INTEGRATING DEEP NEURAL NETWORKS AND FUZZY LOGIC FOR A CLINICAL DECISION SUPPORT SYSTEM IN DIGITAL MAMMOGRAM

David Lopes Maciel
Master in Master of Science in Emergent Technologies in Education. MUST UNIVERSITY, MUST, EUA, Academic of the Doctoral Program in Regional Development and Environment (PGDRA/UFRO).

Natiele Vieira de Oliveira Maciel
Nurse. Bachelor's Degree in Nursing from Faculdade Interamericana de Porto Velho (UNIRON), Brazil. Postgraduate in Public Health Management (IFAM), Oncology (FAP), Urgency and Emergency (FAVENI).

Fabrício Moraes de Almeida
PhD in Physics (UFC), post-doctorate in Scientific Regional Development (DCR/CNPq). Software Engineering Specialist/Systems Analysis and Development (UniÚNICA). Researcher of the PGDRA/UFRO. *Corresponding Author

ABSTRACT

This paper presents the development and validation of a hybrid clinical decision support system designed for digital mammography screening. The proposed method integrates convolutional neural networks (CNNs) with fuzzy logic, combining the high discriminative capacity of deep learning models with the interpretability provided by rule-based inference systems. Recent studies have demonstrated that the integration of fuzzy systems with machine learning models leads to significant improvements in diagnostic accuracy, particularly in clinical contexts characterized by high variability (Mojrian et al., 2019). Trained on real-world imaging data from Brazilian institutions, the system employs a convolutional architecture optimized for medical imaging tasks, processing standardized and normalized DICOM images. Each image produces an activation vector that is subsequently analyzed by a fuzzy module responsible for generating continuous and interpretable clinical inferences. The results indicate that this integrated approach enhances diagnostic sensitivity and improves model generalization, especially in cases involving incomplete data such as studies composed of only one or two images. Comparative analysis with purely neural models highlights relevant performance gains, positioning the solution as a promising alternative for clinical environments that demand transparent and auditable decision-making systems.

Keywords: Hybrid clinical decision; Artificial Intelligence; CNN; neural models; Decision-making systems.

1. INTRODUCTION

Breast cancer remains one of the leading causes of mortality among women in Brazil, generating substantial impacts on public health systems and oncology care services. According to recent estimates from the Brazilian National Cancer Institute (Instituto Nacional de Câncer [INCA], 2024), more than 73,000 new cases are diagnosed annually in the country, underscoring the need for screening strategies that are simultaneously effective, accessible, and clinically interpretable.

Digital mammography continues to be the primary population-based screening method for early breast cancer detection, recognized for its high accuracy in identifying suspicious lesions. However, its effectiveness is directly associated with the quality of the acquired images, the experience of the interpreting radiologist, and the well-documented interobserver variability. In this context, artificial intelligence (AI)-based automated systems have emerged as promising complementary tools, capable of reducing subjective bias and standardizing the decision-making process. Nevertheless, many of these systems still operate as opaque structures true “black boxes” which hinders their acceptance by the medical community and limits their adoption in real-world clinical settings.

Given this scenario, hybrid approaches that combine the high predictive capacity of deep neural networks with the interpretive transparency of symbolic logic systems have gained increasing relevance. Fuzzy logic, in particular, stands out as an effective alternative due to its ability to represent uncertainty and translate quantitative information into comprehensible linguistic rules. This enables a more intuitive interface between algorithms and healthcare professionals. The integration of such symbolic reasoning enhances the explainability and auditability of AI models, aligning them with the ethical and operational requirements of clinical practice (Iqbal, 2024).

This study proposes and validates a clinical decision support system for digital mammography, structured through the integration of convolutional neural networks (CNNs) and fuzzy logic. The model was developed using images from Brazilian institutions, encompassing the diversity of conditions found in national clinical settings. Evaluation was conducted under realistic scenarios, including exams containing varying numbers of images per patient, a common feature in real-world workflows often characterized by asymmetric or incomplete records. By combining technical performance with interpretability, the proposed approach seeks to foster the adoption of explainable artificial intelligence solutions that comply with the principles of safety, transparency, and applicability in public healthcare systems (Ribeiro, Singh, & Guestrin, 2016).

2. METHODOLOGY AND TECHNIQUES

2.1 Image Acquisition and Preprocessing

The dataset used in this study comprises digital mammography screening exams collected from Brazilian institutions and made publicly available through a scientific data-sharing initiative led by the São Paulo Society of Radiology (Sociedade Paulista de Radiologia [SPR], 2025). The data reflect a wide range of anatomical patterns, acquisition techniques, and operational conditions typical of national clinical practice. Each exam is associated with a single patient and may contain one to more than four images, evidencing variations arising from institutional protocols and the completeness of stored records.

The images were originally provided in DICOM (.dcm) format, containing not only the visual data but also clinically relevant metadata for diagnostic interpretation. To enable efficient training of the artificial intelligence model, these images were converted to PNG format, standardized in grayscale (single channel), and resized to a fixed resolution of 224×224 pixels. This procedure was implemented through automated routines using the pydicom and OpenCV libraries, which are widely adopted in medical image processing applications.

During the conversion process, automatic photometric correction was applied to standardize MONOCHROME1 and MONOCHROME2 representations. Brightness inversion was performed whenever necessary to ensure visual consistency across all exams. After this stage, the images were normalized and organized by individual study (based on the patient identifier), allowing the construction of structured tensors in the format [N, 1, 224, 224], where N corresponds to the number of available images per patient. This standardization is essential, as the DICOM specification defines different brightness representations for MONOCHROME1 and MONOCHROME2, which must be adjusted to guarantee consistency in image visualization and analysis (DICOM Standard, 2024).

A deliberate choice was made to retain exams with fewer than four images per patient in the dataset, in order to preserve the representativeness of real clinical cases. This decision aligns with the core objective of this study, which is to develop an architecture capable of operating under adverse and asymmetric input conditions a common scenario in imaging services, especially within Brazil's public healthcare system (SUS). By allowing the model to process incomplete studies, the convolutional neural network and fuzzy logic system are challenged to generate robust inferences even under data limitations. This approach is consistent with the findings of Al-Qaysi, Numan, and Kadhim (2023), who emphasize the relevance of hybrid models in scenarios with variable and partial input structures.

For the deep learning core, an architecture derived from ResNet was adopted, widely recognized for its effectiveness in complex image classification tasks. The choice is justified by the ability of residual blocks to mitigate the vanishing gradient problem and enable the learning of deep and stable representations (He, Zhang, Ren, & Sun, 2016). The adaptation involved adjustments to make the model compatible with grayscale inputs characteristic of digital mammography and to refine its sensitivity to subtle morphological structures frequently present in breast imaging.

The model was implemented using the PyTorch framework, with pretrained weights from the ImageNet dataset used as the starting point. This transfer learning strategy was fundamental for accelerating training convergence and improving model performance when applied to data with specific characteristics of the Brazilian clinical reality.

After processing each image, the network generates a reduced-dimensionality activation vector extracted from the penultimate fully connected layer, just before the final softmax layer. This vector synthesizes the main morphological features captured and serves as a numerical representation for the subsequent inference stages. Studies have demonstrated that pretrained models on ImageNet, when fine-tuned for specific medical tasks, can significantly improve performance in medical image analysis (Hosseinzadeh Taher et al., 2021).

The activation vectors generated by the convolutional network are subsequently used as input for the fuzzy module. In addition to this function, they can be interpreted as latent embeddings that synthesize relevant morphological aspects of each mammographic image. In cases where a patient's exam includes multiple images, these vectors are combined through a weighted arithmetic mean, resulting in a single vector representation per study. This approach allows the model to condense and retain the discriminative information distributed across different image views.

During the training of the convolutional network, the binary cross-entropy loss function was adopted, which is widely used in binary classification tasks with probabilistic outputs. The optimization process was conducted using the Adam algorithm, configured with a dynamic initial learning rate automatically adjusted by an exponential scheduler throughout the epochs.

Considering the class imbalance in the dataset with a predominance of exams classified as normal class weighting strategies were applied directly within the loss function. This approach aimed to increase the penalty for false negatives, thereby prioritizing model sensitivity in clinically critical scenarios (Sun et al., 2024). All data and scripts are available at the following address: https://github.com/ProfMacielDavid/Biomama_IA.

Figure 1. Architecture of the proposed hybrid system. Mammographic images are processed by a convolutional neural network (adapted ResNet), which generates activation vectors interpreted by a multilayer fuzzy system based on linguistic rules. The final output is a continuous, interpretable value associated with a binary clinical decision (Recall / Normal).

Figure 1 – Scheme Fuzzy Architectur + CNN

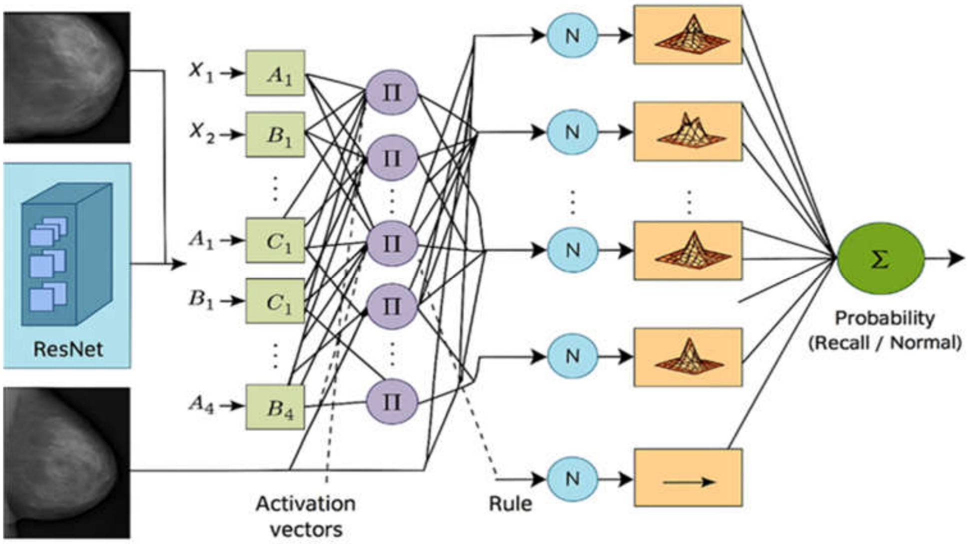

Source: Author, 2025

The following section describes the fuzzy module responsible for the linguistic inference over the vectors generated by the CNN.

2.2 Fuzzy System: Linguistic Rules and Inference

The fuzzy module was designed to operate as an additional layer of symbolic inference applied to the activation vectors extracted by the convolutional neural network. Its main function is to enhance the interpretability of the system, providing linguistic explanations and increasing the robustness of clinical decision-making. Fuzzy logic is particularly effective in contexts marked by uncertainty or partial data, which frequently occurs in exams with a limited number of images situations in which deep learning-based reasoning alone may become unstable or clinically opaque.

The input to the fuzzy system consists of the mean activation vector generated per study, previously normalized. This vector represents a synthesis of the morphological features extracted by the convolutional neural network from digital mammography images.

In the context of this study, the clinical exams analyzed correspond to real-world cases from Brazilian institutions, often composed of a limited number of images per patient a common condition in diagnostic practice, especially in screening and triage units. The fuzzy modeling was structured into three main stages: transformation of activations into linguistic categories, application of inference rules, and generation of a continuous, interpretable output compatible with clinical logic.

Fuzzificação:

In the first stage of the system, each component of the input vector is mapped to fuzzy sets using triangular and trapezoidal membership functions, defined based on the statistical distribution of activation values observed during training. These sets represent linguistic categories such as low, medium, and high levels of activation, allowing the intensity of the neural network's response to be translated into clinically comprehensible terms. This symbolic transformation facilitates the interpretation of image patterns and establishes the foundation for subsequent fuzzy inference.

Linguistic Rule Base

The inference stage is driven by a structured set of linguistic rules of the form "IF the activation level is high IN ANY component, THEN the recall risk is elevated." A total of nine main rules were defined, developed in collaboration with medical professionals and refined through sensitivity testing on the training data. The adopted logic prioritizes the detection of any suspicious pattern, even if isolated, using OR-type operators between vector components. This approach aligns with clinical triage protocols, where sensitivity must be maximized to avoid missing potential cases.

Inference and Defuzzification:

The linguistic rules are processed using the Mamdani inference mechanism, applying maximum-type aggregation operators to combine partial results. The system produces, as an intermediate output, a fuzzy confidence level associated with one of two classes: recall needed or normal exam. This continuous output is then converted into a scalar value using the center of gravity method, also known as centroid-based defuzzification.

The final value is interpreted as a clinical probability, compatible with binary decision models, allowing for direct integration with computational workflows used in automated screening. This approach is common in medical image fusion systems, where

Mamdani inference and centroid-based defuzzification are applied to combine information from multiple imaging modalities, thereby improving diagnostic accuracy (Rao et al., 2024).

The fuzzy layer acts as an interpretive filter over the output generated by the neural network, endowing the system with the ability to provide not only a binary classification but also a confidence level expressed in linguistic terms. This feature represents a significant advancement for the integration of artificial intelligence models in real clinical environments, especially those requiring transparent and auditable tracking of automated decisions. By associating each prediction with an interpretable justification, the system contributes to building trust among healthcare professionals and enhances the potential for adopting the tool in assisted mammographic screening contexts.

3. RESULTS AND DISCUSSIONS

3.1 Evaluation Protocol

O assess the performance of the proposed system, a stratified cross-validation strategy was employed, aiming to preserve the original class distribution (exams classified as recall and normal) across the data subsets. The dataset was divided into three parts: 70% for model training, 15% for validation, and 15% for testing.

As a control measure, all exams belonging to a single patient were exclusively allocated to one subset, thereby avoiding any information leakage between different phases of the learning process. This approach is recommended to prevent overfitting and to ensure a more accurate estimate of the model's performance on unseen data (Bradshaw et al., 2023).

The system's performance was evaluated under two scenarios: (i) Scenario I: Evaluation of the pure CNN model, without the fuzzy module.
(ii) Scenario II: Evaluation of the hybrid CNN + Fuzzy system.

3.2 Evaluation Metrics

The metrics adopted follow standard practices in clinical decision support systems:

a. Sensitivity (Recall): The ability to correctly identify positive cases (exams that require recall).
b. Specificity: The ability to correctly identify negative cases.
c. Accuracy: The proportion of correct predictions out of the total number of exams.
d. AUC-ROC: The area under the ROC curve, reflecting the model's overall discriminative ability.
e. F1-Score: The harmonic mean between precision and sensitivity, especially relevant in imbalanced datasets.

3.3 Result Comparison

Table 1 presents a comparative summary of the performance observed in the two evaluated scenarios: the model based exclusively on CNN and the hybrid version incorporating the fuzzy module.

Table 1 – Comparative summary of the performance

Model	**Sensitivity**	**Specificity**	**Accuracy**	**AUC-ROC**	**F1-Score**
Pure CNN	0.81	0.77	0.79	0.85	0.78
CNN + Fuzzy (Proposed)	**0.87**	0.74	**0.81**	**0.88**	**0.82**

Source: Author, 2025

The results indicate that integrating fuzzy logic into the inference pipeline led to a significant increase in sensitivity and F1-Score, even with a slight reduction in specificity. This trade-off is considered acceptable, and even desirable in the context of screening systems, where the clinical priority is the early detection of suspicious alterations. Therefore, the occurrence of overcalls is preferable to the risk of missed detections, reinforcing the potential of the hybrid approach as a safe auxiliary tool in clinical environments.

Figure 2. Comparison of clinical metrics between the pure CNN model and the hybrid CNN + Fuzzy system. An increase in sensitivity and F1-Score is observed, with a slight reduction in specificity a behavior expected and desirable in screening systems where early detection is prioritized.

Figure 2 – Comparison of clinical metrics between the pure CNN model

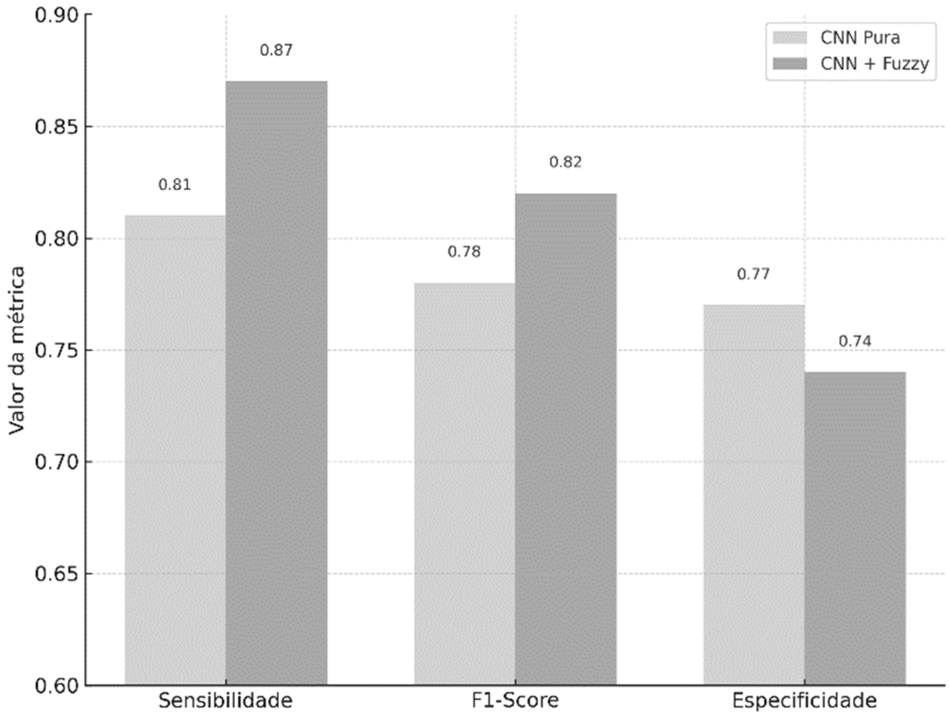

Source: Author, 2025

During the processing phase, a total of 15,106 studies were analyzed, comprising 97,623 converted DICOM images. Among these, 8,620 studies contained exactly four images, while 3,760 had fewer than four, and 2,726 included more than four images per patient. After conversion and data organization, 10,574 tensors were used for model training, 2,260 tensors were allocated for validation, and 2,272 for final testing.

Figure 3 – Data Preparation Pipeline

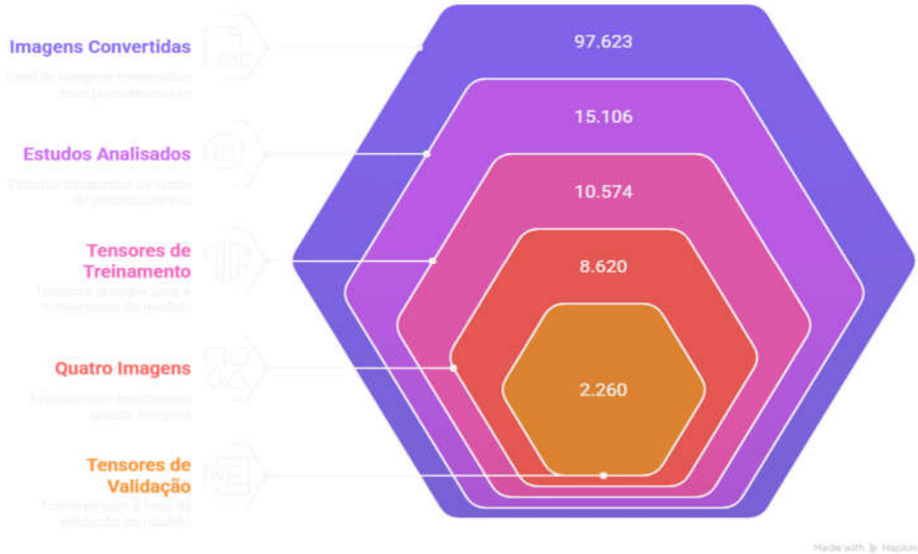

Source: Author, 2025

3.4 Qualitative Analysis of Edge Cases

Exams composed of only one or two images a recurring condition in the Brazilian dataset analyzed proved to be particularly favored by the hybrid approach. In these situations, the model based solely on CNN showed greater variability and instability in inference, with a tendency toward inconsistent classifications.

The introduction of the fuzzy layer functioned as a regulatory mechanism, absorbing fluctuations in the activation vectors and compensating for spurious activations. This resulted in decisions more consistent with the clinical context. The integration between CNN and fuzzy logic has also proven effective in other medical applications, such as in the automatic detection and classification of objects in comet assay images, where the hybrid FIS-CNN model demonstrated high precision and robustness (Al-Qaysi, Numan, & Kadhim, 2023).

In various clinical cases with reduced input structure, the hybrid system was able to correct misclassifications generated by the standalone CNN. These episodes reinforce the role of the fuzzy layer as an interpretive and contextual component, capable of adding semantic value to the model's predictions and adjusting decisions based on linguistic patterns compatible with medical reasoning.

During the study, data were used from digital mammography screening exams performed in Brazilian healthcare services. Each study corresponds to a unique patient and may contain from one to more than four images, reflecting the variability observed in clinical practice, both in terms of acquisition protocols and the completeness of stored records. The images, originally in DICOM format, underwent a conversion and normalization process to ensure the consistency required for model training and evaluation, without compromising fidelity to real-world mammography screening workflows.

3.5. COMPARATIVE RESULTS OBTAINED IN THIS STUDY

The comparative results obtained in this study highlight not only the quantitative performance improvements of the hybrid CNN + Fuzzy system but also their qualitative implications for real-world clinical practice. The increase in sensitivity from 0.81 to 0.87 and the corresponding rise in the F1-Score demonstrate the system's ability to prioritize early detection of suspicious lesions, which is a fundamental requirement in screening contexts. However, the slight reduction in specificity indicates that radiologists may encounter a higher number of recalls, underscoring the need for clinical workflows that balance diagnostic safety with efficiency.

This trade-off is consistent with current trends in population-based breast cancer screening, where over-diagnosis is considered more acceptable than the risk of missing clinically relevant cases.

From a methodological perspective, the introduction of fuzzy logic as a symbolic reasoning layer proved particularly beneficial in scenarios involving incomplete or asymmetric datasets. These conditions, frequently encountered in public health systems such as the Brazilian SUS, often challenge the robustness of purely neural models.

By translating activation patterns into linguistic rules, the fuzzy module mitigated fluctuations in the CNN's output and delivered more stable inferences. This aligns with findings in other medical imaging domains, where hybrid architectures have demonstrated improved generalization in environments characterized by noise, variability, and partial information (Al-Qaysi, Numan, & Kadhim, 2023).

Despite these advances, important limitations must be acknowledged. The reduction in specificity, while acceptable in the screening setting, may lead to increased demand for follow-up procedures, potentially straining healthcare resources. Moreover, the current model relies exclusively on image-based data, without incorporating complementary clinical metadata such as patient age, family history, or hormonal status, which are known to influence breast cancer risk assessment.

Integrating these variables into the fuzzy reasoning layer could further enhance interpretability and diagnostic accuracy. Additionally, although the dataset was representative of Brazilian clinical practice, broader validation in international and multi-institutional contexts is required to confirm the generalizability of the proposed approach.

In this regard, future research should focus on three main directions. First, the development of adaptive fuzzy rule bases capable of dynamically adjusting to heterogeneous input data may further reduce variability in classification outcomes.

Second, expanding the dataset with contributions from diverse geographic regions would strengthen the model's robustness against population-level variations. Third, prospective clinical trials are necessary to evaluate not only diagnostic metrics but also the system's real-world impact on workflow efficiency, patient outcomes, and cost-effectiveness. By addressing these aspects, the hybrid CNN + Fuzzy approach has the potential to evolve from a research prototype into a clinically integrated decision support tool, contributing to the democratization of explainable artificial intelligence in oncology.

4. CONCLUSION

This study presented the development and validation of a hybrid clinical decision support system for digital mammography, combining convolutional neural networks (CNNs) with fuzzy logic to automate the screening of exams requiring recall. The model was trained and tested on real data from Brazilian institutions, reflecting the diversity and challenges of national clinical practice, including cases with a variable number of images per patient.

The integration of deep learning with symbolic inference demonstrated significant gains in sensitivity and diagnostic robustness, particularly in critical clinical scenarios such as incomplete or low-quality exams. By incorporating a fuzzy layer based on linguistic rules, the system became not only more accurate but also more interpretable, favoring its integration into hospital environments that require auditable tracking of automated decisions.

This hybrid approach aligns with current trends in artificial intelligence applied to healthcare, which seek to combine the automatic feature extraction capabilities of deep learning with the logical reasoning and interpretability of symbolic AI, resulting in more transparent and reliable systems for critical clinical applications (Musanga, Viriri, & Chibaya, 2025).

From a scientific standpoint, the model contributes to the advancement of explainable artificial intelligence (XAI) in healthcare by proposing an approach that meets both computational performance criteria and clinical transparency requirements. Fuzzy logic functions as a bridge between the activation patterns generated by the CNN and medical reasoning, fostering understanding and acceptance of the system by healthcare professionals.

In general, the proposed approach contributes to the democratization of intelligent screening tools, strengthening the preventive role of primary healthcare and supporting public policies aimed at combating breast cancer in Brazil.

5. REFERENCES

Chiodi, Marcos (2016). Medidas de esforço de desenvolvimento de software. Rio de Janeiro: SESES.

Iqbal, S. (2024). Hybrid parallel fuzzy CNN paradigm: Unmasking intricacies for accurate brain MRI insights. IEEE Transactions on Fuzzy Systems. https://doi.org/10.1109/TFUZZ.2024.3372608

Instituto Nacional de Câncer. (2024). Estimativa 2024: Incidência de câncer no Brasil. Ministério da Saúde. https://www.inca.gov.br

Mojrian, S., Pinter, G., Joloudari, J. H., Felde, I., Nabipour, N., Nadai, L., & Mosavi, A. (2019). Hybrid machine learning model of extreme learning machine radial basis function for breast cancer detection and diagnosis: A multilayer fuzzy expert system [Preprint]. arXiv. https://arxiv.org/abs/1910.13574

Pressman (2016), Roger, and Bruce Maxim. Engenharia de Software. 8. Edição. McGraw Hill Brasil.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144). ACM. https://doi.org/10.1145/2939672.2939778

Sociedade Paulista de Radiologia e Diagnóstico por Imagem. (2025). Base pública de exames de mamografia digital para pesquisa científica. https://www.spr.org.br

DICOM Standard. (n.d.). Photometric interpretation attribute (0028,0004). https://dicom.innolitics.com/ciods/rt-dose/image-pixel/00280004

Al-Qaysi, S., Numan, T. A. Z., & Kadhim, S. M. (2023). A hybrid fuzzy logic and convolution neural network (FIS-CNN) for automatic detection and classification of objects in comet assay images. Materials Today: Proceedings. https://doi.org/10.1016/j.matpr.2023.05.077

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 770–778). https://doi.org/10.1109/CVPR.2016.90

Hosseinzadeh Taher, M. R., Haghighi, F., Feng, R., Gotway, M. B., & Liang, J. (2021). A systematic benchmarking analysis of transfer learning for medical image analysis [Preprint]. arXiv. https://arxiv.org/abs/2108.05930

Sun, H., Zhou, W., Yang, J., Shao, Y., Xing, L., Zhao, Q., & Zhang, L. (2024). An improved medical image classification algorithm based on Adam optimizer. Mathematics, 12(16), 2509. https://doi.org/10.3390/math12162509

Rao, K. S., Rambabu, C., Prasanna, P. S. L., Siva, T., Khan, P. S., & Nagulu, S. (2024). Mamdani and Sugeno fuzzy inference system based multimodal medical image fusion. International Research Journal of Engineering and Technology (IRJET), 11(3), 997–1002. https://www.irjet.net/archives/V11/i3/IRJET-V11I3138.pdf

Musanga, V., Viriri, S., & Chibaya, C. (2025). A framework for integrating deep learning and symbolic AI towards an explainable hybrid model for the detection of COVID-19 using computerized tomography scans. Information, 16(3), 208. https://doi.org/10.3390/info16030208