CNN for Brain Hemorrhage Classification

Tanmay Basu

Head, Department of Data Science and Engineering
IISER Bhopal

tanmay@iiserb.ac.in

AI for Brain Hemorrhage Detection

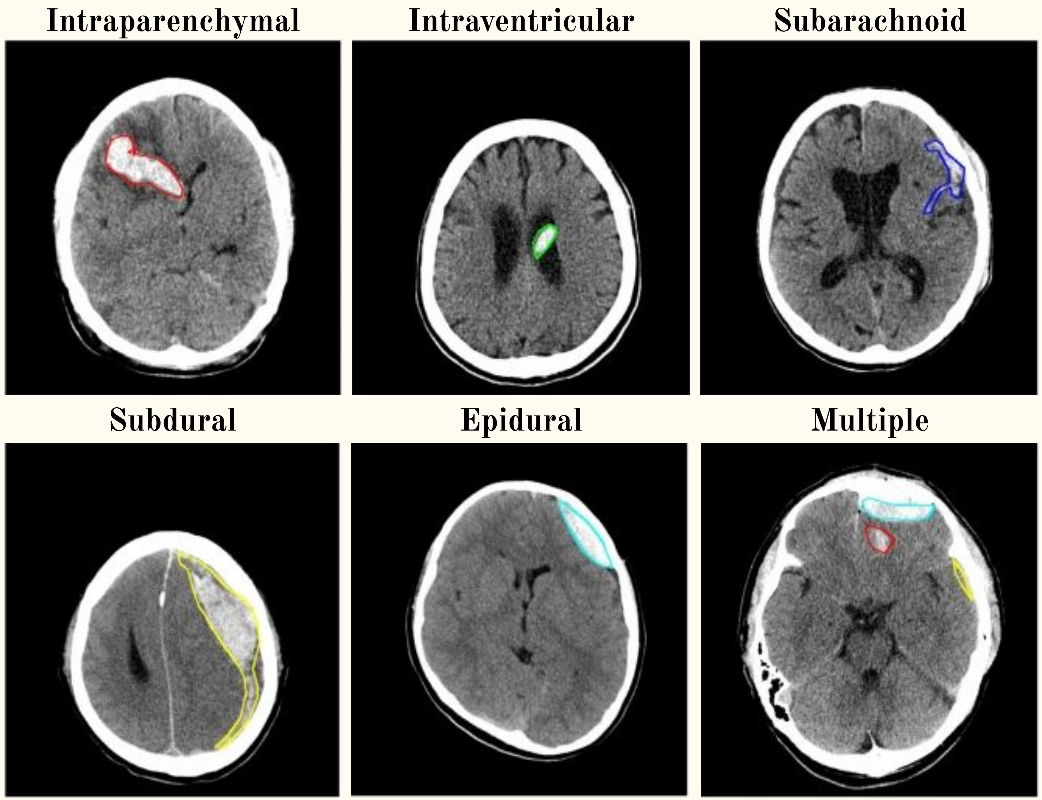

Image Source: https://www.livescience.com/6344-brain-hemorrhage.html
https://www.nature.com/articles/s41598-023-33775-y

Last accessed on May 07, 2025

Overview of Brain Hemorrhage Detection

- **Computed Tomography (CT)** of the head plays a vital role in the rapid assessment of brain hemorrhage.
- **Convolutional Neural Networks (CNNs)** are used to automatically identify and classify the types of brain hemorrhages from the CT scans of patients.

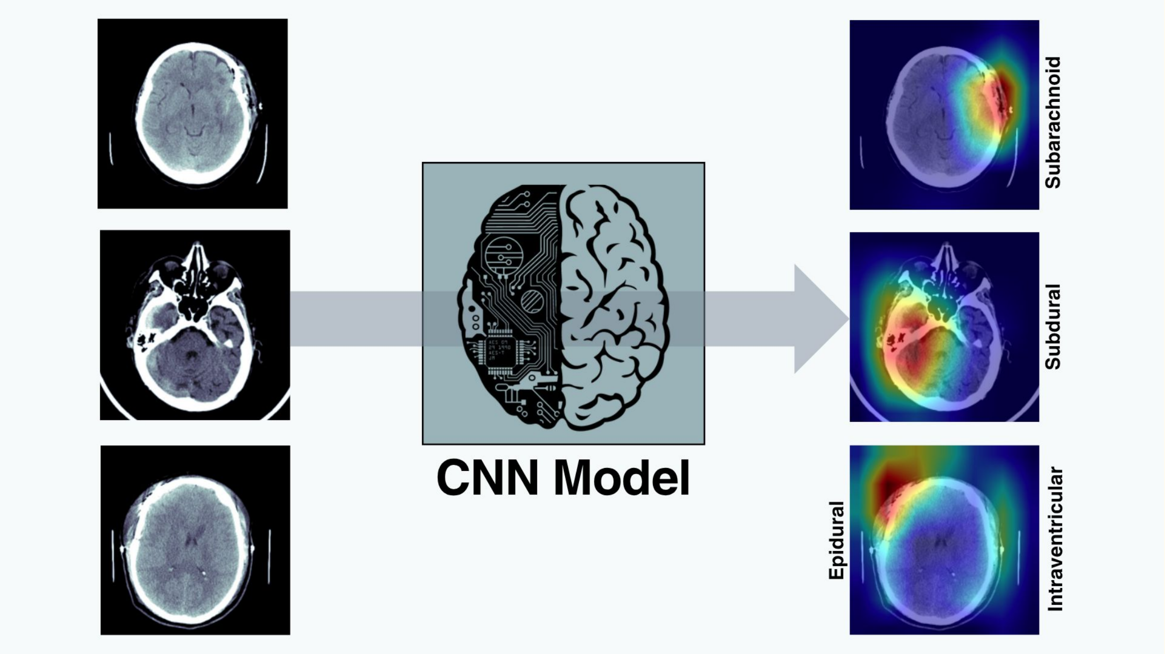

Overview CT Scans

- ❑ CT imaging captures X-rays from multiple angles around the head.
- ❑ A computer processes these signals to create cross-sectional images.
- ❑ Different tissues absorb X-rays differently — blood, bone, and brain have distinct densities, allowing CT to clearly show abnormalities like hemorrhages, tumors, or swelling.

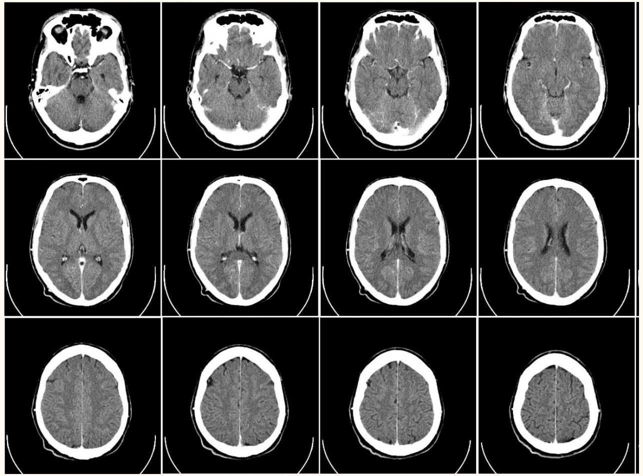

Background of CT Technology

- Computed Tomography (CT) was first introduced in the early 1970s, pioneered by **Sir Godfrey Hounsfield**.
- Early CT scanners used single-detector, first-generation systems that acquired images **slice-by-slice** using a **pencil-beam X-ray and a translate-rotate motion**.
- Image reconstruction relied on simple back-projection and took several minutes per slice, producing low-resolution images with limited clinical applications.

With technological advancements,

- **Helical/spiral CT** introduced continuous rotation and table movement, drastically reducing scan time.
- Multi-slice (MDCT) scanners with hundreds of detector rows allow high-resolution **volumetric imaging**.

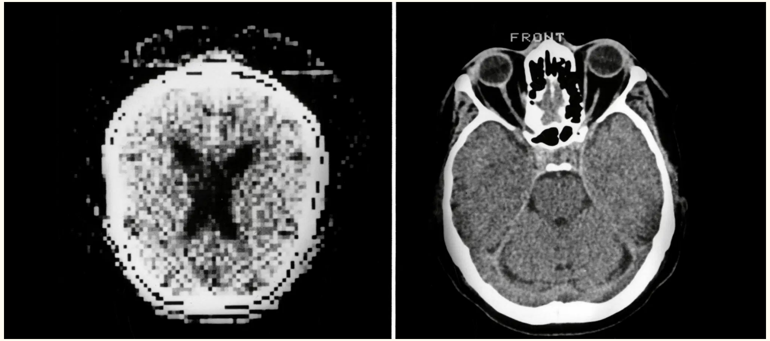

A CT scan from 1972 vs A CT scan using current techniques.

CT Scans - Volume

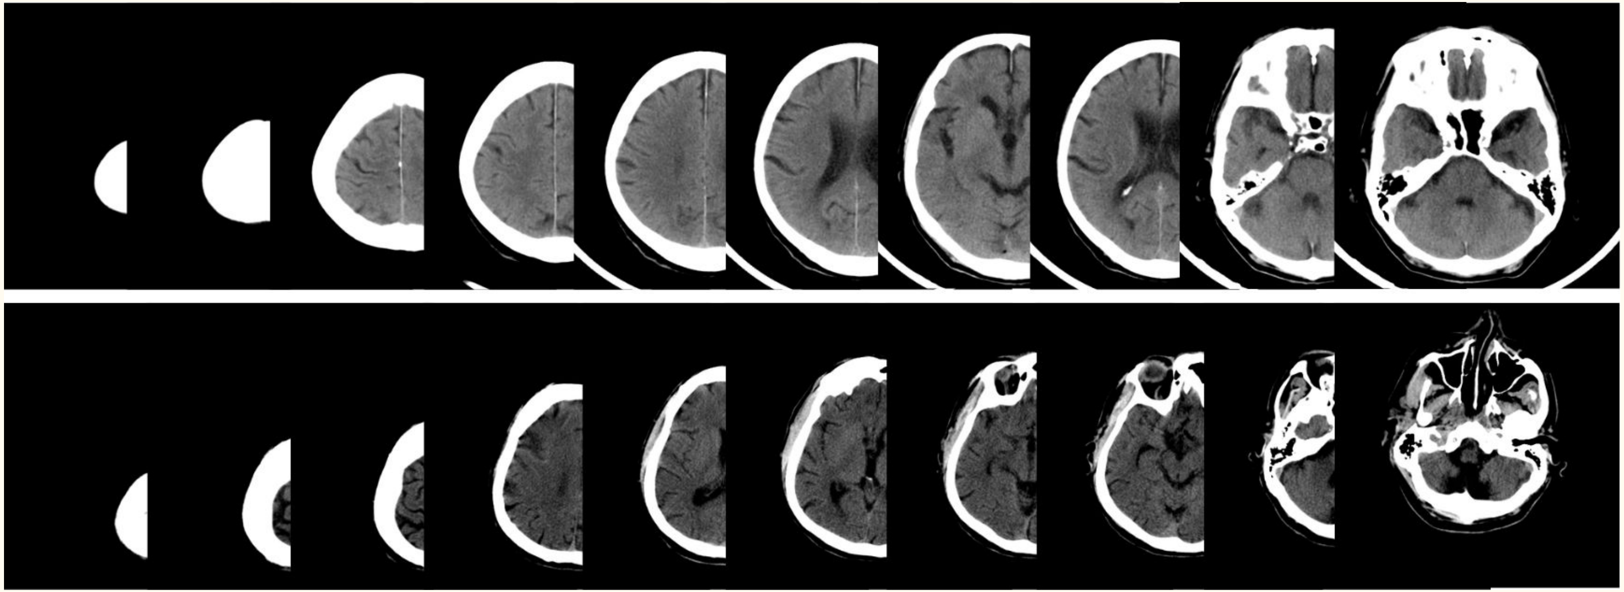

CT Scans of subjects are collected as **Volumes** which are a 3D representation of the brain consisting of multiple individual slices stacked together giving a full anatomical description

CT Scans - Slice

The **volumes** consist of individual **slices**, which refer to a single image on which diagnosis is done. These slices can be of different types:

- ❑ **Thin slices (0.5–1.25 mm)**: The image of each slice taken are very close to each other providing greater granularity and detail for subtle bleeds. High-detail, used for trauma, and vascular imaging.
- ❑ **Thick slices (3–5 mm)**: Images of slices taken at slightly longer distance for better global imaging. It has faster to view with lower noise
- ❑ **Plain (Non-contrast) CT**: Better structural overview, ideal for fractures, stroke, bleeding.
- ❑ **Contrast CT**: For soft tissue detail, vessels, tumors, infections.

THIN SLICE

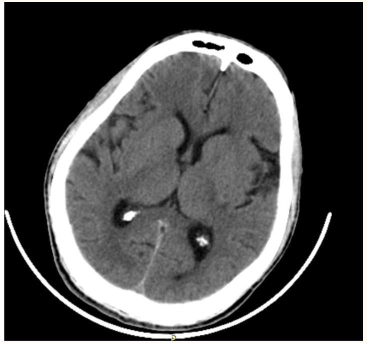

THICK SLICE

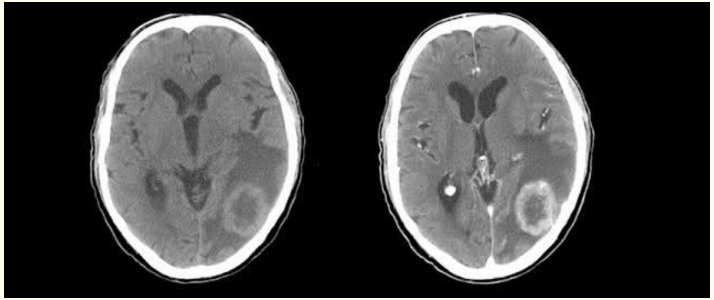

NON-CONTRAST vs CONTRAST

Pros and Cons of CT Scans

Pros

- ❑ Rapid image acquisition and availability in most hospitals.
- ❑ Excellent visualization of bone structures and acute bleeding.

Cons

- ❑ Requires expert interpretation for accurate diagnosis.
- ❑ Involves radiation exposure.

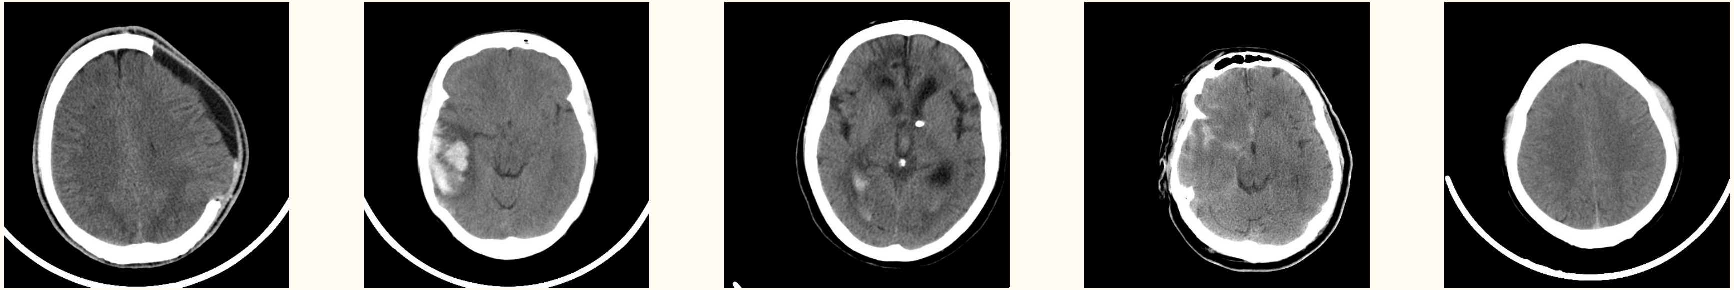

CT Images from RSNA Dataset

Brain Hemorrhage (ICH) Classification

Brain Hemorrhage Types:

- ❑ Epidural **(EDH)**
- ❑ Subdural **(SDH)**
- ❑ Subarachnoid **(SAH)**
- ❑ Intraparenchymal **(IPH)**
- ❑ Intraventricular **(IVH)**
- ❑ **ANY** – indicates presence of at least one hemorrhage type.

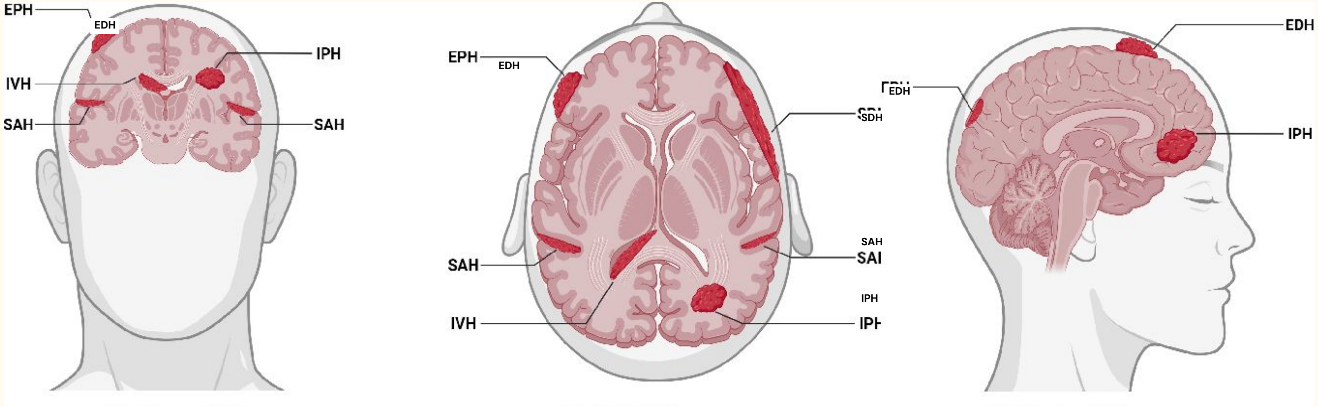

The goal is to **detect** and **classify** the above **hemorrhage types** in CT scans using **AI**

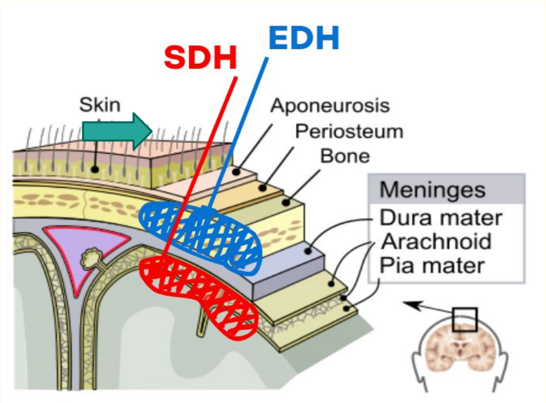

Brain Hemorrhage Sub-Types Details

- **IPH:** Bleeding within the brain tissue itself (cerebral hemispheres, basal ganglia, thalamus, brainstem, or cerebellum). This is the most common type associated with stroke or other ruptures within the brain parenchyma.

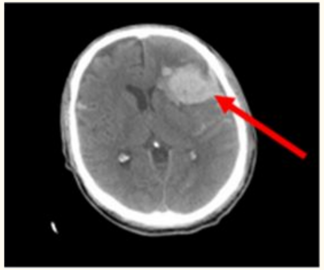

- **IVH:** Bleeding into the brain's fluid-filled cavities (ventricles), as shown in the picture, often as an extension of IPH. Also associated with stroke and other ruptures.

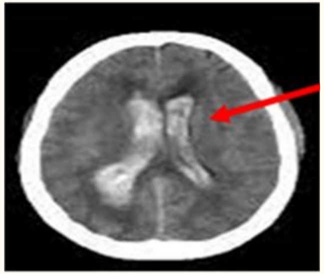

- **SAH:** Bleeding into the space between the brain and the thin tissues covering it (the arachnoid membrane), typically caused by an aneurysm rupture or trauma.

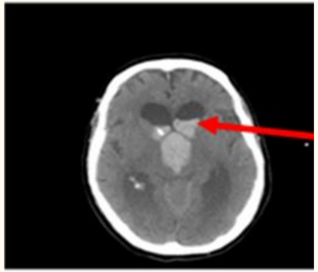

Brain Hemorrhage Sub-Types Details

- **SDH:** A collection of blood on the surface of the brain, under the dura mater (outermost membrane covering the brain). That is why it is called **sub**dural. This is usually caused by trauma that ruptures bridging veins.

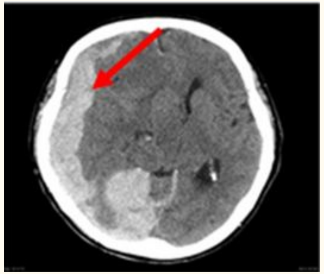

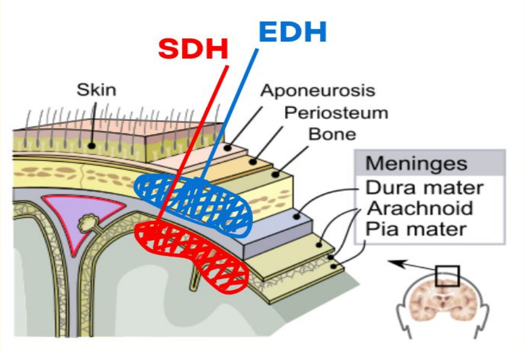

Brain Hemorrhage Sub-Types Details

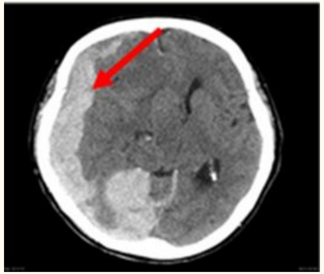

- **SDH:** A collection of blood on the surface of the brain, under the dura mater (outermost membrane covering the brain). That is why it is called **sub**dural. This is usually caused by trauma that ruptures bridging veins.

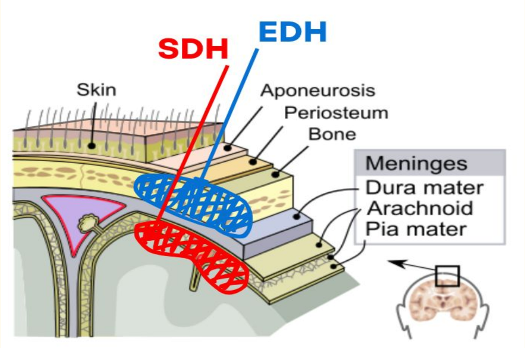

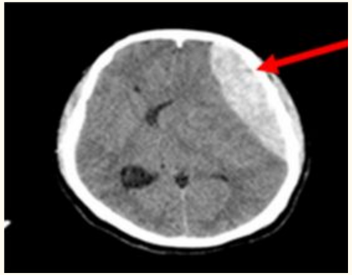

- **EDH:** Bleeding between the skull (bone) and the dura mater. This is usually the result of a traumatic head injury, often with a skull fracture.

Brain Hemorrhage Sub-Types Details

The following are the descriptions of each individual sub-types:

- **IPH**: Bleeding within the brain tissue itself (cerebral hemispheres, basal ganglia, thalamus, brainstem, or cerebellum). This is the most common type associated with stroke.
- **IVH**: Bleeding into the brain's fluid-filled cavities (ventricles), often as an extension of IPH.
- **SAH**: Bleeding into the space between the brain and the thin tissues covering it (the arachnoid membrane), typically caused by an aneurysm rupture or trauma.
- **SDH**: A collection of blood on the surface of the brain, under the dura mater (outermost membrane covering the brain). This is usually caused by trauma that ruptures bridging veins.
- **EDH**: Bleeding between the skull and the dura mater. This is usually the result of a traumatic head injury, often with a skull fracture.

Datasets

RSNA (Radiological Society of North America)[1]

- The data come from three institutions: Stanford University, Universidade Federal de São Paulo (Brazil), and Thomas Jefferson University Hospital.

- It includes **non-contrast head CT studies**, de-identified and anonymised, and each slice (and thus each volume) is labelled for the presence of intracranial hemorrhage and five main sub-types.

- The **RSNA** brain hemorrhage dataset includes **752,803 training images (slices)** from **21,784 volumes** and **121,232 test images (slices)** from **3,528 volumes**, with hemorrhage subtypes distributed as epidural, intraparenchymal, intraventricular, subarachnoid, and subdural in the training/test sets respectively, alongside normal images.

- It is openly available on **Kaggle**[1], and is free-to-use for academic and research purposes.

[1] https://www.kaggle.com/c/rsna-intracranial-hemorrhage-detection/, https://doi.org/10.1148/ryai.2020190211

Overview of DICOM Format

- **DICOM (Digital Imaging and Communications in Medicine)** was introduced in 1985 as ACR–NEMA 1.0, later renamed and expanded into DICOM in the 1990s to standardize the storage, transmission, and interoperability of medical images across devices and vendors.

- **Advantage of DICOM**: Ensures universal compatibility between scanners, PACS systems, and viewing software and stores both the image and rich metadata (patient details, acquisition parameters, slice position, etc.).

- **Why DICOM is still so popular**: It is the global standard—adopted by virtually all imaging vendors (CT, MRI, X-ray, PET, ultrasound).

- **Limitations of DICOM**: Not optimized for deep learning pipelines, often requiring conversion to PNG, JPEG, or NumPy arrays for training.

Pre-Processing

- The dataset consists of subject folders, each containing several CT scan types (e.g., Plain, Contrast, Thin Slice).
- Each scan type has multiple CT slices, and a CSV file records hemorrhage labels.
- For CNN training, labels are assigned to all individual slices to create slice-level supervision.

- Images undergo preprocessing transformations such as:
  - **Resizing to uniform dimensions** (e.g., 512×512)
  - **Converting labels to tensors [Batch Size, No. of classes]**
  - **Normalizing pixel intensity values**

Images are now ready to be fed into CNNs!

CNNs - A Technical Overview

- CNNs consist of specialized layers that progressively learn features:
  - **Convolutional Layers:** Extract patterns such as edges or textures.
  - **Pooling Layers:** Downsample to reduce size and improve invariance.
  - **Fully Connected (FC) Layers:** Combine features for final classification.

- This layered hierarchy allows CNNs to recognize structures from simple to complex forms.

1. Convolution on an IxI image with an KxK filter

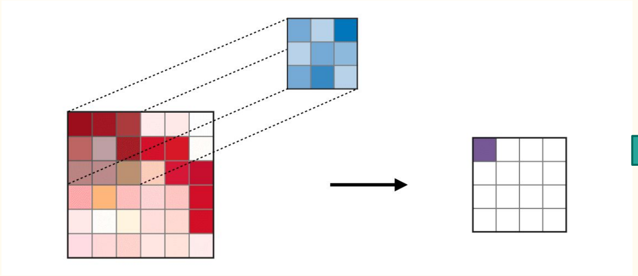

2. Max-Pooling on Feature Maps

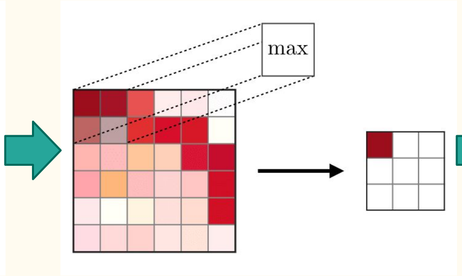

3. Flatten and FC Layer

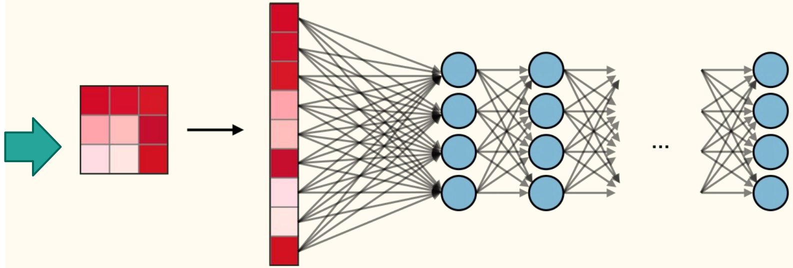

CNN Feature Maps

Suppose, we have an image **I = 5x5,** with **C=3** channels as the input and we apply a filter with **F = 2x2** kernel on it, and stride **S = 1,** with the output channels set to **K,** the convolution steps to obtain feature maps are as follows:

$$X^{(1)} = \begin{bmatrix} 1 & 2 & 3 & 0 & 1 \\ 0 & 1 & 2 & 3 & 1 \\ 1 & 0 & 1 & 2 & 2 \\ 2 & 1 & 0 & 1 & 1 \\ 1 & 2 & 1 & 0 & 1 \end{bmatrix} \quad X^{(2)} = \begin{bmatrix} 0 & 1 & 2 & 1 & 0 \\ 1 & 0 & 1 & 2 & 1 \\ 2 & 1 & 0 & 1 & 2 \\ 1 & 2 & 1 & 0 & 1 \\ 0 & 1 & 2 & 1 & 0 \end{bmatrix} \quad X^{(3)} = \begin{bmatrix} 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 \end{bmatrix}$$

5x5 feature maps of the image, for **3 channels**, stacked to form a **5x5x3** input

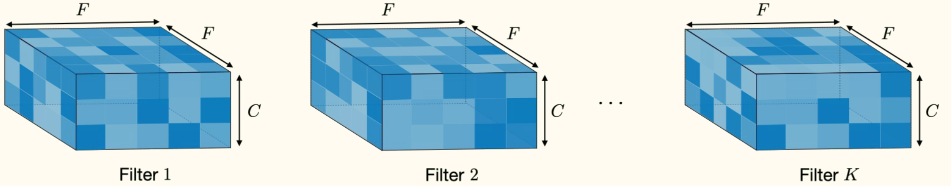

Passed through **K** filters, each of size **2x2x3**

4x4 feature maps of the image, for **k channels**, i.e. a **4x4xk** output*

$$Y_k = \begin{bmatrix} -0.8 & 1.9 & 0.3 & 2.2 \\ 1.1 & -1.4 & 2.6 & 0.5 \\ 0.7 & 3.3 & -0.2 & 1.8 \\ 2.0 & 0.1 & 1.4 & -0.9 \end{bmatrix} \quad \underset{\textbf{(k)}}{\cdots\cdots} \quad Y_1 = \begin{bmatrix} 2.1 & 0.4 & 1.7 & -0.3 \\ 0.9 & 3.2 & -1.1 & 0.8 \\ -0.6 & 1.5 & 2.9 & 0.2 \\ 1.3 & -0.4 & 0.7 & 2.5 \end{bmatrix}$$

*Output dimensions calculated using: $O = \dfrac{I - F}{S} + 1$

CNN Architecture

- The basic CNN model, as shown here, consists of **4 convolutional layers** followed by **2 fully connected layers** for multi-label brain hemorrhage (ICH) classification.
- It is trained using the **Binary Cross Entropy Loss (BCE)** for 10 epochs. The BCE loss measures how well a model's predicted probabilities **[y’]** align with the actual, true outcomes **[y]** in a classification problem **(where the answer is one of two options, "yes" or "no", represented as 1 or 0)**.

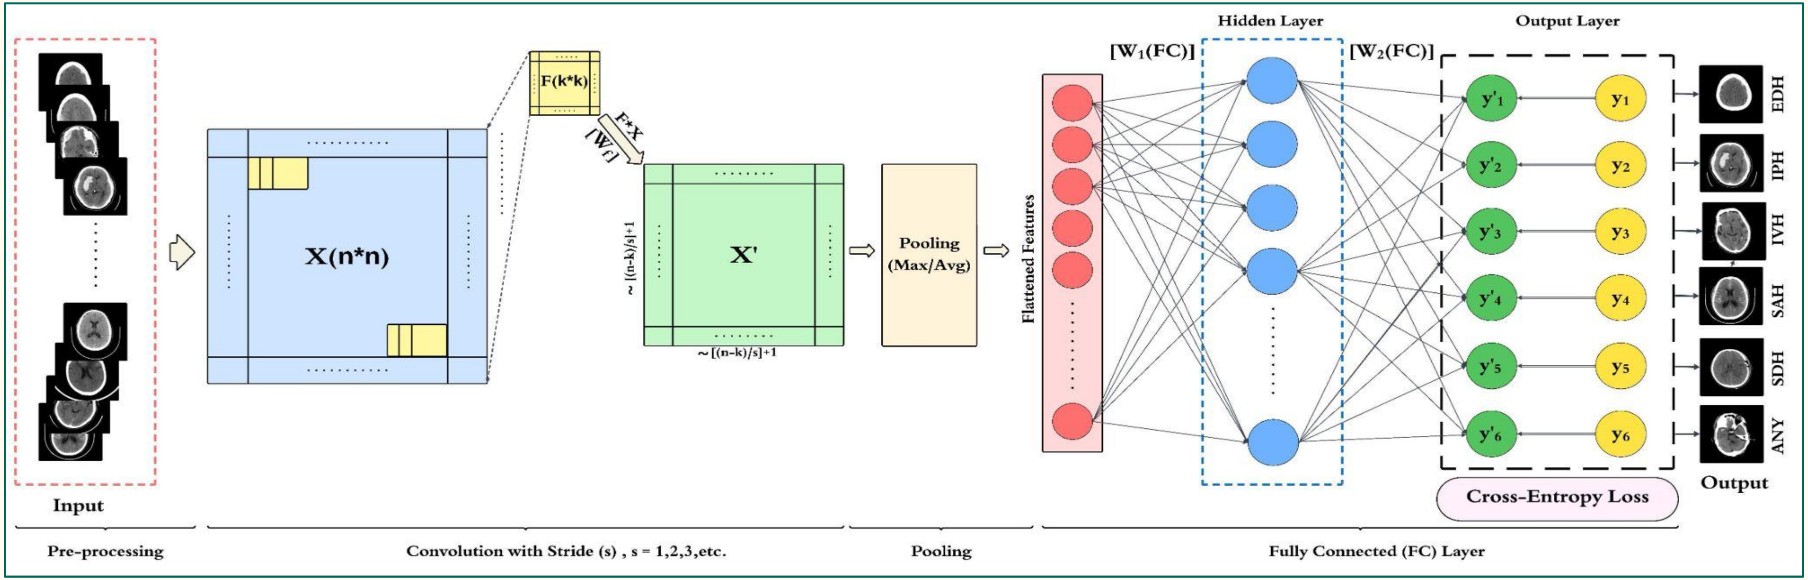

Commonly Used CNN Backbones

- **ResNet (Residual Network):** Uses skip connections to preserve gradients and extract deep hierarchical features efficiently.

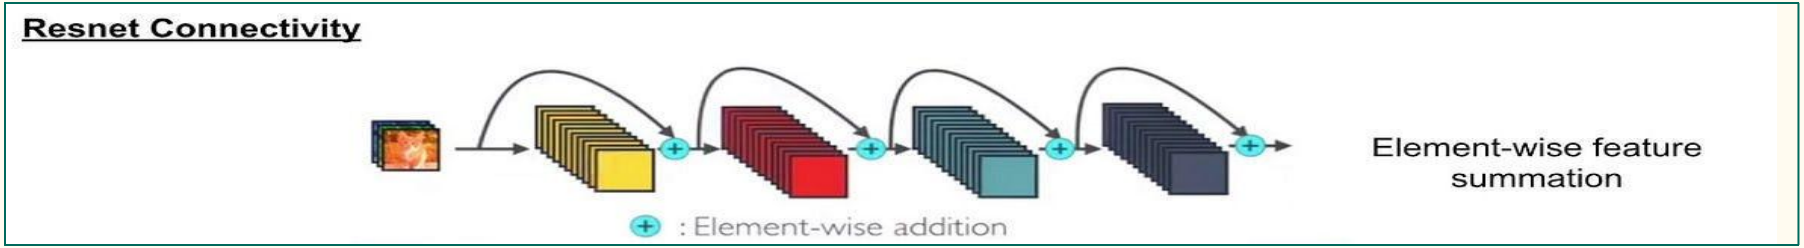

- **DenseNet (Densely Connected Network):** Connects each layer to every other layer, promoting feature reuse and stronger gradient flow.

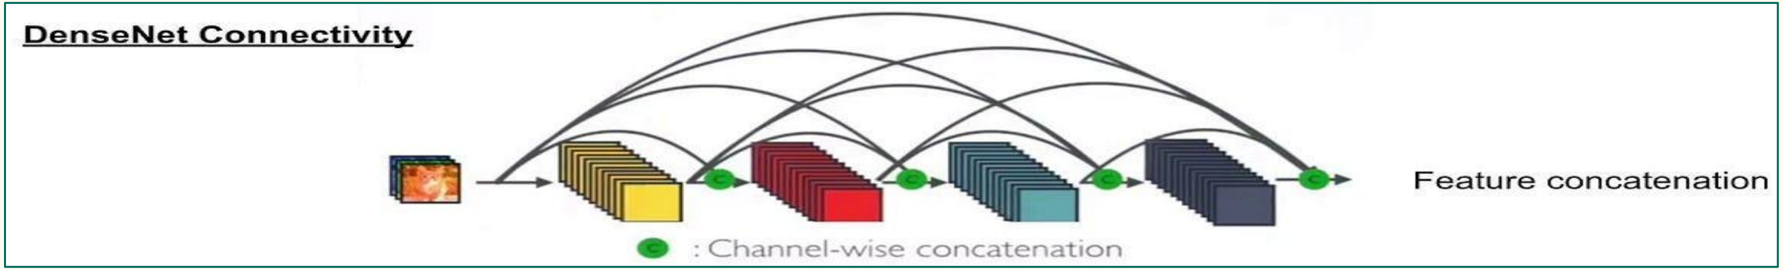

Confusion Matrix

Let us assume there are two classes in a data set, say, *positive* and *negative*. A confusion matrix looks as follows:

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

- TP (True Positive): Number of data points correctly predicted to the positive class.
- FP (False Positive/ Type I Error): Number of data points that actually belong to the negative class, but predicted as positive (i.e., *falsely predicted as positive*).
- FN (False Negative/ Type II Error): Number of data points that actually belong to the positive class, but predicted as negative (i.e., *falsely predicted as negative*).
- TN (True negative): Number of data points correctly predicted to the negative class.

Performance Evaluation Metrics

- **Accuracy** - Represents the proportion of correctly classified samples out of all predictions i.e. out of all predictions, how many did the model get correct?
- **Precision -** Proportion of predicted positives that were actually positive.
- **Recall -** Proportion of actual positives that the model successfully detected.
- **F1-Score -** It is the harmonic mean of precision and recall.
- **AUC (Area Under Curve)** - Measures the probability that the model ranks a randomly chosen positive case higher than a negative one.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \qquad \text{Precision} = \frac{TP}{TP + FP} \qquad \text{Recall} = \frac{TP}{TP + FN} \qquad F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

$$\text{AUC} = \int_0^1 \text{TPR}(\text{FPR})\, d(\text{FPR}) \qquad \text{TPR} = \frac{TP}{TP + FN}, \qquad \text{FPR} = \frac{FP}{FP + TN}$$

Training Metrics

During training, we rely on the following for model performance:

- **Loss Curve:** Tracks change in training and validation loss over epochs.
  - Lowest validation loss indicates optimal model state.
- **Accuracy Curve:** Tracks how well predictions match true labels.
  - Increasing training accuracy with falling validation accuracy indicates overfitting.

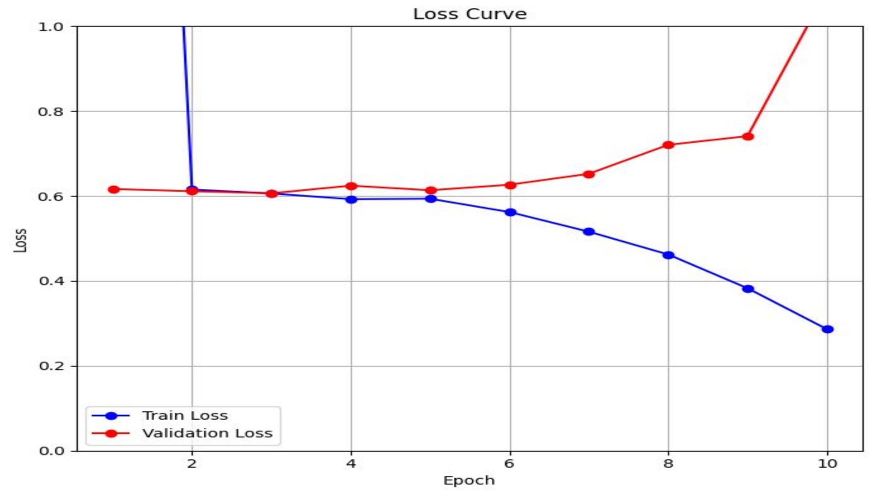

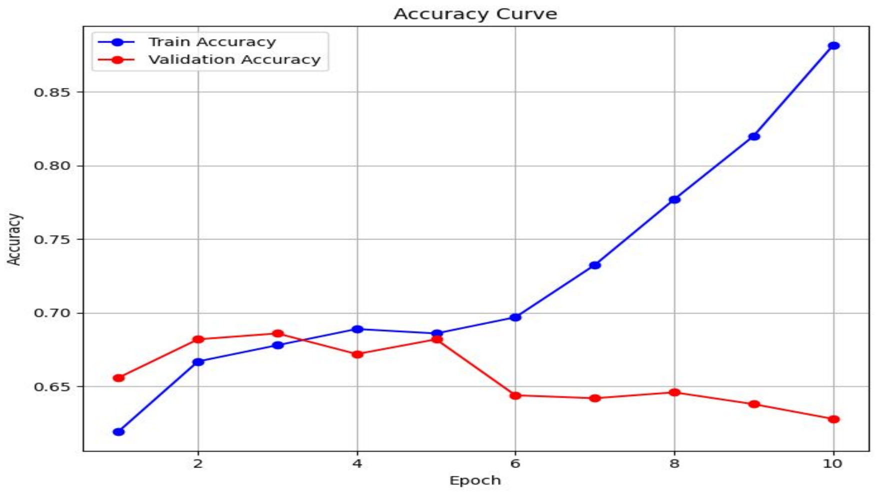

Model Evaluation

After training, we evaluate our model on unseen CT images:

- ❑ **ROC-AUC or Receiver Operating Characteristic Area Under Curve:** Used to measure the model's ability to distinguish between positive and negative classes. It ranges between 0 to 1. We get it by plotting values between **TPR (True Positive Rate)** and **FPR (False Positive Rate)**.

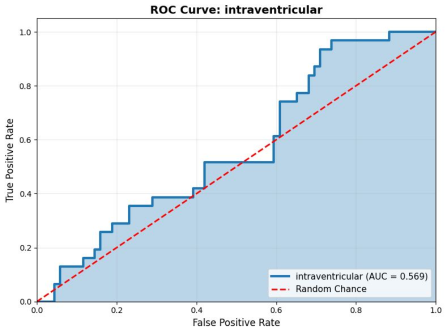

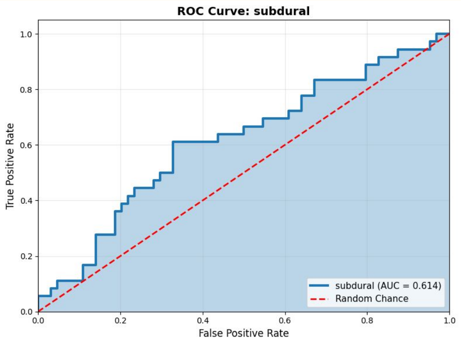

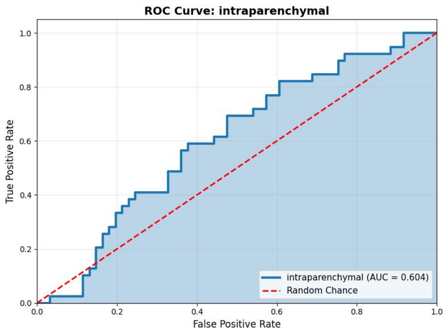

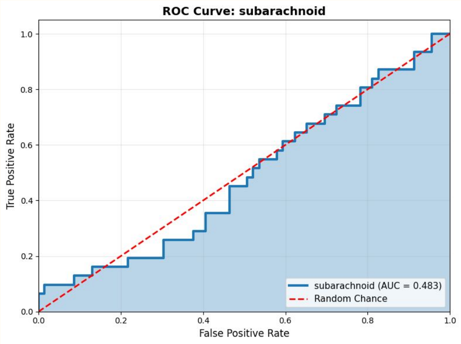

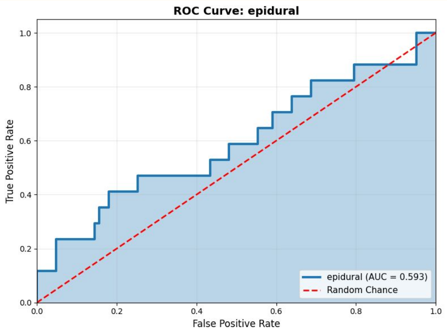

Confusion Matrix of Predictions

- The CNN's predictions on **100 unseen CT images** yield the following confusion matrices, with **TP, TN, FP, FN**:

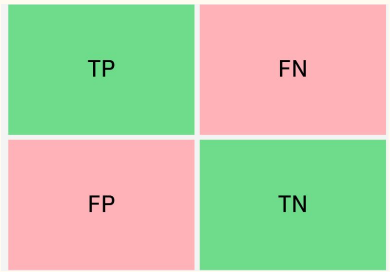

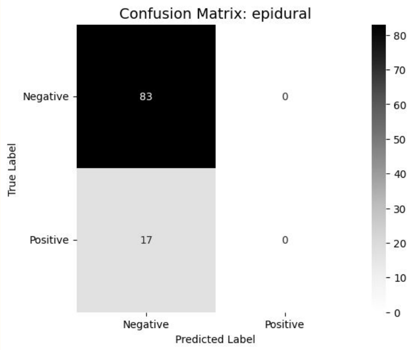

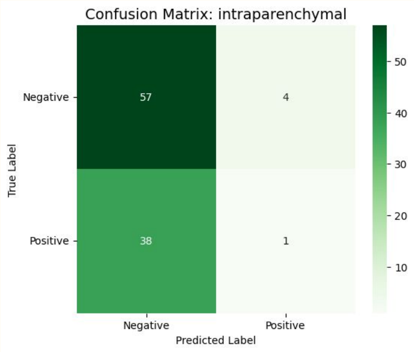

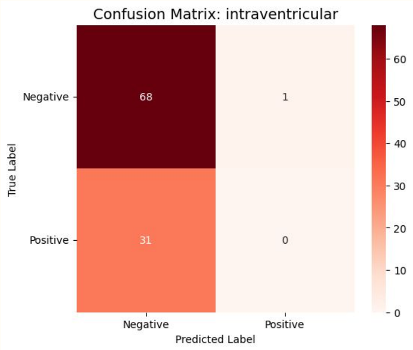

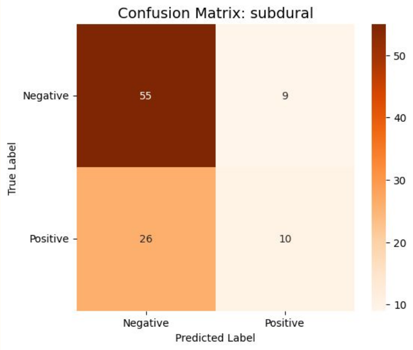

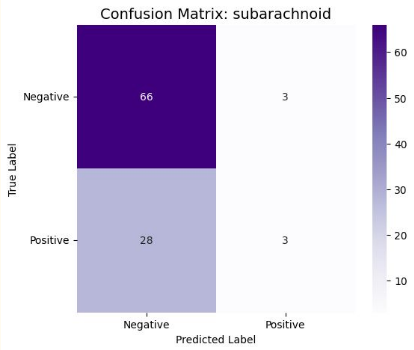

(BDS Lab)
BIOMEDICAL DATA SCIENCE RESEARCH LAB

Acknowledgement

Connect @

Dr. Pragya Kumari
Postdoctoral Researcher

Himadri Sonowal
MS Student

Srutanik Bhaduri
MS Student

Saisab Sadhu
MS Student

Ashim Dhor
MS Student

Dr. Riya Agarwal
Radiation Oncologist

Prateek Sarangi
PhD Scholar

Rasel Mondal
PhD Scholar

Sumit Kumar
PhD Scholar

Dr. Tanmay Basu
PI

Thank You!