

LegalBench.PT: A Benchmark for Portuguese Law

Beatriz Canaverde^{1,2}, Telmo Pessoa Pires³, André F. T. Martins^{1,2,4,5}

¹Instituto de Telecomunicações, ²Instituto Superior Técnico, Universidade de Lisboa

³Equall, ⁴TransPerfect, ⁵ELLIS Unit Lisbon

beatriz.canaverde@tecnico.ulisboa.pt

Abstract

The growing use of LLMs in the legal field has spurred the creation of benchmarks across various jurisdictions and languages, yet none specifically focused on the Portuguese legal system. We present LegalBench.PT, the first comprehensive benchmark for Portuguese law. To develop it, we first collect long-form questions and answers from real law exams, and then use GPT-4o to convert them into multiple-choice, true/false, and matching formats. Once generated, the questions are filtered and processed to improve the quality of the dataset. To ensure accuracy and relevance, we validate our approach by having a legal professional review a sample of the generated questions. Although synthetically generated, we show that their basis in human-created exams and our rigorous filtering and processing methods applied result in a reliable benchmark for assessing LLMs' legal knowledge and reasoning abilities. Finally, we evaluate the performance of leading LLMs on LegalBench.PT and investigate potential biases in GPT-4o's responses. We also assess the performance of Portuguese lawyers on a sample of questions to establish a baseline for model comparison and validate the benchmark.

1 Introduction

Large Language Models (LLMs) have shown impressive capabilities (Anthropic, 2024b; Jiang et al., 2024; OpenAI, 2024; Dubey et al., 2024), driving interest in their legal applications to improve the efficiency and accessibility of legal services. Research has explored developing legal-specific LLMs (Colombo et al., 2024a,b; Junior et al., 2026; Zhou et al., 2024), curating training datasets (Henderson et al., 2022; Niklaus et al., 2024a,b), and building benchmarks for evaluation (Chalkidis et al., 2022a,b; Niklaus et al., 2023; Guha et al., 2023; Fei et al., 2024; Joshi et al., 2024; Stern et al., 2024; Hwang et al., 2022). However, these efforts are often tailored to specific legal systems and ju-

risdictions, limiting their applicability to other contexts. Differences in the systems (e.g., civil law vs common law) and the usual reliance on jurisdiction-specific laws mean that advances in one context are often not transferable to others.

European Portuguese has seen limited research, particularly in the legal field (Rodrigues et al., 2023; Santos et al., 2024; Lopes et al., 2024; Melo et al., 2023), and no standardized benchmark exists yet for Portuguese law. While multilingual legal datasets exist (Chalkidis et al., 2021; Aumiller et al., 2022; Niklaus et al., 2023), they omit many important areas of the Portuguese system and focus only on classification and summarization tasks. To address this, we introduce LegalBench.PT,¹ the first benchmark measuring LLMs' legal knowledge and reasoning across key areas of Portuguese law.

We create LegalBench.PT by defining a taxonomy of the Portuguese law and collecting exams from a leading law school in Portugal. Since these exams mainly feature long-form questions unsuitable for automatic evaluation, we instruct GPT-4o (OpenAI, 2024) to convert them into multiple-choice, true/false, and matching questions. Then, we filter the generated dataset to remove duplicates and undesirable instances, and shuffle the alternative options in multiple-choice and matching questions to minimize biases. The final dataset includes 4,723 questions spanning 31 distinct legal areas. A subset of LegalBench.PT is reviewed by a lawyer, and, as expected with synthetic data, there is some noise: 12% of answers are incorrect, and 15% have suboptimal legal terminology or need rephrasing. We compare leading LLMs on LegalBench.PT, finding that GPT-4o and Claude-3.5-Sonnet (Anthropic, 2024a) are the strongest models, closely followed by Claude-3-Opus (Anthropic, 2024b) and Llama-3.1-405B (Dubey et al.,

¹Data and code available at: <https://github.com/deep-spin/LegalBench.PT>

2024). To test for potential biases from GPT-4o-generated content, we repeat the data creation process with Claude-3.5-Sonnet on a few legal areas and evaluate both models on both datasets, finding no significant differences. Finally, we assess Portuguese lawyers on a sample of questions, and observe that their performance is usually closer to that of the lower-performing Llama-3.1-8B and Mixtral-8x7B. This assessment suggests the presence of ambiguous questions and confirms the previously reported rate of incorrect gold answers.

2 Related Work

Legal evaluation datasets have traditionally focused on tasks that language models learn through fine tuning. These datasets, often derived from public online sources or expert annotations, include tasks such as document review (Hendrycks et al., 2021b; Wang et al., 2023), judgment prediction (Chalkidis et al., 2019; Malik et al., 2021), case summarization (Shen et al., 2022), information extraction (Yao et al., 2022), among others (Chalkidis et al., 2022a; Niklaus et al., 2023; Hwang et al., 2022). Although valuable, they do not fully capture the broader capabilities of LLMs in legal contexts.

Recent efforts have shifted toward LLM-oriented benchmarks. MMLU (Hendrycks et al., 2021a), a general knowledge English multiple-choice benchmark, includes a subset of legal questions useful for preliminary assessments, but not always aligned with specific legal systems or jurisdictions. Professional certification exams offer more tailored evaluations (Zhong et al., 2020; Katz et al., 2024; Junior et al., 2026), but often fall short in comprehensively assessing LLMs’ practical use cases.

LegalBench (Guha et al., 2023) marks the first collaborative effort to benchmark legal reasoning for American law. It integrates existing and expert-crafted datasets to assess practical reasoning skills, such as issue identification, rule applicability, and text interpretation. Their structured approach helps professionals understand models’ utility and limitations. Other benchmarks (Fei et al., 2024; Dai et al., 2025) similarly group tasks to separately measure legal knowledge, inference, and application.

3 Data Collection

Portuguese law can be grouped into 5 main areas, following a taxonomy developed with the support of a legal expert: **1) Public Law**: regulates relationships between private entities and the State; **2)**

Private Law: governs relationships between private entities without State involvement; **3) Public-Private Law**: addresses issues spanning both Public and Private Law, depending on context; **4) Public International Law**: manages relations between states and international organizations; **5) EU and Community Law**: governs interactions between EU member states and EU institutions. Figure 1 shows the full taxonomy adopted. Public International and EU and Community have no subareas.




Figure 1: Full taxonomy of the Portuguese law adopted. Names marked with * are not included in the benchmark due to lack of source data.

3.1 Data Processing

We manually collected 341 exams with solutions from the Faculty of Law at the University of Lisbon² covering the academic years 2021-2024. With a lawyer’s help, we categorized them within

²<https://www.fd.ulisboa.pt/cursos/licenciatura/avaliacao/exames-escritos/>

our taxonomy. The exams, mostly composed of open-ended questions and rarely multiple-choice, required students to thoroughly analyze cases and take well-supported positions on legal issues. Downloaded as PDFs, we extracted the text using PyMuPDF³ and manually segmented the questions, separating independent questions and grouping related ones together. We also removed page numbers and headers, and relocated footnotes that appeared mid-text to more convenient positions.

3.2 Statistics

Our dataset covers a wide range of legal areas, with 13 distinct fields in Public Law and 15 in Private Law. In contrast, the Public-Private group is limited to a single field (Competition Law) due to limited exam availability. Table 1 summarizes the number of distinct exam exercises per area. Although our taxonomy distinguishes several Commercial Law subfields, exams from courses named “Commercial Law” could not be mapped to more specific categories due to their broad content. As a result, Commercial Law is treated as a standalone field rather than an aggregate of its subfields. Conversely, we do not report separate numbers for “Civil Law”, as it is simply an aggregate of its subfields.

4 Benchmark Creation

In preliminary experiments, we evaluated GPT-4o on long-form exam questions following Zheng et al. (2023), where the model graded its own responses against the gold answers. However, this method proved ineffective due to the detailed nature of the answers and lack of clear error penalization criteria. Therefore, we converted the exam questions into more easily assessable formats: multiple-choice, true/false, and matching questions.

As Figure 2 illustrates, our pipeline comprises three phases: 1) providing GPT-4o with exam questions and answers to generate new question-answer pairs; 2) filtering the dataset to remove undesired entries and duplicates; 3) shuffling the options in multiple-choice and matching questions to minimize biases and balance correct answer distributions. Experiments used GPT-4o⁴ (OpenAI, 2024) at a 0.001 temperature, from May to August 2024.

³<https://pypi.org/project/PyMuPDF/>

⁴Version: gpt-4o-2024-05-13.

Fields of Law	Exams	Benchmark
Public	274	2115
Environmental Law	11	52
Administrative Law	39	227
Constitutional Law	58	206
Energy Law	12	21
Public Finance Law	25	208
Financial Law	1	16
Tax Law	19	232
Criminal Law	38	390
Administrative Procedure Law	19	116
Civil Procedure Law	20	272
Criminal Procedure Law	20	274
Labor Procedure Law	2	3
Urban Planning Law	10	98
Private	252	2057
Contract Law	25	212
Family Law	34	220
Law of Obligations	17	218
Property Law	34	502
Succession Law	1	17
Commercial Law	23	240
Banking Law	2	30
Maritime Law	7	8
Corporate Law	12	154
Securities Law	9	45
Transportation Law	15	26
Aviation Law	19	31
Insolvency Law	4	37
Private International Law	18	110
Labor Law	32	207
Public-Private (Competition Law)	12	175
Public International	72	212
EU and Community	28	164
Total	626	4723

Table 1: Distribution of exam and benchmark questions across fields of law.

4.1 Question Generation

To handle the different types of exam exercises, we implemented three generation approaches:

1. Providing the model with a group of short, independent questions, typically asking for topic development or concept differentiation, along with their brief answers. We instruct the model to output new question-answer pairs.
2. Presenting the model one long, complex question-answer pair at a time. These questions are usually long because they include a case analysis statement. We instruct the model to identify the problem statement and generate new question-answer pairs.
3. Feeding the model a set of exam questions and answers sharing a common problem statement. These questions sometimes introduce new assumptions with critical information, as a continuation of the statement. We ask the model to identify the statement and assumptions, and generate new question-answer pairs.




Figure 2: LegalBench.PT construction pipeline.

Within each field of law, questions were divided among these three groups. Each exam question-answer pair was shown to GPT-4o twice: in the second round, it also received the respective question-answer pairs generated in the first. This was useful for approaches 2 and 3 because answers were lengthy, and GPT-4o’s output token limit of 4,096 often caused incomplete coverage of all topics. We designed six prompts in total, two per approach.

We asked the model to generate six question types: 1) **multiple-choice**, with one correct option; 2) **cloze tasks**, fill-in-the-blank multiple-choice; 3) **case analysis**, multiple-choice where the model creates a small case for analysis; 4) **true/false**, classifying sentences as “True” or “False”; 5) **multiple selection**, multiple-choice with more than one correct option; 6) **matching questions**, pairing items between two columns. Case analysis was used only in approach 1, as the exam questions were very theoretical. Matching questions were limited to approaches 2 and 3, since those created with approach 1 tended to test knowledge of specific laws and articles—an issue discussed in Section 4.2. The remaining types were used in all approaches.

Finally, we combined each generated question with its respective statement and assumptions to preserve context. For approach 2, we prepended the statement to the questions. For approach 3, we prepended the statement and all assumptions preceding each question. However, in the filtering and shuffling stages (Sections 4.2 and 4.3), we focus on the questions disregarding statements and assumptions. See Appendix A for examples.

4.2 Filtering

After generation, a Portuguese lawyer reviewed random outputs from the different fields of law and identified two main issues: questions referencing specific articles and laws, and repeated questions.

Articles and Laws. Many questions referenced specific articles and laws by number (e.g., “Artigo 103.º, n.º 2 da Constituição”⁵) to test the model’s knowledge of legal texts. We removed these questions for two reasons. First, laws and articles are frequently updated, and their relevance depends on the timing of events. Since some questions and statements did not specify dates, they could be misleading. Second, lawyers are not expected to memorize legal provisions. We removed all questions matching both of the following patterns in lowercase: `(?<!^)(?<!\n)\d+` and `n\.\d|nº|[^a-z](artigos?|art|reg)[^a-z]`.

Repeated Questions. Many similar questions were found within batches generated from the same exam exercise(s) provided to GPT-4o (first + second rounds). To identify repetitions at scale, within each batch, we first compared questions of the same type, and then compared questions of different types with each other. Multiple-choice, cloze tasks, case analysis, and multiple selection are multiple-choice variants: we compared them with and without answer options. Matching questions were compared before and after a processing step that consisted of removing the first line (usually “Match the items...”) and the letters/numbers identifying the items (Appendix B shows an example). We did not process true/false questions for comparison.

To filter duplicates, we combined lexical and semantic similarity methods. For each pair of processed questions, we selected the appropriate ROUGE-L variant (sentence- or summary-level), we computed precision and recall scores using the Python implementation of ROUGE-L.⁶ If $\min(\text{precision}, \text{recall})$ met or exceeded a given threshold, one question was removed. Additionally, we computed sentence embeddings using a mul-

⁵Article 103.º, n.º 2 of the Constitution.

⁶<https://pypi.org/project/rouge-score/>

tilingual Transformer model trained for semantic similarity detection,⁷ measuring cosine similarity between embeddings. Pairs meeting or exceeding a given threshold were filtered out.

We manually determined the best threshold for each scenario. Before comparison, we randomized the order of questions to avoid biasing the distribution across types. Once a question met the elimination criteria, it was excluded from further comparisons. A summary of the processing steps, thresholds, and methods is provided in Appendix B.

4.3 Option Shuffling

After filtering, we analyzed the frequency of the correct answer options within each multiple-choice variant. We found a highly imbalanced distribution, with GPT-4o exhibiting a marked tendency to generate questions with the second option correct.

To minimize potential evaluation biases and balance the correct options, we randomized the order of answer choices for these question types. Upon inspection, we encountered 18 questions with answer options such as “Both options a) and b) are possible.”, “None of the above.”, and “All of the above.”⁸ In these cases, we randomized the options taking such dependencies into account. We also shuffled the item order in matching questions.

4.4 Dataset Statistics

We obtained a total of 4,723 new questions (down from 10,951 before filtering), as shown in Table 1. The difference between the distribution of exam exercises and benchmark questions is large. It stems from exam questions with long answers being converted into multiple new questions, while short answers tend to produce only one or two.

Out of 31 fields, 14 contain over 200 questions each, and 6 have between 98 and 200 questions, covering nearly all fundamental fields of Portuguese law and representing 94% of the benchmark. The remaining 11 areas have fewer than 60 questions. Of these, 7 contain between 20 and 60 questions, offering a moderate assessment, while 4 have fewer than 20 questions, making them insufficient for a comprehensive evaluation. Nonetheless, we include these for representativeness. See Appendix C for further details.

⁷<https://huggingface.co/sentence-transformers/paraphrase-multilingual-mpnet-base-v2>

⁸“Ambas as opções a) e b) são possíveis.”, “Nenhuma das anteriores.”, “todas as anteriores”.

4.5 Question Analysis

A practicing lawyer reviewed 2 small samples of questions and confirmed their relevance and usefulness in assessing a model’s legal knowledge. A few lower-quality cases were also identified, as discussed below.

Corrections. We randomly selected two groups of questions from different fields⁹ for analysis. Out of 33 questions, 4 contained incorrect answers, and 5 required minor improvements, such as correcting legal terminology or rephrasing for clarity.

“Easy” Questions. In a sample of 64 statement-based questions across three fields¹⁰ (19 questions from the previous sample, 45 randomly chosen), 12 primarily tested interpretation skills rather than legally relevant content. We considered these questions too straightforward and tried to filter them out with GPT-4o, with no success.

5 Evaluation

Evaluations are conducted in a zero-shot setting. All prompts specify the legal area, instructions on the question type (e.g., whether one or multiple options are correct), and the expected answer format. Appendix D shows an example template.

From the outputs, we extract letter options, true/false classifications, and matching pairs according to the prescribed format. We evaluate multiple-choice, cloze tasks, case analysis, and true/false questions separately using balanced accuracy. Multiple selection and matching questions use the F_1 score. We aggregate results from different types and fields using a weighted average.

5.1 Model Performance

We evaluated on LegalBench.PT: GPT-4o¹¹ (OpenAI, 2024), GPT-4o-mini¹² (OpenAI, 2024), Claude-3-Opus¹³ (Anthropic, 2024b), Claude-3.5-Sonnet¹⁴ (Anthropic, 2024a), Llama-3.1-8B, Llama-3.1-70B, Llama-3.1-405B (Dubey et al., 2024)¹⁵, and Mixtral-8x7B (Jiang et al., 2024). We used the instruct versions of the open-source models. We set the temperature to 0.01. Evaluations were conducted in September 2024.

⁹EU and Community, and Civil Procedure Law.

¹⁰EU and Community, Civil Procedure, Competition Law.

¹¹Version: gpt-4o-2024-05-13.

¹²Version: gpt-4o-mini-2024-07-18.

¹³Version: claude-3-opus-20240229.

¹⁴Version: claude-3-5-sonnet-20240620.

¹⁵We used the Llama Turbo versions available at together.ai.

Fields of Law	GPT-4o	GPT-4o-mini	Claude-3-Opus	Claude-3.5-Sonnet	Llama-3.1-8B	Llama-3.1-70B	Llama-3.1-405B	Mixtral-8x7B
Public	84.4	78.3	82.5	84.2	65.8	79.4	81.8	66.5
Environmental Law	94.8	85.0	88.0	89.9	71.3	87.7	89.8	76.4
Administrative Law	86.9	80.4	82.7	88.7	64.4	80.2	81.3	66.8
Constitutional Law	88.0	81.9	86.8	89.2	67.5	85.4	87.1	70.1
Energy Law	94.3	86.0	79.2	81.4	71.5	83.8	80.6	89.2
Public Finance Law	87.6	83.5	87.2	88.8	73.3	86.8	89.3	72.9
Financial Law	95.6	87.2	96.5	99.1	75.3	94.2	95.6	75.6
Tax Law	89.1	77.2	87.8	90.1	58.0	76.0	80.0	60.0
Criminal Law	79.0	71.7	77.9	79.4	62.7	73.0	77.6	61.8
Administrative Procedure Law	80.5	81.1	78.8	80.2	68.5	78.9	81.1	68.2
Civil Procedure Law	79.8	75.2	78.8	78.0	65.2	77.4	78.5	66.6
Criminal Procedure Law	80.9	73.9	76.7	77.7	62.9	74.9	76.6	59.3
Labor Procedure Law	66.7	66.7	66.7	66.7	33.3	66.7	33.3	0.0
Urban Planning Law	93.9	94.2	94.6	95.3	82.4	94.8	95.6	87.1
Private	84.8	80.2	84.3	84.1	69.1	79.9	83.9	70.6
Contract Law	81.2	77.1	81.7	82.0	68.2	77.6	81.3	64.9
Family Law	81.3	72.9	82.6	81.6	64.3	77.4	77.7	65.0
Law of Obligations	84.0	76.3	79.8	80.7	61.1	78.7	80.6	65.8
Property Law	85.1	79.7	84.0	83.8	69.6	79.4	84.2	70.4
Succession Law	60.4	66.9	72.9	75.9	56.9	58.3	62.7	45.9
Commercial Law	88.6	82.5	89.9	86.3	69.8	78.6	85.3	74.7
Banking Law	98.3	97.7	98.3	99.2	95.3	96.6	96.9	90.7
Maritime Law	87.5	60.0	87.5	87.5	68.8	56.2	75.0	87.5
Corporate Law	79.0	79.2	79.3	79.1	67.9	77.1	80.8	69.2
Securities Law	92.9	89.2	84.1	85.7	77.1	89.7	91.2	78.9
Transportation Law	67.2	88.9	70.2	84.7	48.2	74.4	77.0	63.6
Aviation Law	91.5	85.9	89.5	87.0	67.9	90.9	85.1	70.6
Insolvency Law	71.7	85.9	75.4	80.9	73.7	68.5	84.4	66.1
Private International Law	86.6	82.7	90.0	89.1	71.2	83.7	92.1	71.9
Labor Law	93.1	87.3	89.9	89.9	78.1	87.9	91.4	81.5
Public-Private (Competition Law)	93.6	93.7	94.4	95.0	87.9	94.5	95.5	81.2
Public International	89.6	84.5	90.5	90.4	75.4	86.5	90.4	77.2
EU and Community	91.3	83.7	90.4	92.8	69.4	85.3	89.0	77.0
Overall	85.4	80.1	84.3	85.1	68.6	80.7	83.8	69.7

Table 2: Model performance (%) across the different fields of law.

Table 2 reports model performance across legal areas (in percentage). Overall, GPT-4o is the best-performing model, closely followed by Claude-3.5-Sonnet. Notably, Llama-3.1-405B trails behind GPT-4o with a difference of only 1.6%, and Llama-3.1-70B achieves results similar to GPT-4o-mini across most areas. Apart from Llama-3.1-8B and Mixtral-8x7B, whose results are noticeably weaker, likely due to their smaller size, all remaining models show similar performance. This suggests either equivalent capabilities or that LegalBench.PT is not challenging enough to differentiate these models.

Certain legal fields, such as Labor Procedure and Succession Law, consistently show lower scores, likely due to their smaller question sets. Similarly, Maritime, Transportation, Environmental, Energy, Financial, Banking, Aviation, and Insolvency Law have small question sets, making their results less informative. Conversely, areas like Urban Planning, Banking, Labor, and Competition Law consistently show high performance across all models, potentially due to less challenging questions.

Performance in the Public-Private area is not

fully comparable to the broader Public and Private fields due to the smaller and less diverse question set, including only Competition Law. Compared to the Public and Private areas, all models perform better in Public International and EU and Community Law. This likely stems from smaller and less diverse question sets, and coverage of topics better represented in the training data, as they are common to various jurisdictions and languages. Appendix E provides results by question type.

5.2 Investigating GPT-4o Bias

Table 2 shows GPT-4o as one of the top-performing models, with overall results similar to Claude-3.5-Sonnet. Since GPT-4o was used to generate the dataset, one could argue this comparison is unfair. To test this hypothesis, we repeated the data creation process (same prompts, filtering, and shuffling) with Claude-3.5-Sonnet. This time, we focused on 6 key areas¹⁶, generating 1,389 questions, slightly more than the 1,335 questions in the origi-

¹⁶Administrative, Civil Procedure, Family, Commercial, Public International, EU and Community Law.

nal dataset in these areas, and similarly distributed.

Table 3 compares GPT-4o and Claude-3.5-Sonnet on both datasets. We observe that Claude-3.5-Sonnet tends to outperform GPT-4o in both cases, with differences between the models ranging from 0.6% to 3.5% on Claude-generated data, and from 0.3% to 1.8% on GPT-generated data.

Fields of Law	GPT Generated		Claude Generated	
	GPT	Claude	GPT	Claude
Administrative Law	86.9	88.7	82.3	85.8
Civil Procedure Law	79.8	78.0	80.3	81.4
Family Law	81.3	81.6	84.9	87.0
Commercial Law	88.6	86.3	78.8	82.1
Public International Law	89.6	90.4	91.6	92.2
EU and Community Law	91.3	92.8	92.8	91.6
Overall	85.8	85.7	84.5	86.2

Table 3: GPT-4o and Claude-3.5-Sonnet performance (%) on the questions generated by each of these models across different fields of law.

We conducted a Wilcoxon signed-rank test to assess whether the differences between the models’ performances on these fields of law follow a symmetric distribution around zero. For the GPT-generated data, we obtained a p -value of 100%, and for the Claude-generated data, a p -value of 15.6%. In both cases, there is no evidence for considering the performance differences between GPT-4o and Claude-3.5-Sonnet statistically significant.

5.3 Human Evaluation

We assessed Portuguese lawyers on LegalBench.PT to establish a baseline for model comparison and validate the benchmark. Given the high cost of lawyers’ time and the size of our dataset, we randomly selected 1,000 questions from 10 fundamental areas of law (100 questions from each area).¹⁷ To minimize participants’ effort, we randomly divided the questions into 20 groups of 50 questions each, and assigned each lawyer at least one group.

In total, 22 lawyers participated, providing 1,183 answers across 17 distinct groups. Two lawyers answered two groups each, while the remaining answered a single group. For 350 questions across 7 groups, we collected paired answers, meaning each question has responses from two lawyers. The remaining questions received a single response. All groups have 50 answered questions, except for one that is incomplete with only 33 questions.

¹⁷Administrative, Constitutional, Criminal, Civil Procedure, Contract, Family, Commercial, Labor, Public International, and EU and Community Law.

5.3.1 Performance

We analyze each group individually, comparing the lawyers’ performance with that of the top-performing models (GPT-4o, Claude-3.5-Sonnet) and the lower-performing ones (Llama-3.1-8B, Mixtral-8x7B). Table 4 shows that lawyers usually perform similarly to Llama-3.1-8B and Mixtral-8x7B, sometimes even worse than all models (groups 4, 8, 16). Only four people surpass GPT-4o or Claude-3.5-Sonnet (groups 6, 7, 11).

Groups	Person 1	Person 2	GPT -4o	Claude-3.5 -Sonnet	Llama-3.1-8B	Mixtral-8x7B
1	79.0	83.3	94.4	88.5	75.3	77.9
2	73.7	66.9	82.5	81.0	68.3	75.9
3	61.7	80.3	87.8	87.4	60.6	66.6
4	59.1	67.1	91.6	95.3	80.6	81.5
5	66.3	84.2	92.8	88.2	63.4	71.5
6	65.8	91.4	88.1	87.0	61.7	72.2
7	77.1	75.6	74.8	75.9	67.6	74.7
8	56.1	-	71.7	78.2	72.7	69.1
9	67.2	-	80.6	86.4	69.3	73.1
10	57.8	-	95.2	93.5	68.2	59.5
11	82.9	-	82.0	85.8	71.9	72.1
12	83.0	-	88.4	88.2	83.1	77.3
13	64.1	-	75.1	83.5	62.4	54.8
14	61.3	-	73.4	74.3	65.0	56.3
15	77.0	-	79.7	82.4	56.1	59.7
16	60.0	-	90.9	89.0	73.5	78.4
17	78.4	-	88.8	90.8	74.7	72.0

Table 4: Human and model performance (%) on disjoint groups of questions from LegalBench.PT. Person 1 and Person 2 differ across groups. All groups have 50 questions, except for group 16, which has only 33.

Examining groups 1-7, we see some performance disparities between participants within the same group. While groups 1, 2, 4, and 7 show differences of less than 10%, groups 3, 5, and 6 exhibit larger gaps of 18.6%, 17.9%, and 25.6%, respectively. Group 4 stands out for both participants’ low performance, contrasting sharply with the high scores of all models in this group.

We reached out to person 1 from group 5 and person 1 from group 6. Each answered 50 questions, totaling 100. Of these, 32 were answered incorrectly according to our gold standards. After reviewing the correct solutions, participants provided feedback on 13 questions:

- 6 questions were considered ambiguous, allowing different views. Their gold standards are correct, but so are the lawyers’ answers.
- For 7 questions, the gold answers were indeed correct and lawyers answered incorrectly. They said answering multiple lengthy questions consecutively led to overlooked details.

Assuming correctly answered questions are unambiguous, we estimate that 15 questions (14.8%)

out of the 100 considered are ambiguous. Due to limited resources, we were unable to further investigate the causes of the low results.

5.3.2 Agreement

For groups with two lawyers (groups 1-7), we measured the agreement between their answers. For each group, we iterated over multiple-choice, cloze tasks, case analysis, and true/false questions, computing the percentage of answers on which participants agree. For multiple selection and matching questions, we computed the Jaccard similarity between each pair of answers, and averaged across all pairs. The agreement per group is the aggregation of scores across question types using a weighted average. We also computed the percentage of agreed-upon answers matching the gold standards.

Results in Table 5 show relatively low agreement values, averaging 62.4%. They are not surprising given the performance results in Table 4. Group 4 has significantly lower agreement, aligning with the weaker performances previously discussed. Questions on which participants disagree likely represent, at least in part, ambiguous questions.

Groups	Agreement (%)	Accuracy (%)
1	69.7	94.3
2	64.0	79.7
3	57.8	89.7
4	45.8	88.1
5	63.6	91.7
6	63.8	88.2
7	71.8	89.7
Overall	62.4	88.9

Table 5: Agreement rates between participants’ answers and accuracy of the agreed-upon answers on disjoint groups of questions from LegalBench.PT.

The high accuracy scores on the agreed-upon answers suggests the respective gold standards are correct. Groups 1 and 3-7 show error rates between 5.7% and 12%, while group 2 presents a higher rate of 20.3%. These errors may indicate incorrect gold answers. In Section 4.5, the lawyer who reviewed a sample of the generated questions identified 12.1% of incorrect gold standards. This value closely matches the 11.1% average error rate observed here. Since 11.1% was computed from a broader sample, it may be a better estimate of the true rate of incorrect answers in LegalBench.PT.

We also analyzed the questions by legal field rather than by group. Table 6 shows the number

of questions per field, the agreement between participants, and the accuracy of the agreed-upon answers. Each field includes questions from all seven groups, except for Public International Law and EU and Community Law, which only include six. The agreement and accuracy values do not deviate from those observed in Table 5, showing a similar degree of variation across fields as across groups. Assuming the error rate indeed reflects incorrect gold standards, these results suggest that such errors are relatively balanced across fields, with only Administrative Law exhibiting a slightly higher proportion of incorrect answers.

Fields of Law	Total Questions	Agreement (%)	Accuracy (%)
Administrative Law	40	68.3	80.2
Constitutional Law	39	61.5	95.4
Criminal Law	40	56.9	86.0
Civil Procedure Law	34	58.8	87.5
Contract Law	34	54.9	91.3
Family Law	33	57.4	87.0
Commercial Law	35	70.2	89.9
Labor Law	38	71.1	93.1
Public International Law	23	56.1	91.7
EU and Community Law	34	65.2	88.5
Overall	350	62.4	88.9

Table 6: Number of questions analyzed per field, agreement rates between participants’ answers, and accuracy of the agreed-upon answers.

6 Conclusion

We introduced LegalBench.PT, the first legal benchmark for assessing LLMs’ knowledge and reasoning abilities in Portuguese law. Structured around a taxonomy we developed, it enables performance comparisons across key domains. We demonstrated that synthetic data generation, grounded on law school exams and with adequate post-processing and filtering, can produce a good benchmark for model evaluation. Human analysis showed that most questions are relevant and useful, despite the entire dataset being synthetic. Future work may expand underrepresented areas and introduce tasks that assess skills beyond legal knowledge and reasoning, such as contract analysis or case summarization. Filtering out easy, ambiguous, or incorrect items, and expanding human review will further strengthen the benchmark’s reliability.

Limitations

Although derived from real exams, LegalBench.PT does not capture the full scope of the original exam questions. These also assess reasoning and argu-

mentation skills, which cannot be fully captured by multiple-choice and similar question formats. As a synthetic dataset, LegalBench.PT contains some noise, such as incorrect gold answers, improper legal terminology, and ambiguous questions. While our analysis uncovered these issues, the reviewed samples were limited and thus may not represent the entire benchmark.

Ethical Considerations

LegalBench.PT was built from real Portuguese law exam questions, synthetically rewritten into objective formats using GPT-4o. While we took steps to ensure the accuracy and relevance of the generated questions, the dataset may contain biases introduced by GPT-4o or inherent to the Portuguese legal system itself. It is intended solely for research purposes and does not substitute real professional advice. Because it focuses on Portuguese law, results should not be generalized to other legal systems or jurisdictions without caution. Users are encouraged to apply the dataset responsibly, mindful of its limitations and context.

Acknowledgments

This work was supported by Equall, the project DECOLLAGE (ERC-2022-CoG 101088763), and by FCT/MECI through national funds and when applicable co-funded EU funds under UID/50008: Instituto de Telecomunicações. We are very grateful to Leonor Melo Ribeiro, whose expertise as a lawyer contributed to the development of this project, and to all the lawyers who took the time to participate in the survey.

References

- Anthropic. 2024a. [Claude 3.5 Sonnet](#). Accessed 30-Oct-2024.
- Anthropic. 2024b. [The Claude 3 Model Family: Opus, Sonnet, Haiku](#).
- Dennis Aumiller, Ashish Chouhan, and Michael Gertz. 2022. [EUR-lex-sum: A multi- and cross-lingual dataset for long-form summarization in the legal domain](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7626–7639, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. [Neural legal judgment prediction in English](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4317–4323, Florence, Italy. Association for Computational Linguistics.
- Ilias Chalkidis, Manos Fergadiotis, and Ion Androutsopoulos. 2021. [MultiEURLEX - a multi-lingual and multi-label legal document classification dataset for zero-shot cross-lingual transfer](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6974–6996, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. 2022a. [LexGLUE: A benchmark dataset for legal language understanding in English](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4310–4330, Dublin, Ireland. Association for Computational Linguistics.
- Ilias Chalkidis, Tommaso Pasini, Sheng Zhang, Letizia Tomada, Sebastian Schwemer, and Anders Søgaard. 2022b. [FairLex: A multilingual benchmark for evaluating fairness in legal text processing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4389–4406, Dublin, Ireland. Association for Computational Linguistics.
- Pierre Colombo, Telmo Pires, Malik Boudiaf, Rui Melo, Gabriel Hautreux, Etienne Malaboeuf, Johanne Charpentier, Dominic Culver, and Michael Desa. 2024a. [SaulLM-54B & SaulLM-141B: Scaling Up Domain Adaptation for the Legal Domain](#). In *The Thirtieth Annual Conference on Neural Information Processing Systems*.
- Pierre Colombo, Telmo Pessoa Pires, Malik Boudiaf, Dominic Culver, Rui Melo, Caio Corro, Andre F. T. Martins, Fabrizio Esposito, Vera Lúcia Raposo, Sofia Morgado, and Michael Desa. 2024b. [SaulLM-7B: A pioneering Large Language Model for Law](#).
- Yongfu Dai, Duanyu Feng, Jimin Huang, Haochen Jia, Qianqian Xie, Yifang Zhang, Weiguang Han, Wei Tian, and Hao Wang. 2025. [LAIW: A Chinese legal large language models benchmark](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 10738–10766, Abu Dhabi, UAE. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. [The Llama 3 Herd of Models](#).
- Zhiwei Fei, Xiaoyu Shen, Dawei Zhu, Fengzhe Zhou, Zhuo Han, Alan Huang, Songyang Zhang, Kai Chen, Zhixin Yin, Zongwen Shen, Jidong Ge, and Vincent Ng. 2024. [LawBench: Benchmarking legal knowledge of large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7933–7962, Miami, Florida, USA. Association for Computational Linguistics.

- Neel Guha, Julian Nyarko, Daniel E. Ho, Christopher Re, Adam Chilton, Aditya Narayana, Alex Chohlas-Wood, Austin Peters, Brandon Waldon, Daniel Rockmore, Diego Zambrano, Dmitry Talisman, Enam Hoque, Faiz Surani, Frank Fagan, Galit Sarfaty, Gregory M. Dickinson, Haggai Porat, Jason Hegland, and 21 others. 2023. [LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Peter Henderson, Mark Simon Krass, Lucia Zheng, Neel Guha, Christopher D Manning, Dan Jurafsky, and Daniel E. Ho. 2022. [Pile of Law: Learning Responsible Data Filtering from the Law and a 256GB Open-Source Legal Dataset](#). In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. [Measuring Massive Multitask Language Understanding](#). In *International Conference on Learning Representations*.
- Dan Hendrycks, Collin Burns, Anya Chen, and Spencer Ball. 2021b. [CUAD: An Expert-Annotated NLP Dataset for Legal Contract Review](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Wonseok Hwang, Dongjun Lee, Kyoungyeon Cho, Hanuhl Lee, and Minjoon Seo. 2022. [A Multi-Task Benchmark for Korean Legal Language Understanding and Judgement Prediction](#). In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L  lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, and 7 others. 2024. [Mixtral of Experts](#).
- Abhinav Joshi, Shounak Paul, Akshat Sharma, Pawan Goyal, Saptarshi Ghosh, and Ashutosh Modi. 2024. [IL-TUR: Benchmark for Indian legal text understanding and reasoning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11460–11499, Bangkok, Thailand. Association for Computational Linguistics.
- Roseval Malaquias Junior, Ramon Pires, Roseli A. F. Romero, and Rodrigo Nogueira. 2026. Juru: Legal Brazilian Large Language Model from Reputable Sources. In *Intelligent Systems*, pages 121–134, Cham. Springer Nature Switzerland.
- Daniel Martin Katz, Michael James Bommarito, Shang Gao, and Pablo Arredondo. 2024. [GPT-4 passes the bar exam](#). *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 382(2270):20230254.
- Ricardo Lopes, Joao Magalhaes, and David Semedo. 2024. [Gl  ria: A generative and open large language model for Portuguese](#). In *Proceedings of the 16th International Conference on Computational Processing of Portuguese - Vol. 1*, pages 441–453, Santiago de Compostela, Galicia/Spain. Association for Computational Linguistics.
- Vijit Malik, Rishabh Sanjay, Shubham Kumar Nigam, Kripabandhu Ghosh, Shouvik Kumar Guha, Arnab Bhattacharya, and Ashutosh Modi. 2021. [ILDC for CJPE: Indian legal documents corpus for court judgment prediction and explanation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4046–4062, Online. Association for Computational Linguistics.
- Rui Melo, Pedro A. Santos, and Jo  o Dias. 2023. [A Semantic Search System for the Supremo Tribunal de Justi  a](#). In *Progress in Artificial Intelligence: 22nd EPIA Conference on Artificial Intelligence, EPIA 2023, Faial Island, Azores, September 5–8, 2023, Proceedings, Part II*, page 142–154, Berlin, Heidelberg. Springer-Verlag.
- Joel Niklaus, Veton Matoshi, Pooja Rani, Andrea Galassi, Matthias St  rmer, and Ilias Chalkidis. 2023. [LEXTREME: A multi-lingual and multi-task benchmark for the legal domain](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3016–3054, Singapore. Association for Computational Linguistics.
- Joel Niklaus, Veton Matoshi, Matthias St  rmer, Ilias Chalkidis, and Daniel Ho. 2024a. [MultiLegalPile: A 689GB multilingual legal corpus](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15077–15094, Bangkok, Thailand. Association for Computational Linguistics.
- Joel Niklaus, Lucia Zheng, Arya D. McCarthy, Christopher Hahn, Brian M. Rosen, Peter Henderson, Daniel E. Ho, Garrett Honke, Percy Liang, and Christopher Manning. 2024b. [FLawN-T5: An Empirical Examination of Effective Instruction-Tuning Data Mixtures for Legal Reasoning](#).
- OpenAI. 2024. [GPT-4 Technical Report](#).
- OpenAI. 2024. [GPT-4o mini: advancing cost-efficient intelligence](#). Accessed 12-Oct-2024.
- OpenAI. 2024. [GPT-4o System Card](#).
- Jo  o Rodrigues, Lu  s Gomes, Jo  o Silva, Ant  nio Branco, Rodrigo Santos, Henrique Lopes Cardoso, and Tom  s Os  rio. 2023. [Advancing Neural Encoding of Portuguese with Transformer Albertina PT-*](#), page 441–453. Springer Nature Switzerland.

Rodrigo Santos, João Ricardo Silva, Luís Gomes, João Rodrigues, and António Branco. 2024. [Advancing generative AI for Portuguese with open decoder gervásio PT*](#). In *Proceedings of the 3rd Annual Meeting of the Special Interest Group on Under-resourced Languages @ LREC-COLING 2024*, pages 16–26, Torino, Italia. ELRA and ICCL.

Zejiang Shen, Kyle Lo, Lauren Yu, Nathan Dahlberg, Margo Schlanger, and Doug Downey. 2022. [Multi-LexSum: Real-world Summaries of Civil Rights Lawsuits at Multiple Granularities](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 13158–13173. Curran Associates, Inc.

Ronja Stern, Vishvakshen Rasiah, Veton Matoshi, Srinanda Brügger Bose, Matthias Stürmer, Ilias Chalkidis, Daniel E. Ho, and Joel Niklaus. 2024. [One Law, Many Languages: Benchmarking Multilingual Legal Reasoning for Judicial Support](#). In *ICLR 2024 Workshop on Data-centric Machine Learning Research (DMLR): Harnessing Momentum for Science*.

Steven Wang, Antoine Scardigli, Leonard Tang, Wei Chen, Dmitry Levkin, Anya Chen, Spencer Ball, Thomas Woodside, Oliver Zhang, and Dan Hendrycks. 2023. [MAUD: An expert-annotated legal NLP dataset for merger agreement understanding](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16369–16382, Singapore. Association for Computational Linguistics.

Feng Yao, Chaojun Xiao, Xiaozhi Wang, Zhiyuan Liu, Lei Hou, Cunchao Tu, Juanzi Li, Yun Liu, Weixing Shen, and Maosong Sun. 2022. [LEVEN: A large-scale Chinese legal event detection dataset](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 183–201, Dublin, Ireland. Association for Computational Linguistics.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2020. [JEC-QA: A Legal-Domain Question Answering Dataset](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9701–9708.

Zhi Zhou, Jiang-Xin Shi, Peng-Xiao Song, Xiao-Wen Yang, Yi-Xuan Jin, Lan-Zhe Guo, and Yu-Feng Li. 2024. [LawGPT: A Chinese Legal Knowledge-Enhanced Large Language Model](#).

A Question Generation Details

We recall the three approaches implemented for generating new questions:

1. Providing the model with a group of short, independent question-answer pairs.
2. Feeding the model one question-answer pair at a time. These questions usually include a long problem statement.
3. Presenting the model with a set of exam questions and answers related to a common problem statement. These exam questions sometimes present new assumptions, as a continuation of the problem statement, that may contain critical information.

Tables 7 to 9 show English-translated examples of exam questions fed into GPT-4o, along with part of the corresponding outputs, illustrating the different approaches used and the output templates we requested.

For approach 1, we simply requested the outputs to be sequences of tuples, with each tuple identifying the generated question along with its type and answer, as demonstrated in Table 7. Table 8 illustrates a similar template that we requested for approach 2. Since these exam questions usually present long problem statements, namely case analysis or excerpts from court decisions, we instruct the model to first identify the statement and then present the tuples with the new generated questions. Finally, for approach 3, we adopted a slightly different template displayed in Table 9. In this case, the exam questions sometimes present new assumptions, as a continuation of the problem statement, that may contain critical information. We request the model to identify the problem statement, the assumptions, and the tuples with the new questions.

To get the final version of the generated questions: for approach 2, we joined at the beginning of each question the respective statement; for approach 3, we joined at the beginning of each question the respective statement and all assumptions preceding the question in the output. This method allows us to avoid losing important information introduced in previous assumptions.

Table 10 illustrates an example of how we join a statement together with some assumptions. In this example, the dependence on the first assumption to understand the second and the third is clear. For instance, Mário Mendes and his role in the story are

Exam Questions	Output
What distinguishes a federation from a real union? A federation is a composite state, made up of several federated states, which, despite having their own Constitution, do not enjoy sovereignty, being subordinate to the federal Constitution. The real union, in turn, although also corresponding to a composite state, differs from the federation because the states that comprise it retain their autonomy and adopt a common Constitution, along with common bodies. Thus, while in the federation there is an overlap, in the real union there is a fusion. ===== How can it be explained that, in the United States of America, a candidate can be elected President even though they obtained fewer votes than their opponent? Topic to develop: this possibility – which, moreover, has occurred more than once – is explained by the peculiar presidential election system that, for historical reasons, was enshrined in the United States of America and which, for various reasons (including the difficulty of amending the Constitution), is still in force today: the indirect suffrage system. Since the President is elected by an electoral college, which, in turn, is composed of representatives elected by each federated state (the number of representatives per state being larger the greater its population), through a majority electoral system, this means that, more than the total number of votes obtained by each candidate nationwide, what matters most for the outcome of the presidential elections is the number of votes obtained by each candidate in the states that elect a larger number of representatives to the electoral college. ===== <i>Other questions omitted</i>	Type: Multiple Choice Question: What is the main difference between a federation and a real union? a) In the federation, the federated states maintain their sovereignty. b) In the real union, the states retain their autonomy and adopt a common Constitution. c) In the federation, there is a fusion of states. d) In the real union, the states do not have autonomy. Answer: b) In the real union, the states retain their autonomy and adopt a common Constitution. Type: Case Analysis Questions Question: A presidential candidate in the United States of America obtained fewer votes than their opponent but was elected President. What is the correct explanation for this situation? a) The U.S. electoral system is based on direct suffrage. b) The President is elected by an electoral college, composed of representatives elected by each state. c) A candidate with fewer votes cannot be elected President. d) The U.S. electoral system is based on the proportionality of votes. Answer: b) The President is elected by an electoral college, composed of representatives elected by each state. <i>Other questions omitted</i>

Table 7: Constitutional Law: English-translated example of exam question-answer pairs fed into GPT-4o and part of the respective output.

Exam Question	Output
On February 2, 2022, Artur met Benedita at the Vilamoura casino, and he fell in love with her at first sight. Benedita, who has a gambling addiction, ended up owing €80,000.00 to the Algarvian mafia, which now threatens to kill her if she doesn't pay her debt. Taking advantage of the situation, Artur promises to pay the €80,000.00, sparing Benedita from the mafia's retaliation, if she marries him by Valentine's Day. Desperate, Benedita agrees to the request and marries Artur without a preliminary process, due to the urgency. Eight months later, Benedita discovers that Artur is a criminal wanted in Italy for counterfeiting Gucci bags. Assess the validity of the marriage between Artur and Benedita, considering all the described facts. (5 points) The factual matter primarily relates to the causes of marriage annulment. First of all, since the marriage was entered into under the promise of paying Benedita's €80,000.00 debt to the mafia, in order to spare her from retaliation, this is a cause for annulment due to moral coercion (Article 1631(b), 1638(2), and Article 282(1)). However, for the marriage to be annulled on this basis of coercion, it will be necessary to file for annulment (Article 1632) by the persons indicated in Article 1641, within the time frame established by Article 1645. In this case, it seems that the right to annul the marriage on the grounds of moral coercion may have already expired. On the other hand, the fact that Artur is a criminal wanted in Italy may be relevant for invoking annulment based on an error of will (Articles 1631(b) and 1636). For the error of will to be valid under the mentioned terms, the following requirements must be met: i) an error concerning essential qualities of the person; ii) an excusable error; iii) both objective and subjective essentiality. In this case, the error is relevant, as it concerns essential qualities of the other spouse (being a criminal); it is excusable (not a gross error), since Artur is wanted in Italy and not in Portugal, nor is it necessarily a notorious fact, so it is excusable that Benedita did not know; [continues...]	Context: On February 2, 2022, Artur met Benedita at the Vilamoura casino, and he fell in love with her at first sight. Benedita, who has a gambling addiction, ended up owing €80,000.00 to the Algarvian mafia, which now threatens to kill her if she doesn't pay her debt. Taking advantage of the situation, Artur promises to pay the €80,000.00, sparing Benedita from the mafia's retaliation, if she marries him by Valentine's Day. Desperate, Benedita agrees to the request and marries Artur without a preliminary process, due to the urgency. Eight months later, Benedita discovers that Artur is a criminal wanted in Italy for counterfeiting Gucci bags. Type: Multiple Choice Question: Which of the following options is a cause for the annulment of the marriage between Artur and Benedita? a) Lack of true love b) Moral coercion c) Age difference d) Lack of witnesses Answer: b Type: Multiple Selection Questions Question: Which of the following factors can be considered moral coercion in the context of marriage? a) Death threats b) Financial pressure c) Emotional blackmail d) Age difference Answer: a) b) c) <i>Other questions omitted</i>

Table 8: Family Law: English-translated example of exam question-answer pair fed into GPT-4o and part of the respective output.

presented in the first assumption, a fact that is important to understand the following ones. Because this dependence between assumptions might not

always apply, and, in some cases, assumptions may even be contradictory, we follow each assumption, except the last one, by the Portuguese version of

Exam Questions	Output
Adalberto has been running a vegetarian restaurant in Chiado for several years. Tired of the daily grind, he decides to contact Berta to find someone interested in taking over the business. Berta quickly starts working, speaking wonders about the restaurant and convincing Carloto to take over the business. On the agreed date, Adalberto and Carloto sign a contract that states: “the restaurant is sold with all its elements, including the exclusive supply contract with Dioclécio.” Two months later, during a friendly conversation, Carloto convinces Ermelinda to invest in his business in exchange for a share of the profits he may obtain. Answer the following questions, justifying your response: 1. Berta demands compensation from Adalberto for her services. Adalberto claims that, aside from no agreement being made regarding this, he also owes her nothing because Carloto hasn’t paid for the business. Quid iuris? (4 marks) Classification of the contract between Adalberto and Berta as a mediation contract. Explanation of the applicable legal regime, considering the absence of a general regime for this type of contract. Distinction between the mediation contract and other contractual types, particularly agency contracts. Analysis of the right to remuneration: in mediation, unlike in agency contracts, the commission is due regardless of payment by the buyer, as long as the contract is signed. The lack of an agreement between the parties on the amount of remuneration does not prevent payment. 2. Dioclécio is contacted by Carloto, placing several product orders. Dioclécio is surprised because, aside from not knowing who Carloto is, he never agreed to anything with him. Moreover, the commercial conditions agreed with Adalberto stemmed from their long-standing friendship. Quid iuris? (4 marks) Analysis of the situation at hand in light of the effects of business transfer on contracts related to the establishment. Explanation of the key issues and various doctrinal positions, particularly: (i) application of the general regime for assignment of contractual position (Articles 424 and following), (ii) operational legal situations; (iii) existence of a general principle (derived from specific normative examples in Commercial Law) of automatic transfer of contractual positions related to the business establishment; [continued...]	Context: Adalberto has been running a vegetarian restaurant in Chiado for several years. Tired of the daily grind, he decides to contact Berta to find someone interested in taking over the business. Berta quickly starts working, speaking wonders about the restaurant and convincing Carloto to take over the business. On the agreed date, Adalberto and Carloto sign a contract that states: “the restaurant is sold with all its elements, including the exclusive supply contract with Dioclécio.” Two months later, during a friendly conversation, Carloto convinces Ermelinda to invest in his business in exchange for a share of the profits he may obtain. Assumption: Berta demands compensation from Adalberto for her services. Adalberto claims that, aside from no agreement being made regarding this, he also owes her nothing because Carloto hasn’t paid for the business. Type: True/False Question: The lack of an agreement between the parties on the amount of remuneration prevents the payment of the commission in a mediation contract. Answer: False <i>Other questions omitted</i> Assumption: Dioclécio is contacted by Carloto, placing several product orders. Dioclécio is surprised because, aside from not knowing who Carloto is, he never agreed to anything with him. Moreover, the commercial conditions agreed with Adalberto stemmed from their long-standing friendship. Type: Cloze Tasks Question: The automatic transfer of contractual positions related to the business establishment is limited when: a) The assignee offers greater guarantees than the assignor b) The contract is <i>intuitu personae</i> c) The contract is for the supply of products d) The contract is for the provision of services Answer: b) The contract is <i>intuitu personae</i> <i>Other questions omitted</i>

Table 9: Commercial Law: English-translated example of exam question-answer pairs fed into GPT-4o and part of the respective output.

Statement: GELL&CO and SOLIMPA, two companies in the cleaning products manufacturing sector, decided in May 2020 to adapt their production to manufacturing hand sanitizer. On 02/08/2020, both companies were contacted by a representative of a well-known international distributor, JACQUES SILVA, from a Spanish branch, requesting the urgent delivery of 5,000 units of hand sanitizer. GELL&CO and SOLIMPA accepted (each thinking it had been the only one contacted and contracted). GELL&CO shipped merchandise valued at 15,000 euros to the address provided by JACQUES (located in Spain) on 05/09/2020. SOLIMPA shipped its merchandise, also valued at 15,000 euros, on 10/09/2020. At the end of September 2020, having received no payment and unable to contact JACQUES, the companies individually contacted the international distributor, confirming that there was no representative named JACQUES, nor any contract. On 06/10/2020, the Portuguese companies filed a complaint with the Public Prosecutor’s Office. Assumption: On 15/03/2021, following an undercover operation, the Judiciary Police (PJ) arrested MÁRIO MENDES, suspected of impersonating JACQUES SILVA, and the main suspect in several fraud crimes in both Portugal and Spain, in Elvas. On 16/03/2021, MÁRIO was brought before the Criminal Investigation Judge (JIC) for his first interrogation as a detained defendant. [Some questions related to the assumption above, to which you have already answered, were omitted.] Assumption: On 20/09/2021, the Public Prosecutor’s Office charged MÁRIO with two counts of aggravated fraud, in actual competition (Article 218/1 of the Criminal Code), one against GELL&CO and another against SOLIMPA, as well as one count of identity document forgery (Article 256/1 of the Criminal Code), also in actual competition. [Some questions related to the assumption above, to which you have already answered, were omitted.] Now consider the following assumption: Before the trial began, MÁRIO reached an agreement with GELL&CO and SOLIMPA, voluntarily compensating for the damages caused, in accordance with Articles 218/4 and 206/1 of the Criminal Code. Consequently, the Criminal Investigation Judge declared MÁRIO’s criminal liability for the two counts of aggravated fraud extinguished, leaving the case to proceed to trial solely for the crime of document forgery. In his defense, MÁRIO provided evidence that he had used a false name in emails but had never fabricated or used a false document, as he had not used any symbols or identifying words of any registered trademarks. However, the trial court convicted him, invoking the agreement under Article 206 as an implicit confession of the commission of all crimes listed in the indictment.

Table 10: Criminal Procedure Law example: illustration of how we combine a problem statement with three assumptions, all extracted from a GPT-4’s output. The text presented is placed at the beginning of all questions generated by GPT-4o that follow the third assumption.

“Some questions related to the assumption above, to which you have already answered, were omitted.” We did this to mimic a human taking the original exam, who would see all the questions together and have access to all the available information. However, we wanted to avoid the model conditioning its answers on previous questions and responses.

Table 11 shows the English translation of a prompt template used to generate new questions from groups of short and independent exam questions (approach 1, first iteration).

Below, you are presented with a set of questions taken from exams and their respective grading criteria. Your task is to create a new exam to evaluate a group of students, based on the questions and grading criteria presented in 'Questions taken from exams'. You should create various types of questions: Multiple Choice questions, Cloze Tasks, True/False questions, Multiple Selection Questions, Case Analysis Questions where a brief scenario/case is presented and students are asked to choose the best answer from several options. Be creative. The purpose of the questions you will create is to assess how well a student masters the topic. The questions you create should be challenging. The answers to the questions you create should be contained in 'Questions taken from exams'. Avoid creating questions about specific articles or laws. Students are not required to memorize articles and laws and will not have access to them during the exam. The questions you create should focus on assessing legal reasoning. Create at least three to four questions of each of the types described above. The questions you create should all be independent of each other. Your output should follow the following format: <pre> Type: {{insert question type here}} Question: {{insert question here}} Answer: {{insert answer here}} Type: {{insert question type here}} Question: {{insert question here}} Answer: {{insert answer here}} ... Attention: i) For Cloze Tasks, present answer options, multiple-choice style - the answer should be only ONE letter, corresponding to the correct option; ii) In True/False questions, the answer should be only 'True' or 'False'; iii) In Multiple Selection Questions, present the answers in the format 'letter) ...'; iv) In Multiple Choice and Case Analysis Questions, the answer should be only ONE letter, corresponding to the correct option. Questions taken from exams: '{' </pre>

Table 11: English translation of a prompt template used to generate new question-answer pairs.

B Filtering Repeated Questions

Table 12 illustrates the processing applied to matching questions for some comparisons: we removed the first line (usually “Match the items...”) and the letters/numbers identifying options.

We recall that multiple-choice, cloze tasks, case analysis, and multiple selection questions are considered multiple-choice variants. For lexical comparisons, Tables 13 and 14 summarize the ROUGE-L variants, question processing, and thresholds used for filtering within and between the differ-

Original matching question
Match each item on the agenda with the type of competence involved. a) Approval of the Code of Ethics and Conduct of the Municipality of Amadora b) Acceptance of donations of works of art c) Approval of the amendment to the inter-administrative contract for delegation of powers with the Parish of Alfragide 1) Competence of the City Council 2) Competence of the Municipal Assembly 3) Competence of the Parish Council
Processed matching question
Approval of the Code of Ethics and Conduct of the Municipality of Amadora Acceptance of donations of works of art Approval of the amendment to the inter-administrative contract for delegation of powers with the Parish of Alfragide Competence of the City Council Competence of the Municipal Assembly Competence of the Parish Council

Table 12: Illustration of the processing applied to matching questions for lexical comparisons using ROUGE-L at summary level.

ent question types. For semantic comparisons, we considered two questions repeated or similar if the cosine similarity of their sentence embeddings was 0.80 or higher. Table 15 summarizes how we processed questions based on their types. We manually set the best threshold for each scenario.

C More Statistics

Table 16 shows the number of questions in the benchmark by type, and Table 17 presents the distribution of questions within each area of law across the different types, in percentage. The variation in the number of questions across the different types, as shown in Table 16, as well as the differences in distribution across areas, are partly due to the number of exam exercises run with each generation approach. In all approaches, we instructed GPT-4o to generate multiple-choice, cloze tasks, true/false, and multiple selection questions. In the approach using short, independent exam questions, we also instructed the model to generate case analysis questions, as the exam questions were often very theoretical. Conversely, matching questions were generated only in the other two approaches. Some variations in distribution across areas may stem from smaller subsets of questions, as in the case of Labor Procedure Law. Others differences are likely attributed to the filtering phase.

D Evaluation Prompts

Table 18 illustrates the English translation of an evaluation prompt. The “Instruction” field and the

Type	Processing	ROUGE-L Variant	Threshold
Matching questions	Illustrated in Table 12	Summary	0.70
Multiple-choice variants	None	Summary	0.80
Multiple-choice variants	Remove options	Sentence	0.75
True/False	None		

Table 13: Summary of ROUGE-L comparisons within each question type (multiple-choice, cloze tasks, case analysis, multiple selection, true/false, matching questions — the first four referred to as multiple-choice variants): question processing, ROUGE-L variant used, and threshold applied.

Type I	Type II	ROUGE-L Variant	Threshold
Multiple-choice variant (remove options)	Multiple-choice variant (remove options)	Sentence	0.85
	True/False (no processing)		
Multiple-choice variant (no processing)	Multiple-choice variant (no processing)	Summary	0.80
	True/False (no processing)		

Table 14: Summary of ROUGE-L comparisons between questions of different types (multiple-choice, cloze tasks, case analysis, multiple selection, true/false — the first four referred to as multiple-choice variants): question processing, ROUGE-L variant used, and threshold applied. Type I and Type II represent distinct types.

Type I	Type II
Matching question (no processing)	Any type (no processing)
Multiple-choice variant (remove options)	Multiple-choice variant (remove options)
	True/False (no processing)

Table 15: Summary of semantic comparisons: compared types and processing. Type I and Type II may be the same.

Type	Questions
Multiple-choice	1099
Cloze tasks	983
Case analysis	89
True/False	1309
Multiple selection	695
Matching questions	548

Table 16: LegalBench.PT: number of questions by type.

final line specifying the output format vary depending on the question type. The “Statement” and “Assumption” fields are omitted if a question does not have any statement or assumption associated.

E Model Performance Across Types

Table 19 displays the performance of all LLMs evaluated on the benchmark across the different question types. The models do not exhibit striking differences between types, though there is a tendency toward higher results in multiple selection and matching questions. We expected these ques-

tion types to be more challenging than multiple-choice variants, as they require more than simply identifying the most likely correct option. Additionally, they are evaluated using the F_1 score, which equally rewards correct answers and penalizes incorrect ones. Conversely, true/false questions tend to be associated with the lowest scores, which is also surprising, as we expected this type to be the easiest. Differences in the distribution of question types within each field of law may have contributed to the observed variations in performance. A deeper analysis is left for future work.

Fields of Law	Multiple-choice	Cloze tasks	Case analysis	True/False	Multiple selection	Matching questions
Public	22.5	21.5	1.6	28.7	14.3	11.5
Environmental Law	23.1	26.9	5.8	25.0	11.5	7.7
Administrative Law	22.5	22.9	2.2	27.8	12.8	11.9
Constitutional Law	22.8	22.8	3.9	26.2	15.5	8.7
Energy Law	9.5	14.3	19.0	33.3	23.8	0
Public Finance Law	20.2	21.6	2.4	29.3	13.5	13.0
Financial Law	25.0	31.2	0	25.0	0	18.8
Tax Law	21.1	22.4	0	28.9	15.9	11.6
Criminal Law	21.5	22.8	0.3	27.7	14.4	13.3
Administrative Procedure Law	21.6	15.5	1.7	31.9	19.0	10.3
Civil Procedure Law	22.4	21.7	0.7	29.0	13.6	12.5
Criminal Procedure Law	27.0	17.9	0.7	29.6	12.4	12.4
Labor Procedure Law	0	33.3	33.3	33.3	0	0
Urban Planning Law	24.5	21.4	0	31.6	16.3	6.1
Private	23.6	20.5	1.8	26.6	14.9	12.5
Contract Law	24.1	19.8	1.4	26.9	15.1	12.7
Family Law	23.6	17.3	4.1	29.1	14.5	11.4
Law of Obligations	28.4	19.7	0.5	26.1	12.4	12.8
Property Law	23.5	22.1	0.4	25.1	15.7	13.1
Succession Law	29.4	11.8	0	29.4	5.9	23.5
Commercial Law	22.9	19.6	1.7	26.7	15.0	14.2
Banking Law	16.7	30.0	0	26.7	16.7	10.0
Maritime Law	25.0	12.5	25.0	25.0	12.5	0
Corporate Law	24.7	18.2	0.6	29.2	15.6	11.7
Securities Law	20.0	22.2	4.4	22.2	20.0	11.1
Transportation Law	26.9	23.1	11.5	19.2	19.2	0
Aviation Law	22.6	25.8	12.9	22.6	16.1	0
Insolvency Law	27.0	16.2	0	29.7	16.2	10.8
Private International Law	23.6	22.7	1.8	28.2	11.8	11.8
Labor Law	18.4	22.2	2.4	27.1	15.0	15.0
Public-Private (Competition Law)	24.6	21.1	0	25.1	18.3	10.9
Public International	26.9	16.0	6.1	28.8	16.0	6.1
EU and Community	23.8	21.3	3.0	30.5	12.8	8.5
Overall	23.3	20.8	1.9	27.7	14.7	11.6

Table 17: Distribution (%) of questions within each area of law across the different types. Each row sums to 100%.

You are solving an exam on Family Law.
Statement: Nuno and Sofia, both single and without children, got married on October 15, 2023, having previously celebrated, on January 15, 2023, a prenuptial agreement, in which they stipulated the following: a) they adopt the regime of acquired community property, but all assets acquired with the couple's money, even in part (even if minimal), are common assets; b) each spouse may alienate their own property without the need for the other spouse's consent; c) the couple's assets are liable only for debts incurred by both spouses.
Assumption: Nuno and Sofia became parents to João on November 15, 2023.
Instruction: Answer the following multiple-choice question. Only one option is correct.
Question: The paternity of João is, in principle, established by _____. a) declaration of paternity b) presumption of paternity c) judicial recognition d) acknowledgment
The output should only be: "The correct answer is: {letter}"

Table 18: English-translated example of an evaluation prompt.

Models	Multiple-choice	Cloze tasks	Case analysis	True/False	Multiple selection	Matching questions
GPT-4o	84.6	83.6	84.6	83.3	88.6	91.6
GPT-4o-mini	78.0	78.7	81.6	76.6	85.7	88.3
Claude-3-Opus	83.9	86.0	84.6	79.2	87.0	91.3
Claude-3.5-Sonnet	84.7	84.9	83.7	81.8	87.9	91.0
Llama-3.1-8B	69.2	68.3	72.7	64.4	78.4	64.8
Llama-3.1-70B	80.5	80.2	85.5	76.9	84.4	85.6
Llama-3.1-405B	85.2	82.4	81.9	80.6	85.6	89.4
Mixtral-8x7B	69.5	67.5	73.7	69.9	71.6	70.1

Table 19: Model performance (%) across the different types of questions.