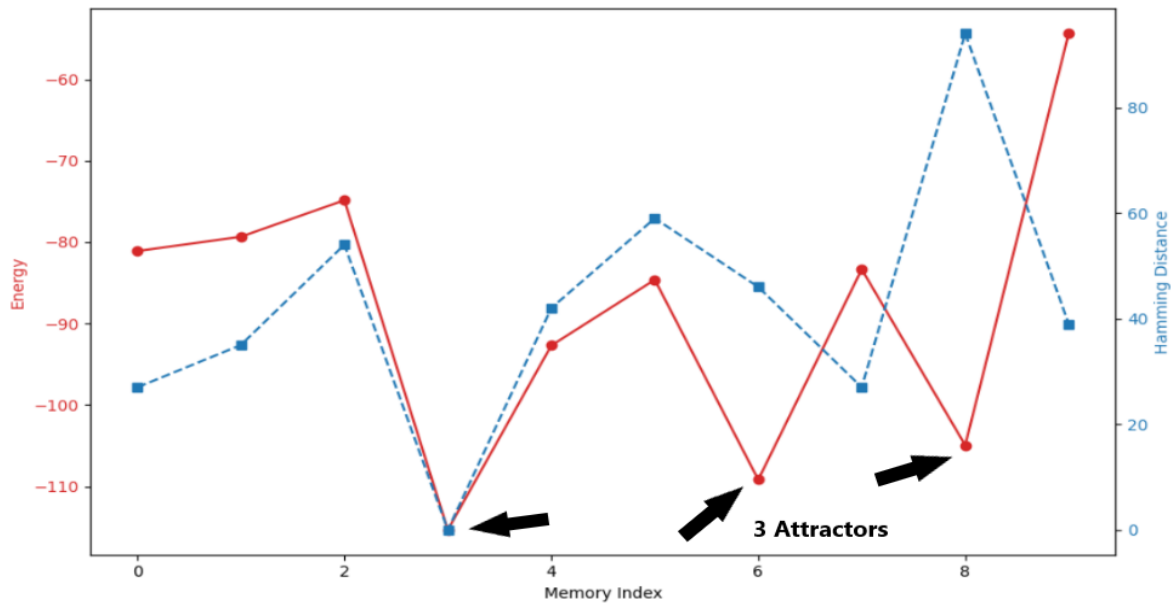



Fig. 10 Ring-down artifact recognition by Hamming distance and energy

By comparing energy curves for ring-down and mirror artifacts, one can evaluate how distinct each is under the learned memory model. The evolution of the Hopfield network energy as a function of iteration for two types of ultrasound artifacts, namely Ring-Down and Mirror artifacts.




Fig. 11 Ring-down artifact recognition by Hamming distance and energy

In both cases, the energy decreases monotonically with successive iterations, indicating stable convergence of the network toward a local minimum. During the initial iterations, the Ring-Down artifact exhibits slightly higher starting energy and a gradual decline, while the Mirror artifact shows a sharper energy drop at the beginning. However, by approximately the second or third iteration, both energy curves converge to nearly the same minimum value and remain constant thereafter. This rapid stabilization suggests that both artifact patterns correspond to strong and well-defined attractor states within the Hopfield network. The similarity in the final energy levels implies that, despite differences in early dynamics, the network ultimately encodes both artifacts with comparable stability, supporting the effectiveness of Hopfield energy as a reliable measure for ultrasound artifact characterization and classification ability.

3.2 Patterns Recognition by K-means clustering

Each pixel in the image is then assigned to the cluster whose centroid is closest to it, forming k groups based on similarity. The centroids are updated by calculating the mean value of all pixels assigned to each cluster. These steps of assigning pixels and updating centroids are repeated iteratively until the centroids no longer change between iterations, indicating convergence. However, since convergence is not always guaranteed, the algorithm is typically stopped after a predetermined maximum number of iterations to prevent infinite loops. This process segments the image into regions of similar pixel intensities, which can be useful to classify the artifacts .

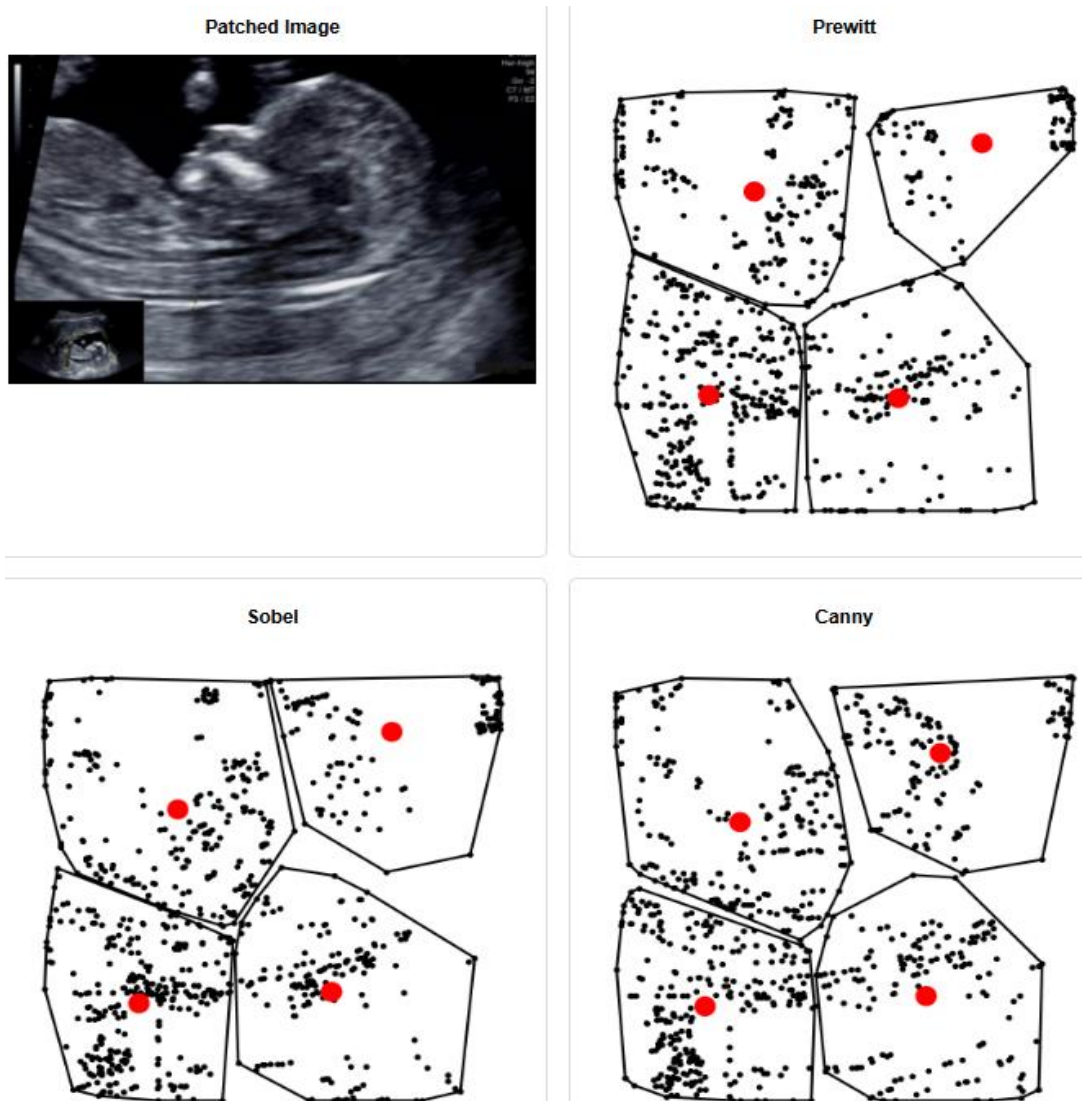


Fig. 12 *K-means clustering of filtered operators* Shenzhen People's Hospital ultrasound images dataset

K-means works by dividing the image into several clusters, where pixels with similar intensity values are grouped together into the same cluster. We choose a non-zero integer $C=4$ which represents the number of clusters or groups (hyperechoic, hypoechoic, anechoic, Isoechoic) pixels. The dataset ultrasound images of Shenzhen People's Hospital show the application of K-means clustering to an ultrasound image using features extracted after different edge detection methods. The original image is first divided into patches to capture local texture information. After applying Prewitt, Sobel, and Canny edge detectors, the detected edge pixels are treated as feature points and grouped using K-means.

```

X = Data set
C = Cluster centroids
P = Partition

K-Means(X, C) → (C, P)

REPEAT
  C_prev ← C

  // Assignment step
  FOR i = 1 TO N DO
    p_i ← FindNearest(x_i, C)
  END FOR

  // Centroid update step
  FOR j = 1 TO k DO
    c_j ← Average of x_i such that p_i = j
  END FOR
UNTIL C = C_prev

```

The black dots represent the data points, while the red dots indicate the cluster centroids learned by the algorithm. With the Prewitt operator, the clustering is relatively scattered due to noise sensitivity, leading to less compact clusters. Sobel produces more structured edge information, allowing K-means to form clearer and more meaningful clusters. The Canny-based result shows the most coherent and compact clusters, as its robust edge detection reduces noise and preserves significant boundaries. The results demonstrate that K-means effectively partitions the image into homogeneous regions, with clustering performance strongly influenced by the quality of the edge detection method, Canny providing the most reliable segmentation.

4. Discussion

Persisting the model in this format enables reproducible inference and further fine-tuning without retraining from scratch. The saved checkpoint can be reloaded for artifact classification on unseen ultrasound images, supporting consistent evaluation across experiments and deployment in downstream analysis pipelines. Using a standardized file path and format also facilitates integration with PyTorch-based workflows and ensures compatibility with future extensions of the ultrasound artifact detection framework. Confusion matrix provides a comprehensive summary of the performance of an ultrasound artifact detection system by comparing the true artifact status with the system's predictions. It is typically represented as a 2×2 table in which rows correspond to the ground truth (actual condition) and columns correspond to the system output (predicted condition). True positives (TP) are images or regions that contain artifacts and are correctly classified as artifacts, false positives (FP) are artifact-free images or regions incorrectly classified as artifacts, true negatives (TN) are artifact-free images or regions correctly classified as artifact-free, and false negatives (FN) are images or regions that contain artifacts but are incorrectly classified as artifact-free. From this confusion matrix, key performance metrics such as recall $\frac{TP}{TP+FN}$, precision $\frac{TP}{TP+FP}$, accuracy $\frac{TP+TN}{TP+TN+FP+FN}$, and specificity $\frac{TN}{TN+FP}$ can be directly derived to quantitatively evaluate the artifact detection performance.

		Predicted Values	
		Positive	Negative
Actual Values	Positive	TP	FN
	Negative	FP	TN

Fig. 13 Confusion Matrix of ultrasound images dataset

Recall and precision are standard performance metrics used to evaluate an ultrasound artifact detection system. Recall measures the system's ability to correctly identify all true artifact cases and is defined as the ratio of true positives to the total number of actual artifact cases, mathematically expressed as

$$\text{Recall} = \frac{TP}{TP+FN} \quad (23)$$

Precision measures how accurately the system identifies artifacts among all detections and is defined as

$$\text{Precision} = \frac{TP}{TP+FP} \quad (24)$$

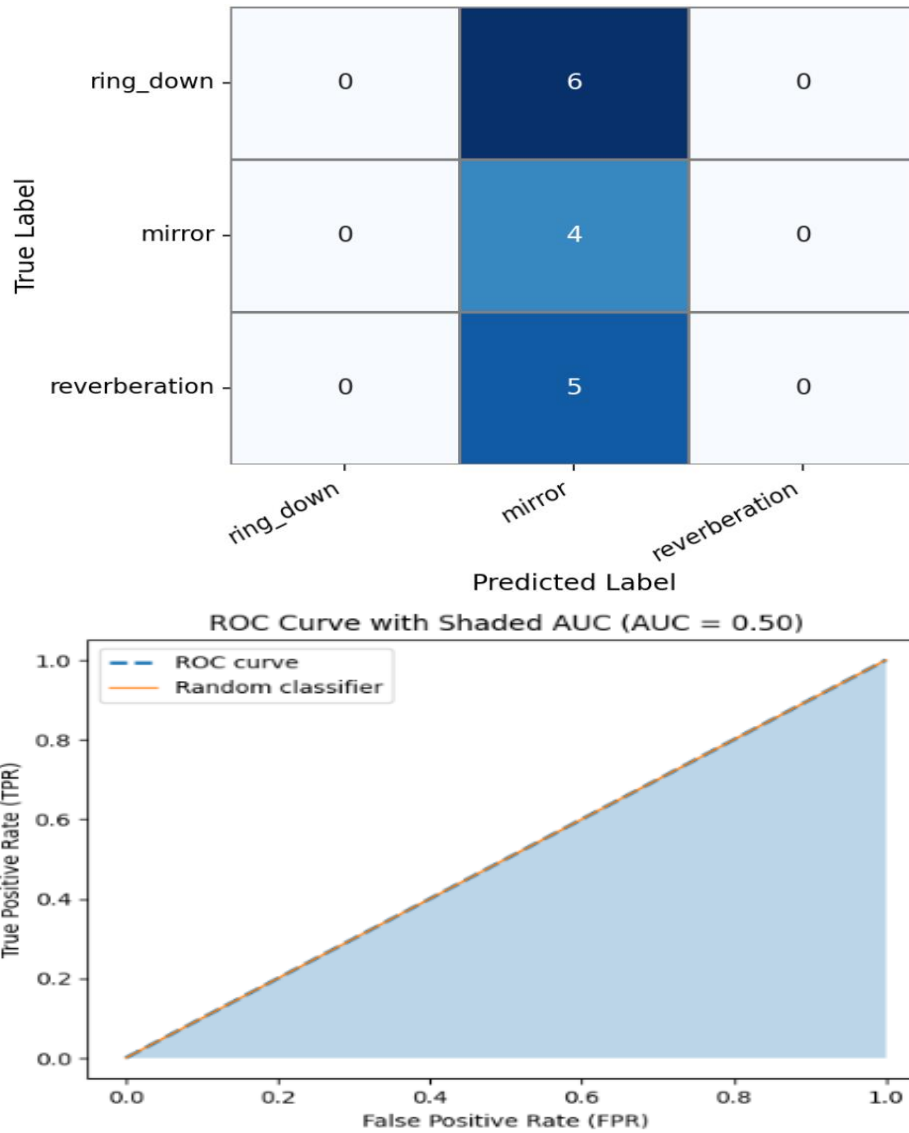


Fig. 14 Ultrasound images dataset Model Examination with AUC–ROC measures

Performance metrics such as the F1-score and the area under the receiver operating characteristic curve (AUC–ROC) are used to further evaluate the effectiveness of the ultrasound artifact detection system on a defined test set. The F1-score provides a balanced measure of performance by combining precision and recall into a single metric, defined as

$$\text{F1-score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (25)$$

and is particularly useful when the dataset is imbalanced. The confusion matrix highlights a severe misclassification issue in the model's performance for ultrasound artifact identification. All samples from the three true classes—ring-down, mirror, and reverberation—are consistently predicted as mirror, as indicated by the nonzero values appearing only in the “mirror” predicted column. This results in zero correct classifications for the ring-down and reverberation classes and no discrimination capability across artifact types. Such behavior suggests that the model has collapsed to a single dominant class, likely due to class imbalance, insufficiently discriminative features, or inadequate training. Consequently, while the model may show high confidence in predicting the mirror class, it demonstrates no generalization ability, rendering it

ineffective for reliable artifact classification. Further improvements are required, including better feature extraction, balanced training data, and possibly more expressive models or supervised learning strategies to achieve meaningful separation between artifact classes.

The AUC–ROC measures the system’s ability to discriminate between artifact and artifact-free images or regions across different decision thresholds; it is computed as the area under the ROC curve, which plots the true positive rate $\frac{TP}{TP+FN}$ against the false positive rate $\frac{FP}{FP+TN}$. Higher F1-scores and AUC–ROC values indicate better overall detection performance on the test set. The ROC curve appears as a diagonal line from (0,0) to (1,1), overlapping the random classifier reference, with a shaded area corresponding to an AUC–ROC value of 0.50. This indicates that the model has no discriminative power and performs no better than random guessing. This behavior occurs because the classifier predicts only a single class (mirror) for all samples, causing both the true positive rate and false positive rate to reach 1, with no intermediate decision thresholds to form a meaningful ROC shape. Consequently, the ROC curve degenerates to the diagonal, confirming that the model does not predict only a single-class strategy, moreover it success to distinguish between different ultrasound artifact types in the same ultrasound image.

5. Conclusions

This study explored the application of Hopfield Neural Networks (HNNs) trained with Hebbian learning for the detection and modeling of ultrasound imaging artifacts, with particular emphasis on ring-down, mirror, and reverberation effects. Ultrasound images were represented as grayscale matrices, and local image patches were extracted, binarized, and stored as associative memories within the Hopfield framework. By combining classical edge detection operators (Prewitt, Sobel, and Canny) with energy-based neural dynamics, the proposed approach aimed to exploit both spatial intensity variations and memory-driven convergence properties for artifact characterization. Theoretical analysis and experimental results confirmed the fundamental stability properties of Hopfield networks. The monotonic decrease of the energy function and convergence toward local minima validated the use of Lyapunov-based energy formulations for associative recall. Both ring-down and mirror artifacts exhibited rapid convergence to stable attractor states, and the similarity of their final energy levels indicated that the network successfully encoded these patterns as strong memories. Hamming distance and energy evolution proved to be meaningful indicators of pattern similarity, demonstrating that HNNs can effectively act as content-addressable memories for noisy or partially corrupted ultrasound data. In parallel, K-means clustering was employed as a complementary, non-associative segmentation method. The results showed that clustering performance was strongly influenced by the quality of edge detection, with Canny-based features yielding the most coherent and compact clusters. While K-means effectively segmented regions of similar echogenicity, it lacked the intrinsic memory and convergence properties of HNNs, highlighting the advantage of associative models for artifact-driven pattern recall.

Despite these strengths, quantitative evaluation using confusion matrices, F1-score, and AUC–ROC revealed a critical limitation in the current classification pipeline. The model collapsed to predicting a single dominant class (mirror artifacts), resulting in zero discriminative capability across artifact types and an AUC–ROC of 0.50, equivalent to random guessing. This behavior underscores the challenges posed by class imbalance, insufficiently discriminative feature representations, and the inherent limitations of unsupervised or weakly supervised associative models when applied to multi-class medical image classification. The findings demonstrate that Hebbian-trained Hopfield Neural Networks provide a robust theoretical and computational framework for modeling ultrasound artifacts, particularly in terms of stability, noise tolerance, and associative recall. However, for reliable multi-class artifact classification, future work must address dataset imbalance, incorporate more discriminative and multi-scale features, and potentially integrate supervised or hybrid deep learning architectures with Hopfield-based energy models. Such extensions could leverage the strengths of biologically inspired associative memory while achieving clinically meaningful diagnostic accuracy in ultrasound imaging.

References

- [1] T. Kohonen, *Associative Memory: A System-Theoretical Approach*, Springer, 1977.
- [2] C. Cui and F. Dong, "Dataset for Fetus Framework," 2022. [Online].
- [3] J. C. Russ, *The Image Processing Handbook*, CRC Press, 2011.
- [4] T. D. M. F. R. & W. P. A. Rossing, *The Science of Sound*, Addison Wesley., 2002.
- [5] A. J. & G. T. M. Bedard, *Atmospheric infrasound*. Physics, 2000, p. 32–37.
- [6] T. L. Szabo, *Diagnostic Ultrasound Imaging: Inside Out*, Academic Press, 2014.
- [7] P. N. T. Wells, *Biomedical Ultrasonics*, Academic Press, 1999.
- [8] F. A. Duck, "Physical Properties of Tissues," in *A Comprehensive Reference Book.*, 1990.
- [9] G. & H. T. Kino, "Interference of acoustic waves in layered media," *Journal of the Acoustical Society of America*, p. 1455–1462, 2005.
- [10] P. R. M. K. & T. A. Hoskins, *Diagnostic Ultrasound: Physics and Equipment*, Cambridge University Press, 2010.
- [11] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," in *Proceedings of the National Academy of Sciences*, 1982.
- [12] J. K. A. & P. R. G. Hertz, *Introduction to the Theory of Neural Computation*, Addison-Wesley, 1991.
- [13] R. C. & W. R. E. Gonzalez, *Digital Image Processing*, Pearson, 2018.

- [14] B. Jähne, *Digital Image Processing*, Springer, 2005.
- [15] W. K. Pratt, "Digital Image Processing:," 2007.
- [16] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, p. 62–66, 1979.
- [17] D. J. Amit, "Modeling Brain Function," *The World of Attractor Neural Networks*, 1989.
- [18] S. Haykin, *Neural Networks and Learning Machines*, Pearson, 2009.
- [19] G. E. & S. T. J. Hinton, "Optimal perceptual inference.," in *In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [20] T. T. Chow, *Ultrasonics: Methods and Applications*, Academic Press., 1968.
- [21] T. L. Szabo, *Diagnostic Ultrasound Imaging: Inside Out*, Academic Press, 2004.
- [22] F. W. Kremkau, *Sonography Principles and Instruments*, Elsevier Health Sciences, 2020.
- [23] F. W. Kremkau, *Diagnostic Ultrasound: Principles and Instruments*, Elsevier., 2015.
- [24] M. K. K. S. & B. M. S. Feldman, *US artifacts. Radiographic*, 2009, p. 1179–1189.
- [25] M. & H. V. Ramalho, "Mirror and Reverberation Artifacts in Ultrasound," *Ultrasound Quarterly*, p. 70–74.
- [26] J. J. Hopfield, "Neurons with graded response have collective computational properties like those of two-state neurons.," *PNAS*, p. 3088–3092, 1984.
- [27] B. Kosko, "Bidirectional associative memories," *IEEE Transactions on Systems, Man, and Cybernetics*, , p. 49–60, 1988.
- [28] K. & D. R. H. Michielsen, "Visualization of interference and diffraction patterns in classical and quantum mechanics," *American Journal of Physics*, p. 400–406, 2002.