Convolutional Neural Network for Detection and Localization of Intracranial Hemorrhages in Brain CT Images

Jorge Luis Hernández Farah[1], Ismael Eliezer Pérez Ruiz [2]

[1] Universidad Modelo, Yucatán NJ 08544, MEX

15221784@modelo.edu.mx

Abstract. Intracranial hemorrhage (ICH) is a life- threatening neurological condition that requires urgent and accurate diagnosis. Computed tomography (CT) is the primary imaging modality used for early detection of ICH, yet manual interpretation of scans remains time-consuming and error-prone, this work proposes a deep learning pipeline based on a U-Net architecture for automated detection and segmentation of ICH in 2D axial brain CT images.
The approach integrates a binary slice level classifier to identify potential hemorrhagic slices, followed by a semantic segmentation to delineate hemorrhagic regions. The system was trained and evaluated on a publicly available dataset of 82 brain CT scans, including more than 2,500 labeled slices across multiple hemorrhage subtypes. All images were preprocessed to 512×512 resolution and normalized to a [0, 1] range.

Keywords: Intracranial hemorrhage, image detection, computed tomography, medical image segmentation, lightweight neural networks.

1 Introduction

Intracraneal hemorrhage (IHC) is a critical condition accounting for about 2 million strakes worlwide[1], it involves bleeding within the brain parenchyma and can result in acute neurological deterioration or death. Computed tomography (CT) is the imaging modality of choice for the early detection of IHC due to its wide availability and high sensitivity to acute blood [2].
However, manual interpretation of CT scans in emergency settings can be error-prone and delayed, especially in the absence of experienced radiologists. To address this challenge, artificial intelligence tools-particularly those based on deep learning-have been proposed to assist in the automated detection of ICH and support clinical workflows [3].
Among the deep learning approaches for biomedical image segmentation, the U-Net architecture has emerged as a particularly successful solution. Originally proposed by Ronneberger et al. (2015), U-Net features a symmetric encoder-decoder structure with skip connections that preserve spatial context, allowing it to produce high-resolution

segmentations even with relatively small datasets. This architecture has been successfully adapted for ICH segmentation in CT scans, with promising clinical implications [4].
To improve processing efficiency—especially in head CT studies containing hundreds of slices—some approaches adopt a two-stage pipeline: an initial binary classification to filter out non-hemorrhagic slices, followed by segmentation only on relevant pathological slices. This strategy reduces computational costs and avoids unnecessary segmentation [5].
This study proposes a fully U-Net-based deep learning system for the automatic segmentation of ICH in axial brain CT images. The pipeline integrates an initial slice level classifier to detect haemorrhage presence, followed by a semantic segmentation to delineate locate haemorrhagic regions.

2 Methodology

In this work it is proposed a deep learning pipeline for the automatic detection and segmentation of intracranial haemorrhages in brain CT images. The methodology consists of two main stages: (1) slice-level binary classification to identify the presence of haemorrhage, and (2) sematic segmentation using a U-Net architecture to locate haemorrhagic regions. The complete workflow is detailed below.

2.1 Dataset description

A dataset of 82 non-contrast brain CT scans was used in this study, comprising 36 scans diagnosed with intracranial hemorrhage (ICH). These positive cases represent various hemorrhage subtypes, including intraventricular, intraparenchymal, subarachnoid, epidural, and subdural hemorrhages. Each scan consists of approximately 30 axial slices with a slice thickness of 5 mm.
A lightweight convolutional classifier was first trained to predict whether a given CT slice contained visible hemorrhage, this step aimed to reduce computational overhead by avoiding unnecessary segmentation of normal slices.
All slices were independently reviewed by two board-certified radiologists. Hemorrhage types and the presence of skull fractures were recorded when identified, hemorrhagic regions were manually delineated, and consensus was reached for all segmentations.
The dataset is organized in two formats: (1) a Patients_CT directory containing 2D axial slices in JPG format, organized by patient ID, and (2) a Raw_Ct_Scans folder with 3D volumetric NIfTI files available for most patients. Within each folder, a subdirectory labeled “brain” contains the individual slices and corresponding binary hemorrhage masks, labeled usings the _HGE_Seg suffix. Accompanying metadata is provided through structured CSV files: patient_demograhics.csv, hemorrhage_diagnosis.csv, and hemorrhage_diagnosis_raw_ct.csv [6].

2.2 Preprocessing

All axial CT slices from the brain window modality were organized per patient and accessed from the Patients_CT directories, where each folder contained sequential JPEG images representing 5 mm axial slices. Each corresponding segmentation mask, when present, was stored in the same folder with filenames ending in _HGE_Seg.jpg. Slices without a corresponding mask were assumed to be negative for hemorrhage.
Each image and mask pair were resized from its original resolution (650x650 pixels) to 512x512 pixels to ensure compatibility with the input layer of the neural network while preserving anatomical proportion. All pixel values were normalized to the [0, 1] range by dividing by 255.
Maks images were binarized post resizing using a simple thresholding approach and saved in PNG format to ensure a consistent 8-bit format. The dataset was then indexed, and a CSV file was constructed to track each image and its binary label: hemorrhagic (1) or non-hemorrhagic (0), based on whether the mask was empty or not.

2.3 Model architecture

It was implemented as a lightweight version of the original U-Net architecture for semantic segmentation of hemorrhagic regions. The architecture consists of an encoder – decoder structure with four level of depth. The encoder is composed of sequential 2D convolutional layers with a kernel size of 3 x 3 with ReLu activation, followed by max pooling of 2 x 2 for spatial down sampling. Each block in the encoder doubles the number of filters (32, 64, 128, 256).
At the bottleneck, two convolutional layers with 256 filters are applied. The decoder mirrors the encoder with transposed convolution layers to progressively restore the spatial solution. At each decoder stage, skip connections are applied by concatenating feature maps from the corresponding encoder block, preserving high-frequency information [7,8].
To maintain a compact architecture, the model omits batch normalization and dropout layers. The final output layer uses a 1x1 convolution with sigmoid activation to produce a pixel wise probability map representing the likelihood of hemorrhagic tissue.
This U-Net variant was implemented using TensorFlow and Keras, with input shape to 512 x 512 for single channel grayscale CT images. The model was compiled with a custom loss function combining binary cross-entropy and Dice loss, suitable for addressing both pixel lev accuracy and segmentation overlap [9].

2.4 Training Strategy

The model was trained using a loss function combining binary cross-entropy and Dice loss, to balance a pixel wise classification accuracy with overlap-based segmentation performance. The Adam optimizer was used with a learning rate of $1x10^{-4}$. Training was conducted for up to 40 epoch with a batch size of 8 with early stopping based on validation Dice coefficient, so it really stopped in the 24^{th} epoch.
To address class imbalance (with hemorrhage or with no hemorrhage), a simple oversampling technique was used to replicate positive (hemorrhagic) samples within the

training set. This significantly reduced the number of false negatives and improved sensitivity. The dataset was split into training (80%) and validation (20%) was set at the patient level to avoid data leakage between slices of the same subject.

3 Results

The model achieved strong performance in binary classification of hemorrhagic versos non-hemorrhagic slices, with a recall of 0.89, precision of 0.50, Dice similarity coefficient of 0.60, and intersection over union (IoU) of 0.44. These results reflect the model´s high sensitivity, which is especially desirable in a clinical screening context to minize the risk of missed hemorrhages.

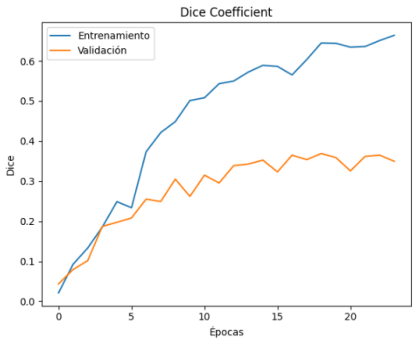

Fig. 3.1 Training and validation Dice coefficient curves. These plots indicate that the model steadily improved its segmentation performance during the initial epochs, reaching a downgrade just to be back up around epoch 15.

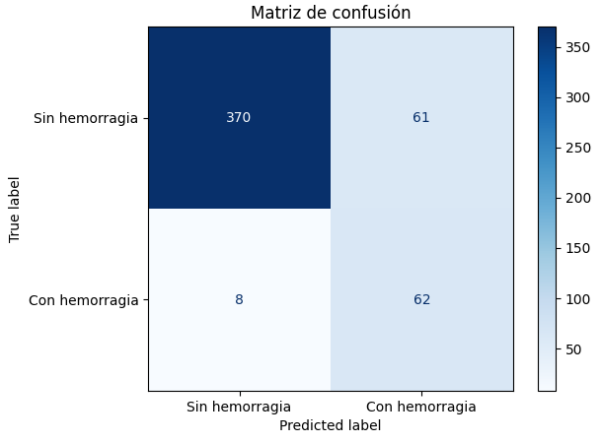

Fig. 3.2 Confusion matrix for the final model. It visually highlights the dominance of true negatives and a reasonably balanced detection of hemorrhagic cases.

The confusion matrix revealed 62 true positives, 370 true negatives, 61 false positives, and only 8 false negatives, confirming that the model was highly effective at detecting actual cases of hemorrhage. The oversampling strategy for hemorrhagic slices significantly improved sensitivity and reduced the number of false negatives compared to earlier training runs.

The model was further evaluated using per-slice binary classification derived from the segmentation masks, thresholding applied to determine presence or absence of hemorrhage. This classification task showed clear separation between negative and positive classes, as validated by high recall and low false negative rate.

Qualitative inspection revealed that the U-Net was capable of correctly segmenting large and small hemorrhagic regions, while maintaining variations in image and anatomical variability. Importantly, predictions on previously unseen slices demonstrated the model´s capacity to generalize beyond the training distribution.

Overall, the results support the effectiveness of a lightweight U.Net model for intracranial hemorrhage segmentation in 2D axial CT images. The combination of strong recall and interpretable prediction in segmentations highlights it's potentially a tool for clinical decision support or radiological triage [10].

To further illustrate the model´s segmentation capabilities, we include visual comparisons of its predictions on both familiar and unseen data. Figure 3.3 presents samples from the training batch, showing the original CT image, the ground truth segmentation mask, and the model´s predicted mask. These examples demonstrate the model´s capacity to localize hemorrhagic regions with high spatial accuracy and minimal oversegmentation.

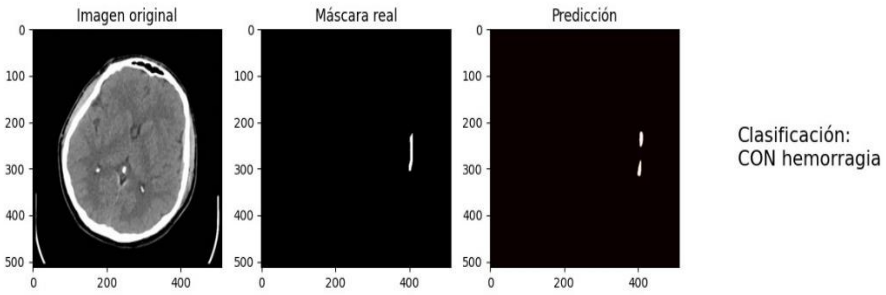

Fig. 3.3 Segmentation performance of the model on slices from the training batch. Each row shows (from left to right): the original brain window CT image, the ground truth hemorrhage mask, and the predicted segmentation mask generated by the U-Net.

Additionally, Figure 3.4 displays a prediction on a CT slice that was not part of the training or validation sets. Despite having no prior exposure to the patient, the model successfully detected the hemorrhagic area, confirming its ability to generalize to new inputs and maintain prediction quality outside the training distribution.

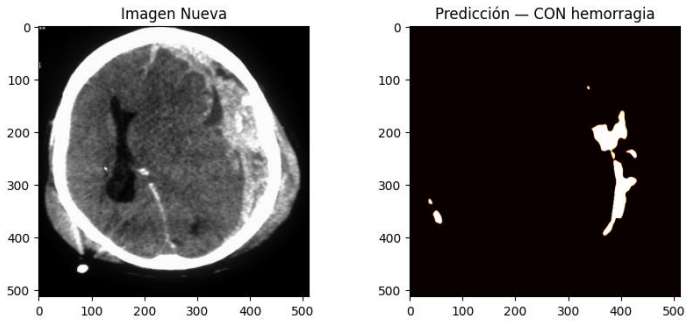

Fig. 3.4 Example of model inference on an unseen CT slice not included in the training or validation sets. Despite no prior exposure to patient data, the model correctly identifies the presence of hemorrhagic regions.

4 Discussion

The promising performance of lightweight U-Net architecture in this study highlights its potential applicability in resource constrained medical environments [11]. Compared to deeper or more computationally intensive models, the streamlined designed approximately 7 million parameters allowed efficient training and inference while maintaining competitive segmentation quality.

A key strength of the model lies in its high recall (0.89), which is of paramount importance in clinical screening workflows. Missing hemorrhage could result in catastrophic consequences, making sensitivity the priority metric. While precision (0.50) was comparatively lower, it is an acceptable trade off given that false positives can be addressed by radiologist review, whereas false negatives may go unnoticed.

5 Conclusion

This study presents an effective and computationally efficient approach for the detection and segmentation of intracranial hemorrhage in axial CT slices using a lightweight U-Net architecture. Despite the inherent challenges posed by limited training data, class imbalance, and the variability in hemorrhage presentation, the model achieved high sensitivity (recall = 0.89) and maintained robust spatial accuracy across diverse scans.

The model´s simplicity, requiring only grayscale brain window images and minimal preprocessing, supports its integration into real time clinical workflows in the future. Its strong generalization to unseen data and interpretable output masks provides an added layer of clinical reliability and transparency.

In conclusion, this work demonstrates the viability and relevance of lightweight deep learning architectures for critical radiological tasks such as hemorrhage detection and contributes to the growing bodies of literature supporting the use of artificial intelligence in neuroimaging diagnostics.

References

1. Qureshi, A. I., Mendelow, A. D. & Hanley, D. F. Intracerebral hemorrhage. Lancet 373, 1632–1634 (2009).
2. Luis C.F., Miguel A.G.N., Juan José E.G., Soledad P.S., Marcin B. (2023). Deep Learning Applied to Intracranial Hemorrhage Detection. PMID: 36826956. https://pmc.ncbi.nlm.nih.gov/articles/PMC9963867/
3. Arbabshirani, M.R., Fornwalt, B.K., Mongelluzzo, G.J., Suever, J.D., Geise, B.D., Patel, A.A., & Moore, G.J. (2018). Advanced machine learning in action: identification of intracranial hemorrhage on computed tomography scans of the head with clinical workflow integration. Medical Image Analysis, 48, 117–126. https://doi.org/10.1016/j.media.2017.07.005
4. Murtadha D. H., Muayad S.C, Aymen D.S., Hassan F.A., Zakaria A.Y., Behnaz G. (2020). Intracranial Hemorrhage Segmentation Using a Deep Convolutional Model. 4-7. https://www.researchgate.net/publication/339009990_Intracranial_Hemorrhage_Segmentation_Using_a_Deep_Convolutional_Model
5. Zhou, S.K., Greenspan, H., Davatzikos, C., Duncan, J.S., Van Ginneken, B., Madabhushi, A., & Summers, R.M. (2020). A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises. 14. https://doi.org/10.3390/data5010014
6. Hssayeni, M. (2019). Computed Tomography Images for Intracranial Hemorrhage Detection and Segmentation (version 1.3.0) PhysioNet. RRID:SCR_007345. https://doi.org/10.13026/w8q8-ky94
7. Bo P., Lianghong C., Qingchuan T., Enhui W., Yanmei Y. (2024). GA-UNet: A Lightweight Ghost and Attention U-Net for Medical Image Segmentation, pp. 1876-1886. https://pubmed.ncbi.nlm.nih.gov/38478188/
8. Badiuzzaman S., Rifat A., Sakib R., Hashem. (2021). CNL-UNet: A novel lightweight deep learning architecture for multimodal biomedical image segmentation with false output suppression (volume 70) ISSN 1746-8094. https://doi.org/10.1016/j.bspc.2021.102959.
9. Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen A.W.M. van der Laak, Bram van Ginneken, Clara I. Sánchez. (2017). A survey on deep learning in medical image analysis, Medical Image Analysis (Volume 42). pp. 60 -88 ISSN 1361-8415. https://doi.org/10.1016/j.media.2017.07.005.
10. Sasank C., Rohit G., Swetha T., Mustafa B., Norbert G.C., Vasantha K.V. (2018). Deep learning algorithms for detection of critical findings in head CT scans: a retrospective study (Volume 393), Issue 101162. https://www.thelancet.com/journals/lancet/article/PIIS0140-6736(18)31645-3/abstract
11. Justin L.W., Hassan F., Hanqi Z., Ali K. I. (2020). Segmentation of Intracranial Hemorrhage Using Semi-Supervised Multi-Task Attention-Based U-Net. *Appl. Sci.* 2020, *10*(9), 3297. https://doi.org/10.3390/app10093297