

BLT: Can Large Language Models Handle Basic Legal Text?

Anonymous ACL submission

Abstract

We find that the best publicly available LLMs like GPT-4, Claude, and PaLM 2 currently perform poorly at basic legal text handling. We introduce a benchmark consisting of tasks that lawyers and paralegals would expect LLMs to handle zero-shot, such as looking up the text at a line of a witness deposition or at a subsection of a contract. LLMs' poor performance on this benchmark casts into doubt their reliability as-is for legal practice. However, fine-tuning for these tasks brings even a smaller model to near-perfect performance on our test set and also raises performance on a related legal task. These results suggest that many simple behaviors needed for a domain may not be present in foundational LLMs, without additional engagement from subject matter experts.

1 Introduction

There has been much excitement about the potential of large language models (LLMs) to handle complex reasoning tasks out of the box. In the legal domain, recent work has focused on LLMs' ability to perform self-contained, high-level tasks such as passing the bar exam, identifying hearsay evidence, or navigating the tax code. These results have led to predictions that LLMs will imminently revolutionize law practice.

However, no prior work has investigated LLMs' ability to handle the most basic text tasks done in daily legal practice. To fill this gap we introduce a benchmark for **Basic Legal Text** (BLT) handling. BLT mimics simple tasks done by lawyers and paralegals, including looking up citations to the record, finding statutory citations containing particular text or definitions, or finding the text in a contract that another lawyer has explicitly cited.

Our code¹ can generate nearly unlimited new examples. We release static train and test sets for a

1: straight that I couldn't do nothing.
2: Q. How did you get to your
3 : girlfriend's house today?
4: A. She picks me up.
5: Q. How long have you guys been
6: together?
7: A. We've been together for four or
8: five years.
9: Q. Okay. Do you currently have
10: any doctors' appointments scheduled?
11: A. I did yesterday, but not today.
12: Q. Like do you have another one
13: this week?
14: A. I have therapy tomorrow.
15: Q. Tomorrow?

What is the exact text of just line 10 above?

A. I did yesterday, but not today.

Figure 1: GPT-4 incorrectly answering a simple question about a page from a line-numbered witness deposition transcript. Line numbers, exactly as shown above, are passed to GPT-4, which incorrectly answers 23% of such one-page, 15-line deposition retrieval prompts.

progression of window sizes (4k, 8k, 16k, 32k, 64k and 128k tokens), as well as code that calls LLM APIs, measures accuracy, and classifies errors.

By design, a paralegal, newly-minted lawyer, or even non-subject matter expert could perform the BLT tasks zero-shot at or near 100%. In contrast, we find that even the most advanced publicly available LLMs show poor zero-shot performance out of the box on BLT. But fine-tuning a less-advanced model, GPT-3.5-turbo, on BLT's training set leads to near-human (i.e., 100%) performance. This indicates that (1) LLM training sets have not been designed with essential text manipulation tasks in mind, but (2) one might enumerate such tasks alongside a domain expert to improve LLMs.

¹ <https://anonymous.4open.science/r/BLT-6C44>

2 Background

2.1 Why Care About Law?

Law is a largely text-based profession, so it is often used to demonstrate advances in natural language processing in general and LLMs specifically. For example, OpenAI’s GPT-4 technical report (OpenAI, 2023b) mentioned only a single benchmark in the abstract itself: GPT-4 had passed the bar exam, with the score in the 90th percentile.

OpenAI’s developer livestream by co-founder Greg Brockman introducing GPT-4 (OpenAI, 2023a) used four examples to show GPT-4’s capabilities. In one, Brockman prompted GPT-4 with several sections of U.S. tax law and had it calculate the taxes of hypothetical taxpayers Alice and Bob. Brockman proclaimed that GPT-4 can “do taxes.”

Subsequent authors have raised doubts about these achievements. Martínez (2023) questions the statistics behind GPT-4 getting in the 90th percentile on the bar exam. Tax lawyers noted that GPT-4 in the livestream actually got the wrong answer for Alice and Bob’s taxes (Zhang, 2023).

2.2 Legal Use of LLMs

Legal NLP is concerned with a diverse range of tasks, reflecting the diversity of tasks lawyers perform. Examples include legal judgment prediction (Chalkidis et al., 2019; Xiao et al., 2018; Medvedeva and McBride, 2023), contract review (Hendrycks et al., 2021), document review (Lewis et al., 2023), and retrieving relevant case law (Kim et al., 2022). There has been extensive discussion of how NLP can benefit the legal system (Zhong et al., 2020; Aidid and Alarie, 2023). LLMs have been deployed for a wide range of legal tasks, including case analysis (Savelka et al., 2023), discovery (Pai et al., 2023), and analyzing contracts (Roegiest et al., 2023).

Many legal benchmarks for LLMs have been created (Chalkidis et al., 2022; Fei et al., 2023), with many incorporated into the broad LegalBench project (Guha et al., 2022). These are all much higher-level tasks than BLT, including identifying testimony to which the hearsay doctrine applies or whether contractual terms impose particular restrictions (Hendrycks et al., 2021).

The SARA (StAtutory Reasoning Assessment) dataset is one of the higher-level tasks in LegalBench (Holzenberger et al., 2020). It consists of nine tax-related sections of the U.S. Code, plus 376 hand-crafted “cases” consisting of facts and

a question that can be unambiguously answered applying the nine sections to the facts. Because SARA is a higher-level task clearly predicated on lower-level text handling (specifically, finding text at a citation), we use it to measure the effectiveness of fine tuning with BLT, discussed in Section 5.

LLMs have seen much of the internet during their training. To evaluate LLMs’ ability to handle novel legal questions, Nay et al. (2023) generate synthetic multiple-choice legal questions. Similarly, to test whether LLMs can handle truly novel legal texts, Blair-Stanek et al. (2023) generate synthetic sections constructed with nonces (phonetically plausible nonsense words) and probed GPT-3’s ability to reason over these synthetic sections. BLT incorporates such synthetic sections, albeit for simpler tasks than statutory reasoning.

2.3 LLMs in Practice

Several LLM-testing benchmarks like BigBench (Srivastava et al., 2023) have included basic word handling tasks like word sorting and returning specified sentences from passages. LLMs can solve quite complicated tasks by being prompted to provide a chain of thought (Kojima et al., 2022), including in the legal domain (Yu et al., 2022). More generally, choosing the appropriate way to prompt LLMs, called prompt engineering, often has a substantial impact on LLM performance (White et al., 2023; Liu et al., 2023b). A complementary approach has been to decompose the task at hand into tasks the LLM can handle (Dua et al., 2022; Khot et al., 2023). This involves appropriate prompting.

LLMs have been trained or otherwise induced to use “Tools” (Schick et al., 2023; Paranjape et al., 2023). For example, an LLM might detect that it needs to call a calculator tool to handle a math problem posed to it in text form. In theory, tools could be written to handle the BLT tasks and then be integrated into LLMs. But the BLT prompts are oversimplified versions of tasks lawyers need LLMs to do seamlessly. A lawyer will not ask an LLM for the citation to the record where the plaintiff says “I have therapy tomorrow.” But a lawyer might expect an LLM to insert a citation to the record that proves a plaintiff receives care for psychological problems, and a basic text-matching tool would not handle that.

Some of BLT’s prompts are quite long, which is realistic because lawyers often handle long texts. Liu et al. (2023a) investigated how LLMs handle long prompts. They found that LLMs’ accuracy

155 followed a U-curve with respect to the informa- 200
156 tion’s position, with information in the middle of 201
157 the prompt used much less than if it were at the start 202
158 or end. They connected this to the “serial-position” 203
159 effect exhibited by humans, who best remember 204
160 material presented near the beginning or end. 205

161 3 The BLT Benchmark 206

162 The BLT benchmark involves three different types 207
163 of legal text, each of which has between two and 208
164 five different tasks run on it. 209

165 3.1 Deposition Transcripts 210

166 In litigation in the U.S., depositions of witnesses 211
167 under oath are a key factfinding tool.² The deposi- 212
168 tions typically occur in lawyers’ offices and allow 213
169 lawyers to ask witnesses questions on virtually any 214
170 topic. Professional court reporters transcribe the 215
171 depositions into transcripts, typically with 25 num- 216
172 bered lines per page,³ often running over 100 pages 217
173 for a single witness deposition. Attorneys must cite 218
174 relevant portions of the resulting transcripts in sub- 219
175 sequent motions, such as those asking the court to 220
176 grant their side summary judgment.⁴ Portions of 221
177 transcripts are cited by page and line number.⁵ 222

178 One basic legal text-handling task a lawyer must 223
179 do, in constructing a motion, is finding the page 224
180 and line of a transcript where particular text ap- 225
181 pears. This motivates the **text**→**cite** task, where 226
182 the prompt consists of one or more pages of ac- 227
183 tual deposition transcript followed by the question, 228
184 “What are the page number and line number of the 229
185 line above with the text “__”?” To ensure there is 230
186 only one clearly correct answer, prompts are never 231
187 constructed asking about lines with less than four 232
188 words, that are subsets of another line, or that are 233
189 too similar to other lines (defined as Levenshtein 234
190 distance under four (Levenshtein et al., 1966)). 235

191 The converse is another basic text-handling task: 236
192 given a citation to a transcript, find the text at the 237
193 cited location. Lawyers must do this basic task 238
194 in order to evaluate the opposing side’s motions. 239
195 Paralegals do it on their side’s own motions before 240
196 submitting them (ProParalegal, 2017). This moti- 241
197 vates the **cite**→**text** task, where the prompt consists 242
198 of one or more pages of a deposition transcript fol- 243
199 lowed by the question “What is the exact text of

just line __ of page __ above?” An example of 200
GPT-4 failing this task appears in Figure 1. 201

202 BLT’s deposition transcript tests are built from 203
204 a novel corpus we constructed of 33,176 lines of 205
206 actual deposition transcripts filed with federal dis- 207
208 trict courts and downloaded via Bloomberg Law 208
209 from the Public Access to Court Electronic Records 209
210 (PACER) system. They are from a variety of 210
211 cases; our main criterion was that the transcripts 211
212 be cleanly OCR’ed and could be fully cleaned with 212
213 regular expressions. This corpus can be extended 213
214 *ad infinitum* by others by downloading further tran- 214
215 scripts and cleaning them, as deposition transcripts 215
216 are likely not copyrightable.⁶ The existing page 216
217 and line numbers are stripped out, and random 217
218 spans of appropriate length are selected, with new 218
219 page numbers and line numbers, in case the LLM 219
220 saw the original numbering during training. 220

221 The size of the prompt is scalable, which was a 221
222 key goal of BLT. For BLT-4k (the version fitting 222
223 within a 4,000 token window), we use a mix of 223
224 1 and 2 page subsets of depositions.⁷ Larger ver- 224
225 sions have larger subsets. For BLT-128k, we use 225
226 a mix of 120 and 140-page subsets of depositions. 226
227 Appendix B details the sizes used. 227

228 3.2 Synthetic Sections 228

229 Lawyers regularly work with hierarchical text, in- 229
230 cluding statutes, contracts, regulations, treaties, 230
231 court rules, and corporate charters. Hierarchical 231
232 text is often organized into sections, subsections, 232
233 paragraphs, subparagraphs, etc. Being able to nav- 233
234 igate such hierarchical text is a basic legal text- 234
235 handling task required of all lawyers, whether they 235
236 are litigators arguing that a statute applies to their 236
237 case or are transactional lawyers negotiating the 237
238 terms of a contract. 238

239 We generate synthetic hierarchical sections, fol- 239
240 lowing the approach of Blair-Stanek et al. (2023). 240
241 They use repeated application of the logical form 241
242 $A \Rightarrow B$, with two parameters: depth, the max- 242
243 imum number of times it is applied; and width, 243
244 the number of times it is applied to each B . The 244
245 terms defined are nonces that are not real words 245
246 but are pronounceable. These synthetic sections 246

²Federal Rule of Civil Procedure 30.
³U.S. Court Reporters Association Manual, Section 18.8
⁴Federal Rule of Civil Procedure 56(c)(1)(A).
⁵See Bluebook Rule B17.1.2 (Bluebook, 2020).

⁶Lipman v. Massachusetts, 311 F. Supp. 593 (D. Mass. 1970); 1 Nimmer on Copyright §5.12[C] (2023 ed.). Even in the unlikely event transcripts are copyrightable, this use of them is likely fair use. 17 U.S.C. §107.
⁷For the 1-page subsets, only the line number is involved (i.e., not any page number). The 1-page subsets are followed by “What is the line number of the line above with the text “__”?” and “What is the exact text of just line __ above?”

Section 5217. Definition of cleight.
 (a) General rule. The term "cleight" means any baitisolist or any roussiont.
 (b) The term "baitisolist" means any ballinated or any caset.
 (c) The term "roussiont" means any dicemercu or any accodpoileare.

What is the exact citation above where the term "roussiont" is defined? (Use standard legal formatting like section 1001(b)(2)).

Section 5217(b)

Figure 2: Example of GPT-4 incorrectly answering **defined**→**cite** question with a 2-deep, 2-wide synthetic section. The correct answer is "section 5217(c)".

can be arbitrarily large, by adjusting the width and depth, thus making tasks based on them scalable to different-sized token windows. For example, BLT-4k has synthetic sections ranging from 2-wide, 2-deep, as in Figure 2, which are very short, up to 3-wide, 4-deep, which takes up much of the 4k token window. At the highest end, BLT-128k has a variety ranging from 60-wide, 2-deep to 5-wide, 5-deep, which (because size is exponential with respect to depth) takes up much of the 128,000-token window. For the full list of sizes in each BLT-*, see Appendix A. Being synthetic ensures they are novel and not seen by LLMs during training. This simulates the challenges faced by lawyers in handling newly drafted contracts, legislation, or other hierarchical text. Nearly unlimited quantities of synthetic sections of any size can be generated.⁸

A basic legal text-processing skill is finding the citation, in a hierarchical text, of the text to which you are pointing a court or another lawyer. This motivates applying the **text**→**cite** task on synthetic sections, where the prompt consists of one synthetic section followed by the question "What is the exact citation above of the text "___"? (Use standard legal formatting like section 1001(b)(2))." The code to generate synthetic sections guarantees there is only a single correct answer.

The converse legal skill is, given a hierarchical citation, finding the text at it. Hence we apply the **cite**→**text** task to synthetic sections, with the prompt consisting of one synthetic section followed by the question "What is the exact text of just sec-

tion ___ above?"

We ask this question only of "leaves" in the statute, meaning they have no subsections underneath them. This ensures there is only a single correct answer. For example, suppose that section 573(a) was not a leaf, perhaps with paragraphs 573(a)(1) and 573(a)(2) underneath it. If you asked for the text of section 573(a), it is ambiguous whether you should also return the text of 573(a)(1) and 573(a)(2) as well. Such ambiguity is avoided by considering only leaves.

We also include two other basic legal text-handling tasks on the synthetic sections. Terms are defined in hierarchical texts and often referenced elsewhere in the same hierarchical text. Lawyers must be able to cite a term's precise definition. With **defined**→**cite**, the prompt is one synthetic section followed by the question "What is the exact citation above where the term "___" is defined? (Use standard legal formatting like section 1001(b)(2))." Conversely, when given such a citation by another lawyer, a lawyer must be able to find the term, which motivates **cite**→**define**. There, the prompt is one synthetic section followed by the question "What is the term defined at section ___?" An example of GPT-4 incorrectly answering a **defined**→**cite** problem appears in Figure 2.

3.3 U.S. Code

The U.S. Code is the official compilation of general and permanent U.S. federal statutes.⁹ The U.S. Code is a large corpus of hierarchical text. We apply to the U.S. Code all four tasks that we applied to synthetic sections: **text**→**cite**, **cite**→**text**, **defined**→**cite**, and **cite**→**defined**. For these four tasks on the U.S. Code, the prompt is the same as for synthetic sections.

During training, LLMs have doubtless seen all of the U.S. Code, which is not copyrighted¹⁰ and is publicly available on multiple websites. To test whether LLMs' familiarity with U.S. Code sections causes errors, we add a fifth test for U.S. Code sections: **cite**→**amended**. In all but one respect, this test is identical to **cite**→**text**, in that it has the text of one or more sections and asks "What is the exact text of just section ___ above?" about a leaf.

The sole difference is that we make a small but semantically-important change to the text in that leaf to see if the LLM returns the original text or the

⁸There are 9000 nonces available, so, for example, there are approximately 10^{197} possible unique synthetic statutes with 50 nonces.

⁹Congress makes it available in XML form at <https://uscode.house.gov/download/download.shtml>.

¹⁰17 U.S.C. §105

changed text (which is the correct answer). This tests a basic legal skill: applying a given newly-amended statute, rather than its old version. If the leaf contains any numbers, we add or subtract one from the last appearing number.¹¹ Otherwise, we tweak the last appearing citation from, say, “(D)” to “(A)”. Otherwise, we toggle the last “and” to “or” or vice versa. Otherwise, we toggle the last “shall” to “may” or vice versa.¹² If none of these changes are available, we insert “unless otherwise provided by section 101,” at the start of the leaf.

For all tasks on the U.S. Code, we do not use sections containing tables, which are not purely text.¹³ We do not use sections with quoted hierarchical text such as model contracts, which are hard for even a human lawyer to read.¹⁴

For text→cite, we do the same test as with transcripts, not using lines that are under four words long, are subsets of any line appearing elsewhere in the prompt, or that have a Levenshtein distance under four from another line in the prompt. For defined→cite, we do not use terms defined in more than one place in the prompt. We never use any of the cites that Congress has sloppily added twice.¹⁵

Unlike synthetic sections, which can be generated in unlimited quantities in arbitrarily large sizes, there are a limited number of U.S. Code sections. But it is a huge corpus, with 43,916 sections that meet the criteria discussed above, 447,037 leaves, and 23,562 unique definitions. Although 94% of sections are under 2,000 GPT-4 tokens, that still leaves 2,602 sections over 2,000 tokens, including 813 sections over 4,000 tokens and 196 sections over 8,000 tokens. When there are insufficient numbers of large enough sections, we can generate prompts of any desired size *ad infinitum* by adding randomly selected other sections of approximately the same size. We randomly shuffle the order of the sections in the prompt so that the target section’s position is not a cue to the model.

Having multiple sections in a prompt resembles how Brockman pasted nine tax-related sections into GPT-4 during his livestream. This is realistic: lawyers handling real-world issues often must

apply several statutes in conjunction, not just one.¹⁶

3.4 General Considerations

For each of the 11 tests, and for each possible size (ranging from BLT-4k to BLT-128k), we generate a training/test split of 1000/100 prompts.¹⁷

Why only 100 test prompts for each test split? Three reasons. First, there are 11 tests, thus 1,100 test prompts for each size of BLT-*. Second, the cost of calling GPT-4 with just 1,100 BLT-8k prompts with around 5,000 tokens per prompt is already nontrivial. Third, any LLM deployed for real-world legal practice really should be at or near 100%, and as accuracy approaches 100% the t-statistic goes to zero.

We have publicly released all test and train files.¹⁸ We retain no “held out” additional test splits, because our code can generate nearly unlimited additional test data simply by changing the random seed and by adding further OCR’ed deposition transcripts.

4 Results and Discussion

We tested three models from OpenAI: GPT-3.5-turbo, GPT-4, and GPT-4-turbo.¹⁹ From Google, we tested two variants of PaLM 2: chat-bison-001 and chat-bison-32k. From Anthropic, we tested Claude-2.1.

Table 1 contains the results on all sizes of BLT-* that each model can handle. All tests were by API call, with temperature set to 0. Each number in Table 1 (other than the mean and numbers in parentheses) corresponds to 100 calls to the relevant LLM’s API. Our code measures accuracy with forgiving rules, ignoring case and whitespace. Our code uses handwritten rules to classify errors, a feature we draw on in the discussion below.

Every model achieves lower accuracy on larger sizes of BLT-*. GPT-4 and GPT-4-turbo generally outperformed others, although Claude-2.1 slightly outperforms GPT-4-turbo on BLT-64k and -128k.

¹⁶E.g., *Home Depot U.S.A., Inc. v. Jackson*, 139 S.Ct. 1743 (2019), involved the interplay of 28 U.S.C. §§1441, 1446, and 1453.

¹⁷Thus on BLT-4k there are $11 \times 100 = 1,100$ test prompts, spread across 11 JSONL files, as well as $11 \times 1,000 = 11,000$ training prompts across 11 JSONL files.

¹⁸BLT-4k fits on Anonymous Github: <https://anonymous.4open.science/r/BLT-6C44/Data/BLT4k/>. Links to all versions will be added upon acceptance.

¹⁹Specifically, GPT-3.5-turbo-1106, GPT-4-0613, and GPT-4-1106-preview respectively.

¹¹We never move from 1 to 2 or from 2 to 1 since that would also require changing singular nouns to plural or vice versa.

¹²In legal phrasing, the distinction between “may” and “shall” is extremely important (*Garner, 2019*).

¹³For example, 26 U.S.C. §1 containing income tax tables is not used. 5,946 sections are excluded for this reason.

¹⁴Examples include 5 U.S.C. §9507 and 25 U.S.C. §5329.

¹⁵E.g., there are two subsection (e)’s in 42 U.S.C. §1397hh.

	model	transc.		synthetic section				U.S. Code					mean
		text→cite	cite→text	text→cite	cite→text	defined→cite	cite→defined	text→cite	cite→text	cite→amended	defined→cite	cite→defined	
BLT-4k	GPT-3.5-turbo	53	32	72	38	83	79	89	52	56 (0)	77	98	66.3
	GPT-4	82	78	88	97	90	100	98	93	93 (0)	98	100	92.5
	GPT-4-turbo	87	88	85	63	76	95	98	84	77 (7)	96	99	86.2
	chat-bison-001	82	18	61	48	95	78	22	61	63 (1)	7	95	57.3
	chat-bison-32k	84	29	7	77	37	92	83	90	89 (1)	81	97	69.6
	Claude-2.1	54	38	74	71	78	85	97	87	87 (1)	96	95	78.4
BLT-8k	GPT-3.5-turbo	24	17	19	20	65	52	63	58	56 (1)	69	89	48.4
	GPT-4	44	26	64	49	82	83	94	74	76 (0)	88	97	70.6
	GPT-4-turbo	57	53	66	45	75	74	94	80	71 (3)	98	99	73.8
	chat-bison-001	36	7	1	2	2	41	15	35	35 (1)	10	66	22.7
	chat-bison-32k	59	6	9	29	48	59	70	83	86 (1)	80	89	56.2
	Claude-2.1	35	11	58	51	70	54	91	81	79 (4)	94	92	65.1
BLT-16k	GPT-3.5-turbo	7	5	25	12	67	57	38	36	36 (1)	50	77	37.3
	GPT-4-turbo	30	20	78	36	83	83	90	64	58 (5)	93	95	66.4
	chat-bison-32k	42	6	36	30	69	57	25	54	52 (2)	43	58	42.9
	Claude-2.1	21	5	64	48	76	61	82	66	65 (4)	85	85	59.8
BLT-32k	GPT-4-turbo	23	11	42	7	64	63	77	38	36 (1)	82	85	48.0
	chat-bison-32k	13	0	8	10	39	32	12	32	36 (0)	30	36	22.5
	Claude-2.1	20	3	45	24	54	38	76	45	43 (3)	71	70	44.5
BLT-64k	GPT-4-turbo	17	4	27	10	60	54	51	19	16 (4)	55	66	34.5
	Claude-2.1	5	1	52	26	56	35	58	31	29 (4)	55	66	37.6
BLT-128k	GPT-4-turbo	9	3	3	0	20	17	34	12	12 (0)	51	56	19.7
	Claude-2.1	10	0	11	3	16	21	47	21	18 (1)	45	60	22.9

Table 1: Accuracy in percent of several models against all the different sizes of BLT. Since each size’s test split consists of 100 prompts (with answers) for each of the 11 tests, all numbers other than the mean are both the number correct and the percent accuracy. Under cite→amended, the number in parentheses is how often the model erred by returning the unamended U.S. Code text rather than the amended text provided to the model in the prompt, results that Subsection 4.5 discusses further.

4.1 GPT-4 on transcript text→cite

GPT-4 got only 82 of the 100 transcript text→cite test cases correct for BLT-4k. To further investigate this poor performance, we generate 1,000 new prompts in the same format (25-lines per page, with half being one-page and half being two-page). GPT-4 achieves 87.5% on these 1,000. Qualitatively, most errors are either identifying the line after the correct one or before the correct one. The full error breakdowns are in Appendix D.

The biggest determinant of performance is whether the transcript was a single page or two pages. GPT-4 correctly answered 91% of single-

page transcript prompts, but just 84% of 2-page transcript prompts. This makes sense, since 2-page transcripts have 50 lines of text, whereas 1-page transcripts have just 25 lines of text. (An example of GPT-4 getting a wrong answer on a 2-page transcript appears in Appendix C.)

To see whether GPT-4 is confused solely by the greater number of lines or by having the text split into two pages, we generate 500 new prompts with single pages but with 50 lines per page. (In other words, the transcript quotation is all one page, but with line numbers starting at "1:" and ending at "50:", followed by the question). We find GPT-4 achieves 84.8% accuracy, nearly identical to the

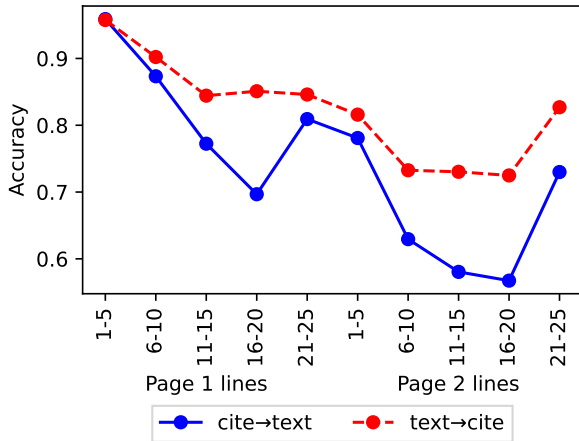


Figure 3: Graph of location of relevant line versus accuracy on both transcript **cite**→**text** and **text**→**cite** on 5,000 prompts to GPT-4.

two-page transcripts, indicating the problem with them is length, not being split into two pages.

To investigate how the location, within the transcripts, of the text impacts accuracy, we generate 5,000 new two-page prompts, and run against GPT-4. The results are the red dashed line in Figure 3. We see a generally downwards trend.

We also perform a sensitivity analysis, trying four question formats other than the standard “What are the page number and line number of the line above with the text “___”?” after the quotations. (Details in Appendix E.) We find the biggest improvement (from 82% to 99%) simply by swapping the question from the end to the beginning and changing “above” to “below”. We tried the same switch – moving the question from the bottom to the top – for all of BLT-4k and re-ran against GPT-4. The results are in Appendix F. It seems the huge improvement in transcript text→cite was anomalous; moving the question from the end to the start actually hurt performance in 7 of 11 tasks. But we do see that GPT-4 is quite sensitive to whether the question is at the start or end of the prompt.

4.2 GPT-4 on transcript cite→text

GPT-4 also performed poorly (78%) on transcript cite→text for BLT-4k. To further investigate, we created 1,000 new prompts in the same format. GPT-4 got 75.7% accuracy on these 1,000. We found little difference between one-page and two-page transcripts, on which GPT-4 got 76.6% and 74.8% respectively. We also created 1,000 new prompts that were one page, with just 15 lines on the page (as opposed to the normal 25). GPT-4 got 77.5% on those shorter transcripts, suggesting

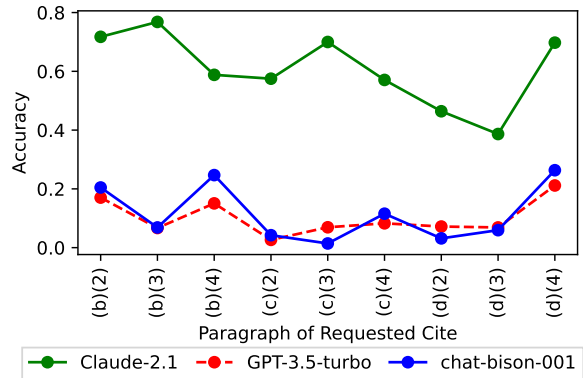


Figure 4: Graph of location of requested cite versus accuracy for 5,000 **cite**→**text** prompts, all using 3-wide, 4-deep synthetic sections, which are 127 lines long. Note that each first subdivision (e.g., (a), (1)) is used for a “General Rule” that has few lines, so such subdivisions are not included in this graph.

a slight accuracy improvement as the number of lines decreases, an intuitive result.

To investigate how accuracy varies with the location of the requested cite, within the transcripts, we generate 5,000 new two-page prompts, and pass them to GPT-4. The results are the solid blue line in Figure 3. We see a trend towards lower accuracy further into the transcript, with higher accuracy near the beginning and end of each page.

4.3 Poor Performance on synthetic cite→text

Several models have their worst performance, among synthetic section tasks, on cite→text. Each BLT-* has a variety of different size sections, shown in Appendix A. Some of the worst BLT-4k performance is on 3-wide, 4-deep synthetic sections. To see if location within the section plays a role in accuracy, we generated 5,000 prompts using 3-wide, 4-deep synthetic sections, and we ran all against three models with particular problems on this task. The results are in Figure 4. There is no U-pattern or other discernable pattern. We found the same lack of discernable pattern in accuracy versus location with all models and all tasks involving either synthetic sections or the U.S. Code.

4.4 GPT-4 on synthetic text→cite

GPT-4 has lackluster performance on synthetic text→cite. The most common error²⁰ is getting a single portion of the hierarchical citation wrong. For example, when the correct answer was “section 2553(b)(2)(B)(i)”, GPT-4 returned “section

²⁰See Appendix H for more detail.

2553(b)(1)(B)(i)”. In the legal context, such mistakes are serious, as these two citations refer to entirely different portions of section 2553.

4.5 Problem Revealed by cite→amended

Recall that for cite→amended we make a minimal but semantically-important change to the subsection of the U.S. Code section being requested. The correct answer is returning the subsection’s text with this amendment. Many of the errors involve returning the text of the wrong subsection. But one type of error is particularly concerning: an LLM returning the subsection *without* the amendment, presumably relying on the original U.S. Code text seen during training. In practice, this error means that, even if a lawyer or paralegal fully pastes in the new version of legislation, the LLM ignores it, giving inaccurate help with the amended text.

The occurrence of this error is in parentheses under the cite→amended column in Table 1, showing that several LLMs make this error a nontrivial number of times. To investigate further, we generated 1,000 new cite→amended prompts. On these, GPT-4-turbo returned the original text 44 times (4.4%), Claude-2.1 did so 17 times (1.7%), and chat-bison-32k did so 4 times (0.4%).

5 Fine-Tuning

We fine-tune the 4,000-token version of GPT-3.5-turbo, since we have not been granted fine-tuning access to more powerful models. Because we are limited to 4,000 tokens, we naturally do fine-tuning with BLT-4k’s training set. For each of the 11 task types, BLT-4k has a training set with 1,000 prompts and answers, for a total of 11,000. Of these, we use 90% for training. We train for two epochs with the default learning rate. The results of the fine tuning are in Table 2. We find that fine-tuning brings GPT-3.5-turbo, which is not OpenAI’s most advanced model, to near the 100% performance expected of lawyers and paralegals.

We tested how this fine-tuned GPT-3.5-turbo performs on SARA, using the 276 cases where the answer is entail/contradict. Each prompt consists of each U.S. Code section(s) mentioned in the case, plus the facts (i.e., the premise) and the hypothesis (i.e., the question). Without fine-tuning, GPT-3.5-turbo’s accuracy was 54.3% (150 / 276), but with our fine-tuning it rises to 60.9% (168 / 276). Qualitatively, the fine-tuned model’s answers focus more clearly on applying just the relevant statutory pro-

text and task	not tuned	fine-tuned
transcript text→cite	53	100
transcript cite→text	32	99
synthetic text→cite	72	98
synthetic cite→text	38	100
synthetic defined→cite	83	100
synthetic cite→defined	79	100
uscode text→cite	89	100
uscode cite→text	52	100
uscode cite→amendedtext	56	100
uscode defined→cite	77	100
uscode cite→defined	98	100

Table 2: Results of fine-tuning GPT-3.5-turbo on 9,900 training samples from BLT-4k. Both numerical columns contain percent accuracy on BLT-4k’s test prompts. Fine-tuning GPT-3.5-turbo improves it to near perfect.

visions, which leads it to answer more questions correctly. An example is in Appendix I.

6 Conclusion

We demonstrate that the best currently available LLMs perform very poorly at many basic legal text-handling tasks when used out of the box. The chief innovation officer at a large international law firm observed to the *New York Times* of LLMs, “At its best, the technology seems like a very smart paralegal.” (Lohr, 2023). We find LLMs are like very sloppy paralegals.

We find poor performance from GPT-3.5-turbo, PaLM 2, and Claude-2.1 on BLT-4k (our smallest test set) and find poor performance from even GPT-4 on finding the text on one or two pages of deposition transcript. But fine-tuning on BLT-4k’s training set brings performance of GPT-3.5-turbo up to the expected human level of performance.

While we focus on law, the BLT tasks are low-level enough that we would expect these findings to be relevant to anyone, regardless of domain. Moreover, the LLMs’ poor performance shows that foundational models cannot be relied on to perform basic domain tasks without fine-tuning.

Our code for BLT allows for the creation of near infinite amounts of training data, but we do not anticipate this benchmark primarily being used to fine-tune. Rather, BLT serves as a representative example of what professionals would expect a capable LLM to handle without examples, in-depth instruction, or expertise in prompt design.

References

Abdi Aidid and Benjamin Alarie. 2023. *The Legal Singularity: How Artificial Intelligence Can Make Law Radically Better*. University of Toronto Press.

Andrew Blair-Stanek, Nils Holzenberger, and Benjamin Van Durme. 2023. [Can GPT-3 Perform Statutory Reasoning?](#) In *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, ICAIL '23, page 22–31.

Bluebook. 2020. *The Bluebook: A Uniform System of Citation, 21st edition*.

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, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. 2022. [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.

Dheeru Dua, Shivanshu Gupta, Sameer Singh, and Matt Gardner. 2022. [Successive prompting for decomposing complex questions](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1251–1265, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Zhiwei Fei, Xiaoyu Shen, Dawei Zhu, Fengzhe Zhou, Zhuo Han, Songyang Zhang, Kai Chen, Zongwen Shen, and Jidong Ge. 2023. [Lawbench: Benchmarking legal knowledge of large language models](#).

Bryan Garner, editor. 2019. *Black's Law Dictionary*.

Neel Guha, Daniel E. Ho, and Julian Nyarko et al. 2022. [Legalbench: Prototyping a collaborative benchmark for legal reasoning](#).

Dan Hendrycks, Collin Burns, Anya Chen, and Spencer Ball. 2021. Cuad: An expert-annotated NLP dataset for legal contract review. *Advances in Neural Information Processing Systems*.

Nils Holzenberger, Andrew Blair-Stanek, and Benjamin Van Durme. 2020. [A dataset for statutory reasoning in tax law entailment and question answering](#). In *Proceedings of the Natural Legal Language Processing Workshop 2020 co-located with the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD 2020), Virtual Workshop, August 24, 2020*, volume 2645 of *CEUR Workshop Proceedings*, pages 31–38. CEUR-WS.org.

Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. 2023. [Decomposed prompting: A modular approach for solving complex tasks](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Mi-Young Kim, Juliano Rabelo, Randy Goebel, Masaharu Yoshioka, Yoshinobu Kano, and Ken Satoh. 2022. [COLIEE 2022 summary: Methods for legal document retrieval and entailment](#). In *New Frontiers in Artificial Intelligence - JSAI-isAI 2022 Workshop, JURISIN 2022, and JSAI 2022 International Session, Kyoto, Japan, June 12-17, 2022, Revised Selected Papers*, volume 13859 of *Lecture Notes in Computer Science*, pages 51–67. Springer.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.

Vladimir I Levenshtein et al. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710. Soviet Union.

David D. Lewis, Lenora Gray, and Mark Noel. 2023. [Confidence sequences for evaluating one-phase technology-assisted review](#). In *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, ICAIL '23, page 131–140.

Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023a. [Lost in the middle: How language models use long contexts](#).

Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023b. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *ACM Comput. Surv.*, 55(9).

Steve Lohr. 2023. [A.I. threatens lawyers? We've heard this before](#). *New York Times*, page B1.

Eric Martínez. 2023. [Re-evaluating gpt-4's bar exam performance](#).

Masha Medvedeva and Pauline McBride. 2023. [Legal judgment prediction: If you are going to do it, do it right](#). In *Proceedings of the Natural Legal Language Processing Workshop 2023*, pages 73–84, Singapore. Association for Computational Linguistics.

John J. Nay, David Karamardian, Sarah B. Lawsky, Wenting Tao, Meghana Bhat, Raghav Jain, Aaron Travis Lee, Jonathan H. Choi, and Jungo Kasai. 2023. [Large language models as tax attorneys: A case study in legal capabilities emergence](#).

OpenAI. 2023a. [Gpt-4 developer livestream](#).

685 OpenAI. 2023b. [Gpt-4 technical report](#).

686 Sumit Pai, Sounak Lahiri, Ujjwal Kumar, Krishanu
687 Baksi, Elijah Soba, Michael Suesserman, Nirmla
688 Pudota, Jon Foster, Edward Bowen, and Sanmitra
689 Bhattacharya. 2023. [Exploration of open large lan-
690 guage models for eDiscovery](#). In *Proceedings of the
691 Natural Legal Language Processing Workshop 2023*,
692 pages 166–177, Singapore. Association for Computa-
693 tional Linguistics.

694 Bhargavi Paranjape, Scott Lundberg, Sameer Singh,
695 Hannaneh Hajishirzi, Luke Zettlemoyer, and
696 Marco Tulio Ribeiro. 2023. Art: Automatic multi-
697 step reasoning and tool-use for large language mod-
698 els. *arXiv preprint arXiv:2303.09014*.

699 Drescher ProParalegal. 2017. [Practice tip sheet](#).

700 Adam Roegiest, Radha Chitta, Jonathan Donnelly,
701 Maya Lash, Alexandra Vtyurina, and Francois
702 Longtin. 2023. [Questions about contracts: Prompt
703 templates for structured answer generation](#). In *Pro-
704 ceedings of the Natural Legal Language Processing
705 Workshop 2023*, pages 62–72, Singapore. Association
706 for Computational Linguistics.

707 Jaromír Savelka, Kevin D. Ashley, Morgan A. Gray,
708 Hannes Westermann, and Huihui Xu. 2023. [Can
709 GPT-4 support analysis of textual data in tasks re-
710 quiring highly specialized domain expertise?](#) In *Pro-
711 ceedings of the 6th Workshop on Automated Semantic
712 Analysis of Information in Legal Text co-located with
713 the 19th International Conference on Artificial Intelli-
714 gence and Law (ICAIL 2023), Braga, Portugal, 23rd
715 September, 2023*, volume 3441 of *CEUR Workshop
716 Proceedings*, pages 1–12. CEUR-WS.org.

717 Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta
718 Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola
719 Cancedda, and Thomas Scialom. 2023. Toolformer:
720 Language models can teach themselves to use tools.
721 *arXiv preprint arXiv:2302.04761*.

722 Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao,
723 Abu Awal Md Shoeb, Abubakar Abid, and
724 Adam Fisch et al. 2023. [Beyond the imitation game:
725 Quantifying and extrapolating the capabilities of lan-
726 guage models](#).

727 Jules White, Quchen Fu, Sam Hays, Michael Sandborn,
728 Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse
729 Spencer-Smith, and Douglas C. Schmidt. 2023. [A
730 prompt pattern catalog to enhance prompt engineer-
731 ing with chatgpt](#).

732 Chaojun Xiao, Haoxi Zhong, Zhipeng Guo, Cunchao Tu,
733 Zhiyuan Liu, Maosong Sun, Yansong Feng, Xianpei
734 Han, Zhen Hu, Heng Wang, and Jianfeng Xu. 2018.
735 [CAIL2018: A large-scale legal dataset for judgment
736 prediction](#). *CoRR*, abs/1807.02478.

737 Fangyi Yu, Lee Quartey, and Frank Schilder. 2022. [Le-
738 gal prompting: Teaching a language model to think
739 like a lawyer](#). *CoRR*, abs/2212.01326.

Libin Zhang. 2023. Tax questions for language models. *Tax Notes Federal*, 179:1699. 740
741

Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang
Zhang, Zhiyuan Liu, and Maosong Sun. 2020. How
does nlp benefit legal system: A summary of legal
artificial intelligence. In *Proceedings of the 58th
Annual Meeting of the Association for Computational
Linguistics*, pages 5218–5230. 742
743
744
745
746
747

A Synthetic Section Sizes 748

Larger versions of BLT have longer and more
749 complicated prompts. Below are the size of syn-
750 thetic sections in each size of BLT. The generated
751 prompts are distributed uniformly among these sec-
752 tion sizes. For example, one-quarter of BLT-16k’s
753 synthetic statutes are 5-wide, 4-deep; one-quarter
754 are 8-wide, 3-deep; and so on. 755

Version	Sizes
BLT-4k	2 wide, 2 deep
	2 wide, 3 deep
	2 wide, 4 deep
	2 wide, 5 deep
	3 wide, 2 deep
	3 wide, 3 deep
	3 wide, 4 deep
	4 wide, 2 deep
BLT-8k	4 wide, 3 deep
	2 wide, 6 deep
	3 wide, 5 deep
	4 wide, 4 deep
BLT-16k	7 wide, 3 deep
	20 wide, 2 deep
	5 wide, 4 deep
	8 wide, 3 deep
BLT-32k	9 wide, 3 deep
	30 wide, 2 deep
	3 wide, 6 deep
	4 wide, 5 deep
	6 wide, 4 deep
	12 wide, 3 deep
BLT-64k	11 wide, 3 deep
	44 wide, 2 deep
	40 wide, 2 deep
	7 wide, 4 deep
	16 wide, 3 deep
BLT-128k	15 wide, 3 deep
	14 wide, 3 deep
	13 wide, 3 deep
	60 wide, 2 deep
	65 wide, 2 deep
	4 wide, 6 deep
	5 wide, 5 deep
8 wide, 4 deep	
9 wide, 4 deep	
20 wide, 3 deep	
80 wide, 2 deep	

756

B Transcript Quotation Sizes

Larger versions of BLT have longer prompts. Below are the number of pages of deposition transcript quotation used in each size BLT model.

Version	Transcript Pages
BLT-4k	1, 2
BLT-8k	5, 10, 15
BLT-16k	25, 40
BLT-32k	30, 60, 80
BLT-64k	100, 130
BLT-128k	120, 140

The generated prompts are distributed uniformly among these page sizes. For example, half of BLT-4k’s prompts have one-page deposition quotations and the other half have two-page deposition quotations.

C Example GPT-4 failing on a 2-page transcript text→cite Prompt

Page 68:

1: after the accident?
2: A. He called me – that same day,
3: he called me at nighttime and asked me how
4: I was.
5: Q. He called you and asked you how
6: you were?
7: A. Yes.
8: Q. What did you tell him?
9: A. I just told him that I wasn’t
10: feeling good, that I was just laying in
11: bed.
12: Q. And what did he say?
13: A. He said take the pills I gave
14: you, you’ll be ready and better by next
15: week. He was just giving me something like
16: a fast way for me to feel up, to cheer up.
17: It wasn’t helping.
18: Q. Going back to your primary care
19: physician. You said she wanted to send you
20: to a specialist?
21: A. Yes.
22: Q. But then you called workers’
23: comp?
24: A. Yeah, that’s when the issue
25: happened and I called workers’ comp.

Page 69:

1: Q. Then who did you treat with
2: next?
3: A. I treated with CityMD.
4: Q. You said CityMD?
5: A. Yes.
6: Q. Was the place that you went to
7: the first time the same place that you went
8: like after you went to your primary care
9: physician?
10: A. Excuse me again?
11: MR. GROSSMAN: You said you
12: went to CityMD. She thinks you went
13: twice, did you go twice?
14: THE WITNESS: No.
15: MR. GROSSMAN: Then you’re not

16: listening to the question. You went
17: to CityMD first, then you went to
18: your primary physician. Now, she
19: wants to know where you went next.
20: THE WITNESS: No, first I went
21: to the primary.
22: MR. GROSSMAN: First primary,
23: and then CityMD and then where did
24: you go?
25: THE WITNESS: I was in CityMD

What are the page number and line number of the line above with the text "your primary physician. Now, she"?

The page number is 69 and the line number is 19.

Note that GPT-4’s answer above is wrong, since the requested text is on line 18, not 19.

D GPT-4 Errors on Transcripts

GPT-4 performs surprisingly poorly on transcripts, even the 1- and 2-page transcripts in BLT-4k. Below are GPT-4’s errors on the large runs used to investigate the problem. Our API-calling code automatically categorizes the errors.

D.1 GPT-4’s errors on transcript text→cite from BLT-4k

Recall that we generated 1,000 new BLT-4k style transcripts with text→cite prompts to investigate GPT-4’s poor performance. Below is a tabulation of the lines returned by GPT-4 relative to the correct ones. We see that the overwhelming majority of errors are returning the line before or line after the correct one.

#	Relative Error
3	prior page
1	19 before
1	3 before
1	2 before
12	1 before (i.e. prior line)
76	1 after (i.e. next line)
1	2 after
1	3 after
1	next page

D.2 GPT-4’s errors on transcript cite→text from BLT-4k

Recall that we generated 1,000 new BLT-4k style transcripts with cite→text prompts to investigate GPT-4’s poor performance on that task. Below is a tabulation of the errors.

#	Type of Error
166	wrong line
48	superset
10	not found
9	subset wrong line
7	multiple wrong lines
3	subset

There was a difference in the type of errors on one-page versus two-page. Almost all (84%) the errors on two-page transcripts was returning the text from the wrong line (almost always the line before or after the correct one). But one-page transcripts were only 51% wrong-line errors, with 35% being superset errors, where GPT-4 returned not only the text of the requested line, but also part of adjacent lines. Here is an example of a superset error that occurred on this dataset:

...

22: MS. WOLVERTON: Objection. Vague. Whose
 23: e-mail records?
 24: Q Ms. Abedin’s e-mail records located. I
 25: apologize.

What is the exact text of just line 25 above?

GPT-4 returned “Q Ms. Abedin’s e-mail records located. I apologize.” which is actually lines 24 and 25, whereas the prompt asks for “the exact text of **just** line 25 above”. These superset errors are much more understandable from a human perspective than an entirely wrong line, but they still fail according to the clear instructions in the prompt, with the word ‘just’.

Of the 1,000 prompts, 166 were “wrong line” errors. Below shows the position, relative to the correct one, from which the text was incorrectly drawn:

#	Relative Error
2	prior page
4	same page, >3 lines before
1	3 lines before
6	2 lines before
81	prior line
39	next line
10	2 lines after
7	3 lines after
6	4 lines after
15	same page, >4 lines after
3	next page

As with text→cite, here we see the overwhelming majority of wrong line errors are returning either

the prior line or next line.

E Sensitivity analysis on GPT-4’s transcript text→cite errors

Recall that GPT-4 performs poorly on text→cite, so we attempted a sensitivity analysis, seeing how performance on the 100 test prompts for the task in BLT-4k changed with changes in the phrasing within the prompt. The results are below.

BLT-4k default What are the page number and line number of the line above with the text " _ "?	82/100
Move question from end to beginning, so question is What are the page number and line number of the line below with the text " _ "?	99/100
Keep question at end, adding “exact” so question is What are the exact page number and the exact line number of the line above with the text " _ "?	84/100
Keep question at end, adding “precise” so question is What are the precise page number and the precise line number of the line above with the text " _ "?	84/100
Keep default question at end, but add the following introduction at the start: Below is a portion of a transcript, with each line starting with a number that is important for referring to that line.	93/100

The improvement from moving the question to the start motivated further experimentation, below.

F Question at Start versus End

Because transcript text→cite against GPT-4 saw large improvements from moving the question from the bottom (which is the standard for all of BLT-* for all tasks) to the top, we attempted the same change for all BLT-4k, running against GPT-4:

text and task	question at...	
	end	start
transcript text→cite	82	99
transcript cite→text	78	85
synthetic text→cite	88	91
synthetic cite→text	97	82
synthetic defined→cite	90	82
synthetic cite→defined	100	98
uscode text→cite	98	96
uscode cite→text	93	82
uscode cite→amendedtext	93	82
uscode defined→cite	98	71
uscode cite→defined	100	100

We see that GPT-4 is quite sensitive to whether the question is at the top or bottom. But moving

the question to the start actually produced worse results for 7 of the 11 tasks. This indicates that the improvement in transcript text→cite by moving the question to the top was an outlier.

G Sensitivity analysis on GPT-4’s transcript cite→text errors

Recall that GPT-4 did poorly on transcript text→cite and that we performed a sensitivity analysis, discussed in Appendix E. Since GPT-4 also did poorly on cite→text, we also do a sensitivity analysis on that.

BLT-4k default What is the exact text of just line _ of page _ above?	78/100
Move question from end to beginning, so question is What is the exact text of just line _ of page _ below?	85/100
Rephrase as Return the exact text of just line 4 of page 59 above?	74/100
Change “exact” to “precise”, so question is What is the precise text of just line _ of page _ above?	79/100
Add clarification, so question is What is the exact text of just line _ of page _ above? Return just the text on that line and return none of the text on the line before or after, even if necessary for a full sentence.	62/100
Keep default question at end, but add the following introduction at the start: Below is a portion of a transcript, with each line starting with a number that is important for referring to that line.	90/100

H GPT-4’s errors on synthetic text→cite

GPT-4 got 12 of 100 wrong on the synthetic text→cite of BLT-4k, which is GPT-4’s worst performance on a non-transcript task. Those errors are in the table below. An actual example of a wrong paragraph is returning “section 2553(b)(1)(B)(i)” when the correct answer is “section 2553(b)(2)(B)(i)”. The first part of a statute is the subsection (e.g. the “(b)” in section 2553(b)), the next part is a paragraph (e.g. the “(1)” in section 2553(b)(1)), then a subparagraph, a clause, a subclause, an item, and a subitem.

#	Type of Error
3	wrong subparagraph
3	wrong paragraph
2	wrong paragraph, subparagraph, & clause
1	wrong subsection, paragraph, & subclause
1	wrong subsection, paragraph, subparagraph, clause
2	not parallel

Above, we see six errors are just one part of the citation wrong, four involve multiple parts wrong (but the correct structure), and two were totally not parallel. An actual answer that is “not parallel” is returning “section 9680(a)(1)(B)(2)(C)(i)(I)” when the correct answer is “section 9680(b)(2)(C)(i)(I)”.

GPT-4 got 36 of 100 wrong on the same task in BLT-8k, and those errors are tabulated below. An actual example of a wrong item is returning “section 3632(c)(2)(B)(ii)(II)(aa)” when the correct answer is “section 3632(c)(2)(B)(ii)(II)(bb)”.

#	Type of Error
7	wrong paragraph
5	wrong subclause
5	wrong subparagraph
5	wrong clause
2	wrong item
2	wrong subsection
3	wrong subsection, wrong paragraph
1	wrong paragraph, wrong subparagraph
2	wrong paragraph, subparagraph, & clause
1	wrong paragraph, clause, & subclause
3	not parallel

Above we again see most errors are just getting the citation wrong in a single component.

I Example of fine-tuned GPT-3.5-turbo performing better

Recall that fine-tuning GPT-3.5-turbo on part of BLT-4k’s training set improved its performance on the SARA dataset (Holzenberger et al., 2020). Below we have an example SARA prompt used, plus both non-tuned GPT-3.5-turbo’s answer and the fine-tuned version’s answer. We see the fine-tuned version sticks much more closely to the relevant statutory text and thus comes to the correct answer.

The prompt (edited with ellipses for brevity):
 We are going to be doing Entailment/Contradiction reasoning applying the statute(s) below:

955 §2. Definitions and special rules
 956 (a) Definition of surviving spouse
 957 (1) In general
 958 For purposes of section 1, the term "surviving spouse" means
 959 a taxpayer-
 960 (A) whose spouse died during either of the two years im-
 961 mediately preceding the taxable year, and
 962 (B) who maintains as his home a household which consti-
 963 tutes for the taxable year the principal place of abode (as a
 964 member of such household) of a dependent (i) who (within
 965 the meaning of section 152) is a son, stepson, daughter, or
 966 stepdaughter of the taxpayer, and (ii) with respect to whom the
 967 taxpayer is entitled to a deduction for the taxable year under
 968 section 151.
 969 ***
 970 (b) Definition of head of household
 971 (1) In general
 972 ***
 973 (2) Determination of status
 974 Notwithstanding paragraph (1),
 975 (A) an individual who is legally separated from his spouse
 976 under a decree of divorce or of separate maintenance shall not
 977 be considered as married;
 978 (B) a taxpayer shall be considered as not married at the
 979 close of his taxable year if at any time during the taxable year
 980 his spouse is a nonresident alien; and
 981 (C) a taxpayer shall be considered as married at the close
 982 of his taxable year if his spouse (other than a spouse described
 983 in subparagraph (B)) died during the taxable year.
 984 ***
 985 Premise: Alice and Bob got married on Feb 3rd, 1992. Alice
 986 died on July 9th, 2014.
 987 Hypothesis: Section 2(b)(2)(C) applies to Bob in 2014.
 988 Answer:

989 **Non-tuned GPT-3.5-turbo's incorrect answer:**
 990 To determine whether Section 2(b)(2)(C) applies to Bob in
 991 2014, we need to analyze the relevant provisions of the tax
 992 code.
 993 Section 2(b)(2)(C) states that a taxpayer shall be considered
 994 as married at the close of his taxable year if his spouse (other
 995 than a spouse described in subparagraph (B)) died during the
 996 taxable year.
 997 In this case, Alice, Bob's spouse, died on July 9th, 2014. To
 998 determine if Section 2(b)(2)(C) applies to Bob in 2014, we
 999 need to consider whether Alice's death during the taxable year
 1000 qualifies Bob as a surviving spouse.
 1001 According to Section 2(a)(1)(A), a surviving spouse is defined
 1002 as a taxpayer whose spouse died during either of the two
 1003 years immediately preceding the taxable year. In this case,
 1004 Alice died in 2014, which is within the two years immediately
 1005 preceding the taxable year.
 1006 Next, we need to consider Section 2(a)(1)(B), which states that
 1007 a surviving spouse must maintain as his home a household that
 1008 constitutes the principal place of abode of a dependent. The
 1009 dependent must be a son, stepson, daughter, or stepdaughter of
 1010 the taxpayer, and the taxpayer must be entitled to a deduction
 1011 for the taxable year under Section 151.
 1012 In this case, we don't have information about whether Bob has
 1013 a dependent who meets the criteria mentioned above. There-
 1014 fore, we cannot determine if Bob meets the requirements of a
 1015 surviving spouse under Section 2(a)(1)(B).
 1016 Since we cannot determine if Bob meets the requirements
 1017 of a surviving spouse, we cannot conclude whether Section
 1018 2(b)(2)(C) applies to Bob in 2014.

Fine-tuned GPT-3.5-turbo's correct answer: 1019
 To determine if Section 2(b)(2)(C) applies to Bob in 2014, we 1020
 need to consider the marital status of Bob and whether his 1021
 spouse died during the taxable year. 1022
 According to the premise, Alice died on July 9th, 2014, and 1023
 she was married to Bob. Therefore, we need to determine 1024
 Bob's marital status at the close of the taxable year 2014. 1025
 Since Bob and Alice got married on Feb 3rd, 1992, and Alice 1026
 died in 2014, Bob was married to Alice at the close of 2014. 1027
 Therefore, Section 2(b)(2)(C) applies to Bob in 2014 because 1028
 his spouse (Alice) died during the taxable year. 1029