

Empowering Digital Pathology Applications through Explainable Knowledge Extraction Tools

Stefano Marchesin^a, Fabio Giachelle^a, Niccolò Marini^{b,g}, Manfred Atzori^{b,c}, Svetla Boytcheva^d, Genziana Buttafuoco^e, Francesco Ciompi^f, Giorgio Maria Di Nunzio^a, Filippo Fraggetta^e, Ornella Irrera^a, Henning Müller^b, Todor Primov^d, Simona Vatrano^e and Gianmaria Silvello^{a,*}

ARTICLE INFO

Keywords:

Clinical Practice
Digital Pathology
Expert Systems
Explainable AI
Knowledge Extraction
Machine Learning

ABSTRACT

Exa-scale volumes of medical data have been produced for decades. In most cases, the diagnosis is reported in free text, encoding medical knowledge that is still largely unexploited. In order to allow decoding medical knowledge included in reports, we propose an unsupervised knowledge extraction system combining a rule-based expert system with pre-trained Machine Learning (ML) models, namely the Semantic Knowledge Extractor Tool (SKET). Combining rule-based techniques and pre-trained ML models provides high accuracy results for knowledge extraction. This work demonstrates the viability of unsupervised Natural Language Processing (NLP) techniques to extract critical information from cancer reports, opening opportunities such as data mining for knowledge extraction purposes, precision medicine applications, structured report creation, and multimodal learning.

SKET is a practical and unsupervised approach to extracting knowledge from pathology reports, which opens up unprecedented opportunities to exploit textual and multimodal medical information in clinical practice. We also propose SKET eXplained (SKET X), a web-based system providing visual explanations about the algorithmic decisions taken by SKET.

SKET X is designed/developed to support pathologists and domain experts in understanding SKET predictions, possibly driving further improvements to the system.

1. Introduction

Exascale volumes of multimodal data have been produced for decades in the biomedical domain. Biomedical data include patient information, clinical data, biological laboratory data, bio-images, bio-signals, instrumental examinations, and genetic data.

Hundred of thousands of reports have been used to describe findings leading to diagnoses, encoding vast medical knowledge. Free-text reporting is the standard for communicating the diagnosis, guiding patients' treatment, and other applications, such as cancer registries. Processing high volumes of free-text reports, usually performed manually, is also required to extract knowledge to train Machine Learning (ML) algorithms.

However, the manual analysis of data becomes an extremely time-consuming process since reports vary widely between institutions, might be written in languages other than English, contain noise, and do not present a standard structure. In this context, Natural Language Processing (NLP) methods are central [1, 2, 3, 4, 5, 6, 7, 8] as they empower the efficient automatic processing of thousands of clinical reports and the extraction of key information for several downstream tasks, such as clinical note mining [9, 10] and structuring [11], risk prediction [12], clinical decision support [13], and precision medicine retrieval [14].

In the context of digital pathology, NLP techniques can drive noticeable advances by exploiting the availability of textual pathology reports paired with digital histopathology images (i.e., Whole Slide Images (WSIs)) in clinical practice. WSIs are used as a gold standard to diagnose cancer cases and related diseases [15, 16]. Within WSIs, tissue patterns and morphology vary depending on the image magnification level – enabling different tasks such as detection, classification, or segmentation [17]. However, the lack of training datasets containing pixel-wise annotations for entire images [18, 19, 20] limits the effectiveness of supervised ML models [21]. Nevertheless, from the textual pathology reports, it is possible to extract key concepts (e.g., the diagnosis outcome) to annotate the associated

*Corresponding author

 <http://www.dei.unipd.it/~marches1, stefano.marchesin@unipd.it> (S. Marchesin)

ORCID(s):

WSIs. Although noisy, the extracted concepts can then serve as weak labels to train prediction models for image classification tasks [22, 23]. However, even though automated solutions involving ML are increasingly being integrated into biomedical domains, NLP applications to digital pathology are less common. Compounding the situation further, the actual use of Artificial Intelligence (AI) algorithms in digital pathology requires a large amount of data annotations by pathologists. However, they are rarely available in a clinical setting [1, 24].

To overcome such limitation, this work aims at proving the viability of unsupervised NLP techniques to automatically extract critical information from pathology reports and use it for different digital pathology applications, such as automatic report annotation, pathological knowledge visualization, and WSI classification. In this regard, we present the Semantic Knowledge Extractor Tool (SKET), an unsupervised hybrid knowledge extraction system that combines an expert system with pre-trained ML models to extract knowledge from pathology reports. In recent years, NLP has shifted from using rules to ML approaches [25, 8], which have the advantage of learning regularities from data and of generalizing to previously unseen patterns.

Moreover, the advent of efficient Neural Language Models (NLMs) [26, 27, 28, 29] paved the way for the pre-training era, where large NLMs trained in a self-supervised fashion on huge datasets are used to develop NLP models for a number of downstream tasks. Nevertheless, similarly to Santus et al. [9], we argue that rule-based techniques capture critical information that should be used together with – and not substituted by – ML to improve performance.

We evaluate SKET effectiveness on entity linking and text classification, considering three different diseases: colon, cervix, and lung cancer. In this regard, we resort on diagnostic reports coming from two medical centers in Italy and The Netherlands. Then, we compare SKET with unsupervised ML approaches to understand the impact that combining rule-based techniques and pre-trained ML models have on the extraction of knowledge from pathology reports. The achieved results highlight the viability of ML methods for information extraction in the pathology domain, but also stress the importance of expert knowledge to reach the high levels of accuracy required to (semi-)automate the clinical practice. Moreover, the applicability of the proposed approach is enhanced by the considered multilingual setting.

Besides effectiveness, we must consider that understanding and explaining decisions and outcomes is crucial in clinical practice.

However, the black-box nature of many ML models, especially those based on Deep Learning, makes it difficult to understand and trace back the underlying decision process. Hence, there is an urgent need for a shift towards eXplainable Artificial Intelligence (XAI) [30, 31, 32, 33].

In the biomedical domain, clinicians and domain experts need to understand *why* a specific output has been produced to trust the system and its predictions; moreover, the explainability of algorithmic decisions is increasingly required for legal reasons [34]. To this end, we propose SKET eXplained (SKET X). This web-based system allows domain experts to interact with SKET and visually comprehend the outcomes, rules, and parameters used in the knowledge extraction process. In addition, SKET X allows users to compare different SKET executions (e.g., with varying parameters of the system) and inspect punctual information, which provides valuable insights into the knowledge extraction process and the contribution of each component and parameter. SKET X aims to support pathologists and domain experts in the interaction with SKET, allowing them to gain an in-depth comprehension of the system decision process – and thus increasing trust and confidence in the system. Beyond explainability, we also report different digital pathology applications where SKET has been successfully integrated as a core system [35, 36]. In particular, we deepen the use of SKET in such applications and the advantages it entails.

SKET source code is publicly available at <https://github.com/ExaNLP/sket>. Besides, SKET can also be deployed as a Docker container. For information about the Docker version of SKET, please refer to <https://github.com/ExaNLP/sket#docker>. SKET X is available at <http://w3id.org/sketx>¹.

The rest of this paper is organized as follows: Section 2 describes the considered data resources. Section 3 presents SKET. Section 4 describes the experimental setup and reports quantitative and qualitative results. Section 5 presents SKET X. Section 6 outlines the digital pathology downstream applications empowered by SKET. Finally, Section 7 draws some conclusions.

2. Material

The data used to develop and evaluate SKET comes from two different medical centers: the Cannizzaro Hospital (AOEC), Catania, Italy and the Radboud University Medical Center (RUMC), Nijmegen, The Netherlands. The AOEC data includes diagnostic reports for colon, cervix, and lung cancer cases, written in Italian and associated with WSIs.

¹Access credentials for reviewing: [demo/demo](#)

Table 1

Data size. For each medical center, we report the number of diagnostic reports associated with each use-case. The “–” symbol represents the lack of reports for a given use-case.

	Colon	Cervix	Lung
AOEC	1,704	1,777	1,902
RUMC	2,065	2,350	–

All data were collected in the clinical workflow and fully anonymized afterwards. Similarly, the RUMC data consists of diagnostic reports and the associated WSIs for colon and cervix cases, written in Dutch – after the use of speech-to-text tools – and anonymized. For both medical centers, the considered reports have been provided directly in digital format and span several diagnostic outcomes. Note that other medical centers may provide reports in non-digital format, thus requiring a digitization step upstream of SKET. Table 1 reports the total number of diagnostic reports for each considered use-case and medical center.

Diagnostic reports contain the results of the analyses performed on specific tissues (or cells) to obtain a pathological-clinical diagnosis – i.e. presence or absence of the disease. AOEC and RUMC diagnostic reports follow the College of American Pathologists (CAP) international guidelines² for pathology reports [37, 38] and contain the patient’s personal and clinical-specific information, the description of how a specimen appears to the naked eye and at the microscope, and provide the final diagnosis.

As mentioned above, AOEC and RUMC diagnostic reports are written in Italian and Dutch, respectively. However, most of the resources required to develop NLP methods that extract concepts from unstructured text are in English. To overcome this limitation, we first translated diagnostic reports in English and then performed data curation over them. We used the open-source, pre-trained Marian Neural Machine Translation (NMT) models [39], which exhibit a Transformer-based [40] encoder-decoder architecture with six layers in each component. Given the complexity of the task, such an automatic approach introduces systematic translation errors that, if propagated, could hamper the effectiveness of the extraction process. For this reason, we performed a data curation step, in which recurring, manually identified translation errors were corrected through the use of handcrafted rules.

We defined an ontology³ for modeling the clinical reports in the digital pathology domain: ExaMode⁴ ontology. Amongst other aspects not relevant for the current work, the ontology specifically defines the key concepts and properties to model the diagnosis of colon, cervix and lung cancer, the anatomical location where the disease might be located, the procedure employed to get the tissue, and the tests conducted on the tissue. Despite many medical ontologies focusing specifically on cancer exist, no single ontology comprehensively models all the diseases related to the cases mentioned above, their anatomical location, topography, and pathology laboratory process.

3. Methods

SKET adopts a combination of pre-trained Named Entity Recognition (NER) models and unsupervised Entity Linking (EL) methods to extract key concepts (entities) from the diagnostic reports and to link them to the reference ontology. The use of pre-trained NER models and unsupervised EL methods makes SKET suitable for weak supervision tasks. In this regard, the pathological concepts extracted from diagnostic reports can serve as weak labels to train prediction models for image classification tasks [22, 23], or as nodes to build report-level knowledge graphs for information retrieval tasks [41].

As reported in Figure 1, SKET consists of four components: (A) Named Entity Recognition, (B) Entity Linking, (C) Data Labeling, and (D) Graph Creation. Components (A) and (B) are sequential, whereas components (C) and (D) are parallel. Below, we describe for each component the different methods and techniques we adopted, expanded, or developed.

²<https://www.cap.org/protocols-and-guidelines>

³<https://w3id.org/examode/ontology/>

⁴ExaMode stands for “Extreme-scale Analytics via Multimodal Ontology Discovery & Enhancement” and is an H2020 project financed by the EU commission. More information can be found at: <http://www.examode.eu/>

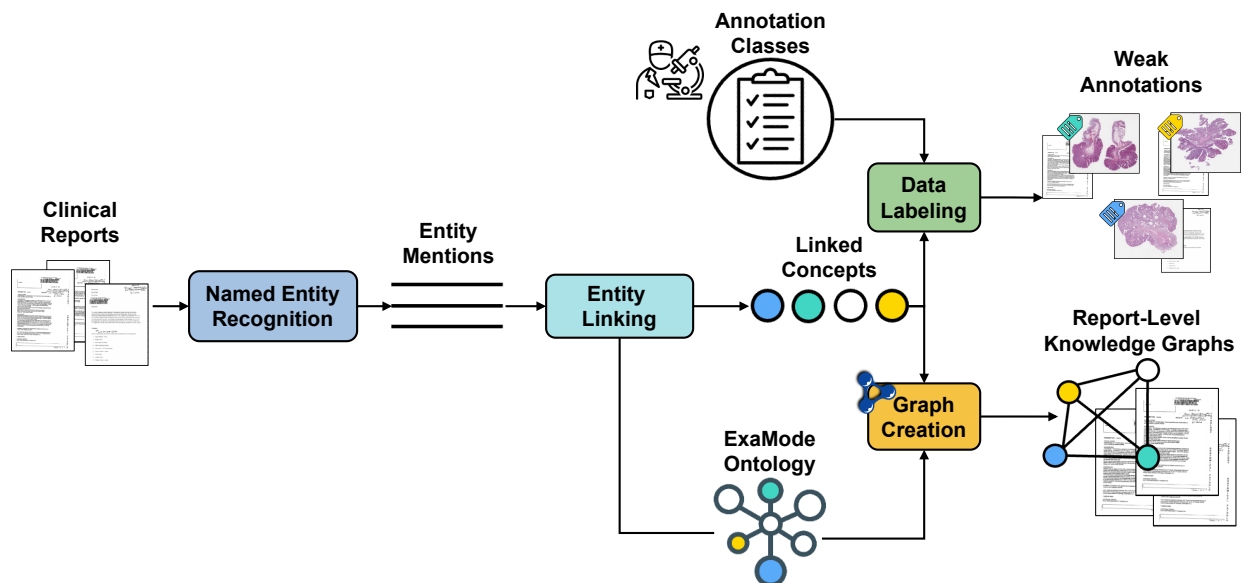


Figure 1: SKET architecture. SKET main components are: (A) Named Entity Recognition, (B) Entity Linking, (C) Data Labeling, and (D) Graph Creation.

3.1. Named Entity Recognition

NER is the task of identifying and categorizing key information – i.e., entities – within text. An entity can be any word or phrase that consistently refers to the same concept or object of the world. Each identified entity is classified into a pre-defined category, such as disease, protein, gene, cell type, etc.

SKET relies on a combination of pre-trained neural models and rule-based techniques to perform NER. At its core, SKET adopts Scispacy models [42], which provide full NER pipelines for biomedical data, comprising large medical vocabularies, and Word2Vec [26] word vectors trained on the PubMed Central Open Access Subset [43]. It is worth mentioning that SKET has been designed to be deployed with any of the core models available at: <https://allenai.github.io/scispacy/>.

Then, SKET extends the Scispacy pipeline with two additional components: Entity Fusion and Negation Detection.

Entity Fusion: SKET extends the NER pipeline with a set of rules used to identify and merge specific entities otherwise regarded as separate by Scispacy. For instance, “transverse” and “colon” are considered as separate entities, whereas we are interested in “transverse colon” as a unique entity. Hence, we developed regular expressions that identify trigger terms indicative of a set of otherwise potentially separate entities. Once a trigger term is identified, SKET matches the entities extracted by Scispacy with the candidate terms associated with the trigger. Depending on the trigger term, the match that SKET performs between extracted entities and candidate terms follows different rules based on directional and positional attributes. Directional attributes specify the set of extracted entities to be matched with the candidate trigger terms, and it can assume three values:

- PRE: match with the entities preceding the trigger entity.
- POST: match with the entities succeeding the trigger entity.
- BOTH: match between the entities both preceding and succeeding the trigger entity.

Positional attributes specify the maximum distance allowed between the trigger entity and the other one, and it can assume two values:

- EXACT: the matched entity must be right before/after the trigger entity.
- LOOSE: the matched entity can be anywhere before/after the trigger entity.

The described set of rules has been developed on a holdout dataset and it is available on the SKET GitHub repository⁵. The dataset consists of 50 diagnostic reports for each use-case and medical center, for a total of 250 diagnostic reports.

Negation Detection: To handle negated entities, we extend the NER pipeline with NegEx [44], a negation detection algorithm evaluating whether extracted entities are negated within text. NegEx uses regular expressions to identify the scope of trigger terms that are indicative of negation, such as “no” or “ruled out”. Then, the entities extracted within the scope of a trigger term are marked as negated. In this way, SKET identifies – and removes from the final list of extracted entities – those entities that NegEx regards as negated. For example, if we consider the phrase “free of dysplasia”, NegEx identifies the trigger term “free of” and marks “dysplasia” as negated, which is then removed by SKET.

3.2. Entity Linking

EL is the task of assigning unique meanings to entities mentioned within text. In other words, the objective of EL is to determine whether a given entity refers to a specific concept or object within a reference ontology.

SKET employs a combination of ad-hoc and similarity matching techniques to link the extracted entities to unique concepts within the ExaMode ontology. Given an extracted entity, SKET first tries to match it using ad-hoc matching and when it fails SKET employs similarity matching.

Ad-Hoc Matching: SKET uses regular expressions to identify trigger terms indicative of a specific ontological concept. Once a trigger term is identified, SKET matches the entity containing the trigger term with the closest ontology concept. For instance, if an extracted entity contains the term “carcinoma”, then SKET links the entity to the ontology concept “colon adenocarcinoma”. As for Entity Fusion, the ad-hoc matching rules have been developed on the holdout dataset and are available on GitHub.

Similarity Matching: SKET performs similarity matching using a combination of string and semantic matching techniques. For string matching, SKET relies on the Gestalt Pattern Matching (GPM) algorithm [45], which computes the similarity of two strings as the number of matching characters divided by the total number of characters in the two strings. Matching characters are those in the longest common subsequence plus, recursively, matching characters in the unmatched region on either side of the longest common subsequence. For semantic matching, SKET exploits the word vectors provided by ScispaCy models [42]. In other words, SKET performs semantic matching as the cosine distance between the vector representations of the extracted entities and the ontology concepts – where vector representations are the mean of the word vectors composing the extracted entities or the ontology concepts.

Both string and semantic matching produce a ranking of ontology concepts ordered by decreasing similarity with a given target entity. To combine the two rankings – and select the concept with the highest rank – SKET performs rank fusion using the CombSUM [46] with min-max normalization. Before selection, a pruning phase is performed on the combined ranking, in which ontology concepts with a similarity score lower than a predetermined threshold are removed. The threshold value has been set empirically to 1.8 using the holdout dataset. The pruning phase aims to increase precision by reducing false positives, which occur when ontology concepts are incorrectly linked to the extracted entities.

3.3. Data Labeling

SKET also provides labels as one of its main outputs. Given the set of concepts extracted from each diagnostic report, SKET maps a clinically relevant subset of such concepts to a set of annotation classes defined by AOEC pathologists. For each use-case, we report below the set of annotation classes.

Colon Annotations: (1) Cancer; (2) Adenomatous polyp - high grade dysplasia; (3) Adenomatous polyp - low grade dysplasia; (4) Hyperplastic polyp; (5) Non-informative.

Cervix Annotations: (1) Cancer - adenocarcinoma in situ; (2) Cancer - adenocarcinoma invasive; (3) Cancer - squamous cell carcinoma in situ; (4) Cancer - squamous cell carcinoma invasive; (5) High grade dysplasia; (6) Low grade dysplasia; (7) HPV infection present; (8) Koilocytes; (9) Normal squamous; (10) Normal glands.

⁵<https://github.com/ExaNLP/sket/tree/main/sket/nerd/rules/>

Lung Annotations: (1) Cancer - non-small cell cancer, adenocarcinoma; (2) Cancer - non-small cell cancer, large cell carcinoma; (3) Cancer - non-small cell cancer, squamous cell carcinoma; (4) Cancer - small cell cancer; (5) No cancer.

Thus, the Data Labeling component produces annotations from diagnostic reports that can be used to perform weakly supervised classification tasks.

3.4. Graph Creation

SKET also builds report-level knowledge graphs using the extracted concepts as nodes and the semantic relations of the ExaMode ontology as edges. The use of ontology concepts and relations to describe diagnostic reports increases the semantic understanding of the underlying data [47]. Once created, report-level knowledge graphs are encoded in a machine-readable format through Resource Description Framework (RDF).

4. Evaluation

4.1. Tasks

We evaluate the effectiveness of SKET on Entity Linking (Task 1) and Text Classification (Task 2). The evaluation of SKET on entity linking also serves as a proxy to validate the quality of the RDF graphs it produces. On the other hand, text classification results help understanding the viability of using SKET as an automatic annotator in weak supervision tasks. Between the two tasks, text classification has a prominent role as it provides weak annotations that can be used to reduce the high costs of training cancer assisted diagnosis tools – which prevent unleashing the full potential of digital pathology applications [36].

4.2. Datasets

Entity Linking (Task 1): We evaluate SKET effectiveness to extract concepts from pathology reports on a subset of the proprietary data described in Section 2. For each use-case and medical center, 250 reports have been manually annotated by experts using the concepts from the ExaMode ontology. Overall, the total number of annotated reports amounts to 1,250. In terms of annotations, all use-cases have been annotated with a large number of different concepts. For colon cancer, the number of different concepts that can be found within reports stands at 19, while for cervix and lung cancer amounts to 21 and 11, respectively. This large number of different concepts highlights the complexity of the task, both for model predictions and human annotation efforts. In particular, the task can be seen as an extreme multi-label classification problem [48, 49], where the goal is to tag a given report with a subset of the relevant concepts from a large concept list.

Text Classification (Task 2): To evaluate the effectiveness of SKET to weakly annotate pathology reports, the proprietary data described in Section 2 has been manually labeled by experts using the annotation classes defined by AOEC pathologists. For each use-case, AOEC and RUMC reports have been annotated with one or more classes, making the task a multi-label classification problem. Table 2 reports the total number of reports annotated for each class in each use-case. Given the multi-label nature of the task, the total number of annotations does not reflect the total number of reports. As a side note, the class imbalance of the datasets reflects a real-case scenario, where certain conditions – e.g., low-grade dysplasia in colon cases – occur more often than others in the clinical routine.

4.3. Baselines

Entity Linking (Task 1): We compare SKET with two unsupervised approaches based on Bio FastText [27, 50] and BioClinical BERT [29, 51] models. For a fair comparison, both approaches adopt the same NER ScispaCy pipeline used by SKET, but without the extensions introduced with it. Then, the approaches perform EL by computing the cosine distance between the vector representations of the extracted entities and the ontology concepts – obtained with FastText in one case and with BERT in the other. The ontology concept closest to the extracted entity is kept and, when appropriate, mapped to the corresponding annotation class. Both methods represent a straightforward approach to perform text classification with lack of annotated data.

Text Classification (Task 2): We compare SKET with the Bio FastText and BioClinical BERT unsupervised approaches described above. Beyond unsupervised approaches, we also use SKET to weakly annotate diagnostic reports and then train FastText and BERT models in a supervised fashion. In this case, we stack a classification layer on top of the pre-trained models and perform end-to-end classification – that is, the models take diagnostic reports

Table 2

Number of annotated diagnostic reports for each use-case. Label counts are independent of each other except for “Non-informative” in colon, “Normal squamous” and “Normal glands” in cervix, and “No cancer” in lung, which only occur when none of the others does.

Colon	
Cancer	495
Adenomatous polyp - high grade dysplasia	510
Adenomatous polyp - low grade dysplasia	841
Hyperplastic polyp	508
Non-informative	1,140
Cervix	
Cancer - adenocarcinoma in situ	125
Cancer - adenocarcinoma invasive	32
Cancer - squamous cell carcinoma in situ	638
Cancer - squamous cell carcinoma invasive	88
High grade dysplasia	1,544
Low grade dysplasia	1,053
HPV infection present	1,221
Koilocytes	86
Normal squamous	1,265
Normal glands	1,266
Lung	
Cancer - non-small cell cancer, adenocarcinoma	961
Cancer - non-small cell cancer, large cell carcinoma	68
Cancer - non-small cell cancer, squamous cell carcinoma	528
Cancer - small cell cancer	144
No cancer	247

as input and directly produce classes as output. Due to the introduction of supervised models, performances on text classification are obtained through 10-fold cross-validation.

4.4. Results

Entity Linking (Task 1): Table 3 reports the results obtained by SKET and the considered baselines on entity linking. Overall, we see that SKET achieves high performances for both micro- and weighted-average F1 measures in each use-case. As for accuracy, the performances vary depending on the use-case, and the lowest score is obtained in colon cancer with a value of 0.6280. In terms of use-cases, the best SKET results are obtained on lung cancer. Compared to colon and cervix cases, lung cancer presents a lower number of concepts to identify, thus reducing the task complexity. On the other hand, colon and cervix use-cases show similar SKET performances, having a comparable number of concepts.

When we compare SKET performances with unsupervised approaches we can see that SKET outperforms them for all measures in each use-case. This result shows the effectiveness of combining ad-hoc rules with ML models, which make SKET both precise and sensitive. Indeed, ad hoc matching makes SKET precise while semantic matching makes it sensitive. To further support this outcome, we observe that the performances of unsupervised baselines – only relying on ML models and semantic matching – have low accuracy values. Given that we consider entity linking as a multi-label task, we resort on subset accuracy – where the set of concepts predicted for a report must exactly match the corresponding set of ground-truth concepts. Thus, accuracy values are more prone to rapidly decreasing with a large number of classes, and less precise models are naturally affected by this behavior.

Text Classification (Task 2): Table 4 reports the results obtained by SKET and the considered baselines on text classification. Overall, we observe that SKET achieves high performance on colon and lung cancer use-cases, whereas it shows low accuracy values on cervix cancer. The motivation behind this drop in performance on cervix reports can

Table 3

Entity linking results on colon, cervix, and lung cancer pathology reports. The considered measures are subset accuracy, micro F1, and weighted F1. **Bold** values represent the highest scores achieved for each measure.

Colon				
Approach	Model	Measures		
		Accuracy	Micro F1	Weighted F1
Unsupervised	SKET	0.6280	0.8861	0.8694
	FastText	0.0660	0.5000	0.6146
	BERT	0.1840	0.3905	0.4527
Cervix				
Approach	Model	Measures		
		Accuracy	Micro F1	Weighted F1
Unsupervised	SKET	0.7020	0.8322	0.8368
	FastText	0.0900	0.2802	0.3439
	BERT	0.0720	0.2715	0.2940
Lung				
Approach	Model	Measures		
		Accuracy	Micro F1	Weighted F1
Unsupervised	SKET	0.8624	0.9375	0.9262
	FastText	0.2510	0.5610	0.6506
	BERT	0.3806	0.6804	0.8395

be attributed to the high number of annotation classes (i.e., ten) and the multi-label setting. We recall that we rely on subset accuracy, which performs exact match between predicted and ground-truth labels – causing performance to drop faster when the number of classes is larger. The higher values for both micro and weighted F1 measures, which do not perform exact match between predicted and ground-truth labels, further support this intuition.

Compared to unsupervised baselines, SKET achieves better performance in both colon and cervix use-cases. In particular, the (relative) performance gap between SKET and baselines varies from 20% to 40% across measures. To confirm SKET effectiveness, we conducted a paired t-test and found that there is a statistical difference (p-value < 0.01) between its performance and that of the baselines on all the considered measures.

This outcome shows the effectiveness of introducing ad-hoc rules at both NER and EL levels, as well as the soundness of combining different matching techniques together. On the other hand, the unsupervised BERT-based approach outperforms both SKET and FastText in lung cancer. In this case, the paired t-test confirmed a statistical difference between BERT performance and that of SKET and FastText. Nevertheless, the performance gap between BERT and SKET never exceeds 5%. This highlights the robustness of SKET across different use-cases and makes it a viable solution in real scenarios, where annotated data are hard and expensive to get (such as in clinical practice). Besides, the lung cancer use-case presents two major differences with colon and cervix ones. First of all, lung annotation classes all revolve around different, but closely related, cancer types. As a consequence, contextualized NLMs (e.g., BERT [29]) – which are able to properly model the small semantic, contextual variations of such classes – achieve competitive results. Secondly, lung cancer data only consists of AOEC reports. The lack of RUMC reports makes the dataset more homogeneous and easier than the others, thus reducing classification inconsistencies for baseline models too.

Regarding weakly supervised models, the results reported in Table 4 demonstrate the effectiveness of using SKET to weakly annotate diagnostic reports and then train FastText and BERT models in a supervised fashion. In this regard, both weakly supervised FastText- and BERT-based approaches outperform their unsupervised counterparts. The only exception is for BERT on lung cancer data, where the unsupervised BERT approach achieves top performance. On the other hand, the weakly supervised BERT obtains the best results overall in both colon and cervix use-cases. Hence, SKET proves to be effective when used to bootstrap supervised models in absence of manual annotations. Following

Table 4

Text classification results on colon, cervix, and lung cancer pathology reports. The considered measures are subset accuracy, micro F1, and weighted F1. The [†] symbol represents the statistical difference of SKET from unsupervised FastText- and BERT-based approaches – verified using a paired t-test with a p-value < 0.01. **Bold** values represent the highest scores achieved for each measure.

Colon				
Approach	Model	Measures		
		Accuracy	Micro F1	Weighted F1
Unsupervised	SKET	0.7525 [†]	0.8386 [†]	0.8373 [†]
	FastText	0.4146	0.5298	0.5514
	BERT	0.5167	0.5697	0.6587
Weakly Supervised	FastText	0.7116	0.8287	0.8276
	BERT	0.7586	0.8432	0.8421
Cervix				
Approach	Model	Measures		
		Accuracy	Micro F1	Weighted F1
Unsupervised	SKET	0.5281 [†]	0.7791 [†]	0.7611 [†]
	FastText	0.2533	0.4882	0.4445
	BERT	0.3066	0.3962	0.4867
Weakly Supervised	FastText	0.4744	0.7542	0.7566
	BERT	0.5397	0.7901	0.7737
Lung				
Approach	Model	Measures		
		Accuracy	Micro F1	Weighted F1
Unsupervised	SKET	0.8137	0.8387	0.8262
	FastText	0.5221	0.7296	0.6853
	BERT	0.8523[†]	0.8630[†]	0.8526[†]
Weakly Supervised	FastText	0.7701	0.8313	0.8247
	BERT	0.8127	0.8375	0.8249

this procedure, supervised models can first be trained on data automatically annotated by SKET and then fine-tuned on small manually annotated batches, thus reducing annotation times and costs.

5. Understanding through Explainability













In recent years, the application of AI algorithms in the biomedical domain has experienced unprecedented growth [52, 53, 54] – especially to perform clinical decision support and diagnostic activities [55, 56, 57]. Therefore, there is an urgent need for eXplainable Artificial Intelligence (XAI) tools that can help clinicians and domain experts understand algorithm predictions and their underlying rationale. In this regard, explainability techniques highlight decision-relevant aspects of algorithms that contribute to specific predictions, thus trying to answer why a model has made a certain decision [58, 30, 59]. Hence, explainability methods are essential for humans – and in particular for clinicians – to decide whether to trust algorithm predictions and the (underlying) models that generated them [32, 33]. Among its different uses, explainability can be employed to understand the rationale of NER and EL outputs – such as the entity mentions and concepts identified by SKET within clinical reports.

However, since most of the data that humans can easily “visualize” regards objects restricted to the two/three-dimensional space, there is an urgent need not only for explainable models but also for explanation interfaces [30].

Dashboard

Checkout your pipelines' information below 

Your pipelines

Pipeline ID	Use case	Status	Start Timestamp	End Timestamp	Description	Params	Overview	Download
Filter...	Filter...	Filter...	Filter...	Filter...	Filter...			
9e393c1...	Colon	SUCCESS	2022-07-30T...	2022-07-30...	Processing of clinical...			Download 
691bcf1...	Colon	SUCCESS	2022-07-30T...	2022-07-30...	Processing of clinical...			Download 
835ac8a...	Colon	SUCCESS	2022-07-30T...	2022-07-30...	Processing of clinical...			Download 

Rows per page: 25 ▾ 1-4 of 4

Figure 2: SKET X dashboard providing information about the executed SKET pipelines - i.e., pipeline id, use case, pipeline status, start timestamp, end timestamp, description, pipeline parameters. Users can view the parameters of each pipeline by clicking on the dedicated button (A). Similarly, users can access pipelines data by clicking on the dedicated button (B). When the execution of a pipeline ends, its outputs become available for download (C).

To this end, we have developed SKET X,⁶ a web-based environment to interact with SKET and get useful insights about the extraction process and the related outputs. Through SKET X, pathologists and domain experts can visually comprehend SKET and the different components activated during the knowledge extraction process – thus getting a pointwise explanation of the outputs obtained for the provided diagnostic reports.

SKET X exploits Visual Analytics (VA) techniques to support domain experts in the visual comprehension of SKET outputs by means of intuitive and interactive interfaces. Such interfaces allow users to inspect and find out non-evident patterns in data and take decisions accordingly [60]. Specifically, VA techniques enable users to visually comprehend the results of an ongoing task, while it advances asynchronously in the background. Because of this, VA techniques are also used to visually adjust the parameters of a model running as a background task to continuously refine its outputs [61, 62].

5.1. SKET X Workflow

SKET X is an interactive Webapp that runs SKET on a set of uploaded reports. SKET X is based on SKET pipelines definable by the user who can customize the parameters and run SKET multiple times to compare the outputs and all the intermediate steps of the process. Each pipeline runs as an asynchronous task, handled by a scheduler with a queue manager. The pipelines are organized for straightforward access in the dashboard interface, shown in Figure 2. The dashboard provides information about the SKET pipelines executed by the users and enables access and download of the SKET outputs.

The execution of a SKET X pipeline consists at most of three *phases*, where the currently selected stage is shown on the top of the interface (see Figure 3.A).

- *Translation*: the reports are automatically translated from their original language to English. Figure 3 reports the information contained in the *Overview* tab of the interface, i.e. the inputs, outputs, and parameters of the translation phase.

⁶<http://w3id.org/sketx> access provided with username and password: demo

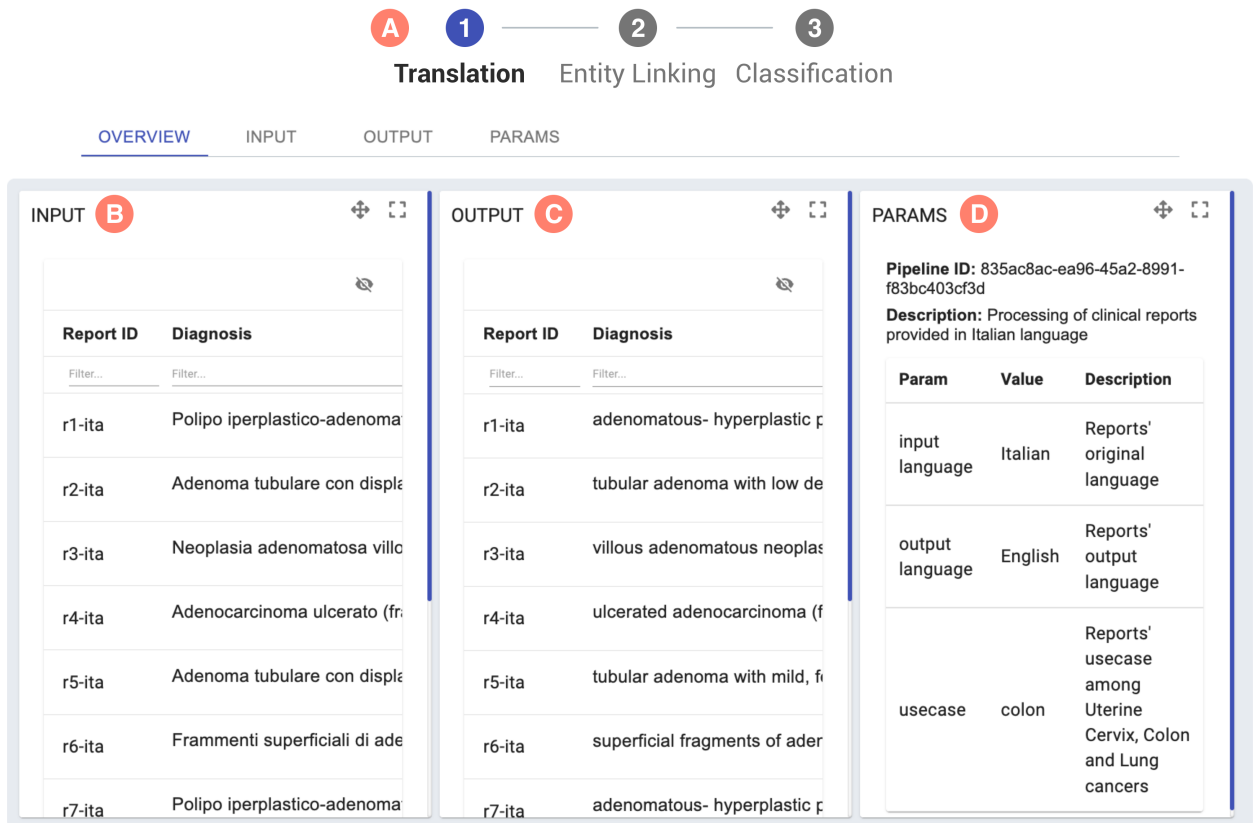


Figure 3: (A) SKET X *Overview* tab for the translation phase, (B) the reports in the original language (input), (C) the translated reports (output) (C), and (D) the parameters and settings for the current phase.

- *Entity Linking*: the entities automatically recognized within the reports are linked to the concepts in the ExaMode ontology. This phase's output consists of the identified mentions and the linked concepts. SKET employs a combination of hand-crafted rules and pre-trained neural models in this phase. The rules relevant to the disease of the given report are shown via a Sankey diagram, where the rules activated for the current report are highlighted. In this context, a rule is activated when one of the identified mentions – e.g., *low degree dysplasia (mild)* – satisfies one rule trigger – e.g., *dysplasia && mild* – that implies a link to a specific concept – e.g., *mild colon dysplasia* – as shown in Figure 4.
- *Classification*: SKET exploits mapping rules to decide the appropriate labels for each report. As for the EL phase, the rules relevant for the disease of the considered report are visualized using a Sankey diagram, where the activated rules are highlighted. A rule is activated when one of the identified concepts – e.g., *Mild Colon Dysplasia* – satisfies one rule trigger – e.g., *dysplasia && mild* – that implies a specific label – e.g., *Adenomatous polyp - low grade dysplasia* – as shown in Figure 5. The mentions and concepts considered in the classification task are regarded as *key mentions/concepts*, whereas the ones not satisfying any rule trigger are regarded as *excluded*, as shown in Figure 5.C and 5.D, respectively. For instance, in Figure 5 we can observe that the key concepts identified are *Colon Hyperplastic Polyp* and *Mild Colon Dysplasia*, whereas the excluded ones are *Biopsy of Colon* and *Colon, NOS* – both related to the same excluded mention *colon biopsy*.

5.2. SKET X Interface

The interface of SKET X consists of six tabs providing different views of the data, according to the selected phase:

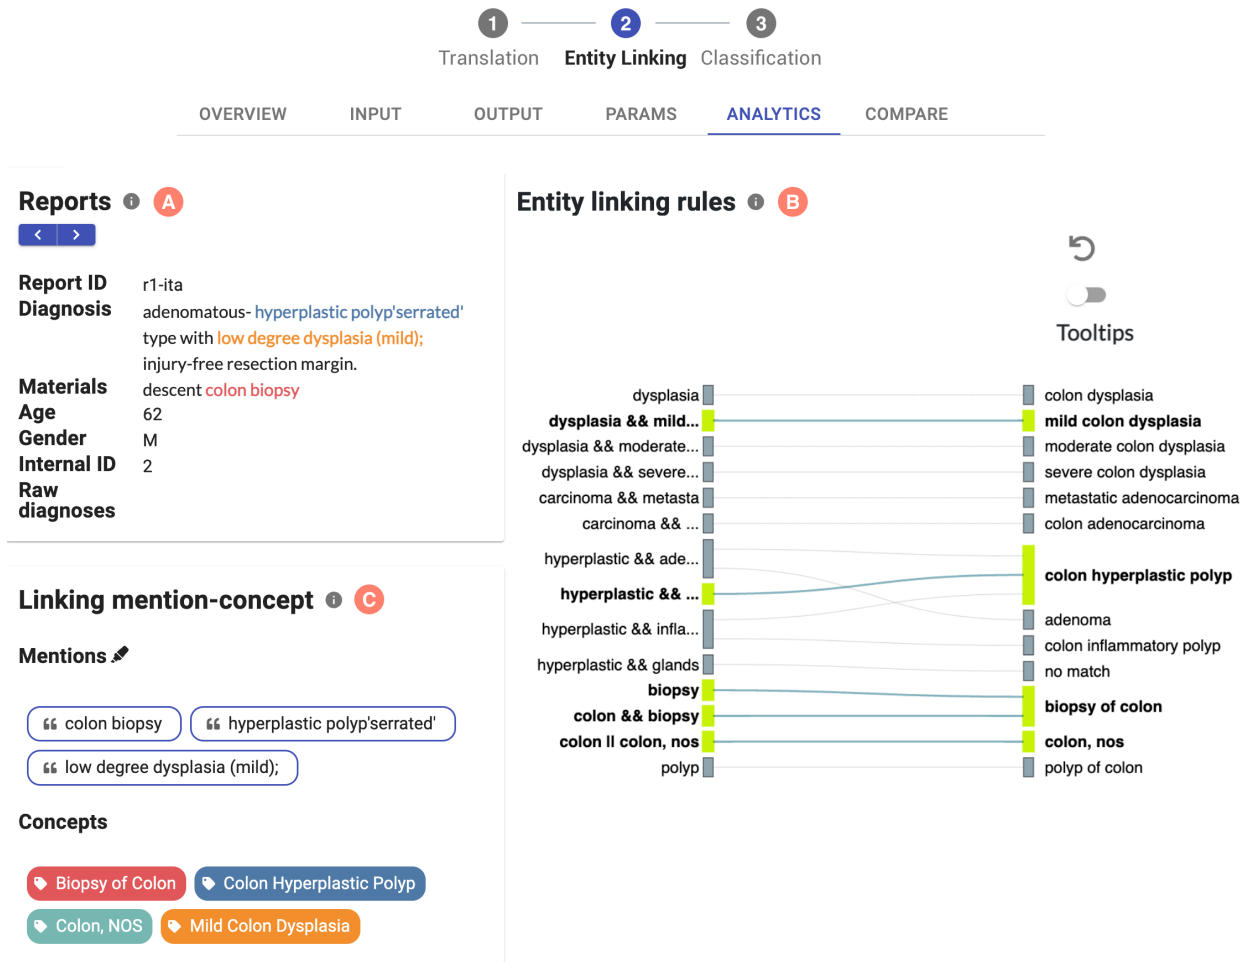


Figure 4: SKET X *Analytics* tab for the EL phase: (A) reports section, the users can change the current report using the left/right buttons; (B) SKET rules for the NER task; and, (C) list of mentions and concepts produced by the knowledge extraction process. Each concept and related mentions are highlighted with the same color in (A) and (C). By clicking/hovering on a specific concept, it is possible to highlight the relevant rules in the Sankey diagram that determined the concept and the related mentions in the report text. On the left side of the Sankey diagram are reported the rules *triggers*, which are boolean expressions tested on each mentioned text. If one or more mentions satisfy a rule trigger, then the related concepts on the right side of the Sankey diagram are highlighted and listed in (C).

- *Overview tab*: overview of the visual outputs available in the other tabs (i.e., *Input*, *Output*, *Params*, *Analytics*) for the current phase. The contents of the *Analytics* tab are shown in the overview only for the EL and classification phases.
- *Input tab*: it reports the input data for the current phase. For instance, if the considered phase is translation, this tab shows the reports in the original language, as shown in Figure 3.B. Instead, if the considered phase is EL, it shows the translated reports. Similarly, the mentions and the concepts extracted for each report are shown in this tab for the classification phase.
- *Output tab*: it reports the output data for the current phase. For translation, this tab shows the reports translated into English, as shown in Figure 3.C. Instead, if the considered phase is EL, it shows the mentions and the concepts extracted for each report. Similarly, the labels generated for each report are shown in this tab for the classification phase.



Figure 5: SKET X *Analytics* tab for the classification phase: (A) reports section to select the current report via left/right buttons; (B) SKET rules for determining the labels visualized with a Sankey diagram; and, (C) list of labels, mentions, and concepts determined by SKET. Each concept and the related mentions are highlighted with the same color in (A) and (C). The Sankey diagram highlights the relevant rules by clicking/hovering on a specific label. On the left side of the Sankey diagram are reported the rules *triggers*. If one or more concepts satisfy a rule trigger, then the related label is highlighted on the right side of the Sankey diagram and also listed in (C). The mentions and concepts involved in the classification task are the *key mentions/concepts* (C), while the *excluded* ones are reported in (D).

- *Params tab*: it reports the parameters for the current phase, as shown in Figure 3.D. For instance, for EL, it shows the methods and models used by SKET to perform the linking process between mentions and related concepts. Another important parameter is the *threshold* used by SKET in the pruning phase to reduce false positives and

thus increase precision, as described in section 3.2. When the phase considered is EL, users can change one or more parameters and then re-run SKET. This is useful to compare two pipelines using different parameters.

- Analytics tab:** it allows the users to analyze the current report's mentions, concepts, and labels in detail. In particular, if the considered phase is EL, users can inspect the identified mentions and concepts concerning the report textual content, as shown in Figures 4.A and 4.C. Moreover, by clicking on a mention, the user can inspect the list of associated concepts. At the same time, a user can also do the reverse - identifying the relevant mentions for a given concept. In addition, if the considered phase is classification, this tab shows the labels determined by SKET and the relations between a label and the concepts from which it derives, as shown in Figure 5. To visually explain the rules used by SKET to determine both the concepts and labels, a Sankey diagram is reported on the right side of the interface as depicted in Figure 4.B and 5.B. On the left side of the Sankey diagram, the rules *triggers* are reported, which are boolean expressions tested on the text of each mention – for the EL phase – and concept – for the classification phase. If one or more mentions/concepts satisfy a rule trigger, then the related concepts/labels on the right side of the Sankey diagram are highlighted.
- Compare tab:** it allows the users to compare the outputs of two different SKET X pipelines in terms of mentions, concepts, and labels identified for the current report. When the users click on the *compare* tab, they are provided with an initial menu that allows them to specify the two pipelines to compare. After the selection, users can click on the *compare* button to visualize the interface dedicated for the comparison, illustrated in Figure 6. The comparison interface is divided into four parts: (A) the reports section displaying information about the current report and two buttons for switching to the next/previous report; (B) the parameters section displaying pipeline information, such as the identifier, the description, and its parameters; (C) first pipeline section showing the outputs for the phase selected (e.g., mentions and concepts) and (D) second pipeline section with the same structure of (C). In particular, if the considered phase is EL, users can compare the concepts and the mentions identified by each pipeline and deduce which parameters have determined the major differences (e.g., the threshold for the NER task). Moreover, by clicking/hovering on each mention, the users can inspect the list of associated concepts (highlighted in different colors) among the two pipelines. On top of that, the common mentions between the two considered pipelines can be highlighted, thus making them and their related concepts easy to identify. Figure 6 shows the outputs of two SKET pipelines that have been executed with different models, where the first pipeline uses only the neural model while the second one uses only GPM. Since the two pipelines considered in Figure 6 use different models, they identify different concepts and mentions. Indeed, the common concepts between the two pipelines – i.e., Biopsy of Colon, Colon Hyperplastic Polyp, Colon NOS, and Mild Colon Dysplasia – have been identified using SKET rules, which are used in both pipelines. On the other hand, the disjoint concepts have been identified using the neural model – for Rectal mucous membrane – and GPM – for Adenoma and Resection – respectively. If the considered phase is classification, users can also compare the labels generated by each pipeline and the *key* concepts associated – that is, the ones from which the labels are determined. For instance, in Figure 7, we can observe that the labels generated by SKET are Adenomatous polyp - high grade dysplasia and Hyperplastic polyp for the first pipeline (C) while only Adenomatous polyp - high grade dysplasia for the second one (D). By clicking/hovering on the Hyperplastic polyp label, users can realize that it derives from the Colon Hyperplastic Polyp concept, which, in turn, is associated with the polyp sigmoid mention. Nevertheless, the latter mention does not suggest the presence of a Colon Hyperplastic Polyp. Thus it is a false positive. Similarly, users can do the same with the Adenomatous polyp - high-grade dysplasia label, discovering that it derives from the Severe Colon Dysplasia concept, which is correctly associated, through a SKET rule, with the severe dysplasia mention. Finally, users can also compare the *excluded* mentions and concepts that are not considered for the label generation process, but that can be a good indicator to determine whether the chosen threshold for models produces noisy concepts.

Hence, using SKET X pathologists and domain experts can visually comprehend why a certain concept/label has been extracted. Moreover, by leveraging both inspection and comparison functionalities, users can also understand the impact of different parameters on the obtained outputs, and thus investigate the advantages of combining ad hoc rules with ML models to improve the overall effectiveness of knowledge extraction systems.

1 ————— 2 ————— 3
Translation Entity Linking Classification

OVERVIEW INPUT OUTPUT PARAMS ANALYTICS **COMPARE**

A

Report ID: r1-ita

Diagnosis: adenomatous- hyperplastic polyp'serrated' type with low degree dysplasia (mild); injury-free resection margin.

Materials: descent colon biopsy

Age: 62

Gender: M

B

Pipeline	Description	Params
691bcf14-7f03-489b-8e1d-cb8a9025162a	Processing of the clinical reports without using GPM (only neural model + rules)	👁
9e393c1b-ba1b-4595-9403-21876f70fa31	Processing of the clinical reports without the neural model (only GPM + rules)	👁

C

Pipeline: 691bcf14-7f03-489b-8e1d-cb8a9025162a

Pipeline description: Processing of the clinical reports without using GPM (only neural model + rules)

Mentions

adenomatous- colon biopsy hyperplastic polyp'serrated' low degree dysplasia (mild);

Concepts

Biopsy of Colon Colon Hyperplastic Polyp Colon, NOS Mild Colon Dysplasia Rectal mucous membrane

D

Pipeline: 9e393c1b-ba1b-4595-9403-21876f70fa31

Pipeline description: Processing of the clinical reports without the neural model (only GPM + rules)

Mentions

adenomatous- colon biopsy hyperplastic polyp'serrated' injury-free resection margin. low degree dysplasia (mild);

Concepts

Adenoma Biopsy of Colon Colon Hyperplastic Polyp Colon, NOS Mild Colon Dysplasia Resection

Figure 6: SKET X *Compare* tab for the EL phase showing the comparison interface for the two pipelines specified for the comparison. The interface is organized in four parts: (A) the reports section displaying information about the current report and two buttons for switching to the next/previous report; (B) the parameters section displaying pipeline information, such as the identifier and its description, and parameters (e.g., the models used for EL phase and the threshold); (C) first pipeline outputs for the phase selected (e.g., mentions and concepts) and (D) second pipeline outputs for the phase selected. The mentions in common, and the related concepts, are highlighted both in the report text (A) and also in the mention/concept lists for each pipeline (C) and (D). Hence, we can observe that there is a mention *injury-free resection margin* and a concept *Resection* that are not highlighted since they have been identified only by the second pipeline (D). Nevertheless, the concepts *Rectal mucous membrane* and *Adenoma* have been identified only by respectively the first pipeline (C) and the second one (D), but since both are associated with the same common mention – i.e., *adenomatous* – they are highlighted as well.

6. Digital Pathology Applications

SKET has been integrated as a core system into different downstream applications for digital pathology. Figure 8 depicts the SKET ecosystem, where SKET UP represents the online access point to interact with SKET, SKET X provides explanations for SKET results, medTAG [35] integrates SKET automatic annotations to support semi-automatic tagging, and ExaNet [35] allows to visualize and explore SKET report-level knowledge graphs. Moreover, SKET labels can also be used to supervise cancer-assisted diagnosis tools [36].

6.1. Automatic Report Annotation

SKET has been integrated as an automatic annotator within MedTAG⁷ [35]. MedTAG is a collaborative biomedical annotation tool that provides four annotation types:

⁷MedTAG is available at <https://github.com/MedTAG/medtag-core/>

1 ————— 2 ————— 3
Translation Entity Linking Classification

OVERVIEW INPUT OUTPUT PARAMS ANALYTICS COMPARE

A

Report ID: r15-ita
Diagnosis: polyp sigmoid: tubular adenoma with severe dysplasia of the glandular epithelium. no evidence of infiltration of the connective vascular axis. free of resection margins.
Materials: polycy-caecum
Age: 74
Gender: M
Internal ID: 1

B

Pipeline	Description	Params
691bcf14-7f03-489b-8e1d-cb8a9025162a	Processing of the clinical reports without using GPM (only neural model + rules)	👁
9e393c1b-ba1b-4595-9403-21876f70fa31	Processing of the clinical reports without the neural model (only GPM + rules)	👁

C

Classification ⓘ

Pipeline: 691bcf14-7f03-489b-8e1d-cb8a9025162a

Pipeline description: Processing of the clinical reports without using GPM (only neural model + rules)

Labels

Adenomatous polyp - high grade dysplasia # Hyperplastic polyp

Key Mentions ✎

polyp sigmoid: severe dysplasia

Key Concepts

Colon Hyperplastic Polyp Severe Colon Dysplasia

Excluded Mentions ✎

glandular epithelium. polycy-caecum tubular adenoma

Excluded Concepts

Colon Inflammatory Polyp Colon Tubular Adenoma Rectal mucous membrane

D

Classification ⓘ

Pipeline: 9e393c1b-ba1b-4595-9403-21876f70fa31

Pipeline description: Processing of the clinical reports without the neural model (only GPM + rules)

Labels

Adenomatous polyp - high grade dysplasia

Key Mentions ✎

severe dysplasia

Key Concepts

Severe Colon Dysplasia

Excluded Mentions ✎

glandular epithelium. polycy-caecum polyp sigmoid: tubular adenoma

Excluded Concepts

Caecum Colon Tubular Adenoma Sigmoid colon

Figure 7: SKET X *Compare* tab for the classification phase showing the comparison interface for the two pipelines specified for the comparison. The interface is organized in four parts: (A) the reports section displaying information about the current report and two buttons for switching to the next/previous report; (B) the parameters section displaying pipeline information, such as the identifier, the description, and its parameters; (C) first pipeline outputs for the phase selected (e.g., mentions and concepts) and (D) second pipeline outputs for the phase selected. The mentions/concepts considered for determining the report labels are regarded as *key mentions/concepts* and are differentiated by the *excluded* ones. Here, two concepts are identified in the first pipeline, namely, *Colon Hyperplastic Polyp* and *Severe Colon Dysplasia*, while in the second one only *Severe Colon Dysplasia* has been identified. Nevertheless, *Colon Hyperplastic Polyp* and *Sigmoid colon* are negligible concepts (i.e., false positives) both associated with the *polyp sigmoid* mention. In contrast, *Severe Colon Dysplasia* is correct since it has been identified using a SKET rule verified by the *severe dysplasia* key mention.

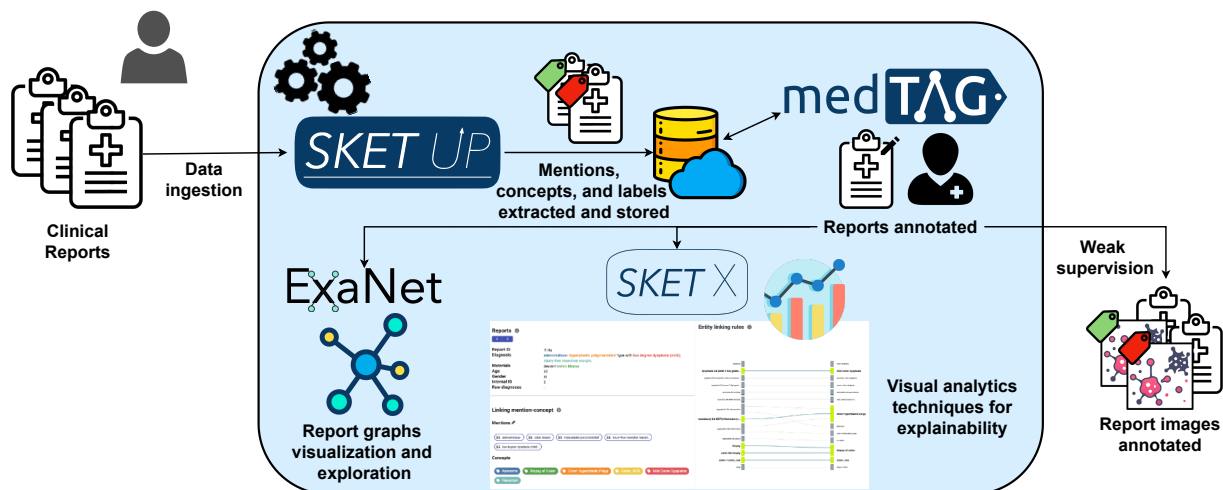


Figure 8: The SKET ecosystem. From clinical reports, a suite of different applications relying on SKET to process (SKET UP), analyze and annotate (medTAG), explore (ExaNET), and explain (SKET X) the knowledge contained within reports – also providing weak supervision to train cancer assisted diagnosis tools.

Table 5

Number of labels, concepts, mentions, and links automatically annotated by SKET within MedTAG. Statistics are reported for each use-case and globally.

Annotation Type	Colon	Cervix	Lung	Total
Labels	9,309	16,033	2,066	27,408
Concepts	11,932	12,936	2,336	27,204
Mentions	10,926	12,070	2,336	25,332
Linking	11,932	12,936	2,336	27,204
Total	44,099	53,975	9,074	107,148

- **Labels:** allows the user to assign, by clicking on the check-boxes, one or more labels to a document. The labels indicate some reports' properties (e.g. "Cancer" label indicates the presence of a cancer-related disease).
- **Concepts:** allows the user to specify which concepts are relevant for a document. Users can take advantage of auto-complete functionalities for searching the relevant concepts to assign to each document.
- **Mentions:** shows the list of the mentions identified by the user in the report text.
- **Linking:** allows the user to link the mentions identified with the corresponding concepts. Users can link the same mention to multiple concepts.

For each annotation type, SKET provides automatic annotations for reports associated with colon, cervix, and lung use-cases. At present, MedTAG has been used by experts to produce more than 7,000 annotations. On the other hand, SKET annotations within MedTAG exceed 100,000 units. Table 5 reports SKET annotation statistics for each annotation type.

6.2. Pathological Knowledge Visualization

The report-level knowledge graphs produced by SKET can be explored with ExaNet⁸. ExaNet is a visual application that allows users (e.g., experts and pathologists) to explore the pathology reports linked data by using an interactive graph visualization tool. ExaNet enables users to explore graph connections by leveraging pan and zoom functionalities. On top of this, ExaNet allows users to visualize an interactive JSON serialization of each pathology report, providing also download capabilities.

⁸ExaNet can be accessed through the "Reports' stats" functionality of MedTAG, under the "Graph" feature associated with each report that has been annotated by SKET.

Table 6

CNN colon cancer performance when trained with SKET weak labels (CNN-SKET) and with manual ones (CNN-GT). Results refer to WSI classification on AOEC and RUMC data. For each considered measure, we report the average obtained through 10-fold cross-validation. **Bold** values represent the highest scores achieved for each measure.

Model	Accuracy	Micro F1	Weighted F1
CNN-SKET	0.6666	0.7741	0.7694
CNN-GT	0.6795	0.7866	0.7800

Conceptually, ExaNet stems from ontology visualization tools. The visualization of ontologies is a fundamental task to assess ontologies and enable users to explore, verify, and understand them and their underlying structures [63, 64, 65, 66]. Nevertheless, compared to ontology visualization, where the focus is primarily on the Terminological Box (TBox) – i.e., definition of classes and properties – ExaNet focuses instead on the Assertional Box (ABox) – that is, individuals and instance data. Furthermore, ExaNet replaces the classes Internationalized Resource Identifier (IRI) with the corresponding literals.

6.3. WSI Classification

The labels produced by SKET are used to reduce supervised-training limitations for colon cancer assisted diagnosis tools [36] – limitations that prevent the full exploitation of digital pathology applications. In other words, SKET labels serve as weak labels to train a deep image classifier.

The proposed model, based on Multiple Instance Learning Convolutional Neural Networks (CNNs), makes multi-class predictions at patch-level and then aggregates them through an attention pooling layer [67, 68] to obtain multi-label WSI predictions. The multi-label setting reflects the very nature of the pathology domain, where images (and reports) can highlight multiple findings for the same sample. Therefore, employing models that produce multi-label predictions allows to better approximate real-world pathology scenarios.

The proposed approach has been trained and tested using data composed of colon WSIs from AOEC and RUMC medical centers. The training set consists of the WSIs associated with the 3769 colon reports reported in Table 1, whereas the test set consists of 227 WSIs from AOEC and 423 from RUMC, for a total of 650 WSIs. Colon cancer was chosen as use-case due to its high social impact and difficulty in diagnosing it. In fact, colon cancer is the fourth most diagnosed cancer in the world [69]. Besides, the need to identify malignant polyps – which are cell agglomerations protruding from the colon surface – makes it problematic to diagnose [69]. Thus, to prove the effectiveness of SKET as a weak annotator, we compared the performance of the image classifier trained with SKET labels against its performance when trained using manual labels. Table 6 reports the results for subset accuracy, micro-, and weighted-average F1 measures, obtained through 10-fold cross-validation.

The obtained results show the effectiveness of SKET when used as a weak annotator. The performance obtained using weak labels are close to those achieved with manual ones. Precisely, the performance difference between the two CNNs does not exceed 1.3%. Furthermore, we performed the Wilcoxon Rank-Sum test and verified that such performance difference is not statistically significant (p -value < 0.05). Thus, SKET allows training cancer diagnostics models for digital pathology without human intervention, paving the way to the use of ML models in the clinical practice.

7. Conclusions and Future Work

In this work, we presented the Semantic Knowledge Extractor Tool (SKET), an unsupervised hybrid knowledge extraction system that combines rule-based techniques with pre-trained Machine Learning (ML) models to extract critical pathological concepts from diagnostic reports. The concepts extracted from diagnostic reports can serve different digital pathology applications, such as automatic annotation, knowledge visualization, discovery, or image classification. A throughout evaluation demonstrated SKET effectiveness in annotating colon, cervix, and lung cancer use-cases – making it a viable solution to reduce pathologists' workload. The results and analyses highlighted the importance of expert knowledge in developing unsupervised systems for specialized medicine. Moreover, the effectiveness of SKET as a weak annotator suggests that it can be used as a first, cheap solution to bootstrap supervised models in the absence of manual annotations.

Together with SKET, we also introduced SKET eXplained (SKET X), a web-based system to support pathologists and domain experts in the understanding of SKET outputs, as well as the role that its different components have on such outputs. Through SKET X, users can comprehend predictions and get valuable insights into the knowledge extraction process. Beyond explainability, SKET has also been used to empower different digital pathology downstream applications. In particular, SKET labels have been used to reduce training limitations for colon cancer-assisted diagnosis tools. The use of SKET for training deep image classifiers without human intervention paves the way to ML models in the clinical practice [36].

As future work, we plan to extend SKET to other emerging but under-researched use-cases, such as celiac disease – whose prevalence has significantly increased over the past 20 years [70].

References

- [1] T. Davenport and R. Kalakota, “The Potential for Artificial Intelligence in Healthcare,” *Future Healthc J.*, vol. 6, no. 2, pp. 94–98, 2019. [Online]. Available: <https://doi.org/10.7861/futurehosp.6-2-94>
- [2] J. M. Buckley, S. B. Coopey, J. Sharko, F. Polubriaginof, B. Drohan, A. K. Belli, E. M. Kim, J. E. Garber, B. L. Smith, M. A. Gadd, M. C. Specht, C. A. Roche, T. M. Gudewicz, and K. S. Hughes, “The Feasibility of Using Natural Language Processing to Extract Clinical Information from Breast Pathology Reports,” *J. Pathol Inform.*, vol. 3, no. 1, p. 23, 2012. [Online]. Available: <https://doi.org/10.4103/2153-3539.97788>
- [3] S. Hassanpour and C. P. Langlotz, “Information Extraction from Multi-Institutional Radiology Reports,” *Artif. Intell. Medicine*, vol. 66, pp. 29–39, 2016. [Online]. Available: <https://doi.org/10.1016/j.artmed.2015.09.007>
- [4] G. Burger, A. Abu-Hanna, N. de Keizer, and R. Cornet, “Natural Language Processing in Pathology: a Scoping Review,” *Journal of Clinical Pathology*, vol. 69, no. 11, pp. 949–955, 2016. [Online]. Available: <https://doi.org/10.1136/jclinpath-2016-203872>
- [5] M. Topaz, L. Murga, K. M. Gaddis, M. V. McDonald, O. Bar-Bachar, Y. Goldberg, and K. H. Bowles, “Mining Fall-Related Information in Clinical Notes: Comparison of Rule-Based and Novel Word Embedding-Based Machine Learning Approaches,” *J. Biomed. Informatics*, vol. 90, 2019. [Online]. Available: <https://doi.org/10.1016/j.jbi.2019.103103>
- [6] T. Oliwa, S. B. Maron, L. M. Chase, S. Lomnicki, D. V. T. Catenacci, B. Furner, and S. L. Volchenboum, “Obtaining Knowledge in Pathology Reports Through a Natural Language Processing Approach With Classification, Named-Entity Recognition, and Relation-Extraction Heuristics,” *JCO Clinical Cancer Informatics*, vol. 1, no. 3, pp. 1–8, 2019. [Online]. Available: <https://doi.org/10.1200/CCI.19.00008>
- [7] K. Kreimeyer, M. Foster, A. Pandey, N. Arya, G. Halford, S. F. Jones, R. Forshee, M. Walderhaug, and T. Botsis, “Natural Language Processing Systems for Capturing and Standardizing Unstructured Clinical Information: A Systematic Review,” *J. Biomed. Informatics*, vol. 73, pp. 14–29, 2017. [Online]. Available: <https://doi.org/10.1016/j.jbi.2017.07.012>
- [8] Y. Wang, L. Wang, M. Rastegar-Mojarad, S. Moon, F. Shen, N. Afzal, S. Liu, Y. Zeng, S. Mehrabi, S. Sohn, and H. Liu, “Clinical Information Extraction Applications: A Literature Review,” *J. Biomed. Informatics*, vol. 77, pp. 34–49, 2018. [Online]. Available: <https://doi.org/10.1016/j.jbi.2017.11.011>
- [9] E. Santus, T. Schuster, A. M. Tahmasebi, C. Li, A. Yala, C. R. Lanahan, P. Prinsen, S. F. Thompson, S. Coons, L. Mynderse, R. Barzilay, and K. Hughes, “Exploiting Rules to Enhance Machine Learning in Extracting Information From Multi-Institutional Prostate Pathology Reports,” *JCO Clinical Cancer Informatics*, no. 4, pp. 865–874, 2020. [Online]. Available: <https://doi.org/10.1200/CCI.20.00028>
- [10] Y. Kim, J. H. Lee, S. Choi, J. M. Lee, J. H. Kim, J. Seok, and H. J. Joo, “Validation of Deep Learning Natural Language Processing Algorithm for Keyword Extraction from Pathology Reports in Electronic Health Records,” *Sci Rep*, vol. 1, no. 10, pp. 1–9, 2020. [Online]. Available: <https://doi.org/10.1038/s41598-020-77258-w>
- [11] P. Giannaris, Z. Al-Taie, M. Kovalenko, N. Thanintorn, O. Kholod, Y. Innokenteva, E. Coberly, S. Frazier, K. Laziuk, M. Popescu, C. R. Shyu, D. Xu, R. Hammer, and D. Shin, “Artificial Intelligence-Driven Structurization of Diagnostic Information in Free-Text Pathology Reports,” *Journal of Pathology Informatics*, vol. 11, no. 1, p. 10, 2020. [Online]. Available: https://doi.org/10.4103/jpi.jpi_30_19
- [12] J. R. Gregg, M. Lang, L. L. Wang, M. J. Resnick, S. K. Jain, J. L. Warner, and D. A. Barocas, “Automating the Determination of Prostate Cancer Risk Strata From Electronic Medical Records,” *JCO Clinical Cancer Informatics*, vol. 1, no. 1, pp. 1–8, 2017. [Online]. Available: <https://doi.org/10.1200/CCI.16.00045>
- [13] A. P. Glaser, B. J. Jordan, J. Cohen, A. Desai, P. Silberman, and J. J. Meeks, “Automated Extraction of Grade, Stage, and Quality Information From Transurethral Resection of Bladder Tumor Pathology Reports Using Natural Language Processing,” *JCO Clinical Cancer Informatics*, vol. 1, no. 2, pp. 1–8, 2018. [Online]. Available: <https://doi.org/10.1200/CCI.17.00128>
- [14] K. Roberts, D. Demner-Fushman, E. M. Voorhees, W. R. Hersh, S. Bedrick, A. J. Lazar, and S. Pant, “Benchmarking Information Retrieval for Precision Oncology: the TREC Precision Medicine Track,” in *AMIA 2018, American Medical Informatics Association Annual Symposium, San Francisco, CA, November 3-7, 2018*. AMIA, 2018. [Online]. Available: <https://knowledge.amia.org/67852-amia-1.4259402/t006-1.4263223/t006-1.4263224/2976780-1.4263306/2970178-1.4263303>
- [15] V. Cheplygina, M. de Bruijne, and J. P. W. Pluim, “Not-so-supervised: A Survey of Semi-Supervised, Multi-Instance, and Transfer Learning in Medical Image Analysis,” *Medical Image Anal.*, vol. 54, pp. 280–296, 2019. [Online]. Available: <https://doi.org/10.1016/j.media.2019.03.009>
- [16] F. Aeffner, K. Wilson, N. T. Martin, J. C. Black, C. L. L. Hendriks, B. Bolon, D. G. Rudmann, R. Gianani, S. R. Koegler, J. Krueger, and G. D. Young, “The gold standard paradox in digital image analysis: Manual versus automated scoring as ground truth,” *Archives of pathology & laboratory medicine*, vol. 141, no. 9, pp. 1267–1275, 2017. [Online]. Available: <https://doi.org/10.5858/arpa.2016-0386-RA>
- [17] O. J. del Toro, S. Otálora, M. Andersson, K. Eurén, M. Hedlund, M. Rousson, H. Müller, and M. Atzori, “Analysis of Histopathology Images: From Traditional Machine Learning to Deep Learning,” in *Biomedical Texture Analysis*. Academic Press, 2017, pp. 281–314. [Online]. Available: <https://doi.org/10.1016/B978-0-12-812133-7.00010-7>

- [18] B. E. Bejnordi, M. Veta, P. J. V. Diest, B. V. Ginneken, N. Karssemeijer, G. Litjens, J. A. W. M. V. D. Laak, M. Hermesen, Q. F. Manson, and M. Balkenhol, "Diagnostic Assessment of Deep Learning Algorithms for Detection of Lymph Node Metastases in Women with Breast Cancer," *JAMA*, vol. 318, no. 22, pp. 2199–2210, 2017. [Online]. Available: <https://doi.org/10.1001/jama.2017.14585>
- [19] A. J. Schaumberg, W. Juarez, S. J. Choudhury, L. G. Pastríán, B. S. Pritt, M. P. Pozuelo, R. S. Sánchez, K. Ho, N. Zahra, and B. D. Sener, "Large-Scale Annotation of Histopathology Images from Social Media," *BioRxiv*, p. 396663, 2018. [Online]. Available: <https://doi.org/10.1101/396663>
- [20] P. Courtiol, E. W. Tramel, M. Sanselme, and G. Wainrib, "Classification and Disease Localization in Histopathology Using Only Global Labels: A Weakly-Supervised Approach," *CoRR*, vol. abs/1802.02212, 2018. [Online]. Available: <http://arxiv.org/abs/1802.02212>
- [21] D. Komura and S. Ishikawa, "Machine Learning Methods for Histopathological Image Analysis," *Computational and Structural Biotechnology Journal*, vol. 16, pp. 34–42, 2018. [Online]. Available: <https://doi.org/10.1016/j.csbj.2018.01.001>
- [22] G. Campanella, M. G. Hanna, L. Geneslaw, A. Mirafior, V. W. K. Silva, K. J. Busam, E. Brogi, V. E. Reuter, D. S. Klimstra, and T. J. Fuchs, "Clinical-Grade Computational Pathology using Weakly Supervised Deep Learning on Whole Slide Images," *Nat Med*, vol. 25, pp. 1301–1309, 2019. [Online]. Available: <https://doi.org/10.1038/s41591-019-0508-1>
- [23] M. A. Carbonneau, V. Cheplygina, E. Granger, and G. Gagnon, "Multiple Instance Learning: A Survey of Problem Characteristics and Applications," *Pattern Recognit.*, vol. 77, pp. 329–353, 2018. [Online]. Available: <https://doi.org/10.1016/j.patcog.2017.10.009>
- [24] A. Dhrangadhariya, S. Otálora, M. Atzori, and H. Müller, "Classification of noisy free-text prostate cancer pathology reports using natural language processing," in *ICPR, Artificial Intelligence for Digital Pathology Workshop (AIDP)*, 2021. [Online]. Available: https://doi.org/10.1007/978-3-030-68763-2_12
- [25] L. Chiticariu, Y. Li, and F. R. Reiss, "Rule-Based Information Extraction is Dead! Long Live Rule-Based Information Extraction Systems!" in *Proc. of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA*. ACL, 2013, pp. 827–832. [Online]. Available: <https://aclanthology.org/D13-1079/>
- [26] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed Representations of Words and Phrases and their Compositionality," in *Proc. of the 27th Annual Conference on Neural Information Processing Systems 2013, NIPS, Lake Tahoe, Nevada, United States, December 5-8, 2013*, 2013, pp. 3111–3119. [Online]. Available: <https://proceedings.neurips.cc/paper/2013/hash/9aa42b31882ec039965f3c4923ce901b-Abstract.html>
- [27] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching Word Vectors with Subword Information," *Trans. Assoc. Comput. Linguistics*, vol. 5, pp. 135–146, 2017. [Online]. Available: https://doi.org/10.1162/tacl_a_00051
- [28] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep Contextualized Word Representations," in *Proc. of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018*. ACL, 2018, pp. 2227–2237. [Online]. Available: <https://doi.org/10.18653/v1/n18-1202>
- [29] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *CoRR*, vol. abs/1810.04805, 2018. [Online]. Available: <http://arxiv.org/abs/1810.04805>
- [30] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, and H. Müller, "Causability and explainability of artificial intelligence in medicine," *WIREs Data Mining Knowl. Discov.*, vol. 9, no. 4, 2019. [Online]. Available: <https://doi.org/10.1002/widm.1312>
- [31] A. Holzinger, "From machine learning to explainable ai," *2018 World Symposium on Digital Intelligence for Systems and Machines (DISA)*, pp. 55–66, 2018. [Online]. Available: <https://doi.org/10.1109/DISA.2018.8490530>
- [32] EU AI HLEG, "A Definition of Artificial Intelligence: Main Capabilities and Scientific Disciplines," <https://digital-strategy.ec.europa.eu/en/library/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines>, 2019, accessed: 2022-09-05.
- [33] —, "Ethics Guidelines for Trustworthy AI," <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>, 2019, accessed: 2022-09-05.
- [34] J. Amann, A. Blasimme, E. Vayena, D. Frey, V. I. Madai, and the Precise4Q consortium, "Explainability for artificial intelligence in healthcare: a multidisciplinary perspective," *BMC Medical Informatics and Decision Making*, vol. 20, no. 1, p. 310, 2020. [Online]. Available: <https://doi.org/10.1186/s12911-020-01332-6>
- [35] F. Giachelle, O. Irrera, and G. Silvello, "MedTAG: a portable and customizable annotation tool for biomedical documents," *BMC Medical Informatics Decis. Mak.*, vol. 21, no. 1, p. 352, 2021. [Online]. Available: <https://doi.org/10.1186/s12911-021-01706-4>
- [36] N. Marini, S. Marchesin, S. Otálora, M. Wodzinski, A. Caputo, M. van Rijthoven, W. Aswolinskiy, J. M. Bokhorst, D. Podareanu, E. Petters, S. Boytcheva, G. Buttafuoco, S. Vatrano, F. Frassetto, J. der Laak, M. Agosti, F. Ciompi, G. Silvello, H. Muller, and M. Atzori, "Unleashing the potential of digital pathology data by training computer-aided diagnosis models without human annotations," *npj Digital Medicine*, vol. 5, no. 1, Jul 2022. [Online]. Available: <http://dx.doi.org/10.1038/s41746-022-00635-4>
- [37] J. R. Srigley, T. McGowan, A. Maclean, M. Raby, J. Ross, S. Kramer, and C. Sawka, "Standardized synoptic cancer pathology reporting: a population-based approach," *J. Surg. Oncol.*, vol. 99, no. 8, pp. 517–524, 2009. [Online]. Available: <https://doi.org/10.1002/jso.21282>
- [38] D. W. Ellis and J. Srigley, "Does standardised structured reporting contribute to quality in diagnostic pathology? the importance of evidence-based datasets," *Virchows Arch.*, vol. 468, no. 1, 2016. [Online]. Available: <https://doi.org/10.1007/s00428-015-1834-4>
- [39] M. Junczys-Dowmunt, R. Grundkiewicz, T. Dwojak, H. Hoang, K. Heafield, T. Neckermann, F. Seide, U. Germann, A. F. Aji, N. Bogoychev, A. F. T. Martins, and A. Birch, "Marian: Fast neural machine translation in C++," in *Proc. of ACL 2018, Melbourne, Australia, July 15-20, 2018, System Demonstrations*. ACL, 2018, pp. 116–121. [Online]. Available: <https://aclanthology.org/P18-4020/>
- [40] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is All you Need," in *Proc. of the 30th Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 2017, pp. 5998–6008. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>
- [41] S. Marchesin, "Case-Based Retrieval Using Document-Level Semantic Networks," in *Proc. of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*. ACM, 2018, p. 1451. [Online]. Available: <https://doi.org/10.1145/3209978.3210221>

- [42] M. Neumann, D. King, I. Beltagy, and W. Ammar, "ScispaCy: Fast and Robust Models for Biomedical Natural Language Processing," in *Proc. of the 18th BioNLP Workshop and Shared Task, BioNLP@ACL 2019, Florence, Italy, August 1, 2019*. ACL, 2019, pp. 319–327. [Online]. Available: <https://doi.org/10.18653/v1/w19-5034>
- [43] S. Pyysalo, F. Ginter, H. Moen, T. Salakoski, and S. Ananiadou, "Distributional Semantics Resources for Biomedical Text Processing," *Proc. of LBM*, pp. 39–44, 2013. [Online]. Available: <https://bio.nlpplab.org/pdf/pyysalo13literature.pdf>
- [44] W. W. Chapman, W. Bridewell, P. Hanbury, G. F. Cooper, and B. G. Buchanan, "A Simple Algorithm for Identifying Negated Findings and Diseases in Discharge Summaries," *J. Biomed. Informatics*, vol. 34, no. 5, pp. 301–310, 2001. [Online]. Available: <https://doi.org/10.1006/jbin.2001.1029>
- [45] J. W. Ratcliff and D. E. Metzner, "Pattern Matching: the Gestalt Approach," *Dr Dobbs Journal*, vol. 13, no. 7, p. 46, 1988. [Online]. Available: <https://www.drdoobbs.com/database/pattern-matching-the-gestalt-approach/184407970>
- [46] J. A. Shaw and E. A. Fox, "Combination of Multiple Searches," in *Proc. of The Third Text REtrieval Conference, TREC 1994, Gaithersburg, Maryland, USA, November 2-4, 1994*, ser. NIST Special Publication, vol. 500-225. NIST, 1994, pp. 105–108. [Online]. Available: <http://trec.nist.gov/pubs/trec3/papers/vt.ps.gz>
- [47] M. Agosti, S. Marchesin, and G. Silvello, "Learning Unsupervised Knowledge-Enhanced Representations to Reduce the Semantic Gap in Information Retrieval," *ACM Trans. Inf. Syst.*, vol. 38, no. 4, pp. 38:1–38:48, 2020. [Online]. Available: <https://doi.org/10.1145/3417996>
- [48] W. C. Chang, H. F. Yu, K. Zhong, Y. Yang, and I. S. Dhillon, "Taming pretrained transformers for extreme multi-label text classification," in *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*. ACM, 2020, pp. 3163–3171. [Online]. Available: <https://doi.org/10.1145/3394486.3403368>
- [49] P. Ruas, V. D. T. Andrade, and F. M. Couto, "Lasige-biotm at MESINESP2: entity linking with semantic similarity and extreme multi-label classification on spanish biomedical documents," in *Proc. of the Working Notes of CLEF 2021 - Conference and Labs of the Evaluation Forum, Bucharest, Romania, September 21st - to - 24th, 2021*, ser. CEUR Workshop Proceedings, vol. 2936. CEUR-WS.org, 2021, pp. 324–334. [Online]. Available: <http://ceur-ws.org/Vol-2936/paper-24.pdf>
- [50] Y. Zhang, Q. Chen, Z. Yang, H. Lin, and Z. Lu, "BioWordVec, Improving Biomedical Word Embeddings with Subword Information and MeSH," *Scientific Data*, vol. 6, no. 52, pp. 1–9, 2019. [Online]. Available: <https://doi.org/10.1038/s41597-019-0055-0>
- [51] E. Alsentzer, J. R. Murphy, W. Boag, W. H. Weng, D. Jin, T. Naumann, and M. B. A. McDermott, "Publicly Available Clinical BERT Embeddings," *CoRR*, vol. abs/1904.03323, 2019. [Online]. Available: <http://arxiv.org/abs/1904.03323>
- [52] S. Marchesin and G. Silvello, "TBGA: a large-scale gene-disease association dataset for biomedical relation extraction," *BMC Bioinform.*, vol. 23, no. 1, p. 111, 2022. [Online]. Available: <https://doi.org/10.1186/s12859-022-04646-6>
- [53] K. Dzobo, S. Adotey, N. E. Thomford, and W. Dzobo, "Integrating artificial and human intelligence: A partnership for responsible innovation in biomedical engineering and medicine." *Omics : a journal of integrative biology*, 2019. [Online]. Available: <https://doi.org/10.1089/omi.2019.0038>
- [54] F. Wang and A. M. Preininger, "Ai in health: State of the art, challenges, and future directions," *Yearbook of Medical Informatics*, vol. 28, pp. 16 – 26, 2019. [Online]. Available: <https://doi.org/10.1055/s-0039-1677908>
- [55] R. T. Sutton, D. Pincock, D. C. Baumgart, D. C. Sadowski, R. N. Fedorak, and K. I. Kroeker, "An overview of clinical decision support systems: benefits, risks, and strategies for success," *npj Digital Medicine*, vol. 3, 2020. [Online]. Available: <https://doi.org/10.1038/s41746-020-0221-y>
- [56] F. Magrabi, E. Ammenwerth, J. B. McNair, N. F. de Keizer, H. Hyppönen, P. Nykänen, M. L. Rigby, P. J. Scott, T. Vehko, Z. S.-Y. Wong, and A. Georgiou, "Artificial intelligence in clinical decision support: Challenges for evaluating ai and practical implications," *Yearbook of Medical Informatics*, vol. 28, pp. 128 – 134, 2019. [Online]. Available: <https://doi.org/10.1055/s-0039-1677903>
- [57] S. Montani, "Exploring new roles for case-based reasoning in heterogeneous AI systems for medical decision support," *Appl. Intell.*, vol. 28, no. 3, pp. 275–285, 2008. [Online]. Available: <https://doi.org/10.1007/s10489-007-0046-2>
- [58] A. Holzinger, "Explainable AI and multi-modal causability in medicine," *i-com*, vol. 19, no. 3, pp. 171–179, 2021. [Online]. Available: <https://doi.org/10.1515/icom-2020-0024>
- [59] A. Holzinger, C. Biemann, C. S. Pattichis, and D. B. Kell, "What do we need to build explainable AI systems for the medical domain?" *CoRR*, vol. abs/1712.09923, 2017. [Online]. Available: <http://arxiv.org/abs/1712.09923>
- [60] J. J. Thomas and K. A. Cook, *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. United States, Department of Homeland Security, 2005. [Online]. Available: <https://www.hsdl.org/?view&did=485291>
- [61] M. Angelini, G. Santucci, H. Schumann, and H. Schulz, "A review and characterization of progressive visual analytics," *Informatics*, vol. 5, no. 3, p. 31, 2018. [Online]. Available: <https://doi.org/10.3390/informatics5030031>
- [62] F. Giachelle and G. Silvello, "A progressive visual analytics tool for incremental experimental evaluation," in *Proc. of the 10th Italian Information Retrieval Workshop, Padova, Italy, September 16-18, 2019*, 2019, pp. 2–5. [Online]. Available: <http://ceur-ws.org/Vol-2441/paper17.pdf>
- [63] S. Lohmann, S. Negru, F. Haag, and T. Ertl, "Visualizing ontologies with VOWL," *Semantic Web*, vol. 7, no. 4, pp. 399–419, 2016. [Online]. Available: <https://doi.org/10.3233/SW-150200>
- [64] S. Lohmann, V. Link, E. Marbach, and S. Negru, "Webvowl: Web-based visualization of ontologies," in *Proc. of the Knowledge Engineering and Knowledge Management - EKAW 2014 Satellite Events, VISUAL, EKM1, and ARCOE-Logic, Linköping, Sweden, November 24-28, 2014. Revised Selected Papers*, ser. LNCS, vol. 8982. Springer, 2014, pp. 154–158. [Online]. Available: https://doi.org/10.1007/978-3-319-17966-7_21
- [65] S. Lohmann, S. Negru, F. Haag, and T. Ertl, "VOWL 2: User-oriented visualization of ontologies," in *Proc. of the 19th International Conference on Knowledge Engineering and Knowledge Management EKAW 2014, Linköping, Sweden, November 24-28, 2014*, ser. LNCS, vol. 8876. Springer, 2014, pp. 266–281. [Online]. Available: https://doi.org/10.1007/978-3-319-13704-9_21
- [66] M. Lanzemberger, J. Sampson, and M. Rester, "Visualization in ontology tools," in *Proc. of the 2009 International Conference on Complex, Intelligent and Software Intensive Systems, CISIS 2009, Fukuoka, Japan, March 16-19, 2009*. IEEE Computer Society, 2009, pp. 705–711. [Online]. Available: <https://doi.org/10.1109/CISIS.2009.178>

- [67] M. Ilse, J. M. Tomczak, and M. Welling, "Attention-based Deep Multiple Instance Learning," in *Proc. of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, ser. Proc. of Machine Learning Research, vol. 80. PMLR, 2018, pp. 2132–2141. [Online]. Available: <http://proceedings.mlr.press/v80/ilse18a.html>
- [68] M. Y. Lu, D. F. K. Williamson, T. Y. Chen, R. J. Chen, M. Barbieri, and F. Mahmood, "Data Efficient and Weakly Supervised Computational Pathology on Whole Slide Images," *CoRR*, vol. abs/2004.09666, 2020. [Online]. Available: <https://arxiv.org/abs/2004.09666>
- [69] A. B. Benson, A. P. Venook, M. M. Al-Hawary, L. Cederquist, Y. J. Chen, K. K. Ciombor, S. Cohen, H. S. Cooper, D. Deming, and P. F. Engstrom, "NCCN guidelines insights: colon cancer, version 2.2018," *J Natl Compr Canc Netw*, vol. 16, no. 4, pp. 359–369, 2018. [Online]. Available: <https://doi.org/10.6004/jnccn.2018.0021>
- [70] J. A. King, J. Jeong, F. E. Underwood, J. Quan, N. Panaccione, J. W. Windsor, S. Coward, J. deBruyn, P. E. Ronskley, and A. A. Shaheen, "Incidence of Celiac Disease is Increasing over Time: a Systematic Review and Meta-Analysis," *Am J Gastroenterol*, vol. 115, no. 4, pp. 507–525, 2020. [Online]. Available: <https://doi.org/10.14309/ajg.0000000000000523>