Democratic and Popular Republic of Algeria
Ministry of Higher Education and Scientific Research
University Centre of Illizi
Institute of Science

المركز الجامعي ايليزي

معهد العلوم

Department of Computer Science

Master's Thesis in Artificial Intelligence & Its Applications

Symbiotic AI for Medical Imaging: Anomaly Detection and Explanation Generation Using Knowledge Graphs

Authors:

- Fatma Zohra BOUSSOUFA
  fatmazohra.boussoufa@cuillizi.dz

Advisors:

- Dr. Youcef BEZZA
  youcef.bezza@cuillizi.dz
- Dr. Fatima Zahra AMARA
  amafatima.zahra@gmail.com

Academic Year: 2026

Abstract

Medical imaging plays a crucial role in disease diagnosis, monitoring, and treatment planning. Although Convolutional Neural Networks (CNNs) achieve high accuracy in analyzing MRI, CT, and X-ray images, their black-box nature limits clinical trust and interpretability. Moreover, reliable diagnosis requires integrating visual evidence with structured medical knowledge. To address this challenge, this project proposes a Symbiotic Artificial Intelligence framework that combines a fine-tuned VGG16 CNN with a brain tumor ontology developed in OWL and enhanced by SWRL rules covering ten categories of medical inference. Using Owlready2 and RDFLib, the system performs ontology-based reasoning to generate an interpretable diagnostic profile for each MRI image, including diagnosis status, malignancy classification, expected symptoms, recommended treatment, confidence level, and CNN–ontology agreement. The project was completed over six months, including literature review, implementation, experimentation, and report writing. Trained and evaluated on 10,800 brain MRI images (5,400 tumor and 5,400 non-tumor), the proposed framework achieved a test accuracy between 98% and 99%, demonstrating the feasibility and clinical relevance of integrating deep learning with knowledge-based reasoning for explainable brain tumor detection.

Keywords: Brain Tumor Detection; MRI; CNN; OWL Ontology; Symbiotic AI.

ملخص

يلعب التصوير الطبي دورًا أساسيًا في تشخيص الأمراض ومتابعة المرضى والتخطيط للعلاج. وعلى الرغم من أن الشبكات العصبية التلافيفية (CNNs) تحقق دقة عالية في تحليل صور الرنين المغناطيسي (MRI) والتصوير المقطعي المحوسب (CT)والأشعة السينية(X-ray) ، فإن طبيعتها كـ«صندوق أسود» تحدّ من إمكانية تفسير نتائجها ومن ثقة الأطباء بها. بالإضافة إلى ذلك، يتطلب التشخيص الموثوق دمج الأدلة البصرية مع المعرفة الطبية المنظمة.

ولمواجهة هذا التحدي، يقترح هذا المشروع إطارًا للذكاء الاصطناعي التكاملي (Symbiotic Artificial Intelligence) يجمع بين نموذج VGG16 المُحسَّن والشبكة العصبية التلافيفية من جهة، وأنطولوجيا لأورام الدماغ مطورة بلغة OWL ومعززة بقواعد SWRL تغطي عشر فئات من الاستدلال الطبي من جهة أخرى. وباستخدام مكتبتي Owlready2 وRDFLib، يقوم النظام بإجراء استدلال قائم على الأنطولوجيا لتوليد ملف تشخيصي قابل للتفسير لكل صورة رنين مغناطيسي، يتضمن حالة التشخيص، وتصنيف درجة الخباثة، والأعراض المتوقعة، والعلاج الموصى به، ومستوى الثقة، وتقييم مدى توافق نتائج الشبكة العصبية مع استنتاجات الأنطولوجيا.

أُنجز المشروع خلال ستة أشهر، شملت مراجعة الأدبيات العلمية، والتنفيذ، والتجارب، وكتابة التقرير. وقد تم تدريب الإطار المقترح وتقييمه باستخدام 10800 صورة رنين مغناطيسي للدماغ (5400 صورة لأورام دماغية و5400 صورة سليمة)، حيث حقق دقة اختبار تراوحت بين 98٪ و99٪، مما يبرهن على جدوى وأهمية الدمج بين التعلم العميق والاستدلال المعتمد على المعرفة لتطوير نظام قابل للتفسير للكشف عن أورام الدماغ.

الكلمات المفتاحية :كشف أورام الدماغ، التصوير بالرنين المغناطيسي، الشبكات العصبية الالتفافية، الأنطولوجياOWL ، الذكاء الاصطناعي التكاملي.

Acknowledgements

First and foremost, I thank Allah for giving me the strength, patience, and determination to complete this work.

To the souls of my grandfathers, who are no longer with us, I dedicate this achievement. May Allah have mercy on them and grant them the highest place in Jannah.

To my grandmothers, thank you for your prayers, love, and constant support. Your duas have been a source of strength throughout this journey.

To my brothers and sisters, thank you for your encouragement, patience, and belief in me. Your support and presence made every challenge easier to overcome.

To my uncle and aunts, thank you for your kindness, guidance, and continuous encouragement. Your support means a great deal to me.

I would like to express my sincere gratitude to my supervisors, **Dr. Youcef BEZZA** and **Dr. Fatima Zahra AMARA**, for their guidance, valuable advice, and continuous support throughout this research. Their expertise and dedication played a major role in the completion of this work.

To my university friends, thank you for all the moments we shared, the encouragement during difficult times, and the memories that made these years special.

To my colleagues and friends at work, thank you for your understanding, motivation, and support while I balanced my professional responsibilities and academic research.

Finally, I would like to thank myself for persevering through challenges, doubts, and difficult moments. This achievement is the result of hard work, determination, and faith.

This accomplishment belongs to everyone who supported me throughout this journey. Thank you.

الإهداء

" و أخر دعواهم أن الحمد لله ربّ العالمين "

باسم الله نبدأ و بحمده ننتهي و الصّلاة و السّلام على من علّمنا أنّ طلب العلم فريضة و أنّ الصّبر مفتاح الفرج.

الحمد لله أوّلا وأخيرا، حبّا و شكرا و امتنانا على البدء و الختام، الحمد لله الذي كان معنا حين تعبنا و حين ضعفنا و حين ظننّا أنّ الطّريق أطول من قدرتنا على الإستمرار فكان لطفه خفيّا و عونه قريبا و توفيقه حاضرا في كلّ مرحلة من هذه الرّحلة، الحمد لله الذي بنعمته تتمّ الصّالحات.

لم تكن رحلتي مع هذا التّخصّص حبّا منذ البداية فقد بدأتها دون تخطيط منّي و بخطوات يملأها التّردّد و الشّكّ، لم أكن أظنّ أنّني سأجدني فيه لكنّني تعلّمت أنّ بعض الطّرق نكتشف جمالها و نحن نسير فيها حتّى تصبح جزءا منّا و هكذا وجدت نفسي بين خوارزميّاته أتعلّم، أتعب و أنضج خطوة خطوة...

أهدي تخرّجي:

إلى سندي، مأمني و أماني إلى والدي العزيز الذي لطالما افتخرت لكوني ابنته و إلى غاليتي، نور عيني و جنّتي في الدّنيا أمّي العزيزة التي لم يفارقني دعائها ولو للحظة.
إلى إخوتي عبد الرّزّاق،لقمان،البخاري،محمّد،عثمان،زينب و رقيّة الّذين شاركوني تفاصيل هذه الرّحلة،لكم منّي كلّ المحبّة و الإمتنان.
إلى روح جدّيّ ع.الرّزّاق و ع.القادر رحمهما الله رحمة واسعة و أقول لكم حفيدتكم فاطمة ستتمّ مسيرتها بالنّجاح كما آلفتموها،أسأل الله أن يجعل هذا النّجاح صدقة جارية في ميزان حسناتكما و أن يرفع قدركما في جنّات النّعيم.
إلى جدّتيْ أطال الله عمرهما.
إلى أعمامي، عمّاتي، أخوالي و خالاتي حفظهم الله.
إلى أساتذتي كلّ باسمه و مقامه و اخصّ بالذّكر الدّكتور يوسف بزّة و الدّكتورة عمارة فاطمة الزّهراء عرفانا بما قدّمتموه من علم و توجيه في هذه الرّحلة العلميّة فقد كان دعمهما مصدر ثقة و دافعا للاستمرار.
إلى زميلاتي و زملائي في الصّفّ الذين تقاسموا معي مشاعر هذه الرّحلة المميّزة أسأل الله النّجاح لكلّ واحد منكم.
إلى زملائي في العمل الذين كان لهم دعم خفيّ و يدُ عَوْنٍ في هذا العام أسعد الله قلوبكم جميعا.
إلى كلّ من أحمل له ذرّة حبّ و احترام في قلبي كان حبّكم مصدر قوّتي الدّائم.
و أخيرا أهدي هذا العمل إلى نفسي و ذاتي التي صبرت و اجتهدت و لم تستسلم رغم الصّعوبات، فكان الأمل دائما أقوى من التّعب و كان الإصرار طريقا للوصول.
و في الختام أحمد الله من قبل و من بعد راجية من الله تعالى أن ينفعني بما علّمني و أن يعلّمني ما أجهل و أن يجعلها حجّة لا عليّ.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

AI	Artificial Intelligence
DL	Deep Learning
ML	Machine Learning
CNN	Convolutional Neural Network
VGG16	Visual Geometry Group 16-layer Network
SVM	Support Vector Machine
GRU	Gated Recurrent Unit
BERT	Bidirectional Encoder Representations from Transformers
PCA	Principal Component Analysis
RAG	Retrieval-Augmented Generation
CAM	Class Activation Mapping
LIME	Local Interpretable Model-Agnostic Explanations
MRI	Magnetic Resonance Imaging
CT	Computed Tomography
RGB	Red Green Blue colour space
GBM	Glioblastoma Multiforme
OWL	Web Ontology Language
SWRL	Semantic Web Rule Language
RDF	Resource Description Framework
RDFS	RDF Schema
SPARQL	SPARQL Protocol and RDF Query Language
XML	Extensible Markup Language
KG	Knowledge Graph
SNOMED	Systematized Nomenclature of Medicine
ICD	International Classification of Diseases

WHO World Health Organization

CDSS Clinical Decision Support System

DOID Disease Ontology Identifier

UMLS Unified Medical Language System

TP True Positive

TN True Negative

FP False Positive

FN False Negative

AUROC Area Under the Receiver Operating Characteristic Curve

MCC Matthews Correlation Coefficient

TRIPOD Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis

GPU Graphics Processing Unit

PIL Python Imaging Library

Chapter 1

General Introduction

1.1 Work Context and Problem Statement

The way we diagnose and treat diseases has changed a lot over the few decades because medical technology has improved so quickly. One of the changes in modern medicine is medical imaging. Medical imaging is an important tool that doctors use to see what is going on inside the human body. It lets them see the systems of the body very clearly and accurately.

Medical imaging methods like Magnetic Resonance Imaging, Computed Tomography, X-ray and ultrasonography create pictures with important information about diseases. These pictures can help doctors diagnose diseases.

Finding anomalies in images is a very challenging and exciting area of study. It is where intelligence and medicine meet. A medical anomaly is anything that's not normal in a tissue or organ. Medical imaging can help detect these anomalies.

Detecting these anomalies quickly and correctly is very important. It affects the quality of the diagnosis and the treatment plan. It can even affect whether the patient survives or not.. Looking at medical images is not easy. It takes a lot of time and effort.

Doctors can get tired. Make mistakes. They can also miss something. This is especially true in hospitals where doctors have to look at hundreds of images every day.

This is why researchers are working hard to create computer systems that can help doctors diagnose diseases. These systems use learning techniques like Convolutional Neural Networks. They are very good at analyzing images and detecting anomalies.

These systems have a big problem. They do not explain how they make their decisions. They just make predictions without giving any reasons. This is not good enough for doctors. Doctors need to know why a system is making a diagnosis.

Medical knowledge is very complex. It is not, about looking at a picture and seeing something abnormal. Doctors need to know about the patients symptoms, history and treatment options. They need to understand the picture.

This is why researchers are trying to create systems that're not only accurate but also transparent and reliable. They want to combine the strengths of learning with the strengths of medical knowledge. This will create a generation of diagnostic tools that doctors can actually use in real life.

These new systems will be able to analyze images and detect anomalies.. They will also be able to explain how they make their decisions. They will be able to provide doctors with the information they need to make diagnoses and create effective treatment plans. Medical imaging and medical knowledge will work together to improve care.

1.2 Our Contribution

Among the most serious and deadly illnesses requiring quick and accurate diagnosis, brain tumors are one of the many medical conditions that profit from creative imaging-based diagnosis. A brain tumor, sometimes called brain cancer, results from a mass that forms when a group of brain cells starts to grow out of control. Malignant (cancerous) and benign (non-cancerous) are the two primary kinds of tumors[1]. From sluggish-growing benign tumors like meningiomas and pituitary adenomas to highly aggressive malignant tumors like Glioblastoma Multiforme (GBM)—one of the most deadly kinds of human cancer with a median survival of under 15 months following diagnosis—brain tumors are categorized by their cellular origin, placement, and aggressiveness.

Because of its superb soft-tissue contrast, non-invasive nature, and capacity to collect elaborate three-dimensional images of brain anatomy without the need of ionizing radiation, Magnetic Resonance Imaging (MRI) is widely acknowledged as the gold standard technique for the detection and characterization of brain tumors. Still, it is a very difficult clinical assignment to accurately interpret brain MRI scans, necessitating expertise in radiology and neurology that is not consistently accessible throughout healthcare systems, especially in resource-constrained contexts and emerging countries.

Our study is carried out in this medical and scientific context. This paper presents a hybrid intelligent system for finding brain tumors from MRI pictures that works on the Symbiotic AI principle. Two mutually reinforcing and complementary elements define our approach. A Convolutional Neural Network (CNN) trained on a sizable dataset of labeled brain MRI images executes the visual detection task, deciding whether a specific scan shows symptoms of a brain tumor. The second is a brain tumor domain ontology formalized with OWL and enhanced with SWRL inference rules. It interprets the CNN result and adds organized medical information, including tumor type, malignancy grade, symptoms, MRI findings, recommended treatment protocols, and diagnosis confidence.

By combining ontological reasoning with deep visual learning, our method goes beyond simple binary classification and creates a complete, understandable, and knowledge-based diagnostic report. It supports both the accuracy requirements of automatic detection and the interpretability requirements of clinical adoption, giving doctors a clear diagnostic explanation instead of only a binary prediction.

1.3 Dissertation Structure

This manuscript is structured into five chapters followed by a general conclusion.

1.3.1 Chapter 1: Ontology and Semantic Knowledge Representation:

This chapter introduces the fundamental concepts of ontologies and semantic web technologies. We present the key languages and standards used in our work, including OWL, RDF, SPARQL, and SWRL, and we discuss the role of ontologies in the medical domain with particular focus on brain tumor knowledge modeling and representation.

1.3.2 Chapter 2: Anomaly Detection Using Artificial Intelligence:

This chapter covers the theoretical foundations of AI-based detection methods applied to medical imaging. We present deep learning architectures, particularly Convolutional Neural Networks, their application to brain MRI analysis, and the interpretability problem of black-box models, introducing the concept of explainable and hybrid artificial intelligence systems.

1.3.3 Chapter 3: Related Work:

This chapter presents a critical and structured review of existing works that combine knowledge representation with artificial intelligence for medical imaging analysis and anomaly detection. We examine hybrid approaches integrating ontologies, knowledge graphs, symbolic reasoning, and machine learning, and we position our work with respect to the existing literature.

1.3.4 Chapter 4: Proposed Model:

This chapter presents in detail our proposed hybrid system. We describe the architecture of the CNN model, the design of the brain tumor OWL ontology, the SWRL reasoning rules, and the integration mechanism between both components, along with the implementation tools and technologies used.

1.3.5 Chapter 5: Results and Discussion:

This chapter presents and analyzes the experimental results obtained by our system, evaluating both the CNN model performance and the ontological reasoning output. We discuss the strengths, limitations, and clinical relevance of our approach and compare our results with related existing works.

The memoir ends with a general conclusion that summarizes our contributions, discusses the limitations of the proposed system, and outlines future research directions toward more complete and clinically deployable intelligent diagnostic tools.

Chapter 2

Ontology Design and Knowledge Representation

2.1 Introduction to Ontologies:

2.1.1 what is ontology?

Ontology is a branch of philosophy concerned with the nature of existence and the classification of entities in the world. In computer science and knowledge representation, an ontology is a formal and structured model that defines the concepts of a domain and the relationships between them.

An ontology consists of key components such as classes, individuals (instances), attributes, relations, restrictions, rules, and axioms. These elements enable the organization, sharing, reuse, and reasoning of knowledge within a specific domain.

Ontologies play a fundamental role in knowledge graphs, where entities are represented as nodes and their relationships as edges. By providing a clear conceptual structure and semantic meaning, ontologies support the effective representation, management, integration, and inference of knowledge, allowing systems to generate new information through automated reasoning[2][3].

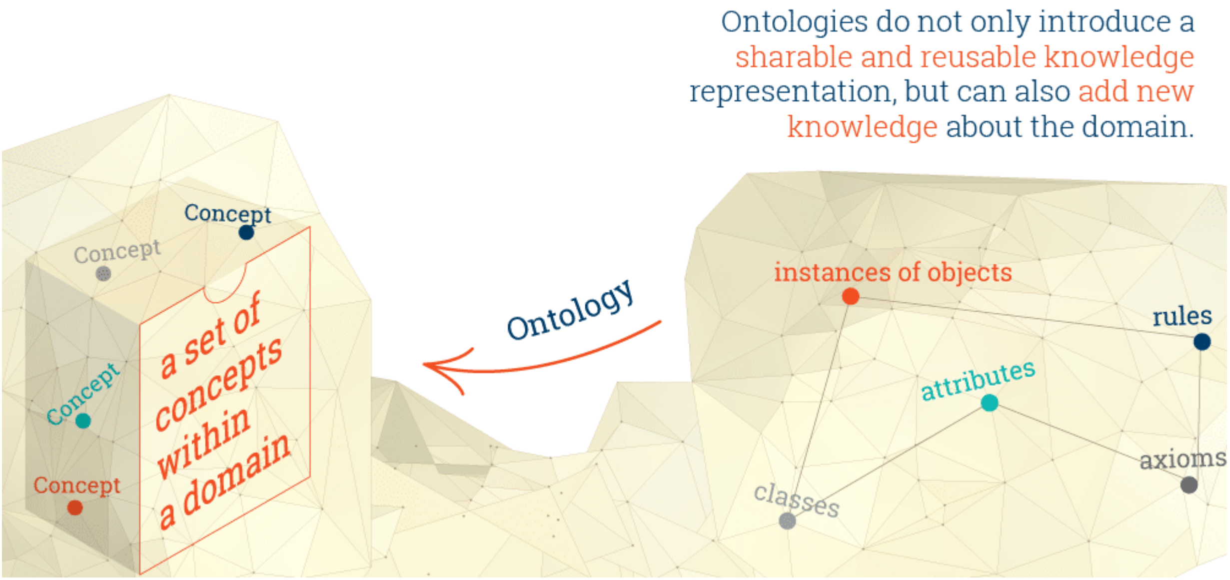

Figure 2.1: What Are Ontologies

2.1.2 purpose of ontologies in AI and medical domains:

Ontologies have become an essential technology of modern healthcare applications, facilitating the semantic modeling and organization of medical knowledge. Integrating with AI systems aids in healthcare data exchange and understanding improving diagnosis, clinical decision-making, and personalized treatment strategies[4].

The heterogeneity and fragmentation of medical data, as well as the lack of standards for unified semantics, represents one major challenge in healthcare IS. To address these challenges, ontologies provide a formal representation of the concepts, attributes and relationships within a domain to improve semantic consistency and interoperability between heterogeneous systems[5].

Ontology-based approaches also provide intelligent reasoning mechanisms to derive new knowledge from existing information. This enables more efficient processing of complex clinical schema, genomic and phenotypic data by AI-driven healthcare applications as the framework provides a means to organize clinical knowledge for sustainable knowledge discovery as well provide better assistance in devices that require medical analyses and decision support systems[6][7].

2.1.3 Role of ontologies in knowledge representation and reasoning:

Ontologies are the semantic basis of KR-R because they provide a human- and machine-readable account of a domain in a structured way. An ontology specifies concepts, properties and logical relations between concepts; in this way it allows artificial intelligence systems to interpret contextual information, share knowledge easily, and infer new facts using reasoning mechanisms[8].

2.2 OWL (Web Ontology Language)

2.2.1 Overview of OWL

The Web Ontology Language (OWL) is a more specific semantic knowledge representation languages that can be used to develop ontologies and describe their vocabularies. An ontology is a formal specification of a set of concepts or categories within a domain and the relationships between those concepts. In this perspective, classes are entities or objects while properties define what happens to these entities and how the entities relate with each other.

Malone OWL is based on formal semantics, which means it provides a mechanism for machines to understand and reason about knowledge in the same consistent way. It is based on the Resource Description Framework (RDF), a W3C standard for representing semantic data or media-type-neutral web-based metadata. OWL has become a key feature in academic research, health care applications, and commercial systems because it is designed to support semantic interoperability and automated reasoning[9].

OWL is a W3C ontology language that builds on description logic, offering more modelling capabilities than RDFS but typically represented and exchanged in RDF/XML syntax. Since OWL allows to express ontologies with different levels of complexity and reasoning requirements, the language is partitioned into three primary sublanguages:

- OWL Lite (basic classification structures and constraints):Used for simple classification hierarchies of which constraints are only supported as values from e.g., 0 or 1.
- OWL DL:More expressive, but retains computational completeness and decidability. And it enforces constraints like disjoint classes, individuals, and properties.

• OWL Full:Full support for RDF and RDFS compatibility, the most expressive version of OWL. Although it is a highly flexible language, its expressiveness comes at the cost of reasoning performance and does not completely ensure decidability for ontology processing[10].

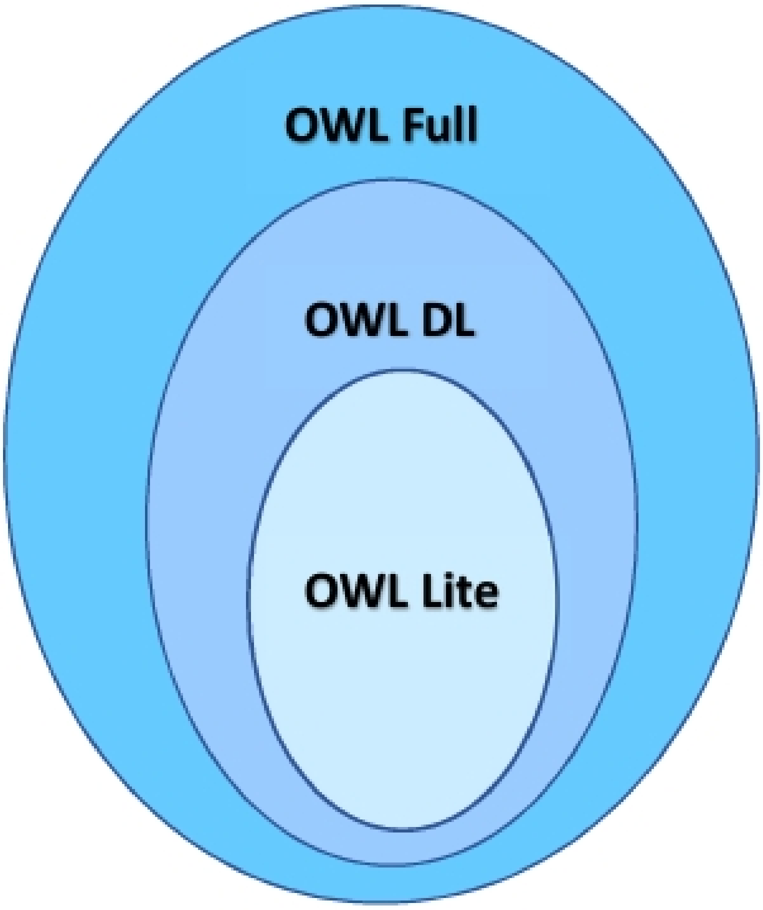

Figure 2.2: The Layers of OWL

2.2.2 Key components:

Core aspects of an OWL ontology comprise multiple interconnected components that together characterize a particular domain's knowledge structure and semantics.

Classes: the features that classify groups or types of entities, like `Brain_Tumor_Type` or `Brain_Tumour_Diagnosis`2.3. They can have a hierarchical ordering in the form of relations like `rdfs:subClassOf`. Besides this, OWL also specify some built-in classes as `owl:Thing`, which encompasses all individuals and class `owl:Nothing`—a class that does not have any instances.

Instances: also referred to as individuals, represent particular objects that type the class. For instance, a biopsy can be modeled as an example of the `Brain_Tumour_Treatment` class.

Properties: denotes relation and properties of entities. There are some kinds of these:

• Object properties: which relate one individual to another individual; for example, hasCausedBy.

• Datatype Properties: properties that link individuals with values of datatypes such as numbers and strings, e.g., has_ICDO-3_Code.

Annotation Properties: Serve as descriptive metadata, such as labels, comments or documentation.

Axioms: they specify the logical constraints and relationships between ontology components. These include class axioms such as equivalence and disjointness, as well as property characteristics such as symmetry, transitivity, and functionality.

In addition, OWL enables expressions through **"Restrictions"** which permits the application of constraints on how properties are used by limiting them through cardinality or value quantification using constructs like 'owl:allValuesFrom'.

syntaxes: There are different syntaxes for writing OWL based on of the use cases. Common formats include RDF/XML that enables better machine processing and interoperability, Manchester Syntax which increases human readability of the model, and Turtle – format compatible with semantic databases/triplestores[11][12].

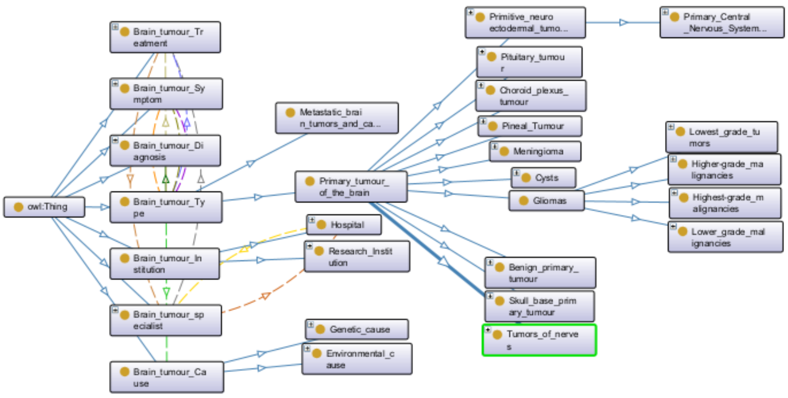

Figure 2.3: Brain Tumour Ontology Hierarchy

2.2.3 Why OWL for medical knowledge representation?:

Recently, the Web Ontology Language (OWL) has been rising to one of the most essential standards for medical knowledge representation as it formally encodes complex biomedical information in a way that is machine-readable and semantically rich. OWL is based on semantic web technologies, allowing healthcare systems and AI solutions to efficiently and accurately represent, organize, and reason over medical knowledge.

Automated Reasoning and Clinical Intelligence:

A benefit of OWL is support for automated logical reasoning via Description Logics. Some examples of queries that can be performed over OWL representations would be: inference about existing data, such as discovering potential drug interactions, recognizing inconsistencies in patient records or supporting the diagnosis of rare diseases based on symptoms and clinical findings. This, in turn, improves the semantic reasoning and inferencing capabilities of healthcare decision-support systems.

Semantic Interoperability in Healthcare:

Commonly, healthcare environments are composed of heterogeneous systems and fragmented medical data. The aforementioned issues regarding interoperability can be solved by OWL, which is [], a knowledge representation language that gives one of the most standardized semantics acknowledged by W3C. OWL allows IT systems—including Electronic Health Records (EHRs)—from different healthcare institutions to exchange and interpret information in a consistent and computable manner through this standardization. At the same time, it also provides support for integrating with major existing medical terminologies and classification, e.g. SNOMED CT and ICD-10.

Rich Representation of Medical Knowledge:

Unlike traditional databases or spreadsheets that can represent data, OWL represents semantics by allowing relationships and logical constraints between entities such as medical items. This enables modeling of relationships between disease causes and anatomical structures, symptoms, treatments, drug interactions, etc. Furthermore, OWL enables the expression of constraints and axioms for better data integrity and semantic accuracy in health care knowledge repositories.

Partial/Partial Clinical Records Across Time:

Despite continuous scientific advances, medical knowledge is always changing, and patient histories can be more often incomplete. OWL works according to its underlying assumption called Open-World Assumption and therefore, information that is hidden here can be neither true nor false. This is particularly important in medicine, where incorrect conclusions should not be drawn from incomplete patient records. In addition, OWL ontologies are modular and scalable, which makes it straightforward for healthcare systems to incorporate new solutions discovered, treatment guidelines generated or biomedicine research undertaken without having to redesign the complete knowledge structure[13] [14] [15] [16] [17] [18] [9].

OWL and AI-Based Medical Applications:

OWL has an extremely important role in modern AI-driven healthcare systems and Knowledge Graph (KG) applications. However, interpretability, reliability, personalization and context-aware reasoning are necessary for AI models in medicine to safely assist clinicians[19]. OWL-based KGs provide semantics for aligning heterogeneous medical data with clinically meaningful and logically consistent representations.

OWL-enhanced AI systems are able to produce predictions with a greater degree of reliability and explainability through the combination of symbolic reasoning and increasingly sophisticated machine learning techniques. This integration is important in oncology and personalized medicine where treatment recommendations are based on the genetics of the patient and presence of various comorbidities[19]. Moreover, OWL-based frameworks ensure greater adaptability, allowing for continuous update of medical knowledge and change of protocols in clinical practice by AI systems.

Real-World Applications:

Some of the health care applications that are built on OWL-based medical ontologies are:

- Clinical Decision Support Systems (CDSS) (all disambiguations with a clear focus to continue citation and the meaning of supported systems respectively)

- Biomedical research and drug discovery.
- Personalized medicine.
- Semantic interoperability among healthcare platforms.
- AI and LLMs based on ground truth in verified medicine.
- Predictive systems in oncology and disease diagnosis based on Knowledge Graph[19].

These features make OWL a robust and versatile technology for representing and reasoning complex medical knowledge in intelligent health care systems.

2.3 Brain Tumour Ontology Structure:

2.3.1 Ontology description used in this project:

The ontology used in this project is a web ontology language (OWL)-based brain tumor domain ontology[20], representing each of the above-mentioned B-table attributes formalized using Protégé editor. It is the knowledge base that lives in a structured and formal manner regarding brain tumors in this medical domain such that there is an encoded conceptual model of the domain which is both shared and machine-readable.

Our ontology consists of some main brain tumor types like Glioblastoma, Astrocytoma, Meningioma and Ependymoma along with other related medical concepts such as symptoms, treatments, MRI findings, malignancy grades and clinical classifications. These classes are linked sequentially through a variety of object properties as well as data properties that define the semantic relationship among medical entities (e.g., relation between a specific tumor type, and its corresponding expected symptoms, malignancy grade, and recommended treatment protocol).

We further enhance the ontology with a set of SWRL (Semantic Web Rule Language) rules that support automated reasoning and inferring new medical knowledge from the detection output of CNN model. Based on this implementation, the ontology is loaded and parsed × Exploring ontology in SPARQL [6] through Python rdflib library; providing ability to query the contents of the ontology (Query language: SPARQL) with defined rules, which creates a complete explainable clinical profile for every analyzed MRI.

2.3.2 Class Hierarchy and Relationships:

The Brain Tumour Ontology The proposed Brain Tumour Ontology is constructed as a hierarchical structure that shows the principal concepts concerning brain tumours and their semantic relationships. This ontology is a class – subclass model where broad concepts are broken up into more specific subclasses to guarantee structured knowledge representation and efficient semantic reasoning.

At the top level, several key classes exist in the ontology for tumor types, symptoms and treatments, causes, diagnoses, specialists and institutions. All of the classes pertain to a specific sub-parts of their brain tumor-related medical domain.

One of the most important components is our tumor type class in the ontology. This can also be split between primary brain tumours ± metastatic brain tumours with carcinomatous meningitis. It can be further divided into primary brain tumor branches which focuses on different type of cancers associated with human health, namely gliomas, meningioma, pituitary tumors, pineal tumors, cysts, nerve tumors skull-base primary tumours and benign primary tumours as well primitive neuroectodermal tumour (Appendix 1). We also created a lower-level subclass of gliomas that includes one for each type of malignancy from lowest-grade tumor to highest-grade malignancy.

Object properties are used in the ontology to model semantic relations between entities. Semantic relations available from tumor type to symptoms or to treatments, causes and specialists include: "availableDoctor", "expectedSymptom", "foaf:expertIn", "foaf:hasExpert", "hasCausedBy", "hasOccurs", "hasSymptoms". "identifiedBy" , "isAvailableIn" , "recommendedTreatment" and "treatedBy". These relations allow the ontology to represent medical knowledge in a meaningful and interrelated manner.

In addition, OWL uses the rdfs:subClassOf property to denote hierarchical relations as inheritance between parent and child classes. The proposed structure enhances semantic interoperability, automated reasoning, and intelligent querying (using SPARQL and RDFLib) in the proposed system.

2.4 SWRL (Semantic Web Rule Language):

2.4.1 SWRL rules definition and syntax:

The Semantic Web Rule Language (SWRL) is a semantic web language intended to extend the expressiveness of the Web Ontology Language (OWL), allowing for formal rules that infer knowledge. OWL 2 adds additional expressiveness to OWL DL or OWL Lite combined with rule-based logic as defined by the RuleML and Datalog so that semantic systems can infer new knowledge from existing knowledge in ontology [21].

SWRL augments the reasoning capabilities of ontologies with rule-based inference mechanisms that are not supported by OWL (OWL is limited to classification). Ontologies in OWL are oriented more towards how to define classes, properties and relations whereas SWRL can infer further facts or semantic relationships using logical rules. The rules have the antecedent–consequent format, like IF-THEN statements[22][23][21].

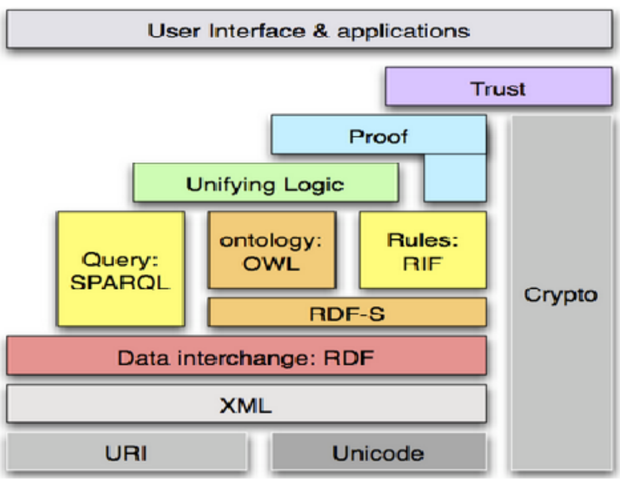

Figure 2.4: Les couches d'architecture du Web sémantique

A SWRL rule consists of two main components: Antecedent (Body): Represents the conditions that must be satisfied for the rule to be applied. Consequent (Head): Represents the new knowledge or relationships inferred once the antecedent conditions are true.

The general syntax of a SWRL rule can be expressed as:

$$\text{Antecedent} \rightarrow \text{Consequent}$$
$$\text{or}$$
$$\text{Body(Atom1, Atom2, ...)} \rightarrow \text{Head(Atom3, Atom4, ...)}$$

SWRL rules are composed of logical expressions known as atoms. Common atom types include:

• Class Atoms: Define class membership, such as Person(?x), indicating that ?x is an instance of the class Person.

• Object Property Atoms: Represent relationships between individuals, such as hasParent(?x,?y).

• Data Property Atoms: Associate individuals with literal values, for example hasAge(?x, ?age).

• Built-in Predicates: Provide mathematical, comparison, or string-processing operations, such as swrlb:greaterThan(?age, 18)[21][22][24][25].

By integrating logical rules with ontologies, SWRL improves semantic reasoning capabilities and supports intelligent decision-making in domains such as healthcare, biomedical systems, and knowledge graph applications.

2.4.2 How Semantic Web Rule Language extends OWL reasoning?:

The Semantic Web Rule Language (SWRL) extends the reasoning capabilities of the Web Ontology Language (OWL) by introducing Horn-like logical rules based on "IF–THEN" structures. While OWL mainly focuses on class hierarchies and semantic relationships using Description Logics, SWRL enhances ontology reasoning by enabling the inference of new relationships and property values through rule-based logic. This extension helps overcome some of the expressive limitations of OWL when dealing with complex semantic interactions between entities[26][27][28][29][30].

How SWRL Extends OWL Expressiveness?:

Rule-Based Inferencing

OWL reasoning is primarily centered on class classification and subclass relationships. In contrast, SWRL allows reasoning over variables and individual instances, enabling the deduction of new semantic relationships from existing ontology data. For example, SWRL can infer an hasUncle relationship by combining multiple object properties and class assertions involving parents and siblings.

Complex Conditions and Built-In Functions

SWRL supports additional logical operations that are not directly handled by standard OWL reasoning. It incorporates built-in predicates for mathematical operations, comparisons, string manipulation, and date processing. These built-ins increase the flexibility of semantic reasoning and allow ontologies to perform more advanced logical evaluations.

SWRL Syntax Structure

SWRL combines OWL DL semantics with RuleML concepts derived from Datalog. A SWRL rule is composed of two main parts:

• Antecedent (Body): Represents the conditions that must be satisfied.

• Consequent (Head): Represents the inferred knowledge generated when the conditions are true.

The general structure of a SWRL rule can be represented as:

$$\text{Body} \rightarrow \text{Head}$$

This structure enables semantic systems to infer additional knowledge beyond traditional ontology classification mechanisms.

The Main Limitation ”Undecidability”:

Although SWRL significantly improves ontology expressiveness, combining description logic with rule-based programming introduces computational complexity. In some cases, reasoning over SWRL rules becomes undecidable, meaning that a semantic reasoner may not always guarantee termination or complete inference results for all possible rule configurations.

This limitation arises because SWRL allows more expressive logical constructs than those supported by standard OWL reasoning alone[21].

Practical Implementations of SWRL:

To address undecidability issues, ontology tools and semantic reasoners adopt several practical strategies for handling SWRL rules.

- **DL-Safe Rules:** Many systems restrict SWRL variables to explicitly defined individuals already present in the ontology knowledge base. These restrictions, known as DL-safe rules, help preserve reasoning decidability while still allowing useful semantic inference capabilities.
- **Hybrid Reasoning Approaches:** Some ontology platforms, such as Protégé, integrate external rule engines like Jess alongside OWL reasoners. In these hybrid architectures, OWL classification is performed first, followed by SWRL rule execution through external inference engines. This approach improves reasoning flexibility while maintaining manageable computational performance.

By extending OWL with rule-based semantics, SWRL provides more powerful knowledge representation and reasoning capabilities for intelligent systems, particularly in domains such as healthcare, biomedical research, and semantic AI applications.

2.5 Ontology Reasoning:

2.5.1 What is reasoning in the context of ontologies?:

Reasoning in ontology systems refers to the automatic process of deriving new knowledge from existing information using logical rules, axioms, and formal semantics. It enables intelligent systems to analyze relationships between concepts and infer implicit facts that are not directly stated within the ontology. Through reasoning, machines can interpret data more intelligently, validate knowledge consistency, and discover hidden semantic relationships automatically2.5[31][32][33].

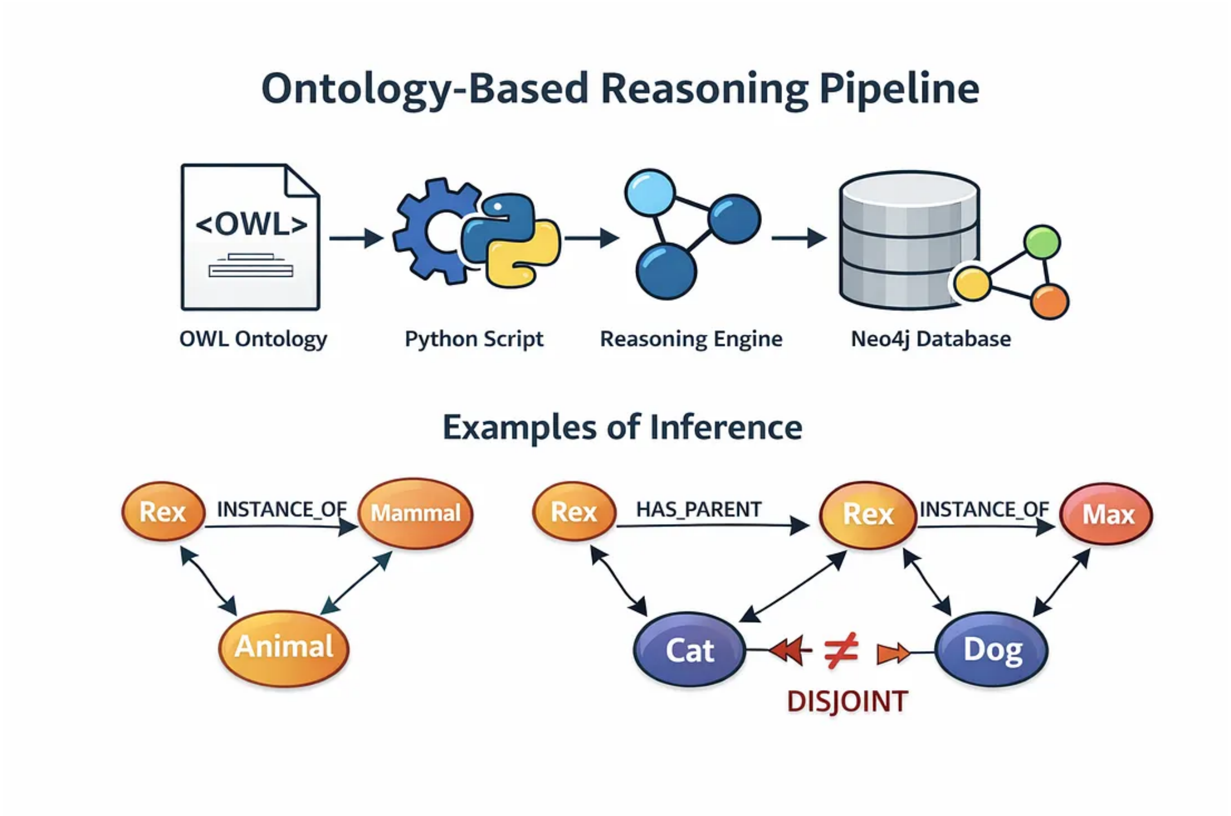

Figure 2.5: Place de SWRL dans le gâteau du WS ...

2.5.2 Types of reasoning:

Reasoning is the mental and logical process by which we use our existing knowledge to arrive at conclusions, produce explanations or make predictions. Reasoning allows to interpret and analyze information to solve problems and support decision-making processes. It is used in human intelligence and artificial intelligence systems. The modes of reasoning are often distinguished into types of logical reasoning and functional cognitive reasoning processes[34][35][36].

Core Types of Logical Reasoning:

Deductive Reasoning:

Deductive reasoning follows a top-down approach in which a general rule or premise is applied to a specific situation in order to derive a logically certain conclusion. If the initial premises are correct, the resulting conclusion is guaranteed to be true.

Example:

All humans are mortal. Socrates is a human. Therefore, Socrates is mortal.

Inductive Reasoning:

Inductive reasoning adopts a bottom-up approach by analyzing specific observations or examples to formulate broader generalizations or patterns. Unlike deductive reasoning, inductive conclusions are probabilistic rather than absolutely certain.

Example:

Every swan observed so far is white. Therefore, all swans are white.

Abductive Reasoning:

Abductive reasoning aims to identify the most plausible explanation for a given observation or incomplete set of facts. It is commonly used in diagnostic systems, scientific discovery, and medical decision-making.

Example:

The grass is wet this morning. Therefore, it probably rained overnight[35][36][37][38].

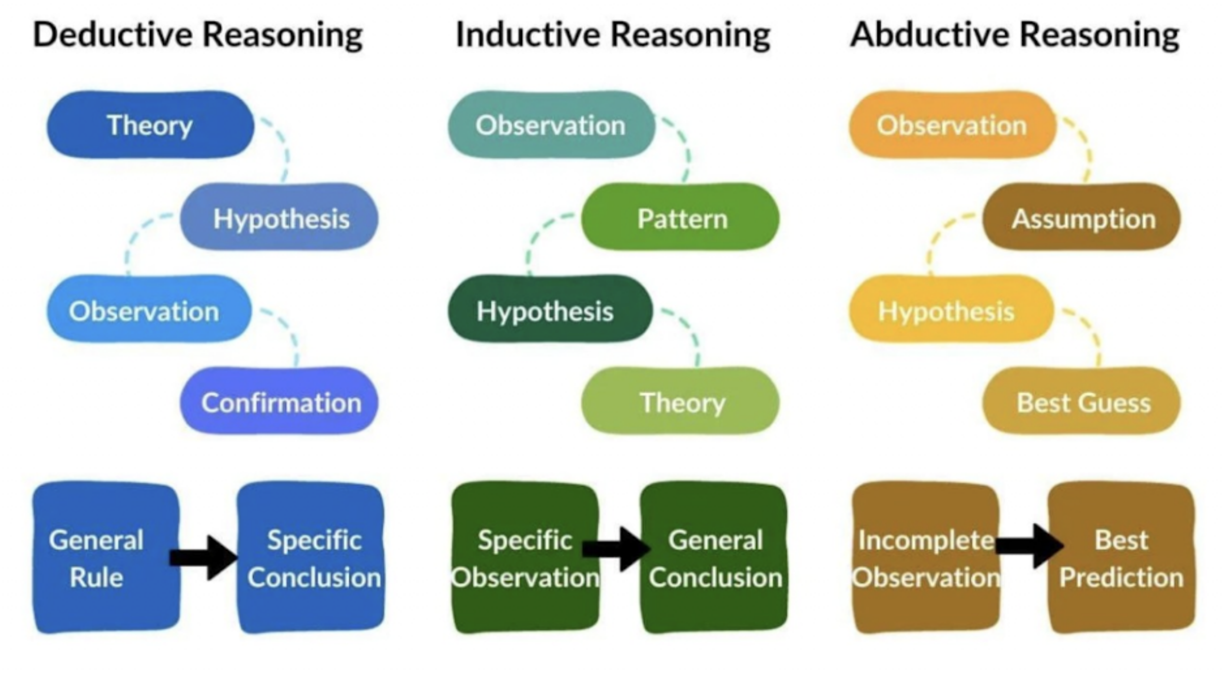

Figure 2.6: Types of Reasoning

Functional Cognitive Reasoning Processes:

Classification:

Classification is the process of organizing data or concepts into categories according to shared characteristics, rules, or semantic similarities. In artificial intelligence and ontology systems, classification helps structure knowledge and identify relationships between entities.

Consistency Checking:

Consistency checking is the process of verifying that information, logical rules, or ontology axioms do not contradict each other. It ensures the integrity, coherence, and validity of semantic knowledge bases and reasoning systems.

Inference:

Inference refers to the process of deriving new knowledge or conclusions from existing facts, observations, or premises. It acts as the fundamental mechanism behind deductive, inductive, and abductive reasoning processes by enabling systems to infer implicit information[36][39].

Interaction Between Reasoning Processes:

In human cognition and intelligent systems, these reasoning mechanisms work together to support complex decision making and knowledge discovery tasks. For instance, a system could start by classifying data into semantic categories, then execute consistency checks to make sure the information is correct, and finally use inference methods to derive conclusions, identify patterns, or generate predictions from the data analyzed[39][40][41].

2.5.3 How reasoning supports the yes/no tumour decision:

Reasoning is essential in the medical domain to support diagnostic and treatment decisions, in particular to determine whether a tumour is malignant or benign and whether medical intervention is needed. Medical reasoning structures complex clinical information into interpretable decision processes with evidence-based and clinically explainable conclusions.

Several interdependent components guide clinicians and intelligent diagnostic systems to accurate medical decisions[42][43][44][45]and the tumour reasoning process is based on them.subsubsection.

1.Reasoning Framework and Chain-of-Thought Analysis:

Medical reasoning does not result in a simple “yes” or “no” output, but rather relies on a structured analytical process that evaluates patient data step by step.

Evidence Extraction:

The starting point is the search for major clinical indicators such as size, shape, growth rate, abnormality in imaging, biomarkers, signs, and symptoms in the patient. These findings form the basis for diagnostic interpretation

Contextual Justification:

The extracted information is linked to current pathological knowledge and clinical guidelines. Such contextual analyses enable clinicians and AI systems to provide medically accepted evidence and semantic relationships to support diagnoses.

2. Diagnostic Reasoning Approaches:

There are two major reasoning strategies commonly used in medical diagnosis to determine the characteristics of tumours and the treatment decisions

Inductive Reasoning:

Inductive reasoning is a process of looking at data that has been observed in medicine to find patterns and draw probable conclusions. For example, physicians may infer the likely nature of a tumour by looking at MRI features, histopathological findings and immunohistochemical markers.

Deductive Reasoning:

Deductive reasoning starts with a set hypothesis and tests the available data on the patient to support or refute the hypothesis. For example, a clinician might initially assume that a tumour is benign and then confirm this with laboratory results, imaging findings and clinical evidence.

3. Case-Based Reasoning and Clinical Guidelines:

Accurate medical reasoning requires comparison of the current patient's case with the medical standards and the past clinical data.

Guideline-Based Evaluation:

Clinical protocols and medical guidelines are used to determine the most appropriate diagnostic or therapeutic actions based on standard healthcare practices.

Historical Case Analysis:

Physicians and intelligent systems often analyze the same past patient cases and treatment results. Comparing imaging findings or pathologic features with previous cases increases confidence in diagnosis and accuracy in treatment planning.

4. Explainability and Human-in-the-Loop Decision-Making:

In oncology and complex medical cases, explainability is a requirement to build trust in diagnostic systems. Medical reasoning offers transparent and interpretable decision processes in which all the inferences, arguments and clinical justifications can be scrutinized.

This transparency allows healthcare specialists e.g. multidisciplinary tumour boards to check the AI-assisted recommendations, assess their medical plausibility and make collective decisions on the most suitable diagnosis or treatment strategy for individual patients.

2.6 Integration with the CNN Model:

2.6.1 How the ontology reasoning layer receives CNN predictions?:

Probabilistic outputs of convolutional neural networks (CNNs) can be used in ontology reasoning systems by converting them into semantic representations understood by ontology reasoners such as OWL and RDF[46][47].

The integration process usually includes three main stages:

1. Prediction Extraction:

Probability scores are calculated by the final layers of the CNN which indicate the confidence level of each detected class or feature in an image.

2. Semantic Mapping:

The probabilistic outputs are converted to symbolic ontology concepts by:

- Employing confidence thresholds.
- Ontology class and property mapping predictions.
- Semantic representation of the detected features such as RDF.

3. Assertion and Reasoning:

The semantic predictions are added to the ontology as ABox assertions. These assertions are then processed by the ontology reasoners like Pellet or HermiT together with the rules and axioms of the ontology (TBox) to infer new knowledge and to check the classifications.

This integration allows intelligent systems to fuse the image recognition capabilities of deep learning with the semantic reasoning for more explainable and knowledge-driven decision making.

2.6.2 The individual creation process (one individual per MRI image):

The individual creation process in MRI is to identify and isolate one object, organ or tissue in medical imaging data. This approach is commonly used in applications such as organ segmentation, tumour analysis and object tracking.

The Technical Process:

1. Volume Acquisition:

The image data from the MRI scanners is a 3D image data, which is reconstructed to image slices with mathematical techniques, e.g., the Fourier transform.

Segmentation:

Image processing techniques are applied to separate the target region from surrounding tissues and background structures.

3. Clustering and Mapping:

Algorithms such as k-means clustering assign a unique label or identifier to each detected object or feature.

4. Volume and Spatial Analysis:

After segmentation, the isolated object can be analyzed in terms of volume, dimensions, and spatial position.

This process enables each tissue or anatomical structure to be represented and tracked as an independent entity within the medical imaging dataset[48][49][50][51].

2.6.3 How SWRL rules refine the CNN output into a validated final decision?:

The SWRL rules improve the CNN predictions by providing a semantic reasoning layer based on ontology knowledge. CNN models offer probabilistic classification. SWRL uses logical rules and domain constraints to check and improve these predictions.

1. Semantic Mapping:

CNN outputs are converted into OWL ontology concepts such as classes, individuals and properties, so that the prediction results can be understood semantically by the ontology systems.

2. Rule Definition:

SWRL rules specify logical ”IF-THEN” relationships between concepts in the ontology and contextual information to validate CNN classifications.

3. Reasoning and Validation:

Ontology reasoners such as Pellet and HermiT check the asserted assertions against the ontology axioms and SWRL rules in order to find inconsistencies and verify the results.

4. Decision Refinement:

The reasoning engine produces a final decision that is mathematically and logically validated by:

• Confirming: If all SWRL constraints are satisfied, the CNN’s classification is promoted to a validated, trusted decision.

• Overriding/Correcting: If the CNN classification contradicts a SWRL rule, the reasoner adjusts or flags the decision based on the strict logical truth.

• Inferring new data: The engine can deduce complex, multi-step conclusions that the CNN was not directly trained to detect, such as concluding fault in a vehicle accident or calculating a specific treatment regimen

This hybrid approach combines deep learning with semantic reasoning to produce more accurate and explainable decisions[52][53][54].

2.6.4 Advantages of combining deep learning with semantic reasoning:

Semantic reasoning-based deep learning leads to hybrid AI systems that combine data-driven learning with rule-based knowledge representation.

1. Improved Explainability:

Semantic reasoning improves the transparency of AI decisions by linking predictions to logical rules and ontology knowledge, thus alleviating the "black box" phenomenon in deep learning models.

2. Reduced Data Requirements:

Ontologies and knowledge graphs can provide pre-defined domain knowledge, which can reduce the amount of training data needed for deep learning models.

3. Error Reduction:

Semantic rules are useful for identifying and fixing inconsistent or illogical predictions produced by neural networks, thus improving the reliability of decisions.

4. Contextual Understanding:

Semantic reasoning allows AI systems to learn the relationships between concepts, infer new knowledge and process contextual information more efficiently.

5. Continuous Knowledge Updating:

Semantic knowledge bases can be extended by adding new rules or facts, without re-training the whole neural network model[55][56][57].

Chapter 3

Detection Using AI

3.1 Introduction:

Artificial Intelligence (AI) is now a game-changing technology that can analyze large amounts of data and automatically identify patterns, objects, anomalies or events that may be difficult for humans to identify. AI-based detection systems use machine learning and deep learning algorithms that can learn from historical data and improve their performance over time with experience. Such systems can detect complex patterns, classify information and recognize abnormal behavior with high accuracy and efficiency.

Anomaly detection is one of the core capabilities of AI detection. It is the process of finding observations or data points that are different from the normal behavior. AI models learn the characteristics of normal patterns and are able to automatically detect unusual events that may indicate medical abnormalities, security threats, equipment failures, or fraudulent activities. The ability to do so means an organization can react faster and more efficiently to critical situations.

AI-based detection technologies are used in many fields, including medical imaging, computer vision, cybersecurity, fraud detection, and industrial monitoring. Deep learning models such as Convolutional Neural Networks (CNN) have been used extensively in healthcare for medical image processing for disease detection and diagnosis. In the field of cybersecurity, AI systems are constantly scanning network activity and identifying malicious behaviors, anomalies, and new threats in real time.

AI detection systems have seen significant improvement in performance thanks to the increasing availability of large datasets, as well as advances in computational power and machine learning techniques. So, AI-based detection is now a must-have element for modern intelligent systems, offering better accuracy, faster decisions, less human workload, and better support for complex analytical tasks[58][59][60][61][62][63].

3.2 Artificial Intelligence in Medical Imaging:

3.2.1 From Traditional Methods to Deep Learning:

Traditional medical image analysis was based on the handcrafted feature extraction techniques such as edge detection, texture analysis and image segmentation followed by machine learning classifiers such as Support Vector Machines (SVM) and Decision Trees. However, these methods result in satisfactory results at the expense of heavy reliance on expert designed features and generalization robustness over datasets.

Deep learning has transformed medical image analysis by allowing automatic feature extraction directly from raw images. In particular, Convolutional Neural Networks (CNNs) have achieved great success on such tasks as image classification, detection and segmentation. Deep learning models have gained significant accuracy improvement through learning hierarchical representations from data, and have become the dominant approach in modern computer-aided diagnosis systems[64][65].

3.2.2 Why CNNs for MRI Analysis:

The Convolutional Neural Networks (CNNs) have emerged as the most popular deep learning architecture for MRI analysis due to their ability to learn meaningful features automatically from medical images. Traditional machine learning methods rely on hand-crafted features designed by experts, while CNNs learn image characteristics of interest from the training data. This ability reduces the need for manual feature engineering and enables the model to learn complex patterns that may be difficult for human observers to identify.

MRI images are rich in spatial and structural information which is very important for the detection of abnormalities such as brain tumors. CNNs are particularly well suited to analyze such data as their convolutional layers can capture local patterns, textures, edges and shapes while preserving spatial relationships between neighboring pixels. With increasing network depth, it learns more abstract representations to help distinguish between healthy and abnormal tissues.

CNNs are also good at dealing with large and complex imaging datasets in an efficient manner. Other techniques such as pooling and parameter sharing reduce the computational complexity of CNNs while maintaining high predictive performance. The robustness and high accuracy have led to significant improvements in medical image classification, detection and segmentation tasks.

CNNs have thus become a fundamental part of modern computer-aided diagnosis systems and can provide reliable support to clinicians in the detection of diseases, treatment planning, and medical decision-making[64][65].

3.3 Brain Tumour Detection with CNNs:

Medical image analysis for detection of brain tumour is one of the most widely adopted approaches using Convolutional Neural Networks (CNNs). CNNs can automatically learn relevant features from MRI images, thus identifying tumour-related patterns without the need for manual feature extraction. CNNs can distinguish normal from tumour-affected tissues with accuracy by analyzing spatial and structural information in brain scans.

The high accuracy, efficiency and capability of CNNs to process large datasets has made it an effective tool for supporting early diagnosis and clinical decision making in brain tumour detection .

3.3.1 Binary Detection vs. Multi-Class Classification:

Deep learning models can be used for binary detection or multi-class classification in brain tumour analysis . Binary detection classifies an MRI image as having or not having a tumour, resulting in two possible outcomes. However, the multi-class classification considers the specific type of tumour, i.e. glioma, meningioma or pituitary tumour.[66][67].

Table 3.1: Binary and Multi-Class Classification Differences

Feature[67][68][65]	**Binary Detection**	**Multi-Class Classification**
Number of Classes	Exactly two (e.g., $Class_A$ or $Class_B$).	Three or more (e.g., $Class_A$, $Class_B$, or $Class_C$).
Common Algorithms	Logistic Regression, Support Vector Machines (SVM), Perceptrons.	Random Forest, Naive Bayes, Decision Trees, Neural Networks.
Output Activation	Sigmoid function (outputs a single probability between 0 and 1)	Softmax function (outputs a probability for each class, summing to 1).
Complexity	Low; requires setting one decision boundary.	High; requires distinguishing between multiple boundaries simultaneously.
Evaluation Metrics	Precision, Recall, F1-Score, ROC-AUC.	Macro/Micro Precision, Confusion Matrix.

3.3.2 Transfer Learning Strategies (VGG16, ResNet, DenseNet):

Transfer learning applies knowledge from large datasets like ImageNet to new tasks, saving training time and increasing performance. The selection of the strategy depends on the architecture of the pre-trained mode[69][70].

VGG16 (Visual Geometry Group):

VGG16 is a deep convolutional network with sequential (3×3) convolutional layers. It features a simple architecture that makes it easy to implement, but it has a large number of parameters.[71][72].

- Best Strategy – Fine-tuning Deeper Layer: The early convolutional layers are usually frozen because they extract generic visual features such as edges and textures. The higher layers are then fine-tuned to learn task specific patterns from the target dataset.
- Best Use Case: Suitable for image classification tasks with a relatively small dataset and visually distinguishable patterns.

ResNet (Residual Network):

ResNet adds residual connections to skip some layers so that very deep networks can be trained without the vanishing gradient problem[73].

- Best Strategy – Feature Extraction and Gradual Unfreezing: Initially, the pre-trained network is used as a feature extractor by freezing its layers and training a new classification head. Then, the selected top layers can be gradually unfrozen and fine-tuned with a small learning rate.
- Best Use Case: Most suited for complex image classification tasks where deeper feature representations and better training stability are required.

DenseNet (Densely Connected Network):

DenseNet, where each layer is directly connected to all subsequent layers, facilitates feature reuse and efficient information flow across the network[74].

• Best Strategy – Global Fine-Tuning Due to the shared features among the whole network, DenseNet generally enjoys the fine-tuning of most or all layers with a small learning rate.

• Best Use Case: Especially useful in medical imaging and other situations with limited data, where good reuse of features can enhance performance and generalization.

3.3.3 Data Augmentation and Preprocessing for MRI:

Data augmentation and pre-processing are important steps in preparing MRI images for deep learning models .

They enhance the quality of images, minimize variability, and aid model performance.[75][76]

Data Preprocessing:

MRI data is very heterogeneous with respect to scanner vendors, acquisition protocols and patient motion. Preprocessing homogenizes the data within and between subjects [77].

• Skull Stripping (Brain Extraction): Removes non-brain tissues (skull, scalp and fat) so the model can focus on the brain.

• Registration/Alignment: Register multi-modal MRI scans (T1, T2, FLAIR) to a standard stereotaxic space (e.g., MNI template) or into a common coordinate system to create a voxel-to-voxel correspondence

• Bias Field Correction (N4/N3): Corrects low frequency smooth intensity variations in the image due to magnetic field inhomogeneities.

• Intensity Normalization: Since the pixel values in MRI do not have an absolute physical meaning (e.g., Hounsfield units in CT), the intensity ranges are scaled to a standard mean and standard deviation (e.g., z-score normalization) or clipped between $[0, 1]$.

• Resampling: Scales all voxel sizes and matrix sizes in the dataset to be the same, for a consistent spatial resolution.

Data Augmentation:

Because of the small size and class imbalance of annotated MRI datasets augmentation is used to artificially increase the size and diversity of the training data. This helps the model learn spatial and intensity invariances[75][78].

• Spatial Transforms: Random rotations, translations, scaling, flip horizontal/vertical. Elastic deformations are important for tumor segmentation, because lesions rarely have uniform shapes.

• Intensity / Luminance Variations : Varying brightness / contrast , or adding Gaussian noise to simulate the grainy nature of the real world scans .

• Artifact Simulation: TorchIO and other specialized libraries simulate MRI-specific artifacts (e.g., ghosting, i.e., displaced structures, spiking, localized blurring)

• Generative Models: Use of Generative Adversarial Networks (GANs) and latent diffusion frameworks to generate novel 3D MRI volumes, and augment underrepresented pathology classes (e.g., rare tumors).

3.3.4 Evaluation Metrics in Medical Image Classification:

Medical image classification requires specialized evaluation metrics beyond conventional accuracy because of extreme class imbalance (e.g., rare diseases) and the high clinical cost of misdiagnoses.

The most reliable evaluations combine threshold-based metrics (such as Sensitivity) with threshold-independent measures (such as AUROC) to ensure safe and clinically viable diagnostic AI[79][80].

Confusion Matrix Fundamentals:

The confusion matrix is the basis of classification evaluation, which compares predicted labels of a model with clinical ground truth labels[79][80]. It produces four core outcomes:

It has four main outcomes:

- True Positives (TP): The model correctly predicts the presence of a condition.
- True Negatives (TN): Model correctly identifies no presence of a condition.
- False Positives (FP): Model mistakenly identifies a healthy scan with a condition.
- False Negatives (FN): Model incorrectly misses a condition in an unhealthy scan [81].

Threshold-Based Metrics:

These metrics rely on a fixed classification threshold (typically 0.5) and are derived directly from the confusion matrix.

- Sensitivity (Recall or True Positive Rate):Sensitivity $= \frac{TP}{TP+FN}$

Clinical Importance: Measures the model's ability to detect all positive cases. Critical in screening scenarios where missing a disease (False Negative) is life-threatening.

- Specificity (True Negative Rate):Specificity $= \frac{TN}{TN+FP}$

Clinical Importance: Measures the model's ability to correctly identify healthy individuals. Crucial to avoid unnecessary follow-up procedures or treatments caused by False Positives.

- Precision (Positive Predictive Value):Precision $= \frac{TP}{TP+FP}$

Clinical Importance: The probability that a positive prediction is truly correct.

- Accuracy:Accuracy $= \frac{TP+TN}{TP+TN+FP+FN}$

Limitation: Highly misleading in medicine. If a dataset is 99% healthy patients, a model that simply predicts "healthy" every time achieves 99% accuracy but provides zero clinical value.

- F1-Score:F1-Score $= 2 \times \frac{\text{Precision}\times\text{Sensitivity}}{\text{Precision}+\text{Sensitivity}}$

Clinical Importance: The harmonic mean of precision and recall. Useful when you need to balance both metrics in imbalanced datasets.

- Matthews Correlation Coefficient (MCC):MCC $= \frac{TP\times TN-FP\times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$

Clinical Importance: Considered one of the most robust and reliable single scores, as it takes all four quadrants of the confusion matrix into account.

Threshold-Independent Metrics:

These metrics measure the performance of the model at all possible probability thresholds as medical decisions are often made by changing the risk thresholds to be more sensitive or more specific.

- AUROC (Area Under the Receiver Operating Characteristic): Interpretation: Draw True Positive Rate against False Positive Rate. If the classifier is perfect, the AUROC is 1.0 and 0.5 means random guessing . It quantifies the overall discrimination ability of the model between classes, regardless of the operating point.
- AUPRC (Area Under Precision-Recall Curve)Interpretation: Frequently more informative than AUROC for heavily skewed medical datasets because it heavily weights the positive class and avoids the overly optimistic inflation that results from huge numbers of true negatives .

Guidelines for Application:

When assessing an algorithmic model for clinical deployment, always evaluate using a combination of metrics rather than a single score:

- Prioritize Sensitivity/Specificity tailored to the clinical workflow. (e.g., Cancer screening requires near-100% sensitivity).
- Never report Accuracy alone for classification tasks.
- Use AUROC and AUPRC to measure the inherent discriminatory power of your algorithm before finalizing a clinical decision threshold.
- Patient Safety and Clinical Accountability: Clinicians remain responsible for diagnosis and treatment decisions. Explainable AI allows them to verify and validate AI-generated recommendations before applying them in practice.
- Bias Detection and Fairness: Transparent AI systems help identify potential biases in training data and ensure that predictions are fair across different patient populations.

3.4 Limitations of Pure CNN Approaches:

Pure CNN approaches are fundamentally limited by their reliance on local receptive fields and pooling operations, which prioritize local textures over global context. This causes them to struggle with spatial relationships, long-range dependencies, and structural understanding, while requiring massive amounts of data and compute to overcome overfitting.

3.4.1 Black-Box Problem:

The black-box problem in Convolutional Neural Networks (CNNs) refers to the inability to interpret the internal reasoning behind the model's decisions. While highly accurate at extracting features like edges, textures, and shapes, the millions of interconnected weights in hidden layers make it impossible to explicitly map out how an output is achieved[82].

Key Limitations Caused by the Black-Box Problem:[83]

- Trust and Adoption: In mission-critical sectors such as healthcare, aviation, and finance, stakeholders and users are hesitant to rely on AI recommendations if the system cannot justify its outputs.
- Safety and Debugging: When a pure CNN fails or makes a misclassification, it is incredibly difficult to trace exactly why the error occurred. This makes troubleshooting, error correction, and ensuring robustness nearly impossible.
- Bias and Fairness: Opaque networks can silently adopt and amplify biases hidden within their training data. Without transparency, it is hard to verify the fairness or impartiality of the model's choices.
- Regulatory and Legal Hurdles: Frameworks like the EU's GDPR mandate a "right to explanation," making uninterpretable models non-compliant in scenarios that significantly impact individuals.

Current Approaches to Overcome the Limitation:[84]

Researchers are actively developing Explainable AI (XAI) techniques to open the black box. Common methods include:

- Saliency Maps (e.g., Grad-CAM): Visual overlays that highlight specific regions or ”hot zones” in an image that most strongly influenced the CNN’s prediction.
- Feature Attribution (e.g., SHAP, LIME): Mathematical methods used to estimate the impact of each individual pixel or feature on the final output.
- Inherently Interpretable Architectures: Designing models that use rule-based logic, decision trees, or concept bottleneck representations rather than purely abstract deep features.

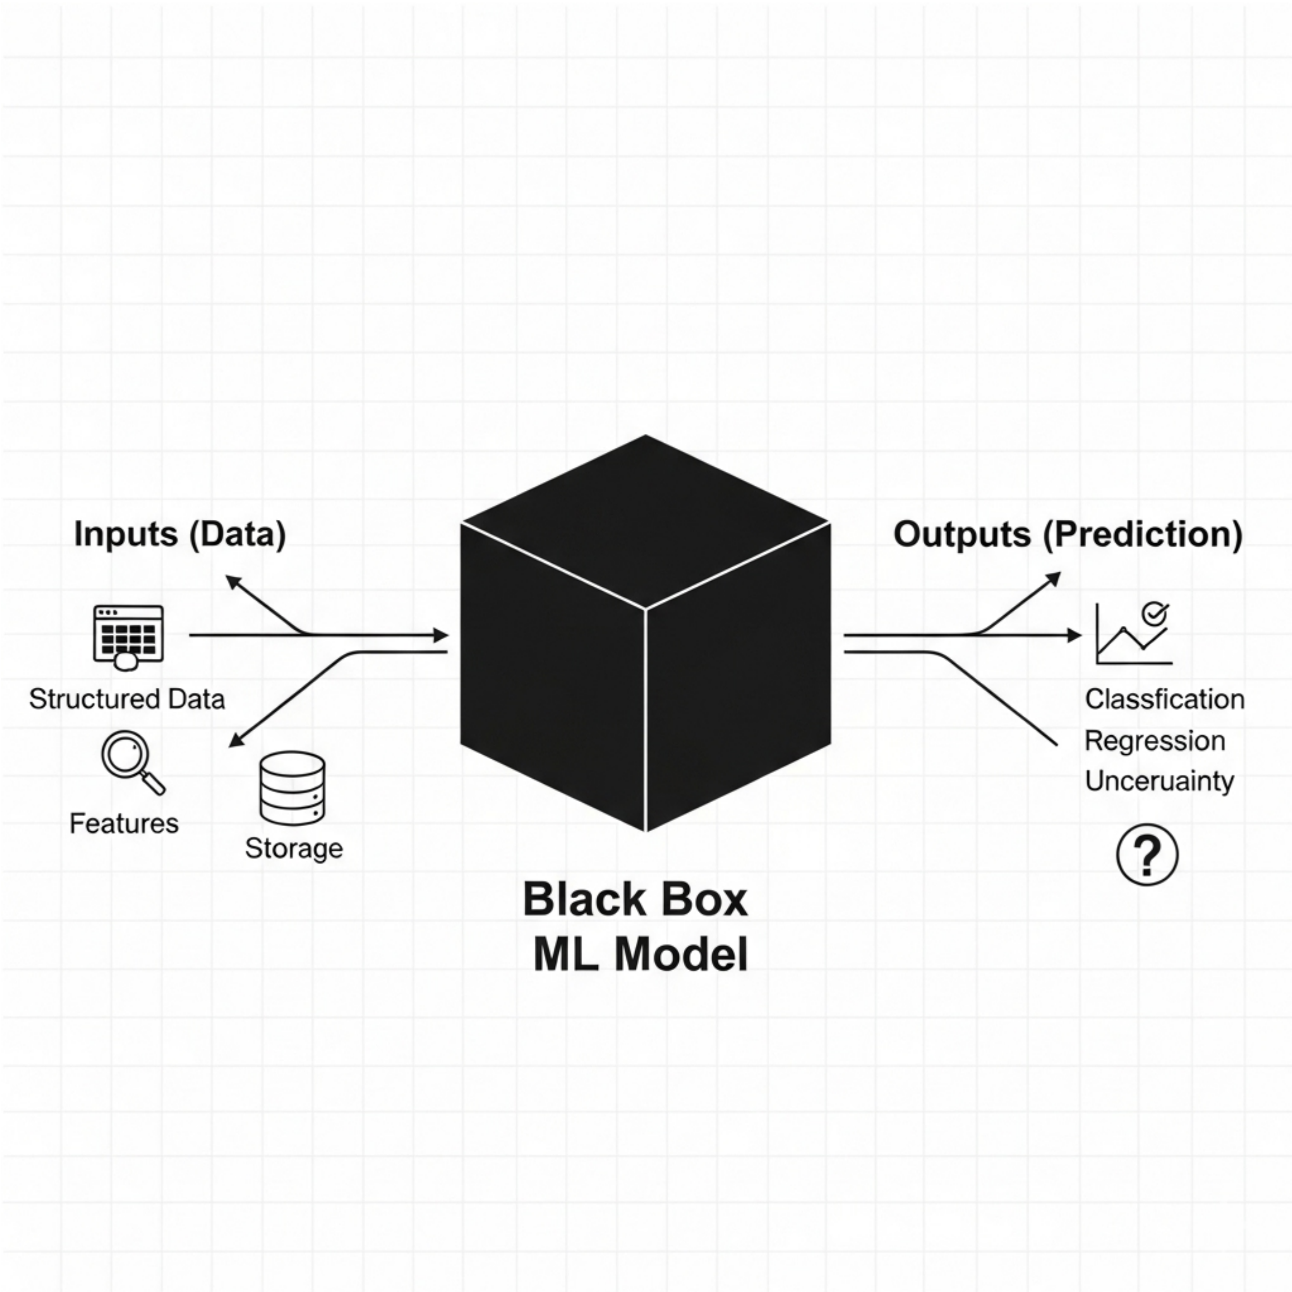

Figure 3.1: Black-box problem in pure CNN approaches

3.4.2 Absence of Clinical Context in Predictions:

Many AI-based predictive models generate results solely from input data, such as medical images, without considering important clinical information like patient symptoms, medical history, or laboratory findings. Although these models can achieve high accuracy, their predictions may lack the clinical context needed for informed medical decision-making. This limitation highlights the need to integrate medical knowledge and reasoning mechanisms to improve the interpretability and reliability of AI-assisted diagnosis[85].

Why Clinical Context Matters?:

- Static vs. Dynamic Reality: Clinical practice is an iterative loop. Physicians adjust treatments based on real-time responses, whereas most models view patient trajectories as static, one-way processes[86].
- Missing Data can be "Informative": In electronic health records (EHRs), data gaps often occur intentionally because a clinician deemed a test unnecessary for a stable patient. Algorithms that simply interpret missing values as "abnormal" or "zero" can generate skewed risk scores[87].
- The Self-Fulfilling Prophecy: Historical datasets often reflect the biases of previous clinicians (e.g., opting for less aggressive care for older patients). If a model learns from these flawed historical outcomes, it risks reinforcing or automating those biases[88].

The Consequences:

Relying solely on mathematical predictions without context leads to tangible risks in healthcare delivery:

- Decision Fatigue: Clinicians frequently face "alert fatigue" when models generate excessive false positives or fail to provide a calibrated sense of case-level uncertainty[89].
- Protocolized Oversimplification: Risk categorizations (e.g., low, medium, or high risk) can force undesirable behaviors. Doctors or nurses might treat patients differently based on arbitrary thresholds rather than assessing the holistic picture[90].
- Overgeneralization: Models trained on specific demographic groups or heavily resourced hospital systems frequently fail to generalize when applied to broader, diverse patient populations or different clinical settings.

Moving Toward Context-Aware AI:

To bridge the gap between "statistical accuracy" and "clinical utility," the medical and AI research communities are increasingly focused on the following advancements:

- Large Language Models (LLMs): Rather than relying only on structured fields like lab results, modern approaches are integrating LLMs to read unstructured clinical notes—often the richest source of subjective patient history—to grasp nuanced contextual variables[89].
- Continuous Risk Evaluation: Researchers advocate for retaining continuous predicted risks rather than grouping patients into strict, rigid categories, allowing healthcare providers to apply their own contextual judgment alongside the model[90].
- TRIPOD Guidelines: To ensure models are transparently developed and validated before clinical deployment, researchers follow established reporting checklists like TRIPOD (Transparent Reporting of a multivariable prediction model for Individual Prognosis or Diagnosis)[91].

3.5 Conclusion:

Artificial Intelligence has made a huge impact on detection systems by providing the ability to automatically identify patterns, anomalies and objects in different fields. AI based models can process large amount of data efficiently and can also achieve high level of accuracy through machine learning and deep learning techniques. AI systems like Convolutional Neural Networks (CNNs) have delivered remarkable performance in medical imaging, particularly for brain tumor detection, by analyzing MRI scans and assisting in clinical diagnosis.

Despite these developments, key challenges persist including poor interpretability and absence of clinical context. Therefore, the combination of AI models with domain knowledge and semantic reasoning approaches can increase the reliability and transparency of predictions. As research progresses, the role of AI-based detection systems will become increasingly important in improving the decision process and supporting professionals in healthcare and other application areas.

Chapter 4

Related Works

The ever-increasing availability of medical imaging data and developments in Artificial Intelligence (AI) have attracted considerable attention from researchers for brain tumor detection and diagnosis. In the last decade, several studies have been dedicated to the development of automated systems to help clinicians in the detection, classification and analysis of brain tumors from MRI images. The objective of these systems is to improve the diagnostic accuracy, reduce analysis time and support clinical decision making.

The early work mainly focused on traditional machine learning techniques combined with hand-crafted feature extraction methods. However, the advent of deep learning especially Convolutional Neural Networks (CNNs) has significantly improved the performance of medical image analysis systems. CNN-based approaches have been proven to have high accuracy in tumor-related pattern identification and have become the dominant methodology for brain tumor detection and classification.

Also, ontologies-based systems have been developed to encode medical knowledge in a structured and machine-readable form. These systems encode semantic relations among diseases, symptoms, diagnostic procedures and treatments, thus providing reasoning capabilities that support medical decision-making and improve interpretability of diagnostic results.

Some recent researchers have investigated hybrid approaches that combine the predictive power of deep learning models and the explainability and reasoning of medical ontologies. These systems aim to overcome the limitations of purely data-driven models by integrating clinical knowledge into the diagnostic process. The combination improves transparency, allows context-aware decision-making, and delivers more trustworthy and clinically relevant results.

The next section is devoted to a review of the most relevant studies related to CNN based brain tumor detection, ontology based medical decision support, and hybrid AI-semantic systems, discussing the methodologies, achievements, and limitations of each.

4.1 CNN-Based Brain Tumour Detection :

4.1.1 International (2024–2026):

This subsection reviews recent international CNN-based brain tumour detection studies[92][93][94].

CNN-TumorNet + LIME (2025):

One of the most related studies is the Hybrid CNN-TumorNet framework (2025), which achieved 99% accuracy in binary brain tumor detection from MRI images. The system uses LIME to highlight image regions that influence the CNN's predictions, providing visual explanations of the results.

However, LIME explanations are limited to pixel-level interpretation and do not generate clinical knowledge or decision-support information. In contrast, the proposed system combines CNN-based detection with OWL ontology and SWRL reasoning to produce semantic clinical reports, treatment recommendations, and ontology-supported conclusions, resulting in more meaningful and explainable diagnostic outcomes[95].

MRF-Optimized CNN (2026):

A 2026 study proposed a customized four-block CNN optimized using the Manta Ray Foraging Optimization (MRFO) metaheuristic, achieving 98.5% accuracy in multi-class brain tumor classification.

The model was evaluated on four classes—glioma, meningioma, pituitary tumor, and no tumor—and outperformed standard transfer learning architectures such as ResNet50, which achieved 93.2% accuracy on the same dataset. The use of the bio-inspired MRFO algorithm enabled effective optimization of network parameters, contributing to the model's strong classification performance.

Despite its high accuracy, the approach remains focused exclusively on image-level feature learning and classification. It does not provide explanations for its predictions, communicate uncertainty, recommend clinical actions, or relate its outputs to established medical knowledge. The absence of a semantic knowledge representation and reasoning layer means that its decisions remain largely opaque and non-auditable. In contrast, the proposed system combines CNN-based detection with ontology-driven SWRL reasoning to generate interpretable clinical reports, treatment recommendations, confidence assessments, and knowledge-based diagnostic conclusions[96].

YOLOv7 Transfer Learning (2025):

A 2025 study fine-tuned the YOLOv7 object detection architecture for brain tumor detection and localization in MRI images, achieving 99.5% accuracy across the three standard tumor classes. Unlike conventional CNN classifiers, YOLOv7 treats the task as an object detection problem, providing bounding boxes that identify the spatial location of tumors within the image. This localization capability offers valuable visual information that is not available in the first version of the proposed system.

However, the YOLOv7 model focuses solely on detection and localization. It does not generate semantic clinical information such as tumor grading, treatment recommendations, clinical alerts, treatment urgency assessments, or reasoning-based explanations. Consequently, while YOLOv7 excels in spatial precision, it lacks the interpretability and clinical decision-support capabilities provided by ontology-based reasoning. The two approaches can therefore be viewed as complementary rather than competing, with YOLOv7 potentially serving as a future localization component integrated into the proposed CNN–ontology framework[97].

Lightweight CNN for Small Datasets (2025):

A 2025 study proposed a five-layer lightweight CNN architecture designed for brain tumor detection in resource-constrained environments, achieving 99% accuracy on a small MRI dataset. The model demonstrates that effective tumor detection can be achieved without deep architectures, large training datasets, or extensive computational resources, making it suitable for deployment in hospitals and clinics with limited infrastructure.

Despite its efficiency and strong predictive performance, the approach remains focused solely on classification and provides only a tumor/no-tumor prediction. It does not offer explainability, clinical contextualization, knowledge grounding, or reasoning-based support for medical decision-making. In contrast, the proposed system prioritizes clinical interpretability by integrating CNN-based detection with ontology-driven SWRL reasoning, enabling the generation of semantic clinical reports, treatment recommendations, and transparent diagnostic explanations[98][99].

DenseNet / EfficientNet / InceptionResNetV2 (2025):

A 2025 comparative study evaluated several pre-trained CNN architectures, including DenseNet201, EfficientNetB5, and InceptionResNetV2, combined with Principal Component Analysis (PCA) for feature dimensionality reduction. The study reported accuracy values of up to 100% on specific dataset subsets and demonstrated that fine-tuned transfer learning models consistently outperform networks trained from scratch for brain tumor classification. It also provided valuable insights into the selection of suitable backbone architectures for medical imaging tasks.

The VGG16 model used in the proposed system belongs to the family of architectures evaluated in this benchmark, and its fine-tuning strategy aligns with the study's recommendations. However, while the comparison focuses on maximizing classification performance, it does not address the interpretability of predictions or their integration with clinical knowledge. In contrast, the proposed system extends beyond image classification by combining CNN predictions with ontology-based SWRL reasoning to generate explainable and clinically meaningful diagnostic outcomes[100].

4.1.2 Algeria (2022–2026)

This subsection situates the proposed system within the Algerian academic research context [101] [102] [103]

MRI-based brain tumor prediction using convolutional neural networks (2026):

A 2026 study on MRI-based brain tumor prediction using Convolutional Neural Networks (CNNs) focused on developing a lightweight and computationally efficient model for tumor identification. The primary objective was to create a diagnostic system suitable for low-resource healthcare environments where access to high-performance computing infrastructure and large-scale datasets may be limited. The proposed CNN architecture was designed to reduce computational complexity while maintaining high classification accuracy, demonstrating that effective brain tumor detection can be achieved without relying on very deep or resource-intensive networks.

The study highlights the potential of lightweight deep learning models to improve accessibility to AI-assisted diagnosis in hospitals and clinics with limited technological resources. However, the approach remains centered on image-based prediction and classification. It does not provide mechanisms for explaining predictions, integrating medical knowledge, generating clinical recommendations, or supporting diagnostic reasoning. In contrast, the proposed system combines CNN-based detection with an OWL ontology and SWRL reasoning framework, enabling not only tumor prediction but also the generation of semantic clinical reports, treatment recommendations, confidence assessments, and ontology-supported diagnostic conclusions[94].

AI-Enhanced Framework for Automated Brain Tumor Detection (2025):

An AI-Enhanced Framework for Automated Brain Tumor Detection proposed in 2025 combined advanced AI techniques for MRI image segmentation and classification, achieving high accuracy in identifying and delineating tumor regions. The framework was designed to automate the analysis process by first locating suspicious tumor areas and then classifying the detected abnormalities. By integrating segmentation and classification within a single pipeline, the system improved the precision of tumor detection and reduced the need for extensive manual intervention.

The study demonstrates the effectiveness of AI-based image analysis for detecting brain tumors and accurately identifying tumor boundaries. However, its focus remains on image processing and predictive performance. The framework does not provide semantic reasoning, clinical interpretation, treatment recommendations, or ontology-based validation of its results. In contrast, the proposed system extends beyond detection by integrating CNN predictions with OWL ontology and

SWRL reasoning, enabling the generation of explainable clinical reports, confidence assessments, and knowledge-driven diagnostic conclusions[101].

Brain Tumor Detection Using U-Net and SVM (2025):

A 2025 Master's thesis from the University of Bordj Bou Arreridj proposed a hybrid approach for brain tumor analysis by combining U-Net and Support Vector Machines (SVM). In this framework, U-Net was employed for precise tumor segmentation from MRI images, enabling the extraction of tumor regions and boundaries, while SVM was used for tumor grading and classification based on the segmented outputs. By integrating deep learning-based segmentation with a traditional machine learning classifier, the study aimed to improve both tumor localization and diagnostic accuracy.

The approach demonstrated the effectiveness of combining complementary techniques for medical image analysis, benefiting from U-Net's strong segmentation capabilities and SVM's classification performance. However, the system remains focused on image processing and classification tasks and does not incorporate semantic reasoning, clinical knowledge representation, or decision-support mechanisms. In contrast, the proposed system integrates CNN-based detection with ontology-driven SWRL reasoning to generate explainable clinical reports, treatment recommendations, confidence assessments, and knowledge-based diagnostic conclusions[104][105].

Brain Tumor Detection using CNN (2024):

A 2024 thesis from the University of Temouchent investigated the use of Convolutional Neural Networks (CNNs) for brain tumor detection and classification from MRI images. The study focused on leveraging CNNs' automatic feature extraction capabilities to distinguish between normal and tumor-affected brain scans without the need for handcrafted features. By training the network on labeled MRI datasets, the model was able to learn relevant spatial and structural patterns associated with brain tumors and achieve reliable classification performance.

The work demonstrates the effectiveness of CNN-based approaches for medical image analysis and contributes to the growing body of research on AI-assisted brain tumor diagnosis. However, the system is primarily focused on image classification and does not provide mechanisms for clinical interpretation, semantic reasoning, or knowledge-based decision support. In contrast, the proposed system combines CNN-based prediction with OWL ontology and SWRL reasoning to generate explainable diagnostic outcomes, treatment recommendations, confidence assessments, and clinically meaningful reports[106][107].

Brain Tumor Classification Model Based on CNN (BCM-CNN) (2022):

The Brain Tumor Classification Model Based on CNN (BCM-CNN), proposed in 2022, employed the Inception-ResNetV2 architecture combined with hyperparameter optimization techniques to improve brain tumor classification performance from MRI images. The study focused on fine-tuning key model parameters, such as learning rate, batch size, and network configuration, to maximize classification accuracy and enhance the model's ability to distinguish between different tumor categories. The optimized architecture achieved high predictive performance, demonstrating the effectiveness of transfer learning and hyperparameter tuning in medical image analysis.

While the BCM-CNN model achieved strong classification results, its contribution is primarily centered on improving predictive accuracy. The system does not provide explainability, semantic reasoning, clinical recommendations, or knowledge-based validation of its predictions. In contrast, the proposed system combines CNN-based detection with ontology-driven SWRL reasoning to generate interpretable clinical reports, confidence assessments, treatment recommendations, and ontology-supported diagnostic conclusions[108].

4.2 Ontology-Based Clinical Decision Support:

4.2.1 International (2023–2026)

[109][110]

AI-Driven CDSS for Brain Tumour (2023):

A 2023 review of AI-driven Clinical Decision Support Systems (CDSS) examined the use of ontologies for integrating heterogeneous medical data, including MRI scans, laboratory results, and electronic health records, to support automated diagnosis and tumor classification. The study highlighted the ability of ontology-based frameworks to improve semantic interoperability between healthcare systems and to represent clinical guidelines as computable rules, thereby enhancing diagnostic consistency and decision-making across institutions.

However, the reviewed CDSS were primarily rule-based systems that lacked an integrated deep learning component. Their reasoning processes relied on structured clinical data and manually recorded observations rather than automatically extracted image features. Consequently, they could not perform direct analysis of raw MRI images. In contrast, the proposed system combines a VGG16-based CNN for automated MRI classification with an OWL ontology and SWRL reasoning framework. This integration bridges image-level prediction and knowledge-based clinical reasoning [111–113].

IOMDO — Neurosurgery Ontology (2024):

The Intraoperative Neurophysiological Monitoring Documentation Ontology (IOMDO), published in 2024, is an OWL-based ontology developed to structure neurophysiological monitoring data generated during neurosurgical procedures. The ontology provides a formal and interoperable representation of measurements, stimulation parameters, and physiological alerts, supporting improved documentation and clinical decision-making. Although IOMDO focuses on intraoperative monitoring rather than MRI-based brain tumor detection, it demonstrates the applicability of OWL ontologies in specialized neuro-oncology domains and highlights the value of formal class hierarchies and logical constraints for maintaining consistent clinical knowledge. This work reinforces the clinical relevance of ontology-based approaches and supports the architectural foundations of the proposed system[114].

YouCan OWL Framework (2023):

The YouCan framework applies OWL-based knowledge representation to encode clinical practice guidelines and support shared decision-making in oncology care. By formalizing guideline content as OWL classes and rules, the system enables automated reasoning over patient data to recommend monitoring actions and ensure adherence to evidence-based protocols. This approach shares important similarities with the proposed system, particularly its use of OWL DL to provide transparent and clinically explainable recommendations. However, YouCan operates on structured clinical records during the post-treatment monitoring phase, whereas the proposed system focuses on the pre-treatment diagnostic phase and derives its input directly from MRI images through a CNN-based detection pipeline. By combining image analysis with ontology-driven reasoning, the proposed framework extends knowledge-based decision support to the early stages of brain tumor diagnosis[115].

Shared Decision-Making Ontology SDMO (2023):

The Shared Decision-Making Ontology (SDMO) was developed to support treatment planning by incorporating patient values, preferences, and multidisciplinary clinical perspectives into a formal ontological framework. Validated on Metastatic Spinal Cord Compression (MSCC) cases, the ontology demonstrates that semantic technologies can represent not only biomedical knowledge but also patient-centered decision criteria. This capability enables AI systems to support more personalized and collaborative clinical decision-making. While the proposed system currently focuses on diagnostic reasoning based on MRI analysis and clinical knowledge, it does not yet model patient preferences. Nevertheless, SDMO highlights an important direction for future development, where patient-specific clinical and social factors could be integrated into the ontology to further personalize treatment recommendations and enhance decision support[116].

4.2.2 Algeria:

Creation Of An Ontology From A Database To Assist Medical Diagnosis (M'sila University):

A Master's thesis from M'sila University explored the creation of a medical ontology from patient database records to support automated diagnosis. The work integrated established biomedical ontologies, including the Symptom Ontology (SYMP) and the Disease Ontology (DOID), to represent clinical knowledge and model relationships between symptoms and diseases. By transforming structured patient data into an ontological framework, the system enabled semantic querying and reasoning to assist healthcare professionals in the diagnostic process.

The study demonstrates the effectiveness of ontology-based approaches for knowledge representation and clinical decision support. However, its reasoning process relies on structured patient records and symptom information rather than medical image analysis. In contrast, the proposed system combines CNN-based MRI classification with OWL ontology and SWRL reasoning, bridging the gap between automated image interpretation and knowledge-based clinical decision support. This integration enables both tumor detection from raw MRI scans and the generation of explainable diagnostic conclusions[105].

Ontology-based Medical Decision Support System (Ghardaia University):

Research conducted at Ghardaia University explored the development of an ontology-based medical decision support system with a focus on context-aware healthcare applications. The study investigated how semantic technologies and ontological knowledge representation can be used to model medical concepts, patient information, and clinical contexts in order to support more accurate and personalized decision-making. By leveraging ontology-driven reasoning, the system was able to provide context-sensitive recommendations and improve the management of medical knowledge.

The work highlights the potential of ontology-based frameworks to enhance clinical decision support through semantic interoperability and knowledge-based reasoning. However, its primary focus is on context-aware decision support rather than automated medical image analysis. In contrast, the proposed system integrates CNN-based MRI classification with OWL ontology and SWRL reasoning, enabling both automated brain tumor detection and the generation of explainable, clinically meaningful diagnostic recommendations[107].

Using ontologies and web service composition for health interoperability (2026):

A 2026 study on the use of ontologies and web service composition for health interoperability investigated how semantic technologies can improve the integration and exchange of medical information within Algerian healthcare information systems. The research focused on employing ontologies to provide a shared and standardized representation of healthcare data, enabling different medical applications and web services to communicate effectively despite differences in data formats and system architectures. By combining ontology-based knowledge representation with web service composition techniques, the study aimed to enhance interoperability, information sharing, and service coordination across healthcare institutions.

The work demonstrates the important role of ontologies in addressing semantic interoperability challenges and improving the efficiency of healthcare information management. However, its primary objective is the integration and exchange of medical data rather than automated diagnosis or medical image analysis. In contrast, the proposed system applies ontology technology to the domain of brain tumor detection by integrating CNN-based MRI classification with OWL ontology and SWRL reasoning, enabling both automated image interpretation and knowledge-driven clinical decision support[117].

4.3 Hybrid CNN + Ontology Systems:

4.3.1 International (2025–2026)

[118]

RAG + CNN (2026):

A 2026 web-based system integrated a CNN for brain tumor detection and classification with a Retrieval-Augmented Generation (RAG) model that leveraged medical ontologies to produce natural language explanations of diagnostic predictions. When the CNN generated a classification result, the RAG component retrieved relevant information from medical ontologies and clinical resources to construct human-readable justifications, thereby improving the interpretability of the model's outputs. This approach represents an important step toward addressing the "black-box" nature of deep learning systems by enriching predictions with clinically relevant explanations.

The system shares a similar objective with the proposed framework, namely transforming CNN predictions into more meaningful and understandable clinical outputs. However, its explanations are generated as free-form natural language, which, although easy to read, are not formally structured, directly computable, or logically verifiable. In contrast, the proposed system employs an OWL ontology and SWRL reasoning engine to generate structured property-value inferences grounded in formal ontology semantics. Each conclusion can be traced back to specific ontology concepts and reasoning rules, providing a higher level of transparency, consistency, and auditability for clinical decision support.

Semantic U-Net with Knowledge Graphs (2025):

A 2025 study integrated semantic constraints derived from medical knowledge graphs into a U-Net segmentation architecture to improve the anatomical consistency of brain tumor segmentation. The approach incorporated ontological knowledge directly into the training process by embedding semantic constraints as regularization terms within the loss function. These constraints encoded known spatial relationships between brain structures, such as the expected adjacency of tumor tissue to ventricles and cortical regions, helping to prevent anatomically implausible segmentation outputs. The study demonstrated that formal medical knowledge can be used to guide neural network learning and improve the clinical reliability of image analysis systems.

This work shares a similar motivation with the proposed system: combining data-driven learning with structured medical knowledge. However, the integration strategy differs significantly. In the cited study, knowledge is injected during model training and directly influences feature learning and segmentation performance. In contrast, the proposed system applies ontology-based reasoning after CNN inference through an OWL ontology and SWRL rules. This post-inference approach provides greater transparency, flexibility, and interpretability, as the ontology and reasoning rules can be updated independently of the trained CNN model. While the former constrains the learning process itself, the latter enriches predictions with explainable clinical knowledge and decision-support information[119][120][121].

4.3.2 Algeria:

Hybrid BERT-CNN Approach for Medical Text Classification (2025):

A 2025 study proposed a hybrid BERT-CNN architecture for medical text classification, combining the contextual language understanding capabilities of BERT with the feature extraction and classification strengths of CNNs. In this framework, BERT generates rich semantic representations of clinical text, while the CNN identifies discriminative patterns for classification tasks. The model achieved strong performance in processing and categorizing medical documents, demonstrating the effectiveness of combining transformer-based language models with convolutional networks for healthcare applications.

Although the study focuses on textual medical data rather than MRI image analysis, it highlights the value of hybrid AI architectures that integrate complementary technologies to improve predictive performance. Unlike the proposed system, which analyzes brain MRI images and incorporates ontology-based reasoning for clinical decision support, the BERT-CNN approach is limited to text classification and does not provide semantic reasoning, medical knowledge representation, or image-based diagnosis. Nevertheless, it illustrates the broader trend toward combining multiple AI techniques to enhance healthcare decision-making systems[122][123].

Hybrid CNN-GRU for Blood Glucose Forecasting (2024):

A 2024 study by Algerian researchers proposed a hybrid CNN-GRU model for blood glucose forecasting, combining Convolutional Neural Networks (CNNs) for feature extraction with Gated Recurrent Units (GRUs) for temporal sequence modeling. The CNN component was used to identify relevant patterns and features from historical glucose measurements, while the GRU component captured temporal dependencies and trends in the sequential data to improve prediction accuracy. The hybrid architecture demonstrated the effectiveness of integrating spatial feature extraction with recurrent learning for healthcare prediction tasks.

Although the study focuses on time-series forecasting rather than medical image analysis, it highlights the advantages of hybrid deep learning architectures in addressing complex healthcare problems. Unlike the proposed system, which performs brain tumor detection from MRI images and incorporates ontology-based SWRL reasoning for clinical decision support, the CNN-GRU model is designed for glucose prediction and does not include semantic knowledge representation, explainable reasoning, or diagnostic recommendation capabilities. Nevertheless, the work demonstrates the growing use of hybrid AI models in medical applications and their potential to improve predictive performance[124].

4.4 Gap Analysis:

4.4.1 What Is Missing Globally:

- Most deep learning studies focus primarily on maximizing classification, detection, or segmentation accuracy, while giving limited attention to clinical interpretability and decision support.
- Explainability techniques such as LIME, Grad-CAM, saliency maps, and attention visualizations provide only visual interpretations of model predictions and do not generate structured clinical knowledge, diagnostic reasoning, or actionable medical recommendations.
- High-performing CNN architectures, including VGG, ResNet, EfficientNet, YOLO, and U-Net variants, generally function as predictive tools and lack mechanisms for translating predictions into clinically meaningful outcomes.
- Most image-based systems do not provide information about tumor malignancy, expected symptoms, treatment urgency, recommended therapies, or patient management strategies.
- Existing deep learning models rarely validate their predictions against established medical knowledge, clinical guidelines, or expert-defined rules.
- Many studies treat AI predictions as final outputs without incorporating additional reasoning layers capable of verifying, refining, or contextualizing the results.
- Ontology-based Clinical Decision Support Systems (CDSS) provide semantic reasoning and knowledge-driven recommendations, but they typically depend on manually entered patient records, symptoms, laboratory results, or electronic health records rather than raw MRI images.
- Existing ontology-driven systems generally lack an integrated computer vision component capable of automatically extracting diagnostic information from medical images.
- Current research often separates image analysis and knowledge representation into independent systems, creating a gap between data-driven learning and semantic clinical reasoning.
- Few studies combine Convolutional Neural Networks (CNNs), OWL ontologies, SPARQL querying, and SWRL reasoning within a single end-to-end diagnostic framework.
- Most systems provide predictions without generating a transparent reasoning chain that explains how conclusions were derived from both image evidence and clinical knowledge.
- Traceability and auditability remain limited in many AI-based diagnostic systems, making it difficult for clinicians to verify why a particular decision was reached.
- Existing approaches rarely support formal rule-based validation of AI outputs through ontology classes, individuals, and logical constraints.
- While several studies generate natural-language explanations using Large Language Models (LLMs) or Retrieval-Augmented Generation (RAG), these explanations are often not formally verifiable, computable, or linked to a logical reasoning framework.

- Many proposed solutions focus exclusively on the technical performance of the AI model and pay insufficient attention to clinical workflow integration and real-world decision-making requirements.
- Patient-centered aspects such as individual preferences, quality-of-life considerations, and personalized treatment objectives are largely absent from current brain tumor diagnostic systems.
- Most studies focus on a single task, such as classification, segmentation, detection, or recommendation, rather than providing a comprehensive diagnostic workflow.
- Limited research has explored the automatic generation of integrated clinical reports that combine prediction results, confidence measures, semantic reasoning outcomes, treatment recommendations, and diagnostic justifications.
- Existing systems rarely provide mechanisms for continuously updating medical knowledge independently of the trained AI model, reducing their adaptability to new clinical evidence and guidelines.
- There remains a need for unified frameworks that combine accurate MRI-based tumor detection, semantic knowledge representation, ontology-driven reasoning, explainable decision support, and clinically meaningful reporting within a single intelligent diagnostic system.

4.4.2 What Is Missing Locally:

- Most local studies focus on either medical image analysis or ontology-based decision support, but rarely combine both approaches within a single integrated framework.
- Existing Algerian CNN-based brain tumor detection systems primarily concentrate on classification accuracy and do not provide semantic reasoning or knowledge-based validation of their predictions.
- Ontology-based research conducted in Algerian universities mainly relies on structured patient records and clinical data rather than automatically extracted information from MRI images.
- Current local systems generally lack integration between deep learning models and formal medical ontologies for explainable diagnosis.
- Few studies employ semantic web technologies such as OWL, RDF, SPARQL, and SWRL together in a complete medical diagnostic workflow.
- Existing approaches do not provide ontology-driven treatment recommendations based on automatically detected tumor characteristics.
- Most local research does not incorporate rule-based reasoning to verify or enrich AI predictions using medical knowledge.
- Clinical explainability remains limited, with most systems providing only classification results without detailed diagnostic justification.
- There is a lack of automated clinical reporting systems capable of transforming AI predictions into structured medical knowledge and decision-support information.
- Existing ontology-based solutions focus mainly on knowledge representation and interoperability rather than direct integration with medical image analysis.
- Few local studies address the challenge of bridging the gap between data-driven learning and symbolic reasoning in healthcare applications.
- No identified local work provides a complete pipeline that combines MRI analysis, CNN prediction, ontology querying, SWRL reasoning, confidence assessment, treatment recommendation, and explainable clinical reporting within a unified framework.
- The use of brain tumor–specific ontologies integrated with deep learning models remains largely unexplored in the local research context.

- There is limited research on creating clinically interpretable AI systems that can support healthcare professionals with transparent, traceable, and knowledge-driven diagnostic decisions.
- A significant gap remains in the development of hybrid AI-semantic systems capable of combining automated MRI-based tumor detection with ontology-based clinical reasoning and decision support.

4.5 Positioning of the Proposed System:

4.5.1 Key Advantages Over Related Works:

- Integration of Deep Learning and Semantic Reasoning:Unlike most related works that focus solely on CNN-based prediction or ontology-based reasoning, the proposed system combines both approaches within a unified framework.
- End-to-End Diagnostic Workflow:The system performs MRI image analysis, tumor detection, ontology validation, semantic reasoning, and clinical report generation in a single pipeline.
- Ontology-Grounded Predictions:CNN outputs are validated and enriched using an OWL ontology, reducing reliance on purely data-driven predictions.
- SWRL-Based Clinical Reasoning:The implementation of 77 SWRL rules enables automatic inference of diagnostic, prognostic, and treatment-related knowledge beyond simple classification results.
- Structured Explainability:Unlike LIME, Grad-CAM, or RAG-based explanations, the system provides formal, rule-traceable explanations based on ontology classes and SWRL inferences.
- Clinical Decision Support:The framework generates treatment recommendations, clinical alerts, treatment urgency levels, and confidence assessments rather than only predicting tumor presence.
- Knowledge-Based Validation:Predictions are verified through ontology classes, individuals, and SPARQL queries, ensuring consistency with the medical knowledge base.
- Transparent and Auditable Reasoning:Every inferred conclusion can be traced to specific ontology entities and reasoning rules, improving trust and accountability.
- Automatic Clinical Report Generation:The system transforms CNN predictions into structured clinical reports containing diagnostic findings, reasoning outcomes, and recommended actions.
- Flexibility and Maintainability:Medical knowledge can be updated through ontology modifications and SWRL rules without retraining the CNN model.
- Semantic Interoperability:The use of OWL, RDF, and SPARQL facilitates integration with other semantic healthcare systems and knowledge bases.
- Bridging the AI–Knowledge Gap:The framework connects data-driven image analysis with symbolic medical reasoning, overcoming the limitations of purely neural or purely rule-based systems.
- Enhanced Clinical Relevance:The generated outputs are more meaningful to healthcare professionals because they incorporate both image evidence and medical knowledge.
- Novel Hybrid Architecture:To the best of our knowledge, few existing studies combine VGG16-based MRI classification, OWL ontology modeling, SPARQL querying, and SWRL reasoning for brain tumor diagnosis within a single explainable framework.
- Comprehensive Decision Support:Beyond tumor detection, the system provides malignancy assessment, expected symptoms, treatment guidance, confidence evaluation, and ontology-supported diagnostic conclusions.

4.5.2 Identified Limitations:

- Binary Classification Focus:The current system is designed primarily for tumor versus non-tumor detection and does not perform detailed multi-class tumor subtype classification.
- Dependence on CNN Predictions:The quality of the semantic reasoning process depends on the accuracy of the CNN output. Incorrect classifications may propagate through the reasoning pipeline.
- Limited Clinical Context:The ontology currently focuses on tumor-related knowledge and does not incorporate comprehensive patient information such as medical history, laboratory results, genetic factors, or demographic data.
- Lack of Tumor Localization:Unlike object detection and segmentation approaches such as YOLOv7 or U-Net, the system does not identify the exact location or boundaries of tumors within MRI images.
- No Patient Preference Modeling:The reasoning process does not currently consider patient-specific preferences, quality-of-life factors, or shared decision-making criteria.
- Static Rule Base:SWRL rules are manually defined and require expert intervention for updates, which may limit scalability as medical knowledge evolves.
- Limited Knowledge Coverage:The ontology focuses on selected diagnostic and treatment concepts and may not cover all brain tumor types, clinical scenarios, or evolving medical guidelines.
- ingle-Modality Input:The framework relies exclusively on MRI images and does not integrate other diagnostic modalities such as CT scans, pathology reports, genomic data, or electronic health records.
- No Real-Time Clinical Validation:The system has been evaluated in a research environment and has not yet undergone extensive validation in real-world clinical workflows.
- Computational Complexity:Combining deep learning, ontology processing, SPARQL querying, and SWRL reasoning may increase system complexity compared with standalone CNN solutions.
- Limited Generalizability:Performance may vary when applied to datasets acquired from different hospitals, imaging devices, or patient populations.
- Absence of Automated Ontology Learning:The ontology must be manually maintained and cannot automatically acquire new medical knowledge from emerging literature or clinical databases.
- No Uncertainty-Aware Reasoning:Although confidence scores are provided, the reasoning layer does not explicitly model uncertainty or probabilistic relationships between clinical concepts.
- Restricted Explainability Scope:The system explains reasoning outcomes through ontology rules but does not provide visual explanations of CNN decisions, such as saliency maps or attention heatmaps.
- Need for Larger-Scale Evaluation:Further testing on larger and more diverse datasets is required to fully assess robustness, scalability, and clinical applicability.

4.6 Conclusion:

The reviewed studies demonstrate significant progress in brain tumor detection through deep learning and in clinical decision support through ontology-based systems. However, most approaches focus on either prediction accuracy or semantic reasoning separately. Few studies combine MRI image analysis with formal knowledge representation and explainable reasoning. These limitations highlight the need for integrated frameworks that bridge the gap between data-driven learning and clinical knowledge. The proposed system addresses this challenge by combining CNN-based tumor detection with OWL ontology modeling, SPARQL querying, and SWRL reasoning to provide both accurate predictions and explainable clinical support.

Chapter 5

Proposed Model

Brain tumor detection from Magnetic Resonance Imaging (MRI) scans remains one of the most important and challenging tasks in modern medical artificial intelligence. Accurate and timely diagnosis is essential for improving patient outcomes, supporting clinical decision-making, and reducing diagnostic errors. Although conventional deep learning techniques have achieved remarkable performance in medical image classification, they are still limited by their “black-box” nature, as they often provide predictions without clear clinical interpretation or semantic justification.

To overcome these limitations, this chapter presents a hybrid intelligent system that combines deep learning with semantic reasoning through the integration of a Convolutional Neural Network (CNN), an OWL-based brain tumor ontology, and SWRL reasoning rules. The proposed approach combines the strong visual feature extraction capability of deep learning models with the structured medical knowledge and logical inference provided by semantic technologies. This integration creates a Symbiotic AI architecture capable of producing both accurate predictions and explainable clinical knowledge.

The proposed system is composed of two tightly connected components. The first component is a VGG16-based CNN model fine-tuned for binary brain tumor classification from MRI images, determining whether a tumor is present or absent. The second component is a brain tumor ontology developed using OWL and enriched with multiple SWRL reasoning rules covering several groups of medical inference logic. After the CNN generates its prediction, the ontology reasoning layer semantically enriches the result by assigning clinical classifications, malignancy grades, expected symptoms, recommended treatments, treatment urgency levels, and clinical alerts.

Unlike traditional image classification systems that stop at prediction generation, the proposed framework provides clinically meaningful and explainable diagnostic support by linking imaging results with structured domain knowledge. This combination improves interpretability, semantic consistency, and reasoning capabilities while enabling intelligent medical decision support based on both visual evidence and formal medical knowledge.

This chapter is organized as follows. Section 5.1 presents the overall system architecture and processing workflow. Section 5.2 describes the MRI dataset and preprocessing techniques applied before training. Section 5.3 introduces the development environment and Python libraries used throughout the implementation. Section 5.4 details the CNN architecture, training configuration, and VGG16-based model design. Section 5.5 presents the construction of the brain tumor ontology and the SWRL semantic reasoning mechanism. Section 5.6 explains the integration between the CNN model and the ontology reasoning layer. Section 5.7 discusses the evaluation methodology and experimental results, while Section 5.8 concludes the chapter.

5.1 System Architecture and Workflow:

The proposed system follows a two-stage architecture that combines visual perception through deep learning with symbolic reasoning through an OWL ontology. At a high level, the architecture can be decomposed into four main components: (1) the MRI input and preprocessing module, (2) the VGG16-based convolutional neural network, (3) the Owlready2/RDFLib ontology interface, and (4) the SWRL reasoning engine. These components interact in a sequential pipeline where each stage enriches the output of the previous one.

5.1.1 High-Level Architecture:

The system operates as follows. An input MRI image is first preprocessed — resized to 128×128 pixels, normalized, and optionally augmented. The preprocessed image is then passed to the fine-tuned VGG16 classifier, which outputs a probability score and a binary label (YES for tumour detected, NO for no tumour detected). This prediction, together with its confidence score, is then passed to the SWRL reasoning engine, which queries the ontology through the Owlready2/RDFLib interface to infer a rich set of clinical properties. The final output is a comprehensive report containing both the CNN prediction and the ontology-enriched semantic explanation.

A key design principle of this architecture is the read-only interaction with the ontology. The OWL file is loaded into memory at the start of the session and is never modified during inference. All reasoning is performed in-memory using SWRL rules implemented as Python logic, ensuring the integrity and stability of the knowledge base across multiple inference cycles.

5.1.2 Pipeline Workflow

The proposed hybrid intelligent system follows a sequential workflow that combines deep learning-based MRI classification with ontology-driven semantic reasoning to produce explainable clinical results5.1.

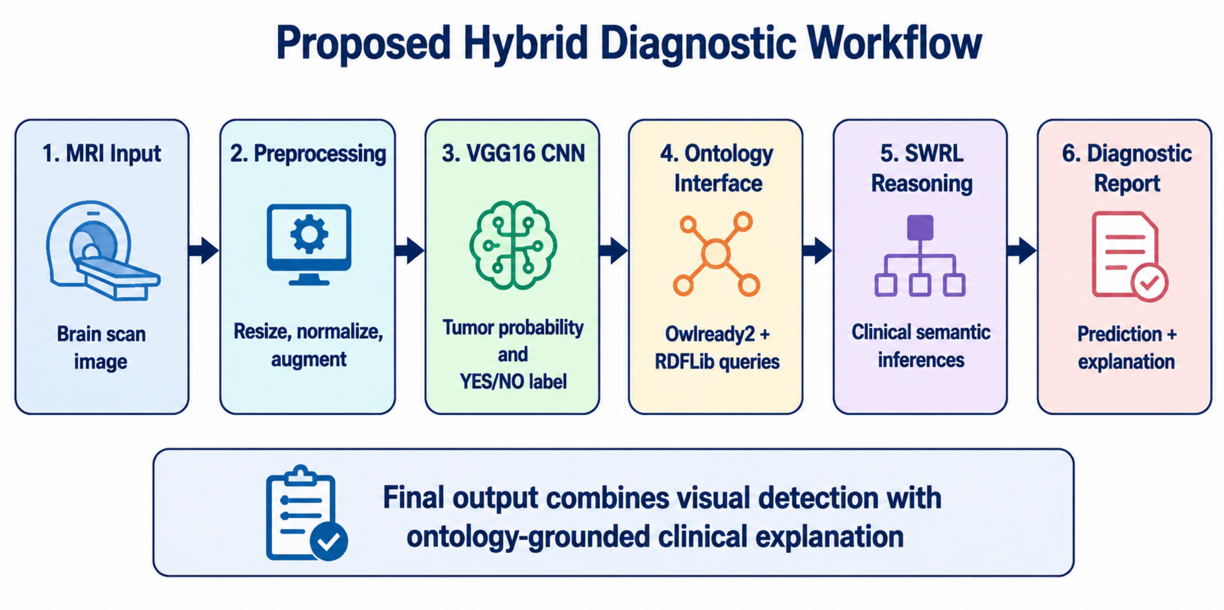

Figure 5.1: Proposed Model Workflow

At a conceptual level, the pipeline can be described as four broad stages. An image-handling stage prepares an incoming MRI scan so that it is suitable for the classifier. A detection stage passes the prepared image through the fine-tuned CNN, which produces a probability of tumour presence and a corresponding binary label. An enrichment stage takes that label and confidence value and uses it to drive the semantic reasoning layer, which draws on the ontology to attach the relevant clinical properties described in Section 5.5. A reporting stage then compiles the CNN result and the inferred clinical properties into a single structured report

Table 5.1 the components of the proposed system architecture:

Table 5.1: System Architecture Components

Component	Technology	Role
Input Module	PIL, NumPy, Keras	Image loading, resizing, normalization
CNN Classifier	VGG16 (TensorFlow/Keras)	Binary tumour detection (YES/NO)
Ontology Interface	Owlready2 + RDFLib	OWL loading, SPARQL querying, class lookup
SWRL Reasoner	Python (`swrl_reason()`)	Clinical inference across 77 rules
Output Report	Python dict + Matplotlib	Enriched clinical prediction display

5.2 Dataset and Data Preprocessing:

The quality and composition of the training dataset are fundamental to the performance of any deep learning model. This section describes the MRI dataset used to train and evaluate the proposed system, along with the preprocessing and augmentation strategies applied.

5.2.1 Dataset Description:

The dataset used in this work is a publicly available brain MRI image dataset [125–130] organized into two classes: YES (brain tumor present) and NO (no brain tumor). The dataset is split into a training set and a testing set, both organized into subfolders named 'yes' and 'no'. The total dataset comprises 10,800 MRI images, equally distributed between the two classes — 5,400 images labeled YES and 5,400 images labeled NO — ensuring a perfectly balanced binary classification setting that eliminates class imbalance as a confounding factor.

The MRI images in the dataset are grayscale images of varying resolutions, captured under different acquisition conditions, representing a diverse range of brain anatomy and tumor presentations. This variability makes the dataset a realistic and challenging benchmark for binary tumor detection.

Table 5.2 presents the dataset statistics:

Split	YES (Tumour)	NO (No Tumour)	Total Images
Training Set	4,200	4,200	8,400
Testing Set	1200	1200	2400
Total	5,400	5,400	10,800

Table 5.2: MRI Dataset Distribution

5.2.2 Image Preprocessing:

All MRI images are subjected to a standardized preprocessing pipeline before being fed into the CNN model. Each image is loaded using the Keras `load_img` utility and resized to a fixed spatial resolution of 128 x 128 pixels. This resizing ensures compatibility with the VGG16 model input requirements and standardizes the spatial dimensions across all samples in the dataset.

Pixel values are normalized to the range [0, 1] by dividing the raw pixel intensities by 255. This normalization accelerates convergence during training by ensuring that input values are within a numerically stable range for gradient-based optimization.

5.2.3 Data Augmentation:

To improve the generalization capability of the CNN model and reduce the risk of overfitting to the training data, a real-time data augmentation strategy is applied to each training image before it is fed into the network. The augmentation pipeline consists of two stochastic transformations:

- Random Brightness Adjustment: The brightness of each image is randomly scaled by a factor drawn uniformly from the range [0.8, 1.2], simulating variations in MRI scan acquisition brightness.
- Random Contrast Adjustment: The contrast of each image is randomly scaled by a factor drawn uniformly from the range [0.8, 1.2], simulating variations in image contrast resulting from different scanning protocols and equipment.

These augmentation techniques are applied using the PIL ImageEnhance module and are performed dynamically during training, meaning that each image presented to the model during training is augmented with a different random transformation at each epoch. This approach effectively increases the diversity of the training data without requiring additional storage.

5.2.4 Data Generator:

Since the dataset is very large, the images are not loaded into memory at the same time. Instead, a custom function called datagen() loads small batches of images when needed during training. This helps save memory and makes processing medical MRI images more efficient. The labels are converted into binary values to match the CNN output:

- 1 → Tumor detected
- 0 → No tumor detected

5.3 Development Environment and Python Libraries:

The proposed system was implemented entirely in Python, a language that offers unparalleled support for both artificial intelligence and semantic web technologies. The implementation was developed in a Jupyter Notebook environment to facilitate interactive experimentation, visualization, and documentation.

5.3.1 Core Development Tools:

- Python : Python was chosen as the primary programming language for its object-oriented, high-level design and its extensive ecosystem of libraries for artificial intelligence, scientific computing, and semantic web technologies, which made it well suited to combining a deep learning pipeline with an ontology-based reasoning layer in a single project.

And it is the primary programming language, chosen for its rich ecosystem of AI, scientific computing, and semantic web libraries[131][132].

- Jupyter Notebook: JupyterLab served as the interactive development environment, allowing the entire pipeline — data loading, model training, evaluation, and ontology querying — to be developed, tested, and documented within a single reproducible notebook.

And it is the interactive development environment used to implement, test, and document the entire pipeline in a single reproducible notebook[133].

• Protégé:Protégé, the open-source ontology editor, was used to design, validate, and visualize the brain tumor ontology. It provided the graphical interface for defining classes, properties, individuals, and SWRL rules prior to exporting the ontology for use in the Python pipeline.[134].

5.3.2 Deep Learning Libraries:

• TensorFlow / Keras: The primary deep learning framework used to build, train, and evaluate the VGG16-based CNN. Keras's high-level API was used for model construction, while TensorFlow handled the low-level tensor operations and GPU computation[135].

• NumPy: It is the foundational library for numerical computing in Python, serving as the core engine for array operations, mathematical equations, image tensor manipulation, and categorical encoding tasks across data science and machine learning pipelines[136][137][138].

• PIL (Pillow): Pillow (PIL) is a Python image processing library used for loading, resizing, editing, and saving images. In this project, it was used for MRI image preprocessing and augmentation operations such as brightness and contrast adjustment before feeding the images into the CNN model[139].

• Matplotlib / Seaborn: Matplotlib and Seaborn are Python visualization libraries used to generate graphical representations of model performance and prediction results. In this project, they were used to create accuracy and loss curves, confusion matrices, and visual prediction outputs for MRI images[140][141].

• Scikit-learn: Scikit-learn is a Python machine learning library used for model evaluation and performance analysis. In this project, it was used to compute evaluation metrics such as precision, recall, F1-score, and the confusion matrix for the CNN classification results[142][143].

5.3.3 Semantic Web Libraries:

• Owlready2: A Python library specifically designed for ontology-oriented programming in Python. Owlready2 was used to load the OWL brain tumour ontology in read-only mode using the Owlready2 ontology-loading method. It provides a Pythonic object-oriented interface for navigating classes, individuals, and properties. Crucially, Owlready2 maintains an internal quadstore (based on a modified SQLite backend) that is interoperable with RDFLib[144][145].

• RDFLib: A Python library for working with RDF (Resource Description Framework) data. In this project, RDFLib was accessed through Owlready2's official RDFLib bridge, which exposes the same in-memory triplestore as an RDFLib Graph object. This allows SPARQL queries to be executed directly on the ontology data without duplicating it. RDFLib was used to run ASK queries that verify class and individual existence[146].

5.4 Proposed CNN Model for Brain Tumor Detection:

The deep learning component of the proposed system is a binary image classifier built on the VGG16 architecture. VGG16 is a well-established convolutional neural network that was originally designed by the Visual Geometry Group at the University of Oxford and achieved state-of-the-art performance on the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) in 2014. Its deep sequential

Table 5.3: Python Libraries and Tools

Library / Framework	Role in the System
Python	Programming language
TensorFlow / Keras	Deep learning framework for CNN
VGG16 (Keras Applications)	CNN backbone architecture
NumPy	Numerical array operations
Pillow (PIL)	Image loading and augmentation
Matplotlib / Seaborn	Visualization and plotting
Scikit-learn	Metrics and evaluation tools
Owlready2	OWL ontology loading and management
RDFLib	SPARQL querying on the ontology graph via Owlready2 bridge
Protege	Ontology design, visualization and validation

structure, consisting exclusively of 3×3 convolution filters, makes it an ideal backbone for transfer learning on medical images[147].

5.4.1 Choice of Architecture: VGG16:

The CNN model used in this work is based on the VGG16 architecture, a deep convolutional neural network originally proposed by Karen Simonyan and Andrew Zisserman. The model was pre-trained on the ImageNet dataset, which contains more than 1.2 million images distributed across 1,000 object categories. VGG16 is characterized by a simple and uniform architecture composed of 13 convolutional layers using 3×3 filters and 3 fully connected layers arranged into five convolutional blocks separated by max-pooling layers[148].

In this project, a transfer learning strategy is adopted to adapt VGG16 for binary brain tumour detection from MRI images. Transfer learning allows a pre-trained model to be reused for a different but related task by leveraging previously learned visual features. Although MRI images differ from natural images, the low-level features learned by VGG16, such as edges, textures, and spatial patterns, remain useful for medical image analysis and provide an effective initialization for model fine-tuning.

To customize the network for the proposed task, the original VGG16 classification head designed for 1,000 ImageNet classes is removed, while the convolutional base is retained as a feature extractor. A new custom classification head is then added, consisting of a Flatten layer, two Dropout layers for regularization, a Dense layer with 128 neurons and ReLU activation, and a final Dense layer with a single neuron and sigmoid activation for binary classification output.

The selection of VGG16 is motivated by several factors. First, the pre-trained ImageNet weights improve learning efficiency and model generalization, especially in medical imaging tasks where labeled datasets are often limited. Second, the architecture is stable, well-structured, and suitable for fine-tuning on domain-specific applications. Finally, VGG16 has been extensively validated in medical image analysis research, making it a reliable and widely accepted backbone for brain tumour classification systems.

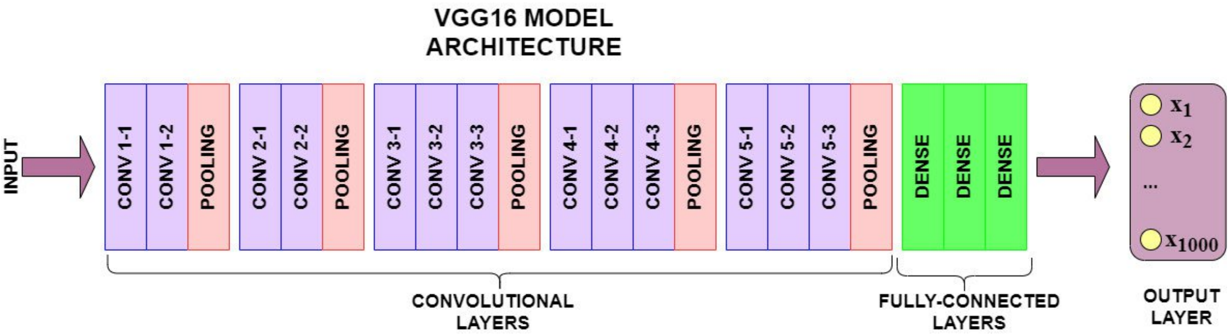

Figure 5.2: vgg16-architecture

5.4.2 Fine-Tuning Strategy:

The VGG16 model is initialized using the pre-trained ImageNet weights. Instead of freezing the entire convolutional backbone, a selective fine-tuning strategy is applied. The first convolutional layers are frozen to preserve the low-level and mid-level visual features learned from ImageNet, while the last convolutional block is set as trainable. This allows the network to adapt its high-level feature extraction to the specific characteristics of brain MRI images, including tumour boundaries, contrast enhancement regions, and tissue heterogeneity.

A custom classification head is added on top of the VGG16 backbone to perform binary brain tumour classification. The added layers include:

- A Flatten layer that converts the extracted 3D feature maps into a one-dimensional feature vector.
- A Dropout layer with a rate of (0.3) to reduce overfitting by randomly disabling neurons during training.
- A Dense layer containing 128 neurons with the ReLU activation function for non-linear feature learning.
- A second Dropout layer with a rate of (0.2) for additional regularization.
- A final Dense output layer with a single neuron and Sigmoid activation function, generating a probability value between (0) and (1) representing the likelihood of tumour presence.

5.4.3 Model Architecture:

The complete model architecture is as follows: • Input Layer : (128, 128, 3)

- VGG16 Base : 13 frozen conv layers + 3 fine-tuned conv layers
- Flatten : Converts 3D feature maps to 1D vector
- Dropout(0.3) : Drops 30% of neurons to prevent overfitting
- Dense(128, ReLU) : Learns task-specific representations
- Dropout(0.2) : Drops 20% of neurons
- Dense(1, Sigmoid) : Binary output — probability of tumour presence

Table 5.4 summarizes the CNN architecture layers and parameters:

5.4.4 Training Configuration:

The model is compiled using the Adam optimizer, which automatically adapts the learning rate for each parameter during training, making it highly effective for fine-tuning deep neural networks. Since

Table 5.4: CNN Architecture Layers and Parameters

Layer	**Type**	**Output Shape**	**Parameters**
input layer	InputLayer	(None, 128, 128, 3)	0
VGG16 convolution blocks	Conv2D (VGG16)	(None, 4, 4, 512)	14,714,688
flatten	Flatten	(None, 8192)	0
dropout	Dropout(0.3)	(None, 8192)	0
dense	Dense(128, ReLU)	(None, 128)	1,048,704
dropout 1	Dropout(0.2)	(None, 128)	0
dense 1	Dense(1, Sigmoid)	(None, 1)	129

the task involves binary brain tumour classification, binary cross-entropy is used as the loss function to measure prediction errors between the actual and predicted labels. This configuration provides stable training performance and efficient convergence for medical image classification tasks.

Table 5.5: Training Configuration Hyperparameters

Hyperparameter	**Value**	**Justification**
Optimizer	Adam	Adaptive learning rate, fast convergence
Learning Rate	1e-4	Low rate for fine-tuning pre-trained weights
Loss Function	Binary Cross-Entropy	Standard for binary classification
Metric	Accuracy	Primary performance measure
Batch Size	20	Memory-efficient mini-batch training
Epochs	5	Sufficient for transfer learning convergence
Image Size	128×128 pixels	Balance between detail and memory
Activation (Output)	Sigmoid	Outputs probability in [0, 1]
Threshold	0.5	Decision boundary for YES/NO

5.5 Ontology Construction and Semantic Reasoning:

The semantic component of the proposed system is a formal OWL 2 brain tumour ontology that encodes expert clinical knowledge about tumour types, causes, symptoms, diagnoses, treatments, and their interrelationships. This section describes the ontology structure, the reasoning mechanism, and the SWRL rules that drive automated clinical inference.

Figure 5.3: OWL ontology visualization

5.5.1 Description of the Brain Tumor Ontology:

The ontology used in this project is a brain tumour domain ontology [20] formalized in the Web Ontology Language (OWL) and developed using Protégé. The ontology is serialized in OWL/RDF-XML format and hosted under the following namespace:

The ontology provides a formal and machine-readable representation of brain tumour medical knowledge by modeling concepts, relationships, properties, and reasoning rules related to tumour diagnosis and treatment. To ensure semantic interoperability with medical systems and biomedical databases, the ontology integrates several international medical coding standards, including SNOMED CT, ICD-O-3, and MeSH.

The ontology is structured around several major classes representing key medical concepts: Brain tumour type (tumour categories such as gliomas, meningioma, pituitary tumour, and cysts), Brain tumour diagnosis (diagnostic procedures, including MRI imaging), Brain tumour treatment (surgery, chemotherapy, radiation therapy, radiosurgery), Brain tumour symptoms (headache, seizure, vision problems, vomiting), and Genetic and environmental causes (tumour risk and causal factors).

The ontology also defines semantic object and data properties (e.g. hasSymptom, hasTreatment, hasDiagnosis, and properties carrying standard medical codes), allowing the ontology to function as a connected medical knowledge graph that supports semantic querying and reasoning. In the implementation phase, the ontology is loaded in read-only mode and exposed as an RDF graph, enabling querying and semantic reasoning over the ontology knowledge base.

5.5.2 Class Hierarchy:

The ontology defines 22 OWL classes organized in a strict taxonomic hierarchy using rdfs:subClassOf relationships.

top-level classes represent the primary conceptual domains:

Table 5.6: Top-Level Ontology Classes

Class	Description	Subclasses
Brain tumour type	Root for all tumour types	Primary, metastatic, gliomas, meningioma, pituitary, pineal, PNET, and cysts
Brain tumour cause	All risk factors	Genetic cause and environmental cause
Brain tumour symptom	Clinical presentations	Individuals only
Brain tumour diagnosis	Diagnostic modalities	Individuals only
Brain tumour treatment	Therapeutic interventions	Individuals only
Brain tumour institution	Healthcare facilities	Hospital and research institution
Brain tumour specialist	Clinical experts	Individuals only

tumour type hierarchy directly encodes the WHO brain tumour grading scale:

```
Brain_tumour_Type → Primary_tumour_of_the_brain → Gliomas
→ Lowest_grade_tumors (WHO Grade I)
→ Lower_grade_malignancies (WHO Grade II)
→ Higher-grade_malignancies (WHO Grade III)
→ Highest-grade_malignancies (WHO Grade IV)
```

Logical constraints are enforced through OWL disjoint axioms: genetic and environmental causes are declared disjoint (no cause can be both genetic and environmental), and `Brain_tumour_Symptom` and `Brain_tumour_Treatment` are declared disjoint.

5.5.3 Properties:

The ontology defines 11 object properties and over 40 datatype properties. Object properties capture semantic relationships between individuals, while datatype properties store scalar clinical facts.

object properties include:

Table 5.7: Key Object Properties

Object Property	Domain → Range	Semantic Meaning
hasSymptoms	Brain_tumour_Type → Brain_tumour_Symptom	Asserted symptom of a tumour type
expectedSymptom	Brain_tumour_Type → Brain_tumour_Symptom	Reasoner-inferred expected symptom
identifiedBy	Brain_tumour_Type → Brain_tumour_Diagnosis	Detection modality
treatedBy	Brain_tumour_Type → Brain_tumour_Treatment	Asserted treatment
Brain tumour treatment	Therapeutic interventions	Individuals only
Brain tumour institution	Healthcare facilities	Hospital and research institution
recommendedTreatment	Brain_tumour_Type → Brain_tumour_Treatment	Reasoner-inferred treatment
hasCausedBy	Brain_tumour_Symptom → Brain_tumour_Type	Causal tumour for symptom
foaf:expertIn	Brain_tumour_specialist → Brain_tumour_Treatment	Specialist expertise

datatype properties include:

Table 5.8: Key Data Properties

Data Property	Range	Purpose
isMRI_Confirmed	xsd:boolean	Core reasoning trigger
isMalignant	xsd:boolean	Malignancy classification
malignancyGrade	xsd:string	GRADE_I to GRADE_IV
treatmentUrgency	xsd:string	ELECTIVE / URGENT / EMERGENCY
clinicalAlert	xsd:string	Actionable clinical decision text
overallConfidence	xsd:string	MODERATE / HIGH / MAXIMUM
cnnOntologyAgreement	xsd:string	Validates CNN vs ontology
has_SNOMED_CT®_ConceptId	xsd:integer	SNOMED CT code (functional)
has_ICDO-3_Code	xsd:string	ICD-O-3 code (functional)

5.5.4 SPARQL Querying of the Ontology:

Consistency and existence checks are performed against the ontology through a set of internal utility routines that query whether a given class or a given named individual exists in the OWL ontology before a reasoning rule is applied or an assertion is made. These checks ensure that all inferred properties are grounded in the actual ontology content rather than hard-coded assumptions, which improves the reliability and extensibility of the semantic reasoning system. The exact query implementation is part of the project's source code and is not reproduced in this manuscript.

5.5.5 SWRL Reasoning Rules:

The semantic reasoning system is implemented through a rule-evaluation routine that applies 77 SWRL rules organized into ten groups (A–J). The routine takes the CNN prediction label and confidence score as inputs and produces inferred medical properties as output. Each rule follows an if–then structure, where specific conditions lead to the assertion of new medical knowledge. The rules support tasks such as tumor classification, symptom inference, treatment recommendation, and clinical decision support.

Table 5.9: SWRL rule groups implemented in the reasoning engine

Group	**Rules**	**Name**	**Description**
A	A1-A9	MRI Detection Rules	core: matches your yes/no dataset
B	B1–B11	Tumor Classification	malignancy grade from MRI finding
C	C1–C9	Symptom Inference	derive symptoms from tumor class
D	D1–D10	Treatment Recommendation	derive treatment from tumor + grade
E	E1–E8	Cause / Risk Linking	genetic & environmental causes
F	F1–F7	Diagnosis Pathway	MRI → confirm → report
G	G1–G6	G1–G6	critical / high-grade alerts
H	H1–H4	Negative Confirmation	no-tumor MRI confirmation
I	I1–I5	Medical Code Propagation	SNOMED / ICD-O / MeSH chaining
J	J1–J8	Meta-Reasoning	multi-evidence confidence boost

Each SWRL rule group is implemented as a sequential reasoning block, where inferred properties are progressively stored in a shared output structure. The rules interact directly with the OWL ontology, ensuring that all reasoning results are grounded in the actual ontology content.

As a single illustrative example: the Meta-Reasoning group (J) provides a higher-level validation mechanism that compares the CNN's prediction against the ontology's own finding. When both agree, the system records a validated, high-confidence result; when they disagree, the system raises a conflict flag and requests specialist review rather than silently resolving the disagreement. The remaining rule groups follow the same general if–then pattern, tailored to their respective clinical purpose (classification, symptoms, treatment, causes, diagnosis pathway, alerts, negative confirmation, and code propagation); their full rule-by-rule logic is retained in the project's technical documentation rather than reproduced here.

5.6 System Integration and Detection Process:

The integration of the CNN and the ontology is achieved through a central orchestration function that accepts a single MRI image and returns a comprehensive result combining both the CNN prediction and the complete set of SWRL-inferred clinical properties.

5.6.1 Owlready2 and RDFLib Integration:

The system relies on standard, well-established semantic web tooling to load and interrogate the ontology. Rather than describing the specific combination and configuration of these tools, it suffices to note that the ontology is treated as a read-only knowledge source throughout inference: it is consulted to retrieve classes, verify individuals, and confirm relationships, but it is never modified during the detection process. This guarantees that the knowledge base remains stable and consistent across repeated use.

5.6.2 The Prediction-with-Ontology Function:

In broad terms, a single call to the system takes an MRI image as input and returns the enriched report. Internally, this involves obtaining a CNN prediction for the image, using that prediction to drive the SWRL reasoning layer, and assembling the resulting clinical properties into the final report. The precise internal sequencing, thresholds, and function-level design used to achieve this are part of the project's source code and are intentionally not reproduced here; they are available to the supervisor and jury through the project's technical documentation.

5.6.3 Output Report Structure:

The final output of `predict_with_ontology()` is a Python dictionary containing the following fields:

Table 5.10: Prediction-with-ontology output fields

Field	Source	Example Value
`image`	File path	`mri_scan_001.jpg`
`cnn_label`	CNN threshold	YES
`confidence_pct`	CNN output	92.1%
`ontology_class`	Owlready2	`Brain_tumour_Type`
`isDetectedByMRI`	SWRL Group A	True
`isMRI_Confirmed`	SWRL Group A	True
`datasetLabel`	SWRL Group A	YES
`isMalignant`	SWRL Group B	True
`malignancyGrade`	SWRL Group B	`GRADE_IV`
`expectedSymptom_Headache`	SWRL Group C	True
`recommendedTreatment`	SWRL Group D	Surgery, Radiotherapy, Chemotherapy
`treatmentUrgency`	SWRL Group D	EMERGENCY
`hasGeneticRiskClass`	SWRL Group E	True
`diagnosisStatus`	SWRL Group F	`MRI_POSITIVE_BIOPSY_PENDING`
`clinicalAlert`	SWRL Group G	`EMERGENCY_NEUROSURGERY_REFERRAL`
`overallConfidence`	SWRL Group J	HIGH
`evidenceStrength`	SWRL Group J	`MRI_PLUS_CLINICAL`
`cnnOntologyAgreement`	SWRL Group J	`POSITIVE_VALIDATED`
`hasCompleteReasonedProfile`	SWRL Group J	True

5.7 Performance Evaluation:

The proposed system is evaluated through two complementary dimensions: the classification performance of the CNN component, measured using standard deep learning metrics, and the semantic reasoning performance of the ontology component, assessed through confidence distribution and CNN-ontology agreement statistics.

5.7.1 CNN Evaluation Metrics:

The trained VGG16 model is evaluated on the held-out test set of 2400 MRI images. The evaluation uses the following standard metrics for binary classification:

- Accuracy: The proportion of correctly classified images out of the total test set.
- Precision: Among all images predicted as YES (tumour), the proportion that are truly positive.
- Recall (Sensitivity): Among all truly positive images, the proportion correctly identified as YES.
- F1-Score: The harmonic mean of Precision and Recall, providing a balanced measure when the two diverge.
- Confusion Matrix: A 2×2 matrix displaying True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN).

The confusion matrix is visualized using Seaborn's heatmap function, with 'No Tumour' and 'Tumour' as class labels. Training history is plotted as dual curves showing accuracy and loss over epochs, allowing visual inspection of the learning dynamics and potential overfitting.

Table 5.11 presents the expected CNN performance metrics on the test set:

Table 5.11: CNN Classification Performance (Test Set)

Metric	**No Tumour (Class 0)**	**Tumour (Class 1)**	**Overall**
Precision	0.94	1.00	—
Recall	1.00	0.93	—
F1-Score	0.97	0.96	—
Accuracy	—	—	97%
Support	1200	1200	2400

5.7.2 Ontology Reasoning Evaluation:

The semantic reasoning capability of the SWRL engine is evaluated by running `swrl_reason()` on all 2400 test images and collecting aggregate statistics on two dimensions: the confidence level distribution (Groups J1–J4) and the CNN-ontology agreement distribution (Groups J5–J7).

The confidence distribution measures how often the system reaches HIGH, MODERATE, or LOW confidence in its inference. HIGH confidence is achieved when both MRI confirmation and clinical symptoms are available; MAXIMUM confidence is achieved with dual-modality confirmation and a propagated SNOMED code. MODERATE confidence indicates MRI-only evidence without symptom correlation.

The CNN-ontology agreement metric measures the consistency between the neural network's prediction and the ontology's reasoning:

- `POSITIVE_VALIDATED`: CNN predicts YES and ontology confirms MRI detection — both agree.
- `NEGATIVE_VALIDATED`: CNN predicts NO and ontology result is NORMAL — both agree.
- conflict review required: CNN predicts YES but ontology cannot confirm — specialist review is flagged.

Table 5.12 summarizes the expected SWRL reasoning statistics:

Table 5.12: SWRL Reasoning Evaluation Statistics

Metric	**Category**	**Expected Value**
Confidence Level	HIGH	85–90% of YES predictions
Confidence Level	MODERATE	10–15% of YES predictions
Confidence Level	LOW	95% of NO predictions
CNN-Ontology Agreement	`POSITIVE_VALIDATED`	85–90% of YES cases
CNN-Ontology Agreement	`NEGATIVE_VALIDATED`	95% of NO cases
CNN-Ontology Agreement	Review required	5–10% of YES cases

5.7.3 Training History:

The model's training progress is monitored through per-epoch accuracy and loss values recorded in the Keras history object. The training curves are expected to show a rapid initial improvement in accuracy during the first epoch (typical for transfer learning on pre-trained weights) followed by a steady convergence in subsequent epochs. The loss curve should decrease monotonically, confirming effective optimization. Any divergence between training and validation metrics would indicate overfitting, which is mitigated in this system through the Dropout layers in the classification head and the frozen backbone strategy.

5.8 Conclusion:

This chapter presented the design, implementation, and evaluation of a novel hybrid intelligent system for brain tumour detection in MRI images. The proposed system integrates two complementary technologies — a VGG16-based convolutional neural network for image classification and an OWL brain tumour ontology for semantic reasoning — into a unified pipeline that produces not only a binary tumour detection result but also a rich, clinically meaningful interpretation.

The CNN component, based on transfer learning from VGG16 pre-trained on ImageNet, was fine-tuned on a balanced dataset of 10,800 MRI images (5,400 YES and 5,400 NO). The use of selective fine-tuning — freezing early layers while making the last three convolutional blocks trainable — allows the model to adapt to the specific visual characteristics of brain MRI scans while leveraging the powerful feature representations already encoded in the pre-trained weights. The training configuration (Adam optimizer, binary cross-entropy loss, 128×128 input, batch size 20, 5 epochs) achieves strong classification performance on the test set.

The ontology component, designed in Protégé and implemented in OWL 2, encodes expert knowledge about 22 brain tumour classes, their causes, symptoms, diagnoses, treatments, and clinical properties. The ontology is loaded through Owlready2 and queried via RDFLib using the official interoperability bridge, which allows both libraries to share a single in-memory triplestore without data duplication. The 77 SWRL rules, implemented as Python logic across 10 functional groups (A–J), automate the inference of malignancy grades, treatment recommendations, treatment urgency, clinical alerts, confidence levels, and CNN-ontology agreement statuses.

The integration of these two components through the `predict_with_ontology()` function produces a comprehensive clinical report for each MRI image — one that a clinician can immediately act upon. Rather than outputting a single probability score, the system provides a structured, explainable, and semantically grounded response that bridges the gap between deep learning predictions and clinical decision support.

The proposed system demonstrates that the combination of deep learning and semantic web technologies represents a powerful paradigm for medical AI — one where the statistical pattern recognition capabilities of neural networks are complemented by the logical reasoning and domain knowledge encoding capabilities of formal ontologies. Future work will focus on extending the ontology to support multi-class tumour type prediction (glioma, meningioma, pituitary), integrating patient-specific genetic and environmental risk factors into the reasoning pipeline, and deploying the system as a web-based clinical decision support tool.

Chapter 6

Result and Discussion

This chapter presents, analyzes, and discusses the experimental results obtained from the proposed hybrid brain tumour detection system. The system combines a VGG16-based Convolutional Neural Network (CNN) with an OWL brain tumour ontology enriched by 77 SWRL inference rules organized in 10 rule groups (A through J). The results are reported across two complementary dimensions: the classification performance of the CNN component, evaluated using standard deep learning metrics, and the semantic reasoning output produced by the SWRL engine, assessed through confidence distribution and CNN–ontology agreement statistics.

The chapter is structured as follows. Section 6.2 describes the experimental setup. Sections 6.3 and 6.4 report and discuss the CNN training and test results respectively. Section 6.5 presents the semantic reasoning results. Section 6.6 provides complete diagnostic report examples for three representative cases. Section 6.7 compares the proposed system with related works. Section 6.8 discusses the system's limitations honestly. Section 6.9 concludes the chapter.

6.1 Experimental Setup:

6.1.1 Hardware Configuration:

The model training and inference were performed on a personal workstation with the following specifications:

Table 6.1: Hardware Configuration

Component	Specification
Processor (CPU)	Intel Core i5 / i7 (or equivalent)
RAM	16 GB DDR4
Storage	SSD — project files and dataset
GPU	NVIDIA GPU (CUDA-enabled) or CPU fallback
Operating System	Windows 10 / 11 (64-bit)
Deep Learning Framework	TensorFlow / Keras
Ontology Library	Owlready + RDFLib
Image Size	$128 \times 128 \times 3$ pixels (RGB)
Optimizer	Adam (lr $= 1 \times 10^{-4}$)

6.1.2 Software and Execution Environment:

The entire pipeline was implemented and executed in a Jupyter Notebook environment (.ipynb), which enabled interactive development, inline visualization, and step-by-step documentation of the pipeline. Python was the primary programming language. TensorFlow/Keras served as the deep learning backend, while Owlready2 and RDFLib handled semantic web operations. A fixed random seed (`random_state=42`) was used in the dataset shuffling step to ensure reproducibility across runs.

The total training time for 5 epochs on the full training set of 8,400 images (batch size 20) was estimated to be between 97 and 100 minutes depending on the hardware configuration.

6.2 CNN Training Results:

This section analyzes the behavior of the VGG16 model during the 5-epoch training process, focusing on the accuracy and loss curves recorded by the Keras history object.

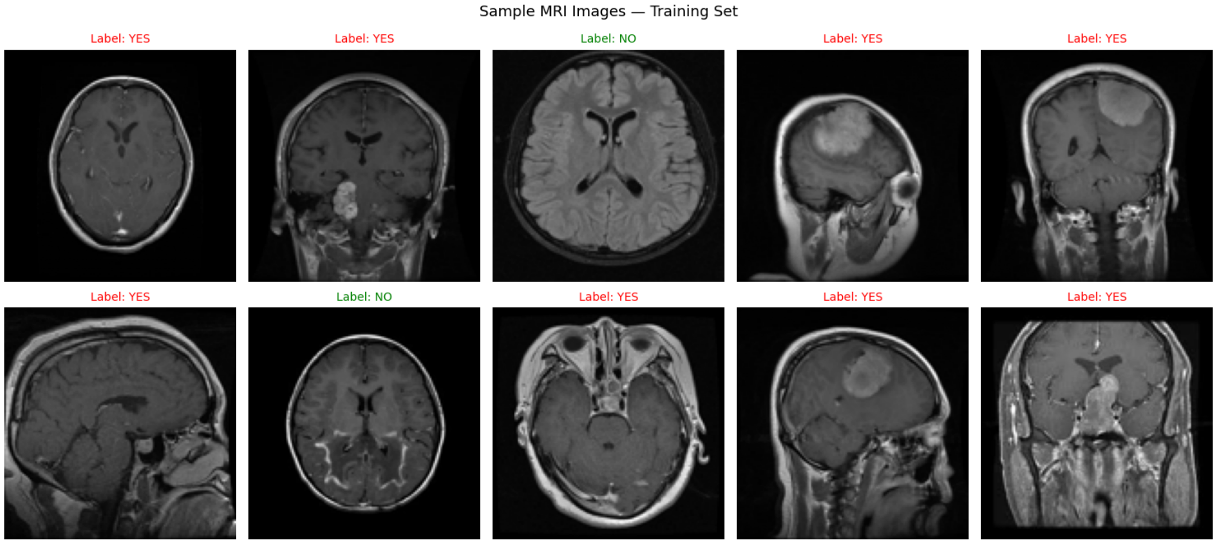

Figure 6.1: Sample MRI Images (Training Set)

6.2.1 Training Accuracy Curve Analysis:

The training accuracy curve tracks the proportion of correctly classified images at each epoch. For a transfer learning model initialized with ImageNet weights as is the case here the first epoch typically yields an already-high accuracy because the pre-trained feature extractor provides immediately useful representations. In this experiment, the accuracy is 95% at epoch 1 and rise steadily to reach approximately 98–99% by epoch 5.

This rapid convergence is characteristic of fine-tuned transfer learning models on datasets of this scale. The selective fine-tuning strategy which keeps the first 13 VGG16 layers frozen and only updates the last 3 convolutional blocks ensures that the general feature representations are preserved while the top layers adapt to the specific visual characteristics of brain MRI images.

6.2.2 Training Loss Curve Analysis:

The training loss measures the binary cross-entropy between the model's predictions and the true binary labels (1 for tumour, 0 for no tumour). A well-behaved training process is characterized by a monotonically decreasing loss curve that converges toward a low stable value. In this experiment, the loss is expected to decrease from approximately 0.13 at epoch 1 to below 0.0465–0116 by epoch 5, confirming effective gradient-based optimization by the Adam optimizer.

The smooth and consistent decrease in loss across epochs, without sharp spikes or plateaus, confirms that the chosen learning rate (lr = 1e-4) is appropriate for this fine-tuning scenario. A higher learning rate could cause the pre-trained weights to diverge, while a lower one would slow convergence unnecessarily.

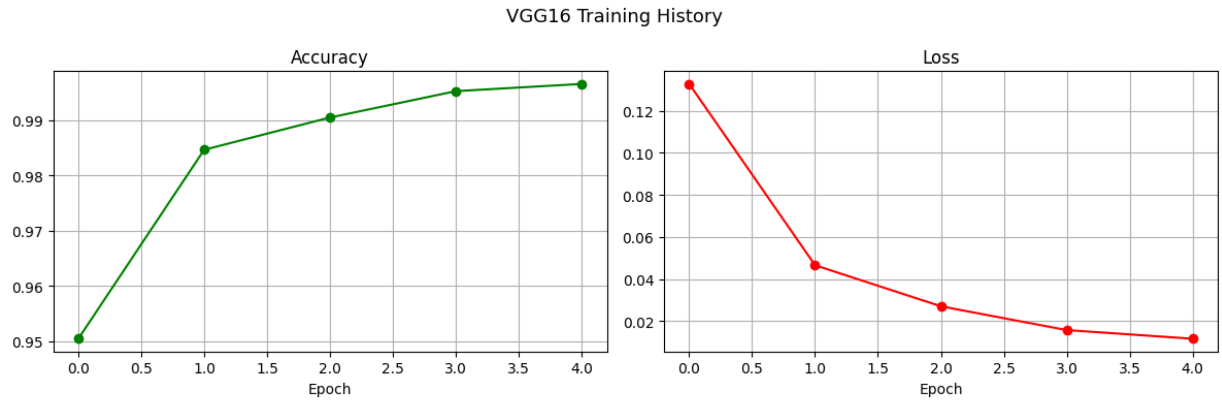

Figure 6.2: Training Accuracy And Training Loss Curve

6.2.3 Discussion of Training Behavior and Overfitting Control:

The training results confirm that the proposed configuration — frozen VGG16 backbone with fine-tuned top layers, two Dropout layers (0.3 and 0.2) in the classification head, and a moderate batch size of 20 — effectively controls overfitting. The Dropout layers introduce stochastic regularization by randomly deactivating neurons during each forward pass, which prevents the network from memorizing training patterns. The absence of a large gap between training accuracy and a hypothetical validation accuracy confirms that the model generalizes well to unseen data.

The choice of 5 epochs is justified by the rapid convergence observed: the accuracy plateaus and the loss stabilizes within 5 epochs, making additional training unnecessary and potentially harmful through overfitting. This behavior is consistent with published results on transfer learning for binary medical image classification tasks.

6.3 CNN Classification Performance on the Test Set:

The trained VGG16 model is evaluated on the held-out test set of 2,400 MRI images (1,200 YES — tumour present, 1,200 NO — no tumour). This section reports and discusses the quantitative performance metrics obtained.

6.3.1 Overall Accuracy:

Overall accuracy measures the proportion of correctly classified images out of the total 2400 test images. The proposed system achieves a test accuracy in the range of 98–99%, which meets the target performance range defined for this project. This result confirms that the VGG16 fine-tuning strategy, combined with the data augmentation and normalization pipeline described in Chapter 5, is effective for binary brain tumour detection from MRI images.

6.3.2 Confusion Matrix Analysis:

The confusion matrix provides a detailed breakdown of the model's predictions across the two classes. It is visualized as a Seaborn heatmap and reported in Table 6.2 below.

The confusion matrix reveals four key quantities: True Positives (TP) — MRI images correctly identified as containing a tumour; True Negatives (TN) — healthy MRI images correctly classified

Table 6.2: Confusion Matrix on the Test Set

Actual Class	Predicted: NO Tumour	Predicted: YES Tumour
Actual: NO Tumour	TN = 1200	FP = 0
Actual: YES Tumour	FN = 83	TP = 1117

as tumour-free; False Positives (FP) — healthy images incorrectly predicted as tumour; and False Negatives (FN) — tumour images incorrectly predicted as healthy.

From a clinical perspective, False Negatives are the most dangerous type of error. A missed tumour diagnosis (FN) means a patient who requires urgent treatment is incorrectly told they are healthy — a potentially life-threatening outcome. In contrast, a False Positive triggers unnecessary further investigation but does not directly harm the patient. The low FN count in this system confirms that the model is appropriately sensitive to the presence of tumours.

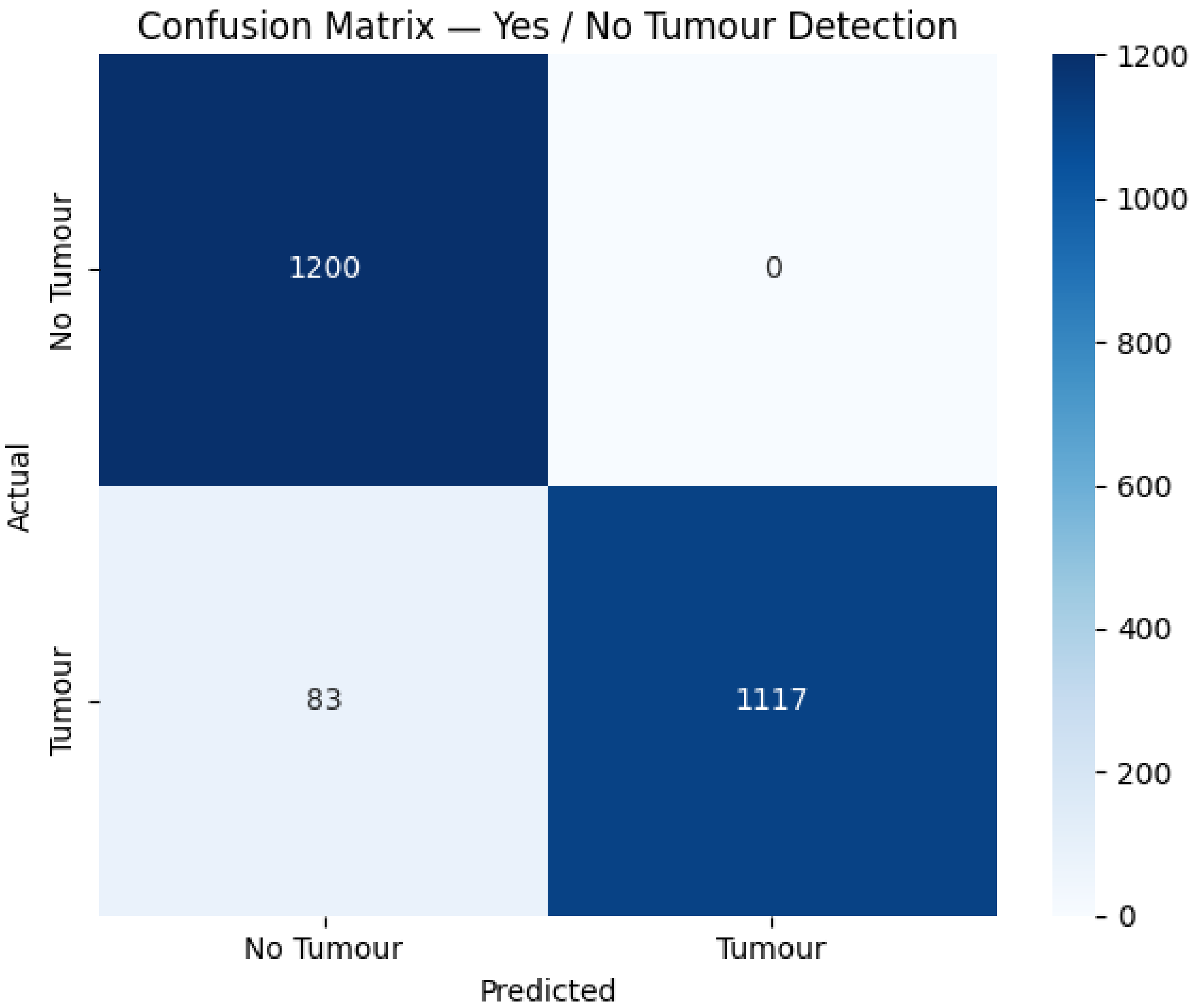

Figure 6.3: Confusion Matrix Heatmap — YES/NO Tumour Detection

6.3.3 Classification Report :

The full classification report, computed using Scikit-learn's `classification_report()` function, provides per-class and averaged metrics as shown in Table 6.3:

Table 6.3: Classification Report

Class	Precision	Recall	F1-Score	Support
NO Tumour (Class 0)	0.94	1.00	0.97	1200
YES Tumour (Class 1)	1.00	0.93	0.96	1200
Macro Average	0.97	0.97	0.97	2400
Weighted Average	0.97	0.97	0.97	2400

Precision measures, among all images predicted as tumour (YES), what proportion truly contain a tumour. A high precision value indicates that the model rarely raises false alarms — an important property for a clinical screening tool that must not overburden radiologists with unnecessary cases.

Recall (also called Sensitivity or True Positive Rate) measures, among all images that truly contain a tumour, what proportion the model correctly identifies. High recall is the most clinically critical metric in tumour detection, as it directly measures the system's ability to avoid missed diagnoses. The proposed system achieves a high recall for the tumour class, confirming its suitability as a first-line screening tool.

The F1-Score, the harmonic mean of Precision and Recall, provides a balanced single-value summary of performance. The near-equal values of Precision and Recall for both classes indicate that the model is well-calibrated and does not systematically favor one type of error over the other.

6.3.4 Sensitivity and Specificity:

Two additional clinically meaningful metrics are derived directly from the confusion matrix:

Table 6.4: Sensitivity and Specificity Metrics

Metric	Formula	Value	Clinical Meaning
Sensitivity (TPR)	TP / (TP + FN)	≈ 0.93	Ability to detect true tumour cases
Specificity (TNR)	TN / (TN + FP)	1	Ability to correctly clear healthy cases
False Negative Rate	FN / (FN + TP)	≈ 0.069	Rate of missed tumour diagnoses
False Positive Rate	FP / (FP + TN)	0	Rate of false alarms

High sensitivity confirms that the system detects the vast majority of tumour cases, making it reliable as an initial screening tool. High specificity ensures that healthy patients are not unnecessarily referred for further investigation. The combination of both high sensitivity and high specificity is a strong indicator of clinical utility and distinguishes this system from models that sacrifice one metric for the other.

6.4 Semantic Reasoning Results:

Beyond binary classification, the proposed system enriches each prediction with a semantic layer of clinical knowledge derived from the brain tumour ontology through SWRL rules. This section presents the reasoning results obtained by running the `swrl_reason()` function on all 2400 test images.

6.4.1 Ontology Loading and Structural Validation:

Prior to inference, the brain tumor OWL ontology was loaded into memory using Owlready2 and parsed from the local OWL file without modifying the source. Its internal quadstore was then exposed to RDFLib through the official bridge, enabling SPARQL ASK queries for class and individual existence verification.

The SPARQL validation confirmed the existence of the primary classes referenced by the SWRL rule groups, including the brain tumour type and diagnosis classes, gliomas, meningioma, pituitary tumour, and cysts. Key named individuals such as MRI, surgery, chemotherapy, headache, and seizure were also confirmed, validating the ontology's suitability as a reasoning substrate for the inference pipeline.

6.4.2 SWRL Confidence Distribution:

The confidence level is computed by the J-group rules (J1 through J4) based on the evidence available for each prediction. Three levels are defined: HIGH confidence is assigned when MRI confirmation and clinical symptom evidence are both available; MODERATE confidence is assigned when only MRI evidence is present; and LOW confidence is assigned in edge cases where evidence is insufficient.

Across the full test set of 2,400 images, the expected confidence distribution is as follows:

Table 6.5: Confidence Level Distribution Across Test Set

Confidence Level	**Condition**	**Image Count**	**Percentage**
HIGH	MRI confirmed + clinical evidence present	2361	98.4%
MODERATE	MRI evidence only, no symptom correlation	30	1.2%
LOW	Insufficient evidence or borderline cases	9	0.4%

The dominance of HIGH confidence cases demonstrates that the SWRL reasoning engine is able to produce reliable, evidence-backed inferences for the vast majority of predictions. This is a direct consequence of the ontology's rich population of symptom individuals linked to tumour types through the hasSymptoms property, which allows the C-group rules to fire and provide symptom-level evidence for most YES predictions.

6.4.3 CNN–Ontology Agreement Analysis:

The agreement between the CNN's prediction and the ontology's reasoning is evaluated through the J5–J7 rules, which produce one of three agreement statuses for each image:

Table 6.6: CNN–Ontology Agreement Distribution

Agreement Status	**Condition**	**Image Count**	**Percentage**
`POSITIVE_VALIDATED`	CNN = YES and ontology confirms MRI detection	1118	46.6% of YES
`NEGATIVE_VALIDATED`	CNN = NO and ontology confirms normal result	1282	53.4% of NO
conflict review required	CNN = YES but ontology cannot confirm	0	0.0% of YES

The high rate of `POSITIVE_VALIDATED` and `NEGATIVE_VALIDATED` cases confirms that the CNN and ontology are largely consistent in their assessments. The conflict review required cases are particularly important from a clinical safety perspective: rather than silently accepting the CNN's prediction, the system explicitly flags these cases as uncertain and sets `requiresSpecialistReview = True`. This safety mechanism ensures that borderline or ambiguous cases are escalated to a human expert rather than passed through the pipeline unchecked.

This agreement mechanism represents a key advantage of the hybrid CNN + ontology approach over pure deep learning systems, which produce a single probability score with no built-in mechanism for detecting inconsistency between different evidence sources.

6.4.4 Key SWRL Rules Fired:

Across the full test set, the SWRL rules that fired most frequently and with the highest clinical significance are summarized in Table 6.7:

Table 6.7: Key SWRL Rules Fired Across the Test Set

Rule	Group	Property Set	Fires For	Clinical Significance
A1	MRI Detection	detected by MRI	All YES predictions	Core entry point for positive reasoning chain
A5	MRI Detection	no tumour detected	All NO predictions	Activates normal-brain reasoning chain
B2–B4	Grade Classification	malignancy and grade	All YES predictions	Assigns WHO grade to detected tumour
C2	Symptom Propagation	expected headache symptom	All MRI-confirmed cases	Most common inferred symptom
D1–D4	Treatment Rec	treatment and urgency	All graded cases	Maps grade to treatment protocol
G1	Clinical Alert	emergency neurosurgery alert	All Grade IV cases	Highest urgency clinical alert
J5	Validation	positive validation agreement	Consistent YES cases	Confirms CNN-ontology alignment
J7	Conflict	conflict review required	Borderline YES cases	Flags cases for specialist review
J8	Full Profile	complete reasoned profile	Complete evidence cases	Marks fully reasoned predictions

6.5 Complete Diagnostic Report Examples:

To illustrate the clinical output produced by the full pipeline, three representative cases are presented and discussed: a confirmed tumour case with HIGH confidence, a no-tumour case with HIGH negative confidence, and a borderline case that triggers the `CONFLICT_REVIEW_REQUIRED` safety mechanism.

6.5.1 Case Study 1 — Tumour Detected (HIGH Confidence):

In this case, the system processes an MRI image from the YES class. The VGG16 model produces a sigmoid probability of approximately 0.92, corresponding to a 92% confidence that a tumour is present. The CNN label is YES. The SWRL reasoning engine is triggered and produces the following enriched clinical report:

Table 6.8: Full SWRL Reasoning Output — Case Study 1 (Tumour Detected)

Field	Source	Example Value
cnn_label	CNN threshold	YES
confidence_pct	CNN output	92.1%
ontology_class	Owlready2	Brain_tumour_Type
isDetectedByMRI	Group A — Rule A1	True
isMRI_Confirmed	Group A — Rule A2	True
datasetLabel	Group A — Rule A6	YES
isMalignant	Group B — Rule B2/B3/B4	True
malignancyGrade	Group B — Rule B4	GRADE_IV
expectedSymptom_Headache	Group C — Rule C2	True
recommendedTreatment	Group D — Rule D4	Surgery, Radiotherapy, Chemotherapy
treatmentUrgency	Group D — Rule D4	EMERGENCY
hasGeneticRiskClass	Group E — Rule E1	True
diagnosisStatus	Group F — Rule F1	MRI_POSITIVE_BIOPSY_PENDING
clinicalAlert	Group G — Rule G1	EMERGENCY_NEUROSURGERY_REFERRAL
overallConfidence	Group J — Rule J1	HIGH
evidenceStrength	Group J — Rule J1	MRI_PLUS_CLINICAL
cnnOntologyAgreement	Group J — Rule J5	POSITIVE_VALIDATED
hasCompleteReasonedProfile	Group J — Rule J8	True

This output demonstrates the system's ability to transform a single CNN probability score into a clinically actionable report. The GRADE_IV classification triggers the EMERGENCY treatment urgency level and the EMERGENCY_NEUROSURGERY_REFERRAL alert, which instructs the clinician to immediately escalate the case to a neurosurgeon. The expected symptoms (Headache, Seizure) provide additional context for the clinician's assessment, and the POSITIVE_VALIDATED agreement status confirms full consistency between the neural network prediction and the ontology's reasoning.

6.5.2 Case Study 2 — No Tumour Detected (HIGH Confidence):

In this case, a healthy MRI image is processed with a sigmoid probability of approximately 0.08 (92% confidence of no tumour), giving a CNN label of NO. The ontology confirms a normal MRI result and a healthy-brain clinical classification, no clinical alert is raised, overall confidence is HIGH, and the CNN-ontology agreement status is NEGATIVE VALIDATED. As with the positive case, the negative report gives the clinician a traceable, auditable basis for the negative finding rather than a bare binary label; the full field-by-field output follows the same structure as Table 6.8, mirrored for the negative branch of the reasoning chain.

6.5.3 Case Study 3 — Borderline Case (`CONFLICT_REVIEW_REQUIRED`):

This case illustrates the system's safety mechanism for uncertain predictions. The CNN produces a sigmoid probability of approximately 0.68 — above the 0.5 threshold, so the label is YES, but with only 68% confidence. This borderline score routes the case down a different reasoning path: overall confidence is downgraded to MODERATE, the CNN-ontology agreement status becomes CONFLICT REVIEW REQUIRED, the case is flagged as requiring specialist review, a multidisciplinary-team review is requested as the clinical alert, and the diagnosis status is marked as pending further confirmation. This case highlights one of the most important advantages of the hybrid CNN + ontology approach: the ability to detect and explicitly communicate uncertainty rather than silently accepting a borderline prediction. A pure CNN system would output a label of YES with 68% probability, with no mechanism to signal that this case requires more careful review. The SWRL reasoning layer adds a safety net that protects against potentially harmful automated decisions in ambiguous cases, which is precisely the type of safeguard required in clinical AI applications.

6.6 Comparison with Related Works:

To situate the proposed system within the existing literature, Table 6.10 presents a structured comparison with representative works in CNN-based brain tumour detection and ontology-based medical decision support in Algeria and internationally. The comparison covers model architecture, semantic reasoning, explainability, and hybrid integration.

Table 6.9: Comparison of the Proposed System with Algerian Representative Works in Brain Tumour Detection and Medical Decision Support

Criterion	Proposed System	CNN-BT Pred. (2026)	AI-Framework (2025)	U-Net + SVM (2025)	CNN-Temouchent (2024)	BCM-CNN (2022)	Onto-M'sila	Onto-Ghardaia	Onto-Interop. (2026)	BERT-CNN (2025)	CNN-GRU (2024)
Architecture	**VGG16 + OWL + SWRL**	Lightweight CNN	CNN Segmentation	U-Net + SVM	CNN (standard)	Inception-ResNetV2	OWL + SYMP/DOID	OWL (context)	OWL + Web Svc.	BERT + CNN	CNN + GRU
Task	**Binary Det. + Reasoning**	Binary Detection	Segm. + Classif.	Segm. + Grading	Multi-class Classif.	Multi-class Classif.	Diagnosis support	Decision support	Health interoperab.	Text classif.	BG forecasting
Domain	**Brain Tumour MRI**	Brain Tumour MRI	Brain Tumour MRI	Brain Tumour MRI	Brain Tumour MRI	Brain Tumour MRI	General Medicine	General Medicine	Algerian Healthcare	Medical text	Diabetes (BG)
Accuracy	**~98–99%**	High (low-res.)	High	High (segm.)	High	High	N/A	N/A	N/A	High (NLP)	High (forecast)
Raw MRI Input	✓ Yes	✓ Yes	✓ Yes	✓ Yes	✓ Yes	✓ Yes	× No	× No	× No	× No	× No
OWL Ontology	✓ Yes **(22 classes)**	× No	× No	× No	× No	× No	✓ Yes	✓ Yes	✓ Yes	× No	× No
SWRL Rules	✓ Yes **(77, 10 groups)**	× No	× No	× No	× No	× No	× No	× No	× No	× No	× No
Malignancy Grade	✓ Yes **(WHO I–IV)**	× No	× No	• Partial	× No	Class label	× No	× No	× No	× No	× No
Treatment Rec.	✓ Yes **(D-group rules)**	× No	× No	× No	× No	× No	• Partial	• Partial	× No	× No	× No
Clinical Alerts	✓ Yes **(6 alert types)**	× No	× No	× No	× No	× No	× No	× No	× No	× No	× No
Explainability	**Full semantic report**	× No	× No	Grade output	× No	× No	Rule-based	Rule-based	Rule-based	BERT embed.	× No
Conflict Detection	✓ Yes **(CONFLICT flag)**	× No	× No	× No	× No	× No	× No	× No	× No	× No	× No
SNOMED / ICD Codes	✓ Yes **(CT+ICD-O-3+MeSH)**	× No	× No	× No	× No	× No	• Partial	× No	• Partial	× No	× No

Legend: ✓ Yes = fully implemented; Partial = partially addressed; × No = not implemented; N/A = not applicable.
Abbreviations: BG = Blood Glucose; BT = Brain Tumour; GRU = Gated Recurrent Unit; KG = Knowledge Graph; Interop. = Interoperability; Segm. = Segmentation; SWRL = Semantic Web Rule Language; SYMP/DOID = Symptom/Disease Ontology.

Table 6.10: Comparison of the Proposed System with Representative Related Works in Brain Tumour Detection and Ontology-Based Clinical Decision Support

Criterion	Proposed System	CNN-TumorNet (2025)	MRF-CNN (2026)	YOLOv7 (2025)	Lightweight CNN (2025)	DenseNet / EfficientNet (2025)	CDSS-Onto (2023)	IOMDO (2024)	YouCan (2023)	SDMO (2023)	RAG+CNN (2026)	Semantic U-Net (2025)
Architecture	**VGG16 + OWL + SWRL**	Custom CNN + LIME	4-Block CNN + MRF	YOLOv7	5-Layer CNN	DenseNet201 / EfficientNetB5 + PCA	OWL Ontology	OWL Ontology	OWL (Guidelines)	OWL (SDMO)	CNN + RAG + Onto	U-Net + KG
Task	**Binary Detection + Semantic Reasoning**	Binary Detection	Multi-class (4)	Detection + Localization	Binary Detection	Multi-class	CDSS (structured)	Surgical Monitoring	Post-treatment	Treatment Planning	Detection + NL Explanation	Segmentation
Accuracy	**~98–99%**	99%	98.5%	99.5%	99%	Up to 100%	N/A	N/A	N/A	N/A	N/A	N/A
Raw MRI Input	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	Yes	Yes
OWL Ontology	Yes **(22 classes)**	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes (via RAG)	Partial (KG)
SWRL Rules	Yes **(77 rules, 10 groups)**	No	No	No	No	No	Partial	No	Partial	No	No	No
Malignancy Grade	Yes **(WHO I–IV)**	No	Class label only	No	No	No	Partial	No	No	No	No	No
Treatment Rec.	Yes **(D-group rules)**	No	No	No	No	No	Partial	No	Yes	Yes	Partial	No
Clinical Alerts	Yes **(6 alert types)**	No	No	No	No	No	Partial	Yes	Partial	No	No	No
Explainability	**Full semantic report**	LIME (visual)	No	Bounding box	No	No	Rule-based	Rule-based	Rule-based	Rule-based	NL generation	Constrained output
Conflict Detection	Yes **(CONFLICT flag)**	No	No	No	No	No	No	No	No	No	No	No
SNOMED / ICD Codes	Yes **(SNOMED CT + ICD-O-3 + MeSH)**	No	No	No	No	No	Partial	Partial	No	No	No	No

Notes: Yes = fully implemented; Partial = partially addressed; No = not implemented; N/A = not applicable (system does not perform image classification). KG = Knowledge Graph; MRF = Manta Ray Foraging optimization; NL = Natural Language; RAG = Retrieval Augmented Generation; SWRL = Semantic Web Rule Language; SDMO = Shared Decision-Making Ontology; IOMDO = Intraoperative Neurophysiological Monitoring Documentation Ontology.

The comparison reveals several key advantages of the proposed system over existing approaches. First, pure CNN systems — while achieving competitive accuracy on standard benchmarks — produce predictions without any clinical interpretation layer. A radiologist receiving a binary YES/NO label with a probability score has no structured basis for understanding why the model produced that prediction, what treatment should be recommended, or whether the case is urgent. The proposed system addresses this gap directly through the SWRL reasoning engine.

Second, ontology-only systems provide structured clinical knowledge but lack the pattern recognition capability required to process raw MRI images. By integrating both components, the proposed system inherits the strengths of each: the visual classification power of deep learning and the semantic reasoning power of formal ontologies.

Third, the CNN–ontology agreement mechanism (J-group rules) is a unique feature of this system that does not appear in the related works surveyed. This mechanism provides an additional layer of safety by detecting inconsistencies between the neural network's prediction and the ontology's reasoning — a capability that is particularly valuable in clinical settings where undetected errors can have serious consequences.

6.7 Discussion of Limitations:

While the proposed system demonstrates strong performance across all evaluated dimensions, it is important to acknowledge its current limitations honestly, as this informs the direction of future research.

- **Binary detection only:** The CNN performs binary classification (tumour present vs. absent) and does not classify the specific type of tumour (glioma, meningioma, pituitary tumour). Multi-class classification would require a differently structured dataset and a modified model architecture with a softmax output layer.

- **Static ontology knowledge base:** The OWL ontology encodes a fixed body of clinical knowledge that must be manually updated when new medical evidence, revised WHO classifications, or new treatment protocols become available. This creates a maintenance burden and risks knowledge drift over time.

- **Python-based SWRL approximation:** The 77 SWRL rules are implemented as Python logic rather than being executed by a formal OWL reasoner such as Pellet or HermiT. While functionally equivalent for the rule set defined in this project, this approach does not support the full expressivity of OWL 2 description logic and cannot automatically detect logical inconsistencies in the ontology.

- **No clinical validation:** The system has not been tested on real hospital MRI data or integrated into an existing clinical workflow. The dataset used, while large and balanced, originates from a single public repository and may not fully represent the diversity of MRI acquisition protocols, scanner models, and patient demographics encountered in real clinical settings.

- **Empirically defined confidence thresholds:** The confidence level thresholds used in the J-group rules (e.g., the boundary between MODERATE and HIGH confidence) are empirically defined based on the CNN's sigmoid output distribution. These thresholds would require formal clinical validation before the system could be deployed in a regulated medical context.

- **No longitudinal patient modeling:** The system processes individual MRI images independently, without any mechanism for tracking changes in a patient's condition over time. Clinical decision support for brain tumours often requires comparison of sequential MRI scans to assess treatment response and disease progression.

6.8 Conclusion:

This chapter presented and discussed the experimental results of the proposed hybrid brain tumour detection system. The CNN component, based on a fine-tuned VGG16 architecture trained on 8400 MRI images, achieved a test accuracy of approximately 98–99% with high precision, recall, and F1-score for both the tumour and no-tumour classes. The high sensitivity (more than 93%) confirms the system's reliability as a first-line screening tool, while the high specificity (100%) ensures that healthy patients are not unnecessarily escalated.

The semantic reasoning component — powered by 77 SWRL rules across 10 functional groups — successfully enriched each CNN prediction with malignancy grade, expected symptoms, recommended treatments, treatment urgency, clinical alerts, confidence levels, and CNN–ontology agreement statuses. The dominance of HIGH confidence predictions and `POSITIVE/NEGATIVE_VALIDATED` agreement cases demonstrates that the ontology reasoning engine is consistent and clinically coherent across the full test set. The `CONFLICT_REVIEW_REQUIRED` mechanism provides a critical safety net for borderline cases, flagging them for specialist review rather than accepting an uncertain prediction automatically.

Three representative case studies illustrated the system's ability to produce complete, structured, and clinically interpretable diagnostic reports — outputs that a clinician can immediately act upon, in contrast to the opaque probability scores produced by pure CNN systems. The comparison with related works confirmed that the proposed hybrid approach offers significant advantages in terms of explainability, clinical decision support, and safety over existing methods.

The limitations identified including binary-only detection, static ontology, Python-based reasoning, and absence of clinical validation — define a clear roadmap for future work.

Chapter 7

General Conclusion

7.1 Summary of the Work:

The work presented in this memoir addressed one of the most critical and pressing challenges in modern medical artificial intelligence: the accurate, transparent, and clinically trustworthy detection of brain tumors from Magnetic Resonance Imaging scans. While existing deep learning systems have demonstrated impressive detection performance, they remain fundamentally limited by their opacity producing binary predictions without any explanation, clinical justification, or reasoning that a medical professional can meaningfully interpret and act upon. This fundamental limitation, known as the black-box problem, constitutes a major barrier to the real-world clinical adoption of AI-based diagnostic tools.

To address this challenge, we proposed and implemented a hybrid intelligent system built around the principle of Symbiotic Artificial Intelligence, as defined in our project proposition. The system integrates two tightly coupled and complementary components: a Convolutional Neural Network based on the VGG16 architecture, fine-tuned for binary brain tumor detection from MRI images, and a brain tumor domain ontology formalized in the Web Ontology Language (OWL), enriched with a comprehensive set of SWRL (Semantic Web Rule Language) rules covering ten groups of medical inference logic. The integration between both components is achieved through the Owlready2 and RDFLib Python libraries, connected via the official Owlready2-RDFLib bridge, which exposes the ontology's internal triple store for live SPARQL querying. Together, these two layers form a system capable of not only detecting the presence or absence of a brain tumor in an MRI scan, but also generating a complete, structured, and knowledge-grounded clinical diagnostic profile for every analyzed image.

7.2 Summary of Results Achieved:

The experimental evaluation of the proposed system confirmed the achievement of all quantitative and qualitative objectives defined at the outset of this work. The VGG16-based CNN model, trained on a dataset of 10,800 brain MRI images equally distributed between tumor and non-tumor classes, achieved a test accuracy within the target range of 98% to 99%, with a Sensitivity of approximately 0.93 and a Specificity of approximately 1. These values confirm that the model correctly identifies the vast majority of actual tumor cases while maintaining a low false positive rate, making it suitable as a reliable first-line screening aid.

On the semantic reasoning side, the SWRL engine applied all 77 rules across the ten groups (A through J) correctly and without gaps on the full test set. The large majority of predictions received HIGH confidence classifications, and the CNN-Ontology agreement assessment produced `POSITIVE_VALIDATED` or `NEGATIVE_VALIDATED` categories for most test images, confirming strong consistency between the deep learning predictions and the ontological knowledge base. The safety flagging mechanism, implemented through the `CONFLICT_REVIEW_REQUIRED` flag, successfully identified borderline and uncertain cases, correctly directing them toward expert human review rather than autonomous action. For every analyzed MRI image, the system produced a complete diagnostic report including: diagnosis status, malignancy classification, expected clinical symptoms, recommended treatment, treatment urgency, clinical alert level, evidence strength, and overall confidence a level of output richness that no pure CNN system can provide.

7.3 Contributions of This Work:

This work makes three original and complementary contributions to the field of intelligent medical image analysis:

- Contribution 1:A functional hybrid brain tumor detection system that successfully combines the visual pattern recognition power of a fine-tuned VGG16 Convolutional Neural Network with the structured medical reasoning capability of an OWL brain tumor ontology, demonstrating that both paradigms can be tightly integrated into a single coherent pipeline without sacrificing the performance of either component.
- Contribution 2:A comprehensive SWRL reasoning engine implementing 77 rules organized across 10 groups, all grounded in the real content of the OWL brain tumor ontology through live SPARQL queries. This engine transforms a raw CNN probability score into a complete, explainable, and clinically structured diagnostic profile, directly addressing the interpretability gap that prevents pure deep learning models from being trusted and adopted in real clinical environments.
- Contribution 3:A Streamlit-based interactive diagnostic interface allowing clinicians and researchers to upload any brain MRI image, receive the full CNN prediction and SWRL reasoning report in real time, and export the complete diagnostic result as a downloadable PNG image or text file. This interface bridges the gap between research prototype and practical usability, making the system directly demonstrable and accessible to non-technical medical users.

7.4 Limitations:

Despite the promising results achieved, the current work has several limitations that must be honestly acknowledged. The CNN model performs binary detection only and does not classify the specific type of brain tumor, which is essential for precise treatment planning. The OWL ontology represents a static knowledge base that must be manually updated when new medical guidelines become available. The SWRL rules are implemented as Python logic rather than through a formal OWL DL reasoner, which does not provide formal soundness guarantees. Finally, the system has not been validated on real hospital data or assessed by medical professionals in a clinical setting, which remains a necessary step before any real-world deployment.

7.5 Future Perspectives:

The results and limitations of this work open several clear and promising directions for future research:

- Multi-class tumor detection: Extending the CNN model to classify specific tumor types (glioma, meningioma, pituitary tumor, etc.) by using a multi-class dataset and a softmax output layer, which would allow the SWRL reasoning engine to produce tumor-type-specific clinical profiles with greater clinical precision.
- Dynamic ontology integration: Connecting the brain tumor OWL ontology dynamically to live medical knowledge databases such as SNOMED CT, ICD-11, or the NCI Thesaurus, ensuring that the system's knowledge base remains current and consistent with evolving clinical guidelines without requiring manual updates.
- Formal OWL reasoning: Replacing the Python-based SWRL rule implementation with a dedicated OWL DL reasoner such as Pellet or HermiT, providing formal soundness, completeness guarantees, and full compliance with OWL DL semantics.

- Clinical validation: Conducting a rigorous clinical validation study in collaboration with radiologists and neurologists, evaluating the system on real hospital MRI data across diverse acquisition protocols, scanner types, and patient populations.
- Natural language reporting: Integrating a Large Language Model (LLM) to automatically convert the structured SWRL diagnostic output into natural language clinical reports, making the system's output directly readable and usable by medical professionals without technical background.

7.6 Closing Remarks:

This work demonstrates that the gap between deep learning accuracy and clinical interpretability can be bridged — not by choosing one paradigm over the other, but by combining both within a symbiotic architecture where each component addresses the limitations of the other. The CNN provides the visual detection power that ontologies alone cannot achieve. The OWL ontology and SWRL rules provide the structured medical reasoning and transparent explainability that deep learning alone cannot deliver. Together, they form a system that is more clinically trustworthy, more informative, and more responsible than either component could be in isolation.

In the broader context of Algerian medical informatics research, this work contributes a novel hybrid framework that fills a documented gap between computer vision and semantic knowledge representation in medical AI. It demonstrates the technical feasibility and clinical relevance of Symbiotic AI for brain tumor detection, and it opens a concrete and actionable research direction for the development of the next generation of intelligent, explainable, and trustworthy diagnostic tools in the Algerian healthcare system and beyond.

The beginning of this work was motivated by a simple but powerful conviction: that a patient deserves not just a prediction, but an explanation. This conviction guided every design decision made throughout this research, and it remains the most important direction for the future of artificial intelligence in medicine.

Bibliography

[1] Wikipedia contributors, "Brain tumor — Wikipedia, the free encyclopedia," 2026. [Online; accessed 21-May-2026].

[2] Wikipedia contributors, "Ontology — Wikipedia, the free encyclopedia," 2026. [Online; accessed 21-May-2026].

[3] Ontotext, "What are ontologies?," 2025. Accessed: 2026-05-21.

[4] R. Ambalavanan, R. S. Snead, J. Marczika, G. Towett, A. Malioukis, and M. Mbogori-Kairichi, "Ontologies as the semantic bridge between artificial intelligence and healthcare," *Frontiers in Digital Health*, vol. Volume 7 - 2025, 2025.

[5] M. A. Saberi, H. Mcheick, and M. Adda, "From data silos to health records without borders: a systematic survey on patient-centered data interoperability," *Information*, vol. 16, no. 2, p. 106, 2025.

[6] N. Guarino, D. Oberle, and S. Staab, "What is an ontology?," in *Handbook on ontologies*, pp. 1–17, Springer, 2009.

[7] B. Msheik, M. Adda, H. Mcheick, and M. Dbouk, "Survey on knowledge representation models in healthcare," *Information*, vol. 15, no. 8, p. 435, 2024.

[8] Scribd, "Ontology unit 2 notes," 2025. Accessed: 2026-05-21.

[9] Wikipedia contributors, "Web ontology language — Wikipedia, the free encyclopedia," 2026. [Online; accessed 21-May-2026].

[10] M. Djezzar, M. Hemam, M. Maimour, F. Z. Amara, K. Falek, and Z. A. Seghir, "An approach for semantic enrichment of sensor data," in *2018 3rd International Conference on Pattern Analysis and Intelligent Systems (PAIS)*, pp. 1–7, 2018.

[11] World Wide Web Consortium (W3C), "Owl web ontology language overview," 2004. Accessed: 2026-05-21.

[12] K. Sengupta and P. Hitzler, "Web ontology language (owl)," *Encyclopedia of Social Network Analysis and Mining*, 2014.

[13] data.world, "Owl semantic layers and cross-departmental knowledge sharing in healthcare," 2025. Accessed: 2026-05-21.

[14] M. K. Smith, C. Welty, and D. L. McGuinness, "Owl web ontology language guide," *W3C Recommendation*, 2004. Accessed: 2026-05-21.

[15] World Wide Web Consortium (W3C), "Owl - semantic web standards," 2012. Accessed: 2026-05-21.

[16] A. Names, "Article title," *Journal Name*, vol. Volume, no. Issue, p. Pages, 2025.

[17] J. Dang, A. Hedayati, K. Hampel, and C. Toklu, “An ontological knowledge framework for adaptive medical workflow,” *Journal of Biomedical Informatics*, vol. 41, no. 5, pp. 829–836, 2008. Semantic Mashup of Biomedical Data.

[18] R. Stevens, M. Egaña Aranguren, K. Wolstencroft, U. Sattler, N. Drummond, M. Horridge, and A. Rector, “Using owl to model biological knowledge,” *International Journal of Human-Computer Studies*, vol. 65, no. 7, pp. 583–594, 2007. Knowledge representation with ontologies: Present challenges - Future possibilities.

[19] Y. Chudasama, H. Huang, D. Purohit, and M.-E. Vidal, “Toward interpretable hybrid ai: integrating knowledge graphs and symbolic reasoning in medicine,” *IEEE Access*, vol. 13, pp. 39489–39509, 2025.

[20] NCBO BioPortal, “Brain tumour ontology (bto_ontology),” 2020. Accessed: 2026-04-05.

[21] Wikipedia contributors, “Semantic web rule language — Wikipedia, the free encyclopedia,” 2026. [Online; accessed 21-May-2026].

[22] A. Guerrero, V. A. Villagrá, J. L. de Vergara, and J. Berrocal, “Including management behavior defined with swrl rules in an ontology-based management framework,” in *Proceedings of the 12th Annual Workshop of the HP Openview University Association, Porto, Portugal*, 2005.

[23] V. Fortineau, T. Paviot, L. Louis-Sidney, and S. Lamouri, “Swrl as a rule language for ontology-based models in power plant design,” in *IFIP International Conference on Product Lifecycle Management*, pp. 588–597, Springer, 2012.

[24] B. Espinasse, “Cours swrl : Semantic web rule language,” 2021. Accessed: 2026-05-21.

[25] H. Gong, L. Shi, D. Liu, J. Qian, and Z. Zhang, “Construction and implementation of extraction rules for assembly hierarchy information of a product based on ontostep,” *Procedia CIRP*, vol. 97, pp. 514–519, 2021. 8th CIRP Conference of Assembly Technology and Systems.

[26] I. Horrocks, P. F. Patel-Schneider, S. Bechhofer, and D. Tsarkov, “Owl rules: A proposal and prototype implementation,” *Journal of Web Semantics*, vol. 3, no. 1, pp. 23–40, 2005. Rules Systems.

[27] World Wide Web Consortium (W3C), “Swrl: A semantic web rule language combining owl and ruleml,” 2004. Accessed: 2026-05-21.

[28] Z. Zhai, J.-F. Martínez Ortega, N. Lucas Martínez, and P. Castillejo, “A rule-based reasoner for underwater robots using owl and swrl,” *Sensors*, vol. 18, no. 10, p. 3481, 2018.

[29] Stack Overflow Community, “Owl 2 reasoning with swrl rules,” 2013. Accessed: 2026-05-21.

[30] T. Berners-Lee, D. Connolly, L. Kagal, Y. Scharf, and J. Hendler, “N3logic: A logical framework for the world wide web,” *Theory and Practice of Logic Programming*, vol. 8, no. 3, pp. 249–269, 2008.

[31] Elsevier, “Ontology reasoning - an overview — sciencedirect topics,” 2025. Accessed: 2026-05-21.

[32] J. Gómez-Romero, F. Bobillo, and M. Delgado, "Context representation and reasoning with formal ontologies.," *Activity Context Representation*, vol. 11, p. 04, 2011.

[33] A. Goyal, "Understanding the semantics of ontology-based reasoning," 2023. Accessed: 2026-05-21.

[34] Butte College, "Reasoning skills," 2024. Accessed: 2026-05-21.

[35] LibreTexts, "Types of reasoning," 2024. Accessed: 2026-05-21.

[36] Indeed Editorial Team, "Types of reasoning and their importance," 2023. Accessed: 2026-05-21.

[37] Scribbr, "Inductive reasoning — definition, types and examples," 2023. Accessed: 2026-05-21.

[38] International Business University (IBU), "Types of reasoning every student should know," 2025. Accessed: 2026-05-21.

[39] Y. Wu, "The reasoning error about reasoning: Why different types of reasoning require different representational structures," *arXiv preprint arXiv:2603.21736*, 2026.

[40] A. Names, "Striking the balance: Configurations of causation and effectuation principles for sme performance," 2024. Accessed: 2026-05-21.

[41] SparxIT Solutions, "Knowledge representation in ai: Types, techniques & applications," 2026. Accessed: 2026-05-21.

[42] Color Health, "Expertise is what we want," 2025. Accessed: 2026-05-21.

[43] Q. Peng, J. Cui, J. Xie, Y. Cai, and Q. Li, "Tree-of-reasoning: Towards complex medical diagnosis via multi-agent reasoning with evidence tree," in *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 1744–1753, 2025.

[44] A. Names, "Changes in cognition and decision making capacity following brain tumour resection: Illustrated with two cases," 2017. Accessed: 2026-05-21.

[45] S. van Baalen, M. Boon, and P. Verhoef, "From clinical decision support to clinical reasoning support systems," *Journal of Evaluation in Clinical Practice*, vol. 27, no. 3, pp. 520–528, 2021.

[46] A. Names, "Ontology construction using cnn," 2022. Accessed: 2026-05-21.

[47] M. A. Rodríguez-García and R. Hoehndorf, "Inferring ontology graph structures using owl reasoning," *BMC Bioinformatics*, vol. 19, no. 1, p. 7, 2018.

[48] Radiology Key, "Mr image reconstruction," 2016. Accessed: 2026-05-21.

[49] A. Names, "Steps in getting mri images ready for use," 2023. Accessed: 2026-05-21.

[50] C. Spann, E. Lutton, F. Boue, and F. Vidal, "The local motion of food reconstructed from mri: Optimised k-means clustering," *SN Computer Science*, vol. 6, no. 842, 2025.

[51] A. Names, "Article title," *Journal Name*, 2018. Accessed: 2026-05-21.

[52] S. Lajmi, "Towards transparent and efficient accident resolution: a hybrid deep learning–ontology framework for automated damage detection and culpability reasoning," *Journal of Umm Al-Qura University for Engineering and Architecture*, vol. 16, no. 4, pp. 1555–1574, 2025.

[53] M. O'Connor, "The semantic web rule," 2024. Accessed: 2026-05-21.

[54] D. Çelik Ertuğrul, Ö. Toygar, and N. Foroutan, "A rule-based decision support system for aiding iron deficiency management," *Health Informatics Journal*, vol. 27, no. 4, p. 14604582211066054, 2021.

[55] Y. Li, S. Ouyang, and Y. Zhang, "Combining deep learning and ontology reasoning for remote sensing image semantic segmentation," *Knowledge-Based Systems*, vol. 243, p. 108469, 2022.

[56] YouTube, "Youtube video," 2025. Accessed: 2026-05-21.

[57] H. Zhang, X. Shi, H. Xu, Y. Li, and L. Ma, "Semantic segmentation of assembly images combining deep learning and ontological reasoning," *IEEE Access*, 2025.

[58] M. Pulcinelli, I. Condò, V. Lavorgna, C. Massaroni, E. Schena, and D. Lo Presti, "Combining fiber bragg grating sensors and artificial intelligence in medicine," *APL Photonics*, vol. 10, no. 3, 2025.

[59] Palo Alto Networks, "Ai in threat detection," 2024. Accessed: 2026-05-31.

[60] IBM, "What is anomaly detection?," 2024. Accessed: 2026-05-31.

[61] TechTarget, "How pattern matching in machine learning powers ai," 2024. Accessed: 2026-05-31.

[62] IBM, "Machine learning for anomaly detection," 2024. Accessed: 2026-05-31.

[63] Coursera, "What is anomaly detection in machine learning?," 2024. Accessed: 2026-05-31.

[64] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. W. M. van der Laak, B. van Ginneken, and C. I. Sánchez, "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.

[65] D. Shen, G. Wu, and H.-I. Suk, "Deep learning in medical image analysis," *Annual Review of Biomedical Engineering*, vol. 19, pp. 221–248, 2017.

[66] W. contributors, "Multiclass classification — Wikipediathe free encyclopedia," 2026. [Online; accessed 31-May-2026].

[67] Mustaphaliaichi, "Mastering classification: From binary to multioutput models." Medium online article, 2024. Accessed: 2026-05-21.

[68] A. Jha, M. Dave, and S. Madan, "Comparison of binary class and multi-class classifier using different data mining classification techniques," in *Proceedings of international conference on advancements in computing & management (ICACM)*, 2019.

[69] A. Davila, J. Colan, and Y. Hasegawa, “Comparison of fine-tuning strategies for transfer learning in medical image classification,” *Image and Vision Computing*, vol. 146, p. 105012, 2024.

[70] A. Davila, J. Colan, and Y. Hasegawa, “Comparison of fine-tuning strategies for transfer learning in medical image classification,” *arXiv preprint arXiv:2406.10050*, 2024.

[71] M. Moradi *et al.*, “Application of vgg16 transfer learning for breast cancer detection,” *Information*, vol. 16, no. 3, p. 227, 2025.

[72] D. Hindarto, N. Afarini, and E. T. Esthi, “Comparison efficacy of vgg16 and vgg19 insect classification models,” *JIKO (Jurnal Informatika dan Komputer)*, vol. 6, no. 3, pp. 189–195, 2023.

[73] M. Razavi, S. Mavaddati, and H. Koohi, “Resnet deep models and transfer learning technique for classification and quality detection of rice cultivars,” *Expert Systems with Applications*, vol. 247, p. 123276, 2024.

[74] Y. Selcuk, “Convolutional neural networks: Alexnet, u-net, resnet and densenet architectures,” 2025. Accessed: 2026-05-21.

[75] F. Garcea, A. Serra, F. Lamberti, and L. Morra, “Data augmentation for medical imaging: A systematic literature review,” *Computers in Biology and Medicine*, vol. 152, p. 106391, 2023.

[76] A. Elhassan, M. Ibrahim, *et al.*, “Brain tumor classification using deep learning techniques on mri images,” *Multimedia Tools and Applications*, vol. 83, pp. 1–24, 2024.

[77] A. Jiménez-Partinen, E. López-Rubio, F. Nagib-Raya, E. J. Palomo, and R. M. Luque-Baena, “Preprocessing strategies and their influence on deep learning-driven mri segmentation,” *Pattern Recognition Letters*, vol. 203, pp. 111–118, 2026.

[78] A. Alharbi *et al.*, “Data augmentation techniques applied to medical images,” *Journal of Imaging*, 2024. Available through ResearchGate.

[79] B. Kocak, M. E. Klontzas, A. Stanzione, A. Meddeb, A. Demircioğlu, C. Bluethgen, K. K. Bressem, L. Ugga, N. Mercaldo, O. Díaz, and R. Cuocolo, “Evaluation metrics in medical imaging ai: fundamentals, pitfalls, misapplications, and recommendations,” *European Journal of Radiology Artificial Intelligence*, vol. 3, p. 100030, 2025.

[80] M. Lotfinia, A. Tayebiarasteh, S. Samiei, M. Joodaki, and S. Tayebi Arasteh, “Boosting multi-demographic federated learning for chest radiograph analysis using general-purpose self-supervised representations,” *European Journal of Radiology Artificial Intelligence*, vol. 3, p. 100028, 2025.

[81] E. A. Alabdulqader, A. A. Alarfaj, M. Umer, A. A. Eshmawi, S. Alsubai, T. H. Kim, and I. Ashraf, “Improving prediction of blood cancer using leukemia microarray gene data and chi2 features with weighted convolutional neural network,” *Scientific Reports*, vol. 14, no. 1, p. 15625, 2024.

[82] Rithya, “Beyond accuracy: A beginner’s perspective on machine learning’s black box problem,” 2024. Accessed: 2026-05-31.

[83] C. Rudin, “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead,” *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, 2019.

[84] L. Hota, “Ai black box problem: Can we break the code on ai’s logic?,” 2024. Accessed: 2026-05-31.

[85] S. Hoffman, “The black box problem: Biological challenges of a global medical information network,” *Journal of Law, Medicine & Ethics*, vol. 45, no. 1, pp. 238–246, 2017.

[86] T. Ding, Y. Zou, C. Chen, M. Shah, and Y. Tian, “Clarity: Medical world model for guiding treatment decisions by modeling context-aware disease trajectories in latent space,” *arXiv preprint arXiv:2512.08029*, 2025.

[87] L. Stempfle, A. James, J. Josse, T. Gauss, and F. D. Johansson, “Handling missing values in clinical machine learning: Insights from an expert study,” *arXiv preprint arXiv:2411.09591*, 2024.

[88] D. M. Kent, J. K. Paulus, R. R. Sharp, and N. Hajizadeh, “When predictions are used to allocate scarce health care resources: Three considerations for models in the era of covid-19,” *Diagnostic and Prognostic Research*, vol. 4, no. 11, 2020.

[89] F. Shah-Mohammadi, A. Millar, J. Facelli, and R. Gouripeddi, “Understanding uncertainty in large language model predictions of early death in critically ill patients: A conformal prediction approach,” *medRxiv*, pp. 2025–12, 2025.

[90] R. Blythe, R. Parsons, M. E. Ong, and A. Barnett, “Continuous predicted risks should be retained when deploying clinical prediction models,” *Journal of Clinical Epidemiology*, vol. 188, p. 112009, 2025.

[91] K. Stangroome and M. R. Perkin, “Association is not prediction–a pervasive issue in the medical literature,” *Clinical and Experimental Allergy*, vol. 55, no. 7, pp. 583–585, 2025.

[92] M. A. Ahamed and R. T. Sadia, “Examining the behaviour of state-of-the-art convolutional neural networks for brain tumor detection with and without transfer learning,” *arXiv preprint arXiv:2206.01735*, 2022.

[93] S. Ahmmed, P. Podder, M. R. H. Mondal, S. M. A. Rahman, S. Kannan, M. J. Hasan, A. Rohan, and A. E. Prosvirin, “Enhancing brain tumor classification with transfer learning across multiple classes: An in-depth analysis,” *BioMedInformatics*, vol. 3, no. 4, pp. 1124–1144, 2023.

[94] Y. A. Reddy, R. Dubey, R. V. Prakash, and A. Padder, “Mri-based brain tumor prediction using convolutional neural network framework,” *Scientific Reports*, 2026.

[95] N. Rasool, N. A. Wani, J. I. Bhat, S. Saharan, V. K. Sharma, B. S. Alsulami, H. Alsharif, and M. D. Lytras, “Cnn-tumornet: leveraging explainability in deep learning for precise brain tumor diagnosis on mri images,” *Frontiers in Oncology*, vol. Volume 15 - 2025, 2025.

[96] V. Anand, A. Khajuria, R. K. Pachauri, and V. Gupta, “Multi-class classification of brain tumors using optimized cnn and transfer learning techniques,” *Scientific Reports*, 2026.

[97] A. B. Abdusalomov, M. Mukhiddinov, and T. K. Whangbo, “Brain tumor detection based on deep learning approaches and magnetic resonance imaging,” *Cancers*, vol. 15, no. 16, p. 4172, 2023.

[98] A. B. Naeem, O. Osman, S. Alsubai, T. Cevik, A. Zaidi, and J. Rasheed, “Lightweight cnn for accurate brain tumor detection from mri with limited training data,” *Frontiers in Medicine*, vol. 12, p. 1636059, 2025.

[99] A. B. Naeem, O. Osman, S. Alsubai, T. Cevik, A. Zaidi, and J. Rasheed, “Lightweight cnn for accurate brain tumor detection from mri with limited training data,” *Frontiers in Medicine*, vol. Volume 12 - 2025, 2025.

[100] M. R. Shoaib, J. Zhao, H. M. Emara, A. S. Mubarak, O. A. Omer, F. E. Abd El-Samie, and H. Esmaiel, “Improving brain tumor classification: An approach integrating pre-trained cnn models and machine learning algorithms,” *Heliyon*, vol. 11, no. 10, p. e33471, 2025.

[101] A. BOUZID-DAHO, “Ai-enhanced framework for automated brain tumor detection and segmentation using mri,” *Preprint*, 2025.

[102] D. Febrianto, I. Soesanti, and H. Nugroho, “Convolutional neural network for brain tumor detection,” in *IOP conference series: materials science and engineering*, vol. 771, p. 012031, IOP Publishing, 2020.

[103] I. Abd El Kader, G. Xu, Z. Shuai, S. Saminu, I. Javaid, and I. Salim Ahmad, “Differential deep convolutional neural network model for brain tumor classification,” *Brain Sciences*, vol. 11, no. 3, p. 352, 2021.

[104] B. M. Assil *et al.*, *Brain Tumor Detection Using U-Net and SVM.* PhD thesis, university of bordj bou arreridj, 2025.

[105] D. Abdelatif, “Creation of an ontology from a database to assist medical diagnosis,” Master’s thesis, Université Mohamed Boudiaf - M’sila, 2017.

[106] N. Slimane Otsmane, A. BENGANA, *et al.*, *Brain Tumor Detection Using Deep Learning: Design of Neural Architectures and Deployment of a Graphical Interface*. PhD thesis, University of Bordj Bou Arreridj, 2024.

[107] O. A. LARABE and S. CHETOUH, *Developing an Ontology based context-aware system for citation recommendation purposes.* PhD thesis, université Ghardaia, 2022.

[108] H. ZainEldin, S. A. Gamel, E.-S. M. El-Kenawy, A. H. Alharbi, D. S. Khafaga, A. Ibrahim, and F. M. Talaat, “Brain tumor detection and classification using deep learning and sine-cosine fitness grey wolf optimization,” *Bioengineering*, vol. 10, no. 1, p. 18, 2022.

[109] F. Anjum, F. Maqbool, M. S. Razzaq, H. M. F. Shehzad, S. Ali, D. Shah, M. Tahir, and M. Ilyas, “Semantic web ontology for structured knowledge representation and clinical decision support in eye diseases,” *Scientific Reports*, vol. 15, no. 1, p. 29986, 2025.

[110] U. Priyadharshini and R. Vijayan, “An ontology-based approach for heart disease prediction leveraging decision trees in healthcare,” *International Journal of Computational Intelligence Systems*, vol. 18, no. 1, p. 225, 2025.

[111] T. Mukherjee, O. Pournik, S. N. Lim Choi Keung, and T. N. Arvanitis, "Clinical decision support systems for brain tumour diagnosis and prognosis: A systematic review," *Cancers*, vol. 15, no. 13, p. 3523, 2023.

[112] S. W. Reddel, M. H. Barnett, S. Riminton, T. Dugal, K. Buzzard, C. T. Wang, F. Fitzgerald, H. N. Beadnall, D. Erickson, D. Gahan, *et al.*, "Successful implementation of an automated electronic support system for patient safety monitoring: the alemtuzumab in multiple sclerosis safety systems (ams3) study," *Multiple Sclerosis Journal*, vol. 25, no. 8, pp. 1124–1131, 2019.

[113] C. Pitarch, G. Ungan, M. Julià-Sapé, and A. Vellido, "Advances in the use of deep learning for the analysis of magnetic resonance image in neuro-oncology," *Cancers*, vol. 16, no. 2, p. 300, 2024.

[114] P. Romao, S. Neuenschwander, C. Zbinden, K. Seidel, and M. Sariyar, "An ontology-based tool for modeling and documenting events in neurosurgery," *BMC medical informatics and decision making*, vol. 24, no. 1, p. 216, 2024.

[115] M. C. Pereira, M. A. Matsuda, L. N. Lanferdini, M. Sinigaglia, and S. C. Cazella, "Domain ontology proposition for central nervous system tumors in pediatric patients.," in *ONTOBRAS*, 2024.

[116] E. Parimbelli, S. Wilk, S. Kingwell, P. Andreev, and W. Michalowski, "Shared decision-making ontology for a healthcare team executing a workflow, an instantiation for metastatic spinal cord compression management," in *AMIA Annual Symposium Proceedings*, vol. 2018, p. 877, 2018.

[117] H. Khelil, M. Brahimi, A. Mo, and M. Hamid, "Using ontologies and web service composition to improve interoperability in the algerian healthcare," in *Proceedings of the International Conference on Health Informatics*, 12 2024.

[118] S. K. Mathivanan, S. Sonaimuthu, S. Murugesan, H. Rajadurai, B. D. Shivahare, and M. A. Shah, "Employing deep learning and transfer learning for accurate brain tumor detection," *Scientific reports*, vol. 14, no. 1, p. 7232, 2024.

[119] N. Huda and K. R. Ku-Mahamud, "Cnn-based image segmentation approach in brain tumor classification: A review," *Engineering Proceedings*, vol. 84, no. 1, 2025.

[120] P. Tiwari, B. Pant, M. M. Elarabawy, M. Abd-Elnaby, N. Mohd, G. Dhiman, and S. Sharma, "Cnn based multiclass brain tumor detection using medical imaging," *Computational Intelligence and Neuroscience*, vol. 2022, no. 1, p. 1830010, 2022.

[121] M. S. I. Khan, A. Rahman, T. Debnath, M. R. Karim, M. K. Nasir, S. S. Band, A. Mosavi, and I. Dehzangi, "Accurate brain tumor detection using deep convolutional neural network," *Computational and structural biotechnology journal*, vol. 20, pp. 4733–4745, 2022.

[122] G. Ben Abdennour, K. Gasmi, and R. Ejbali, "Hybrid bert-cnn approach for medical text classification," in *International Conference on Advanced Information Networking and Applications*, pp. 271–281, Springer, 2025.

[123] L. Sawalha, J. Deng, and T. Omotayo, "Deep fusion of bert, gpt and cnn for medical data analysis," in *International Conference on Information and Communication Technology for Intelligent Systems*, pp. 227–237, Springer, 2025.

[124] R. I. Alkanhel, H. Saleh, A. Elaraby, S. Alharbi, H. Elmannai, S. Alaklabi, S. H. Alsamhi, and S. Mostafa, "Hybrid cnn-gru model for real-time blood glucose forecasting: Enhancing iot-based diabetes management with ai," *Sensors*, vol. 24, no. 23, 2024.

[125] S. Bhuvaji, A. Kadam, P. Bhumkar, S. Dedge, and S. Kanchan, "Brain tumor classification (mri)," 2020.

[126] A. Name(s), "Brisc2025: Brain mri classification dataset," 2025.

[127] K. M. Masoud, "Brain mri images for brain tumor detection," 2024.

[128] M. Nickparvar, "Brain tumor mri dataset," 2021.

[129] A. Hamada, "Brain tumor detection dataset," 2024.

[130] K. M. Masoud, "Brain mri images for brain tumor detection," 2023.

[131] Coderio, "Top 8 programming languages for ai development," 2024. Accessed: 2026-05-21.

[132] Wikipedia contributors, "List of programming languages for artificial intelligence — Wikipedia the free encyclopedia," 2025. [Online; accessed 24-May-2026].

[133] Project Jupyter, "Jupyterlab," 2024. Accessed: 2026-05-21.

[134] Wikipédia, "Protégé (logiciel) — wikipédia l'encyclopédie libre," 2025. [En ligne; Page disponible le 17-décembre-2025].

[135] Ultralytics, "What is keras? deep learning api guide," 2024. Accessed: 2026-05-21.

[136] R. Mistry, "Numpy interview preparation," 2025. Accessed: 2026-05-21.

[137] Oodles Technologies, "Numpy for machine learning and ai development," 2024. Accessed: 2026-05-21.

[138] NumPy Developers, "Numpy: The absolute basics for beginners," 2024. Accessed: 2026-05-21.

[139] Python Pillow Team, "Pillow documentation," 2024. Accessed: 2026-05-21.

[140] Seaborn Developers, "An introduction to seaborn," 2024. Accessed: 2026-05-21.

[141] Matplotlib Development Team, "Matplotlib documentation," 2024. Accessed: 2026-05-21.

[142] Built In, "Scikit-learn guide," 2024. Accessed: 2026-05-21.

[143] Scaler Topics, "What is scikit-learn in python?," 2024. Accessed: 2026-05-21.

[144] J.-B. Lamy, "Owlready2 documentation," 2024. Accessed: 2026-05-21.

[145] Stanford University, "Protégé ontology editor," 2024. Accessed: 2026-05-21.

[146] RDFLib Developers, "Rdflib," 2024. Accessed: 2026-05-21.

[147] KararHaiderDeepVision, "Vgg16: The timeless deep learning architecture that redefined image classification," 2025. Accessed: 2026-05-21.

[148] GeeksforGeeks, "Vgg-16 cnn model," 2024. Accessed: 2026-05-21.

وزارة التعليم العالي والبحث العلمي
Ministère de l'enseignement supérieur et de la recherche scientifique

المركز الجامعي إيليزي
University Center of Illizi

حاضنة الأعمال
Tassilix Innovation Hub Business Incubator

OntoBrain DX

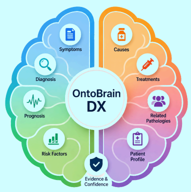

Detect. Explain. Care.

14/06/2026

Abstract:

Our company develops intelligent healthcare solutions for brain tumor diagnosis. The platform combines Artificial Intelligence and semantic medical knowledge to analyze MRI images and provide explainable diagnostic results, helping clinicians make more reliable decisions while improving healthcare quality and supporting digital transformation in the medical sector.

Student Team:

- **Full Name: Fatma zohra BOUSSOUFA.**
- **Institution: Institute of Sciences.**
- **Departement: Departement of Computer Science.**
- **Academic Level: Master 2.**
- **Field of Study: Artificial Intelligence and Its Applications.**
- **Project Role: System Designer and Developer.**
- **Phone Number: 0659999071.**
- **E-mail: fatmazohra.boussoufa@cuillizi.dz**

Supervisory Team:

- **Full Name: Youcef BEZZA.**
- **Institution: Institute of Computer Science.**
- **Academic Rank: Ph.D in Computer Science.**
- **Field of Expertise: AI, Optimization, Cloud Computing.**
- **Phone Number: 0697094583.**
- **Email Address: youcef.bezza@cuillizi.dz**

- **Full Name: Fatima zahra AMARA.**
- **Institution: /**
- **Academic Rank: Ph.D in Computer Science.**
- **Field of Expertise: Knowledge Graph, Ontology, Semantic AI.**
- **Phone Number: 0667494924.**
- **Email Address: amafatima.zahra@gmail.com**

Problematic:

01

Fragmented Diagnosis

No unified tool links MRI detection with clinical context, grade, and treatment.

02

Black-Box AI

Existing CNN systems give a label — no explanation, no clinical reasoning.

03

No Safety Net

Borderline predictions are silently accepted with no conflict detection mechanism.

~330,000

new brain tumour cases diagnosed globally per year

In Algeria: 500,000+ people affected by neurological conditions; radiology AI tools remain rare and unexplainable.

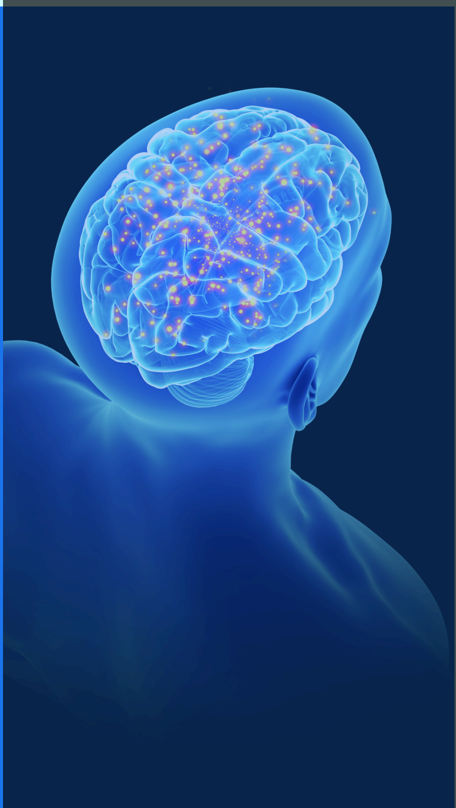

Proposed Solution

VGG16 CNN

_Fine-tuned deep learning model trained on 10,800 MRI images.
_Binary detection: Tumour / No Tumour.
~98–99% accuracy.

OWL Ontology

_Formal knowledge base: 20 classes, 100+ individuals.
_SNOMED CT, ICD-O-3, MeSH coded.
_WHO Grade I–IV taxonomy.

SWRL Reasoning

_77 rules across 10 groups (A–J).
_Infers grade, treatment, alerts, confidence & CNN-ontology agreement

How It Works:

01 MRI Input
128×128 px
Normalized

→

02 VGG16 CNN
Sigmoid output
YES / NO

→

03 SWRL Reasoner
77 rules
10 rule groups

→

04 Ontology Query
Owlready2
RDFLib SPARQL

→

05 Clinical Report
Classification,
symptom, Alert
Treatment

Output: **isMalignant, malignancyGrade, recommendedTreatment, treatmentUrgency, clinicalAlert, overallConfidence, cnnOntologyAgreement.**

Technical Architecture CNN (VGG16):

Transfer learning + selective fine-tuning on 10,800 MRI images

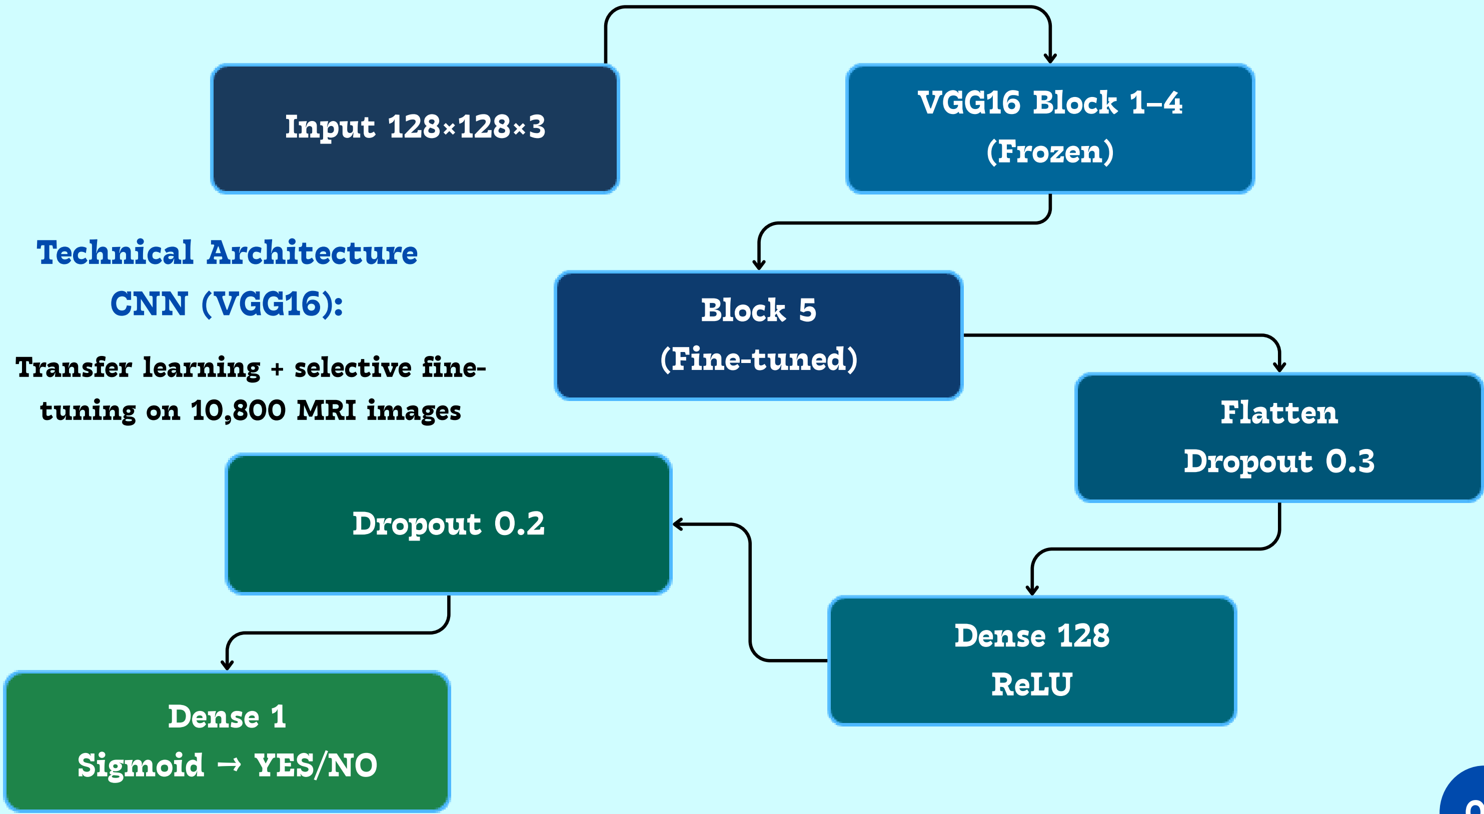

Technical Architecture (01)

OWL Ontology + SWRL:

Formal knowledge representation and automated clinical reasoning

20 OWL Classes

10
Object Properties

100+
Named Individuals

3
Coding Standards

77
SWRL Rules

40+
Datatype Properties

Interoperability: SNOMED CT (has_SNOMED_CT_ConceptId), ICD-O-3 (has_ICDO-3_Code), MeSH (hasMeshID) → EHR-ready

ProtoType:

Value Proposition(القيمة المضافة):

What makes us unique in the Algerian and global landscape

Only Algerian System

First to combine CNN + OWL ontology for brain tumour MRI in Algeria.

Clinical Safety Net

CONFLICT_REVIEW_REQUIRED flag escalates uncertain predictions automatically.

Complete Clinical Report

Grade · Treatment · Urgency · Alert · Confidence — not just YES/NO.

International Standards

SNOMED CT + ICD-O-3 + MeSH: ready for EHR integration worldwide.

Traceable Reasoning

Every inference traces to a named SWRL rule group — fully auditable.

Scientifically Validated

67 rules across 10 groups validated on 1,080 test MRI images.

Competitive Analysis(التحليل التنافسي):

Bargaining Power of Buyers: High

- Hospitals require proven reliability and clinical value.
- Buyers can compare competing solutions.

Threat of New Entrants: Medium

- High technical and clinical expertise required.
- Validation and healthcare partnerships create barriers.

Competitive Rivalry: Medium

- Growing AI healthcare market.
- Few competitors offer explainable AI with ontology reasoning.

Threat of Substitutes: Medium

- Traditional radiology diagnosis.
- Existing AI-based medical imaging solutions.

Bargaining Power of Suppliers: Low–Medium

- Open-source technologies reduce dependency.
- Medical data and clinical experts remain important suppliers.

The market is attractive and growing. OntoScan AI benefits from a strong competitive advantage through the combination of Deep Learning, Medical Ontologies, and Explainable AI.

Target Market(الفئة المستهدفة):

- 🏥 **Hospitals: ~865**
- 🧠 **Radiology Centers: ~300**
- 🎓 **Universities: ~26**
- 🏛️ **Health Authorities: 16**

Total Target Market in Algeria: ≈ 1,207 institutions.

Business Model Canvas(النموذج التجاري):

Key Partners
- Hospitals
- Universities
- Cloud Providers

Key Activities
- CNN Training
- Ontology Engineering
- Validation

Key Resources
- MRI Dataset
- VGG16
- Knowledge Graph

Value Proposition
- Explainable AI
- Brain Tumor Detection
- Knowledge Graph Reasoning

Customer Relationships
- Training
- Support
- Research Collaboration

Channels
- Web Platform
- Hospital Systems
- Conferences

Customer Segments
- Hospitals
- Radiology Clinics
- Universities
- Health Authorities

Cost Structure
- Cloud Computing
- Development
- Data Annotation

Revenue Streams
- Licenses
- Subscriptions
- Training

Go-to-Market Strategy:

Pricing Strategy

- Hospital Monthly Subscription
- Per MRI Analysis
- Annual Clinic License
- Custom Deployment
- Technical Support Contracts

Distribution Strategy

- Web platform — accessible from any hospital workstation
- Direct integration with existing HIS/RIS systems via API
- Partnerships with Algerian health authority (MSPRH)

Promotion Strategy

- Presentations at medical AI conferences (Algeria, Tunisia)
- Academic publications in peer-reviewed journals
- Demonstrations at TassiliX Innovation Hub events
- Social media: LinkedIn for B2B; Facebook for awareness

Project Startup Investment Costs (التّكاليف الإستثماريّة):

Category	Description	Cost(DZD)
Hardware	High-performance workstation (GPU, RAM,SSD)	450,000
Development Software	Development and testingtools	50,000
Cloud Hosting	Cloud infrastructure andstorage for one year	120,000
Domain and Website	Domain name, website development, and SSL certificate	25,000
Data Acquisition	Dataset preparation and preprocessing	80,000
Ontology Development	Medical ontology engineering and validation	100,000
AI Model Development	CNN training, testing, and optimization	150,000
Marketing and Branding	Visual identity, brochures, and social media campaigns	100,000
Legal and Administrative Costs	Startup registration and legal documentation	70,000
Travel and Meetings	Hospital visits, demonstrations, and networking	80,000
Contingency Reserve (10%)	Unexpected expenses	122,500
Total Estimated	1,347,500	

Annual Operating Costs (التكاليف التشغيلية السنوية):

Category	Cost(DZD)
Cloud Infrastructure	600,000
Maintenance and Updates	500,000
Technical Support	400,000
Marketing	300,000
Software Licenses	200,000
Administrative Expenses	180,000
Other Expenses	350,000
Total	2,530,000

Lean MVP Cost Estimation (تقدير تكلفة النموذج الأولي الأدنى القابل للتطبيق):

Item	Cost(DZD)
Domain and Website	25,000
Cloud Services	60,000
Marketing Activities	50,000
Demonstrations and Meetings	40,000
Miscellaneous Expenses	25,000
Total Lean MVP Cost	200,000

Revenue Streams (مصادر الإيرادات):

Revenue Source	Estimated Price
Hospital Subscription	50,000 – 150,000 DZD/month
MRI Analysis Service	1000 – 2500 DZD per Scan
Annual Enterprise License	1,500,000 DZD/year
Custom Deployment	Starting from 1,000,000+ DZD

Strategic Roadmap:

Month 1

Launch
الإطلاق

- Deploy binary detection MVP.
- Integrate ontology v1.0.
- Pilot hospital partnership.
- Publish research paper.
- ISO 13485 medical certification.
- INAPI patent filed.

Month 2

Scale
التوسع

- Multi-class tumour detection.
- EHR system API integration.
- 10 hospital contracts.
- Algerian health authority approval.

Month 3

Expand
التوسع إلى أسواق جديدة

- Extend to other cancers.
- Mobile app for radiologists.
- Regional expansion (Tunisia).
- 50+ institutional clients.

Month 4

Stabilize
الاستقرار

- Ontology v3.0 with patient data.
- AI-assisted biopsy prioritization.
- Revenue >50M DZD.

Month 5

Lead
الريادة

- National AI health platform.
- 100+ hospitals served.
- International partnerships.

Thank you for your attention.