



## Short Communication

## PIRO: A web-based search platform for pathology reports, leveraging large language models to generate discrete searchable insights

Scott Robertson<sup>a,\*</sup>, Venkata Koppireddy<sup>a</sup>, Jeremy Cumbo<sup>a</sup>, Hooman Rashidi<sup>b</sup>, Samer Albahra<sup>a</sup><sup>a</sup> Center for Diagnostics and Artificial Intelligence, Cleveland Clinic Foundation, Department of Pathology and Laboratory Medicine, 9500 Euclid Ave, Cleveland, OH 44195, USA<sup>b</sup> Computational Pathology and AI Center of Excellence (CPACE), University of Pittsburgh School of Medicine, Department of Pathology, 200 Lothrop Street, Pittsburgh, PA 15261, USA

## ARTICLE INFO

## Keywords:

Search  
Pathology reports  
Full-text  
Unstructured text  
Large language model

## ABSTRACT

Pathologists rely on access to historical diagnostic case texts for research, education, and peer learning. However, many laboratory information systems (LIS), including Epic Beaker, lack optimized search tools tailored to pathology-specific text queries. To address this need, we developed PIRO (Pathology Information Retrieval Optimizer), a web-based platform enabling efficient text searches of diagnostic archives. Built using FastAPI, Angular, and Apache Solr, PIRO supports both basic and advanced search functionalities, faceted filtering, and data extraction, while ensuring compliance with institutional privacy protocols. PIRO's capabilities extend to case cohort building, search result export, and secure access control within the institutional network. In an 8-month study, we observed significantly higher PIRO adoption rates (67 %) among pathologists compared to Epic Beaker's SlicerDicer (9 %), underscoring PIRO's usability and relevance. Additionally, we implemented a large language model (LLM) to annotate reports with a "Malignancy Risk" label, enhancing search precision and enabling future expansion of automated annotations. Ongoing work focuses on integrating PIRO with our digital pathology platform, enabling direct access to digital slides from case results. PIRO's adaptable design makes it applicable across institutions, advancing search and retrieval efficiency in pathology archives and enhancing support for pathology research and education.

## 1. Introduction

Pathologists benefit from the ability to search their institution's archive of diagnostic case texts. This capability is essential for research, education, and other academic activities. In research, most pathology-based projects begin with querying diagnostic text to identify candidate cases. For educational purposes, the archive serves as a valuable resource, especially for resident and fellow trainees who can refer to past case texts as guides when documenting current cases. This resource also benefits staff pathologists, allowing them to engage to gain insights from colleagues' diagnostic reports. Reviewing past cases involving rare or unusual diagnostic entities is often helpful when drafting similar case texts.

Despite the clear need for text search, many laboratory information systems (LIS) provided by vendors are not optimized for this purpose. Our institution uses Epic Beaker LIS, whose built-in analytics tool, SlicerDicer, supports cross-domain analytics across various Epic data models, not limited to pathology. However, SlicerDicer is not tailored for pathology-specific text searches, and usage among our staff remains low.

To address this gap, we developed PIRO (Pathology Information Retrieval Optimizer), a web-based platform designed to enable rapid

searches of the diagnostic case archive. PIRO is tailored for non-technical users, allowing pathologists and trainees to use the platform effectively with minimal training. Furthermore, various safeguards ensure compliance with our institution's IRB, HIPAA, and legal requirements.

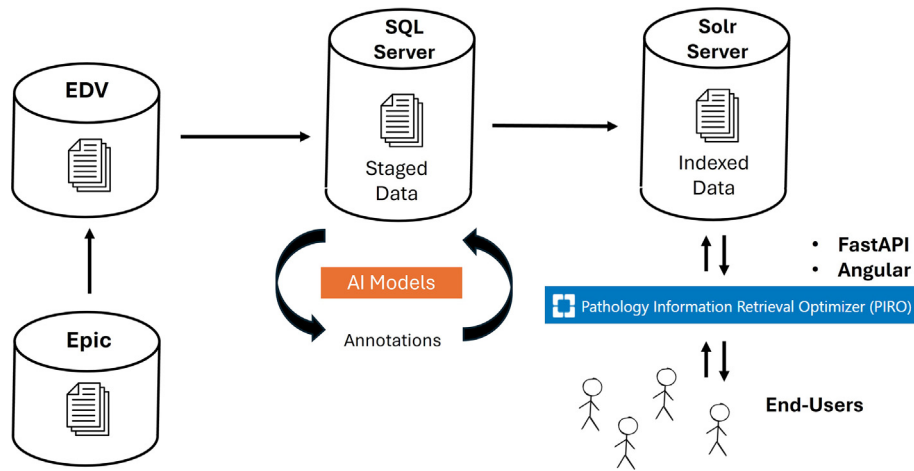
## 2. Implementation

## 2.1. Application stack

PIRO uses FastAPI (v0.104.1), a high-performance RESTful web framework that handles API requests, database connections, authentication, and server-side logic (Fig. 1). On the client side, PIRO utilizes the Angular framework (v15.2) to render HTML, process user inputs, manage UI components, and handle browser interactions. PIRO's search capabilities are powered by Apache Solr (v8.11.2), an open-source library that provides advanced full-text search. Solr indexes the case data, optimizing it for search, and supports features like faceted searching and search term highlighting, which are prominently featured in PIRO. The database for PIRO is built on Microsoft SQL Server 2019, hosted in our institution's data center.

\* Corresponding author.

E-mail address: [roberts10@ccf.org](mailto:roberts10@ccf.org) (S. Robertson).



**Fig. 1.** PIRO's Schema. Data originates in Epic, flows into the Enterprise Data Vault (EDV), and then into PIRO's staging database in SQL Server. The data is indexed and copied into PIRO's Solr server. Users interact with PIRO's web-based front end (FastAPI, Angular). Search queries are directed to the Solr server. Finally, AI models can be applied to staged data in the SQL Server database to generate discrete, searchable insights, which are saved back into the SQL Server database.

## 2.2. Data ingestion

PIRO is populated with data from two systems: Epic Beaker, our current LIS, and Cerner Copath Plus v2014, our previous LIS. Epic Beaker data is ingested into our Enterprise Data Vault (EDV) nightly. The EDV, a central repository for enterprise data from various operational systems, uses a schema similar to Clarity (Epic's reporting database), with most tables corresponding 1:1 to Clarity tables. We developed ETL scripts to extract data from the EDV, transform it into PIRO's simplified schema, and load it into PIRO's SQL Server. PIRO uses an LIS-agnostic schema that is less complex than Epic's Clarity.

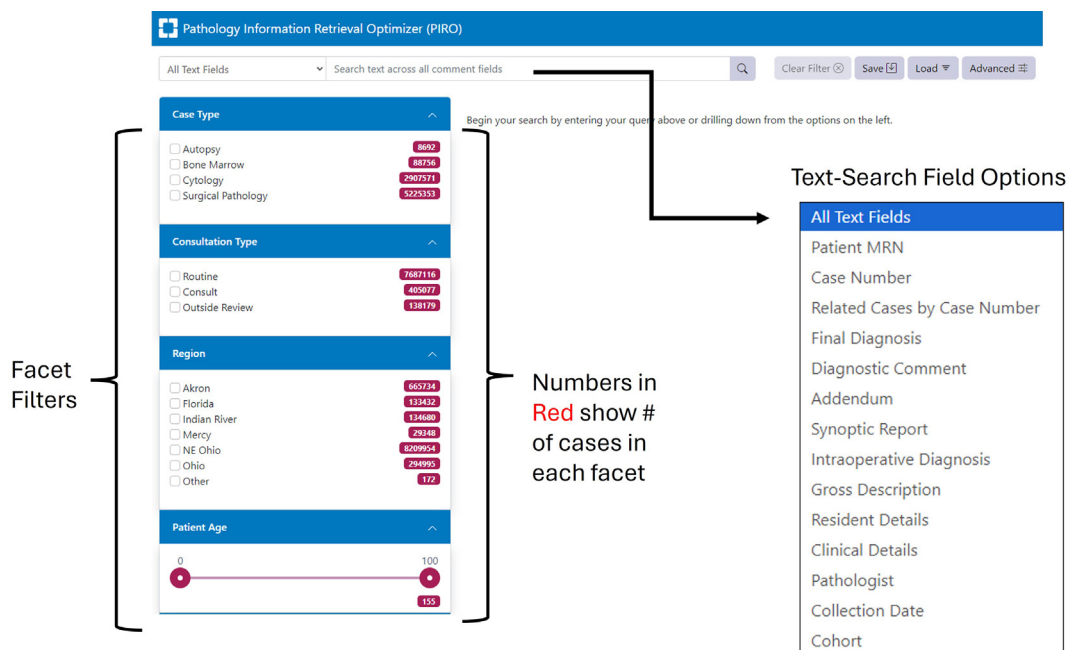
We also performed a one-time data load from Cerner Copath, our previous LIS. While, we transitioned from Copath to Beaker in February 2022, our instance of Copath is still accessible. Therefore, we designed an ETL to extract all legacy Copath reports, transform them into PIRO's schema, and load them into our SQL Server database.

Currently, PIRO holds a total of 8.2 million cases, including 5.2 million surgical pathology cases, 2.9 million cytopathology cases, 88,990 bone marrow biopsies, and 8710 autopsy cases. Of these, 4.9 million cases were extracted from Copath, with the remaining cases were sourced from Epic.

## 3. Features

### 3.1. Basic search

PIRO's main screen includes a text input box along the top, allowing users to search across different text components within a case, such as final diagnosis, diagnostic comment, addendum(s), intraoperative diagnosis, gross description, resident diagnosis, synoptic, and clinical information (Fig. 2). By default, PIRO searches across all these fields ("All Text Fields"), but users can specify a particular field using the drop-down menu. Users can



**Fig. 2.** PIRO's home page. This is the default page that users see after logging in. There is a text input box along the top, which can be configured to search various fields. The left side of the screen displays facet filters, showing the number of cases matching those filters. Users can click on a filter, triggering an immediate search with the selected filter.

The screenshot shows the PIRO Advanced Search Interface. At the top, there are buttons for 'AND' and 'OR' logic, and '+ Rule' and '+ Ruleset' buttons. The main search area contains several rows of search criteria. The first row is 'Final Diagnosis' contains 'adenocarcinoma'. Below it, there is a nested group using OR logic, containing three rows: 'Final Diagnosis' contains 'liver', 'Final Diagnosis' contains 'pancreas', and 'Final Diagnosis' contains 'bile duct'. Below the nested group, there is a row for 'Pathologist' contains 'Goldblum'. At the bottom, there is a row for 'Collection Date' >= '01/01/2023'. Each row has a red minus button to its right.

**Fig. 3.** PIRO's advanced search interface. This figure shows PIRO's Advanced Search Interface, configured for a complex search. AND logic is used at the top level of the search, requiring the Final Diagnosis to contain “adenocarcinoma,” the Pathologist to be “Goldblum,” and the Collection Date to be later than 1/1/2023. However, there is a nested group using OR logic that matches if “liver,” “pancreas,” or “bile duct” is present in the Final Diagnosis.

also search by pathologist, MRN, case number (accession number), and collection date. Additionally, the “Related Cases by Case Number” feature allows users to input a case number and retrieve all pathology cases for that patient.

Many case aspects are categorical, and for these elements, PIRO provides faceted searching with clickable checkboxes, letting users select among various categorical data elements. The number of matching cases is displayed in real-time, allowing users to quickly review and refine their searches as needed.

3.2. Advanced search

In default mode, PIRO allows only one search term per text field, so users can search for only a single string in fields like Final Diagnosis. To enable more complex queries, PIRO offers an Advanced Search mode, which allows users to enter multiple search terms and use all available filters

(Fig. 3). For example, a user could search for “adenocarcinoma” AND “lung” AND “biopsy” in the Final Diagnosis field. Users can also negate search terms with a “not” operator, matching cases in which the search term is not present. Finally, users can create nested sub-groups with user-specified AND/OR logic.

3.3. Results pane

Search results are displayed in the main window, sorted by the most recent cases (Fig. 4). Each page shows 5 search results, and the user can navigate forwards and backwards through the pages by clicking the arrow buttons. Each result contains a header, which shows the case number, patient name, age, sex, laboratory, and pathologist. Below the header, the Final Diagnosis, Diagnostic Comment, are displayed, with search terms highlighted in yellow. The accession number is a link, and clicking it will open up the full details of the case, with all of the available text fields (final diagnosis, diagnostic comment, addendum(s), gross description, intraoperative diagnosis, synoptic, clinical details). Furthermore, the user can see patient-level information (name, MRN, sex, DOB, age, vital status) and case-level details (collection date, accession date, sign-out date, laboratory location, subspecialty, case type, pathologist).

3.4. Saving searches, data extraction workflow and data governance

PIRO lets users save searches to their profile. After designing a search, users can click “Save” to store the parameters. To perform the same search later, users simply click “Load” and select the saved search from a drop-down menu. Saved searches are also the entry-point for PIRO's data extraction workflow, allowing users to request a Microsoft Excel file of their search results. This workflow was designed in close collaboration with our IRB, HIPAA and legal stakeholders. Our goal was to make user-driven, efficient process that also complies with our institutions data governance policies. To request a data extract, the user clicks the Data Request Form button. Users are taken to a page where they provide information about their request. Most importantly, the users select whether their request is for research, or not for research. If their request is for research, the user must also supply their IRB protocol number, and an IRB-verified

The screenshot shows the PIRO results interface. At the top, there is a navigation bar with 'Home', 'Search', 'Cohort', and 'Data Request Form'. Below the navigation bar, there is a search bar with 'All Text Fields' selected and 'Search text across all comment fields'. To the right of the search bar are buttons for 'Clear Filter', 'Save', 'Load', and 'Advanced'. Below the search bar, there is a 'FILTERED BY' section showing 'Helicobacter pylori gastritis' and 'Routine'. On the left side, there is a 'Facet Numbers Updated After Search' section with three categories: 'Case Type' (Autopsy: 0, Bone Marrow: 0, Cytology: 0, Surgical Pathology: 4274), 'Consultation Type' (Routine: 4274, Consult: 0, Outside Review: 0), and 'Region' (Akron: 328, Florida: 1, Indian River: 46, Mercy: 20, NE Ohio: 3877, Ohio: 2, Other: 0). The main results area shows two case entries. The first entry is for Jane Doe, 76/F, S23-123456, Cleveland Clinic Main Hospital Lab, NE Ohio, with pathologist LAPINSKI, JAMES E. The final diagnosis is 'Stomach, biopsy: - Helicobacter pylori gastritis'. The second entry is for Jane Doe, 23/F, S23-234567, Cleveland Clinic Main Hospital Lab, NE Ohio, with pathologist LAPINSKI, JAMES E. The final diagnosis is 'A. Stomach, biopsy: - Chronic active Helicobacter pylori gastritis. - Negative for intestinal metaplasia or dysplasia.' The total number of matching cases is displayed in the upper right as 4274 Cases.

**Fig. 4.** PIRO's results interface. The example shows the results of searching All Text Fields for the term “*Helicobacter pylori gastritis*.” Additionally, the Routine facet is clicked, filtering for in-house (non-consult) cases. For matching cases, a thumbnail of the case information is displayed, with matching text highlighted in yellow. The total number of matching cases is displayed in the upper right (4274 cases). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

copy of their data collection sheet, which lists all of the variables for which they are approved to collect. After the request is submitted, the request is routed to an independent analyst. For research requests, the analyst checks the requestor's data collection sheet against the details of the search request. With this check complete, the analyst executes the search, and a Microsoft Excel file is generated. The file is then securely delivered back to the requestor. Most requests are fulfilled within 24 h of submission.

### 3.5. Cohorts

PIRO enables users to build cohorts by uploading lists of patients or accession numbers. This feature is helpful for cases where an investigator has a list of patients with a specific disease or a set of accession numbers compiled elsewhere (e.g., Copath, Epic or research databases). Users can create a cohort by uploading MRNs or case numbers via an Excel template. Once uploaded, PIRO identifies matching cases, and cohorts become active facets in the user's PIRO user interface. Selecting the facet enables users to perform searches within their custom cohort.

### 3.6. Security

PIRO is accessible only within our institution's firewall. Access is requested through our IT departments service platform (ServiceNow) and granted by adding users to the PIRO's Active Directory security group. PIRO tracks search activity, logging user access to specific cases, thus providing a means for auditing search history and user access to patient records.

### 3.7. AI-generated annotations

Pathology reports typically contain unstructured data. Large Language Models (LLMs) excel at extracting discrete data elements from unstructured text.<sup>1,2</sup> We used an LLM to annotate surgical pathology cases with a "Malignancy Risk" label to flag cases asserting the presence of malignancy or high-grade dysplasia.

For the security of our PHI, we used open-source LLMs that could be deployed locally, behind our institution's firewall. Furthermore, we used Ollama, which is a LLM management platform, for locally-hosted LLMs. Ollama allows the user to easily download LLM models, and switch

between them quickly. It offers a web-based UI for fast prototyping, and a API for longer production jobs. We tested many different models for our task, ultimately, selecting Llama3:8b-instruct-fp16.<sup>3</sup> The details of our testing procedure is reported in the results section below. Ollama is hosted on our internal High Performance Cluster (HPC) which features an AMD EPYC 2.65GHz 24 Core CPU and four Nvidia A100 80 GB GPUs.

We prompted Llama3:8b-instruct-fp16 with:

PROMPT:

You are assisting a pathologist. You will read a pathology report and determine if the report asserts the presence of cancer-associated diagnoses. You will respond in JSON format with one field "malignancy\_risk", which will be either "POSITIVE" OR "NEGATIVE".

Specifically, return "POSITIVE" if the report asserts the presence of cancer, malignancy, high-grade dysplasia, carcinoma, sarcoma, mesothelioma, lymphoma or leukemia. Otherwise, report "NEGATIVE". For cases in which the diagnosis is not clearly malignant or dysplastic, and contains the word "atypical", report "NEGATIVE".

The pathology report may have many diagnoses. Analyze the most severe diagnosis.

Do not include any backslashes "\ in the JSON output.

##QUESTION##

Interpret this pathology report:

{FINAL DIAGNOSIS INSERTED HERE}

The Final Diagnosis would be concatenated to the end of this zero-shot prompt. The LLM returns the results in JSON format, in this case for a single data element, Malignancy Risk, which can only be POSITIVE or NEGATIVE. Next, the annotations are added to PIRO's SQL Server database, then indexed into Solr. Finally, this data element is surfaced in PIRO's UI as a selectable facet, in which the user can easily select positive or negative cases (Fig. 5).

### 3.8. Study period: comparing PIRO and SlicerDicer usage

PIRO usage was compared to SlicerDicer usage from January 2024 through August 2024. PIRO data was collected from PIRO's SQL database, while SlicerDicer data was provided by our Epic Cogito reporting team, filtered for anatomic pathology users to ensure comparability. From each query, we determined the query's datetime, the user's name and the execution time. For both platforms, the definition of query was similar; a unique

The screenshot displays the PIRO web application interface. At the top, there's a navigation bar with links for Home, Search, Cohort, and Data Request Form. Below this, a search bar is populated with 'Pathologist' and 'GOLDBLUM, JOHN'. To the right of the search bar are buttons for 'Clear Filter', 'Save', 'Load', and 'Advanced'. Below the search bar, there are filters for 'Case Type' and 'Consultation Type'. The 'Case Type' filter includes options for Autopsy, Bone Marrow, Cytology, and Surgical Pathology. The 'Consultation Type' filter includes options for Routine, Consult, and Outside Review. A 'Malignant' filter is prominently displayed at the bottom left, with a red flag icon and a count of 8131 cases. The main content area shows a case thumbnail for 'Jane Doe', a 76/F patient with ID 523-123456. The thumbnail includes patient information, lab location (Cleveland Clinic Main Hospital, Lab, NE Ohio), and a final diagnosis: 'A. Soft tissue, right lower quadrant abdomen, subcutaneous lesion, ultrasound-guided biopsy: - Extensively necrotic malignant neoplasm, possibly carcinoma. See comment.' A red flag icon in the upper right corner of the thumbnail indicates the Malignant filter status.

Fig. 5. PIRO's LLM-generated Malignant filter. In this example, the user is searching for Routine (in-house) cases signed out by Dr. Goldblum. The Malignant filter is checked as Positive. A red flag in the upper right of the case thumbnail denotes this status. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



query would be recorded each time the user adjusted the query parameters, in any way. Therefore, for both platforms, a single user session would usually result in many individual queries, as the user adds, modifies or deletes search parameters.

4. Results

4.1. Usage analytics

During the study period from January 2024 through August 2024, PIRO averaged 12,288 queries per month compared to 384 queries submitted through SlicerDicer (Fig. 6). Over this period, 168 unique users performed a PIRO search, whereas only 11 unique users conducted a SlicerDicer search. Our anatomic pathology group consists of 119 staff pathologists, of whom 80 (67 %) performed at least one PIRO search, while only 11 (9 %) used SlicerDicer.

We categorized each PIRO user as staff, trainee (resident or fellow), or other (e.g., QA specialists, administrative assistants, analysts, clinical researchers). Staff pathologists averaged 7522 (59 %) queries per month, trainees averaged 4683 (37 %) queries per month, and other users averaged 503 (4 %) queries per month (Fig. 7).

The average execution time for PIRO queries was 0.14 s (standard deviation 0.34 s) compared to 409.3 s (standard deviation 810.6 s) for SlicerDicer queries.

4.2. PIRO and SlicerDicer user experience survey

Our analytics clearly indicated that PIRO is used much more than SlicerDicer, however, we wanted to understand what factors may explain this difference (Table 1). Therefore, we conducted a survey that was sent out to all our AP staff (n = 119) and trainees (n = 50). We received 55 responses, from 37 staff and 18 trainees. Interestingly, there was large difference in the number of users who used each tool. 50 respondents used PIRO while only 10 had ever used SlicerDicer. Users were asked to rate PIRO and SlicerDicer on a scale of 1 to 10 with respect to three questions. How easy was it to learn PIRO/SlicerDicer? How would you rate the search capabilities of PIRO/SlicerDicer? How would you rate the query speed of

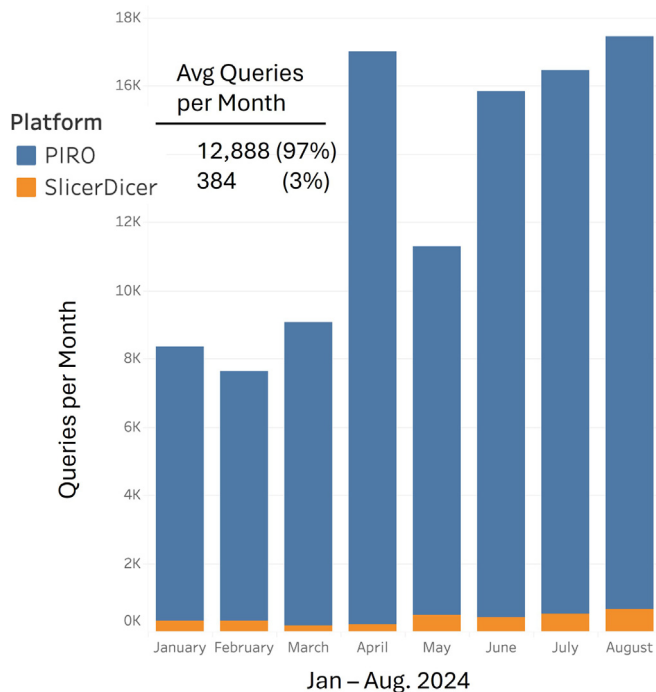


Fig. 6. Bar chart comparing PIRO usage to SlicerDicer usage over an 8-month study period.




Fig. 7. Bar chart showing PIRO usage by staff, trainees, and other personnel.

PIRO/SlicerDicer. PIRO scored better than SlicerDicer for all three questions. Respondents were also asked to list the use-cases for which they use each tool (Fig. 8). The options were: Real-time knowledge retrieval (learning from the text of previous cases, sometimes to aid in the diagnosis of a current case), case-finding (for education, conferences, or presentations), feasibility testing prior to research, generating data extracts for research, operational analytics, QA/QC projects. PIRO was used most frequently for real-time knowledge retrieval and case-finding, while SlicerDicer was used most frequently for operational analytics.

4.3. LLM test results for malignancy risk


A random sample of 1000 surgical pathology reports from 2023 was retrieved from PIRO. Each report was manually scored by one author (SR) as positive or negative for Malignancy Risk, defined as cases indicating the presence of malignancy, or high-grade dysplasia.

The sample was processed using several models (see Methods), with Llama3 models yielding the best performance. Llama3:70b-instruct-fp16 achieved 99 % accuracy, 96 % sensitivity, and 100 % specificity. Llama3:8b-instruct-fp16 demonstrated 99 % accuracy, 98 % sensitivity, and 99 % specificity. Given its smaller memory footprint, the 8b parameter model was selected for production. Consequently, Llama3:8b-instruct-fp16 was used to process our entire corpus of surgical pathology reports (~5.2 million), averaging 0.11 s per report. This resulted in approximately 850,000 reports (16 %) flagged as positive for Malignancy Risk.

Table 1  
PIRO and SlicerDicer User Experience Survey.

Overall Survey Respondents (N = 55)		
Staff = 37 (67 %)		
Trainees = 18 (33 %)		
Respondent has used:	PIRO N = 50 (91 %)	SlicerDicer N = 10 (9 %)
Questions with a 1–10 scale*:		
How easy was it to learn?	9.1	3.7
How would you rate its search capabilities?	8.9	4.2
How would you rate its query speed?	9.3	5.4

\* Values are an average across respondents.



**Fig. 8.** Bar chart showing how frequently PIRO and SlicerDicer were used for various use-cases. The options were: Real-time knowledge retrieval (learning from the text of previous cases, sometimes to aid in the diagnosis of a current case), case-finding (for education, conferences, or presentations), feasibility testing prior to research, generating data extracts for research, operational analytics, QA/QC projects.

5. Discussion

The Laboratory Information System (LIS) plays a fundamental role in managing workflows and supporting timely clinical decision-making.<sup>4</sup> However, beyond these essential tasks, the LIS should also facilitate database queries essential for generating reports, monitoring QA/QC metrics, and supporting research and educational case-finding. Our current LIS, Epic Beaker, shows gaps in case-finding capabilities for end-users. User adoption of SlicerDicer, Epic Beaker's built-in query tool, has been limited among our staff, with only 9 % of anatomic pathology staff using it during our eight-month study period. Although all staff pathologists have access to SlicerDicer and receive onboarding tutorials, PIRO has seen greater adoption, with 67 % of our staff using it for ad hoc data queries, indicating PIRO's effectiveness and preference among users. Our user-survey was conducted to determine why PIRO is used more frequently than SlicerDicer. Overall, users found that PIRO was easier to use, had superior search capabilities and that the query speed was superior, compared to SlicerDicer.

PIRO is not the first tool designed to enhance anatomic pathology search functionality. Several custom-built tools exist at various institutions for this purpose.<sup>5,6</sup> However, PIRO introduces significant advancements. Firstly, it leverages Apache Solr, a powerful full-text search engine optimized for handling large volumes of text data through inverted indexing. Unlike SQL databases, which excel with structured data, Solr efficiently processes unstructured text data, enhancing speed and flexibility. Solr also includes faceted search, enabling users to dynamically filter and group results—features that SQL can replicate only through complex and often inefficient queries.

Secondly, PIRO is highly adaptable across different medical centers, especially those using Epic Beaker. We are currently extracting pathology data from our enterprise data vault (EDV) using a Clarity-like schema, and are developing ETL scripts that will enable direct extraction from Clarity, streamlining deployment for Epic Beaker institutions. A great feature of PIRO is its LIS-agnostic design, supporting data ingestion from other systems, as we demonstrated by integrating data from both Epic Beaker and our previous LIS, Cerner Copath.

About a third of all search activities are conducted by trainees, most of which occur during a trainee's preview period. We encourage our trainees to preview as many cases as possible and record their findings in the LIS.

During this preview time, our trainees utilize PIRO as a learning tool to help them formulate diagnostic text. The ability to craft diagnostic text is an important skill for pathologist to develop. The diagnostic text serves as our primary method of communication with medical practitioners, and becomes part of the permanent medical record, driving clinical decision making.

Moreover, pathology diagnoses are not always conclusive, requiring pathologists to choose their words carefully to convey the appropriate level of certainty. A well-crafted diagnosis contributes to effective patient care, while a poorly phrased one can create confusion and potentially lead to inappropriate treatment. Unfortunately, there are limited resources available to help trainees develop this specific skill. While numerous books support trainees with morphologic and clinicopathologic descriptions of various pathological entities, formulating or dictating the diagnosis often requires a process of trial and error. For pathology trainees, this can be an intimidating task with a steep learning curve. The capability to quickly search the diagnostic archive is invaluable in this context, and PIRO has become a crucial resource on which our trainees rely.

6. Future directions

As our institution rolls out a new digital pathology (DP) platform, we aim to integrate it with PIRO to support clinical and research needs more seamlessly. The integration will allow users to search PIRO for cases with available digital slides, accessible through a clickable facet. Post-search, each case's result will include a direct link to the DP platform, enabling rapid access to digital slides—streamlining workflows and eliminating the need to retrieve glass slides from the archive.

We also plan to broaden PIRO's use of large language models (LLMs) to convert unstructured text into specific, searchable annotations. Currently, we generate a single Malignancy Risk for surgical pathology reports. Next, we plan to extend this annotation to cytology and bone marrow specimens—though, the accuracy of the model will need to be tested in these contexts. Furthermore, we plan to expand the number of annotations per case. To start, we plan to generate several discrete data elements for oncologic cases, like cancer type, primary site and grade. Also, oncologic cases often contain the results of biomarker and molecular tests, like mismatch repair proteins and PD-L1. An LLM can pick out these results from the free text, and save them as discrete data.

However, scaling-up LLM usage presents computational challenges. For instance, generating the Malignancy Risk annotation took 0.11 s per report, totaling 6.6 days to process 5.2 million surgical pathology reports. With a vast range of potential annotations, further scaling requires strategies to optimize computational efficiency. Our first approach is to target LLM inference more precisely, reducing processing to cases that meet specific criteria. For example, we plan to use the Malignancy Risk annotation to pre-filter cases for extracting oncologic data elements (eg cancer type, primary site, grade, etc). This will allow us to focus the annotation job on the 16 % of cases that assert the presence of malignancy.

A second approach is to batch multiple annotations within a single LLM call, such as extracting malignancy type, primary site, and grade in one LLM call. This approach could substantially reduce computational demand per annotation. However, we will have to determine how batching annotations impacts accuracy. Similarly, we will need to determine the optimal balance between query complexity and processing speed, as both input and output length can significantly affect LLM inference time.<sup>7,8</sup>

## 7. Conclusion

PIRO demonstrates an effective, scalable solution for enhancing pathology text search capabilities in laboratory information systems. By leveraging a combination of FastAPI, Angular, Apache Solr, and LLMs, PIRO provides a robust, user-friendly platform that meets both the complex and day-to-day needs of pathologists. The tool's successful adoption—evidenced by significantly higher user engagement compared to SlicerDicer. Furthermore, PIRO's flexible design, allowing for deployment across various LIS infrastructures, makes it well-suited for broader institutional adoption. Future developments, including integration with Digital Pathology and expanded LLM annotations, aim to further enhance PIRO's clinical

and research applications, offering a promising framework for continuous innovation in pathology informatics.

## Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author(s) used ChatGPT (OpenAI) for spelling and grammar checking. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

## References

1. Hart SN, Hoffman NG, Gershkovich P, et al. Organizational preparedness for the use of large language models in pathology informatics. *J Pathol Inform* 2023;14, 100338. <https://doi.org/10.1016/J.JPI.2023.100338>.
2. Sciannameo V, Pagliari DJ, Urru S, et al. Information extraction from medical case reports using OpenAI InstructGPT. *Comput Methods Programs Biomed* 2024;255. <https://doi.org/10.1016/J.CMPB.2024.108326>.
3. Dubey A, Jauhri A, Pandey A, et al. The Llama 3 Herd of Models. Published online July 31, 2024. Accessed November 4, 2024: <https://arxiv.org/abs/2407.21783v2>.
4. Henricks WH. Laboratory information systems. *Clin Lab Med* 2016;36(1):1-11. <https://doi.org/10.1016/J.CLL.2015.09.002>.
5. Robertson SA. A novel web application for rapidly searching the diagnostic case archive. *J Pathol Inform* 2020;11(39).
6. Arvisais-Anhalt S, Lehmann CU, Bishop JA, et al. Searching full-text anatomic pathology reports using business intelligence software. *J Pathol Inform* 2022;13, 100014. <https://doi.org/10.1016/J.JPI.2022.100014>.
7. Fu Q, Cho M, Merth T, Mehta S, Rastegari M, Najibi M. LazyLLM: Dynamic Token Pruning for Efficient Long Context LLM Inference. Published online July 19, 2024. Accessed November 5, 2024: <https://arxiv.org/abs/2407.14057v1>.
8. Yang Y, Xu Y, Jiao L. A Queueing Theoretic Perspective on Low-Latency LLM Inference with Variable Token Length. *WiOpt* 2024. Published online July 7, 2024. Accessed November 5, 2024: <https://arxiv.org/abs/2407.05347v1>.