

PEDANTIC: A Dataset for the Automatic Examination of Definiteness in Patent Claims

Valentin Knappich^{1,2}, Annemarie Friedrich², Anna Hättö¹ and Simon Razniewski³

¹Bosch Center for AI

²University of Augsburg

³ScaDS.AI & TU Dresden

Abstract

Patent claims define the scope of protection for an invention. If there are ambiguities in a claim, it is rejected by the patent office. In the US, this is referred to as indefiniteness (35 U.S.C § 112(b)) and is among the most frequent reasons for patent application rejection. The development of automatic methods for patent definiteness examination has the potential to make patent drafting and examination more efficient, but no annotated dataset has been published to date.

We introduce PEDANTIC (Patent Definiteness Examination Corpus), a novel dataset of 14k US patent claims from patent applications relating to Natural Language Processing (NLP), annotated with reasons for indefiniteness. We construct PEDANTIC using a fully automatic pipeline that retrieves office action documents from the USPTO and uses Large Language Models (LLMs) to extract the reasons for indefiniteness. A human validation study confirms the pipeline’s accuracy in generating high-quality annotations. To gain insight beyond binary classification metrics, we implement an LLM-as-Judge evaluation that compares the free-form reasoning of every model-cited reason with every examiner-cited reason. We show that LLM agents based on Qwen 2.5 32B and 72B struggle to outperform logistic regression baselines on definiteness prediction, even though they often correctly identify the underlying reasons. PEDANTIC provides a valuable resource for patent AI researchers, enabling the development of advanced examination models. We release the dataset and code at <https://github.com/boschresearch/pedantic-patentsemtech>.

Keywords

Patent AI, Patent Examination, Patent Definiteness, Patent Clarity, Patent Classification

1. Introduction

The patent system plays a crucial role in fostering innovation by granting inventors exclusive rights to their inventions. Central to this system are patent claims, which are concise and precise statements that define the metes and bounds of the protected invention. The process of obtaining a patent involves rigorous examination by patent offices to ensure that the application meets specific criteria, including novelty, non-obviousness, and, critically, *definiteness*. The latter requires that every claim is sufficiently clear and unambiguous to enable a person skilled in the art (called *the Person of Ordinary Skill in the Art*, or *POSITA*) to understand the scope of the invention. In American patent law, this is defined in 35 U.S.C. § 112(b) (comparable to *clarity* in the EU¹), which states that “*the specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.*” The Manual of Patent Examining Procedure (MPEP) [1] provides detailed instructions for the examination. We present the most common categories of indefiniteness in Table 1.

Ensuring definiteness is challenging for patent attorneys and examiners. Patent applications typically undergo multiple rejection-response cycles, averaging 26 months from filing to disposition [2]. With increasing application volumes [3], AI-powered methods are needed to improve examination efficiency

6th Workshop on Patent Text Mining and Semantic Technologies (PatentSemTech) 2025

*Corresponding author.

†These authors contributed equally.

✉ valentin.knappich@de.bosch.com (V. Knappich); annemarie.friedrich@uni-a.de (A. Friedrich); anna.haetty@de.bosch.com (A. Hättö); simon.rzniewski@tu-dresden.de (S. Razniewski)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://www.epo.org/en/legal/epc/2020/a84.html>

https://www.epo.org/en/legal/guidelines-epc/2024/f_iv_4.html

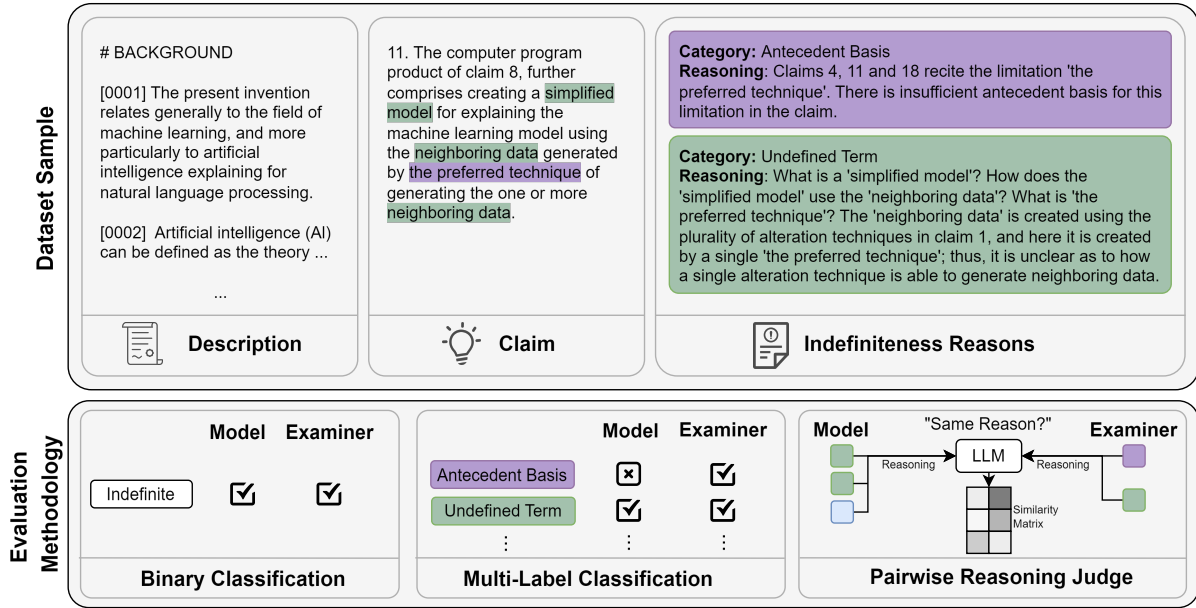


Figure 1: PEDANTIC example and evaluation methodology. Every sample includes a claim, the patent description, and a list of rejection reasons that cause indefiniteness, each with a free-form reasoning, indefiniteness category, and recited claim spans. We evaluate models on (1) the prediction of the binary label definite or indefinite, (2) the identification of the correct indefiniteness categories as multi-label classification, and (3) the generation of correct free-form reasoning for indefiniteness reasons using an LLM-as-Judge that compares all pairs of examiner-cited reason and model-cited reason.

and consistency. Automating examination could yield significant cost savings. For instance, it could assist examiners, aid attorneys in drafting more robust applications, and improve accessibility for those without extensive patent expertise. It could also provide feedback for training and evaluating automatic drafting systems [4].

Indefiniteness is one of the most common reasons for rejection [5, 6, 7], yet its automatic examination has received very little attention in prior work compared to novelty and non-obviousness. To bridge this gap, we specifically target the automatic examination of definiteness. We argue that predicting definiteness as a binary label is insufficient for practical applications. For patent drafters and examiners, knowing that a claim is indefinite does not provide actionable information. Rather, understanding *why* it is indefinite is crucial, as it enables targeted improvements to the claim. For this purpose, we introduce PEDANTIC, a dataset of patent claims annotated with detailed justifications for rejection due to indefiniteness. In addition to the claims, PEDANTIC includes fine-grained indefiniteness categories (such as '*antecedent basis*') along with a free-form reasoning, and the affected ranges of the claim for every indefiniteness reason. We create it fully automatically, leveraging Large Language Models (LLMs) to parse USPTO office actions into a structured format that includes these annotations. In total, PEDANTIC includes 14k claims from 3k utility patent applications relating to Natural Language Processing (NLP) filed after 2014. It facilitates a nuanced evaluation, distinguishing between systems that identify indefiniteness based on superficial clues and those that are able to pinpoint the correct underlying issue. In addition to established classification metrics, we implement a reference-based LLM-as-Judge [8] evaluation that compares each model-cited rejection reason with every examiner-cited rejection reason. We release our dataset and code publicly to facilitate further research.

2. Related Work

Automated Patent Examination. Most related work on automatic patent examination has focused on novelty and non-obviousness. Datasets for novelty assessment [9, 10, 11, 12] pair patent claims

either with short passages [9, 10] or large chunks [11, 12] from prior art. They use citations marked as novelty-destroying by the examiner as positive samples and obtain negative samples either from other citations [9, 10, 11] and/or from related patents [10, 12]. Novelty has been evaluated using BERT-based models [9, 10, 11, 12, 13, 14], graph neural networks [15, 16], and LLMs [11, 17]. Lee et al. [18] furthermore present a dataset for the prediction of patent edits following rejections based on novelty and non-obviousness.

Patent Clarity and Definiteness. Hido et al. [19] predict patentability using linguistic features including syntactic complexity and word age. Kong et al. [20] model patent readability using § 112(a) (lack of disclosure) with linguistic features, finding university-issued patents more readable than corporate-issued ones. Lo and Chu [7] frame patentability prediction as a multi-label classification task, including definiteness, using BERT-based models, but do not report definiteness-specific results. Ashtor [21] trains a definiteness classifier based on linguistic features and reaches 68% AUROC. They use this classifier as proxy for clarity and show that clarity has improved over time through policy changes regarding definiteness rejections.

LLM Explanations. While the primary purpose of a classifier is to predict the most likely class, practical applications often demand more transparency and insight into the reason behind a prediction. There is a large body of work concerning explainable AI (XAI) [22] that has developed various techniques to produce explanations for a classifier’s prediction, including feature attribution methods, counterfactual explanations, rule extraction, and example-based reasoning. Recently, LLMs have been employed to directly generate explanations along with their predictions. Such self-explanations have been shown to perform on par with traditional explainability methods [23], but also have limited faithfulness [24, 25]. In this work, we propose to use the free-form reasons written by the patent examiner in the rejection full-text as ground truth to evaluate whether a model predicts indefiniteness for the right reasons.

3. PEDANTIC Dataset

In this section, we describe the creation of our dataset, including the retrieval of seed patent applications, the retrieval of rejection notices, document parsing, and dataset splitting.

3.1. Seed Patent Applications

Our dataset creation pipeline first requires a set of seed patent applications. We focus on patent applications in the area of NLP, i.e., applications with the CPC class “G06F40” (“Handling natural language data”), but our pipeline is agnostic to this and can be readily used with seed applications from any other field. We query the USPTO Open Data Portal (ODP)² API for applications from 2014 to date with at least one filed rejection notice. We do not consider applications prior to 2014 because the requirements for definiteness changed significantly after the US Supreme Court’s ruling on *Nautilus vs. Biosig* [26]. Before 2014, a claim was only rejected for indefiniteness if it was “insolubly ambiguous”, whereas afterward, indefiniteness was interpreted as a claim “[failing] to inform, with reasonable certainty, those skilled in the art about the scope of the invention”. Ashtor [21] finds empirical evidence that this ruling has significantly increased the fraction of claims being rejected for indefiniteness and the average claim clarity. To avoid exploitable biases, we filter out applications filed before 2014.

3.2. Document Download

For all retrieved seed patents, we download the claim, specification and office action documents from the ODP File Wrapper API³. We only consider the first office action, as it generally establishes the core

²<https://data.uspto.gov/>

³<https://data.uspto.gov/apis/patent-file-wrapper/search>

Reason	MPEP Section	Description
Antecedent Basis	2173.05(e)	Claim contains a term referencing an element lacking a clear prior introduction, creating ambiguity what it references.
Undefined Term	2173.05(a)	A term lacks a clear, accepted and/or unambiguous meaning to a POSITA, making the claim’s scope uncertain.
Relative Term	2173.05(b)	A relative term (e.g., ‘thin,’ ‘substantial’) is used without providing a clear point of comparison, rendering the claim’s scope indefinite.
Exemplary Phrasing	2173.05(d)	Claim uses ‘such as,’ or similar phrasing, making it unclear whether the listed items are exhaustive or merely examples, leading to indefiniteness.
Functional Claiming	2173.05(g)	Claim recites ‘means for’ or ‘step for’ without disclosing adequate corresponding structure, material, or acts in the specification, as required under 35 U.S.C. 112(f) or pre-AIA equivalent.
Contradicting Limitations	2173	Claim includes an element that contradicts or is inconsistent with other claim limitations, making the claim’s scope unclear.
Omission of Essential Elements or Steps	2172.01	Claim fails to recite an element, step, or cooperative relationship between elements/steps that is essential to the invention as disclosed.
Dependence	-	Claim depends on an indefinite claim.
Other	-	Catch-all category for indefiniteness reasons not covered above.

Table 1

Brief descriptions of indefiniteness categories and their corresponding MPEP sections. The same descriptions were used during parsing to help the model identify the correct reason. Greyed-out reasons were used only during parsing and filtered out for the final dataset.

grounds for rejection. While subsequent office actions may address remaining or newly introduced issues, the initial rejection often provides the most comprehensive overview of the examiner’s concerns. We select the latest claim and specification documents preceding this office action and download all three documents in XML format. We parse the XML documents into Markdown and JSON for easier further processing.

3.3. Office Action Parsing

We parse the office actions’ full-text into a structured representation using Gemma 3 27B [27]. First, we select the sections related to indefiniteness by filtering for sections with headings containing “112”. Next, we prompt the LLM with the selected sections and instruct it to extract indefiniteness reasons in a JSON schema. In particular, each rejection contains the text snippet arguing why the claim is indefinite, a category from Table 1, and a list of recited phrases. We instruct the LLM to extract the free-form argumentation and recited phrases verbatim and not to paraphrase, extend or explain them. The prompt is attached in Appendix A. Lastly, we use fuzzy matching to find the occurrences of recited phrases in the claim text. An example from our dataset is shown in Figure 1. A claim can have multiple rejection reasons, but most have only one (see Section 3.7).

3.4. Sampling Definite Claims

To ensure robust evaluation, we balance the dataset between definite and indefinite claims. We use claims from applications whose office actions do not contain the term “112(b)”, i.e., applications whose claims are all definite. This makes sure that all claims labelled as definite are indeed not rejected for indefiniteness; using claims from the same applications as the previously extracted indefinite claims

		Dataset Splits					
		All		Train		Test	Validation
Claims	Total	14536	(100.00%)	8730	(60.06%)	4362	(30.01%)
	Definite	7268	(50.00%)	4363	(49.98%)	2178	(49.93%)
	Indefinite	7268	(50.00%)	4367	(50.02%)	2184	(50.07%)
	Independent	3339	(22.97%)	1996	(22.86%)	1026	(23.52%)
	Dependent	11197	(77.03%)	6734	(77.14%)	3336	(76.48%)
Applications	Total	3710	(100.00%)	2226	(60.00%)	1113	(30.00%)
	Definite	1855	(50.00%)	1107	(49.73%)	559	(50.22%)
	Indefinite	1855	(50.00%)	1119	(50.27%)	554	(49.78%)
Rejection Reason	Total	9215	(100.00%)	5598	(60.75%)	2765	(30.01%)
	Antecedent Basis	3350	(36.35%)	1994	(35.62%)	1015	(36.71%)
	Undefined Term	3394	(36.83%)	2093	(37.39%)	1035	(37.43%)
	Relative Term	910	(9.88%)	568	(10.15%)	258	(9.33%)
	Exemplary Phrasing	121	(1.31%)	97	(1.73%)	19	(0.69%)
	Functional Claiming	850	(9.22%)	488	(8.72%)	270	(9.76%)
	Contradicting Limitation	442	(4.80%)	288	(5.14%)	113	(4.09%)
	Omission of Essential Elements or Steps	148	(1.61%)	70	(1.25%)	55	(1.99%)

Table 2

Dataset statistics across splits. In “Total” sections, the percentages sum to 100% per row, otherwise per column.

could introduce noise if the parsing pipeline does not detect all rejections. To also balance the number of applications in each class, we first compute the average number of indefinite claims per application included in the dataset. We iterate through the definite applications in random order and sample the same number of claims until there are as many definite as indefinite claims. While sampling claims from the applications, we round the number of sampled claims down or up depending on whether there are currently more or less claims per application than in the indefinite samples.

3.5. Dataset Splits

We randomly split the resulting dataset into train (60%), test (30%), and validation (10%). To avoid data leakage, we perform this split on the application-level, such that claims from the same application are always in the same split.

3.6. Human Validation Study

To validate the quality of our extracted annotations, we manually inspect 50 randomly sampled claims from PEDANTIC. Among these claims, 24 are definite and 26 are indefinite, with a total of 27 reasons for indefiniteness (one indefinite claim has two reasons while the rest have one). For each claim, we compare the rejection document with the extracted annotations and analyze the correctness of the binary label, the extracted reasoning texts, and the assigned categories. We find that all binary labels are correct. The extracted free-form reasoning is also correct in all samples except one, where it says “*The underlined lacks antecedent basis*”, but this formatting is not translated into Markdown. This validates the reliability of our automatically extracted binary labels and reasoning texts. However, we find substantial noise in the assigned categories, with 19 out of 27 reasons having the correct one. In six out of eight incorrect categories, none of the proposed categories would have been a good fit. The LLM assigned the category “*undefined term*”, instead of “*other*” as instructed. In the remaining two out of eight incorrect category assignments, the correct category would have been “*antecedent basis*”, but the LLM assigned “*undefined term*”. In both cases, the term “*antecedent basis*” was not mentioned explicitly, unlike in most reasoning texts of this category.




Figure 2: Definite vs. indefinite claims in our dataset over time. Left y-axis: stacked absolute values. Right y-axis: fraction of indefinite claims. The fraction remains around 50%, therefore models cannot determine definiteness by inferring the date.

3.7. Dataset Statistics

Table 2 shows dataset statistics over the splits. In total, the dataset contains 14.5k claims from 3.7k patent applications, half of which are indefinite. It is thus a balanced binary classification dataset, and a multi-label classification dataset considering the presence of each category as a binary label. It consists of roughly 77% dependent⁴ claims and 22% independent claims. The fraction of definite claims, the fraction of independent claims, and the distribution of rejection reasons are roughly the same in each split, i.e., there are no substantial distribution shifts.

Indefiniteness Categories. Table 2 also shows the distribution of indefiniteness categories. Missing or ambiguous antecedent bases and unclear definitions dominate the dataset, constituting a combined 73% of all indefiniteness reasons. Relative terms and functional claiming are also commonly cited, each constituting about 9% of indefiniteness reasons. Contradicting limitations, exemplary phrasing and omission of essential elements or steps are each cited in less than 5% of indefiniteness reasons.

Characteristics of definite vs indefinite claims. If superficial characteristics differ between definite and indefinite claims, a trained classifier will likely use them as a shortcut. We therefore report several such characteristics across the classes. As shown in Figure 2, the fraction of indefinite claims remains around 50% over time, i.e., models gain no advantage from the filing date, even if they infer it from the specific technology and terminology in the claim. As shown in Table 3, the indefinite claims are more frequently independent claims than definite claims. In other words, in our dataset, independent claims are more likely to be rejected due to indefiniteness than dependent claims. A plausible reason is that independent claims are longer, i.e., they introduce more features that could be indefinite. We consequently observe that indefinite claims are longer on average in terms of characters, words, and features.

4. Definiteness Prediction

In this section, we present two baseline approaches to predict indefiniteness: logistic regression and an LLM agent. The available input is the claim in question and the accompanying patent description.

4.1. Logistic Regression

We include logistic regression in our experiments because it is computationally efficient and interpretable, and because Ashtor [21] has shown it to achieve non-trivial performance predicting indefiniteness.

⁴A dependent claim is a claim that refers back to another claim, incorporating its features and narrowing its scope of protection.

Independent	Definite			Indefinite		
	✓+✗	✓	✗	✓+✗	✓	✗
% Samples	50.0%	7.8%	42.2%	50.0%	15.0%	35.0%
# Parents	1.2	0.0	1.4	1.0	0.0	1.5
# Characters	415.4	931.8	319.4	597.9	1080.8	391.5
# Words	64.9	143.2	50.3	93.0	166.6	61.5
# Features	2.7	6.4	2.0	3.6	6.8	2.2

Table 3

Claim characteristics across definite and indefinite claims, further disaggregated over independent and dependent claims. The fraction of independent claims is larger among indefinite claims than among definite claims, and indefinite claims are longer than definite claims on average. The latter is largely, but not entirely, a side-effect of the former, since independent claims are naturally much longer than dependent claims.

We use TF-IDF features and a number of handcrafted linguistic features, similar to Ashtor [21]. The latter include the claim length, the claim length relative to the description length, readability metrics, trigger word flags, and a flag indicating whether the claim is independent. The full list is visualized in Figure 3. For all features sets, we train separate classifiers for the binary classification and the multi-label classification.

4.2. LLM Agent

LLMs have shown impressive performance on many tasks related to patents [4], yet they have not been evaluated on definiteness prediction. We implement a zero-shot LLM agent to identify the issues causing indefiniteness in the claims. The prompt (see Appendix B) includes instructions and the claim in question. The agent is equipped with two tools that allow it to search for relevant information in the remaining document. First, we add a tool that returns a claim given its claim number, allowing the agent to analyze parent claims, and other claims deemed relevant. Second, we implement a TF-IDF search tool that allows the agent to retrieve paragraphs containing certain key words or phrases from the patent’s description section. We choose this tool-based approach because pasting the entire description quickly fills up the context window, while pre-selecting the relevant parts is restrictive and requires hand-crafting selection criteria and retrieval mechanisms. Using these tools, the agent can flexibly retrieve whatever information it deems relevant. The agent analyses the claim and performs tool calls in an interleaved fashion until it arrives at a final prediction.

Rather than generating the binary label directly (definite/indefinite), we instruct the LLM to use a verbalized expression of likelihood among a set of possible options ranging from “almost no chance” to “almost certain”. Confidence scores, if reliable, should allow users to configure a level of strictness for issue detection. Each expression is converted to a numerical probability according to the empirically determined human perception of probability expressions by Fagen-Ulmschneider [28]. We choose this approach because Tian et al. [29] and Xiong et al. [30] show that asking the LLM for a confidence value is equally or more reliable than logit-based confidence estimation.

Lastly, we instruct the LLM to format the result as JSON according to a fixed schema. The final output contains a prediction of the likelihood of the claim being rejected, and a list of potential reasons for the indefiniteness, where each reason contains a confidence score (using the same verbalized expressions as above), a free-form reasoning, one of the categories in Table 1, and a list of claim recitations. Thus, the output is in the same format as the structured representations extracted from the office actions, with the additional confidence scores. This allows quantitative evaluation of the binary classification, multi-label classification, and the correctness of textual reasons.

5. Evaluation Metrics

In this section, we propose a set of metrics to evaluate claim indefiniteness prediction models. First, since the core task is binary classification, we use established metrics that measure how well the system can predict the binary label “definite” or “indefinite”. While this binary label is already interesting by itself, practical applications demand a more fine-grained and explainable approach. It is equally or even more important to know *why* a claim is indefinite to give actionable feedback to drafters. To that end, we implement an LLM-as-Judge [8] approach that determines whether an examiner-cited reason and a model-cited reason point to the same essential issue in the claim. Lastly, we evaluate the task as a multi-label classification, where the overlap in the categories of cited reasons also indicate whether a model decided to accept or reject a claim for the right reasons.




Figure 3: Feature importance scores of the logistic regression classifier with linguistic features. Positive values (red) indicate indefiniteness, negative values (green) indicate definiteness. The classifier performs similarly on train and test, i.e., the feature weights are not overfitted.

5.1. Binary Classification

We evaluate binary classification performance for identifying indefinite claims using the following metrics:

- **Precision:** The proportion of true positives (i.e., claims correctly predicted as indefinite) out of all claims predicted as indefinite by the model.
- **Recall:** The proportion of true positives out of all claims found indefinite by the examiner.
- **F1-score:** The harmonic mean of precision and recall, providing a balanced measure of both.
- **AUROC:** The area under the Receiver Operating Curve (ROC), which plots the true positive rate against the false positive rate for varying classification thresholds.
- **Accuracy:** The proportion of correctly classified claims (as either definite or indefinite) out of all claims.

These metrics provide a general understanding of the system’s ability to distinguish between definite and indefinite claims. For all models, we compute these metrics for a confidence threshold that balances the predictions, i.e., we first determine the threshold with which half of the claims from the validation set are predicted as indefinite.

5.2. Pairwise Reasoning Judge

To assess the quality of the model’s reasoning for indefiniteness, we employ a reference-based LLM-as-Judge approach [8]. We prompt Gemma 3 27B to evaluate whether an examiner-cited reason and a

model-cited reason point to the same essential issue in the claim on a scale from 1 (worst) to 5 (best) (we normalize the final scores to the range $[0, 100]$). The few-shot prompt is attached in [Appendix C](#). To allow the model to analyze the ground truth and the predicted reason before settling on a grade, we first prompt it to find similarities and differences, then ask for a numerical grade in a follow-up message. Following latest research in LLM-as-Judge systems [31, 32], we use the probability-weighted mean as the similarity score of a reason-pair:

$$\text{norm}(s) = 100 \cdot \frac{s - 1}{4} \quad \text{sim} = \text{norm} \left(\frac{\sum_{i=1}^5 i \cdot p_i}{\sum_{i=1}^5 p_i} \right)$$

where i is a grade and p_i is the LLM-determined probability of that grade (i.e., the exponential of the logit of the token corresponding to the grade).

We compute this similarity for every pair of examiner-cited reason and model-cited reason. Hence, for a sample with n examiner-cited reasons and m model-cited reasons, we obtain a similarity matrix $S \in \mathbb{R}^{n \times m}$. From this matrix, we compute precision, recall and F1, each in a thresholded and a soft variant. In the thresholded variant, we consider a model-cited reason correct if there is an examiner-cited reason with which it has a similarity score of 75 (corresponding to “*The reasons are closely related and largely address the same issue in the claim*”) and above, and vice versa.

$$\begin{aligned} P_{\geq 75} &= \frac{1}{m} \sum_{j=1}^m \max_i \mathbb{1}_{S_{ij} \geq 75} \\ R_{\geq 75} &= \frac{1}{n} \sum_{i=1}^n \max_j \mathbb{1}_{S_{ij} \geq 75} \end{aligned}$$

We report the macro average (i.e., the average of all average per-claim scores) and the micro average (i.e., the average of all per-reason scores). For the soft variant, we do not set a fixed threshold but directly average over the maximum similarity in the respective dimension.

$$P = \frac{1}{m} \sum_{j=1}^m \max_i S_{ij} \quad R = \frac{1}{n} \sum_{i=1}^n \max_j S_{ij}$$

In both cases, F1 is computed as harmonic mean of P and R.

5.3. Multi-Label Classification

We also assess the system as a multi-label classification task, where each label serves as a binary indicator denoting the presence or absence of a specific category within the list of indefiniteness reasons. A high score means that the model has classified a claim for reasons belonging to the same or overlapping indefiniteness categories as the examiner, and thus also represents a kind of explanation for the binary classification. We report the micro and macro average of the category-wise F1 scores. As with the binary classification, we balance the predictions using the confidence threshold per category under which the fraction of positive predictions matches the distribution in the validation set.

6. Results

[Table 4](#) shows the classification results. The best overall model achieves an AUROC score of 60.3. Thus, its predictions are considerably better than random, but there is substantial room for improvement. LLM agents do not outperform the logistic regression baseline in terms of binary classification performance. Despite this, we find that the agents correctly identify many of the reasons for indefiniteness that were also noted by the examiner. In the following, we further analyze the classification performance, classifier calibration, and the LLM-as-Judge results.

Model	% indef	Binary					Multi-Label	
		P	R	F1	AUC	Acc	Macro F1	Micro F1
Random	50.4%	50.9	51.4	51.1	51.3	50.9	6.6	11.9
<i>Logistic Regression</i>								
TF-IDF	46.2%	53.2	49.3	51.2	54.8	53.0	9.3	18.5
Ling. Features	49.2%	57.2	56.3	56.8	59.4	57.2	11.8	21.1
All Features	49.2%	56.8	55.9	56.3	59.5	56.7	12.1	22.0
<i>LLM Agents</i>								
Qwen 2.5 32B	45.9%	54.7	50.3	52.4	55.2	54.4	13.7	22.8
+ Log. Reg.	48.3%	55.8	54.0	54.9	58.8	55.7	13.7	22.7
Qwen 2.5 72B	33.7%	57.8	39.0	46.6	56.8	55.3	13.5	21.7
+ Log. Reg.	48.9%	59.3	58.2	58.8	60.3	59.2	14.1	22.9

Table 4

Classification Results. The best value per column is bolded. All predictions use the prediction threshold that balances the label distribution on the validation set. P=Precision, R=Recall, Acc=Accuracy, AUC=AUROC. The random baseline samples from a uniform distribution.



Figure 4: Distributions of claim-level indefiniteness probability estimates for different binary classifiers and the relation of estimated probability and accuracy. The accuracy per probability bin is plotted as red dots. For bins with less than 1% of samples, the dot is faded. The logistic regression model shown here uses the full feature set.

Logistic Regression. Among the three proposed variants of logistic regression, the one with linguistic features and the one with all features perform comparably, while using only TF-IDF features performs substantially worse. This indicates that the handcrafted features are indeed helpful to differentiate definite and indefinite claims, and that TF-IDF features provide only limited additional information.

We show the feature importance values in Figure 3. The most important feature indicating definiteness is the IOU of words between claim and description, i.e., samples with a high word overlap between claim and description are less likely to be rejected for indefiniteness. The highest-weighted features indicating indefiniteness are the number of unique word stems, the type token ratio and the number of stopwords. Among the binary features showing whether a trigger word is contained in the claim, “step for” is the most indicative of indefiniteness. The flag whether the claim is independent indicates definiteness, which is in line with our dataset statistics, albeit with moderate importance. Many of the text complexity and readability metrics have high importance scores. A higher text complexity is associated with definiteness (with Gunning Fog, Automated Readability Index, and Flesch Kincaid Grade having a negative feature weight) and better readability is associated with indefiniteness (with Flesch Reading Ease and Dale Chall Readability Score having a positive feature weight). While this might seem counterintuitive when relating definiteness to clarity or understandability, it makes sense when considering definiteness as the absence of ambiguity, as removing ambiguities often adds complexity through additional specification.

LLM Agent. The LLM agents using Qwen 2.5 32B and 72B both make use of the provided tools, each making at least one tool call for 98% of the samples. Notably, the larger model makes more tool calls (3.3 calls per sample on average) than the smaller model (2.3 calls per sample on average). This indicates that

		Soft			Thresholded		
		P	R	$F1$	$P_{\geq 75}$	$R_{\geq 75}$	$F1_{\geq 75}$
Qwen 2.5 32B	macro	14.6	43.9	21.9	11.2	17.9	13.8
	micro	10.8	43.6	17.3	7.7	35.4	12.7
Qwen 2.5 72B	macro	9.8	33.5	15.1	7.4	12.3	9.2
	micro	6.1	33.6	10.3	4.0	25.0	6.9

Table 5

LLM-as-Judge results using Gemma 3 27B as judge. P=Precision, R=Recall. The thresholded variant considers reasons with scores above 75% to be correct.




Figure 5: Distributions of reason-level confidence scores and the relation of confidence and correctness according to LLM-as-Judge. The reported judge scores are the soft micro-averaged precision scores, and thus represent the average max similarity per model-cited reason. For bins with less than 1% of samples, the dot is faded. The dashed line is a linear regression that ignores the faded points. ρ corresponds to the pearson correlation between confidence and judge score.

the larger model is more thorough in searching for relevant information in the document, which could be a reason for the slightly higher AUROC and accuracy. Recall and F1 are notably worse for the larger model, caused by the imbalance of its predictions. Both metrics favor classifiers that frequently predict indefiniteness: F1 is per definition $\frac{2}{3}$ when predicting indefiniteness for every sample. The distribution of confidence scores produced by the 72B model makes it impossible to find a threshold that moves the fraction of indefinite predictions close to 50%. We show the distribution of confidence scores for both LLM agents and the logistic regression in Figure 4. The LLM agents are typically highly confident in either extreme, whereas the logistic regression produces a smoother distribution with a peak in the middle. In addition, as visible in Figure 4, there is no clear relationship between confidence and accuracy in the LLMs’ binary predictions. This is unlike the predictions of the logistic regression model, where the accuracy drops notably around the classification threshold. Overall, the results indicate that the LLMs confidence estimates are not well-calibrated, and that better calibration could enhance the LLMs’ classification performance.

Multi-Label Classification. LLM agents and logistic regression achieve only moderate F1 scores. All models achieve a much better micro F1 than macro F1, indicating that the performance varies between categories. All models perform much better on the well-represented categories than on rare ones. Given the substantial noise in the labels for the category “*undefined term*” (see Section 3.6), additional investigations are necessary to draw further conclusions.

Ensemble. We achieve the best overall performance using an ensemble between Qwen 2.5 72B and the logistic regression. To create the ensemble, we average the predicted probability of indefiniteness for every sample. This leads to moderate improvements, indicating that the predictions are complementary to some degree.

Judge Evaluation. Our LLM-as-Judge evaluation provides further insight into LLM agents’ performance by directly analyzing the quality of their identified reasons for indefiniteness. Table 5 shows the

Examiner	Qwen 2.5 32B	Judge Similarity
1 The term 'right before' in claim 9 is a relative term which renders the claim indefinite. The term 'right before' is not defined by the claim, the specification does not provide a standard for ascertaining the requisite degree , and one of ordinary skill in the art would not be reasonably apprised of the scope of the invention. The amount of time necessary for an input to constitute being obtained 'right before' the first input is unclear. Furthermore and for example, a person of ordinary skill in the art may consider an arbitrary number of previous inputs, such as three, to constitute inputs obtained 'right before' a first input.	The term 'right before' is a relative term that lacks a clear, objective standard for determining the exact input used to determine the first domain.	100%
2 The claims recite the limitation 'ranking the one or more groups' There is insufficient antecedent basis for this limitation in the claim, it appears likely that this claim should instead depend from claim 2 .	The term 'one or more groups' in claim 3 lacks antecedent basis as it is not introduced or defined in claim 1.	75%
3 As per Claim 10, 'each highlight in the second pane in the second interface' lacks antecedent basis to the extent that 'each' implies that there is more than one highlight in the second pane. Claim 9 only recites where a single highlight is dropped into the second pane and Claim 1 recites where selected ones of the highlights are associated with the second pane of the second interface in a sequence (associating does not imply putting the selected highlights into the second pane of the second interface, and associating in a sequence can be where a sequence in the second interface includes elements referencing to the highlights themselves without including the highlights themselves).	The term 'highlight panel' is not clearly defined in the claims or the description, leading to uncertainty about its scope.	25%

Table 6

Examples of examiner-cited reasons and best matching model-cited reasons for indefiniteness. Important parts are highlighted with matching colors between examiner and model if applicable.

results for Qwen 2.5 32B and 72B. Interestingly, the smaller model clearly outperforms the larger one in all judge metrics, despite exhibiting slightly lower binary classification performance. There seems to be a disconnect between the models' ability to identify potential reasons for indefiniteness and its ability to accurately decide whether a claim is indefinite. The 32B model is able to correctly identify 35.4% of the examiner-cited reasons (micro $R_{\geq 75}$, 25.0% for 72B). Both models list many reasons, even if they consider them unlikely to cause indefiniteness, causing the precision to be lower than the recall.

Figure 5 also shows the relation between reason-level confidence and the soft micro-averaged precision. Both models exhibit a positive trend; as their confidence in a reason increases, the likelihood of the examiner citing the same or a similar reason also increases. The Pearson correlation between confidence and precision is positive and moderate ($\rho = 0.432$ for 32B and $\rho = 0.514$ for 72B). That is, the reason-level confidence appears to be better calibrated than the claim-level confidence. Future work should further investigate the relation between these two sub-tasks and develop methods to bridge this disconnect.

7. Discussion

We here compare example reasons for indefiniteness generated by Qwen 2.5 32B with those written by the human examiner. Table 6 lists three examples. In Example 1, the LLM agent correctly identified the reason for indefiniteness, albeit with a less verbose explanation. In Example 2, the LLM agent also found the correct underlying issue. However, its explanation is not equally helpful as the one provided by the human examiner, which additionally suggests that the claims might accidentally depend from the wrong base claim. In Example 3, the model-generated reason points to similar problematic phrases, but without identifying the underlying issue. The LLM-as-Judge correctly identifies that there is little

substantial overlap and assigns a lower score of 25%. Generally, the model-generated reasons seem to be less detailed, and often lack the deep analysis and reasoning performed by human examiners.

While our dataset creation process is domain-agnostic, our study focuses on patents from the NLP domain. Therefore, the claim language and terminology is relatively homogeneous. Other domains might differ in terms of terminology, distribution of claim types and categories of indefiniteness, and legal conventions. Future work should generalize our study design and verify that our findings transfer to other domains.

8. Conclusion and Outlook

In this work, we tackle the task of automatic patent definiteness examination. We present PEDANTIC, the first publicly available dataset for this task to enable reproducible experiments with different examination models. We conduct first experiments on PEDANTIC using logistic regression and LLMs. All models are able to perform better than random, but there is substantial room for improvement. LLMs are able to identify many of the reasons for indefiniteness also cited by the human examiner, yet their binary classification performance fails to clearly outperform logistic regression with hand-crafted features. We show that poor calibration is one of the issues of current LLM-based setups. Promising directions for future research include better calibration methods, fine-tuning on task-specific data using supervised and/or reinforcement learning, and using the rejection full-text for in-depth evaluation of other patent requirements like novelty and non-obviousness.

Acknowledgments

We would like to thank the patent attorneys Philipp Mangold and Charlotte Hellmann for insightful discussions and helpful pointers.

Declaration on Generative AI

During the preparation of this work, the authors used Gemini 2.0 Flash in order to⁵: grammar and spelling check, improve writing style, paraphrase and reword. After using these tool(s)/service(s), the authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

References

- [1] USPTO, Manual of Patent Examining Procedure (MPEP), 2024. URL: <https://www.uspto.gov/web/offices/pac/mpep/index.html>.
- [2] USPTO, Pendency | Patents Dashboard | USPTO, 2025. URL: <https://www.uspto.gov/dashboard/patents/pendency.html>.
- [3] World Intellectual Property Organization, World Intellectual Property Indicators 2024, World Intellectual Property Organization, 2024. doi:10.34667/TIND.50133.
- [4] V. Knappich, S. Razniewski, A. Hättig, A. Friedrich, Pap2pat: Towards automated paper-to-patent drafting using chunk-based outline-guided generation, arXiv preprint arXiv:2410.07009 (2024).
- [5] Q. Lu, A. Myers, S. Beliveau, USPTO Patent Prosecution Research Data: Unlocking Office Action Traits, 2017. doi:10.2139/ssrn.3024621. arXiv:3024621.
- [6] The Most Common Rejections: 102, 103, and 112(b), 2019. URL: <https://blog.juristat.com/most-common-rejections>.

⁵<https://ceur-ws.org/GenAI/Taxonomy.html>

- [7] H.-C. Lo, J.-M. Chu, Pre-trained Transformer-based Classification for Automated Patentability Examination, in: 2021 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE), 2021, pp. 1–5. doi:10.1109/CSDE53843.2021.9718474.
- [8] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. P. Xing, H. Zhang, J. E. Gonzalez, I. Stoica, Judging llm-as-a-judge with mt-bench and chatbot arena, in: Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23, 2023.
- [9] J. Risch, N. Alder, C. Hewel, R. Krestel, PatentMatch: A Dataset for Matching Patent Claims & Prior Art, 2020. doi:10.48550/arXiv.2012.13919. arXiv:2012.13919.
- [10] K. Vowinckel, V. D. Hähnke, SEARCHFORMER: Semantic patent embeddings by siamese transformers for prior art search, World Patent Information 73 (2023) 102192. doi:10.1016/j.wpi.2023.102192.
- [11] M. Blume, G. Heidari, C. Hewel, Comparing complex concepts with transformers: Matching patent claims against natural language text, in: 5th Workshop on Patent Text Mining and Semantic Technologies (PatentSemTech, volume 3775, CEUR-WS.org, 2024, pp. 72–76.
- [12] A. Parikh, S. Dori-Hacohen, ClaimCompare: A data pipeline for evaluation of novelty destroying patent pairs, in: 5th Workshop on Patent Text Mining and Semantic Technologies (PatentSemTech, volume 3775 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024, pp. 61–66.
- [13] V. Stamatis, M. Salampasis, K. Diamantaras, A novel re-ranking architecture for patent search, World Patent Information 78 (2024) 102282. doi:10.1016/j.wpi.2024.102282.
- [14] J. Shan, Q. Zhang, C. Shi, M. Gui, S. Wang, U. Naseem, Structural Representation Learning and Disentanglement for Evidential Chinese Patent Approval Prediction, in: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, ACM, 2024, pp. 2014–2023. doi:10.1145/3627673.3679766.
- [15] Y. Shen, Z. Lin, PatentGrapher: A PLM-GNNs Hybrid Model for Comprehensive Patent Plagiarism Detection Across Full Claim Texts, IEEE Access 12 (2024) 182717–182725. doi:10.1109/ACCESS.2024.3508762.
- [16] T. Wei, D. Feng, S. Song, C. Zhang, An extraction and novelty evaluation framework for technology knowledge elements of patents, Scientometrics (2024). doi:10.1007/s11192-024-04990-9.
- [17] H. Ikoma, T. Mitamura, Can AI Examine Novelty of Patents?: Novelty Evaluation Based on the Correspondence between Patent Claim and Prior Art, 2025. doi:10.48550/arXiv.2502.06316. arXiv:2502.06316.
- [18] R. Lee, A. Spangher, X. Ma, PatentEdits: Framing Patent Novelty as Textual Entailment, 2024. arXiv:2411.13477.
- [19] S. Hido, S. Suzuki, R. Nishiyama, T. Imamichi, R. Takahashi, T. Nasukawa, T. Id[^]acute;, Y. Kanehira, R. Yohda, T. Ueno, A. Tajima, T. Watanabe, Modeling Patent Quality: A System for Large-scale Patentability Analysis using Text Mining, Journal of Information Processing 20 (2012) 655–666. doi:10.2197/ipsjjip.20.655.
- [20] N. Kong, U. Dulleck, A. B. Jaffe, S. Sun, S. Vajjala, Linguistic metrics for patent disclosure: Evidence from university versus corporate patents, Research Policy 52 (2023) 104670. doi:10.1016/j.respol.2022.104670.
- [21] J. H. Ashtor, Modeling patent clarity, Research Policy 51 (2022) 104415. doi:10.1016/j.respol.2021.104415.
- [22] M. Mersha, K. Lam, J. Wood, A. K. AlShami, J. Kalita, Explainable artificial intelligence: A survey of needs, techniques, applications, and future direction, Neurocomputing 599 (2024) 128111. URL: <https://www.sciencedirect.com/science/article/pii/S0925231224008828>. doi:<https://doi.org/10.1016/j.neucom.2024.128111>.
- [23] S. Huang, S. Mamidanna, S. Jangam, Y. Zhou, L. H. Gilpin, Can large language models explain themselves? a study of llm-generated self-explanations, 2023. URL: <https://arxiv.org/abs/2310.11207>. arXiv:2310.11207.
- [24] A. Madsen, S. Chandar, S. Reddy, Are self-explanations from Large Language Models faithful?, in: L.-W. Ku, A. Martins, V. Srikumar (Eds.), Findings of the Association for Computational Linguistics:

- ACL 2024, Association for Computational Linguistics, 2024, pp. 295–337. doi:10.18653/v1/2024.findings-acl.19.
- [25] C. Agarwal, S. H. Tanneru, H. Lakkaraju, Faithfulness vs. plausibility: On the (un)reliability of explanations from large language models, 2024. URL: <https://arxiv.org/abs/2402.04614>. arXiv:2402.04614.
 - [26] Nautilus, Inc. v. Biosig Instruments, Inc, 572 U.S. 898 (2014), <https://supreme.justia.com/cases/federal/us/572/898/>, 2014.
 - [27] A. K. et al., Gemma 3 technical report, 2025. URL: <https://arxiv.org/abs/2503.19786>. arXiv:2503.19786.
 - [28] W. Fagen-Ulmschneider, Perception of Probability Words, <https://waf.cs.illinois.edu/visualizations/Perception-of-Probability-Words/>, 2025.
 - [29] K. Tian, E. Mitchell, A. Zhou, A. Sharma, R. Rafailov, H. Yao, C. Finn, C. Manning, Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback, in: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, 2023, pp. 5433–5442. URL: <https://aclanthology.org/2023.emnlp-main.330/>. doi:10.18653/v1/2023.emnlp-main.330.
 - [30] M. Xiong, Z. Hu, X. Lu, Y. Li, J. Fu, J. He, B. Hooi, Can LLMs express their uncertainty? an empirical evaluation of confidence elicitation in LLMs, in: The Twelfth International Conference on Learning Representations, 2024. URL: <https://openreview.net/forum?id=gjeQKFxFpZ>.
 - [31] V. Wang, M. J. Q. Zhang, E. Choi, Improving llm-as-a-judge inference with the judgment distribution, 2025. URL: <https://arxiv.org/abs/2503.03064>. arXiv:2503.03064.
 - [32] M. Yasunaga, L. Shamis, C. Zhou, A. Cohen, J. Weston, L. Zettlemoyer, M. Ghazvininejad, Alma: Alignment with minimal annotation, 2024. URL: <https://arxiv.org/abs/2412.04305>. arXiv:2412.04305.

A. Office Action Parsing Prompt

```
1 ### TASK
2
3 Your task is to extract data related to claim rejections under 35 U.S.C. 112(b) or pre-AIA 35 U.
  S.C. 112, second paragraph (indefiniteness).
4
5 You will be provided with a snippet of text from a USPTO Office Action. You will identify claim
  rejections based on indefiniteness (35 U.S.C. 112(b) or pre-AIA 35 U.S.C. 112, second
  paragraph) and represent them in a JSON format.
6
7 ### OUTPUT FORMAT
8 Strictly adhere to the following JSON schema and reply with nothing but the JSON:
9 ```json
10 {schema}
11 ```
12 ### GUIDELINES
13
14 Here is a detailed guide on how to extract the information:
15
16 1. Identify the section that describes rejections regarding 35 U.S.C. 112(b).
17 2. Identify which claims were rejected based on 35 U.S.C. 112(b). Note this in the "
   rejectedClaims" field. You can either list individual claims as integers or inclusive claim
   ranges (e.g., 1-3 equals 1,2,3) as strings (ranges are preferred if applicable).
18 3. Next, extract the individual reasons for rejection based on 35 U.S.C. 112(b) and list them
   under "rejectionReasons". For each reason, determine the following properties:
19 3.1 "reasonText": The full text describing the rejection. Copy this verbatim. Do not include
   statements regarding interpretation ("For purpose of examination ...").
20 3.2 "reasonContext": The reasoning context. Use this if the reasonText reads in the lines of
   "Claim X is rejected for similar reasons as claim Y", and add the reason of claim Y here.
   Make sure enough context is given to make the rejection understandable on its own. Leave
   this field empty if the reasonText does not refer to another reason.
21 3.3 "claims": The mentioned claims. You can either list individual claims as integers or
   inclusive claim ranges (e.g., 1-3 equals 1,2,3) as strings (ranges are preferred if
   applicable).
22 3.4 "reasonCategory": The category of 112(b) rejection. Choose the most appropriate among
   the following categories. You can use only those exact values! {rejection_categories}
23 3.5 "claimRecitations": The parts of the claim(s) in question that the examiner cited as
   causing the indefiniteness. They are often, but not always, included in quotes. Copy
   verbatim. Every entry here should be a substring of the claim text, i.e., do NOT include
   the reasoning around recitations. Include EVERY SINGLE recitation in the reasonText, even
   if they are repeated or overlapping!.
24
25 Some further notes:
26
27 * **Focus solely on 35 U.S.C. 112(b) (or pre-AIA 35 U.S.C. 112, second paragraph) rejections
   .** Ignore other types of rejections (e.g., 112(a), 102, 103).
28 * **Extract information verbatim. Do not paraphrase or summarize any text. Do not fix spelling
   or grammar errors.** Ignore all newlines and consistently use single quotes.
29 * **Do not mix multiple reasons in a single "rejectionReasons" entry.** Extract rejection
   reasons as fine-grained as possible.
30 * If unsure about the 'reasonCategory', use 'other'.
31 * Output only valid JSON. Do *not* include any explanatory text before or after the JSON. If
   there are no 35 U.S.C. 112(b) claim rejections to extract from the provided text, leave "
   rejectionReasons" empty.
32
33 From now on, I will only send an office action document and you will respond with the
   corresponding JSON representation. Do you understand?
34
```

Listing 1: Prompt used to parse office action into the described JSON schema. schema and rejection_categories are replaced with their respective values.

B. Indefiniteness Examination Prompt

```
1 Examine this patent claim with respect to definiteness.
2
3 ### Guidelines
4
5 - Carefully dissect the claim and its features and search for common patterns that cause
  indefiniteness.
6 - Think step by step and reason about potential issues! You can correct yourself at any point in
  time, only the final verdict counts.
7 - Be very thorough and list all potential issues you can find!
8 - Ultimately, estimate the likelihood of the claim begin rejected due to indefiniteness. Note
  that a single issue renders the entire claim indefinite.
9 - Completely ignore all other aspects like novelty or non-obviousness, focus entirely on
  indefiniteness.
10
11 ### Categories of Indefiniteness
12
13 {indefiniteness_categories}
14
15 ### Tools
16
17 - You are given access to the remaining patent application using simple tools.
18 - You can get other claims (e.g. parent claims) using 'get_claim' and retrieve relevant
  description paragraphs using 'search_description'.
19 - If you want to search for multiple terms, invoke 'search_description' separately for each term
  . This will yield much more specific results than using a combined query.
20 - You can perform as many tool calls as you like; make sure you have all the necessary
  information.
21
22 ### Claim
23
24 {claim}
25
```

Listing 2: Prompt used to examine a given claim with respect to indefiniteness. `indefiniteness_categories` and `claim` are replaced with their respective values.

C. LLM-as-Judge Prompt

```
1 <instruction>
2 You will evaluate the performance of an AI system used to identify indefiniteness issues in
   patent claims. You will be given 2 short text snippets that mention an issue in a claim
   that renders it indefinite, one written by a human examiner, one written by the AI. Your
   task is to determine whether the two text snippets refer to the same issue or not. Ignore
   differences in phrasing and evaluate only whether they pinpoint the same issue. You can
   always assume that both text snippets talk about the same claim. Use the scale shown below.
3 </instruction>
4 <scale>
5 1: The reasons are completely unrelated and address different concerns.
6 2: The reasons address distinct aspects of the claim with minimal overlap.
7 3: The reasons overlap in some areas but also have notable differences.
8 4: The reasons are closely related and largely address the same issue in the claim.
9 5: The reasons are essentially identical and address the same issue in the claim.
10 </scale>
11 <examples>
12   <example>
13     <text-1>
14 The phrase 'like' (stated in all the claims) renders the claim(s) indefinite because the claim(s)
   ) include(s) elements not actually disclosed (those encompassed by 'like'), thereby
   rendering the scope of the claim(s) unascertainable. See MPEP 2173.05(d).
15     </text-1>
16     <text-2>
17 The claim does not specify how these 'component variables' are utilized or how they relate to
   the method described in Claim 13, which is critical to understanding the claim's scope.
18     </text-2>
19     <score>          1          </score>
20   </example>
21   <example>
22     <text-1>
23 The phrase 'like' (stated in all the claims) renders the claim(s) indefinite because the claim(s)
   ) include(s) elements not actually disclosed (those encompassed by 'like'), thereby
   rendering the scope of the claim(s) unascertainable. See MPEP 2173.05(d).
24     </text-1>
25     <text-2>
26 The use of 'like' in 'example component variables of solution automation & interface analysis (
   like 'solution automation workflow variables')' suggests that these might be examples
   rather than exhaustive lists, leading to potential ambiguity.
27     </text-2>
28     <score>          5          </score>
29   </example>
30   <example>
31     <text-1>
32 Claim 1 recites the limitation 'a corresponding link instruction' (line 2) and 'at least one
   link instruction' (line 12). It would be unclear to one having ordinary skill in the art
   whether the above limitations are intended to be identical to, common to, or distinct from
   one another.
33     </text-1>
34     <text-2>
35 The term 'link instruction' is not clearly defined in the specification, leading to ambiguity in
   the scope of the claim.
36     </text-2>
37     <score>          4          </score>
38   </example>
39 </examples>
40 Apply this scheme, as shown in the examples, to the text snippets shown below.
41 <text-1>{reason_1}</text-1>
42 <text-2>{reason_2}</text-2>
43 Before you settle on a score, summarize the similarities and differences between the two reasons
   for indefiniteness.
```

Listing 3: Prompt used to evaluate the similarity of two reasons for indefiniteness. reason_1 and reason_2 are replaced with their respective values.