Universitat de Girona
Escola Politècnica Superior

Final Grade Project

Study: Grau en Enginyeria Biomèdica

Titol: Application of an Image Segmentation Method for Intracerebral Hemorrhage Images

Document: Memory

Student: Andreu Bofí Pla

Tutor: Arnau Oliver i Malagelada

Cotutor: Xavier Lladó Bardera

Department: Arquitectura i Tecnologia de Computadors

Area: Arquitectura i Tecnologia de Computadors

Call(month/year): June/2023

1 Introduction

Intracerebral hemorrhage (ICH) is a serious medical condition characterized by the accumulation of blood within the brain tissue. It can occur as a result of various factors, including trauma or spontaneous rupture of blood vessels, commonly referred to as a hemorrhagic stroke. The blood can either be confined within the brain tissue or extend into the cerebral ventricles or the subarachnoid space.

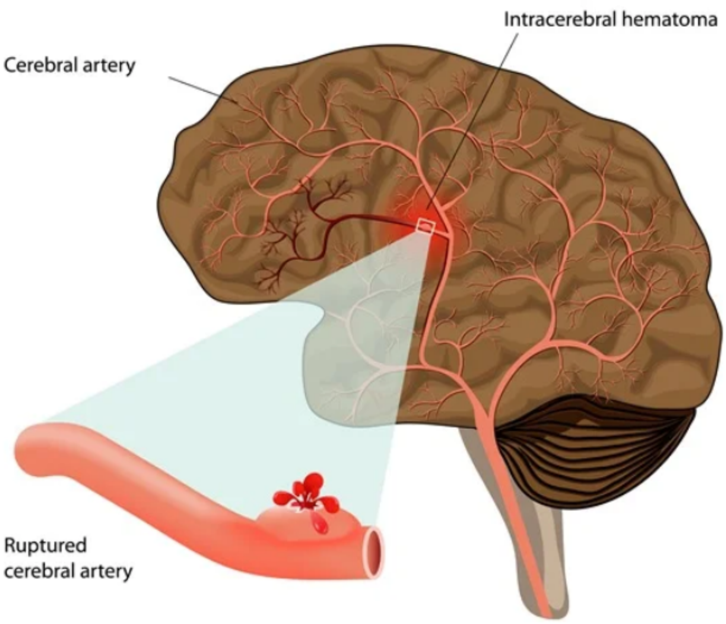

Figure 1: Cerebral hemorrhage

Cerebral hemorrhages carry a poor prognosis, with an approximate mortality rate of 40% ([2], [13]). Even for those who survive, the damage caused by prolonged increased cerebral pressure can often be irreversible. Medical treatment focuses on halting the bleeding, removing any clots, and alleviating pressure on the brain. In some cases, the brain may naturally reabsorb the blood, while in others, surgical intervention becomes necessary to save lives or minimize brain damage.

Computed tomography (CT) imaging is the primary diagnostic tool used by doctors to guide their decision-making process.

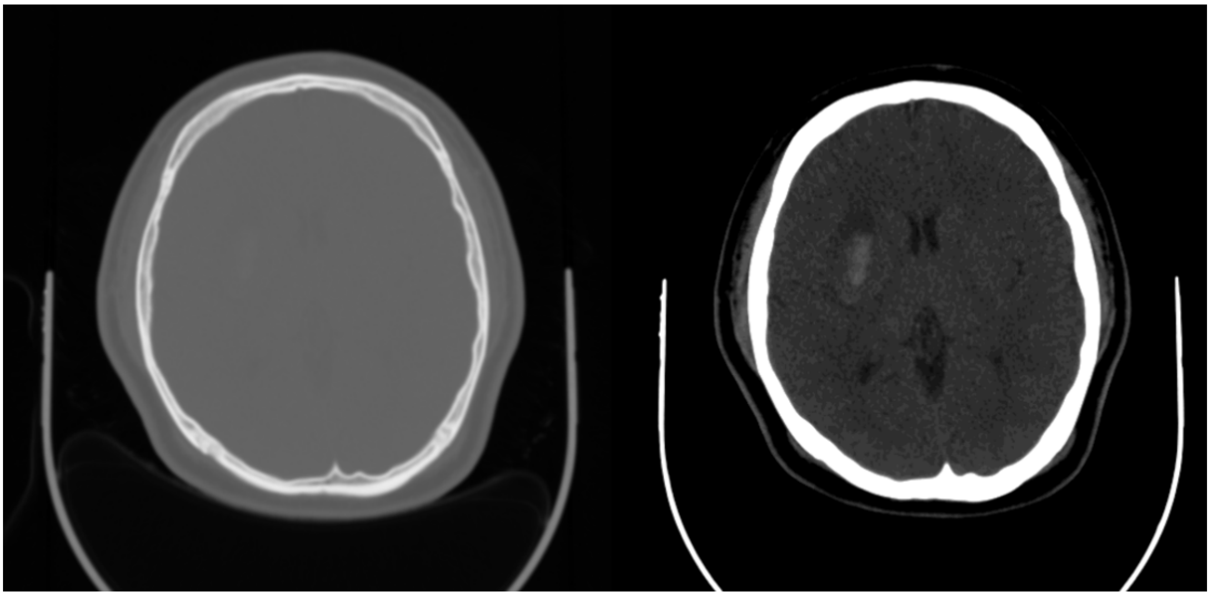

Figure 2: Comparison between the default image and a manual contrast correction using Mango

The development of an algorithm for the segmentation of intracerebral hemorrhages in CT images holds significant potential to support and enhance the decision-making process for medical professionals. This algorithm aims to improve efficiency, accuracy, and consistency in identifying and delineating hemorrhages, thereby facilitating treatment planning. Additionally, it contributes to medical education and training by providing standardized and objective results for teaching and improving diagnostic skills.

The main objective of this Bachelor's Thesis (TFG) is to develop an algorithm for the segmentation of cerebral hemorrhages, with a focus on facilitating subsequent decision-making in treatment by medical professionals. The algorithm will be based on a convolutional neural network (CNN) architecture, a deep learning technique that has shown great success in image analysis tasks.

By employing a CNN-based algorithm for hemorrhage segmentation, the research aims to achieve accurate and reliable results. This will contribute to improving the speed and precision of diagnosis, treatment planning, and patient care in cases of intracerebral hemorrhages.

To conduct this research, the publicly available INSTANCE 2022 database will be utilized. This database consists of 100 CT images of cerebral hemorrhages, accompanied by corresponding medical annotations. These annotations will enable the evaluation of the algorithm's automatic segmentation results using measures of overlap. The implementation of the algorithm will be carried out using the Python programming language, leveraging the deep learning library "Tensorflow" for efficient and effective development.

2 Previous Concepts

Before delving into the specifics of intracerebral hemorrhage (ICH) segmentation and its significance in medical decision-making, it is essential to establish a foundation by discussing the key concepts associated with this field. This introduction will provide an overview of the relevant concepts, including intracerebral hemorrhage, hemorrhagic stroke, medical imaging techniques, and deep learning algorithms.

Intracerebral hemorrhage refers to the accumulation of blood within the brain tissue, which can occur due to various factors such as trauma or spontaneous rupture of blood vessels. This condition poses a significant medical challenge as it can lead to severe neurological damage and even mortality. A particular type of intracerebral hemorrhage is a hemorrhagic stroke, characterized by bleeding in the brain resulting from the rupture of blood vessels. The prognosis for patients with cerebral hemorrhages is generally poor, with a substantial mortality rate. Survivors may experience irreversible damage due to increased cerebral pressure over an extended period.

Medical imaging techniques play a crucial role in diagnosing and assessing intracerebral hemorrhages. Computed tomography (CT) is one of the primary imaging modalities used by medical professionals to visualize and analyze brain abnormalities. CT scans provide detailed cross-sectional images of the brain, allowing doctors to identify and locate hemorrhages, evaluate their size and extent, and guide subsequent treatment decisions. These images serve as a valuable resource in understanding the nature and severity of intracerebral hemorrhages.

Deep learning algorithms, specifically convolutional neural networks (CNNs), have revolutionized image analysis in recent years. CNNs are powerful machine learning models capable of automatically learning intricate patterns and features directly from raw data. They have demonstrated remarkable success in various computer vision tasks, including medical image analysis. By leveraging CNNs, it is possible to develop algorithms that can accurately and efficiently segment intracerebral hemorrhages in CT images, enabling improved diagnosis, treatment planning, and overall patient care.

2.1 Hemorrhagic Stroke

Hemorrhagic stroke is a critical medical condition that falls under the broader category of stroke. It occurs when a blood vessel in the brain ruptures or leaks, leading to bleeding within the brain. This bleeding can cause damage to brain tissue, impairing its function and potentially leading to severe neurological complications.

Understanding the basic concepts related to stroke is crucial. Stroke refers to a sudden disruption of blood flow to the brain, resulting in the deprivation of oxygen and nutrients. Ischemic stroke, the most common type, is caused by a blockage in a blood vessel supplying the brain. On the other hand, hemorrhagic stroke occurs due to bleeding within the brain tissue.

To comprehend hemorrhagic stroke, it is necessary to familiarize oneself with the following concepts:

Blood vessels: Blood vessels are responsible for transporting oxygenated blood to the brain and removing deoxygenated blood. They consist of arteries, which carry blood away from the heart, and veins, which return blood to the heart. Any abnormalities or weaknesses in blood vessel walls can increase the risk of hemorrhagic stroke.

Intracerebral hemorrhage (ICH): This is the most common type of hemorrhagic stroke and occurs when a blood vessel ruptures within the brain, leading to bleeding in the surrounding tissue. The bleeding can cause compression and damage to brain cells, resulting in neurological deficits. This is the one that I will work on.

Subarachnoid hemorrhage (SAH): SAH refers to bleeding that occurs in the space between the brain and the surrounding membranes (meninges). It is often caused by the rupture of an aneurysm, which is a weakened area in a blood vessel wall. SAH can cause a sudden and severe headache, and if left untreated, it can result in brain damage or even death.

Epidural Hemorrhage: This type of hemorrhagic stroke occurs when bleeding accumulates between the inner surface of the skull and the outer layer of the brain's protective covering, known as the dura mater. Epidural hemorrhages are usually caused by a severe head injury that ruptures an artery, leading to bleeding in the space between the skull and the dura mater.

Subdural Hemorrhage: Subdural hemorrhage refers to bleeding that occurs between the dura mater and the brain's surface. It typically happens when blood vessels tear and blood collects in the space between these layers. Subdural hemorrhages often result from head trauma but can also occur spontaneously due to underlying conditions such as chronic alcoholism or the use of blood-thinning medications.

Risk factors: Several risk factors contribute to the development of hemorrhagic stroke. These include hypertension (high blood pressure), arteriovenous malformation (AVM), cerebral aneurysms, certain blood disorders, trauma, and drug abuse. Understanding these risk factors helps identify individuals who may be more susceptible to hemorrhagic stroke.

Symptoms: Hemorrhagic stroke presents with various symptoms, including a sudden and severe headache, nausea, vomiting, vision disturbances, weakness or numbness in the face, arm, or leg (often on one side of the body), difficulty speaking or understanding speech, loss of coordination, and loss of consciousness. Prompt recognition of these symptoms is crucial for timely medical intervention.

Diagnosis and treatment: The diagnosis of hemorrhagic stroke involves a combination of medical history assessment, physical examination, neuroimaging (such as computed tomography or magnetic resonance imaging), and sometimes additional tests to determine the underlying cause. Treatment typically includes stabilizing the patient's condition, managing blood pressure, surgically repairing aneurysms or AVMs, and providing supportive care to prevent further complications.

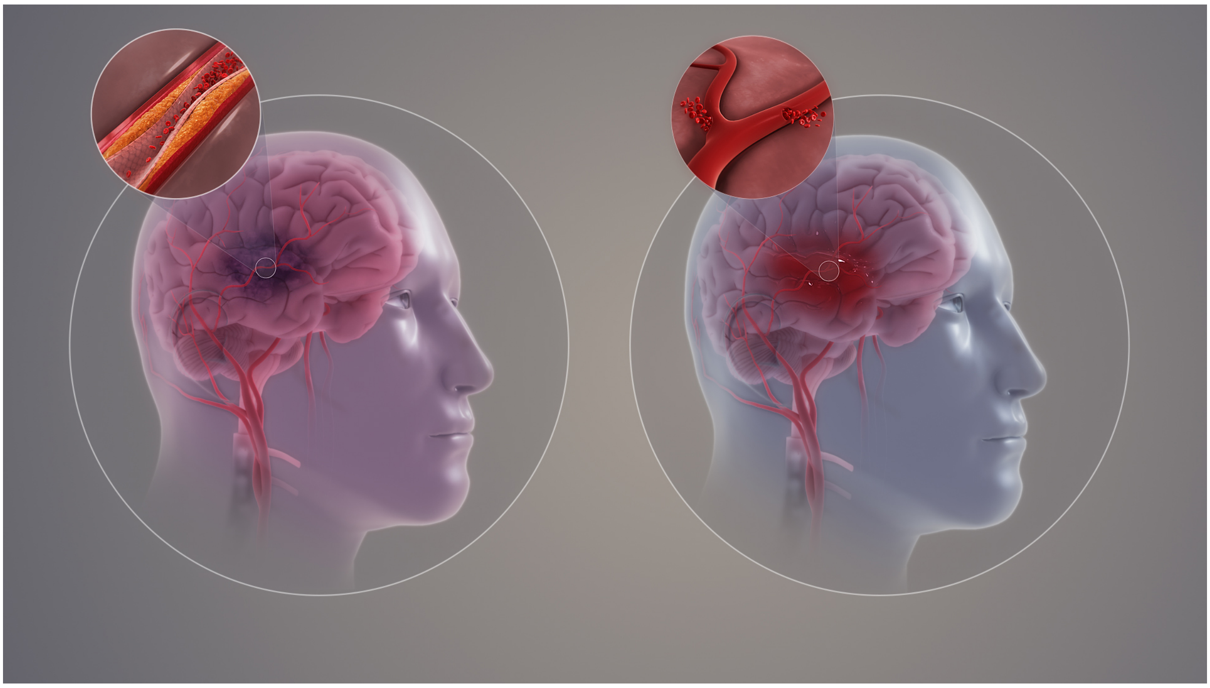

Figure 3: Blood clot stops blood flow to an area of the brain in Ischemic (L) whereas blood leaks into brain tissue in case of Hemorrhagic(R).

2.2 Intracerebral Hemorrhage

Intracerebral hemorrhage (ICH) is a complex medical condition that involves the accumulation of blood within the brain tissue. This condition can arise from various causes, including trauma, high blood pressure, abnormalities in blood vessels, or the use of blood-thinning medications. When a blood vessel ruptures or leaks within the brain, it leads to bleeding which can result in significant neurological damage.

The consequences of intracerebral hemorrhage can be severe and life-threatening. As blood accumulates within the brain tissue, it creates pressure and compresses surrounding structures. This compression can disrupt normal brain function and lead to a range of neurological symptoms, such as severe headaches, loss of consciousness, paralysis, difficulty speaking, or impaired vision. The severity of symptoms depends on the location, size, and extent of the hemorrhage.

Hemorrhagic stroke is a specific type of intracerebral hemorrhage that occurs when a blood vessel within the brain ruptures and causes bleeding. It is distinct from an ischemic stroke, which is caused by a blockage in a blood vessel that cuts off the blood supply to a part of the brain. Hemorrhagic strokes account for approximately 15% of all strokes but tend to have higher mortality rates and a greater likelihood of long-term disability.

Diagnosing intracerebral hemorrhage and understanding its characteristics are crucial for effective medical intervention. Medical imaging techniques, such as computed tomography (CT) scans, play a vital role in visualizing and identifying the presence of hemorrhages in the brain. CT scans provide detailed cross-sectional images that allow medical professionals to accurately locate and assess the size and extent of the bleeding. This information is essential for determining the appropriate treatment approach and predicting the prognosis for the patient.

Managing intracerebral hemorrhage involves several treatment strategies aimed at stopping the bleeding, reducing pressure within the brain, and minimizing further damage. Medical interventions may include medication to control blood pressure, surgical procedures to remove blood clots or repair ruptured blood vessels, and interventions to manage increased intracranial pressure.

In summary, intracerebral hemorrhage is a complex medical condition characterized by the accumulation of blood within the brain tissue. It can result from various causes and can have severe neurological consequences. Accurate diagnosis through medical imaging techniques, such as CT scans, is crucial for effective treatment and management. Understanding the nature of intracerebral hemorrhage is essential for developing algorithms that can assist medical professionals in accurately segmenting and analyzing hemorrhages, ultimately improving treatment decision-making and patient outcomes.

2.3 Image Processing and Segmentation

Image segmentation is a fundamental task in computer vision that involves partitioning an image into meaningful and semantically coherent regions or objects. The goal of image segmentation is to assign a label or class to each pixel or region in the image, effectively dividing it into distinct and meaningful components.

There are different types of image segmentation techniques, each with its strengths and limitations. Here are some commonly used approaches:

Thresholding:

Thresholding is a simple and commonly used technique in image segmentation. It involves selecting a threshold value and classifying each pixel based on its intensity or color. Pixels with values above the threshold are assigned to one class, while those below the threshold are assigned to another class. Thresholding works well when the foreground and background intensities are distinctly separable, but it may struggle with images containing complex or overlapping regions.

Region-based Segmentation:

Region-based segmentation algorithms group pixels into regions based on specific criteria such as color, texture, or intensity similarity. These algorithms typically start with an initial set of regions and iteratively merge or split them based on predefined rules. Examples of region-based segmentation methods include the watershed algorithm, mean-shift clustering, and graph-based segmentation.

Boundary-based segmentation:

Boundary-based segmentation separates objects or regions within an image by detecting their boundaries or edges. It involves two main steps: edge detection and edge linking. Edge detection identifies and highlights the boundaries or edges in the image, while edge linking connects these edges to form closed contours or boundaries around objects. By detecting and linking boundaries, this technique allows for the separation and identification of different objects or regions within an image, making it useful in various computer vision applications.

The above-mentioned methods can segment without a priori information of the image. However, in medical imaging, it is common to have a set of previous examples to guide the segmentation algorithm. This is usually referred to as supervised segmentation. The general approach involves training a model to classify each pixel in an image into different predefined classes or categories. During the training process, the algorithm is presented with a set of input images along with their corresponding pixel-level annotations, which serve as the ground truth for segmentation. The algorithm learns to identify patterns and features in the images that are associated with each class label. Once the model is trained, it can be used to segment new, unseen images by predicting the class labels for each pixel. The algorithm applies the learned knowledge to classify pixels in the input image, producing a segmentation mask where each pixel is assigned a specific class label. In this way, the segmentation process is transformed into a classification problem, and techniques based on artificial intelligence can be applied to solve the segmentation.

2.4 Introduction to Artificial Intelligence

Artificial Intelligence (AI) refers to the simulation of human intelligence in machines, enabling them to perform tasks that typically require human intelligence, such as learning, reasoning, problem-solving, and decision-making, in this case, the correct segmentation of a CT ICH.
AI encompasses a wide range of techniques, algorithms, and technologies that enable machines to mimic and automate intelligent behavior.

At its core, AI aims to develop computer systems that can perceive, understand, and interact with the world in a manner similar to humans. It involves creating models and algorithms that allow machines to process and interpret data, learn from experience, and make informed decisions or predictions. AI systems often rely on large amounts of data, sophisticated algorithms, and computational power to perform complex tasks.

The working principle behind AI involves training models and algorithms using large datasets that are representative of the task or domain at hand. The models are exposed to the data, and through iterative processes, they learn to recognize patterns, make predictions, or generate responses. The learning process involves adjusting the internal parameters of the models based on feedback and evaluation metrics. The more data and feedback the models receive, the better they become at performing the desired tasks.

AI systems are typically designed to be adaptable, capable of learning from new data or experiences and improving their performance over time. They can handle complex and unstructured data, reason through uncertain scenarios, and generalize knowledge to make decisions in new situations.

2.4.1 Deep Learning

Deep learning is a subfield of machine learning that focuses on training artificial neural networks with multiple layers, also known as deep neural networks, to learn hierarchical representations of data. It is inspired by the structure and functioning of the human brain, specifically the interconnected network of neurons.

Deep learning has gained significant attention and popularity due to its ability to automatically learn and extract complex features from raw data. Unlike traditional machine learning approaches that rely on handcrafted features, deep learning models can learn intricate representations directly from the data. This capability makes deep learning particularly effective in handling high-dimensional and unstructured data such as images, text, and audio.

One of the most widely used deep learning architectures is the Convolutional Neural Network (CNN). CNNs are designed specifically for processing grid-like data, such as images. They have revolutionized computer vision tasks and achieved state-of-the-art performance in various image-related tasks, including image classification, object detection, image segmentation, and more.

The key components of a CNN are convolutional layers, pooling layers, and fully connected layers:

- Convolutional Layers: Convolutional layers are the building blocks of a CNN. They consist of learnable filters or kernels that convolve across the input image to extract local features. Each filter performs a convolution operation by computing dot products between its weights and a small receptive field of the input. The outputs of multiple filters are stacked to create feature maps that highlight different visual patterns or features in the input.
- Pooling Layers: Pooling layers help reduce the spatial dimensions of the feature maps while retaining the most relevant information. The pooling operation (typically max pooling) downsamples the feature maps by selecting the maximum value within each local region. This downsampling helps make the representations more invariant to small translations and reduces the computational complexity of the network.
- Fully Connected Layers: Fully connected layers are typically placed at the end of the CNN architecture. They take the flattened feature maps from the previous layers and connect every neuron to every neuron in the subsequent layers. These layers help in making high-level predictions or classifications based on the learned representations.

During training, CNNs use a technique called backpropagation to adjust the weights of the filters and fully connected layers based on the error between the predicted output and the ground truth labels. This optimization process aims to minimize the difference between predicted and actual outputs, improving the model's ability to make accurate predictions.

2.4.2 U-Net

A U-Net is a type of convolutional neural network (CNN) architecture that is widely used for image segmentation tasks, including medical image segmentation. It was originally introduced by researchers at the University of Freiburg in 2015, and its name comes from its U-shaped architecture.

The U-Net architecture is designed to address the challenges of accurately segmenting objects or regions of interest within images while preserving spatial information. It is particularly effective when the objects of interest are small or have complex shapes. The U-Net architecture has been successfully applied in various medical imaging tasks, including the segmentation of organs, tumors, and lesions.

The key feature of the U-Net architecture is its symmetric encoder-decoder structure. The encoder part of the network consists of multiple convolutional layers that progressively downsample the input image, capturing high-level feature representations. This encoding pathway is responsible for capturing context and extracting abstract features.

The decoder part of the network is the symmetric counterpart of the encoder and performs the upsampling operation. It consists of upconvolutional layers that progressively increase the spatial resolution of the feature maps. This pathway enables the network to localize and refine the segmented regions accurately.

The U-Net architecture also incorporates skip connections, which are direct connections between corresponding layers in the encoder and decoder. These skip connections help in preserving fine-grained spatial information by concatenating feature maps from the encoder to the decoder layers. This allows the network to combine high-level semantic information with detailed spatial information, facilitating accurate segmentation.

During training, U-Net models are typically trained using a pixel-level loss function, such as the Dice coefficient or cross-entropy loss, to measure the similarity between the predicted segmentation masks and the ground truth masks. The model's parameters are optimized through gradient-based methods, such as stochastic gradient descent (SGD) or its variants, to minimize the loss function.

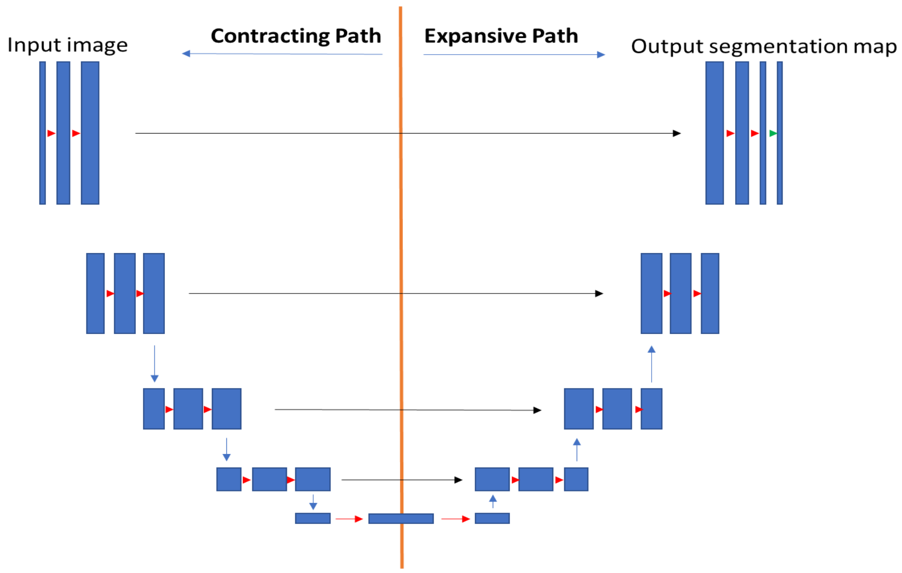

Figure 4: Used U-net exemplification. The left part is the contracting or encoding path, while the right one is the expansive or decoding part. The inner part of the net (bottom layer) is also known as the bottleneck.

3 State of the Art

Intracranial Hemorrhage (ICH) detection and segmentation continues to evolve with advancements in medical imaging and artificial intelligence (AI) techniques.

AI-based approaches, particularly deep learning methods, have shown promise in ICH detection and segmentation tasks. Convolutional Neural Networks (CNNs) and related architectures have been widely employed for automated ICH detection from medical imaging scans, such as computed tomography (CT) and magnetic resonance imaging (MRI). These models can learn to identify and localize ICH regions based on patterns and features learned from large labeled datasets.

Moreover, research efforts are focused on developing advanced algorithms to perform precise ICH segmentation, which involves delineating the exact boundaries of the bleeding. Various techniques, such as U-Net architectures, have been applied to tackle this task by leveraging annotated imaging data for training. These models can generate pixel-level segmentation masks, aiding in quantitative analysis, treatment planning, and monitoring of ICH progression.

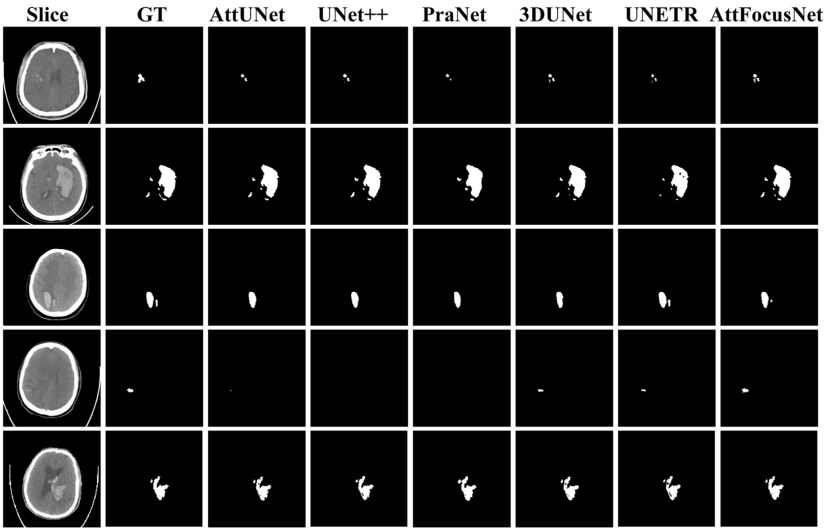

Figure 5: Comparison of the segmentation results of different segmentation methods.

Researchers in the field of ICH detection and segmentation are focused on improving the accuracy of models used for this task. They are exploring novel architectures, incorporating multiple types of medical imaging data, and utilizing advanced techniques such as attention mechanisms and graph neural networks. The aim is to develop models that are more precise and robust in identifying and delineating ICH regions in medical images.

In addition to accuracy, the explainability and interpretability of AI, models are crucial in the medical domain. Efforts are being made to enhance these aspects specifically for ICH detection and segmentation models. The goal is to enable clinicians to understand the reasoning behind the model's decisions and provide transparent explanations for the predicted results. This helps build trust in the models and facilitates their adoption in clinical practice.

Ensuring the generalizability of ICH detection and segmentation models across diverse patient populations, imaging protocols, and scanner types is a significant challenge. Researchers recognize the importance of developing models that can perform well in various scenarios. They are actively working on creating more comprehensive and diverse datasets to improve the performance and generalizability of the models. By training models on a wide range of data, they aim to enhance their ability to accurately detect and segment ICH across different situations.

Real-time ICH detection and segmentation are particularly critical in emergency situations where immediate decision-making is required. Researchers are focused on developing models that can operate efficiently on resource-constrained devices, allowing for real-time analysis of medical images. Moreover, integrating these models seamlessly into clinical workflows is a key area of investigation. The aim is to ensure that the developed models can be effectively deployed and used by healthcare professionals in real-world clinical settings.

These areas of research - improved accuracy, explainability and interpretability, generalization and dataset diversity, and real-time and clinical integration, collectively contribute to advancing the field of ICH detection and segmentation, ultimately improving patient care and outcomes in cases of intracranial hemorrhage.

3.1 The INSTANCE grand-challenge

The "2022 Intracranial Hemorrhage Segmentation Challenge on Non-Contrast head CT (NCCT)" was a specific challenge organized to address the task of segmenting intracranial hemorrhage (ICH) regions in non-contrast head CT scans. It aimed to promote the development and evaluation of advanced algorithms and techniques for automatic ICH segmentation, which is crucial for timely and accurate diagnosis and treatment planning.

The challenge provided a dataset of non-contrast head CT scans with annotated ICH regions, allowing participants to develop and train their algorithms on this data. The participants were then evaluated based on the performance of their algorithms in segmenting the ICH regions accurately. The evaluation metrics typically included measures such as the Dice coefficient, sensitivity, specificity, and other relevant metrics.

By organizing such challenges, the aim is to foster innovation and collaboration among researchers and practitioners in the field of medical image analysis. It encourages the development of state-of-the-art algorithms and techniques to improve the accuracy and efficiency of intracranial hemorrhage segmentation, ultimately benefiting patient care and outcomes.

Most of those challenges have a dual scheme. Firstly, they open it and allow participation until one date, where the results of the different teams are put in common and the best teams are selected. However, to foster more competition and innovation in the field, challenges may remain open allowing participation after the competition day has been done. Mainly, they provide the datasets and the ranking is regularly updated. Unluckily, that was not the case of the INSTANCE, that only allows access to the training dataset. Hence, even though I did not participate in the challenge, I used this dataset to perform this project.

4. Hypotheses and Objectives

4.1. Research question

My question is the following:
Can an artificial intelligence algorithm detect and segment an intracranial hemorrhage from a CT image?

4.2. Hypothesis

My hypothesis is that an artificial intelligence algorithm can detect and segment an intracranial hemorrhage from a CT image

4.3. Objective

The aim of this project is the development of an artificial intelligence algorithm, based on a convolutional neural network, that has the mission of detecting and segmenting intracranial hemorrhages in CT images.

5 Materials and Methods

5.1. Regulations and legal aspects

In order to train the model, I will use “The 2022 Intracranial Hemorrhage Segmentation Challenge on Non-Contrast head CT (NCCT)” dataset, composed by 100 training 3D volumes, each of them corresponding to a different patient, with refined labeling from 10 experienced radiologists.

All the data used for this project was completely anonymous for the sake of the patients involved and their privacy.

5.2 Data

The INSTANCE challenge dataset consisted of 100 3D CT images, although I only used 95 of them to work on the project. Since I trained a convolutional neural network from a 2D U-Net, and most of the images had no lesion, I decided to preprocess the 3D volumes and save only the slices with a big enough lesion, the idea is to obtain a model that detects better the lesions with a smaller train time.

In all the experiments, I used 80% of the data (78 patients) for training and the remaining 20% (17 patients) for testing. The algorithm also does the validation step during the training (the validation dataset is included in the training one).

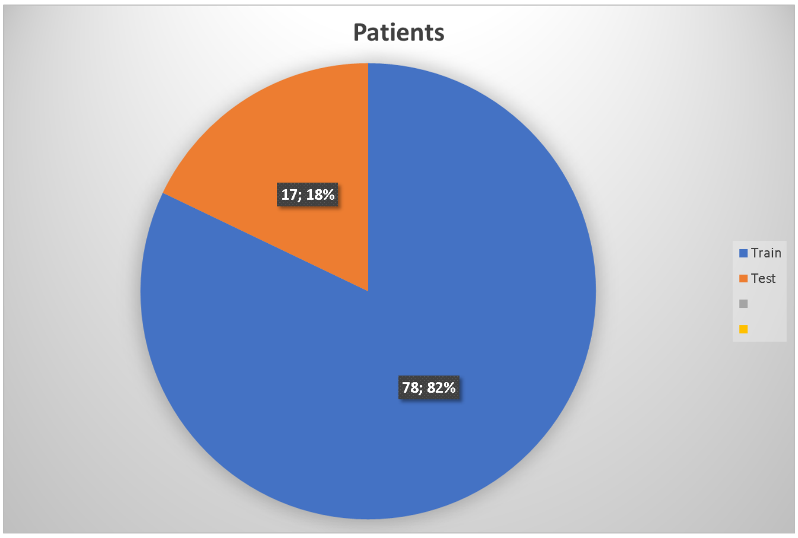

Figure 6: Dataset organization graphic

5.3 Preprocessing

The first task was to identify the lesions in the images. In order to do this, first of all I used the software ITK-Snap to visualize the images, I overlapped the main images and the masks. From this point, it was quite easy to identify the range of intensity of the lesions in the images.

Once I knew which ranges of intensity corresponded to the skull, brain, and lesion, I tried to eliminate the skull, lowered the intensity of the brain, and enhance the intensity of the lesion obtaining images where it was much easier to detect the lesion.

After this, I converted all the volumes into slices, but I only saved the slices which had a big enough lesion, this means that the images without lesions, or with really small lesions were not used to train the model, because these images could confuse the model make it detect lesions where the only thing there, is noise. All the masks of the chosen slices were also stored in order to train them.

5.4 Model training

When I had all the slices ready, I divided the slices by patients, making a separation between the training and the testing patients. As explained, 5 patients were completely excluded due to the lack of lesions in their images, 78 were used as training, and 17 patients were used as testing (80% training, 20 % testing)

From the 78 patients, 563 images were used for the training. As I said before, I only wanted to use images with big enough lesions to avoid as many false positives as possible.

I trained a CNN model with a 2D U-Net architecture. That means that I defined the U-Net model, initialized the model's parameters, and selected an appropriate loss function for the segmentation task.

5.4.1 Optimizer

An optimizer is an algorithm or method used to update the parameters of a model during the training process. It determines how the model's parameters are adjusted based on the computed gradients of the loss function. The optimizer's goal is to minimize the loss function by iteratively updating the model's parameters. Common optimizers include 'Adam', SGD with momentum, RMSprop, and Adagrad, among others.

In this case, I used the optimizer ‘Adam’ which is a popular optimization algorithm used in machine learning, particularly for training deep neural networks. It stands for Adaptive Moment Estimation. Adam combines the concepts of adaptive learning rates and momentum-based optimization.

Here are the reasons behind this choice:

→ Adaptive Learning Rates: 'Adam' adapts the learning rate for each parameter based on its historical gradients. This adaptiveness allows it to handle different magnitudes of gradients and adjust the learning rates accordingly. This can be beneficial when training complex models like U-Net, where different layers and parameters may have varying sensitivities to updates.

→ Momentum: 'Adam' incorporates momentum, which helps accelerate the learning process by maintaining a decaying average of past gradients. This can help the optimizer move faster in relevant directions, which is particularly useful when training deep models with many layers and complex feature representations.

→ Robustness and Convergence: 'Adam' has been shown to be robust and efficient in many deep learning applications. It often achieves quick convergence, especially when combined with appropriate learning rate schedules, regularization techniques, and suitable loss functions.

5.4.2 Loss

The loss function, also known as the objective or cost function, evaluates how closely the model's predictions align with the true target values during training. It measures the discrepancy between the predicted outputs and the actual targets. By providing a single scalar value, the loss function quantifies the model's performance on the training data, given a specific set of parameters.

For U-Net models employed in image segmentation tasks, a commonly used loss function is pixel-wise binary cross-entropy. This loss function compares each predicted pixel value to its corresponding target pixel value and calculates the cross-entropy loss. It is particularly suitable for evaluating the accuracy of pixel-level segmentation.

Alternatively, other loss functions such as Dice coefficient loss or Jaccard loss can also be utilized in U-Net models, depending on the specific requirements and characteristics of the segmentation task. These alternative loss functions offer different properties and may be better suited for certain segmentation scenarios.

In this project, I used the binary cross-entropy because is the most commonly used.

5.4.3 Validation Metrics

In a U-Net model or any other machine learning model, metrics are evaluation measures used to assess the performance of the model during training and evaluation.

There are different metrics good for Unet training, such as IoU, Dice coefficient, Precision and Recall, and Accuracy. In this case, I used the accuracy, but the Dice coefficient could have been also a good choice.

5.5 Testing

For testing the algorithm I used the 17 patients' 3D images. We first converted all the slices to 2D images that will be the input to the algorithm, obtaining a total of 523 images. We have to take into account that the vast majority of the images do not have lesions, but they can present false positives and hence, decrease the performance of the algorithm.

5.5.1 Evaluation Metrics

In order to evaluate the results I reconstruct the 3D images using the predicted slices. Subsequently, I calculate the Dice coefficient and the percentage of false positives between the Ground Truth volume and the predicted 3D volume for each patient.

The Dice coefficient, also known as Dice similarity coefficient or Dice accuracy, is a statistical measure, in my case, used to evaluate the similarity between two images.

Its coefficient formula $Dice = (2 * TP)/(2 * TP + FP + FN)$ is derived from the concept of precision and recall, where TP represents the true positives, FP represents the false positives, and FN represents the false negatives.

→ TP refers to the number of true positive predictions, which are the cases where the predicted label and the true label are both positive.

→ FP refers to the number of false positive predictions, which are the cases where the predicted label is positive, but the true label is negative.

→ FN refers to the number of false negative predictions, which are the cases where the predicted label is negative, but the true label is positive.

This results in a value between 0 and 1, where a value of 1 indicates perfect overlap or similarity between the predicted and true labels, and a value of 0 indicates no overlap, in my case I used percentages.

In terms of region-based, we evaluate the number of false positives per image. It is a similar analysis of the Dice, but based on regions. A region is considered a TP if there is minimum overlap between the region in the test image and the lesion in the ground-truth. Otherwise, will be considered as a FP. And consequently for TN and FN.

All the measures are computed according the 3D reconstruction of the segmentation.

6 Results

After preparing all the processes, it’s time to evaluate the results, both from preprocessing and the segmentation of the model.

In the preprocessing, the most important goal was to remove the skull in order to facilitate the model job, meanwhile, the model had the task to segment correctly the lesions starting from the pre-processed images.

6.1 Preprocessing

The first results I got, were in the post-processing section, where I make slices of each 3D image, and I try to remove the skull and enhance the lesion.

Now, I will display examples of the results showing the evolution of the image from which I started to the preprocessed image

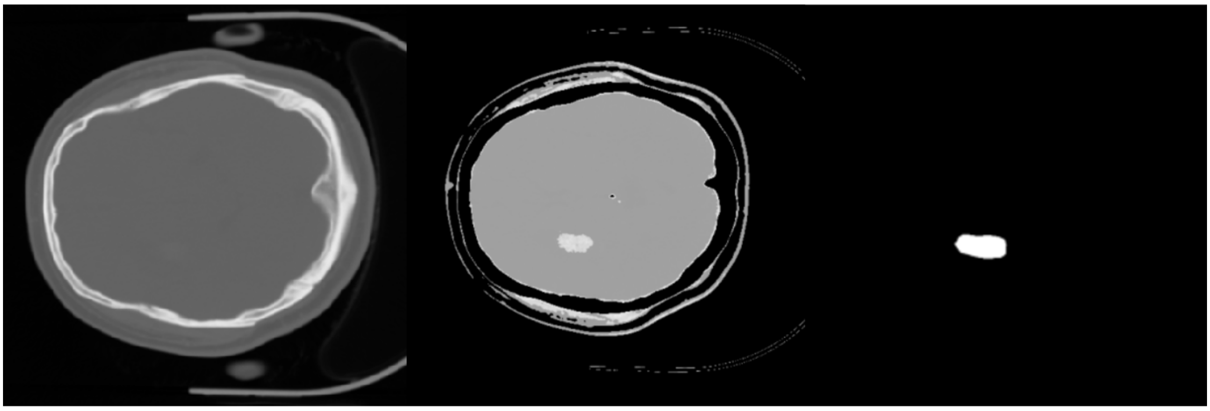

Figure 7: Good preprocessing example

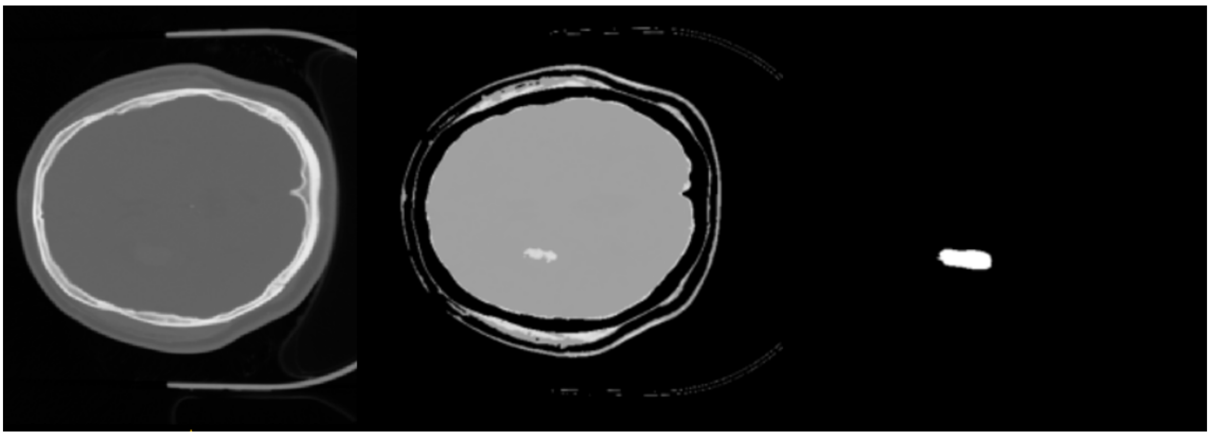

Figure 8: Poor preprocessing example

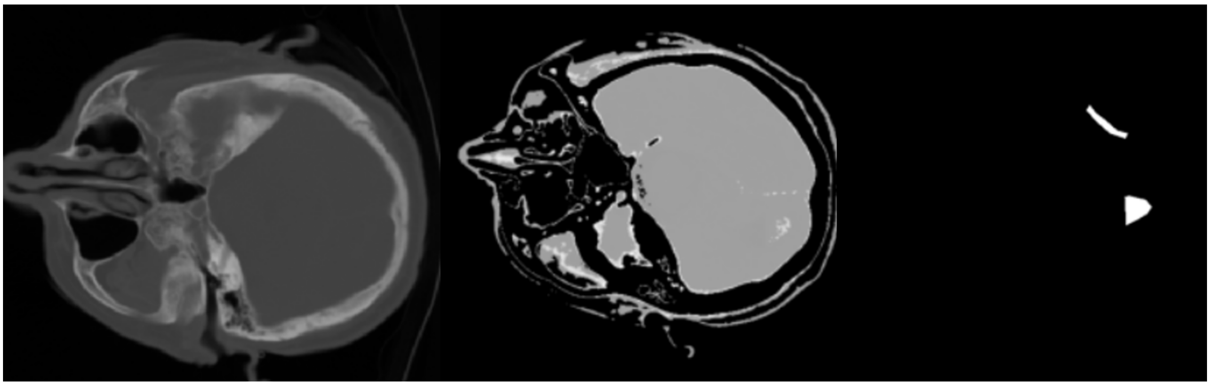

Figure 9: Bad preprocessing example

The task of removing the skull seems to be well done since it removes almost the entire skull in all images, this means that since the intensity of the skull is much higher than the intensity in the brain tissue, is really easy to eliminate through thresholding.

The problem comes when enhancing the lesion, here the usage of thresholding doesn't seem to be the most consistent idea, because each lesion from each patient has different intensity values, and since these values are not as far between each other as the distance between the skull and brain tissue, it's really hard to delimit which intensity range belongs to injuries and not to healthy tissue.

The first thing that we have to take into account is that the preprocessing of the lesions like I did is not a consistent preprocessing, because not always archives the goal of enhancing the lesion. Even though, it is pretty robust when it comes to removing the skull.

This happens due to the variability between lesions of different patients as I mentioned earlier. Since the way I enhance the lesion is by using a range of intensities if the lesion is not within the range I specify, it will not be enhanced, and if you create a range that is too wide, it will enhance any part of the brain regardless of its state.

Adding to this the variability between the lesions of the patients in terms of lesions and healthy tissue makes the task incredibly hard. For example, we can see in Figure 9 how I enhance parts of the brain that are not a lesion and enhance almost nothing of the target. On the other hand, there are lesions, like the one that I show in Figure 7, where the lesion is almost perfectly enhanced.

The problem when not delimitating enough the intensity range that belongs to the injuries is that it can lead to a result like the following.

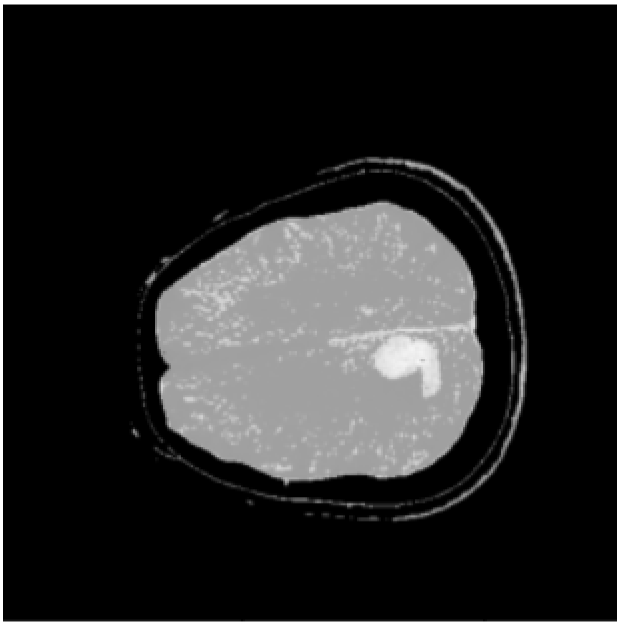

Figure 10: The intensity range belonging to the lesion in this image was too large, making the preprocessing wrong.

Here in Figure 10, we can see how due to the large range of intensity used to detect the lesions, I ended up enhancing part of the healthy tissue of the brain, making it harder for the model to correctly identify the lesions.

With these results, although the enhancement of the lesion is mostly good and the main objective of removing the skull to facilitate the training of the U-net has been achieved, there are some false positives.

That is why I would qualify this preprocessing as sufficient for the task at hand.

6.2 Model results

Once I trained the model with the slices of lesions of 78 patients I began with the testing. For that, as I already showed to you, I used all the slices of the other 17 patients.
This means that I will be having slices without lesions, thing that I didn't have in the training, but this shouldn't be a problem since the model learns from positive and negative inputs, and all the slices that I used to train, even the ones with bigger lesions, have regions without lesions (negatives), this should help the model understand when there shouldn't be a lesion.

My dice for this test was Dice = 0.5482 ± 0.2740, and the FP = 12.5 ± 7.35

Here I will display some of my results.

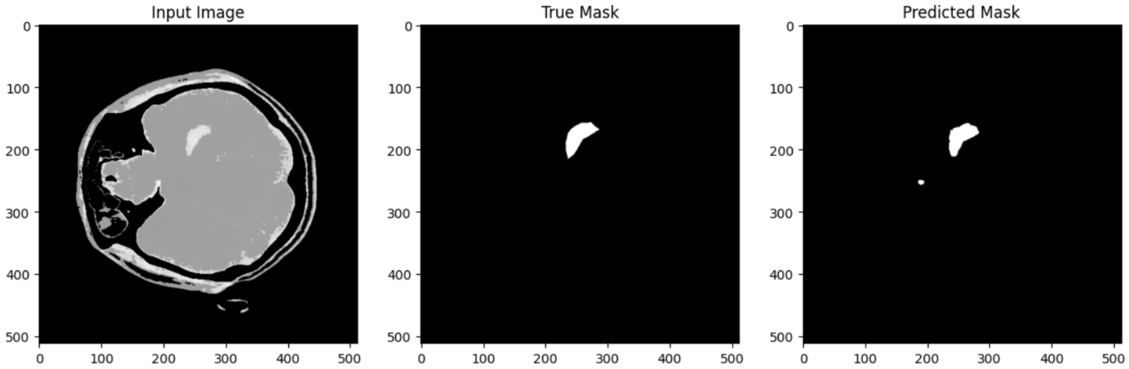

Figure 11: Acceptable example of segmentation

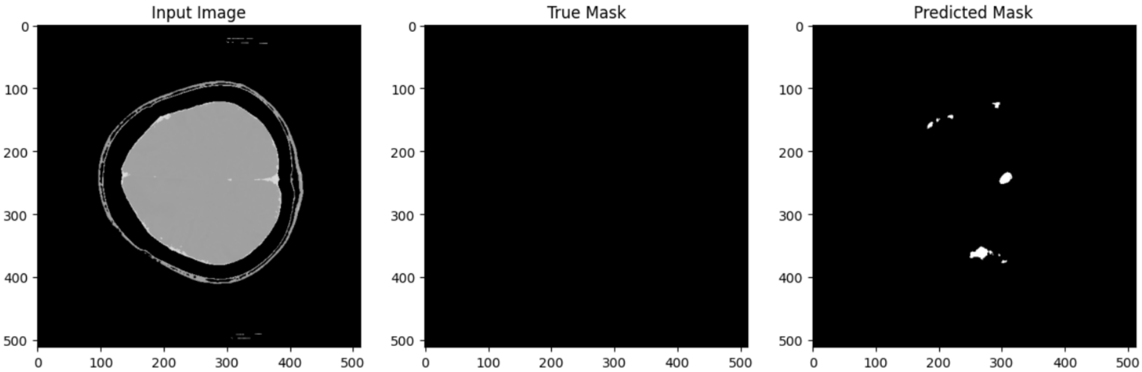

Figure 12: Bad example of segmentation false negatives

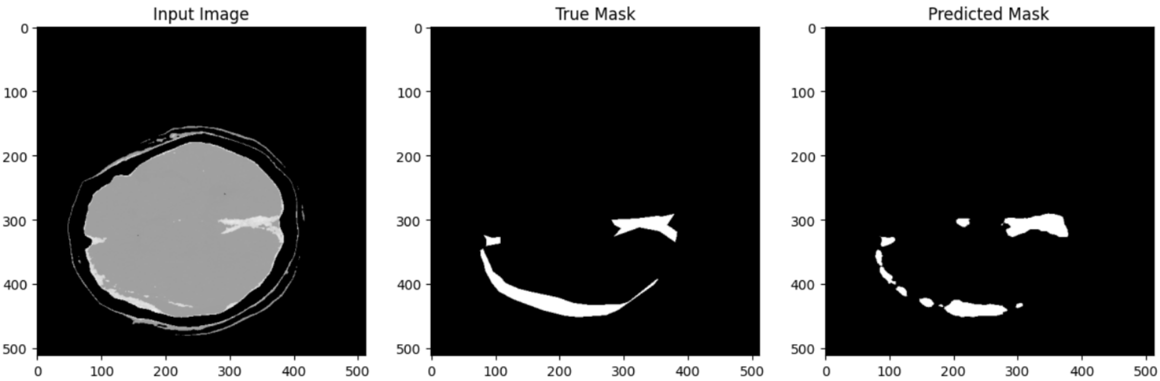

Figure 13: Example of an elongated lesion at the edge of the skull not rightfully segmented.

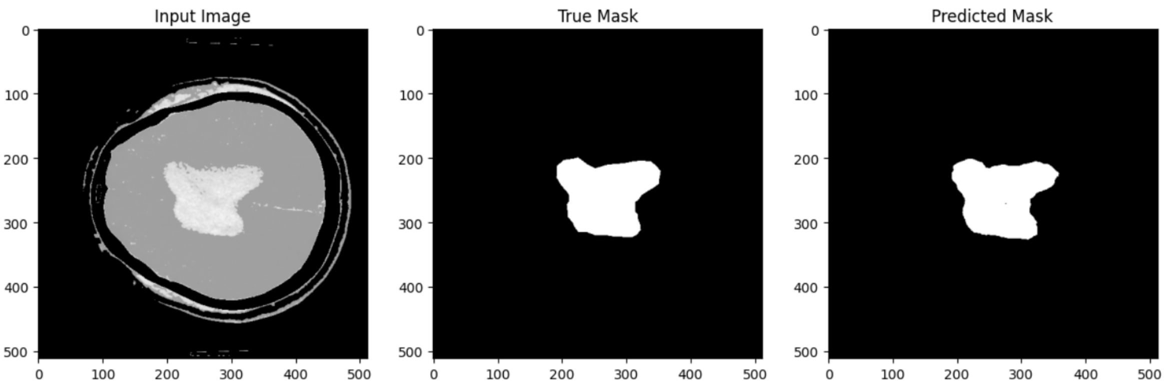

Figure 14: Example of segmentation of a big lesion without elongated shape

From my point of view, the dice vary depending on the type of lesion, with that, I am referring to the position, shape, and size. From what I've seen, the lesions with a similar size or bigger than the one we can see in Figure 11, often have a good result, but there are often false positives with a similar size to these lesions in a lot of slices.

Continuing with the false positives, as we can see in Figure 12, there are slices where the model detects lesions. These lesions are often really small, and they can be easily post-processed, but making the sacrifice of losing the detection of some small lesions.

It also gives me the impression that in the event that the image has a lesion, the algorithm is less likely to show as many false positives as in images where there is no lesion.

The position and shape seem to be something to take into account. Thus, the more centralized and circular lesions are easier to detect, and the more externalized and elongated lesions are more complex to segment. As we can see in Figure 13 the lesion is divided in smaller lesions with more circularity, meanwhile, in Figure 14, the lesion is perfectly segmented.

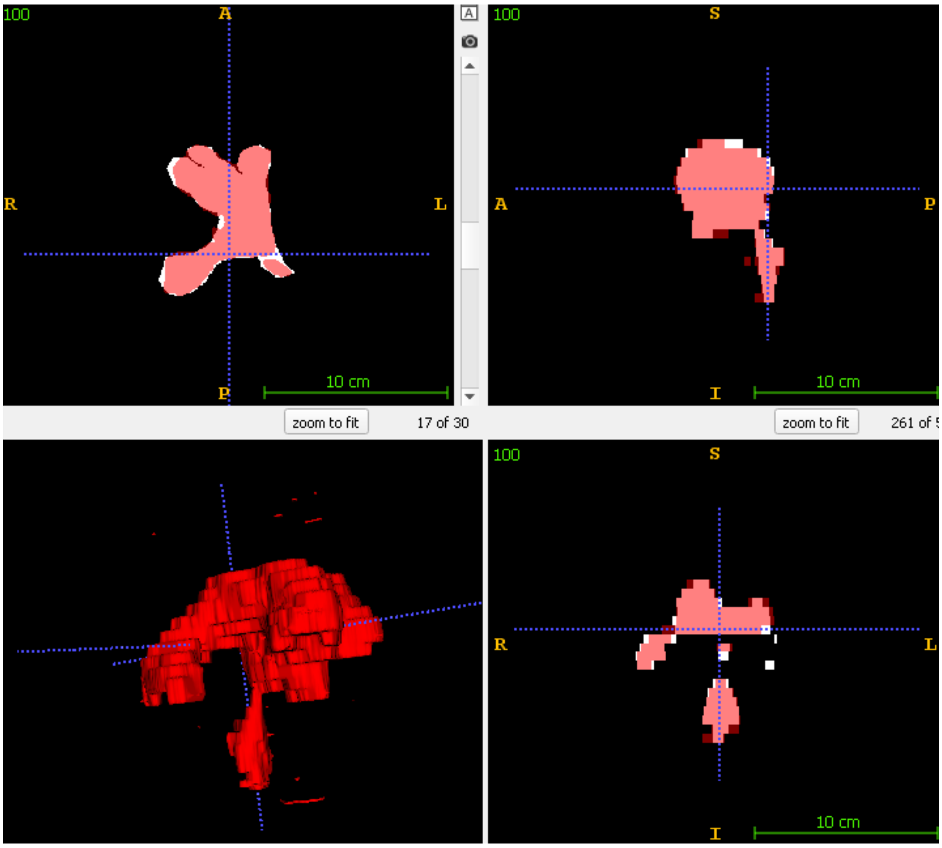

Figure 15: Segmentation reconstruction visualization with ITK-Snap of a good segmentation

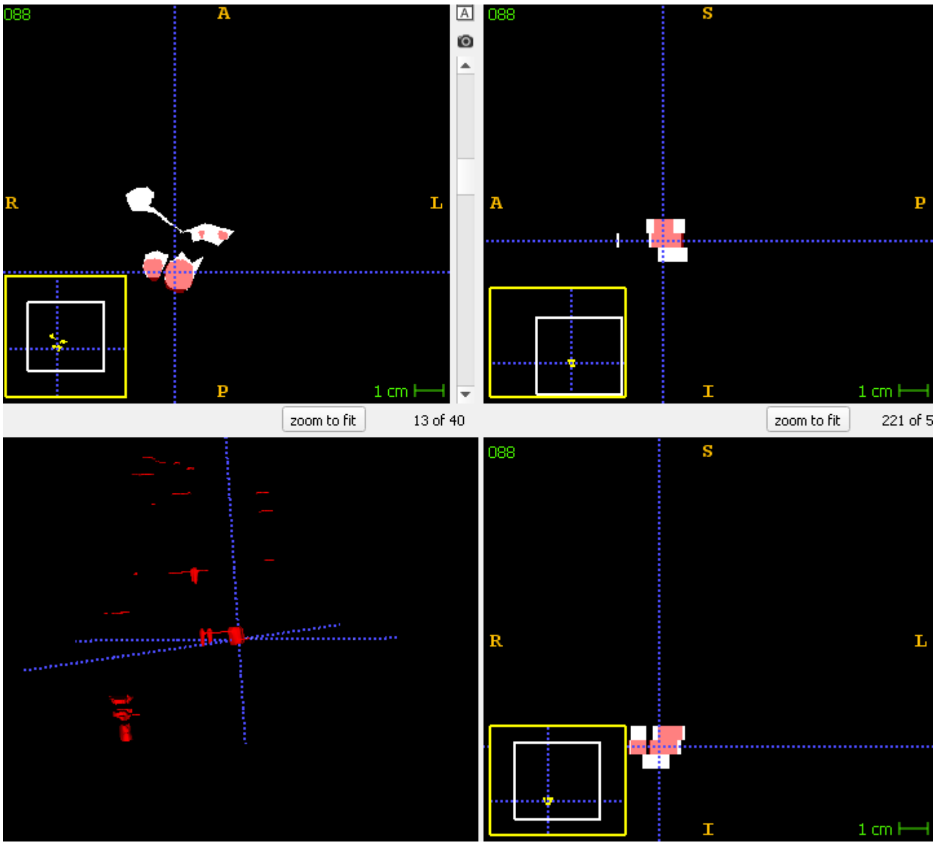

Figure 16: Segmentation reconstruction visualization with ITK-Snap of a bad segmentation

For example, as you can see in Figure 15 the reconstruction of the lesion is very similar to the real lesion and there are not a lot of false positives.

Meanwhile, in Figure 16, the reconstruction is not even close to the real segmentation, from my point of view this is due to the small shape of the lesion.

Challenge comparison

In the challenge, the results obtained from position 1 to 296 in terms of Dice were from 0.7953 to 0.4132 and the deviation standard variated from 0.1713 to 0.3761, not the best Dice had the best deviation standard and nor vice versa, but it was very close. Based on my Dice score of 0.5482 ± 0.2740, I would be ranked around 284. However, there is not a smaller standard deviation since place 244.

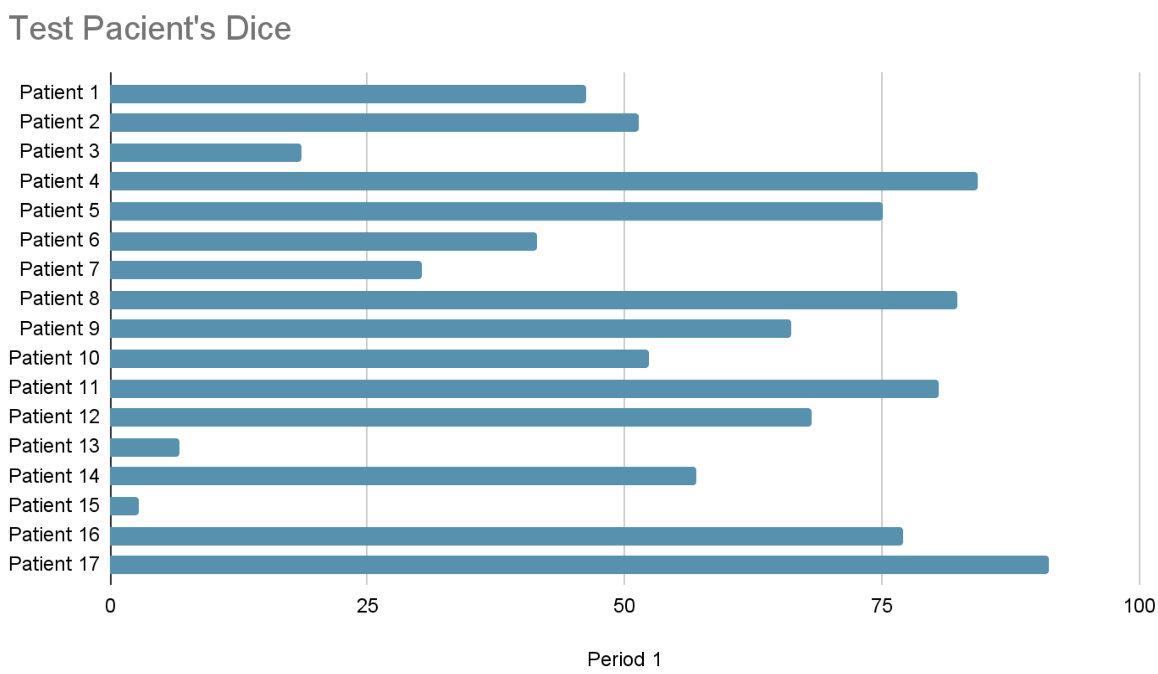

Figure 17: Bar graph representing the Dice coefficient of the final segmentation of each patient in the test

In the previous bar graph we can see that even though the mean Dice is 54,82%, 64% of the patients score above 50%, and 35% of these above the 75% Dice score, leaving only a 17% of patients below the 25% of Dice Score.

Taking into account that the model makes big mistakes when segmenting the patient 13 and 15, if I recalculate the dice without this patients, the mean Dice amounts to a 61.48% with a estandard deviation of 21.28%. This would place the algorithm at the 266th position in the challenge.

In terms of false positives, even though at first glance they seem more numerous, these are the results that I have obtained.

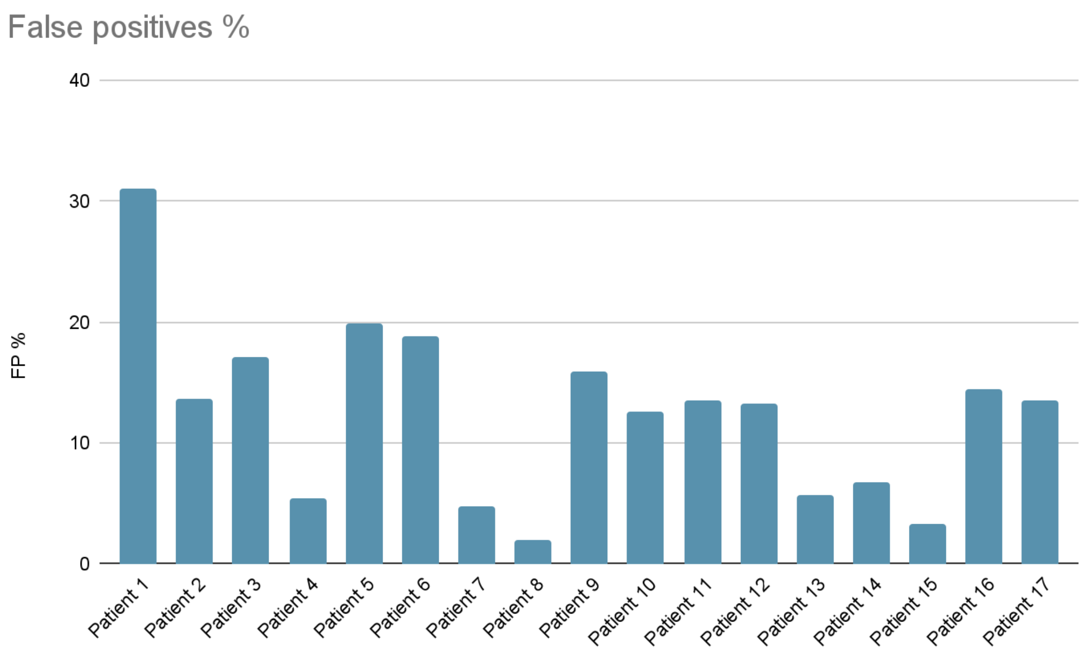

Figure 18: Bar graph representing the percentage of false positives for each patient

Here we can see the percentage of false positives with a mean of 12.5%, this is relatively high taking into account that a false positive rate of 12.5% means that approximately 12.5% of the algorithm's predictions would be incorrect positives, indicating the presence of an intracranial hemorrhage when there isn't one.

While it's important to balance the trade-off between sensitivity and specificity in medical imaging applications, a false positive rate of 12% might be considered too high for many clinical scenarios. In ICH segmentation, false positives can lead to unnecessary interventions, increased healthcare costs, and patient anxiety.

7 Discussion

In this project we look for developing an artificial intelligence algorithm utilizing a convolutional neural network. The primary objective is to detect and accurately segment intracranial hemorrhages in CT images.

7.1 Implementation details

Regarding the implementation steps done, I organized the data using the INSTANCE 2022 dataset, which consisted of 100 test CT scans along with their corresponding masks, as well as 30 CT scans for the challenge evaluation. Since I couldn't use the 30 evaluation scans for testing my model, I excluded them from further processing.

For storing the data, I utilized Google Drive since I didn't have access to a GPU for model training. Although it was a bit challenging due to my internet connection and server limitations, I appreciated the opportunity to work within this environment.

Next, I tackled the issue of 3D volumes having numerous slices without lesions or with very small lesions. To address this, I developed a Python script that divided the volumes into individual slices, analyzed them, and selectively stored only the slices with sufficiently large lesions.

Once the data was curated, I proceeded to split it into training and testing sets for model training. During this process, I encountered a challenge where some images from the same patient appeared in both the training and testing sets. This raised concerns of potential bias and cheating, as the model could exploit similarities in lesions across different slices of the same patient. Moreover, since I didn't test the model with all slices of the volumes, I couldn't assess its performance in the absence of lesions.

To mitigate these issues, I implemented another data division strategy based on patient numbers. I ensured that all slices belonging to a specific patient were included exclusively in the testing or training set, but not both.

In terms of preprocessing, my primary goal was to remove the skull to aid in model training. To achieve this, I utilized thresholding, which proved to be an effective technique. However, I also wanted to enhance the lesions for better segmentation results. To accomplish this, I carefully examined the majority of the images to determine the intensity range of the lesions. This range was then utilized during the preprocessing step for the lesions.

The model training itself was relatively straightforward, and my main concern was the execution time. Fortunately, by utilizing a GPU on Google Colab, the training time was reduced to approximately 3 hours, allowing for efficient model development.

7.2 Results discussion

I obtained a Dice = 0.5482 ± 0.2740 which would be position 284, and the number of false positives of FP = 12.5 ± 7.35. However, after excluding patients 13 and 15, the Dice coefficient improved to 0.6148 ± 0.2128, placing the algorithm at position 266.

In my UNET model, the imbalanced class distribution between ICH and non-ICH instances posed a challenge. The rarity of ICH instances compared to non-ICH instances led to a bias towards the majority class, resulting in a higher number of false positives. To address this, I recognize that including non-ICH images in the training data could have helped improve the model's accuracy. By incorporating a more balanced representation of both classes, the model may have been able to learn more effectively and achieve better performance. In my testing, when the model was evaluated only on ICH images and not the entire range of patient images, the accuracy reached around 60%.

Additionally, I identified the limited capture of global or long-range dependencies as a limitation in the UNET model architecture I utilized. The encoder-decoder structure of UNET has a restricted receptive field, which means it may struggle to capture relationships between distant regions in the image. This lack of contextual information can hinder accurate segmentation, as it is essential to understand how different parts of the brain relate to each other in the presence of ICH. To overcome this limitation, I recognize the potential of utilizing a 3D U-net architecture. By incorporating the additional dimension of depth, a 3D U-net can capture spatial dependencies more effectively and provide a better understanding of the entire volume of brain images. This approach would likely improve the model's ability to capture global or long-range dependencies and subsequently enhance segmentation accuracy for ICH lesions.

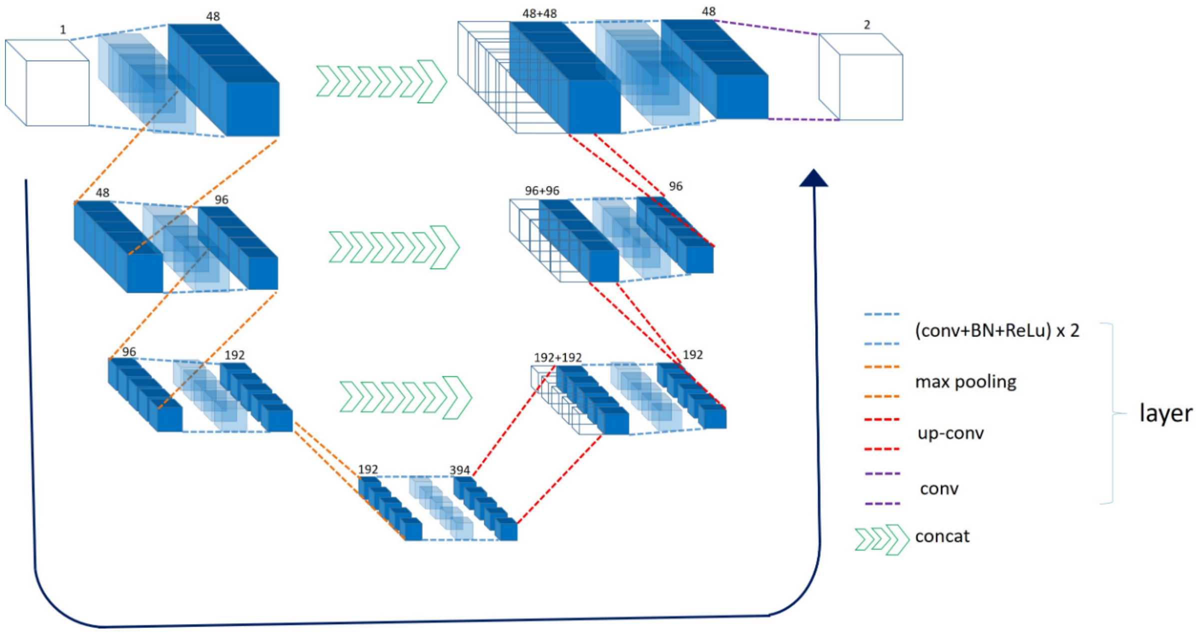

Figure 19: This is the scheme of a 3D U-net

7.3 Limitations

When using a 2D U-Net for intracranial hemorrhage (ICH) detection, there are certain limitations to consider. The 2D U-Net architecture, while effective for image segmentation tasks, has inherent limitations that can impact its performance in detecting ICH accurately.

One significant limitation is the lack of consideration for volumetric information. The 2D U-Net operates on individual 2D image slices independently, disregarding the spatial context and information present in the neighboring slices. ICH can exhibit complex shapes and extend across multiple slices, making it challenging for a 2D U-Net to capture the full extent of the hemorrhage accurately.

Additionally, the reliance on 2D slices limits the ability of the model to capture depth information. ICH detection can benefit from the analysis of multiple image slices along the z-axis, which provides important spatial cues for accurate segmentation. By working with 2D slices only, the 2D U-Net may struggle to differentiate ICH from other structures or artifacts, potentially leading to false positives or negatives.

In addition, this model has only been trained to detect lesions from CTs done with a unique machine, nowadays there are a lot of different CT machines that can do a CT and the result will be different, which makes the model not robust.

Regarding the materials needed to carry this project forward, I have to say that since Google Colab provides GPU service, I used their servers and stored all the data in Google Drive. Although this made the job a bit more tedious due to the internet connection and delay between instructions, plus errors due to packet loss, this was the most affordable option for me since I couldn't use my own GPU for training the model, and this could have greatly lengthened the testing period.

About the knowledge necessary to carry out this project, I had no background in deep learning and the use of a U-net. That is why I took some courses on artificial intelligence and deep learning to introduce myself to the project “Neural Networks and Deep Learning” and “Convolutional Neural Networks” from DeepLearning.AI. This caused the project schedule to slow down more than expected as well but provided me with a good background to start with.

7.4 ODS Contributions

The sole objective of this project is the third ODS, to ensure the achievement of a healthy life and promote well-being for people of all ages.

This is because the proper implementation of this project could help doctors to detect faster, and more accurately the ICH of a patient, opening the possibility to a faster operation, which potentially could save the life of the patient.

8 Conclusions

In conclusion, this project successfully accomplished the primary objective of segmenting intracranial hemorrhage (ICH) lesions using a U-net model, yielding a mean Dice score of 0.5482 ± 0.2740.

The preprocessing step effectively achieved its goal of skull removal, although it did not significantly contribute to the easy detection of lesions due to the limitations of using only thresholding. While thresholding proved to be a quick technique for identifying lesions, it did not always produce precise lesion shapes. However, it successfully identified the presence of lesions.

The trained U-Net model demonstrated promising performance in cases involving large, circular, and centrally located lesions, resulting in moderate Dice scores. However, its performance was suboptimal for cases involving small, elongated, and non-centrally located lesions. To address this limitation, post-processing techniques could be employed to further improve the segmentation results. By incorporating advanced algorithms for refinement and contouring, the accuracy and precision of lesion segmentation could be enhanced.

These findings highlight the challenges in accurately detecting ICH lesions using the proposed approach. Further research and improvements are necessary to enhance the model's performance, particularly in identifying smaller and non-centrally located lesions. Additionally, alternative techniques or modifications to the preprocessing step could be explored to better facilitate lesion detection. Overall, this study provides valuable insights into the strengths and limitations of the thresholding-based skull removal and U-Net-based lesion detection approach for ICH segmentation.

8.1 Possible improvements

This project indeed presents several areas for improvement. One of the crucial enhancements would be the adoption of a 3D U-net architecture. By incorporating volumetric information, the model would gain a better understanding of the three-dimensional context of the brain. This would enable more accurate detection of lesions, as the model could leverage the volumetric relationships and spatial dependencies between slices. The utilization of a 3D U-net could potentially reduce the number of false positives, resulting in improved segmentation performance.

Regarding preprocessing, handling the variations in intensity range across different patients can be challenging. One possible approach to enhance lesion visibility is to personalize the intensity range for each patient. This can be achieved by calculating the mean intensity of each patient's brain tissue and mapping the pixel intensities accordingly. By normalizing the intensity range based on individual patient characteristics, it would be possible to improve the visibility of lesions.

Additionally, incorporating post-processing techniques can be beneficial in reducing false positives. After the initial segmentation, applying methods like morphological operations (e.g., erosion, dilation) or connected component analysis can help refine the results and eliminate spurious detections. Postprocessing steps enable the model to have a better delineation of lesions and improve the overall accuracy of the segmentation.

In conclusion, by implementing a 3D U-net architecture, personalizing the intensity range during preprocessing, and incorporating post-processing techniques, significant improvements can be made in the accuracy and robustness of the ICH segmentation model. These enhancements would address some of the major limitations and challenges encountered in the project, potentially leading to more precise and reliable results.

9 Bibliography

1. An, S. J., Kim, T. J., & Yoon, B.-W. (2017). Epidemiology, risk factors, and clinical features of intracerebral hemorrhage: An update. Journal of Stroke, 19(1), 3–10. https://doi.org/10.5853/jos.2016.00864

2. Flaherty, M. L., Haverbusch, M., Sekar, P., Kissela, B., Kleindorfer, D., Moomaw, C. J., Sauerbeck, L., Schneider, A., Broderick, J. P., & Woo, D. (2006). Long-term mortality after intracerebral hemorrhage. Neurology, 66(8), 1182–1186. https://doi.org/10.1212/01.wnl.0000208400.08722.7c

3. Hemphill, J. C., III, Bonovich, D. C., Besmertis, L., Manley, G. T., & Johnston, S. C. (2001). The ICH score: A simple, reliable grading scale for Intracerebral hemorrhage. Stroke; a Journal of Cerebral Circulation, 32(4), 891–897. https://doi.org/10.1161/01.str.32.4.891

4. Hssayeni, M. D., Croock, M. S., Salman, A. D., Al-khafaji, H. F., Yahya, Z. A., & Ghoraani, B. (2020). Intracranial hemorrhage segmentation using a deep convolutional model. Data, 5(1), 14. https://doi.org/10.3390/data5010014

5. Hu, Y., & Zheng, Y. (2019). A GLCM embedded CNN strategy for computer-aided diagnosis in intracerebral hemorrhage. En arXiv [cs.CV]. http://arxiv.org/abs/1906.02040

6. Islam, M., Sanghani, P., See, A. A. Q., James, M. L., King, N. K. K., & Ren, H. (2019). ICHNet: Intracerebral Hemorrhage (ICH) Segmentation Using Deep Learning. En Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries (pp. 456–463). Springer International Publishing.

7. Kaur, A., Kaur, L., & Singh, A. (2021). GA-UNet: UNet-based framework for segmentation of 2D and 3D medical images applicable on heterogeneous datasets. Neural Computing & Applications, 33(21), 14991–15025. https://doi.org/10.1007/s00521-021-06134-z

8. Li, X., Luo, G., Wang, K., Wang, H., Liu, J., Liang, X., Jiang, J., Song, Z., Zheng, C., Chi, H., Xu, M., He, Y., Ma, X., Guo, J., Liu, Y., Li, C., Chen, Z., Siddiquee, M. M. R., Myronenko, A., … Li, S. (2023). The state-of-the-art 3D anisotropic intracranial hemorrhage segmentation on non-contrast head CT: The INSTANCE challenge. En arXiv [eess.IV]. http://arxiv.org/abs/2301.03281

9. Mehendale, D. N., Gupta, P., Rajadhyaksha, N., Dagha, A., Hundiwala, M., Paretkar, A., Chavan, S., & Mishra, T. (2022). A graphical approach for brain haemorrhage segmentation. En arXiv [eess.IV]. http://arxiv.org/abs/2202.06876

10. Montaño, A., Hanley, D. F., & Hemphill, J. C., 3rd. (2021). Hemorrhagic stroke. Handbook of Clinical Neurology, 176, 229–248. https://doi.org/10.1016/B978-0-444-64034-5.00019-5

11. Nguyen, N. T., Tran, D. Q., Nguyen, N. T., & Nguyen, H. Q. (2020). A CNN-LSTM architecture for detection of intracranial hemorrhage on CT scans. En bioRxiv. https://doi.org/10.1101/2020.04.17.20070193

12. Rajpurkar, P., Chen, E., Banerjee, O., & Topol, E. J. (2022). AI in health and medicine. Nature Medicine, 28(1), 31–38. https://doi.org/10.1038/s41591-021-01614-0

13. van Asch, C. J., Luitse, M. J., Rinkel, G. J., van der Tweel, I., Algra, A., & Klijn, C. J. (2010). Incidence, case fatality, and functional outcome of intracerebral haemorrhage over time, according to age, sex, and ethnic origin: a systematic review and meta-analysis. Lancet Neurology, 9(2), 167–176. https://doi.org/10.1016/S1474-4422(09)70340-0

14. Yu, N., Yu, H., Li, H., Ma, N., Hu, C., & Wang, J. (2022). A robust deep learning segmentation method for hematoma volumetric detection in intracerebral hemorrhage. Stroke; a Journal of Cerebral Circulation, 53(1), 167–176. https://doi.org/10.1161/STROKEAHA.120.032243

ANNEXES

ANNEXE A.

PLANIFICATION

Step 1. Project Initiation and Preparation

Estimated duration: 10 weeks

1. Define the project and its objectives.
2. Acquire a solid background in deep learning and convolutional neural networks.
3. Determine the appropriate software and programming language for the project.
4. Research the state-of-the-art techniques and methodologies in the field.
5. Obtain the necessary dataset from the challenge organizers.

During the initial phase of the project, I collaborated with my tutor to define the project's objectives and hypothesis. However, this process took some time as the data required for the project was not readily available, and I had to request the dataset from the challenge organizers.

While waiting for the dataset, I recognized the importance of having a strong foundation in deep learning and convolutional neural networks (CNNs) to successfully tackle the project. Therefore, I proactively sought guidance from my tutor and decided to enhance my knowledge in this area. To accomplish this, I enrolled in two relevant courses on Coursera: "Neural Networks and Deep Learning" and "Convolutional Neural Networks."

Simultaneously, we discussed and finalized the software and programming language to be utilized for the project. This decision was crucial for ensuring seamless implementation and compatibility with the chosen deep learning frameworks.

Furthermore, I conducted extensive research to gain insights into the state-of-the-art techniques and methodologies employed in similar projects. This helped me acquire a comprehensive understanding of the current advancements and best practices in the field.

Once I received confirmation from the challenge organizers regarding the availability of the image dataset, I proceeded to download it, marking an essential milestone in the project's preparation phase.

Step 2. Image Distribution and Preprocessing

Estimated duration: 3 weeks

1. Convert 3D images to slices.
2. Select suitable images for training the model.
3. Divide the dataset into training and testing sets.
4. Apply preprocessing techniques such as skull-stripping and contrast enhancement.
5. Prepare the dataset for model training and segmentation.

In this phase of the project, I focused on the distribution and preprocessing of the images to prepare them for training the model. The following tasks were performed:

Conversion of 3D Images to Slices: Since the original dataset consisted of 3D volumetric images, I developed a script to divide these volumes into individual slices. This allowed for a more granular analysis and processing of the images.

Selection of Training Images: Considering the objective of the project, I carefully selected the images that would be used for training the model. This involved identifying and prioritizing the slices that contained substantial lesions, as these were crucial for training the model to accurately segment intracranial hemorrhages (ICH).

Division into Training and Testing: To evaluate the performance of the trained model, I divided the dataset into training and testing sets. This ensured that the model could be trained on a sufficient amount of data while still having unseen samples for evaluation.

Preprocessing Techniques: To enhance the quality of the images, I applied preprocessing techniques. Firstly, I implemented a skull-stripping method to remove the skull from the images, which aided in isolating the regions of interest and reducing noise. Additionally, I employed contrast enhancement techniques to improve the visibility and distinction of the lesions, facilitating their accurate segmentation.

Step 3. Model Training

Estimated duration: 3 weeks

1. Organizing the training and testing images
2. Training the model

During the model training, I focused on setting up the necessary folders and directories to streamline the training process. This involved creating dedicated folders for organizing training images and testing images. After this the only thing left was training the model.

Step 4. Discussion

Estimated duration: 2 weeks

1. Analyzing the results of the trained model.
2. Drawing conclusions based on the results.
3. Discussing the performance of the model.
4. Identifying the strengths and weaknesses of the model.
5. Identifying the limitations of the study.

During the discussion phase, I analyzed the results, drew conclusions, and discussed the model's performance, strengths, weaknesses, and limitations.

Step 5. Memory

Estimated duration: 2 weeks

1. Memory redaction

The report drafting was completed as the final step in chronological order

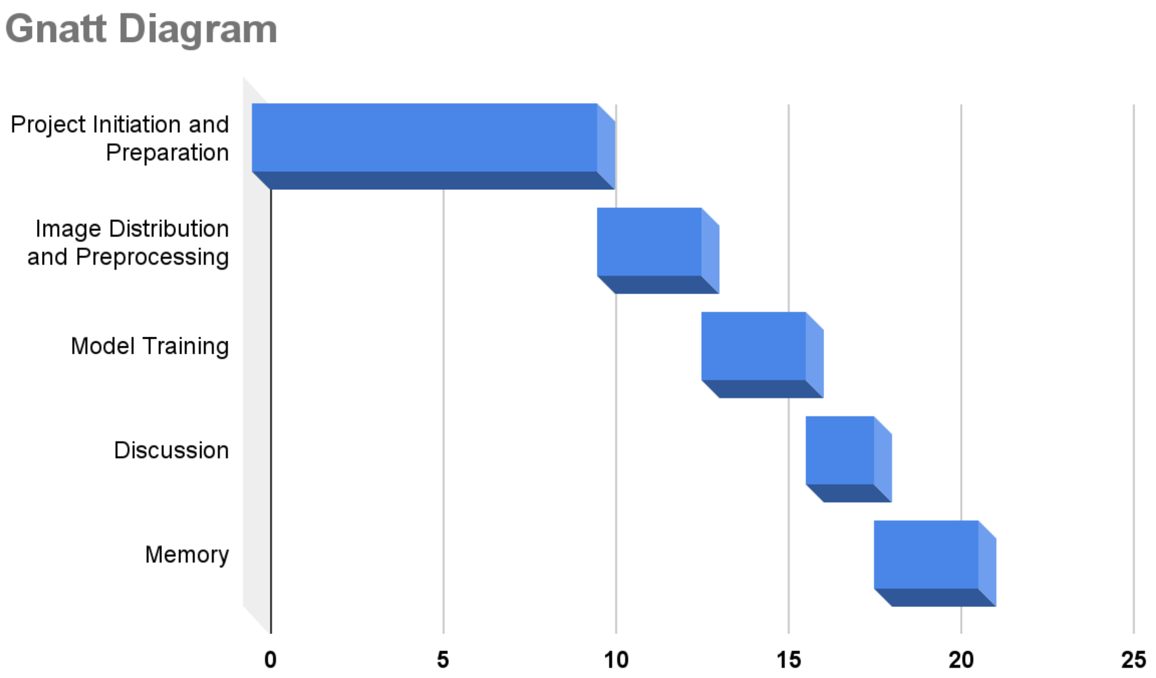

Figure 20: Gnatt diagram representing the project duration

ANNEXE B.

CODE

Because this project includes code, a GitHub repository has been created to make it accessible to the scientific community. You can find the repository at the following link: https://github.com/xBofi/Segmentation-Method-for-Intracerebral-Hemorrhage

ANNEXE C.

PRESSUPOST

In this section I will show the estimated costs of carrying out this project.

C.1. Labour

Considering that the final degree project entails a substantial investment of approximately 375 hours of dedicated effort and the remuneration for an artificial intelligence engineer is estimated to be around 40 €/hour.

Staff	Salary	Total
Artificial intelligence engineer	40 €/hour	15.000 €

C.2. Materials

To complete the project I used an HP Pavillion Laptop with a Windows OEM key.
The programming software was free to use.

Concept	Salary	Total
Hardware	HP Pavillion Laptop	900 €
Software	Windows OEM Key	19 €

C.2. Total

Description	Hours	Price/Hour	Quantity	Total
Artificial intelligence engineer	375 hours	40 €/hour	1	15.000
HP Pavillion Laptop	-	-	1	900 €
Windows OEM Key	-	-	1	19 €
				15.919 €