CS 491 - Senior Design Project I

Project Specification Document

T2515 - PatchMatch

Team Members:

Bertan Uran - 22102541

Ekin Köylü - 22103867

Elif Lara Oğuzhan - 22203138

Emre Yazıcıoğlu - 22201668

İlke Latifoğlu - 22203818

Table of Contents

1. Introduction

1.1 Description

Our project develops an intelligent whole-slide image retrieval system that empowers pathologists to quickly find visually and clinically similar cases from large-scale digital pathology databases. By leveraging state-of-the-art computer vision and scalable search methods, our system transforms high-resolution pathology slides into efficient, comparable representations, enabling instant case-based retrieval and visualization. This innovation supports faster, more consistent, and explainable diagnoses while reducing workload for specialists, enhancing research reproducibility, and improving overall diagnostic confidence in medical practice.

1.2 High Level System Architecture & Components of Proposed Solution

The architecture that will be explained can be seen in Figure 1. The Presentation Layer contains all user interfaces and is responsible for how users interact with the system. The “Viewer Interface” allows users to visualize whole-slide images (WSIs), zoom in/out, and access generated heatmaps. The “Navigation Interface” lets users search between retrieved images. The “Settings Interface” allows users to change configuration such as theme or display preferences.The “Sharing Interface” lets users share the retrieved results and annotations with other users. The “Annotation Interface” provides marking regions of interests (ROIs), adding comments, and seeing the saved annotations. Hence, these components ensure that the system is clinically aligned and efficient for everyday use, and it passes all functional requests to the Application Layer.

The Application Layer implements the core logic and processing capabilities of the system. The “Retrieval” service performs image-to-image, text-to-image, and text-image pair searches by querying stored embeddings. The “Preprocessing” service standardizes the input by generating tissue masks, applying stain normalization and removing background. The “Patch Extraction” subsystem divides whole slide images into smaller, manageable patches which are used for embedding generation. The “Viewer Backend” component streams multi-resolution tiles to the Presentation Layer to enable smooth zoom operations. Therefore, the “User Management” component handles authentication, user profiles, access permissions, and retrieval history. Together, this layer constitutes the logic of the system, by coordinating the flow of data between the Presentation Layer and the Database Layer.

The Database Layer stores all information required by the platform. The “WSI & Report Storage” component holds the original whole-slide images and their associated pathology reports. The “User/Session Storage” maintains user accounts, settings, annotation data, and interaction history. The “Embedding Storage” component keeps all generated feature vectors as slide-level, patch-level, and text embeddings that are used during retrieval operations. This layer ensures data integrity, efficient access, and scalability, providing the Application Layer with quick retrieval of stored information while securing sensitive medical data.

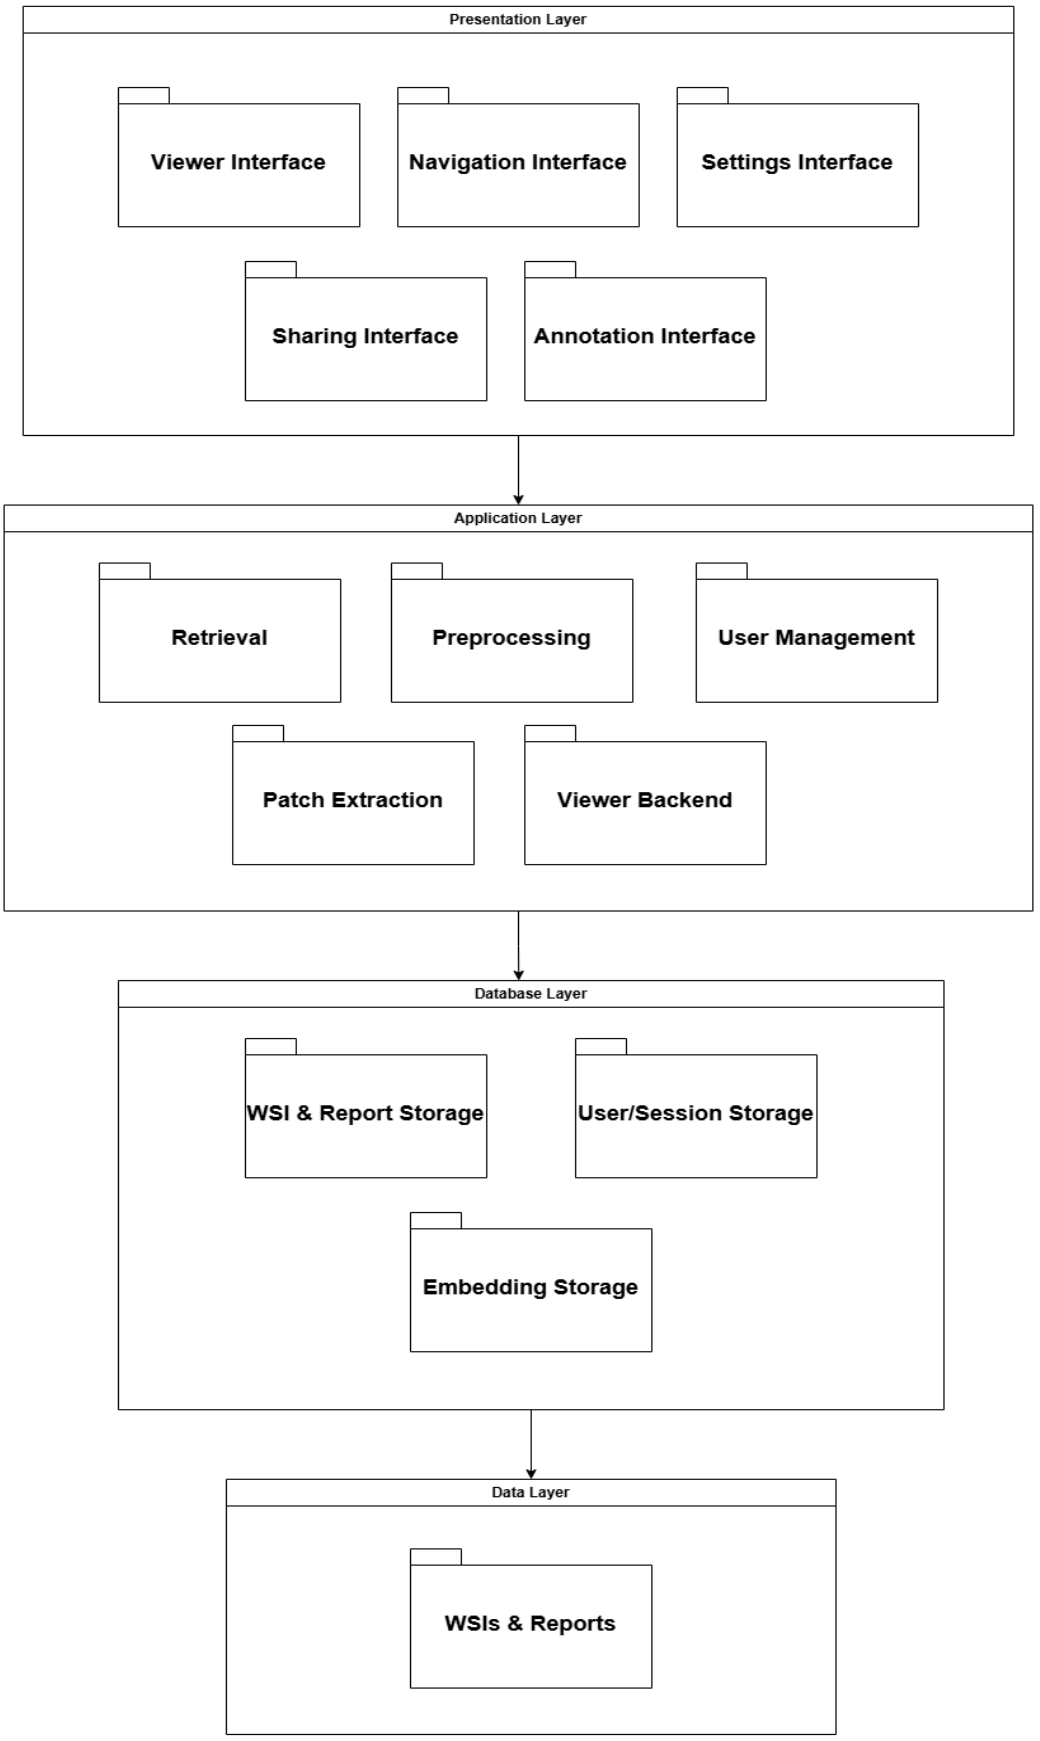

Figure-1 High Level System Diagram

1.3 Constraints

1.3.1 Implementation Constraints

- PatchMatch handles whole-slide images (WSIs) that contain 50,000 x 50,000 pixels. Regular computers cannot load these images into memory. The system partitions them into smaller tiles and supports multi-resolution viewing.
- The AI models that generate embeddings require GPUs to run at practical speeds. Processing thousands of images takes computing power and time. The search database must quickly retrieve similar images when indexing millions of patches from slides.
- Medical scanners save images in proprietary formats (Aperio .svs, Hamamatsu .ndpi). The system requires libraries like OpenSlide to read them. For hospital integration, PatchMatch should support healthcare data exchange standards like DICOM [2] and HL7 [3], though digital pathology hasn't adopted these yet.

1.3.2 Economic Constraints

- Running PatchMatch requires servers with GPUs, storage for images and databases, and reliable infrastructure. Costs scale with the number of images and users.
- Development costs include acquiring training data, potentially licensing AI models, and building a medical-grade system. Hospitals adopting the system may need to upgrade their scanners and IT infrastructure.
- This is an academic project with limited access to commercial resources and datasets. We'll use publicly available datasets, which may not represent the diversity of real-world medical cases.

1.3.3 Ethical Constraints

- Medical images contain patient information. All images must be anonymized before use. The system must comply with KVKK (Law No. 6698) [4] in Turkey and GDPR (Regulation EU 2016/679) [5] in the European Union.
- The system must avoid biases. If certain tissue types or patient populations are underrepresented in training data, performance will suffer for those cases. We must document dataset composition and acknowledge limitations.

- Users must understand that the system retrieves visually similar cases, not necessarily cases with identical diagnoses. The system should indicate confidence levels and explain why cases are considered similar.

1.4 Professional and Ethical Issues

PatchMatch operates in a medical setting where errors have serious consequences. Considering this, the system assists clinical decision-making but does not diagnose patients. Doctors remain responsible for diagnostic conclusions. The system presents information clearly without biasing clinical judgment. Showing AI suggestions prematurely could influence how clinicians evaluate a case.

Also, medical images require the same confidentiality as patient records. The system needs access controls, activity logs, and secure data transmission. Inter-hospital case sharing requires consent and data agreements that comply with GDPR [5] and KVKK [4].

Other than that, we will document how models were trained, what data was used, and where the system underperforms to be honest about limitations. If performance is weak for certain tissue types or diseases, users will be informed. Post-deployment monitoring will track performance and address issues.

While PatchMatch doesn't diagnose, retrieving misleading similar cases could misinform clinical decisions. Therefore, the system must undergo testing before clinical deployment and include mechanisms for users to report problems.

Moreover, digital pathology tools often remain available only to well-resourced hospitals. This project can't solve access inequality, but we use open standards and documentation to lower barriers for future adoption.

1.5 Standards

PatchMatch operates in a regulated medical environment. Designing for these standards now will make clinical deployment feasible later.

1.5.1 Software Development Standards

“IEC 62304” [6] provides guidance on how to develop, test, and maintain medical device software. It categorizes software into three safety classes (A, B, C) based on

potential harm from failures. PatchMatch supports diagnostic decisions, so we'll need Class B documentation, which tracks everything from initial requirements through testing and maintenance.“IEC 82304-1” [7] covers standalone health software. This standard addresses validation, usability testing, and release procedures for web-based systems.

1.5.2 Medical Imaging Standards

DICOM (Digital Imaging and Communications in Medicine) [2] is the standard for medical images in radiology. Digital pathology uses proprietary formats instead. Most scanners save images as Aperio .svs or Hamamatsu .ndpi. We'll use OpenSlide, an open-source library that reads these formats. “DICOM Supplements 122 and 145” [8] define standards for whole-slide pathology images. We should design PatchMatch to support DICOM eventually so hospitals can integrate our system with their PACS systems.

“HL7” [3] standards govern the exchange of clinical data, including patient information and pathology reports. This matters for hospital system integration.

1.5.3 Quality and Risk Management Standards

“ISO 13485” [9] defines the requirements for managing quality throughout the lifecycle of medical device companies. Following these principles (documentation, testing protocols, and change control) will make our work more rigorous. “ISO 14971” [10] requires identifying hazards, estimating risks, implementing controls, and assessing what remains. For PatchMatch, we need to consider misleading retrieval results, system downtime during clinical hours, and misinterpretation of explainability features. IEC 62304 requires ISO 14971 compliance.

1.5.4 Regulatory Requirements

PatchMatch qualifies as a medical device if it supports clinical diagnosis, which means it needs regulatory approval. In the European Union, “Medical Device Regulation (EU) 2017/745” [11] covers software that provides information used in diagnostic decisions. Turkey also follows the EU MDR. This means CE marking is required for both Turkey and the European Union. We'd fall into Class I or IIa under MDR Rule 11. Class I devices are self-certified. Class IIa devices need notified body review. The “EU Artificial Intelligence Act (Regulation EU 2024/1689)” [12] requires transparency and documentation for high-risk AI systems in medical use. In the United States, the FDA classifies digital pathology software as Class II medical devices. This requires “510(k) premarket notification” [13] showing the

device is substantially equivalent to one already approved. Medical device classification depends on marketing and use. A system marketed as a research or educational tool that doesn't affect clinical decisions might not qualify as a medical device. We should consult regulatory experts as the project develops.

1.5.5 Data Protection Requirements

“GDPR (Regulation EU 2016/679)” [5] and “KVKK (Law No. 6698)” [4] regulate medical image handling:

- Processing requires informed consent or research interest
- Security through encryption, access controls, and audit logs
- Patients can access, correct, delete, or port their data
- Research use requires anonymization without individual consent

GDPR/KVKK compliance is built into PatchMatch's architecture because the system stores and retrieves medical images.

1.5.6 Documentation and Modelling Standards

UML 2.5.1 [14] is used for modelling architecture and workflow diagrams. IEEE 830-1998 [15] guides the structure and clarity of the software requirements specification. For instance, the functional and non-functional requirements are separated, and written in a clear language.In addition, the non-functional requirements are written in a testable way as much as possible.

2. Design Requirements

2.1 Functional Requirements

Whole Slide Image Management

The system must allow authorized users to upload, store, and manage whole slide images (WSIs) and their metadata.

Viewer: Navigation, Zoom and Selection

The system must provide an interactive viewer that lets users open a WSI, zoom at multiple magnifications, and select one or more regions of interest (ROIs) for further analysis or retrieval.

Feature Extraction of Patches

The system must extract patches from images and compute feature embeddings using deep learning models.

Multi-Resolution Embeddings

The system must generate multi-resolution slide representations that capture both cellular and tissue-level information to support retrieval at different magnifications.

Indexing and Database

The system must index feature embeddings in a vector database to enable efficient similarity search over large image piles.

Image-to-Image Retrieval (Whole Slide)

Given a query WSI, the system must retrieve the top-k most visually similar WSIs from the database, and display them with thumbnails, key metadata, and similarity scores.

ROI-to-ROI Retrieval

Given a query ROI selected on a WSI, the system must retrieve visually similar ROIs or patches from other WSIs and present them to the user with their parent slide, location, and similarity scores.

Text-to-Image Retrieval

The system must allow users to enter text queries such as “dense nuclei”, “cribriform pattern”, “high-grade tumor”; and retrieve WSIs or ROIs whose visual content and associated reports match the semantic meaning of the query.

Image + Text Paired Retrieval

The system must allow users to specify a paired query consisting of a WSI/ROI and a short text description such as “focus on cell size”, and it must use both inputs jointly to retrieve results that are visually similar while satisfying the text-specified constraint.

Multimodal Fusion with Reports

The system must combine image embeddings with text embeddings from the reports so that similarity and retrieval can consider both visual features and semantic information.

Explainable Retrieval

For each retrieved result, the system must provide an explanation of similarity, such as:

- highlighting heatmaps on query and result images
- listing shared visual or semantic features such as similar nucleus size, pattern, or report keywords.

The explanation should address both image-based features and report-based information.

User Annotation

The system must allow users to annotate images. The users should be able to draw regions, write comments and the system must keep these annotations in the database where users can access later.

Data Sharing

The system must allow users to share their images, results and notes to other users.

2.2 Non-Functional Requirements

The non-functional requirements ensure that PatchMatch works well in the real world to assist pathologists. They help establish a system that is easy to use, reliable, fast, scalable, and supportable.

2.2.1 Usability

- PatchMatch should not disrupt the workflow of the pathologists, so it should be easy to learn and integrate into the existing workflow in under 30 minutes. The pathologists should be able to utilize the application following the instructions and tips provided by the application without formal training.
- The terminology given in the graphical user interface should align with the pathological terms.
- The graphical user interface should be easy to interact with, the user should be able to find similar images in no more than 5 interactions.
- The application should provide feedback after user operations within a second to indicate success or failure of operations.
- The retrieved images should be displayed in a visually easy to interpret and explainable manner.
- The images should be examined by zooming in and out just like the real-world procedure.

2.2.2 Reliability

- The system should be available during the clinical hours more than 99% percent of the time. This way, healthcare services will not halt.
- If an error happens and the system shuts down, it should save the progress up to the point of the error. The progress may consist of uploading the data, annotating text on the data, selecting the region of interest.
- In the case of errors the user should see an informative message.
- The data of WSIs should be preserved in a non-corrupted way. If data corruption happens, the administrator and the uploader of the file should be notified.
- In case of a system failure, the system should be operable back in under six hours.

2.2.3 Performance

- The initial view of a WSI should be available upon selection within 5 seconds for 95% of the requests.
- Zooming in or out should provide the new image view within 0.5 seconds on average.
- For a selected region of interest, the most similar top-K images should be displayed under 7 seconds for 90% of the queries.
- For a text query, the most associated top-K images should be displayed under 7 seconds for 90% of the queries.

- The system should encode the tiles in a fast manner to create the embeddings and make them accessible.

2.2.4 Supportability

- The components like the retrieval engine, embedding pipeline should be modular to be updated or repaired without affecting the other components.
- The code should be documented, and it should be tracked using a version control system.
- The system should be monitored through logging mechanisms to identify the issues.
- Operational parameters such as the “K” in the top-K most similar image retrievals should be adjustable.
- Users should access support resources such as FAQs and troubleshooting guides.
- Different medical file formats should be handled ensuring compatibility.

2.2.5 Scalability

- The retrieval engine should search images from a few thousands of WSIs to a much larger dataset without compromising latency.
- Multi-resolution embeddings should be stored using compression or vector representations to decrease the amount of required storage space. This way, more embeddings can be stored with an increased amount of data.
- Horizontal scaling can be considered to increase the overall capacity without redesigning the system.
- Different AI models should be added and used without breaking the working logic of the previous models or requiring to reprocess all existing WSIs.

2.2.6 Security

- The system should require authenticated access to see patient data.
- The WSIs, embeddings, and associated reports should be stored in an encrypted manner.
- Professionals should be provided access based on which data they can see. The scope of the available data should be determined based on the user.
- The communication through the network should be done in an encrypted way using a protocol like TLS.

3. Feasibility Discussions

3.1 Market & Competitive Analysis

3.1.1 Market Overview

Digital pathology has been growing significantly as most of the hospitals have been digitizing their workflow. Many market research illustrates that the global digital pathology market size is expected to reach $2-2.8 millions dollar by 2030, increasing at a compounded annual growth rate (CAGR) of around 10% from today [16][17][18]. Moreover, AI in the pathology market is growing much faster with more than a quarter percent CAGR expectancy between 2025 and 2030 [19][20]. This demonstrates that there is a demand for tools that will support the domain in ways such as facilitating diagnosis process, improving efficiency and reducing the workload. Hence, healthcare systems are ready to adopt intelligent tools to enhance the pathology field in this environment.

3.1.2 Target Customer Segments

PatchMatch serves three main groups: hospitals and pathologists, educational institutions, and research organizations.

3.1.2.1 Hospitals and Pathologists

In hospitals, pathologists' daily workload involves reviewing large numbers of biopsy samples, often under time pressure [21]. They need quick and reliable ways to compare current cases with previous ones to look for similar patterns and to support their diagnostic decisions. PatchMatch helps them work more efficiently by enabling fast retrieval of visually similar slides or regions of interest, reducing manual search time and improving confidence during case review.

3.1.2.2 Educational Institutions (Medical Schools & Training Programs)

Wide exposure to cases is a significant part of medical student and pathology trainee education, but accessing diverse examples can be highly limited and time-consuming. PatchMatch enables instructors and learners to quickly identify relevant slides and find similar cases based on visual features or keywords. This makes the learning process interactive and enhances a student's understanding of how a particular condition looks under the microscope.

3.1.2.3 Research Institutions (Academic and Clinical Research Centers)

Researchers dealing with large collections of pathology slides often have to identify specific cases with predefined features or analyze variations from sample to sample. PatchMatch simplifies this process by enabling researchers to search large image archives efficiently and allows discovery of patterns that might not be obvious through manual review. Therefore, it speeds up scientific research in pathology and related fields.

3.1.3 Competitors & Differentiation

The digital pathology ecosystem includes enterprise platforms, research prototypes, and AI assistance tools, but none of them fully address the need for fast, explainable, and flexible similarity retrieval. Enterprise systems, such as Sectra [22] and Iron Mountain [23], focus on storage, viewing, and workflow integration but do not offer advanced image-based search. Research-oriented tools like Lagotto [24], SMILY [25], Yottixel [26], and FastPathology [27] present isolated aspects of retrieval or model deployment, however, they fall short of clinical readiness, region-level search, explainability, or multi-modality. Solutions like Ibex [28] specialize in task-specific detection rather than case similarity. Our system combines the state-of-the-art in whole-slide and ROI-level retrieval, natural-language search, explainability, multi-resolution embeddings of HIPT, personalized sessions, and multimodal integration-appropriately positioning itself as the uniquely comprehensive solution designed for clinical, educational, and research workflows.

3.1.4 Barriers to Entry & Risks

3.1.4.1 Data Access and Privacy

Acquiring large datasets of WSIs of various types of biopsies and their associated diagnostic reports involves complex privacy and access challenges. For example, in Türkiye, KVKK [4] regulates the collection, storage, and processing of personal data, which includes WSIs and reports as well. In order to ensure compliance with KVKK and equivalent regulations in other jurisdictions such as GDPR [5] for the European Union, we should negotiate data use agreements, provide anonymization of the information of patients, manage consent processes, and implement data governance [21]. These steps slow down dataset acquisition processes.

3.1.4.2 Regulatory Requirements for Clinical Deployment

Because PatchMatch is intended to support clinical diagnosis, regulatory oversight becomes a substantial barrier to entry. In the United States, digital pathology tools are

treated as Class II medical devices by the FDA in general and may be subject to a 510(k) submission supported by technical performance data and clinical validation [13]. Similar software may be subject to the MDR [11] in the European Union, which requires demonstrating safety, performance, risk management, and post-market monitoring. Türkiye follows the same medical device rules as the EU MDR. This means that any software used to support diagnosis must go through a formal approval process and receive CE marking before it can be used in clinics. These steps require extra testing, documentation, and quality-management activities, which can slow down development and make the regulatory process more complex.

3.1.4.3 Heterogeneity of Slides

WSIs vary significantly in scan resolution, stain type, scanner model, lab protocols, and patient population. Performing accurate similarity checking and retrieving images using these different types of WSIs is technically challenging. If not addressed, it can lead to inconsistent search results, reduced trust from clinicians, and limited adoption across different hospitals or datasets.

3.1.4.4 Integration

Hospitals operate in tight and highly regulated IT environments. Connecting external software to hospital networks requires meeting strict cybersecurity, data-access, and firewall requirements. Hospitals may need to perform security reviews, whitelist domains or IPs, and assess data-flow risks before allowing staff to use an external platform. These steps, combined with slow procurement processes and legacy IT infrastructure, can significantly extend the time required for deployment and adoption.

3.2 Academic Analysis

In designing the WSI data retrieval system, we conducted research into digital pathology and WSI analysis. We analyzed algorithmic indexing methods and deep learning models as CBIR methods. Finally, we examined models that provide users with necessary explanations about similar images, and recent advances in multimodal AI that combine the image and text.

3.2.1 WSI Acquisition and Preprocessing

Improvements in digital pathology have replaced traditional microscope studies with the digital field. This transition allows WSI data to be stored and analyzed at multiple

resolutions. However, the large size of WSIs makes it impossible to store or preprocess them directly in GPU memory. Additionally, it is not practical to annotate such large images at the pixel level. Therefore, academic studies have focused on partitioning WSIs into smaller segments and focusing on more informative regions for storage and analysis. Campanella et al. [29] introduced weakly supervised deep learning on WSI. This technique divides WSIs into smaller non-overlapping patches and manually annotates each WSI as cancerous or non-cancerous, achieving clinical-level diagnostic accuracy. Koohbanani et al. [30] introduced a tissue region extraction module to exclude back-ground, blurred, and non-tissue regions. This improves the representation quality of the WSIs and prevents the wasting capacity of the GPU. Complementary to all preprocessing methods, Tellez et al. [31] demonstrated that data augmentation techniques increase the generalization ability of the model to differences in color, orientation and brightness by quantitatively analyzing the impact of different augmentation strategies on WSI performance.

3.2.2 Feature Extraction and Representation

After preprocessing of the WSI patches, the next step is to extract feature representations. These representations convert each patch into a numerical embedding that enables similarity computation. Since the aim of our project is image retrieval rather than classification, using models that produce semantically meaningful embeddings helps better represent the relationship between similar images [32]. He et al. [33], proposed ResNet which is widely used in histopathological patches among CNN architectures. Residual connections of ResNet enables very deep architectures to learn hierarchical representations without suffering from vanishing gradient problems. Due to the large size of WSI images, it is difficult for a model to learn both cell- and tissue-level information simultaneously. Chen et al. [34] proposed HIPT, a transformer-based model that is effective for multi-resolution feature extraction. In addition to multi-resolution representation of HIPT, Lu et al. [35] introduced CLAM which determines the weights of the patches using the attention mechanism. In this way, it is possible to select important regions according to diagnostic relevance.

3.2.3 Similarity Computation in WSI Retrieval

After extracting the feature embeddings for each patch, the next step is computing image similarity. Two main approaches exist for measuring similarity: traditional distance-based metrics and learned similarity methods. Babenko et al [36]. introduced a technique for using deep convolutional neural networks for image retrieval without training. In this work, the CNN feature vectors are extracted from the intermediate fully connected layers of a pretrained AlexNet model [37], which is why they capture high-level semantic

features. When cosine or Euclidean distance is applied to these semantic features, high-quality similarity can be captured. Kalra et al. [38] proposed Yottixel, which is a search engine designed for content-based retrieval of digital pathology slides. In this work, the feature vector of each patch is extracted using DenseNet [39], and then it is converted into a binary barcode using MinMax binarization. All barcodes from one slide are grouped to form its BoB index. These indices are then compared using Hamming distance, which provides fast and scalable search across extensive WSI archives.

In contrast to distance-based methods, learning-based similarity approaches train models that directly learn semantic relationships between image embeddings. Chen et al. [40] introduce SimCLR, which is a self-supervised contrastive learning framework. In this work, two types of pairs are used during training. The first one is a positive pair, which consists of two augmented views of the same image, while all other images in the batch construct the negative pairs. The contrastive loss (InfoNCE) is employed to train the model, encouraging positive pairs to be pulled closer together and negative pairs to be pushed farther apart. As a result, patches with similar visual characteristics tend to cluster closely together in the embedding space.

3.2.4 Explainability in Digital Pathology

Pathologists need to understand the primary factors contributing to image similarity. That's why explainability of the system is essential in digital pathology systems. In our system, two complementary forms of explainability are considered. The first one is patch-level explainability, which shows which pixel regions inside a patch the model used to compute its embedding. Grad-CAM [41], introduced by Selvaraju et al., generates a heatmap over the patches and categorizes them based on their importance. This heatmap highlights the diagnostically relevant structures, and users can understand the reason for the similarity. The second one is slide-level explainability which aims to show which region across the entire WSI is mostly affecting the retrieval process. Lu et al., introduces an attention-based mechanism CLAM. It produces heatmaps over the entire WSI by computing weighted similarity scores through the self-attention layers.

3.2.5 Multimodal Integration

In digital pathology, pathologists also rely on the textual content of WSI reports to support their diagnostic reasoning. Radford et al. [42] introduced CLIP as a contrastive learning model. The CLIP is trained on image-text pairs and utilizes the InfoNCE loss, where the model aims to increase the similarity of matching pairs and decrease the similarity of mismatched pairs. In this way, it enables text-guided retrieval. Wang et al. [43] introduced

PLIP, which extends this framework to histopathology. It uses pathology image-text pairs during training which include pathology reports, diagnostic labels, region-level annotations.

4. Glossary

MDR: Medical Device Regulations
FDA: Food and Drug Administration
GDPR: General Data Protection Regulation
KVKK: Kişisel Verileri Koruma Kanunu
ISO: International Organization for Standardization
IEC: International Electrotechnical Commission
CBIR: Content-based Image Retrieval
WSI: Whole Slide Image
ROI: Region of Interest

5. References

[1] draw.io (diagrams.net), “Online diagram editor and flowchart tool,” 2025. [Online]. Available: https://www.draw.io/

[2] DICOM Standards Committee, “Digital Imaging and Communications in Medicine (DICOM),” NEMA. [Online]. Available: https://www.dicomstandard.org/

[3] Health Level Seven International, “HL7 Standards.” [Online]. Available: https://www.hl7.org/

[4] Kişisel Verileri Koruma Kanunu (KVKK), Law No. 6698, 2016. [Online]. Available: https://www.kvkk.gov.tr/

[5] European Parliament and Council, *Regulation (EU) 2016/679 (General Data Protection Regulation – GDPR)*, 2016.

[6] IEC 62304:2006/AMD1:2015, *Medical device software — Software life cycle processes*, IEC, 2015.

[7] IEC 82304-1:2016, *Health software — Part 1: General requirements for product safety*, IEC, 2016.

[8] DICOM Standards Committee, *DICOM Supplement 145: Whole Slide Microscopic Image IOD and SOP Classes*, 2010.

[9] ISO 13485:2016, *Medical devices — Quality management systems — Requirements for regulatory purposes*, ISO, 2016.

[10] ISO 14971:2019, *Medical devices — Application of risk management to medical devices*, ISO, 2019.

[11] European Parliament and Council, *Regulation (EU) 2017/745 (Medical Device Regulation – MDR)*, 2017.

[12] European Parliament and Council, *Regulation (EU) 2024/1689 (Artificial Intelligence Act)*, 2024.

[13] U.S. Food and Drug Administration (FDA), “Premarket Notification 510(k).” [Online]. Available: https://www.fda.gov/medical-devices/premarket-submissions

[14] Object Management Group, *OMG Unified Modeling Language (UML), Version 2.5.1*. OMG, Dec. 2017. [Online]. Available: https://www.omg.org/spec/UML/2.5.1/

[15] IEEE, *IEEE Recommended Practice for Software Requirements Specifications (IEEE Std 830-1998)*. IEEE Computer Society, 1998. [Online]. Available: https://standards.ieee.org/standard/830-1998.html

[16] Mordor Intelligence, "Digital Pathology Market Size & Share Analysis – Growth Trends and Forecast (2025–2030)", Mordor Intelligence, 2025. [Online]. Available: https://www.mordorintelligence.com/industry-reports/digital-pathology-market. [Accessed: 27-Nov-2025].

[17] MarketsandMarkets, "Digital Pathology Market by Product (Scanner, Software, Storage System), Type (Human, Veterinary), Application (Teleconsultation, Training, Disease Diagnosis, Drug Discovery), End User (Pharma & Biotech, Academia, Hospitals) – Global Forecast to 2030", MarketsandMarkets, Mar. 2025. [Online]. Available: https://www.marketsandmarkets.com/PressReleases/pathology.asp. [Accessed: 27-Nov-2025].

[18] Grand View Research, "Digital Pathology Systems Market Size, Share & Industry Analysis – Global Forecast to 2033", Grand View Research, 2024. [Online]. Available: https://www.grandviewresearch.com/industry-analysis/digital-pathology-systems-market. [Accessed: 27-Nov-2025].

[19] Grand View Research, "AI In Pathology Market (2025–2033) Size, Share & Trends Analysis Report by Component, Technology, Application, End Use, and Region," Grand View Research, Report ID: GVR-4-68040-758-8, Sept. 2025. [Online]. Available: https://www.grandviewresearch.com/industry-analysis/ai-pathology-market-report.

[20] MarketsandMarkets, "AI in Pathology Market by Neural Network (GAN, CNN, RNN), Function (Diagnostic, Image Analysis, CDSS, Data Management, Analytics), Use Case (Drug Discovery, Clinical Workflow), End User (Hospitals, Labs, Pharma/Biotech), & Region – Global Forecast to 2030," MarketsandMarkets, May 2025. [Online]. Available: https://www.marketsandmarkets.com/Market-Reports/ai-in-pathology-market-86647266.html.

[21] E. Özoğul, [Pathologist], personal interview, 08/11/2025.

[22] Sectra, "Sectra Digital Pathology Solution," Sectra Medical Systems. [Online]. Available: https://medical.sectra.com/product/sectra-digital-pathology-solution/

[23] Iron Mountain, “Digital Pathology Solutions,” Iron Mountain. [Online]. Available: https://www.ironmountain.com/industries/healthcare-services/pathology-solutions

[24] Huron Digital Pathology, “Lagotto — content-based image retrieval solution for pathology,” Huron Digital Pathology. [Online]. Available: https://www.hurondigitalpathology.com/lagotto/

[25] N. Hegde *et al.*, “Similar image search for histopathology: SMILY,” *NPJ Digital Medicine*, vol. 2, no. 1, p. 56, 2019.

[26] S. Kalra, H. R. Tizhoosh, C. Choi, S. Shah, P. Diamandis, C. J. V. Campbell, and L. Pantanowitz, “Yottixel – An Image Search Engine for Large Archives of Histopathology Whole Slide Images,” *Medical Image Analysis*, vol. 65, p. 101757, 2020. doi: 10.1016/j.media.2020.101757.

[27] A. Pedersen, M. Valla, A. M. Bofin, J. P. de Frutos, I. Reinertsen, and E. Smistad, “FastPathology: An open-source platform for deep learning-based research and decision support in digital pathology,” *arXiv preprint* arXiv:2011.06033, 2020.

[28] Ibex Medical Analytics (Ibex), “Home – Ibex Medical Analytics,” Ibex AI. [Online]. Available: https://ibex-ai.com/

[29] G. Campanella *et al.*, “Clinical-grade computational pathology using weakly supervised deep learning on whole slide images,” *Nature Medicine*, vol. 25, pp. 1301–1309, Jul. 2019.

[30] N. A. Koohbanani, B. Unnikrishnan, S. A. Khurram, P. Krishnaswamy, and N. Rajpoot, “Self-Path: Self-Supervision for Classification of Pathology Images With Limited Annotations,” *IEEE Transactions on Medical Imaging*, vol. 40, no. 10, pp. 2845–2856, Oct. 2021, doi: 10.1109/TMI.2021.3056023.

[31] D. Tellez, G. Litjens, P. Bándi, W. Bulten, J.-M. Bokhorst, F. Ciompi, and J. van der Laak, “Quantifying the effects of data augmentation and stain color normalization in convolutional neural networks for computational pathology,” *Medical Image Analysis*, vol. 58, p. 101544, Dec. 2019, doi: 10.1016/j.media.2019.101544.

[32] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A Simple Framework for Contrastive Learning of Visual Representations,” *arXiv preprint* arXiv:2002.05709, 2020.

[33] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778, 2016.

[34] R. J. Chen, M. Ding, A. Lu, D. Williamson, and F. Mahmood, “Scaling vision transformers to gigapixel images via hierarchical self-supervised learning,” Advances in Neural Information Processing Systems (NeurIPS), 2022.

[35] M. E. H. Lu, D. J. Williamson, A. Bagci, and A. E. Shady, “Data-efficient and weakly supervised computational pathology on whole-slide images,” Nature Communications, vol. 12, no. 1, pp. 1–13, 2021.

[36] A. Babenko, A. Slesarev, A. Chigorin, and V. Lempitsky, “Neural Codes for Image Retrieval,” *arXiv preprint* arXiv:1404.1777, 2014.

[37] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, vol. 25, 2012, pp. 1097–1105.

[38] S. Kalra, H. R. Tizhoosh, C. Choi, S. Shah, P. Diamandis, C. J. V. Campbell, and L. Pantanowitz, “Yottixel – An Image Search Engine for Large Archives of Histopathology Whole Slide Images,” *Medical Image Analysis*, vol. 65, p. 101757, Oct. 2020, doi: 10.1016/j.media.2020.101757.

[39] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, “Densely Connected Convolutional Networks,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2261–2269.

[40] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A Simple Framework for Contrastive Learning of Visual Representations,” *arXiv preprint* arXiv:2002.05709, 2020.

[41] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization,” *International Journal of Computer Vision*, vol. 128, no. 2, pp. 336–359, 2019, doi: 10.1007/s11263-019-01228-7.

[42] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, “Learning Transferable Visual Models From Natural Language Supervision,” *arXiv preprint* arXiv:2103.00020, 2021.

[43] Z. Huang, F. Bianchi, M. Yuksekgonul, T. J. Montine, and J. Zou, “A visual–language foundation model for pathology image analysis using medical Twitter,” *Nature Medicine*, vol. 29, no. 9, pp. 2307–2316, Sep. 2023, doi: 10.1038/s41591-023-02504-3.