Extracting memorized pieces of (copyrighted) books from open-weight language models

Anonymous Authors[1]

Abstract

Plaintiffs and defendants in copyright lawsuits over generative AI often make sweeping, opposing claims about the extent to which large language models (LLMs) have memorized plaintiffs' protected expression. Drawing on adversarial ML and copyright law, we show that these polarized positions dramatically oversimplify the relationship between memorization and copyright. To do so, we leverage a recent probabilistic extraction technique to extract pieces of the `Books3` dataset from 13 open-weight LLMs. Through numerous experiments, we show that it's possible to extract substantial parts of at least some books from different LLMs. This is evidence that the LLMs have memorized the extracted text; this memorized content is copied inside the model parameters. But the results are complicated: the extent of memorization varies both by model and by book. With our specific experiments, we find that the largest LLMs don't memorize most books—either in whole or in part. However, we also find that LLAMA 3.1 70B memorizes some books, like *Harry Potter* and *1984*, almost entirely. We discuss why our results have significant implications for copyright cases, though not ones that unambiguously favor either side.

1. Introduction

In the dozens of pending copyright suits over training LLMs, the opposing parties have tended to present the technical operation of models in simplified terms. Plaintiffs say LLMs are just giant (infringing) copy machines that store their works and recombine them in their outputs (Kadrey v. Meta Platforms, Inc.). Defendants say LLMs merely contain linguistic relationships—"statistical correlations" (Concord Music Group, Inc. v. Anthropic PBC)—and don't copy the plaintiffs' works. The situation is more complicated than either side suggests.

Appreciating why requires a deeper understanding of training-data extraction, training-data memorization, and the relationship between the two (Section 2). While extraction refers to recovering specific training data from a model's generated outputs, memorization is broader: it involves reconstructing specific training data by examining the model "through any means" (Cooper et al., 2023, Glossary). It is an uncontroversial statement in machine learning (ML) to say that the extraction of a piece of training data implies that the model has memorized that piece of training data (Carlini et al., 2023b; 2021; Schwarzschild et al., 2024)—that extraction is evidence of memorization *inside* the model. As Carlini explains, when a sufficiently large and unique piece of training data is extracted, "the only possible explanation is that the model has somewhere internally stored [that piece of training data]. There just is no other explanation; it can't be due to chance" (Carlini, 2025). As Cooper and Grimmelmann note, "in order to be able to extract memorized content from a model at generation time, that memorized content must be encoded in the model's parameters. There is nowhere else it could be. A model is not a magical portal that pulls fresh information from some parallel universe into our own" (Cooper & Grimmelmann, 2024, p. 25).

Memorization may have significant consequences for ongoing copyright litigation. Extraction generates a "copy" of training data, but it also demonstrates the existence of a "copy" of that training data is memorized inside the model itself. The model being a "copy" (in a technical sense that copyright cares about) has important implications (Sections 3 & 5). Notably, the models themselves could be deemed infringing copies of the training data they've memorized (Cooper & Grimmelmann, 2024; Lee et al., 2023b). Copyright law offers the destruction of infringing materials as a remedy. So, just as courts have ordered the destruction of bootleg DVDs, a court could order the destruction of infringing models (Lee et al., 2023b; Samuelson, 2023; Wilf-Townsend, 2024).

When seeking such remedies in their lawsuit complaints, plaintiffs often cite ML research papers that demonstrate

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on ICML 2025 Workshop on Reliable and Responsible Foundation Models. Do not distribute.

Prompt (prefix)	Target (suffix)	Generations
They were careless people, Tom and Daisy — they smashed up things and creatures and then retreated	back into their money or their vast carelessness, or whatever it was that kept them together, and let other people clean up the mess they had made.	back into their money or their vast carelessness, or whatever it was that kept them together, and let other people clean up the mess they had made .
They were careless people, Tom and Daisy — they smashed up things and creatures and then retreated	back into their money or their vast carelessness, or whatever it was that kept them together, and let other people clean up the mess they had made.	back into their money or their vast carelessness, or whatever it was that kept them together, and let other people clean up the mess they had made.
They were careless people, Tom and Daisy — they smashed up things and creatures and then retreated	back into their money or their vast carelessness, or whatever it was that kept them together, and let other people clean up the mess they had made.	into their money or extraness or whatever it was and let other people clean up the mess they had made. F. Scott Fitzgerald 1

training-data extraction from open-weight models and production systems. First, citing a research paper that extracts, for example, an alphabetized list of U.S. states from ChatGPT (Nasr et al., 2023) is *not* evidence that ChatGPT has memorized a plaintiff's own work (Carlini, 2025; Cooper & Grimmelmann, 2024). Most training data are not memorized—especially not in high-quality, contemporary LLMs. Second, while most ML research reports average extraction rates for a given extraction attack, averages are not necessarily relevant to specific copyright infringement claims (Section 3). Instead, as others have also noted (Lee et al., 2023b; Cooper & Grimmelmann, 2024), it may be more relevant to show the extent to which the *specific* model in question has memorized the plaintiff's *specific* copyrighted work. Third, adversarial extraction of the type we employ here may require thousands or even millions of efforts to generate a given 50-token sequence extracted from a given book. That number is greater than chance would suggest, so it is evidence for memorization; but it is *not* evidence that any user is actually likely in practice to generate even 50tokens from that book, much less most or all of that book.

the (now notorious) torrented corpus of nearly $200,000$ books. `Books3` is in The Pile (Gao et al., 2020)—an LLM pre-training dataset that has been the focus of ongoing litigation (Reisner, 2023; Kadrey et al. v. Meta Platforms, Inc.). In summary:

- We discuss a recent probabilistic extraction method (Hayes et al., 2025), which is a more suitable tool than prior approaches for revealing information about memorization that is useful for copyright (Section 2).
- We then connect extraction risk (and what this reveals about probabilistic copies of training data inside of models) to U.S. copyright law. To do so, we also explain why *average* extraction rates—the typical metric in ML papers—don't provide sufficient information for *specific* claims about how much an LLM has memorized *specific* works (Section 3).
- We then develop a simple and precise procedure for the probabilistic extraction of specific memorized text, which we apply to a set of 56 books in `Books3` (Section 4). Through extensive experiments, we find that while the largest models don't memorize most books—either in whole or in part—LLAMA 3.1 70B memorizes some books, like *Harry Potter* and *1984*, almost entirely.
- Finally, we discuss how our results have significant implications for copyright cases, though not ones that unambiguously favor either side (Section 5).

2. Quantifying memorization as a probabilistic phenomenon

Training-data extraction and memorization are related concepts, but differ in subtle ways. Following Cooper & Grimmelmann (2024), "[m]ost narrowly, when a user intentionally and successfully prompts a model to generate an output that is an exact or near-exact copy of a piece of training data, that is **extraction**."[1] More broadly, "…when an exact or near-exact copy of a piece of training data can be reconstructed by examining the model 'through any means,' that is **memorization**" (Cooper & Grimmelmann, 2024, Part II.A). This distinction summarizes, at a high level, a large amount of work in adversarial machine learning that runs **extraction attacks** to reveal *privacy* (as opposed to copyright (Carlini, 2025)) risks in LLMs.

However, it's incorrect to think that extraction is only interesting as a generation-time phenomenon. It's evidence of memorization of training data *inside the model* (Cooper & Grimmelmann, 2024; Carlini, 2025; Schwarzschild et al., 2024; Lee et al., 2023b; Feldman, 2020). In fact, successful extraction is one of the most common forms of evidence in ML for quantifying memorization in LLMs (Nasr et al., 2025; Hayes et al., 2025; Prashanth et al., 2024; Carlini et al., 2023a; Lee et al., 2022; Zhang et al., 2023). The standard metric in many frontier model-release reports (Gemini Team et al., 2024; Team et al., 2024; Biderman et al., 2023; Grattafiori et al., 2024) measures **discoverable extraction**: one takes a piece of text from the training data, splits it into a **prefix** and a **target suffix**, uses the prefix as a prompt to the LLM, and counts the example as extracted if the LLM generates text that *exactly* matches the target suffix (Carlini et al., 2021; 2023b; Nasr et al., 2023). For instance, consider the quote from *The Great Gatsby* (Fitzgerald, 1925) in our headline figure on page 1. The middle row in

[1] **Regurgitation** is extraction without user intent (Cooper & Grimmelmann, 2024). We focus on extraction.

the figure indicates the verbatim generation of the target suffix in response to the prefix prompt.

Cooper and Grimmelmann explain why such measurements are incomplete when it comes to the kinds of information relevant to copyright. These measurements capture a simple binary (yes-or-no) outcome of *if* a particular piece of training data was extracted (and thus memorized). But LLMs are probabilistic; they can generate different outputs for the same input—just like in the headline figure on page 1. So, instead, a copyright-relevant claim about extraction will likely

> take the following form: a model, when (a) given a particular type of input, will (b) produce a particular type of memorized output, (c) with a particular probability. That probability could be .01 (i.e., a 1% chance), it could be .35 (i.e., a 35% chance) .... The issue for copyright law ...is what to do with this knowledge[:] ...what to do with the fact that element (c)—the probabilistic element—is inescapable (Cooper & Grimmelmann, 2024, Part II.C).

We elaborate on this point in Sections 3 and 5. For now, we note that Hayes et al. (2025) developed a probabilistic extraction methodology that achieves precisely this. In contrast to the standard *yes-or-no* measurements of discoverable extraction, they propose a measure of **probabilistic discoverable extraction**, which quantifies the *probability* (between 0 and 1) that the LLM, under a specific decoding scheme, outputs the verbatim target suffix when prompted with a given prefix. For a training example $\boldsymbol{z}$ that is $a+k$ tokens long, they divide $\boldsymbol{z}$ into an a-length prefix $\boldsymbol{z}_{1:a}$ and a k-length suffix $\boldsymbol{z}_{a+1:a+k}$, and compute the probability $p_{\boldsymbol{z}}$ of extracting example $\boldsymbol{z}$. That is, they compute

$$p_{\boldsymbol{z}} = \prod_{t=a+1}^{a+k} p(z_t \mid \boldsymbol{z}_{1:t-1}), \tag{1}$$

where the prompt is prefix $\boldsymbol{z}_{1:a}$ and the target suffix is $\boldsymbol{z}_{a+1,a+k}$. This equation just captures that the probability of generating the exact target suffix $\boldsymbol{z}_{a+1:a+k}$ is the product of the conditional probabilities of each token z_t in the suffix, conditioned on all preceding tokens. For each position t from $a+1$ to $a+k$, we compute the probability of token z_t given the preceding sequence $\boldsymbol{z}_{1:t-1}$, which includes both the prompt prefix $\boldsymbol{z}_{1:a}$ and any earlier tokens in the suffix. The overall probability of the suffix is simply the product of these per-token probabilities.[2] As an example, on the left in Figure 1, we plot $p_{\boldsymbol{z}}$ for various models and the "careless people" quote.

In general, even for relatively short suffixes, $p_{\boldsymbol{z}}$ should be very small; it involves repeatedly multiplying together probabilities (i.e., numbers smaller than 1). For example, let's consider that each token z_t in a 50-token suffix has probability 0.9—a very high probability, considering that the entire token vocabulary is large (e.g., $32,000$ for Llama models). This means all of the other tokens combined (e.g., $31,999$ for Llama) share the remaining probability of 0.1. The probability of this whole suffix would be 0.9 multiplied 50 times: $0.9^{50} \approx 0.005$—just a 0.5% chance that it would be generated verbatim! But in Figure 1 (left), the probability of the target suffix is larger than this *for every model*. For LLAMA 1 30B, $p_{\boldsymbol{z}}$ is over 35%—greater than $1/3$ generations with this prompt will result in the *verbatim* suffix, which is more often than illustrated in the headline figure!

Such high probabilities *are*, by definition, memorization. This is precisely what makes memorization interesting: *any* arbitrary sequence should have low probability, so the fact that certain sequences have high probability—high enough for the model to generate the verbatim suffix with meaningful likelihood (Section 3)—"can't be due to chance" (Carlini, 2025). These probabilities reflect patterns, "statistical correlations" (Concord Music Group, Inc. v. Anthropic PBC) that the model has learned from the training data. But when these probabilities are unusually large, "the pattern *is* the memorized training data" (Cooper & Grimmelmann, 2024, Part II.C), copied inside the model.

Other ways to understand $p_{\boldsymbol{z}}$. It turns out that $p_{\boldsymbol{z}}$ is very simple and efficient to compute in practice for verbatim extraction; we defer the details to Appendix B. Beyond directly comparing $p_{\boldsymbol{z}}$ values—for the same example across LLMs, across multiple examples for the same LLM, etc.—Hayes et al. (2025) offer an intuitive alternative. In their work, they emphasize that, once one has $p_{\boldsymbol{z}}$, with some simple math one can determine how many times n one would need to prompt the model with the prefix, in order to guarantee that with probability p that the LLM outputs the exact target suffix. (This is where their metric derives its name, **(n,p)-discoverable extraction**.) For a given $p_{\boldsymbol{z}}$, one can pick a probability threshold p and get the corresponding number of prompts n (or vice versa) with

$$1-(1-p_{\boldsymbol{z}})^n \geq p \quad \Rightarrow \quad n \geq \frac{\log(1-p)}{\log(1-p_{\boldsymbol{z}})}. \tag{2}$$

Since $p_{\boldsymbol{z}}$ is the probability of generating the suffix with 1 prompt to the model, the probability of *not* generating it is $(1-p_{\boldsymbol{z}})$. The probability of *not* generating it in n prompts is $(1-p_{\boldsymbol{z}})^n$, and so the probability of generating it in n prompts is $1-(1-p_{\boldsymbol{z}})^n$. As an intuition, think about a fair coin flip coming up heads. Here, $p_{\boldsymbol{z}} = 0.5$ (the probability of heads in one flip). With 2 actual flips ($n=2$), the probability of flipping heads *at least once* is $1-(1-0.5)^2 = 0.75 = p$; when $n=10$, $p=0.9990$. Figure 1 (right) visualizes how n changes for different settings of p for the "careless people" quote. In practice, for LLMs, we can compute $p_{\boldsymbol{z}}$ *without* prompting n times (Appendix B).

[2] For numerical stability, we compute (1) as the exp of the summed per-token log probabilities (Appendix B).

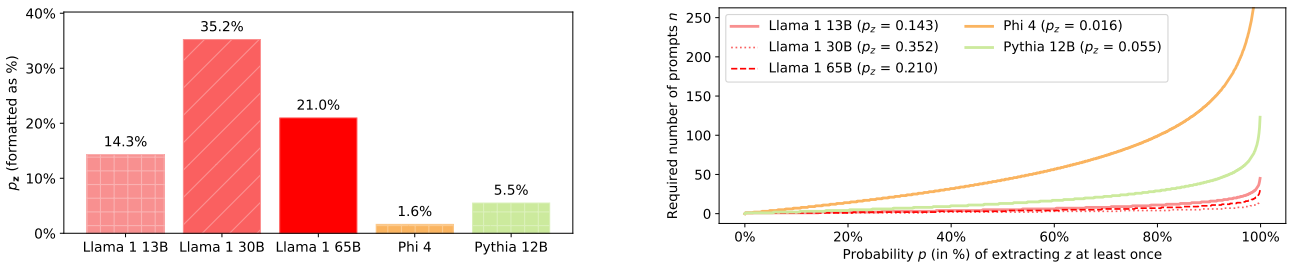

Figure 1: Plotting extraction probability $p_{\boldsymbol{z}}$ for the "careless people" quote from *The Great Gatsby* (see page 1). (**left**) $p_{\boldsymbol{z}}$ (1) for different models. (**right**) Translating $p_{\boldsymbol{z}}$ into how many prompts n it would take to extract the example z with probability p (2). See Appendix A for details.

3. Beyond average extraction rates: Extraction risk, copies, and copyright

There are dozens of pending copyright suits based on the training and output of LLMs (Chat GPT Is Eating the World). Those copyright disputes generally present three interrelated issues: (1) whether training an LLM on copyrighted material is **fair use** (i.e., limited use of the copyrighted material can, under certain circumstances, be used without permission from the copyright owner); (2) whether the model itself is a **copy** or **derivative work** (Section 5) of the works on which it is trained; and (3) whether the model outputs copyrighted material. Some suits present only one issue, while others present all three questions. Some suits are based on content owned by a single company, while others are class action lawsuits purporting to represent all book authors. There are other copyright issues, as well (Lee et al., 2023b; Samuelson, 2023; Sag, 2024a; Goodyear, 2025; Sobel, 2024; Cooper et al., 2024), for example, whether the outputs of generative-AI systems are themselves copyrightable (Lee et al., 2023b; United States Copyright Office).

Memorization and probabilistic copies. Our paper is not about training and fair use. Those issues have been discussed extensively elsewhere (Lemley & Casey, 2021). Nor is our paper about the outputs of LLMs and copyright infringement (Cooper et al., 2023; Lee et al., 2023b; Henderson et al., 2023; Lee et al., 2024; Cyphert, 2024; Sag, 2024b; Bracha, 2024; Charlesworth, 2025), although we use extraction of content from models at generation time as evidence for memorization (Section 2). This paper is about memorization (Section 2). We plan to discuss additional nuances about extraction in future work.

Memorization may matter for copyright law in two ways. First, if a model memorizes all or a substantial portion of a copyrighted work (near-)verbatim, the model itself may be an infringing copy or derivative work (Lee et al., 2023b; Cooper & Grimmelmann, 2024). As discussed above (Sections 1 & 2), this is because memorized training data are encoded inside the model. Others have argued that encoding the work in the form of model weights satisfies the technical definition of "copy" in the U.S. Copyright Act (Cooper & Grimmelmann, 2024). That copy, like the copy used as training data during training, might be fair use, but the analysis of fair use would look somewhat different than internal use in training. This is particularly true for open-weight models like different DeepSeek and Llama models, which are not merely used by the developer but are themselves shared with others (Section 5). Second, memorization of all or part of a particular work may increase the likelihood that the output of the model may be substantially similar to the copyrighted work. That output (at generation time) will be judged separately from model (which is also an artifact that is independent of its outputs), and is less likely to be a fair use (Henderson et al., 2023; Lee et al., 2023b). Two important observations about memorization and extraction follow from these points.

Average extraction rates aren't all you need. Numerous studies have shown that models memorize certain amounts of the works they were trained on (e.g., Carlini et al., 2021; Nasr et al., 2023; Hayes et al., 2025; Lee et al., 2022; Biderman et al., 2023; Gemini Team et al., 2024; Team et al., 2024). In this research, the extent of memorization is typically quantified through an overall extraction rate—much like Figure 2 (left). That is, for some (part of a) training dataset, researchers draw (typically at random) examples of a specified length (e.g., 100 tokens), prompt with the first half of the example (i.e., the example prefix), and count extraction as successful if the resulting generation matches the target suffix (Section 2). The extraction rate is computed as the number of attempted extractions that succeeded, relative to the total number of attempts. These reported averages are generally small, just as we observe in Figure 2; this often forms the basis of defendants in infringement suits calling memorization a rare "bug" (Cooper & Grimmelmann, 2024; OpenAI, 2024).

However, while average extraction rates are useful for estimating *overall* memorization of `Books3` text, they clearly tell an incomplete story about the degree to which specific, underlying pieces of text have or haven't been memorized by the model. Low extraction rates signal that models likely

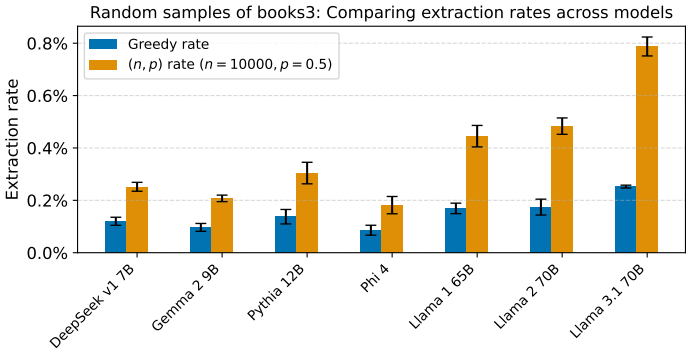

Model	Result	% of book with $\geq$ the given $p_{\boldsymbol{z}}$		
		0.75	**0.50**	**0.01**
LLAMA 3.1 70B	*Harry Potter*	15.10%	42.42%	91.14%
	Sandman Slim	0.00%	0.13%	0.38%
Llama 1 65B	*Harry Potter*	1.69%	4.40%	25.48%
	Sandman Slim	0.09%	0.10%	0.27%
PYTHIA 12B	*Harry Potter*	0.00%	0.08%	0.40%
	Sandman Slim	0.22%	0.26%	0.34%

Figure 2: (**left**) *Average* extraction rates of `Books3` text, comparing the traditional (greedy-sampled) discoverable extraction rate (blue) and the probabilistic extraction rate for a given setting of n and p (orange; see Equation 2) for different models (Appendix F). Greedy-sampled extraction underestimates the true amount of memorization (Hayes et al., 2025). (**right**) For different minimum probabilities of extraction $p_{\boldsymbol{z}}$ (0.75, 0.50, 0.01), we show how much of the whole books (in %) can be extracted with (at least) that $p_{\boldsymbol{z}}$ for *Harry Potter* (Rowling, 1998) and *Sandman Slim* (Kadrey, 2009), which is one of the books by plaintiff Richard Kadrey in *Kadrey et al. v. Meta, Inc.* (Kadrey et al. v. Meta Platforms, Inc.). Nearly *all* of *Harry Potter* is memorized by LLAMA 3.1 70B (measured with respect to minimum extraction probability $p_{\boldsymbol{z}} = 0.01$).

don't memorize *most* text in `Books3`, but it's possible that *specific* pieces of text are highly memorized. We observed this already with Figure 1 for the "careless people" quote, and we can similarly see this in the Table in Figure 2. Some books, like Richard Kadrey's *Sandman Slim* (Kadrey, 2009) are hardly memorized at all. In contrast, *Harry Potter and the Sorcerer's Stone* (Rowling, 1998) is highly memorized in Llama models, especially LLAMA 3.1 70B (Section 4). Over 90% of the book can reconstructed with $p_{\boldsymbol{z}} = 0.01$; nearly half of the book can be reconstructed with $p_{\boldsymbol{z}} = 0.50$—just two prompts (in expectation) per example.

Using average extraction rates, prior work also observes that larger models exhibit larger amounts of memorization (Carlini et al., 2023b; Hayes et al., 2025). We find that these claims (reasonably) do not cleanly generalize to the extent to which models memorize specific examples. In Figure 1, the smaller LLAMA 1 30B exhibits a significantly higher $p_{\boldsymbol{z}}$ than the larger LLAMA 1 65B—0.352 compared to 0.21 (shown in % in the figure). Similarly, using a minimum $p_{\boldsymbol{z}} = 0.75$, we estimate for *Sandman Slim* that LLAMA 3.1 70B memorizes 0% of the text, while LLAMA 1 65B and PYTHIA 12B memorize 0.09% and 0.22%, respectively. (We detail how we compute these estimates in Appendix G.)

Reasoning about "meaningful" extraction probabilities. We show clear evidence that LLAMA 3.1 70B memorizes almost all of *Harry Potter and the Sorcerer's Stone*. For this model, there exists a prefix/suffix combination such that almost every piece of text in the entire book has a high extraction probability $p_{\boldsymbol{z}}$ (Figure 3).[3] These probabilities are too high to be due to random chance; they are clearly memorization (Cooper & Grimmelmann, 2024; Carlini et al., 2021). While it's straightforward to support the claim that LLAMA 3.1 70B memorizes (with respect to probabilistic extraction) nearly all of *Harry Potter*, most cases aren't this clear. In contrast, $p_{\boldsymbol{z}}$ being close to 0 indicates that an example isn't memorized. This is the case for PYTHIA 12B on nearly every example in *Harry Potter* (Section 4).

What is less clear, however, is where to draw a line between the enormous range in between these two extremes: is there a sufficiently small value of $p_{\boldsymbol{z}}$, such that generating the target suffix for $\boldsymbol{z}$ isn't meaningful for copyright—i.e., that it is effectively an instance of "a monkey at the typewriter" (Borel, 1913; The Simpsons: Last Exit to Springfield (Fox television broadcast Mar. 11, 1993))? While we don't address this in this work, we conservatively only consider probabilities $p_{\boldsymbol{z}}$ ($\geq 1 \times 10^{-6}$) that are so high that they aren't by happenstance. However, it's also worth noting that, as others have observed, LLMs aren't like monkeys randomly outputting tokens (Cooper & Grimmelmann, 2024). The model has learned to generate structured, grammatically correct sentences. That means the actual distribution of possible tokens an LLM would reasonably output is considerably smaller than completely random output would suggest. That possibility, which we hope to explore in further work, bears on the question of how high a $p_{\boldsymbol{z}}$ we should consider to be relevant as evidence of memorization of a specific book rather than, say, learning that in most sentences verbs follow subjects and precede objects.

[3]Recall from Section 2 that $p_{\boldsymbol{z}} = 0.005 \approx 0.9^{50}$ corresponds to each token in a 50-token sequence having probability of 0.9; $p_{\boldsymbol{z}} = 0.01$ means these probabilities are even higher.

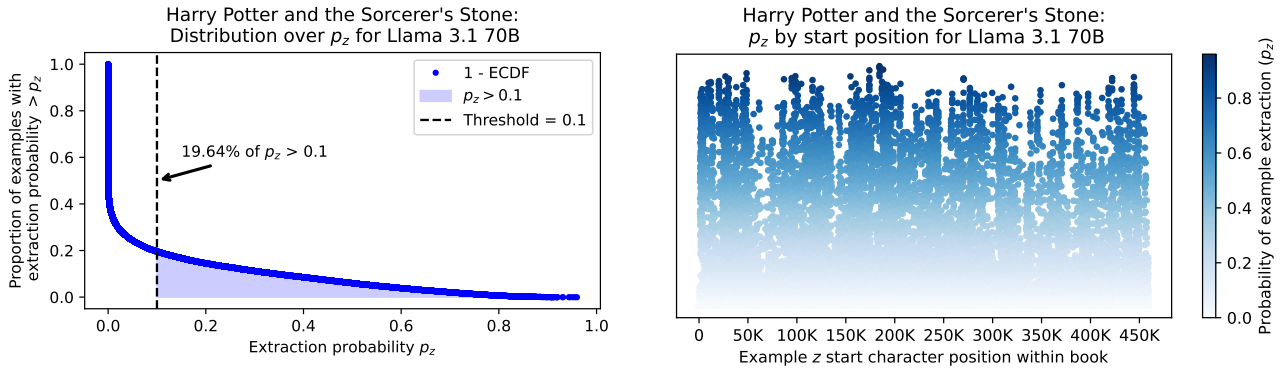

Figure 3: Visualizing memorization of *Harry Potter* in LLAMA 3.1 70B. (**left**) The distribution over $p_{\boldsymbol{z}}$, plotted as the complement of the empirical CDF. This shows the proportion of examples with $p_{\boldsymbol{z}}$ (y-axis) greater than the a given probability (x-axis), e.g., 19.54% of examples have $p_{\boldsymbol{z}} > 0.1$. (**right**) $p_{\boldsymbol{z}}$ for each example according to its start position within the overall book. White gaps along the x-axis indicate that an example (with that start position) was not extracted (i.e., has 0 probability); a continuous horizontal band of blue (associated with the probability on the y-axis) effectively indicates that we can extract the entire book with that associated probability. See Appendix G.1.

4. Quantifying memorization for pieces of individual books

We now dig deeper into average extraction rates, showing the presence and extent of memorization in open-weight models for a sample of books drawn text from `Books3`. Overall, we show how memorization varies across models, across books, and within individual books; memorization for some (but not all) books and for some (but not all) models is surprisingly high. We describe our experiments for identifying memorization "hot-spots" in a selection of books drawn from `Books3`, and how these experiments help us narrow in on follow-up experiments that let us compare the memorization of different pieces of books text within a model, and memorization of a given piece of text across models.

Setup. Similar to Hayes et al. (2025), we use top-k decoding as the sampling algorithm, with temperature $T = 1$ and $k = 40$. The experiments in this section are limited to 5 models: PYTHIA 12B (Biderman et al., 2023), PHI 4 (a 14B model) (Abdin et al., 2024), LLAMA 1 13B, LLAMA 1 65B (Touvron et al., 2023a), and LLAMA 3.1 70B (Grattafiori et al., 2024). Pythia was trained on The Pile (Gao et al., 2020) (which contains the torrented `Books3` dataset (Reisner, 2023)), and Touvron et al. (2023a) explicitly notes that Llama-1 models were trained on `Books3`. PHI 4 was trained predominantly on synthetic data. We also include the newer-generation LLAMA 3.1 70B. We describe each of these choices and important background context more in Appendix G.1. We ran experiments across `Books3` for computing average extraction rates, but otherwise limit ourselves to a selection of 56 books (out of the nearly $200,000$) that are in `Books3`. These books include those listed with the associated plaintiffs in the (amended) class action complaint of Kadrey et al. v. Meta (Kadrey et al. v. Meta Platforms, Inc., pp. 4-5), as well as generally popular books (e.g., Rowling, 1998; Tolkien, 1937; Camus, 1955; Heller, 1961)), and some academic books (e.g., Zittrain, 2008; Barolini, 2006)).

We suspected that some popular books would exhibit high degrees of memorization (due to duplicated text from other sources), and that more obscure books wouldn't. Overall, our results shouldn't be read as a complete account of memorization across the entire `Books3` dataset. Nor did we select the books at random; we deliberately tried to capture variation across the dataset (Appendices E–G). In the future, we'll expand to a much wider set of models and books. In general, we refer to the Appendix, which contains results from nearly one thousand experiments on 13 models.

Identifying memorized book snippets. To identify regions of memorization, we take the following "panning for gold" approach. For a given book, we start at the beginning of the text file in `Books3`. We sample a chunk of text that is sufficiently long to contain 100 tokens of corresponding tokenized text, slide 10 *characters* forward in the book text and repeat this process. We do this for the entire length of the book, which results in approximately one example every 10 characters, e.g., *The Great Gatsby* has $270,870$ characters, which results in roughly $27,000$ examples. As a result, even with our very efficient implementation of probabilistic extraction, these experiments are expensive. We explore more efficient alternatives in Appendix G.2, and intend to pursue this further in future work.

These 100-token examples overlap significantly; *this is deliberate*. Since we don't often know how open-weight models were trained, it's not exactly clear how we should

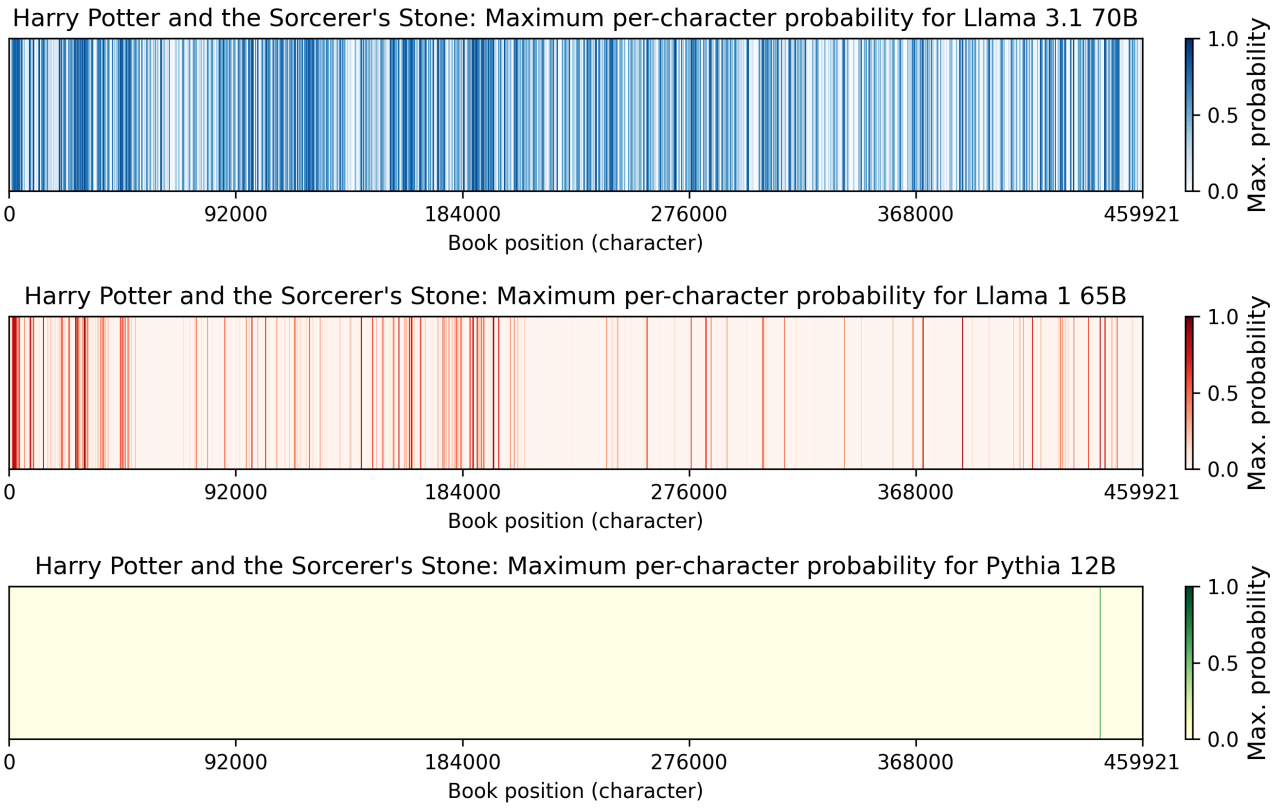

Figure 4: Heatmaps for LLAMA 3.1 70B (**top**), LLAMA 1 65B (**middle**), and PYTHIA 12B (**bottom**) for *Harry Potter* (Rowling, 1998), each showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we plot the *highest* extraction probability at the particular character location. The goal is to identify high-probability "hot-spots." LLAMA 3.1 70B memorized most of *Harry Potter*, reflected in the number of high-probability regions. PYTHIA 12B effectively memorizes none of the same book.

break up books text into examples when attempting extraction. By testing overlapping examples, we expect to surface high-probability *regions* of memorized content within a book, which we can then explore more precisely in follow-up experiments, discussed below. We *don't* expect this approach to result in high per-book extraction rates, as this strategy should result in many 0-probability sequences (Appendix G.1). This isn't our goal: we're trying to surface as much total memorization possible. We can then use these memorization 'hot-spots' to identify longer memorized sequences (Figure 5).

At each point on the heatmap, there exists (at least one example) at that location in the book that is extracted with the shown probability. For LLAMA 3.1 70B, there exists is an example that can be extracted with high probability *almost every location in the book. This doesn't mean (almost) the whole book can be extracted in one continuous segment at generation-time; this is a claim about how much of the book we estimate to be memorized in the model's parameters.* As an alternate view, we can visualize the same data as distributions over all of the examples' $p_{\boldsymbol{z}}$ (Figure 3, Appendix G.1).

Comparing books within a model, a given book across models. With these heatmaps, and the underlying information they contain about $p_{\boldsymbol{z}}$ for each example (plotted in Figure 3, right), we can make comparisons about the extent to which different models memorize different books, and how memorization for different books varies for a given model. We show the first type of comparison in Figure 4. (Appendix G provides numerous examples of the second across 36 books for 5 models.) LLAMA 3.1 70B clearly has memorized more of *Harry Potter* (and with higher probability) than the earlier generation, similarly-sized LLAMA 1 65B. Even though both are Meta models, they clearly exhibit important differences in how much they memorize of specific books (Figure 2, right)—not just how much they memorize across `Books3` more generally (Figure 2, left). PYTHIA 12B is a smaller model that was also trained on *Harry Potter* (Biderman et al., 2023; Gao et al., 2020); however, it effectively memorizes *none* of the book. Together, these three heatmaps underscore the importance of how training choices, not just the mere existence of examples within a training dataset, as having an impact on memoriza-

tion (Lee et al., 2023a; 2022; 2023b; 2024).

Last, aligning with the overall low average extraction rates we observe for models on Books3 (Section 3), we note that most books—regardless of the model we test—have heatmaps that resemble PYTHIA 12B's, above. This is true for most of the books we tested written by plaintiffs in *Kadrey et al. v. Meta, Inc.* (Kadrey et al. v. Meta Platforms, Inc.)—with important exceptions (e.g., Coates, 2017). Very popular books, like *Harry Potter* exhibit higher degrees of memorization for Llama models of all sizes. PYTHIA 12B and PHI 4, which was trained predominantly on synthetic data (i.e., is likely to not directly contain Books3), are smaller models; we expect them to memorize less than larger ones like LLAMA 3.1 70B. However, they also tend to memorize less than LLAMA 1 13B—a model in the same size class (Appendix G).

5. Takeaways for memorization and copyright

Our results complicate the traditional narrative both plaintiffs and defendants typically use in copyright cases in describing how LLMs work. The evidence supports the positions of plaintiffs in some respects and of defendants in other respects. More generally, we show that the extent of memorization in models varies with model size, the specific choice of model, the book tested, and even within individual books (Section 3 & 4). We see three primary implications of our results for copyright disputes.

(1) There is definitely some memorization of books in many models. And, in the case of some models, there is quite a lot of memorization of some books, though most books are not memorized at all (with respect to our specific extraction methodology)—either in whole or in part. Such evidence of memorization in the model matters for the dispute over whether the models themselves are derivative works, an argument that courts have thus far not been receptive to. A work is not a derivative work unless it's "substantially similar" in significant part to the original work (e.g., Litchfield v. Spielberg, 736 F.2d 1352, 1357 (9th Cir. 1984)). And models are not, as plaintiffs sometimes contend, mere copies of the works on which they trained. But our data bolsters plaintiffs' argument that at least some models may be derivative works of at least some books because the model has memorized a significant amount of protectable expression from the book. The law does not require that the entire work be included in the derivative; it is enough that the derivative incorporates a substantial amount of protectable expression.

That doesn't end the inquiry; the model, like the training dataset, may be protected under copyright's fair use doctrine (Lemley & Casey, 2021). But the fair use analysis of the model itself may be different than the analysis of the training dataset. A training-data example is an intermediate copy that isn't itself sold, and is only used internally in the course of producing outputs—many of which won't be substantially similar to any training data (Henderson et al., 2023; Lee et al., 2023b). A model, too, is an intermediate work in one sense; its use is to produce outputs, and those outputs are overwhelmingly not copyright infringements. But for companies that sell or release their models to others under open source licenses, the model itself is the product, and sometimes one that is being sold directly for commercial gain. That may make it harder to rely on the cases that justify training as fair use (e.g., Author's Guild v. Google, Inc., 804 F.3d 202 (2d Cir. 2015); Sega Enters. Ltd. v. Accolade, Inc., 977 F.2d 1510 (9th Cir. 1992), as amended (Jan. 6, 1993); Sony Computer Ent., Inc. v. Connectix Corp., 203 F.3d 596 (9th Cir. 2000)) (Section 3).

The consequences of a finding that the model itself was a copy of some copyrighted works, and that **distribution** of the model was thus a distribution of a copy of those works, could be dramatic for the AI industry. LLAMA 3.1 70B was downloaded 105,029 times last month on HuggingFace (Appendix, Figure 7)—far less from the height of its release in 2024. If we say conservatively that the model has been downloaded 1 million times since its release, then those 1 million downloads could be seen as 1 million, potentially infringing distributions of reproductions of *Harry Potter and the Sorcerer's Stone*—as well as any other books in copyright for which a more than *de minimis* amount of copyrighted expression has been memorized. With $150,000 (potentially) per infringing work, even if 3% of Books3 were to be found to be infringing copies in the model, that is a potential statutory damages award of nearly $1 billion dollars, and profit disgorgement awards could be substantially greater. If the model itself is a copy and is not protected by fair use, courts might also order that it not be distributed or even that copies of the model itself be destroyed (Copyright Law of the United States, 2010). It is not an exaggeration to say that the risk of damages across all memorized books and of orders of destruction are an existential threat to the companies training these models.

(2) The extent of memorization varies in several important ways. Our experiments show that the extent of memorization varies widely from model to model and, within a model, even from work to work in the Books3 dataset. This means it's hard to make any sort of **class**-wide (in the class-action-lawsuit sense) general assessment of whether a particular model copied a particular work and whether, for that model, infringing output based on memorization is even possible. Indeed, we show that memorization rates vary for many of the actual named authors in pending lawsuits, and even from book to book for the same plaintiffs (e.g., Ta-Nehisi Coates, see Appendix G). Because many of the pending cases are proceeding as class actions (e.g.,

Prefix (25 tokens): `\'t carry on.\n\n"That Quidditch Cup\'ll have our name on it this year," said Wood happily as they tr`

Suffix (275 tokens):

`udged back up to the castle. "I wouldn\'t be surprised if you turn out better than Charlie Weasley, and he could have played for England if he hadn\'t gone off chasing dragons."\n\nPerhaps it was because he was now so busy, what with Quidditch practice three evenings a week on top of all his homework, but Harry could hardly believe it when he realized that he\'d already been at Hogwarts two months. The castle felt more like home than Privet Drive ever had. His lessons, too, were becoming more and more interesting now that they had mastered the basics.\n\nOn Halloween morning they woke to the delicious smell of baking pumpkin wafting through the corridors. Even better, Professor Flitwick announced in Charms that he thought they were ready to start making objects fly, something they had all been dying to try since they\'d seen him make Neville\'s toad zoom around the classroom. Professor Flitwick put the class into pairs to practice. Harry\'s partner was Seamus Finnigan (which was a relief, because Neville had been trying to catch his eye). Ron, however, was to be working with Hermione Granger. It was hard to tell whether Ron or Hermione was angrier about this. She hadn\'t spoken to either of them since the day Harry\'s broomstick had arrived.\n\n"Now, don\'t forget that nice wrist movement'`

Llama 3.1 70B

T=1, k=40

p_z = 36.1%

Figure 5: One of many sequences from *Harry Potter* extracted from LLAMA 3.1 70B. This 300-token example has extraction probability of $p_z = 36.1\%$—in expectation),this suffix is generated verbatim more than $1/3$ of the time with this prompt. See Appendix H for more extracted sequences.

Kadrey et al. v. Meta Platforms, Inc.), plaintiffs will have to demonstrate that *all* book owners have sufficient common legal and fact issues, such that it makes sense for a court to certify the class and treat them all together (Fed. R. Civ. P. 23(a)(2)). That may be impossible because, as is clear from our results, basic questions about whether the model actually incorporated any significant expression from any individual plaintiff's book can't be generalized. Some plaintiffs may be able to show copying, but others won't. Courts generally deny class certification in such circumstances (e.g., Wal-Mart Stores, Inc. v. Dukes, 564 U.S. 338, 349–50 (2011)).

(3) There is no deterministic path from model memorization to outputs of infringing works. While we've used probabilistic extraction as proof of memorization, to actually extract a given piece of 50 tokens of copied text often takes hundreds or thousands of prompts. Using the adversarial extraction method of Hayes et al. (2025), we've proven that it can be done, and therefore that there is memorization in the model (Cooper & Grimmelmann, 2024; Carlini et al., 2021). But this is where, even though extraction is evidence of memorization, it may become important that they are not identical processes (Section 2). Memorization is a property of the model itself; extraction comes into play when someone *uses* the model (Cooper & Grimmelmann, 2024). This paper makes claims about the former, not the latter.

Nevertheless, it's worth mentioning that, in the average case, it's unlikely anyone in the real world would actually use the model in practice with this extraction method to deliberately produce infringing outputs, because doing so would require huge numbers of generations to get non-trivial amounts of text in practice (Appendix A). For the majority of cases we've observed—where models don't exhibit extreme degrees of memorization of examples (e.g., $p_z < 0.1$)—that makes output infringement much less of a real-world problem, even for models that exhibit a lot of memorization overall. On the other hand, in several cases (e.g., *1984*, *Harry Potter*), it doesn't take thousands or even hundreds of generations to get large amounts of text; it takes only a handful. Altogether, this also may affect who courts view as directly liable for any output infringement; someone (like some plaintiffs' lawyers) who runs hundreds or thousands of queries to try to get one that is infringing seems much more like the direct volitional actor misusing the model (CoStar Grp., Inc. v. LoopNet, Inc., 373 F.3d 544 (4th Cir. 2004)). We'll address this further in future work.

The fact that extraction is probabilistic and difficult to do also matters to the "is training infringement" issue—though it isn't determinative, because it means use of the training data to create **superseding copies** (i.e., a copy that can stand in for/ replace the work) is unlikely in practice. Nonetheless, the fact of memorization creates a point of distinction from *Author's Guild v. Google* (Author's Guild v. Google, Inc., 804 F.3d 202 (2d Cir. 2015)), which considered and rejected an argument that hackers could access the plaintiffs' works through Google's internal database (804 F.3d, at 227-28). Our results show that it is possible to recreate some of the content of some books by "hacking" the model itself. We think it unlikely as a practical matter that people will use the model in this way; there are easier and more effective ways to pirate a book. But at a minimum, it complicates the copyright fair use analysis.

6. Conclusion

Building on the recent work of Hayes et al. (2025)—a novel probabilistic extraction method—we show that the extent of verbatim memorization of books from the `Books3` dataset is more significant than previously described. We also show that memorization varies widely from model to model and from book to book within each model, as well as varying in different parts of individual books. Our results complicate current disputes over copyright infringement, both by rejecting easy claims made by both sides about how models work and by demonstrating that there is no single answer to the question of how much a model memorizes. There is much more work to do; we only ran book-specific, verbatim-extraction experiments for 56 books in `Books3`. In future work, we intend to explore non-verbatim extraction, as well as possible explanations for this variability and for the high level of memorization of some specific works by particular models.

Impact Statement

Our work has significant implications for current debates about memorization and copyright, with respect to LLMs. In particular, we add nuance to existing debates by re-envisioning how memorization could be measured to have relevance for copyright. Our work further clarifies points first raised in Lee et al. (2023b): it isn't possible to make sweeping statements about all of generative AI (or even all LLMs) and copyright. Models and systems differ significantly; the amount they memorize within the same book can vary significantly, as well as between books. Even if memorization is rare overall, the underlying amounts of memorization for specific books tell a very different story. All of this is relevant for copyright. Developing knowledge in this area can help us build and use models more responsibly.

References

804 F.3d.

Abdin, M., Aneja, J., Behl, H., Bubeck, S., Eldan, R., Gunasekar, S., Harrison, M., Hewett, R. J., Javaheripi, M., Kauffmann, P., Lee, J. R., Lee, Y. T., Li, Y., Liu, W., Mendes, C. C. T., Nguyen, A., Price, E., de Rosa, G., Saarikivi, O., Salim, A., Shah, S., Wang, X., Ward, R., Wu, Y., Yu, D., Zhang, C., and Zhang, Y. Phi-4 Technical Report, 2024. URL https://arxiv.org/abs/2412.08905.

amongglue/books3-subset-raw, 2025. URL https://huggingface.co/datasets/amongglue/books3-subset-raw.

Anthropic. Introducing 100K Context Windows, May 2023. URL https://www.anthropic.com/index/100k-context-windows/.

Author's Guild v. Google, Inc., 804 F.3d 202 (2d Cir. 2015).

Barolini, T. *Dante and the Origins of Italian Literary Culture*. Fordham University Press, 2006.

Biderman, S., Schoelkopf, H., Anthony, Q. G., Bradley, H., O'Brien, K., Hallahan, E., Khan, M. A., Purohit, S., Prashanth, U. S., Raff, E., et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pp. 2397–2430. PMLR, 2023.

Borel, E. La mecanique statique et l'irreversibilite. *J. Phys. Theor. Appl.*, 3(1):189–196, 1913. doi: 10.1051/jphystap:019130030018900. URL https://doi.org/10.1051/jphystap:019130030018900.

Bracha, O. The Work of Copyright in the Age of Machine Production, January 2024. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4581738.

Bret, D. *Elizabeth Taylor*. Greystone Books, 2011.

Brown, D. *The Da Vinci Code*. Doubleday, 2003.

Brown, G., Bun, M., Feldman, V., Smith, A., and Talwar, K. When is memorization of irrelevant training data necessary for high-accuracy learning? In *Proceedings of the 53rd annual ACM SIGACT symposium on theory of computing*, pp. 123–132, 2021.

Camus, A. *The Myth of Sisyphus*. Hamish Hamilton Publishers, Trinity Press, 1955. Justin O'Brien (translator).

CANBERT/pile_books3_text, 2025. URL https://huggingface.co/datasets/CANBERT/pile_books3_text.

Carlini, N. What my privacy papers (don't) have to say about copyright and generative AI, 2025. URL https://nicholas.carlini.com/writing/2025/privacy-copyright-and-generative-models.html.

Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pp. 2633–2650, 2021.

Carlini, N., Hayes, J., Nasr, M., Jagielski, M., Sehwag, V., Tramèr, F., Balle, B., Ippolito, D., and Wallace, E. Extracting Training Data from Diffusion Models, 2023a.

Carlini, N., Ippolito, D., Jagielski, M., Lee, K., Tramèr, F., and Zhang, C. Quantifying Memorization Across Neural Language Models. In *International Conference on Learning Representations*, 2023b.

Carroll, L. *Alice's Adventures in Wonderland*. Macmillan, 1865.

Charlesworth, J. Generative AI's Illusory Case for Fair Use. *Vanderbilt Journal of Entertainment and Technology Law*, 27, 2025.

Chat GPT Is Eating the World, 2024. URL https://chatgptiseatingtheworld.com.

Coates, T.-N. *The Beautiful Struggle*. One World, 2009.

Coates, T.-N. *We Were Eight Years in Power: An American Tragedy*. One World, 2017.

Coates, T.-N. *The Water Dancer: A Novel*. One World, 2019.

Cogburn, J. *Dungeons and Dragons and Philosophy*. Open Court Publishing, 2012.

Concord Music Group, Inc. v. Anthropic PBC. 3:23-cv-01092 (M.D. Tenn.).

Cooper, A. F. and Grimmelmann, J. The Files are in the Computer: Copyright, Memorization, and Generative AI. *arXiv preprint arXiv:2404.12590*, 2024.

Cooper, A. F., Lee, K., Grimmelmann, J., Ippolito, D., Callison-Burch, C., Choquette-Choo, C. A., Mireshghallah, N., Brundage, M., Mimno, D., Choksi, M. Z., Balkin, J. M., Carlini, N., Sa, C. D., Frankle, J., Ganguli, D., Gipson, B., Guadamuz, A., Harris, S. L., Jacobs, A. Z., Joh, E., Kamath, G., Lemley, M., Matthews, C., McLeavey, C., McSherry, C., Nasr, M., Ohm, P., Roberts, A., Rubin, T., Samuelson, P., Schubert, L., Vaccaro, K., Villa, L., Wu, F., and Zeide, E. Report of the 1st Workshop on Generative AI and Law. *arXiv preprint arXiv:2311.06477*, 2023.

Cooper, A. F., Choquette-Choo, C. A., Bogen, M., Jagielski, M., Filippova, K., Liu, K. Z., Chouldechova, A., Hayes, J., Huang, Y., Mireshghallah, N., Shumailov, I., Triantafillou, E., Kairouz, P., Mitchell, N., Liang, P., Ho, D. E., Choi, Y., Koyejo, S., Delgado, F., Grimmelmann, J., Shmatikov, V., Sa, C. D., Barocas, S., Cyphert, A., Lemley, M., danah boyd, Vaughan, J. W., Brundage, M., Bau, D., Neel, S., Jacobs, A. Z., Terzis, A., Wallach, H., Papernot, N., and Lee, K. Machine Unlearning Doesn't Do What You Think: Lessons for Generative AI Policy, Research, and Practice. *arXiv preprint arXiv:2412.06966*, 2024.

Copyright Law of the United States. 17 U.S. Code § 503 - Remedies for infringement: Impounding and disposition of infringing articles, December 2010. URL https://www.law.cornell.edu/uscode/text/17/503.

CoStar Grp., Inc. v. LoopNet, Inc., 373 F.3d 544 (4th Cir. 2004).

Cyphert, A. B. Generative AI, Plagiarism, and Copyright Infringement in Legal Documents. *Minnesota Journal of Law, Science & Technology*, 25, 2024.

DeepSeek-AI et al. DeepSeek LLM: Scaling Open-Source Language Models with Longtermism, 2024. URL https://arxiv.org/abs/2401.02954.

Doctorow, C. *Down and Out in the Magic Kingdom*. Tor Books, 2003.

Díaz, J. *Drown*. Riverhead Books, 1996.

Díaz, J. *The Brief Wondrous Life of Oscar Wao*. Riverhead Books, 2007.

EleutherAI/the_pile_deduplicated, 2025. URL https://huggingface.co/datasets/EleutherAI/the_pile_deduplicated.

Farnsworth, C. *Blood Oath*. Penguin Publishing Group, 2010.

Farnsworth, C. *The President's Vampire*. G.P. Putnam's Sons, 2011.

Fed. R. Civ. P. 23(a)(2). (requiring questions of law or fact common to the class before it can be certified).

Feldman, V. Does learning require memorization? a short tale about a long tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, pp. 954–959, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369794.

Finder, J. *Killer Instinct*. St. Martin's Press, 2006.

Fitzgerald, F. S. *The Great Gatsby*. Charles Scribner's Sons, 1925.

Gao, L., Biderman, S., Black, S., Golding, L., Hoppe, T., Foster, C., Phang, J., He, H., Thite, A., Nabeshima, N., et al. The Pile: An 800GB Dataset of Diverse Text for Language Modeling. *arXiv preprint arXiv:2101.00027*, 2020.

Gemini Team, Savinov, N., Teplyashin, D., Lepikhin, D., Lillicrap, T., Alayrac, J.-b., Soricut, R., Lazaridou, A., Firat, O., Schrittwieser, J., et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.

Golden, C. *Poison Ink*. Delacorte Press, 2008.

Golden, C. *Dead Ringer*. St. Martin's Press, 2016.

Golden, C. *Ararat*. St. Martin's Press, 2017.

Golden, C. *Blood of the Four*. Harper Voyager, 2018.

Goodyear, M. P. Artificial Infringement. In *Proceedings of the 2025 Symposium on Computer Science and Law*, CSLAW '25, pp. 26–38, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400714214. doi: 10.1145/3709025.3712208. URL https://doi.org/10.1145/3709025.3712208.

Grattafiori, A. et al. The Llama 3 Herd of Models, 2024. URL https://arxiv.org/abs/2407.21783.

Greer, A. S. *The Path of Minor Planets*. Picador, 2002.

Greer, A. S. *The Confessions of Max Tivoli*. Picador, 2005.

Hayes, J., Swanberg, M., Chaudhari, H., Yona, I., Shumailov, I., Nasr, M., Choquette-Choo, C. A., Lee, K., and Cooper, A. F. Measuring memorization in language models via probabilistic extraction. In Chiruzzo, L., Ritter, A., and Wang, L. (eds.), *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 9266–9291, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-189-6. URL https://aclanthology.org/2025.naacl-long.469/.

Heller, J. *Catch-22*. Simon & Schuster, 1961.

Henderson, P., Li, X., Jurafsky, D., Hashimoto, T., Lemley, M. A., and Liang, P. Foundation Models and Fair Use, 2023. URL https://arxiv.org/abs/2303.15715.

Hwang, D. H. *M. Butterfly*. Plume, 1988.

Hwang, D. H. *Golden Child*. Theatre Communications Group, 1998.

Hwang, D. H. *Trying to Find Chinatown*. Theatre Communications Group, 1999.

Jacobs, B. E. *All the Onions*. Storey Publishing, LLC, 1999.

Joyce, J. *Ulysses*. Paris-Shakespeare, 1922.

Kadrey, R. *Sandman Slim: A Novel*. Harper Collins, 2009.

Kadrey, R. *Aloha from Hell*. Harper Voyager, 2011.

Kadrey, R. *Kill the Dead*. Harper Collins, 2014.

Kadrey, R. *The Wrong Dead Guy*. Harper Voyager, 2017.

Kadrey et al. v. Meta Platforms, Inc. URL https://www.courthousenews.com/wp-content/uploads/2025/02/kadrey-vs-meta-third-amended-complaint.pdf?utm_source=chatgpt.com. No. 3:23-cv-03417-VC (Third Amended Consolidated Complaint).

Kadrey v. Meta Platforms, Inc. No. 3:23-cv-03417 (N.D. Cal. July 7, 2023).

Klam, M. *Sam the Cat*. Vintage, 2001.

Klam, M. *Who Is Rich?* Random House, 2017.

Lee, K., Ippolito, D., Nystrom, A., Zhang, C., Eck, D., Callison-Burch, C., and Carlini, N. Deduplicating Training Data Makes Language Models Better. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, volume 1, pp. 8424–8445, 2022.

Lee, K., Cooper, A. F., Grimmelmann, J., and Ippolito, D. AI and Law: The Next Generation. *SSRN*, 2023a. http://dx.doi.org/10.2139/ssrn.4580739.

Lee, K., Cooper, A. F., and Grimmelmann, J. Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain. *arXiv preprint arXiv:2309.08133*, 2023b.

Lee, K., Cooper, A. F., and Grimmelmann, J. Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain (The Short Version). In *Proceedings of the Symposium on Computer Science and Law*, CSLAW '24, pp. 48–63, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400703331. doi: 10.1145/3614407.3643696. URL https://doi.org/10.1145/3614407.3643696.

Lemley, M. and Casey, B. Fair Learning. *Texas Law Review*, 99:743, 2021.

Lippman, L. *After I'm Gone: A Novel*. Harper Collins, 2014.

Lippman, L. *Sunburn*. William Morrow, 2018.

Litchfield v. Spielberg, 736 F.2d 1352, 1357 (9th Cir. 1984). ("a work is not derivative unless it has been substantially copied from the prior work.").

Ma, I., Domingo, I., Krone-Martins, A., Baldi, P., and Lopes, C. V. Memorization: A Close Look at Books, 2025. URL https://arxiv.org/abs/2504.12549.

Martin, G. R. *A Game of Thrones*. Voyager Books, 1996.

Meyer, S. *Twilight*. Little, Brown and Company, 2005.

Nasr, M., Carlini, N., Hayase, J., Jagielski, M., Cooper, A. F., Ippolito, D., Choquette-Choo, C. A., Wallace, E., Tramèr, F., and Lee, K. Scalable Extraction of Training Data from (Production) Language Models. *arXiv preprint arXiv:2311.17035*, 2023.

Nasr, M., Rando, J., Carlini, N., Hayase, J., Jagielski, M., Cooper, A. F., Ippolito, D., Choquette-Choo, C. A., Tramèr, F., and Lee, K. Scalable Extraction of Training Data from Aligned, Production Language Models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=vjel3nWP2a.

OpenAI. OpenAI and journalism, January 2024. URL https://openai.com/blog/openai-and-journalism.

Orwell, G. *Nineteen-Eighty Four*. Harcourt, Brace and Company, 1949.

Prashanth, U. S., Deng, A., O'Brien, K., au2, J. S. V., Khan, M. A., Borkar, J., Choquette-Choo, C. A., Fuehne, J. R., Biderman, S., Ke, T., Lee, K., and Saphra, N. Recite, Reconstruct, Recollect: Memorization in LMs as a Multifaceted Phenomenon, 2024. URL https://arxiv.org/abs/2406.17746.

Reisner, A. These 183,000 Books Are Fueling the Biggest Fight in Publishing and Tech. *The Atlantic*, September 2023. URL https://www.theatlantic.com/technology/archive/2023/09/books3-database-generative-ai-training-copyright-infringement/675363/.

Rouighi, R. *The Making of a Mediterranean Emirate*. University of Pennsylvania Press, 2011.

Rowling, J. *Harry Potter and the Sorcerer's Stone*. Scholastic, 1998.

Sag, M. Copyright Safety for Generative AI. *Houston Law Review*, 61, 2024a.

Sag, M. Fairness and Fair Use in Generative AI. *Fordham Law Review*, 92, 2024b.

Samuelson, P. Generative AI meets copyright. *Science*, 381 (6654):158–161, 2023.

Sandberg, S. *Lean In*. Alfred A. Knopf, 2013.

SaylorTwift/the_pile_books3_minus_gutenberg, 2025. URL https://huggingface.co/datasets/SaylorTwift/the_pile_books3_minus_gutenberg.

Schwarzschild, A., Feng, Z., Maini, P., Lipton, Z. C., and Kolter, J. Z. Rethinking llm memorization through the lens of adversarial compression. *arXiv preprint arXiv:2404.15146*, 2024.

Sega Enters. Ltd. v. Accolade, Inc., 977 F.2d 1510 (9th Cir. 1992), as amended (Jan. 6, 1993).

Shelley, M. *Frankenstein*. Lackington, Hughes, Harding, Mavor, & Jones, 1818.

Silverman, S. *The Bedwetter*. Harper Collins, 2010.

Snyder, R. L. *No Visible Bruises*. Bloomsbury Publishing, 2019.

Sobel, B. Elements of Style: Copyright, Similarity, and Generative AI. *Harvard Journal of Law & Technology*, 38, 2024.

Sony Computer Ent., Inc. v. Connectix Corp., 203 F.3d 596 (9th Cir. 2000).

Stephenson, N. *The Diamond Age*. Bantam Spectra, 1996.

Team, G. et al. Gemma 2: Improving Open Language Models at a Practical Size, 2024. URL https://arxiv.org/abs/2408.00118.

TerKeurst, L. *Unglued: Making Wise Choices in the Midst of Raw Emotions*. Zondervan, 2012.

TerKeurst, L. *Embraced*. Thomas Nelson, 2018.

The Simpsons: Last Exit to Springfield (Fox television broadcast Mar. 11, 1993), 1993. ("This is a thousand monkeys working at a thousand typewriters. Soon they'll have written the greatest novel known to man . . . [l]et's see. 'It was the best of times, it was the blurst of times?' You stupid monkey!").

Tolkien, J. *The Hobbit*. George Allen and Unwin, 1937.

Tolkien, J. *Tree and Leaf*. George Allen & Unwin, 1964.

Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., and Lample, G. LLaMA: Open and Efficient Foundation Language Models, 2023a. URL https://arxiv.org/abs/2302.13971.

Touvron, H. et al. Llama 2: Open Foundation and Fine-Tuned Chat Models, 2023b. URL https://arxiv.org/abs/2307.09288.

United States Copyright Office. Copyright and Artificial Intelligence (Part 2: Copyrightability), 2025. URL https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf.

Wal-Mart Stores, Inc. v. Dukes, 564 U.S. 338, 349–50 (2011). (holding that class certification was inappropriate because commonality requires that class members be capable of showing that they have suffered the same injury).

Wharton, K. *Another Bloody Chapter in an Endless Civil War*. Helion and Company, 2016.

Wilf-Townsend, D. The Deletion Remedy. *North Carolina Law Review*, 103, 2024. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4933011. Forthcoming 2025.

Woodson, J. *After Tupac and D Foster*. Nancy Paulsen Books, 2008.

Woodson, J. *Brown Girl Dreaming*. Penguin Group, 2014.

Woodson, J. *Another Brooklyn*. Amistad Press, 2016.

Woolf, V. *Orlando*. Hogarth Press, 1928.

Zhang, C., Ippolito, D., Lee, K., Jagielski, M., Tramèr, F., and Carlini, N. Counterfactual memorization in neural language models. *Advances in Neural Information Processing Systems*, 36:39321–39362, 2023.

Zittrain, J. *The Future of the Internet and How to Stop It*. Yale University Press, 2008.

Appendix Table of Contents

A Additional notes on the introduction . . . 17

B Additional notes on memorization and extraction . . . 19

B.1 Metrics . . . 19

B.2 Computing sequence probabilities in one forward pass . . . 19

C Additional notes on extraction risk, probabilistic copies, and U.S. copyright . . . 23

D Testing our measurement pipeline . . . 25

E Details on experimental settings . . . 27

F Additional information on extraction rates . . . 30

G Additional results for per-book extraction experiments . . . 33

G.1 Sliding-window experiments . . . 33

G.1.1 *After I'm Gone*, Lippman . . . 37

G.1.2 *Alice's Adventures in Wonderland*, Carroll . . . 38

G.1.3 *Another Brooklyn*, Woodson . . . 40

G.1.4 *Ararat*, Golden . . . 42

G.1.5 *Brown Girl Dreaming*, Woodson . . . 42

G.1.6 *Dante and the Origins of Italian Culture*, Barolini . . . 44

G.1.7 *Dead Ringer*, Golden . . . 45

G.1.8 *Down and Out in the Magic Kingdom*, Doctorow . . . 46

G.1.9 *Drown*, Díaz . . . 48

G.1.10 *Dungeons and Dragons and Philosophy*, Cogburn . . . 49

G.1.11 *Embraced*, TerKeurst . . . 51

G.1.12 *Harry Potter and the Sorcerer's Stone*, Rowling . . . 53

G.1.13 *The Future of the Internet and How to Stop It*, Zittrain . . . 55

G.1.14 *M. Butterfly*, Hwang . . . 57

G.1.15 *No Visible Bruises*, Snyder . . . 59

G.1.16 *Sandman Slim*, Kadrey . . . 60

G.1.17 *Sunburn*, Lippman . . . 61

G.1.18 *The Beautiful Struggle*, Coates . . . 62

G.1.19 *The Bedwetter*, Silverman . . . 63

G.1.20 *The Brief Wondrous Life of Oscar Wao*, Díaz . . . 65

G.1.21 *The Confessions of Max Tivoli*, Greer . . . 67

G.1.22 *The Great Gatsby*, Fitzgerald . . . 68

G.1.23 *The Hobbit*, Tolkien . . . 70

G.1.24 *The Making of a Mediterranean Emirate*, Rouighi . . . 72

G.1.25 *The Myth of Sisyphus*, Camus . . . 73

G.1.26 *The President's Vampire*, Farnsworth . . . 75

G.1.27 *The Water Dancer*, Coates . . . 76

G.1.28 *Unglued*, TerKeurst . . . 77

G.1.29 *We Were Eight Years in Power*, Coates . . . 79

G.1.30 *Who Is Rich*, Klam . . . 81

G.1.31 *A Game of Thrones*, Martin . . . 82

G.1.32 *Lean In*, Sandberg . . . 84

G.1.33 *Nineteen-Eighty Four*, Orwell . . . 86

G.1.34 *The Da Vinci Code*, Brown . . . 88

G.1.35 *Twilight*, Meyer . . . 90

G.1.36 *Ulysses*, Joyce 92
G.2 Random subsampling experiments 94
G.3 Additional results on *Harry Potter* 97
H Selection of extracted examples 100
I Discussion of extended results 115

A. Additional notes on the introduction

The headline figure we include after the abstract is meant to be illustrative and catchy—an example, not a general statement about memorization and extraction. We repeat this example Figure 6a with a caption. The example is drawn from *The Great Gatsby*—one of the most famous quotes from the book. It is also a very short example (with the Llama 1 tokenizer, it is 57 tokens in total). We do not focus on such short suffixes in our analysis, for reasons that are discussed both in the main text and in our experimental details below. (In this case we also clip the prefix length to be quite short.) In the introduction, we limit ourselves to something short so that it does not take up too much space.

It is also the case the very famous quotes like this are duplicated in many locations in training-data sources (e.g., blogs, essays, personal web pages, etc.) We are not claiming (here, or in general) that this sequence can be extracted *because* a given book (here, *The Great Gatsby*) was included in `Books3`; it could be extracted (or more extractable) because it is duplicated in many places in the training data. Of course, one of those copies is in `Books3`, if the entirety of `Books3` is included in the model's training data. (And we have reason to believe that it is for the Llama 1 family (Touvron et al., 2023a): "We include two book corpora in our training dataset: the Gutenberg Project, which contains books that are in the public domain, and the Books3 section of The Pile (Gao et al., 2020)(Gao et al., 2020), a publicly available dataset for training large language models," (p. 2).)

We also include a second example in Figure 6b, which is on-par with the length of examples that we extract in experiments (100 tokens in our sliding window and average-extraction-rate experiments). This example was chosen at random from sampling from *The Great Gatsby*, and ensuring that the example began at the beginning of a sentence. (In this case, we sampled a 100-token sequence, and then backed up to the beginning of the sequence, which is what accounts for the additional 7 tokens; the suffix start index was also sampled at random, from indexes 50 through 60 of the full example.) Note, for LLAMA 1 13B (the model we used for these examples), the probability of extracting the second example is much higher than for the famous quote we include after the abstract.

In this work, these are the only two examples for which we actually produce generations. For every other experiment, we compute sequence probabilities directly from the logits we obtain from running sequences through the LLM. **These two approaches are functionally equivalent: the computed sequence probabilities correspond to the statistically expected frequencies of verbatim extracted outputs we would observe if we were to generate a large number of sequences for the same prompt.** Hayes et al. (2025), the authors of this extraction approach, confirm this with ample experimental evidence.

In future work, we intend to generate (not just process logits to compute probabilities), as this is required for computing *non*-verbatim extraction metrics. We limit ourselves to verbatim extraction in this paper; it is significantly computationally cheaper to do so, as operating on logits (as opposed to producing multiple generations per prompt) requires only one forward pass through the model. (See Section 2 and Appendix B for more details.) As a result, our work here (in several respects) only scratches the surface: there are various different ways to instantiate this extraction methodology in practice, we only explore a limited set of open-weight models, and we only test low-level extraction attacks on 36 of the nearly $200,000$ books in `Books3`.

Prompt (prefix)	Target (suffix)	Generations
They were careless people, Tom and Daisy — they smashed up things and creatures and then retreated	back into their money or their vast carelessness, or whatever it was that kept them together, and let other people clean up the mess they had made.	back into their money or their vast carelessness, or whatever it was that kept them together, and let other people clean up the mess they had made .
They were careless people, Tom and Daisy — they smashed up things and creatures and then retreated	back into their money or their vast carelessness, or whatever it was that kept them together, and let other people clean up the mess they had made.	back into their money or their vast carelessness, or whatever it was that kept them together, and let other people clean up the mess they had made.
They were careless people, Tom and Daisy — they smashed up things and creatures and then retreated	back into their money or their vast carelessness, or whatever it was that kept them together, and let other people clean up the mess they had made.	into their money or extraness or whatever it was and let other people clean up the mess they had made. F. Scott Fitzgerald 1

(a) 14.3% probability of extracting the verbatim target suffix from LLAMA 1 13B using top-k sampling (prefix: 25 tokens; suffix: 32 tokens; sampling configuration: $k = 40, T = 1$).

Prompt (prefix)	Target (suffix)	Generations
Conduct may be founded on the hard rock or the wet marshes, but after a certain point I don't care what it's founded on. When I came back from the East last autumn I felt that I wanted the world	to be in uniform and at a sort of moral attention forever; I wanted no more riotous excursions with privileged glimpses into the human heart. Only Gatsby, the man who gives his name to this book, was exempt from my reaction	to be in uniform and at a sort of moral attention forever; I wanted no more riotous excursions, no going back to the nursery. There was a tradition that intermarried with this feeling. … It occurred to me that if a man were to
Conduct may be founded on the hard rock or the wet marshes, but after a certain point I don't care what it's founded on. When I came back from the East last autumn I felt that I wanted the world	to be in uniform and at a sort of moral attention forever; I wanted no more riotous excursions with privileged glimpses into the human heart. Only Gatsby, the man who gives his name to this book, was exempt from my reaction	to be in uniform and at a sort of moral attentio forever; I wanted no more riotous excursions wit privileged glimpses into the human heart. Only Gatsby, the man who gives his name to this book, was exempt from my reaction
Conduct may be founded on the hard rock or the wet marshes, but after a certain point I don't care what it's founded on. When I came back from the East last autumn I felt that I wanted the world	to be in uniform and at a sort of moral attention forever; I wanted no more riotous excursions with privileged glimpses into the human heart. Only Gatsby, the man who gives his name to this book, was exempt from my reaction	to be in uniform and at a sort of moral attention forever; I wanted no more riotous excursions, or plunges, or brutalities. This has been one of my prized days. I have been searching for images, as usual, and I

(b) 57.97% probability of extracting the verbatim target suffix from LLAMA 13B using top-k sampling (prefix: 51 tokens; suffix: 56 tokens; sampling configuration: $k = 40, T = 1$).

Figure 6: Two examples of issuing a given prompt (left column) from *The Great Gatsby* to LLAMA 1 13B three times. In each figure, the middle row shows an instance of verbatim extraction; the top and bottom rows show instances where the generations do not match the target suffix.

B. Additional notes on memorization and extraction

We provide more details on the key metric in this paper: (n, p)-discoverable extraction (Appendix B.1), as well as our implementation (roughly speaking) for this implementation (Appendix B.2). Please also refer to Section 2 of the main paper text.

B.1. Metrics

Hayes et al. (2025) provide a more formal definition for the metric we compute in this paper:

Definition 1 ((n, p)**-discoverable extraction**, from Hayes et al. (2025))**.** Given a training example $\boldsymbol{z}$ that is split into an a-length prefix $\boldsymbol{z}_{1:a}$ and a k-length suffix $\boldsymbol{z}_{a+1:a+k}$, $\boldsymbol{z}$ is (n, p)-*discoverably extractable* if

$$\Pr\Big(\cup_{w \in [n]} \left(g_\phi \circ f_\theta\right)_w^k (\boldsymbol{z}_{1:a}) = \boldsymbol{z}_{1:a+k} \Big) \geq p,$$

where $(g_\phi \circ f_\theta)_w^k(\boldsymbol{z}_{1:a})$ represents the w-th (of n) independent execution of the autoregressive process of generating a distribution over the token vocabulary, sampling a token from this distribution, and adding the token to the sequence $k > 0$ times, starting from the same initial sequence $\boldsymbol{z}_{1:a}$.

Note that this definition can be easily adapted to non-verbatim extraction by testing if a generation is within ϵ distance from the target suffix for a given distance metric, rather than testsing for equality with $\boldsymbol{z}_{1:a+k}$. Also note that, in practice for verbatim extraction, computing this metric involves calculating sequence probabilities $p_{\boldsymbol{z}}$ as follows (rather than issuing n independent prompts to the model):

$$1 - (1 - p_{\boldsymbol{z}})^n \geq p \quad \Rightarrow \quad n \geq \frac{\log(1-p)}{\log(1-p_{\boldsymbol{z}})}, \quad \text{where } \; p_{\boldsymbol{z}} = \exp\Big(\sum_{t=a+1}^{a+k} \log p(z_t \mid \boldsymbol{z}_{1:t-1}) \Big) \tag{3}$$

We compute $p_{\boldsymbol{z}}$ as the $\exp$ of the sum $\log$ probabilities of the conditional token sequences, as this is more numerically stable that multiplying together the probabilities of the conditional token sequences. (See Section 2.)

In our work, we re-frame this definition to emphasize the quantity $p_{\boldsymbol{z}}$ that we actually compute: the probability of extracting a sequence $\boldsymbol{z}$ (for the given LLM, hyperparameter-configured decoding scheme, and suffix start index location). This is useful also for plotting distributions over $p_{\boldsymbol{z}}$ for a given book—to see how these probabilities vary for examples across a book, or for a given example across different models. We provide some examples of these comparisons in Section 4.

We also find that our version of explaining the metric is more accessible to broader audiences, which we intend to reach with this work. Hayes et al. (2025) specifically addresses ML and NLP audiences. We hope that our work supplements theirs in showing just how useful their metric is for capturing a probabilistic notion of extraction, and the underlying probabilistic copies that this reflects are memorized within model parameters.

B.2. Computing sequence probabilities in one forward pass

Given that the core algorithm for (n, p)-discoverable extraction is so elegantly concise, we include a (simplified) version of the code in this appendix. Our implementation differs from Hayes et al. (2025), which we confirm in discussion with the authors of that work. In that work, the authors generated the suffix a token at a time and summed up the per-generated-token $\log$ probabilities. Here, we observe that this is not necessary. We can get the logits for each token in the sequence with just one forward pass through the model, and effectively only one line of code. (See Listing 1, line 28.) For a 50-token suffix, this results in up to a $50\times$ decrease in compute—an enormous savings (it is of course less with KV caching enabled during generation, but still nonzero). This also makes computing this metric cheaper than traditional greedy-sampled discoverable extraction (Carlini et al., 2021), which generates the greedy-sampled sequence (i.e., for a 50-token suffix, there are 50 forward passes through the model, even if those passes involve caching).

Note that this is the *entire* interaction that our experiments have with the underlying model: for a given example, we perform inference on the GPU, get the logits tensor, and move the logits to the CPU for processing. That is it. We are not making *any* changes to the model in our extraction measurements; we are just getting the raw logits from running a given training-data example through the model. (As shown in Listing 1, we do this for a batch of examples at a time, yielding a batch of logits.)

The logits tensors we get are of shape `[batch_size, sequence_length, vocabulary_size]`. This is a three-dimensional tensor:

- Dimension 1: We have a slot for each example in the batch (with `batch_size` examples overall).
- Dimension 2: We have a position for each index in the token sequence of the example (with `sequence_length`$= 100$ being the number of total positions in most of our experiments).
- Dimension 3: We have a position for each token in the whole token vocabulary for this model (with `vocabulary_size` being the overall number of positions, e.g., this is $32,000$ for the Llama 1 family).

So, at position `[i, j, k]` in this tensor, we have the logit (unnormalized probability) value for the i-th example in the batch, at the `j-th` position in the example's token sequence, for the `k-th` token in the whole token vocabulary. That is, for a given token sequence in the batch (at `[i,j]`), we have the logit values for all of the next-possible tokens in the whole token vocabulary (a `vocabulary_size`-length list of numbers, one for each token in the whole vocabulary); these values reflect the probability distribution over the next token in the sequence.

```
import torch
"""
Runs inference to get logits for examples z, which we will use
to compute (n,p)-discoverable extraction (i.e., to compute p_z).
  - dataloader: a torch.utils.DataLoader wrapping a torch.utils.data.Dataset containing
    the tokenized examples z

  - model: loaded transformers.AutoModelForCausalLM (in this snippet, assumed to have
    layers distributed across GPUs with device_map="auto", though for smaller models we
    use data parallelism)

  - returns all_logits (logits indexed by batch number). In practice, we don't do this. We
     asynchronously do computations on logits on CPU to save log probs.
"""
def compute_logits_for_batch(dataloader, model):

    all_logits = {}
    with torch.inference_mode():
      for batch_idx, batch in enumerate(dataloader):
        input_ids = batch["input_ids"]
        attention_mask = batch["attention_mask"]
        example_metadata_batch = batch["metadata"]

        device = model.device

        inputs = {
            "input_ids": input_ids.to(device),
            "attention_mask": attention_mask.to(device)
        }

        # Run inference; shape: [batch_size, seq_len, vocab_size]
        logits_batch = model(**inputs).logits.detach().cpu()

        input_ids = input_ids.detach().cpu()
        attention_mask = attention_mask.detach().cpu()

        all_logits[batch_idx] = logits_batch
  return all_logits
```

Listing 1: Computing logits with one forward pass through the model. This is the key speedup in our implementation and this is the only code that interacts with a given model that we study.

We then implement our sequence probability computations (i.e., compute $p_{\boldsymbol{z}}$) as a post-processing operation on the logits we get from the model. (See Listing 2.) That is, for any given decoding scheme, we can transform the logit distribution—for example, change the shape of the distribution according to temperature, truncate the distribution to the top-k tokens and re-normalize the distribution to have its probability still sum to 1, etc.

This computation is written out on the right side of Equation (3). The key point here is that, for a given sequence that we are evaluating, all we do (starting at the first index in the suffix) is sum up the $\log$ probability of the *actual* next token in the sequence, one at a time. That is, from the logits, we get the $\log$ probabilities (Listing 2, line 19). We work with the $\log$ probabilities (adding them together), rather than the probabilities (and multiplying them together), as this is more numerically stable to compute. However, note that these computations are equivalent. For a given token in what we are considering as our suffix, all we are doing here is *manually* finding the ($\log$) probability associated with the *actual* next token in the sequence (Listing 2, lines 32-40). For example, consider the sequence `"the orange cat"` (where, for simplicity, we will assume each word is a single token). When processing `"the"`, the logits reflect the probabilities of the next token. We find the logit value in the distribution that corresponds to `"orange"` and add the associated $\log$ probability, and then continue on (doing the same for `"orange"`, where we then get the $\log$ probability for the next token being `"cat"`).

In essence, all we are doing is getting the probability of the whole sequence (by adding together the $\log$ probabilities) of the actual sequence; this is $p_{\boldsymbol{z}}$, the probability of the model (under the given decoding scheme we are using) generating that sequence conditioned on the prior context (the prefix). This yields information about the relationships the model has learned from being trained on natural language—patterns in that language.

Any given sequence should be low probability. That is, for a suffix length of 50 tokens, this computation is effectively multiplying together 50 per-token probabilities (by adding together $\log$ probabilities). This number *should* be really small. For a sequence where each token has really high probability, this can still mean the overall sequence has low probability. Consider that each token in the suffix has probability 0.9. This is *really* high; it means that a single token in the logit distribution (at each position, for the probability of the next token) has 90% of the probability mass; for a token vocabulary of size $32,000$, that means the remaining $31,999$ tokens all share the remaining 10%! For this sequence, the probability of generating it is $0.9^{50} \approx 0.005$—i.e., a 0.5% chance of generation. This is a relatively small number in the scheme of things; but it is also a really large number in general, if we consider the details of the multiplication we just did: all other possible 50-token sequences in this example are much less likely. In some cases in this paper, we observe sequences that have ***over 90% probability***—an enormous number, especially considering what the underlying computation is!

This is what memorization is: unusually high-probability sequences (where it is reasonable to consider 0.5%, or even smaller, to be unusually high). As Cooper & Grimmelmann (2024) say, "the pattern [the model has learned from the training data] *is* the memorized training data."

Also note that, since top-k sampling truncates the logit distribution to *only* the top-k tokens (e.g., for a vocabulary of $32,000$ tokens, there are $32,000$ logits, but top-k with $k = 40$ decoding will only consider the 40 highest-valued logits). This means that, if a given sequence has a token at a given position that is *not* in the top-k of the logit distribution, then we will not be able to complete the probability computation. This is what the `has_impossible_token` check iss referring to in Listing 2, lines 36-38. In this case, the given sequence has 0 probability; it is not extractable with *any* probability.

Note that, in practice, our code deviates from Listing 2 for efficiency reasons. (The core algorithm is nevertheless the same as what we show here.) As we process $\log$ probabilities, we save a lot of metadata in order to produce the plots that we include in this paper. For additional speedups (which, ultimately, total around $200\times$ prior work on (n, p)-discoverable extraction), as we compute logits on the GPU (Listing 1), we use parallel coding strategies to delegate our post-processing code to happen in parallel on CPU.

```python
import torch, torch.nn.functional as F, math
"""
For a batch of logits, computes the log probability (and probability)
of the suffix given the prefix for a batch of examples.
  -inputs: a batch of examples
  -logits: a batch of logits (computed doing inference, above)
  -suffix_start_idx: index in examples where suffix begins
  -transforms base distribution for temperature (temp) and top_k (k).
  -returns list of logprobs (and probs)
"""
def compute_logprob_batch(inputs, logits, suffix_start_idx, temp, k):
  input_ids = inputs_batch["input_ids"]
  attention_mask = inputs_batch["attention_mask"]
  if temp != 1.0 or k is not None:
    logits = apply_temperature_and_topk(logits, temp, k)

  # across the vocabulary, all possible next-token predictions at each
  # sequence position. shape: [batch_size, seq_len, vocab_size]
  log_probs = F.log_softmax(logits, dim=-1)
  # log_probs[b, i, j]: for example at batch index b, log probs of
  # predicting token j at position i in the sequence
  log_probs_list = [], prob_list = [], batch_size = input_ids.shape[0]
  for b in range(batch_size):
    input_ids_b = input_ids[batch_b]
    # shape: [seq_len, vocab_size]
    log_probs_b = log_probs[b]
    real_seq_len = attention_mask[b].sum().item()
    total_log_prob = 0.0, has_impossible_token = False
    # Compute log prob for each token in suffix conditioned on prefix
    for i in range(suffix_start_idx - 1, real_seq_len - 1):
      # get actual next token in suffix
      actual_next_token_id = input_ids_b[i + 1]
      # log prob for actual_next_token i+1 based on context up to i
      log_prob = log_probs_b[i, actual_next_token_id].item()
      # possible -inf from top-k filtering
      if not math.isfinite(log_prob):
        has_impossible_token = True
        break

      total_log_prob += log_prob

    if has_impossible_token:
      total_log_prob = float("-inf")
      prob = 0.0
    else:
      prob = math.exp(total_log_prob)
    log_probs_list.append(total_log_prob)
    prob_list.append(prob)

  return log_probs_list, prob_list

def apply_temperature_and_topk(logits, temp, k):
  if temperature != 1.0:
    logits = logits / temperature
  if top_k is not None:
    # top-k values from logits; shape: [batch, seq_len, top_k]
    top_k_vals = torch.topk(logits_batch, top_k, dim=-1)
    # k-th element in top-k values for each item in batch
    kth_vals = top_k_vals.values[:, :, -1].unsqueeze(-1)
    mask = logits < kth_vals
    logits = logits.masked_fill(mask, float("-inf"))
  return logits
```

Listing 2: Computing logprobs

C. Additional notes on extraction risk, probabilistic copies, and U.S. copyright

Not all examples are extractable. We can see this by looking at the maximum probabilistic extraction rate. (See Appendix F.) Because there is a logarithmic relationship between n and the p—that is, $n \propto \log(1-p)$—there is a maximum total extraction rate that we can get. One can see this by looking Equation (3):

Proof. Let $p_{\boldsymbol{z}} \in [0, 1)$ denote the probability of extracting example $\boldsymbol{z}$ in a single trial, i.e., one prompt. Then the probability of not extracting $\boldsymbol{z}$ in a single trial is $1 - p_{\boldsymbol{z}}$, and the probability of never extracting $\boldsymbol{z}$ after n independent trials is:

$$(1 - p_{\boldsymbol{z}})^n$$

Therefore, the probability of extracting $\boldsymbol{z}$ at least once in n trials is:

$$\mathbb{P}^{\boldsymbol{z}}_{\text{extracted}}(n) = 1 - (1 - p_{\boldsymbol{z}})^n$$

Now consider two cases:

- If $p_{\boldsymbol{z}} > 0$, then $\lim_{n\to\infty}(1 - p_{\boldsymbol{z}})^n = 0$, and thus $\lim_{n\to\infty} \mathbb{P}^{\boldsymbol{z}}_{\text{extracted}}(n) = 1$
- If $p_{\boldsymbol{z}} = 0$, then $\mathbb{P}^{\boldsymbol{z}}_{\text{extracted}}(n) = 0$ for all n

Now suppose we have a finite set of examples $\mathcal{Z}$, each with its own extraction probability $p_{\boldsymbol{z}} \in [0, 1)$. (This is how extraction rates are computed in practice, for such a finite set of examples $\mathcal{Z}$.) The expected number of *distinct* examples extracted at least once over $\boldsymbol{z}$ trials is:

$$\mathbb{E}[\text{\# unique extractions}] = \sum_{\boldsymbol{z}\in\mathcal{Z}} [1 - (1 - p_{\boldsymbol{z}})^n]$$

Taking the limit as $n \to \infty$, we get:

$$\lim_{n\to\infty} \mathbb{E}[\text{\# unique extractions}] = \sum_{\boldsymbol{z}\in\mathcal{Z}} \mathbb{I}[p_{\boldsymbol{z}} > 0]$$

where $\mathbb{I}[p_{\boldsymbol{z}} > 0]$ is the indicator function, equal to 1 when $p_{\boldsymbol{z}} > 0$, and 0 otherwise.

Therefore, when some examples may be unextractable (i.e., $p_{\boldsymbol{z}} = 0$), the maximum number of unique extractions is given by:

$$|\mathcal{Z}_{\text{extractable}}| := |\{\boldsymbol{z} \in \mathcal{Z} \mid p_{\boldsymbol{z}} > 0\}|$$

And so, even with an infinite number of trials, the number of distinct examples ever extractable is bounded by $|\mathcal{Z}_{\text{extractable}}| \leq |\mathcal{Z}|$. This means, for this distinct set, there is a maximum probabilistic extraction rate, which is independent of the choice of p, in (n, p)-discoverable extraction. □

We can see this in practice for plots that examine the maximum extraction rate, e.g., see Appendix D.

Unextractable sequences in practice. There are various ways that a given example $\boldsymbol{z}$ could have $p_{\boldsymbol{z}} = 0$, thereby limiting the total extraction rate. For one, if we use top-k sampling (as we do in this paper, with $k = 40$), not all tokens in the vocabulary are reachable during sample (e.g., for Llama models, only 40 out of $32,000$ tokens would be able to be generated in a given sampling iteration). When computing (n, p)-discoverable extraction (see 2), if the actual token in the target suffix that we are summing over (summing their $\log$ probabilities) is not in the top 40 tokens, then it becomes impossible to generate the target suffix. In other words, that example is not extractable (with respect to this model and decoding scheme).

For another, numbers are represented in computers with a finite number of bits. This means that there is a minimum number that can be represented (> 0) before it gets rounded down to 0 in the computer. For instance, imagine a token vocabulary

of $32,000$ tokens (as with Llama models). If the model just generated tokens completing randomly (a true monkey at the typewriter (The Simpsons: Last Exit to Springfield (Fox television broadcast Mar. 11, 1993); Borel, 1913)), then at each iteration every token would have probability $\frac{1}{32000}$. When generating 71 such tokens in a row (i.e., a 71-token sequence), the probability of any such sequence is $\frac{1}{32000}^{71} \approx 1.3626 \times 10^{-320}$. Once we try to generate a 72nd token, the probability for the whole sequence becomes $\frac{1}{32000}^{72} = 0$, in terms of how the probability gets represented in the computer. There is underflow (i.e., rounding down to 0). As a result, even with temperature sampling alone, not every sequence (for long enough sequences) is extractable in practice.

Extractability and copyright. Even though not every sequence is extractable, this does not mean that sequences that have very low extraction probabilities are memorized in a sense that is of use to copyright. The important point to distinguish is if the sequence $\boldsymbol{z}$ in question is more likely to be generated than is expected from the base distribution of the model (and the given decoding scheme). While we don't attempt to draw such a line in this paper, we note that Hayes et al. (2025) does experiments to validate their (n, p)-discoverable extraction metric. They compare the extraction of training data to the rate of, by chance, generating test data (unseen data, that cannot by definition have been memorized). Put differently, they empirically test the "monkey at a typewriter" problem, to see if this is having an effect on their measurements—to confirm that their measurements of extraction are catching true instances of memorization.

They find that it takes orders of magnitude more queries to generate test data (at very large n) than it does to reliably extract training data, indicating that the metric is capturing valid instances of memorization. Of course, this also suggests that there exists a large enough n at which it becomes challenging to distinguish truly extracted training data versus spontaneous generation of test data. In this paper, we avoid being in this regime by choosing a very conservative minimum value of $p_{\boldsymbol{z}}$ (and thus, a capped maximum expected n) in our plots. Every number we report in the main paper is for $p_{\boldsymbol{z}} \leq 1 \times 10^{-6}$. We show a few instances of extractions that have smaller $p_{\boldsymbol{z}}$ than this in Appendix H, for illustrative purposes.

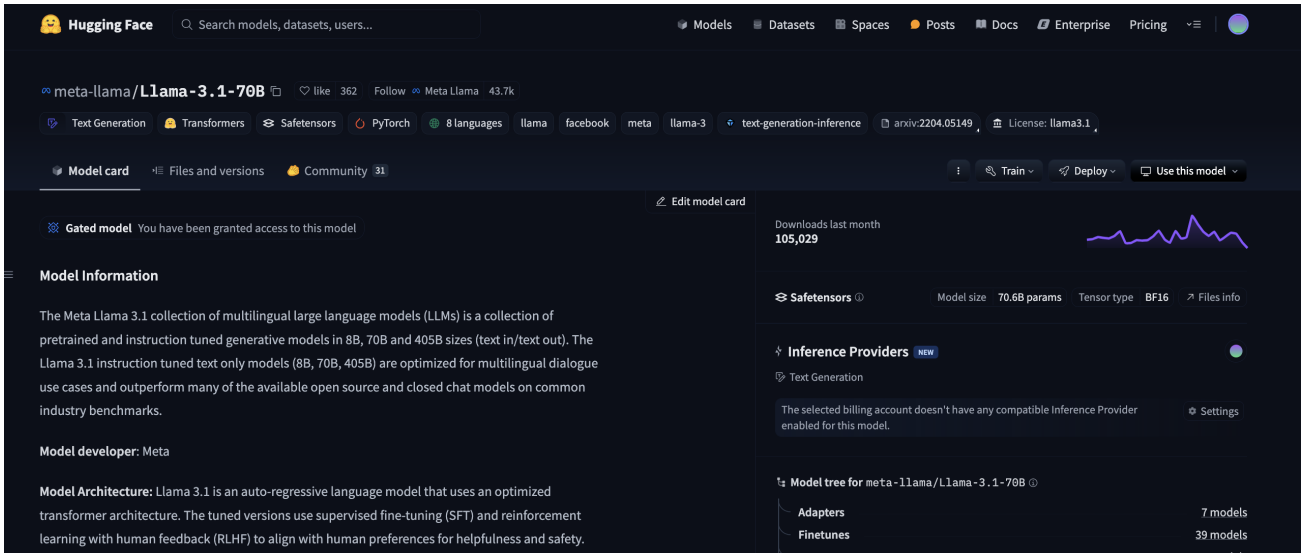

Figure 7: Screenshot of the month's downloads of LLAMA 3.1 70B on HuggingFace, taken by the authors in May 2025.

D. Testing our measurement pipeline

To confirm the accuracy of our implementation of (n, p)-discoverable extraction, as well as to benchmark its efficiency in comparison to the implementation in Hayes et al. (2025), we re-run the experiments from Hayes et al. (2025) for small Pythia models on the Enron dataset. We discussed the original run times and sample of Enron with the authors of Hayes et al. (2025) to ensure consistency. We provide some brief details about these experiments as to show due diligence for our implementation in our copyright-relevant setting.

We ran these experiments with both `float16` and `float32`, and report results for `float32` to align with Hayes et al. (2025). (However, we note no large differences, an observation we defer to separate work). In Table 1, we show the runtime for our experiments on the $10,000$, 100-token examples Hayes et al. (2025) drew from Enron. In Table 2, we report the same metrics as Hayes et al. (2025): the greedy-sampled discoverable extraction rate and (as a point of comparison) the maximum (n, p)-discoverable extraction rate. We also include extraction curves in Figure 8, which show the same settings as the same experiments in Hayes et al. (2025). We run these experiments on the same 4 A100s as all of our other experiments.

	Time (mm:ss)	**Parallelism**	**Batch size**
Pythia 1B	02:30.42	Data (4 GPUs)	250
Pythia 2.8B	03:36.35	Data (4 GPUs)	250

Table 1: Runtime for different Pythia models on a $10,000$-example dataset drawn from Enron.

Pythia 1B on Enron subset	**Hayes et al. (2025) (No BOS token)**	**With BOS token**
Greedy extraction rate	0.76%	0.74%
Max. (n, p)-discoverable extraction rate	5.27%	5.52%
Pythia 2.8B on Enron subset	**Hayes et al. (2025) (No BOS token)**	**With BOS token**
Greedy extraction rate	1.3 %	1.82%
Max. (n, p)-discoverable extraction rate	9.04%	9.47%

Table 2: Greedy discoverable extraction and maximum (n, p)-discoverable extraction rates for Pythia models on a $10,000$-example subset from Enron. We replicate Hayes et al. (2025), which did not include the beginning of sequence (BOS) token in the examples they extract for their experiments using Pythia models. We compare to the extraction rates where the BOS token is included at the start of the example. Note that the (n, p)-discoverable extraction rate increases with the presence of the BOS token. See also Figure 8.

Note that, in the process of running these experiments, we identified a small bug in Hayes et al. (2025). The Pythia tokenizer does *not* add a beginning of sequence (BOS) token by default to the sequences it tokenizes. This is reflective of GPT-2-style tokenizers. However, for our setting, it is best practice to include this token at the start, as we will be submitting these sequences as input for inference to LLMs (that are trained to expect this token at the beginning of the sequence). When we manually prepend this token and re-run the experiments from Hayes et al. (2025), we observe elevated extraction rates compared to Hayes et al. (2025). See Table 2. We similarly will prepend the BOS token in our new experiments using Pythia models.

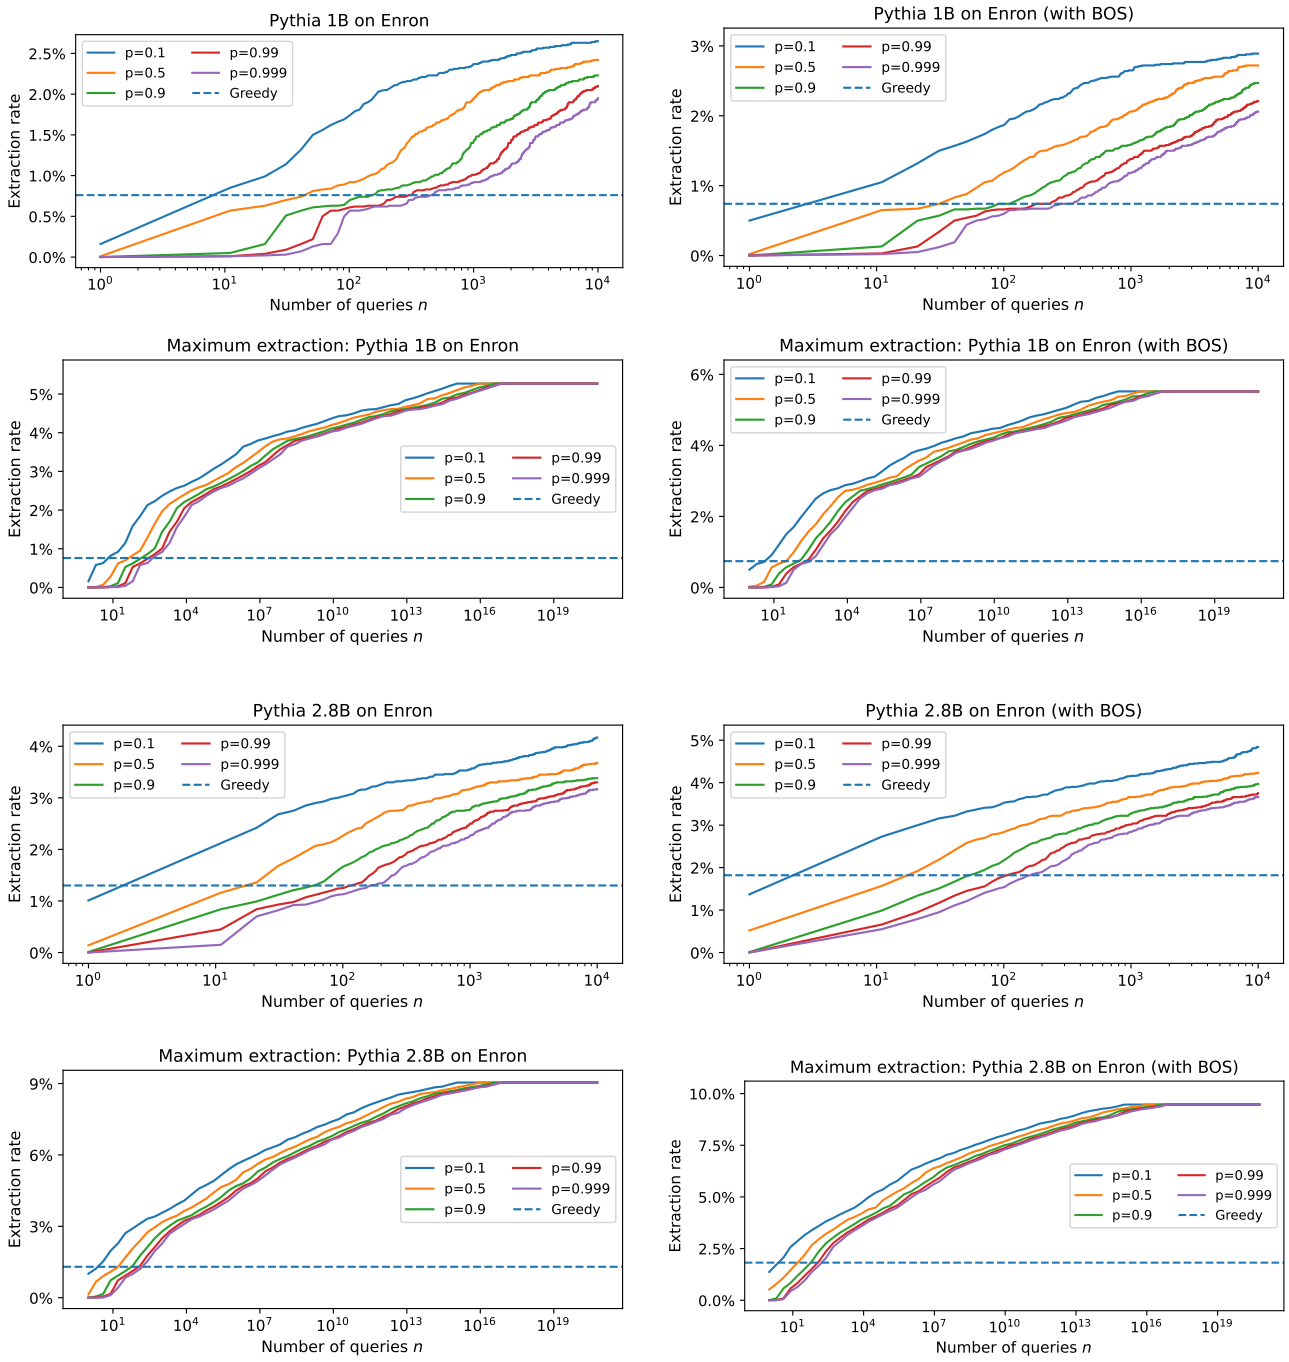

Figure 8: Replicating Hayes et al. (2025): Extraction curves for PYTHIA 1B and PYTHIA 2.8B on a $10,000$-example sample of Enron, where each example is 100 tokens ($50/50$ prefix/suffix split). These results are for temperature $T = 1$, top-k with $k = 40$ sampling. We show to sets of plots for each model. The left column is identical to Hayes et al. (2025); the right column shows the same set of curves, but for the version of the experiment where we manually prepend the BOS token to each example.

E. Details on experimental settings

In this Appendix, we discuss overall experimental setup details. We describe additional experimental details in different appendies about those experiments.

Data. In all of our experiments for this project, we draw our prompts (prefixes) from `Books3`; we verify extraction against the corresponding suffixes drawn from `Books3`. We obtained this dataset from a previous (2022) download of The Pile (Gao et al., 2020), which is stored on a university cluster and which we use for research purposes on language modeling.

The status of `Books3` as a research artifact (as well as a training corpus for language models) remains unresolved. (See Section 3.) We note that this dataset remains widely available in the research community; for example, found in the HuggingFace-hosted version of The Pile (EleutherAI/the_pile_deduplicated), as well as in various other data repositories (e.g., SaylorTwift/the_pile_books3_minus_gutenberg; amongglue/books3-subset-raw; CANBERT/pile_books3_text). It also features as part of a benchmark task in the popular HELM evaluation suite.

We sample this dataset at random (with caveats) for our overall extraction rate experiments in Section 3. We detail these experiments in Appendix F. We work with a specific set of books for our book-specific experiments. We list these books in Tables 3 and 4, according to how we use these books in different types of experiments. In Table 3, we list the books in our sliding-window experiments. (See Section 4 and Appendix G.1.) In Table 4, we list the books we use in a smaller set of random subsampling experiments. (See Section G.2.) In this second set of experiments, we attempted (and confirmed) that we would obtain reasonable approximations of per-book extraction rates with a cheaper sampling strategy. The sliding-window approach remains the most fruitful for completeness of identifying memorization hot-spots within books.

Book selection. We began our project with *The Great Gatsby* (Fitzgerald, 1925). This choice stems from a conversation in 2023 while writing a different paper, when Anthropic announced its 100K-length context for Claude with the example of prompting with the entirety of *The Great Gatsby* (Anthropic, 2023).

We then continued our selection process with the books listed with the associated plaintiffs in the (amended) class action complaint of Kadrey et al. v. Meta (Kadrey et al. v. Meta Platforms, Inc., pp. 4-5). This is how we sourced the first 13 books from `Books3`. We expanded to additional books by these plaintiffs, and also added in some generally popular books that we chose among the team (e.g., Rowling (1998); Tolkien (1937); Camus (1955); Heller (1961)) and some (less publicly popular) academic books based on our personal preferences (e.g., Zittrain (2008); Barolini (2006)). (This team, after all, is composed of academics.) We deliberately include selections from the public domain (e.g., Carroll (1865); Shelley (1818); Woolf (1928)), and an example of a book that was published under a permissive CC license (Doctorow (2003)). We added in Carroll (1865) because of concurrent work that was published on arXiv just as we were about to post our own: Ma et al. (2025). To round out the list, we also selected books at random from the manifest file that accompanies the `Books3` archive (e.g., Jacobs (1999), Rouighi (2011), Cogburn (2012)). Once we observed that certain mainstream-popular books were highly memorized by LLAMA 3.1 70B, we added in 5 additional such books to test: Martin (1996), Brown (2003), Meyer (2005), Orwell (1949), and Sandberg (2013).

Compute resources. We run all of our experiments in a slurm cluster environment, using the same node with 4 A100 GPUs.

	Author	Title	Year	Status	Books3 path
1	Laura Lippman	*After I'm Gone* (Lippman, 2014)	2014	©	the-eye.eu/public/Books/Bibliotik/A/After I'm Gone - Laura Lippman.epub.txt
2	Lewis Carroll	*Alice's Adventures in Wonderland* (Carroll, 1865)	1865	PD	the-eye.eu/public/Books/Bibliotik/A/Alice's Adventures in Wonderlan - Lewis Carroll.epub.txt
3	Jacqueline Woodson	*Another Brooklyn* (Woodson, 2016)	2016	©	the-eye.eu/public/Books/Bibliotik/A/Another Brooklyn - Jacqueline Woodson.epub.txt
4	Christopher Golden	*Ararat* (Golden, 2017)	2017	©	the-eye.eu/public/Books/Bibliotik/A/Ararat - Christopher Golden.epub.txt
5	Jacqueline Woodson	*Brown Girl Dreaming* (Woodson, 2014)	2014	©	the-eye.eu/public/Books/Bibliotik/B/Brown Girl Dreaming - Jacqueline Woodson.epub.txt
6	Teodolinda Barolini	*Dante and the Origins of Italian Literary Culture* (Barolini, 2006)	2006	©	the-eye.eu/public/Books/Bibliotik/D/Dante and the Origins of Italia - Barolini, Teodolinda;.epub.txt
7	Christopher Golden	*Dead Ringer* (Golden, 2016)	2016	©	the-eye.eu/public/Books/Bibliotik/D/Dante and the Origins of Italia - Barolini, Teodolinda;.epub.txt
8	Cory Doctorow	*Down and Out in the Magic Kingdom Cory Doctorow* (Doctorow, 2003)	2003	CC-BY-NC-SA	the-eye.eu/public/Books/Bibliotik/D/Down and Out in the Magic Kingd - Cory Doctorow.epub.txt
9	Junot Díaz	*Drown* (Díaz, 1996)	1996	©	the-eye.eu/public/Books/Bibliotik/J/Junot Diaz - Drown.epub.txt
10	Jon Cogburn	*Dungeons and Dragons and Philosophy* (Cogburn, 2012)	2012	©	the-eye.eu/public/Books/Bibliotik/2/2012 Jon Cogburn - Dungeons and Dragons and Philosophy - Raiding the Temple of Wisdom˙Rsnl.epub.txt
11	Lysa TerKeurst	*Embraced* (TerKeurst, 2018)	2018	©	the-eye.eu/public/Books/Bibliotik/E/Embraced - Lysa Terkeurst.epub.txt
12	J.K. Rowling	*Harry Potter and the Sorcerer's Stone* (Rowling, 1998)	1998	©	the-eye.eu/public/Books/Bibliotik/H/Harry˙Potter˙and˙the˙ ˙orcerers˙Stone-Rowling.epub.txt
13	Jonathan Zittrain	*The Future of the Internet and How to Stop It* (Zittrain, 2008)	2008	©	the-eye.eu/public/Books/Bibliotik/J/Jonathan Zittrain - The Future of the Internet.epub.txt
14	David Henry Hwang	*M. Butterfly* (Hwang, 1988)	1988	©	the-eye.eu/public/Books/Bibliotik/M/M. Butterfly - David Henry Hwang.epub.txt
15	Rachel Louise Snyder	*No Visible Bruises* (Snyder, 2019)	2019	©	the-eye.eu/public/Books/Bibliotik/N/No Visible Bruises - Rachel Louise Snyder.epub.txt
16	Richard Kadrey	*Sandman Slim* (Kadrey, 2009)	2009	©	the-eye.eu/public/Books/Bibliotik/S/Sandman Slim - Richard Kadrey.epub.txt
17	Laura Lippman	*Sunburn* (Lippman, 2018)	2018	©	the-eye.eu/public/Books/Bibliotik/S/Sunburn - Laura Lippman.epub.txt
18	Ta-Nehisi Coates	*The Beautiful Struggle* (Coates, 2009)	2009	©	the-eye.eu/public/Books/Bibliotik/T/The Beautiful Struggle - Ta-Nehisi Coates.epub.txt
19	Sarah Silverman	*The Bedwetter* (Silverman, 2010)	2010	©	the-eye.eu/public/Books/Bibliotik/T/The Bedwetter - Sarah Silverman.epub.txt
20	Junot Díaz	*The Brief Wondrous Life of Oscar Wao* (Díaz, 2007)	2007	©	the-eye.eu/public/Books/Bibliotik/T/The Brief Wondrous Life of Oscar Wao.epub.txt
21	Andrew Sean Greer	*The Confessions of Max Tivoli*	2005	©	the-eye.eu/public/Books/Bibliotik/T/The Confessions of Max Tivoli - Andrew Sean Greer.epub.txt
22	F. Scott Fitzgerald	*The Great Gatsby* (Fitzgerald, 1925)	1925	PD	the-eye.eu/public/Books/Bibliotik/T/The Great Gatsby - F. Scott Fitzgerald.epub.txt
23	J.R.R. Tolkien	*The Hobbit* (Tolkien, 1937)	1937	©	the-eye.eu/public/Books/Bibliotik/T/The Hobbit (Houghton Mifflin Harcourt) (75th Anniversary Edition) [Epub] - J.R.R. Tolkien.epub.txt
24	Ramzi Rouighi	*The Making of a Mediterranean Emirate* (Rouighi, 2011)	2011	©	the-eye.eu/public/Books/Bibliotik/T/The Making of a Mediterranean E - Rouighi, Ramzi;.epub.txt
25	Albert Camus (Justin O'Brien, translator)	*The Myth of Sisyphus* (Camus, 1955)	1955	©	the-eye.eu/public/Books/Bibliotik/2/2013(orig1942) Albert Camus - The Myth of Sisyphus[Transl Justin O'Brien]˙Ral.epub.txt
26	Christopher Farnsworth	*The President's Vampire* (Farnsworth, 2011)	2011	©	the-eye.eu/public/Books/Bibliotik/T/The President's Vampire - Christopher Farnsworth.epub.txt
27	Ta-Nehisi Coates	*The Water Dancer: A Novel* (Coates, 2019)	2019	©	the-eye.eu/public/Books/Bibliotik/T/The Water Dancer - Ta-Nehisi Coates;.epub.txt
28	Lysa TerKeurst	*Unglued: Making Wise Choices in the Midst of Raw Emotions* (TerKeurst, 2012)	2012	©	the-eye.eu/public/Books/Bibliotik/L/Lysa TerKeurst, Unglued.epub.txt
29	Ta-Nehisi Coates	*We Were Eight Years in Power: An American Tragedy* (Coates, 2017)	2017	©	the-eye.eu/public/Books/Bibliotik/W/We Were Eight Years in Power - Ta-Nehisi Coates.epub.txt
30	Matthew Klam	*Who Is Rich?* (Klam, 2017)	2017	©	the-eye.eu/public/Books/Bibliotik/W/Who Is Rich˙ - Matthew Klam.epub.txt
31	George R.R. Martin	*A Game of Thrones* (Martin, 1996)	1996	©	the-eye.eu/public/Books/Bibliotik/A/A Game of Thrones - George R. R. Martin.epub.txt
32	Sheryl Sandberg	*Lean In* (Sandberg, 2013)	2013	©	the-eye.eu/public/Books/Bibliotik/L/Lean In - Sheryl Sandberg.epub.txt
33	George Orwell	*Nineteen-Eighty Four* (Orwell, 1949)	1949	©	the-eye.eu/public/Books/Bibliotik/N/Nineteen Eighty-Four (The Annotated Edition) - George Orwell.epub.txt
34	Dan Brown	*The Da Vinci Code* (Brown, 2003)	2003	©	the-eye.eu/public/Books/Bibliotik/D/Dan Brown - The Da Vinci Code.epub.txt
35	Stephenie Meyer	*Twilight* (Meyer, 2005)	2005	©	the-eye.eu/public/Books/Bibliotik/T/Twilight - Stephenie Meyer.epub.txt
36	James Joyce	*Ulysses* (Joyce, 1922)	1922	PD	the-eye.eu/public/Books/Bibliotik/U/Ulysses - James Joyce - Penguin Group - 2000.epub.txt

Table 3: The 36 Books (drawn from Books3) that we sampled for our book-specific sliding window extraction experiments. See Section 4 and Appendix G.1.

	Author	Title	Year	Status	Books3 path
37	Jacqueline Woodson	*After Tupac and D Foster* (Woodson, 2008)	2008	©	the-eye.eu/public/Books/Bibliotik/A/After Tupac & D Foster - Jacqueline Woodson.epub.txt
38	Betty E.M. Jacobs	*All the Onions* (Jacobs, 1999)	1999	©	the-eye.eu/public/Books/Bibliotik/A/All the Onions - Betty E. M. Jacobs.epub.txt
39	Richard Kadrey	*Aloha from Hell* (Kadrey, 2011)	2011	©	the-eye.eu/public/Books/Bibliotik/A/Aloha from Hell - Richard Kadrey.epub.txt
40	Ken Wharton	*Another Bloody Chapter in an Endless Civil War* (Wharton, 2016)	2016	©	the-eye.eu/public/Books/Bibliotik/A/Another Bloody Chapter In an Endless Civil War, Volume 1 - Ken Wharton.epub.txt
41	Christopher Farnsworth	*Blood Oath* (Farnsworth, 2010)	2010	©	the-eye.eu/public/Books/Bibliotik/B/Blood Oath - Farnsworth, Christopher.epub.txt
42	Christopher Golden	*Blood of the Four* (Golden, 2018)	2018	©	the-eye.eu/public/Books/Bibliotik/B/Blood of the Four - Christopher Golden.epub.txt
43	Joseph Heller	*Catch-22* (Heller, 1961)	2018	©	the-eye.eu/public/Books/Bibliotik/C/Catch-22 - Joseph Heller.epub.txt
44	David Bret	*Elizabeth Taylor* (Bret, 2011)	2011	©	the-eye.eu/public/Books/Bibliotik/E/Elizabeth Taylor - David Bret.epub.txt
45	Mary Shelley	*Frankenstein* (Shelley, 1818)	1818	PD	the-eye.eu/public/Books/Bibliotik/M/Mary Shelley (1818) Frankenstein (Wodsworth Classics) [2011 UK retail].epub.txt
46	David Henry Hwang	*Golden Child* (Hwang, 1998)	1998	©	the-eye.eu/public/Books/Bibliotik/G/Golden Child - David Henry Hwang.epub.txt
47	Richard Kadrey	*Kill the Dead* (Kadrey, 2014)	2014	©	the-eye.eu/public/Books/Bibliotik/K/Kill the Dead - Richard Kadrey.epub.txt
48	Joseph Finder	*Killer Instinct* (Finder, 2006)	2006	©	the-eye.eu/public/Books/Bibliotik/K/Killer Instinct - Joseph Finder.epub.txt
49	Virginia Woolf	*Orlando* (Woolf, 1928)	1928	PD	the-eye.eu/public/Books/Bibliotik/O/Orlando (Canongate) - Virginia Woolf (retail).epub.txt
50	Christopher Golden	*Poison Ink* (Golden, 2008)	2008	©	the-eye.eu/public/Books/Bibliotik/C/Christopher Golden - Poison Ink [retail].epub.txt
51	Matthew Klam	*Sam the Cat* (Klam, 2001)	2001	©	the-eye.eu/public/Books/Bibliotik/S/Sam the Cat - Matthew Klam.epub.txt
52	Neal Stephenson	*The Diamond Age* (Stephenson, 1996)	1996	©	the-eye.eu/public/Books/Bibliotik/S/Stephenson, Neal - The Diamond Age (1995; Penguin, 1996).epub.txt
53	Andrew Sean Greer	*The Path of Minor Planet* (Greer, 2002)	2002	©	the-eye.eu/public/Books/Bibliotik/T/The Path of Minor Planets - Andrew Sean Greer.epub.txt
54	Richard Kadrey	*The Wrong Dead Guy* (Kadrey, 2017)	2017	©	the-eye.eu/public/Books/Bibliotik/T/The Wrong Dead Guy - Richard Kadrey.epub.txt
55	J.R.R. Tolkien	*Tree and Leaf* (Tolkien, 1964)	1964	©	the-eye.eu/public/Books/Bibliotik/T/Tree and Leaf - J. R. R. Tolkien.epub.txt
56	David Henry Hwang	*Trying to Find Chinatown* (Hwang, 1999)	1999	©	the-eye.eu/public/Books/Bibliotik/T/Trying to Find Chinatown - David Henry Hwang.epub.txt

Table 4: The 20 books (drawn from `Books3`) that we sampled for our less expensive, book-specific random sampling experiments. See Appendix G.2.

F. Additional information on extraction rates

For our experiments that produce average extraction rates over random samples of Books3, we do 3 separate runs on $40,000$ examples on 13 different models: LLAMA 1 7B, LLAMA 1 13B, LLAMA 1 65B (Touvron et al., 2023a), LLAMA 2 7B, LLAMA 2 13B, LLAMA 2 70B (Touvron et al., 2023b), LLAMA 3.1 8B, LLAMA 3.1 70B (Grattafiori et al., 2024), PYTHIA 6.9B, PYTHIA 12B (Biderman et al., 2023), DEEPSEEK V1 7B (DeepSeek-AI et al., 2024), GEMMA 2 9B (Team et al., 2024), and PHI 4 (Abdin et al., 2024) (which is a 14B model). We group these models (roughly) into three sizes: small, medium and large.

Sampling procedure. For each run, we sample $4,000$ books without replacement from Books3. Then, from each of these books, we sample 10 non-overlapping 100-token examples, for a total of $40,000$ examples. We make sure that the start token is a space, to ensure that we are not starting in an "odd" spot—one that could lead to unusual tokenization. We run these $40,000$ examples through each of the 13 models to compute (n, p)-discoverable extraction metrics. We follow this procedure 3 separate times (i.e., for a total of $120,000$ examples across $12,000$ different books) to get a sense of the variance across extraction metrics for different samples. We report results for an example size of 100 tokens (50/50 prefix/suffix split) for greedy-sampled discoverable extraction and (n, p)-discoverable extraction with temperature $T = 1$ and top-k sampling with $k = 40$.

We summarize the results in Table 5. This table shows the maximum possible (n, p)-discoverable extraction rate, compared to the greedy-sampled discoverable extraction rate for each model. We show the average and ± 1 standard deviation over the 3 different trials. We also plot these average rates (and a third point—the (n, p)-discoverable extraction rate for $n = 10000$ and $p = 0.5$) in Figure 9. This new (intermediate) point of comparison is the extraction rate of examples that have at least a 50% chance of being extracted at least once within $10,000$ independent prompts with the same prefix. We show more detailed extraction-rate curves (as a function of n and for various settings of p) in Figure 10. Each run of $40,000$ examples took between approximately 8 minutes and 45 minutes on the 4 A100s we use for all of our experiments. We omit detailed timing results for brevity.

High-level takeaways. As we expected (and discuss in Section 3), *average* extraction rates are relatively low. The minimum (n, p)-discoverable extraction rate that we observe (for $n = 10000$, $p = 0.5$) is for PHI 4 ($0.18 \pm 0.03\%$); the largest is for LLAMA 3.1 70B ($0.79 \pm 0.04\%$).

Aligning with prior work, larger models tend to memorize more than smaller ones (Hayes et al., 2025; Carlini et al., 2023b; Nasr et al., 2023; 2025). For our results, this is true within the same model family, and also across model families. PHI 4 is the exception; it is a 14B model with an extraction rate that is on par with the small (7-9B) models that we test. Unlike the other models we test, PHI 4 was trained predominantly on synthetic data; based on the PHI 4 technical report (Abdin et al., 2024), Books3 should not have been included in its training data. We are still able to extract data that is in Books3, and we investigate this further in Appendix H.

Later generations of Llama model families tend to memorize more Books3-contained text. For just one example, LLAMA 3.1 70B memorizes more on average than LLAMA 3 70B, which memorizes more on average than LLAMA 1 65B. Based on more detailed results (Appendix G.1), it seems that later generations of Llama models may have been trained for (significantly) more epochs. We defer additional, related observations about LLAMA 3.1 70B to Appendix I.

Model	**Small (7–9B)**		**Medium (12–14B)**		**Large (65–70B)**	
	Greedy rate	**Max (n, p) rate**	**Greedy rate**	**Max (n, p) rate**	**Greedy rate**	**Max (n, p) rate**
Llama 1	$0.09 \pm 0.01\%$	$0.77 \pm 0.02\%$	$0.11 \pm 0.01\%$	$0.98 \pm 0.02\%$	$0.17 \pm 0.02\%$	$1.78 \pm 0.03\%$
Llama 2	$0.10 \pm 0.02\%$	$0.84 \pm 0.03\%$	$0.12 \pm 0.02\%$	$1.03 \pm 0.03\%$	$0.17 \pm 0.03\%$	$1.94 \pm 0.06\%$
Llama 3.1	$0.12 \pm 0.02\%$	$0.89 \pm 0.02\%$	–	–	$0.25 \pm 0.01\%$	$2.39 \pm 0.02\%$
Pythia	$0.13 \pm 0.03\%$	$0.67 \pm 0.04\%$	$0.14 \pm 0.03\%$	$0.79 \pm 0.03\%$	–	–
DeepSeek v1	$0.12 \pm 0.02\%$	$0.84 \pm 0.01\%$	–	–	–	–
Gemma 2	$0.10 \pm 0.02\%$	$0.63 \pm 0.00\%$	–	–	–	–
Phi 4	–	–	$0.09 \pm 0.02\%$	$0.62 \pm 0.02\%$	–	–

Table 5: Greedy and maximum probabilistic extraction rates (in %) per model family and size. See Figures 9 & 10 for more detailed extraction-rate curves.

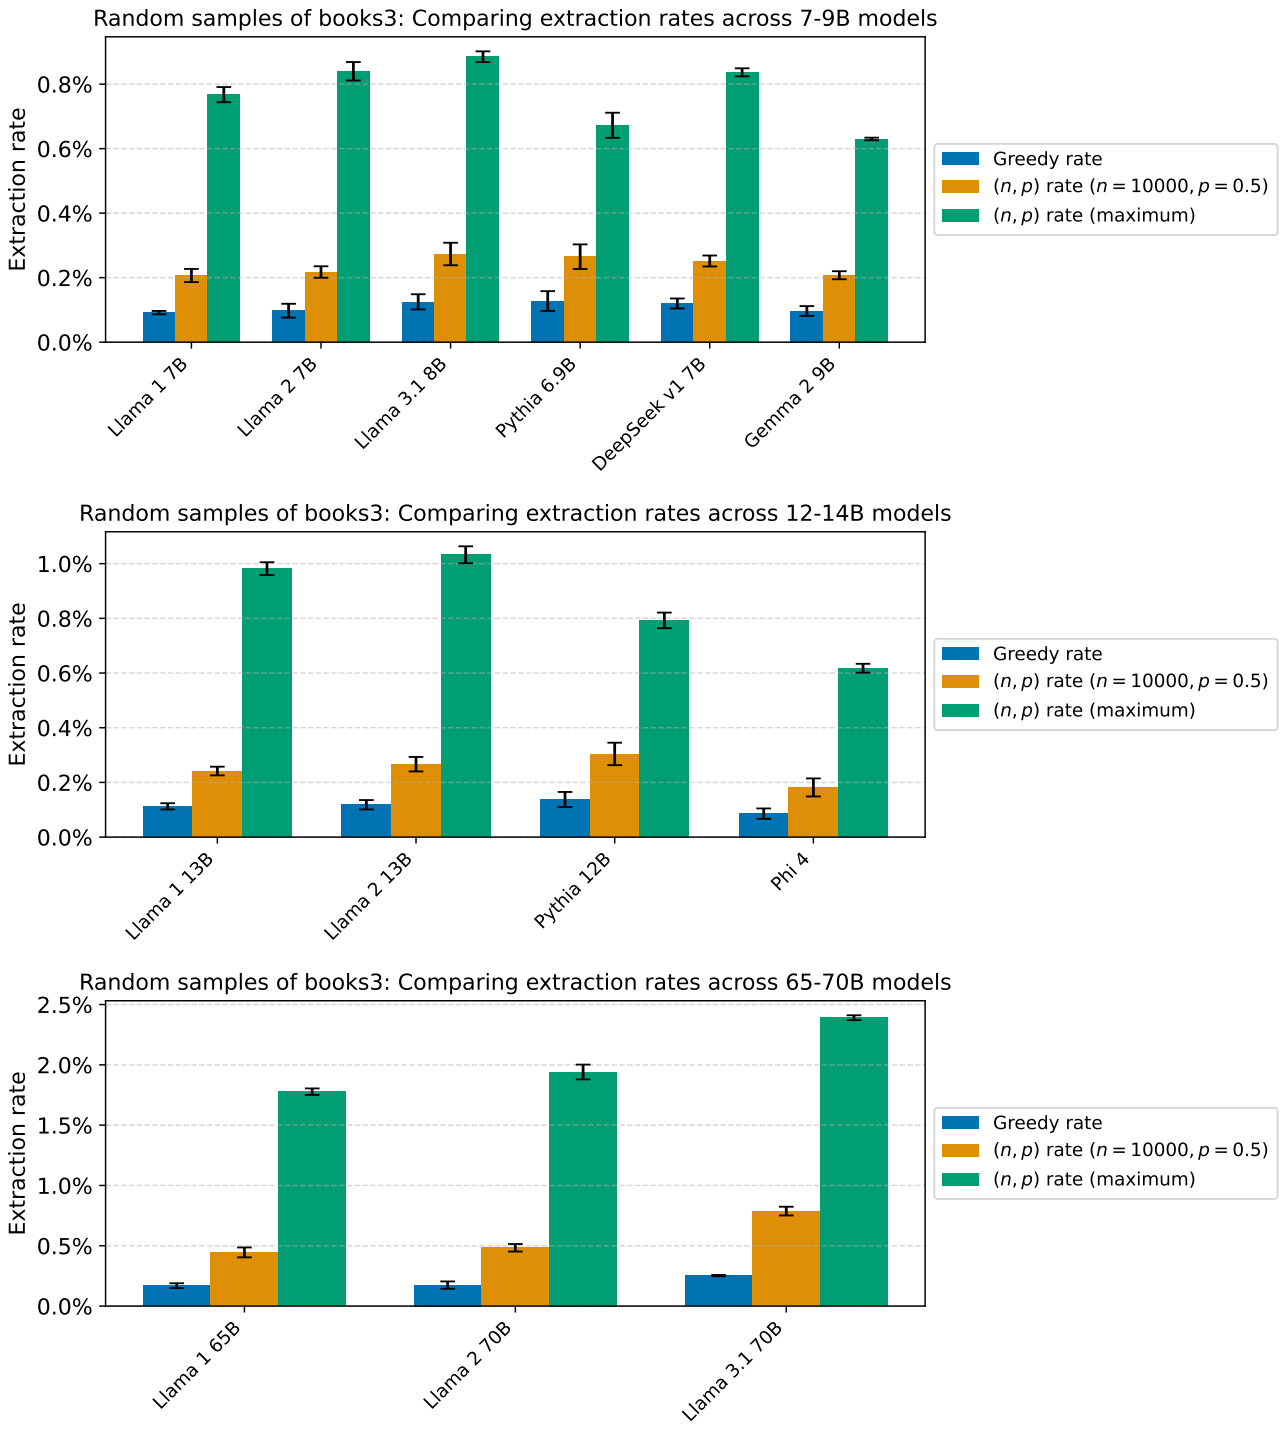

Figure 9: Comparing extraction rates for small (7-9B), medium (12-14B), and large models (65-70B) As expected, *average* extraction rates on `Books3` are relatively low for all models. Aligning with prior work, we observe that larger models memorize more. For our setting, this is true for both within and across model families. PHI 4, a 14B model, is an exception: its extraction rates are on par with smaller models. More recent generations of Llama-family models memorize more than previous ones. Error bars here indicate variance across the 3 samples of $40,000$ examples that we run.

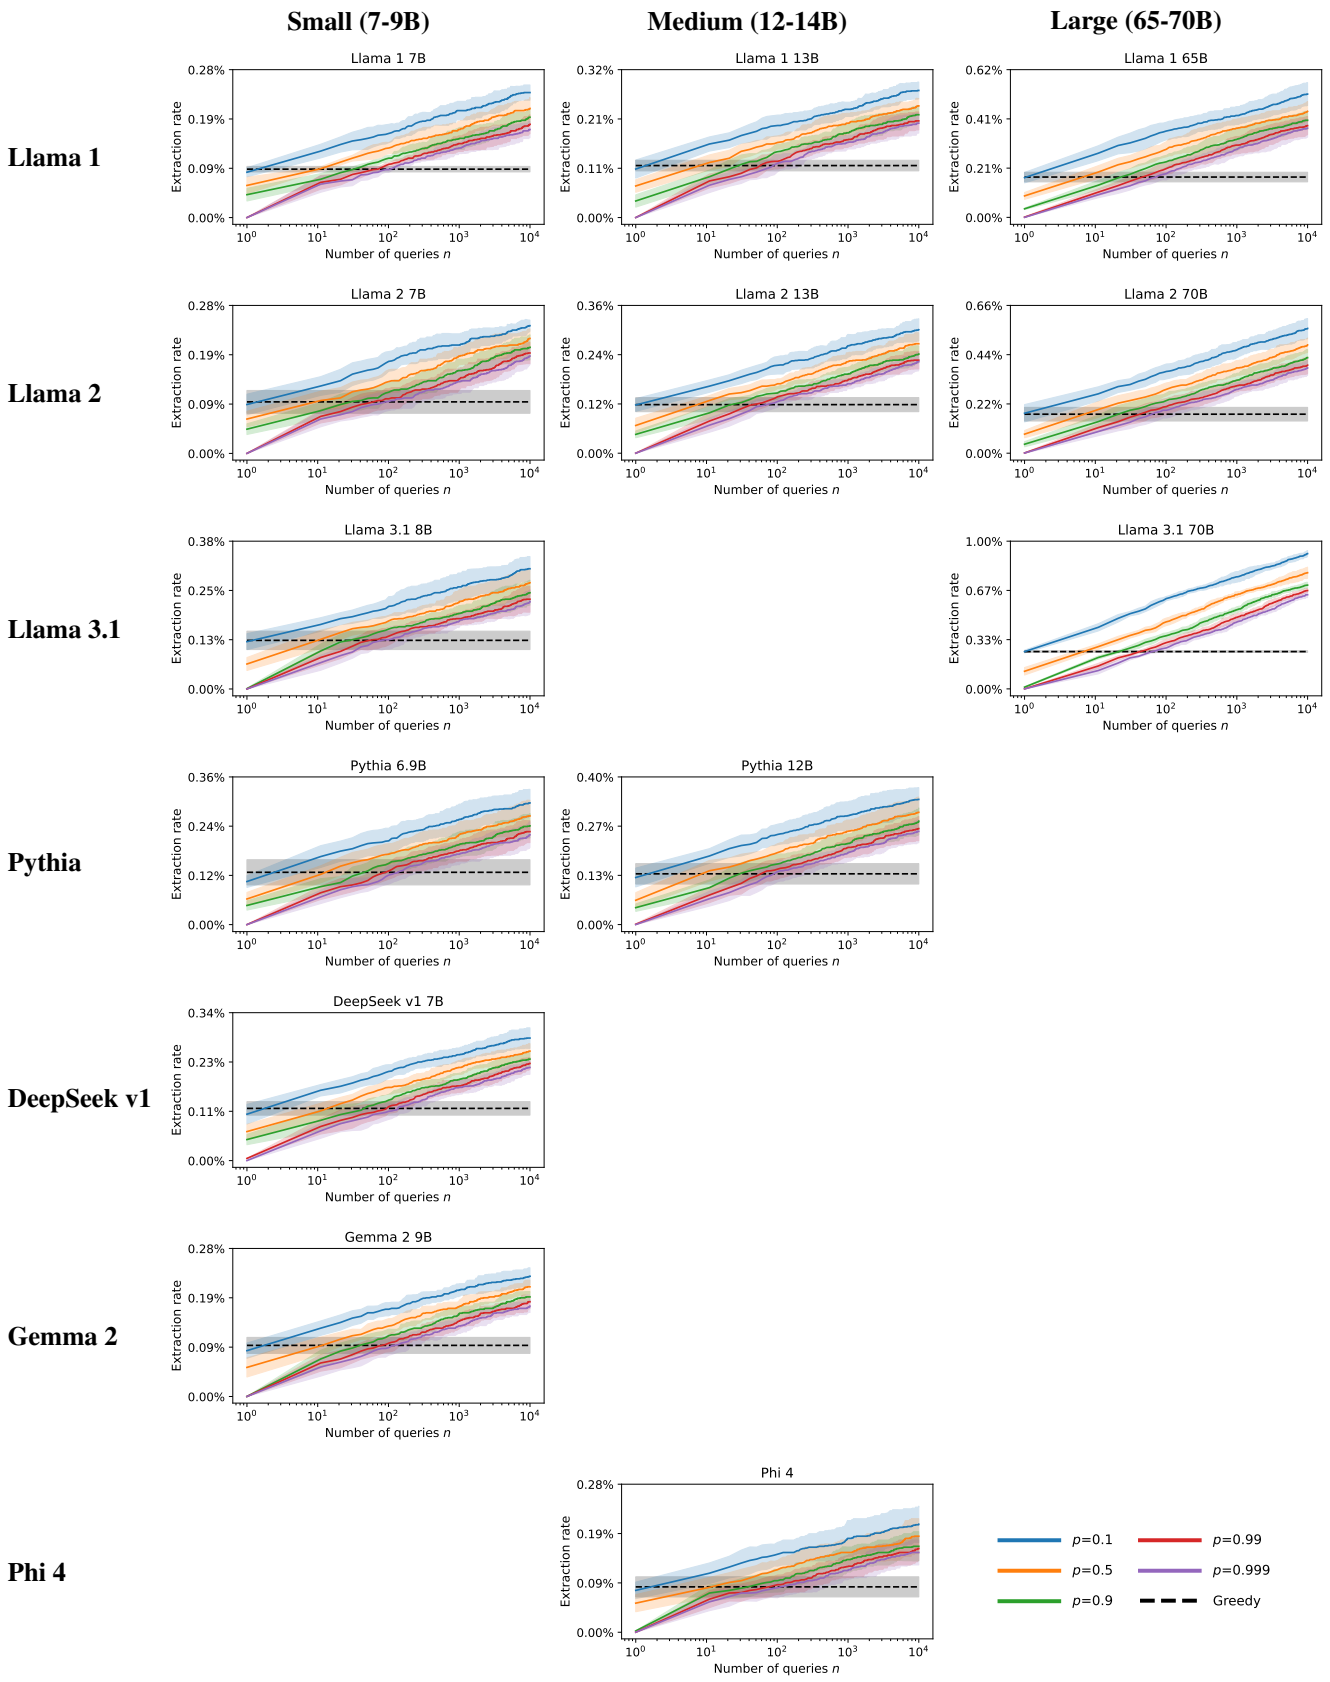

Figure 10: Extraction rates for different models averaged over 3 different samples of 100-token examples drawn from `Books3`. See Appendix F main text, Table 5, and Figure 9 for more details.

G. Additional results for per-book extraction experiments

In this appendix, we provide extensive details on our "sifting through sand" extraction experiments. That is, these experiments take the approach of either building examples out of `Books3` books (to attempt to extract) using a) a sliding-window approach (Appendix G.1) or b) an approach that randomly subsamples (with some small caveats, discussed below) examples from each book (Appendix G.2).

The sliding-window experiments are significantly more expensive than the random-sampling experiments, both because they involve more examples and because we run them with larger models. We discuss experiment-specific setup details and results in each subsection.

G.1. Sliding-window experiments

For each of the 36 books for which we run sliding-window experiments. (See Table 3 for the full list in one summarized format.) These experiments are meant to help us, effectively, "pan for gold": to identify regions within specific books where there are high-probability stretches of memorized content. The approach that we take reveals the position (in character space) in each book where these regions occur (and if they occur at all).

We discuss the sampling configurations we use for each of these experiments, then detail the sampling procedure. In our discussion of the sampling procedure, we take care to explain what the intention is behind these experiments, and how the associated results should be interpreted. For all of these experiments (as with all of the experiments in this paper), we use the same node on a slurm cluster and distribute models across the 4 A100 GPUs on that node. For these experiments, we only use model parallelism.

Sampling configuration Similar to Hayes et al. (2025), after some initial tests with different sampling algorithm hyperparameters and example lengths, we use temperature $T = 1$ and top-k sampling with $k = 40$. We believe that setting $T = 1$ makes the most sense for studying memorization, as this reflects the LLM's base probability distribution. We clip that distribution to the top-40 tokens to only consider sequences of higher-probability sampled tokens.

For these experiments, we use 100-token-length examples $\boldsymbol{z}$, with a prefix (prompt) length of 50 tokens and a suffix (target generation) length of 50 tokens. We pick 100 tokens and a 50/50 split because this is accepted in the literature as a reasonable minimum for being confident that extracted sequences are reflective of memorization (Carlini et al., 2023b; 2021; Lee et al., 2022; Hayes et al., 2025). Some work uses shorter sequences (e.g., 64 tokens and a 32/32 token split (Biderman et al., 2023)). While there is a reasonable argument to be made that such shorter sequences are still sufficiently long to be reflective of memorization, we opt for the more-standard choice of slightly longer prefixes and suffixes. Note that because our implementation of (n, p)-discoverable extraction makes computing $\log$ probabilities a post-processing operation on the output model logits (see Appendix B.1, in particular Listing 2), we also collect results for greedy-sampled discoverable extraction (Lee et al., 2022; Carlini et al., 2021; Nasr et al., 2023) by setting $T = k = 1$.

We also collect metrics for varied prefix/suffix splits (within the 100 token examples): we look at prefix lengths of 25 (with corresponding suffix length 75) and 75 (with corresponding suffix length 25). Using a longer suffix always decreases how much content we measure as extracted (since there are 75 token probabilities to multiply, where each probability is less than 1, as opposed to 50 such tokens). Increasing the prefix length is expected to increase extraction, since there is more context that we are conditioning on (Carlini et al., 2023b). This configuration (75-token prefix and 25-token suffix) is a sanity check, not something for which we intend to draw conclusions about memorization. We are not increasing the suffix length in these experiment (since we are limiting to 100 tokens overall), which also should increase extractability for extractable examples; however this lower token count is (arguably) too small to make as strong conclusions about memorization, which is our focus.

For the plots that follow, we are examining the 50/50 prefix/suffix regime (with some notable exceptions in the discussion section, see Appendix I).

Sampling procedure. To identify regions of memorization, we take the following "panning for gold" approach. For a given book, we start at the beginning of the text file in `Books3`. We sample a chunk of text that is sufficiently long to contain 100 tokens of corresponding tokenized text: to do this, we take a chunk of 2000 characters, tokenize, take the first 100 tokens of the resulting tokenized sequence as the example and discard the rest. (Each 100 token sequence is typically, but not always, reflective of 300-400 characters of text. Rare words and formatting make this vary considerably. This is why we use a character-chunk size of 2000—to make sure we have sufficient head-room to always end up with a 100 token

sequence, even if this is generally wasteful in the average case in terms of discarded tokens.) We then shift 10 *characters* forward in the book text and repeat this process; we do this for the entire length of the book, which results in approximately one example every 10 characters, i.e., `len(book_characters)/10` total examples $\boldsymbol{z}$, for which we compute $p_{\boldsymbol{z}}$. For example, *The Great Gatsby* has $270,870$ characters, which results in roughly $27,000$ examples, for which we compute the $p_{\boldsymbol{z}}$ quantity from (n, p)-discoverable extraction. (See Section 2 and Appendix B.1.)

This means, of course, that the 100-token examples overlap significantly. (Since generally speaking, 100 tokens covers 300-400 characters, shifting only 10 characters means there is a lot of overlap.) The point of doing this is that we do not know how different models were trained; *a priori*, it is not clear exactly where we should begin examples (or where they should end). With this approach, for extractable content, we expect to surface high-probability *regions* of memorized content, which we can then explore in more detail and more precisely in follow-up experiments. We picked a sliding window of 10 characters by running initial experiments with 1-, 10-, 50-, and 100-character windows. 10 characters exhibited no loss in extraction signal (at $10\times$ cheaper cost than 1 character). We intend to investigate reducing this cost further in the future.

We do *not* expect this approach to result in high per-book extraction rates. This is because, even though this approach *should* surface memorization hotspots, it will do so in a worst-case fashion. (This is precisely our goal, not deriving some procedure that maximizes the average extraction rate.) Since we are *indiscriminately* sliding the window 10 *characters* each time, we expect that many prompts will begin in "odd" locations—locations with "broken" or otherwise "strange" tokenizations, for which we expect to yield near-zero probability of extraction. For an intuition, consider this fake example. Imagine starting a prompt on `"ion"` instead of `"tion"` for `"determination"`, or any other word that ends in `"tion"`. The token for the word `"ion"` is likely/reasonably very different in the embedded token space from the token for `"tion"`, a very common suffix. Starting on `"ion"` (for a sentence containing a word that is something like `"determination"`, but with this context cropped out) should mean the probability of the continuing tokenized sequence should be low when run through the model.[4] That is true even of specialized names that might otherwise seem unique. Thus, while certain words may be more likely to follow `"Harry Potter"` in observed sentences, the same words would presumably not follow if the token begins with `"otter"` instead of `"Potter"`.

In summary, because many of the generated sequences should have near-zero extraction probability—due not only to the "broken" tokenization problem discussed above, but also to because memorization (on average) is relatively rare in high-quality models, like those we study here—we expect this procedure to yield a relatively low average, book-specific extraction rate, even in books that contain substantial amounts of memorized content.

With many zero-probability sequences (resulting not just from this "broken" tokenization, but also the fact that, on average, memorization is typically not the majority case for high-quality models, like those we study), this should keep the overall extraction rate (taken with respect to the examples we use from this procedure) relatively low, even for books where there are large amounts of memorization.

Nevertheless, in the subsections that follow, we show per-book extraction rates across the different models we test with this example-sampling procedure. We find this useful to get a very rough sense of the relative number of per-book memorization hot-spots, according to specific models.

We also provide this long explanation because, for LLAMA 3.1 70B on some (very popular) books, we observe (unexpectedly) enormous extraction rates, which we will also revisit in the discussion in Appendix I), and which we allude to in Section 4, where we show some of the results that exhibit this.

Models tested with this procedure. For these experiments, we look at results for 5 models: PYTHIA 12B (Biderman et al., 2023), PHI 4 (a 14B model) (Abdin et al., 2024), LLAMA 1 13B, LLAMA 1 65B (Touvron et al., 2023a), and LLAMA 3.1 70B (Grattafiori et al., 2024). We pick these models for the following reasons: we want to include (the largest) Pythia-family model as a baseline, since we know for certain that these models were trained on The Pile (Gao et al., 2020) (which contains `Books3`), based on EleutherAI's extensive documentation. We also want to compare on a model that (should not) have been trained on `Books3`, based on its described training procedure. We pick PHI 4 for these reasons, since Abdin et al. (2024) was trained predominantly on synthetic data: "Synthetic data constitutes the bulk of the training data for phi-4 and is generated using a diverse array of techniques, including multi-agent prompting, self-revision workflows, and instruction reversal." (PHI 4 was also trained on "web content, *licensed* books, and code repositories to extract seeds for the synthetic

[4]For example, common suffixes like `"tion"` tend to occupy clusters in embedding space because they appear in similar contexts. Fragments or rare (sub)words like `"ion"` tend to be more scattered/ not semantically coherent units. Similarly, this is true `"otter"` and `"Potter"`.

data pipeline," emphasis added.) Further, PHI 4 is 14B parameters, which is a similar size to PYTHIA 12B, which facilitates cross-model comparisons of memorization. (This is because the amount of memorization, reasonably, also varies according to model size; see also Carlini et al. (2023b).) We include LLAMA 1 13B because it is (again) a similar size to PYTHIA 12B and PHI 4, and because we know the Llama 1 family was trained on `Books3` (Touvron et al., 2023a). And last, we include two larger models as another point of comparison. Again, we pick a Llama 1 model because we know it was trained on `Books3`: LLAMA 1 65B. We then pick a newer and more capable model of a similar size, LLAMA 3.1 70B.

In the future, we will expand this to a much wider set of models. For this version of the project, we ran as many experiments as we could with the compute and people-power available. We are particularly interested in expanding to larger models trained by a wider range of organizations and companies.

Visualizations in the subsections that follow. As discussed above, for each book we include a per-book, cross-model comparison of average extraction rate. This is just to provide a high-level summary of relative extraction across these models. Since we are using the same underlying population of samples, these are valid comparisons, even if we do not expect the overall extraction rates to be high.

We then provide a table with the exact extraction rates that we plot: the greedy-sampled discoverable extraction rate, the (n, p)-discoverable extraction rate for $n = 1000$ and $p = 0.5$ (i.e., the rate of probabilistic extraction with 1000 independent prompts that yields at least a 50%/ coin-flip chance that the example is extracted), and the maximum-possible (n, p)-discoverable extraction rate. (The maximum rate captures both valid instances of memorization, and for very large n, cases that could be argued to not necessarily be reflective entirely of "meaningful" memorization. We set aside the debate about what is "meaningful" to future work.) In this table, we also include the proportion of examples for which $p_{\boldsymbol{z}}$ is greater than 0.000001 (i.e., the probability of the sequence is 1 in 1 million), the proportion of examples for which $p_{\boldsymbol{z}}$ is greater than 0.01 (i.e., the probability of the sequence is 1 in 100), and the runtime of the experiment (using the 4 A100s).

We also include plots (where possible) over the distribution of number of independent prompts n that each extracted example would require in order to be extraction (with probability $p = 0.5$). These distributions show the counts of examples (per model, for each book) that require n independent prompts to attain this level of probabilistic extraction.

Lastly, for some books, we include heatmaps that show the worst-case extraction by location (by character) in the book, in order to visualize where (if any) memorization hot-spots occur. That is, for a given character, for the 100 token sequences in which that character is included, we show the highest possible extraction probability (associated with a given 100-token sequence). Again, the point of this aligns with our overarching goal: visualizing memorization hot-spots, for which this type of worst-case visualization is useful for eliciting how high-probability extraction regions ramp up and ramp down within the book.

Tracking down longer extracted sequences. The results we obtain from these experiments—the probabilities of extraction $p_{\boldsymbol{z}}$ and the associated metadata we save (i.e., the position, in characters, of where each example is located within the book)—is useful for future experiments. Notably, by visualizing memorization "hot-spots," we can use this information to try to identify longer "runs" of memorized text.

The intuition is as follows. We take the full list of $p_{\boldsymbol{z}}$ for all examples for a given book and model experiment, along with their start position (character) in the book) and end position. (They are sorted in order by start position.) We filter the list to only contain especially high probability elements (e.g., $p_{\boldsymbol{z}} > 0.1$), and then iterate through this filtered list and merge together elements that overlap in terms of their position (taking the minimum start position as the overall start, and the maximum end position as the overall end). From here, we can then expect runs of regions in the book that have extractable examples with $p_{\boldsymbol{z}}$ greater than the threshold we filtered by (in this example, 0.1). We can then run extraction experiments on those larger runs, using the start and end locations that we've identified through merging (in character space) adjacent examples that are high-probability.

We attempted this simple methodology for *Harry Potter* (Rowling, 1998) and *Nineteen Eighty-Four* (Orwell, 1949) for LLAMA 3.1 70B, given the extent to which that model memorized those books. There is a lot more work to do in order to try to extract longer sequences, but we provide two examples of successes of this process in Appendix H. One of the examples is in Figure 5, where we show extracting an example that is 300 tokens long, with a 25-token prefix and a 275-token suffix.

Estimating memorization of a whole book. The above identification of high-probability runs can also be used to estimate the proportion of a book that is extractable (with a minimum probability for $p_{\boldsymbol{z}}$). We show instances of the following in

Figure 2 (right) for *Harry Potter* (Rowling, 1998) and *Sandman Slim* (Kadrey, 2009). Again, we pick a minimum threshold we are interested in for $p_{\boldsymbol{z}}$ (in Figure 2, we do this three times, picking $p_{\boldsymbol{z}} \in \{0.01, 0.50, 0.75\}$). We filter the whole (sorted by start position) list of $p_{\boldsymbol{z}}$ to only contain examples whose extraction probability is greater than or equal to that threshold, and we merge (in terms of there location) adjacent examples to find the minimum start position and maximum end position. We then end up with a merged list of "runs" of longer sequences that contain examples where we extracted text with at least $p_{\boldsymbol{z}}$ probability. We can compute the length of these "run" (in terms of total characters), sum them all up, and then divide by the total number of characters in the whole book. This gives us an estimate of the proportion of the whole book (in characters) for which there exist examples $\boldsymbol{z}$ (in tokens) that can be extracted with at least probability $p_{\boldsymbol{z}}$. This doesn't mean we can extract the whole book in one go, or that we've identified the minimal set of examples to try to extract. This is just a procedure for estimating (in a simple way) the proportion of total memorized text.

Discussion. In each subsection, we provide some short remarks about the book. Namely, we clarify why we omit heatmaps for some books. In general, we omit heatmaps if there is essentially no observed extraction for the book (i.e., the heatmaps are empty, showing zero-probability everywhere), or in (some) cases where high-probability extraction regions are only at the very beginning or the very end of the book. We show some of these heatmaps for some books as examples, and we still dig into the details of these high-probability regions. However, to conserve some space, we do not show these for every book that exhibits this pattern. We defer further discussion to Appendix I.

G.1.1. *After I'm Gone*, LIPPMAN

All high-probability extracted sequences are at the beginning of the book. We omit heatmaps (which are mostly empty).

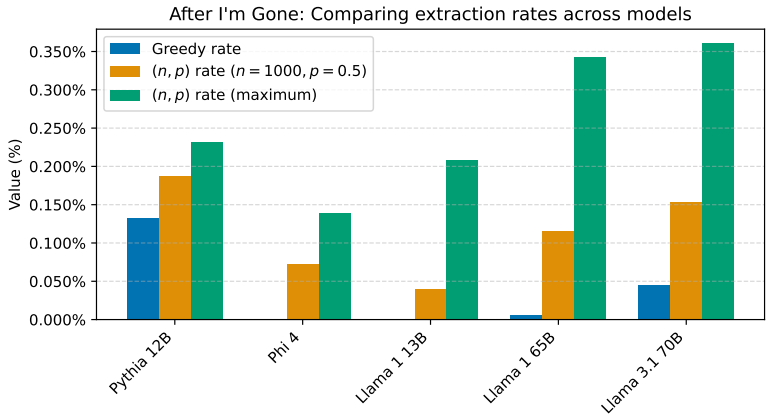

Figure 11: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

After I'm Gone						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.13%	0.19%	0.23%	0.21%	0.17%	00:33:12.10
Phi 4 (14B)	0.00%	0.07%	0.14%	0.11%	0.01%	01:06:48.44
Llama 1 13B	0.00%	0.04%	0.21%	0.10%	0.03%	00:24:18.57
Llama 1 65B	0.01%	0.12%	0.34%	0.18%	0.06%	01:05:17.58
Llama 3.1 70B	0.04%	0.15%	0.36%	0.22%	0.13%	01:58:42.71

Table 6: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

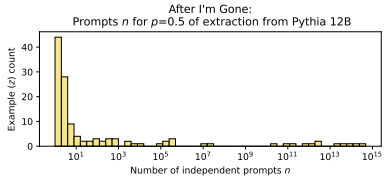

(a) Pythia 12B

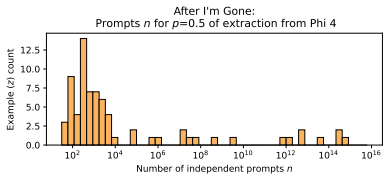

(b) Phi 4 (14B)

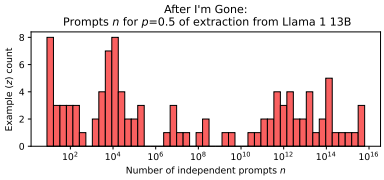

(c) Llama 1 13B

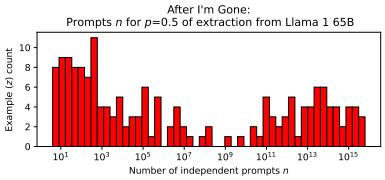

(d) Llama 1 65B

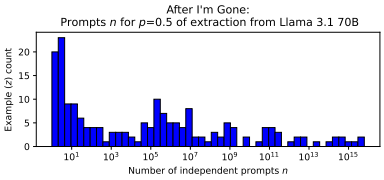

(e) Llama 3.1 70B

Figure 12: *After I'm Gone*, Lippman (2014): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.2. *Alice's Adventures in Wonderland*, CARROLL

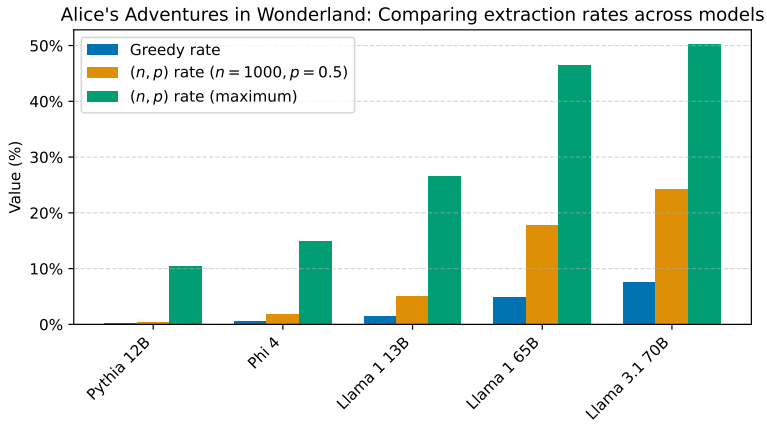

Figure 13: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Alice's Adventures in Wonderland						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.16%	0.43%	10.40%	1.36%	0.25%	00:42:22.36
Phi 4 (14B)	0.65%	1.93%	14.90%	5.07%	1.10%	01:24:06.58
Llama 1 13B	1.49%	5.13%	26.56%	11.30%	2.84%	00:32:19.81
Llama 1 65B	4.96%	17.79%	46.52%	30.49%	11.34%	01:26:03.29
Llama 3.1 70B	7.49%	24.25%	50.30%	37.38%	16.93%	02:26:16.45

Table 7: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

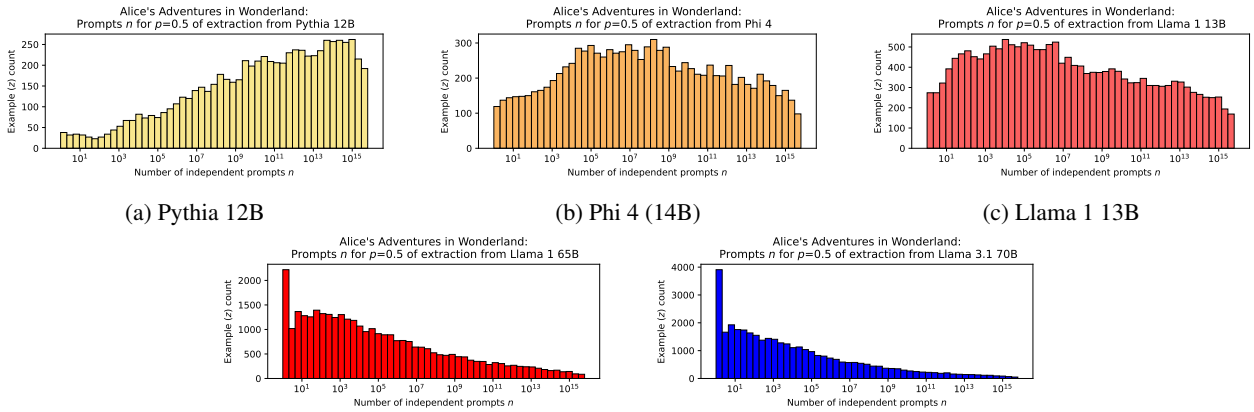

(a) Pythia 12B (b) Phi 4 (14B) (c) Llama 1 13B

(d) Llama 1 65B (e) Llama 3.1 70B

Figure 14: *Alice's Adventures in Wonderland*, Carroll (1865): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

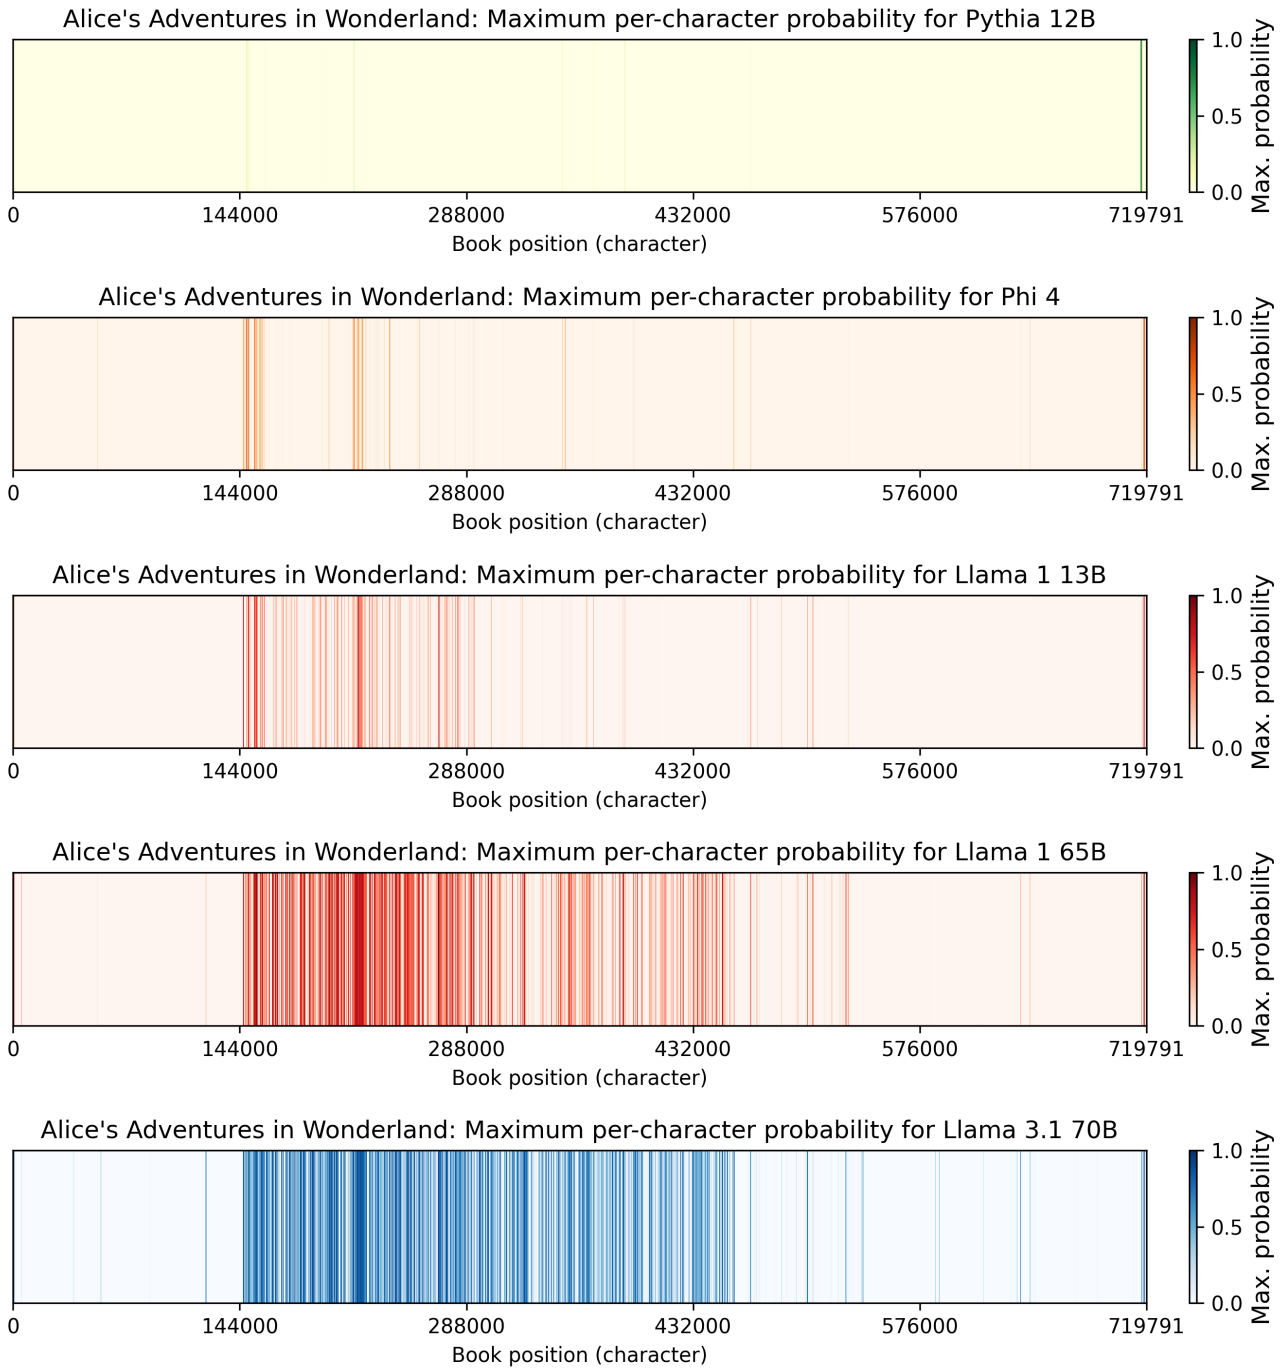

Figure 15: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.3. *Another Brooklyn*, WOODSON

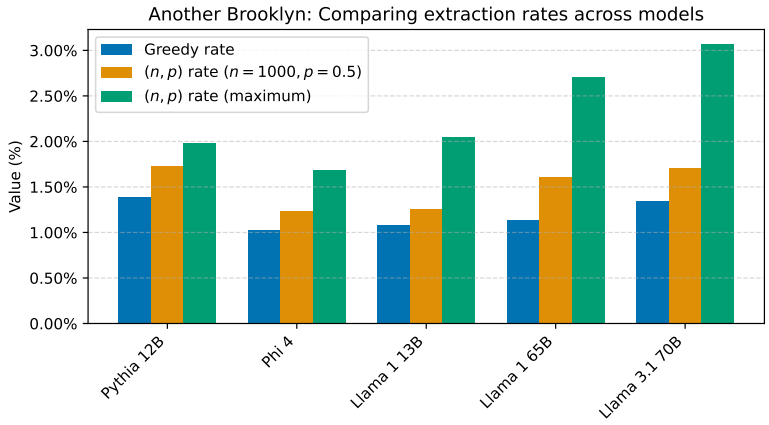

Figure 16: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Another Brooklyn						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	1.39%	1.72%	1.98%	1.91%	1.68%	00:04:24.36
Phi 4 (14B)	1.03%	1.24%	1.69%	1.36%	1.12%	00:09:55.09
Llama 1 13B	1.09%	1.26%	2.05%	1.57%	1.17%	00:04:19.46
Llama 1 65B	1.13%	1.60%	2.71%	1.93%	1.33%	00:15:29.82
Llama 3.1 70B	1.34%	1.71%	3.08%	2.15%	1.58%	00:34:06.59

Table 8: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

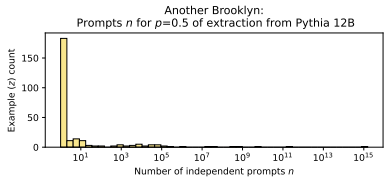

(a) Pythia 12B

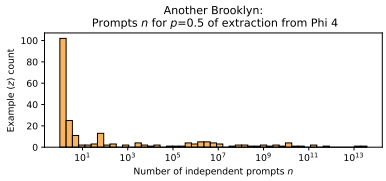

(b) Phi 4 (14B)

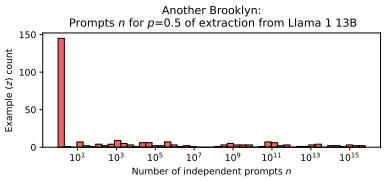

(c) Llama 1 13B

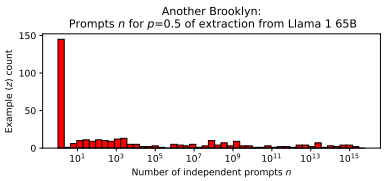

(d) Llama 1 65B

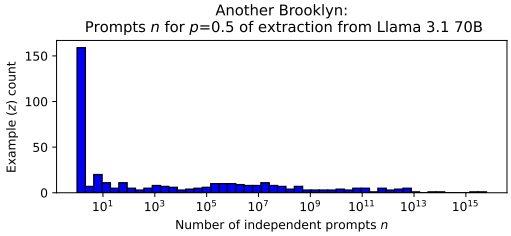

(e) Llama 3.1 70B

Figure 17: *Another Brooklyn*, Woodson (2016): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

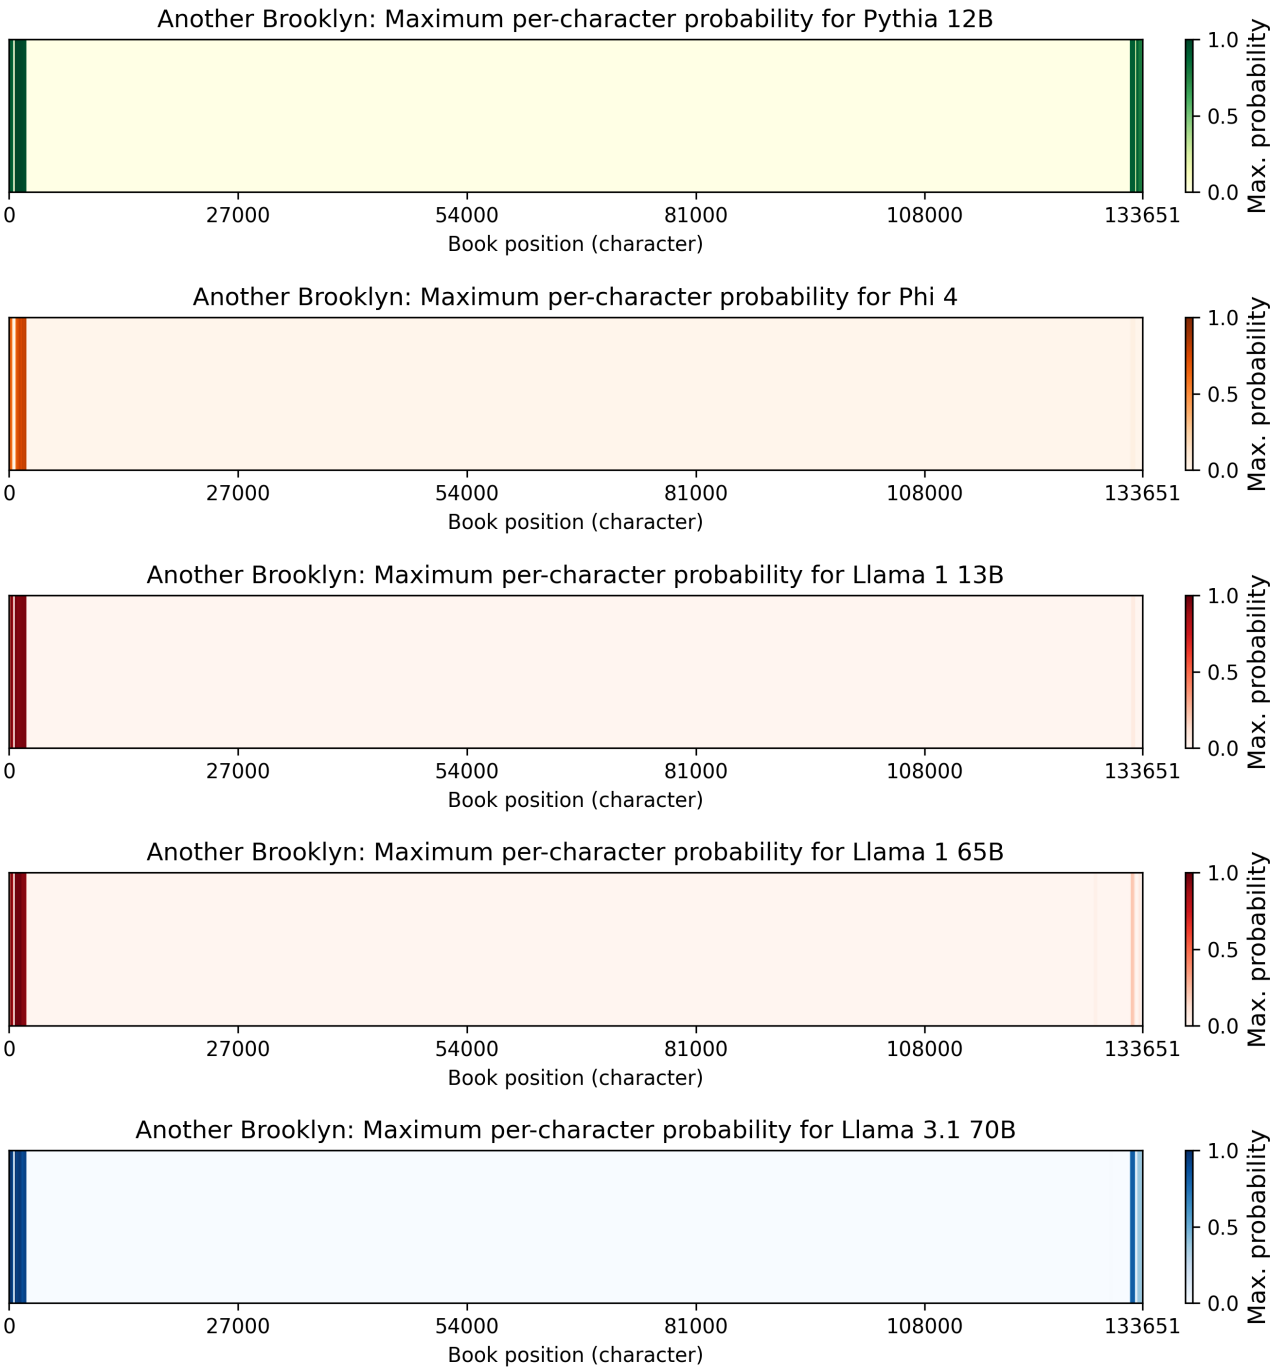

Figure 18: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.4. *Ararat*, GOLDEN

All high-probability extracted sequences are at the beginning of the book. We omit heatmaps (which are mostly empty).

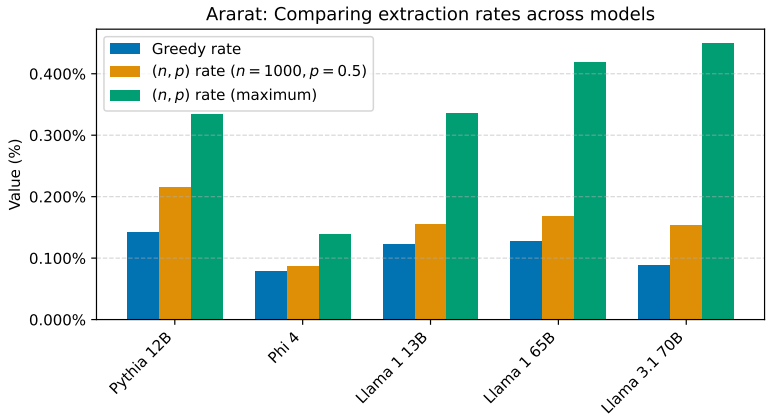

Figure 19: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Ararat						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.14%	0.22%	0.33%	0.28%	0.17%	00:30:36.62
Phi 4 (14B)	0.08%	0.09%	0.14%	0.11%	0.18%	01:02:21.74
Llama 1 13B	0.12%	0.16%	0.34%	0.18%	0.14%	00:22:18.90
Llama 1 65B	0.13%	0.17%	0.42%	0.23%	0.15%	01:02:33.58
Llama 3.1 70B	0.09%	0.15%	0.45%	0.23%	0.13%	01:50:43.33

Table 9: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

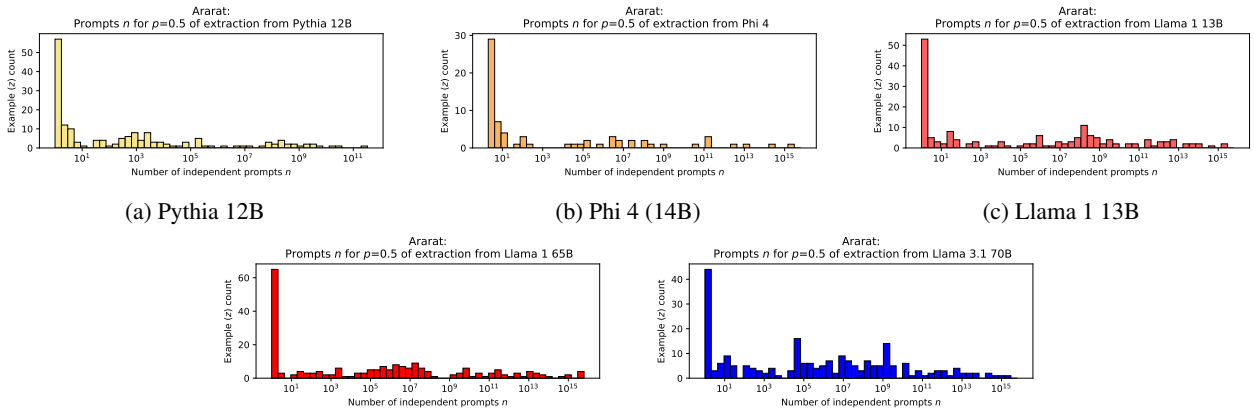

(a) Pythia 12B (b) Phi 4 (14B) (c) Llama 1 13B

(d) Llama 1 65B (e) Llama 3.1 70B

Figure 20: *Ararat*, Golden (2017): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.5. *Brown Girl Dreaming*, WOODSON

All high-probability extracted sequences are at the beginning of the book. We omit heatmaps (which are mostly empty).

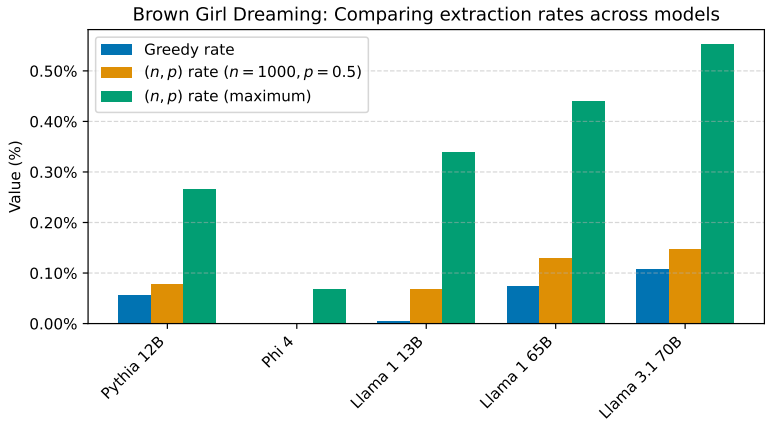

Figure 21: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Brown Girl Dreaming						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.06%	0.08%	0.27%	0.18%	0.06%	00:10:39.32
Phi 4 (14B)	0.00%	0.00%	0.07%	0.03%	0.00%	00:23:42.67
Llama 1 13B	0.01%	0.07%	0.34%	0.13%	0.04%	00:08:01.68
Llama 1 65B	0.07%	0.13%	0.44%	0.23%	0.08%	00:22:08.44
Llama 3.1 70B	0.11%	0.15%	0.55%	0.18%	0.11%	00:47:38.86

Table 10: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

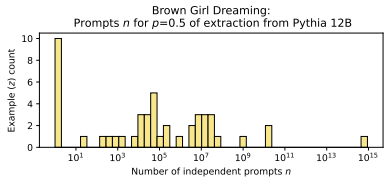

(a) Pythia 12B

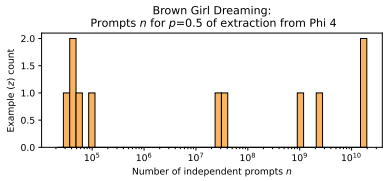

(b) Phi 4 (14B)

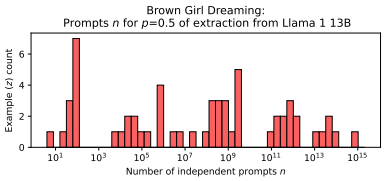

(c) Llama 1 13B

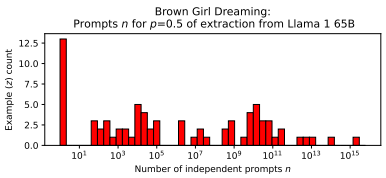

(d) Llama 1 65B

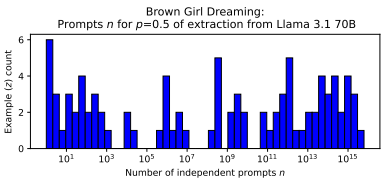

(e) Llama 3.1 70B

Figure 22: *Brown Girl Dreaming*, Woodson (2014): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.6. *Dante and the Origins of Italian Culture*, BAROLINI

There are some low-probability sequences that contribute to a high maximum (n, p)-discoverable extraction rate. We omit heatmaps since these sequences are not very visible in the color gradient, but we do investigate these specific sequences in Appendix H.

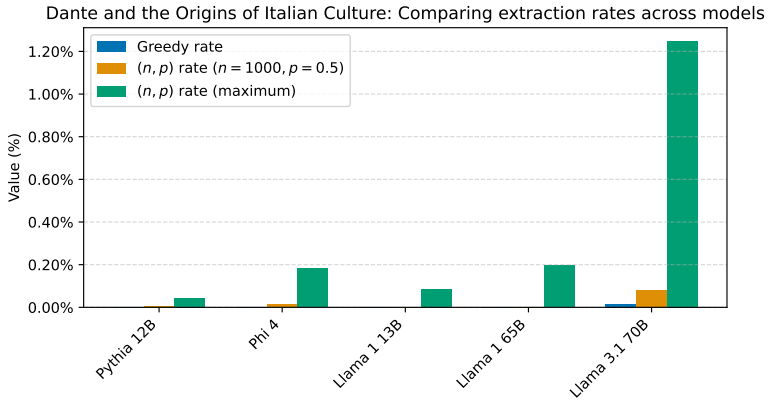

Figure 23: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Dante and the Origins of Italian Culture						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.00%	0.01%	0.04%	0.01%	0.01%	01:09:53.43
Phi 4 (14B)	0.00%	0.01%	0.18%	0.03%	0.00%	02:18:03.72
Llama 1 13B	0.00%	0.00%	0.09%	0.01%	0.00%	00:50:57.34
Llama 1 65B	0.00%	0.00%	0.20%	0.01%	0.00%	02:20:26.40
Llama 3.1 70B	0.02%	0.08%	1.25%	0.22%	0.03%	03:55:26.69

Table 11: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

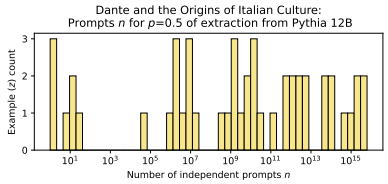

(a) Pythia 12B

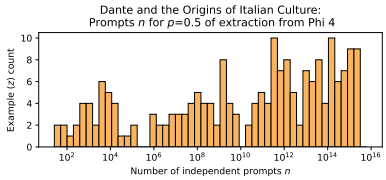

(b) Phi 4 (14B)

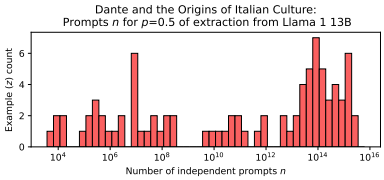

(c) Llama 1 13B

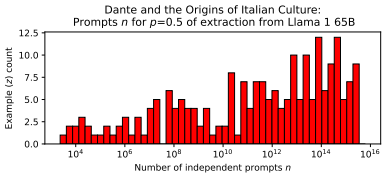

(d) Llama 1 65B

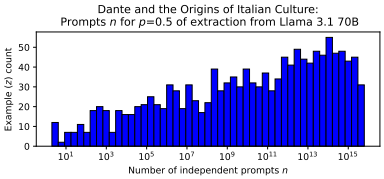

(e) Llama 3.1 70B

Figure 24: *Dante and the Origins of Italian Culture*, Barolini (2006): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.7. *Dead Ringer*, GOLDEN

All high-probability extracted sequences are at the beginning of the book. We omit heatmaps (which are mostly empty).

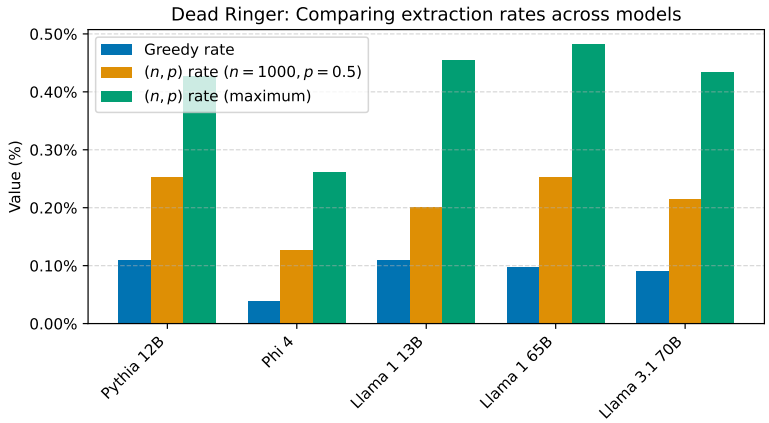

Figure 25: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Dead Ringer						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.11%	0.25%	0.43%	0.34%	0.19%	00:16:42.57
Phi 4 (14B)	0.04%	0.13%	0.26%	0.19%	0.08%	00:30:56.08
Llama 1 13B	0.11%	0.20%	0.45%	0.28%	0.15%	00:17:39.28
Llama 1 65B	0.10%	0.25%	0.48%	0.36%	0.22%	00:58:41.88
Llama 3.1 70B	0.09%	0.22%	0.43%	0.32%	0.16%	01:28:54.09

Table 12: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

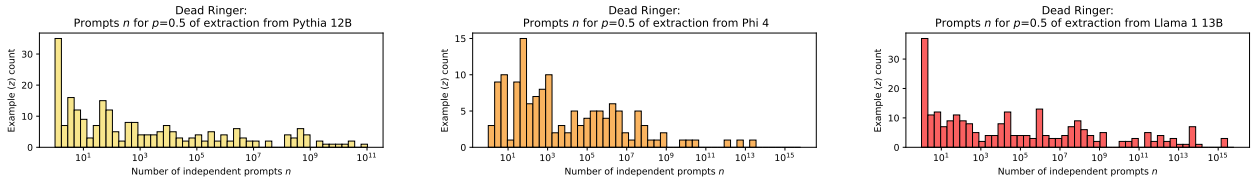

(a) Pythia 12B

(b) Phi 4 (14B)

(c) Llama 1 13B

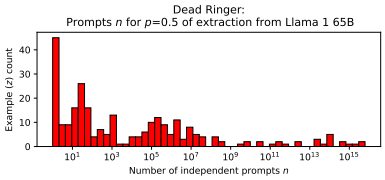

(d) Llama 1 65B

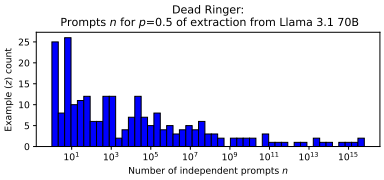

(e) Llama 3.1 70B

Figure 26: *Dead Ringer*, Golden (2016): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.8. *Down and Out in the Magic Kingdom*, DOCTOROW

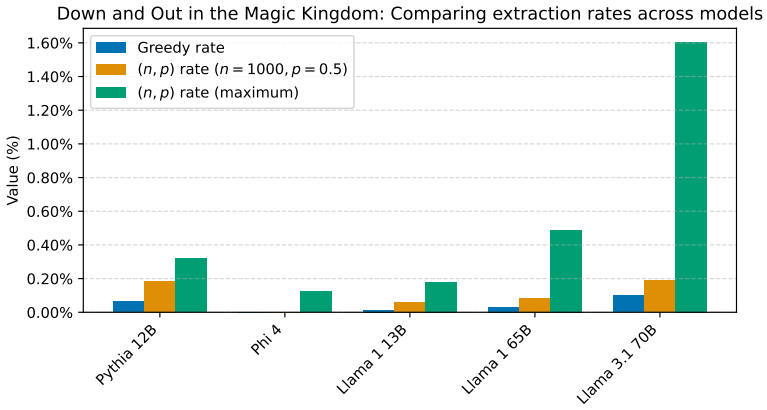

Figure 27: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Down and Out in the Magic Kingdom						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.06%	0.19%	0.32%	0.32%	0.12%	00:16:40.70
Phi 4 (14B)	0.00%	0.00%	0.12%	0.00%	0.00%	00:35:27.50
Llama 1 13B	0.01%	0.06%	0.18%	0.08%	0.03%	00:12:37.05
Llama 1 65B	0.03%	0.09%	0.49%	0.17%	0.07%	00:33:25.06
Llama 3.1 70B	0.10%	0.19%	1.61%	0.43%	0.13%	01:07:30.19

Table 13: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

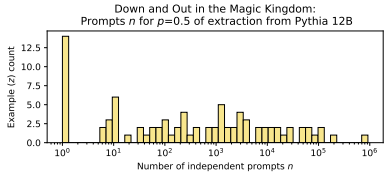

(a) Pythia 12B

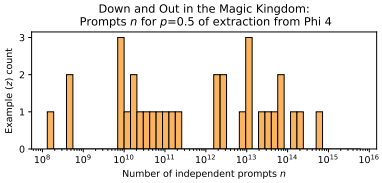

(b) Phi 4 (14B)

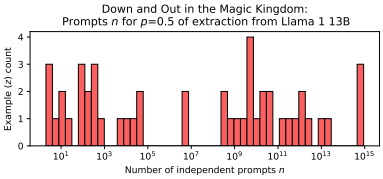

(c) Llama 1 13B

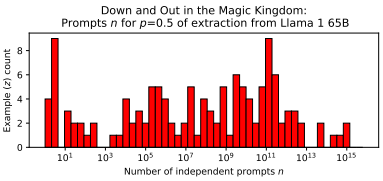

(d) Llama 1 65B

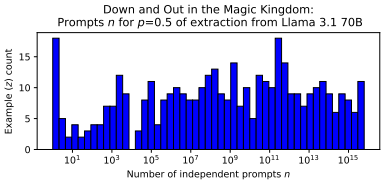

(e) Llama 3.1 70B

Figure 28: *Down and Out in the Magic Kingdom*, Doctorow (2003): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

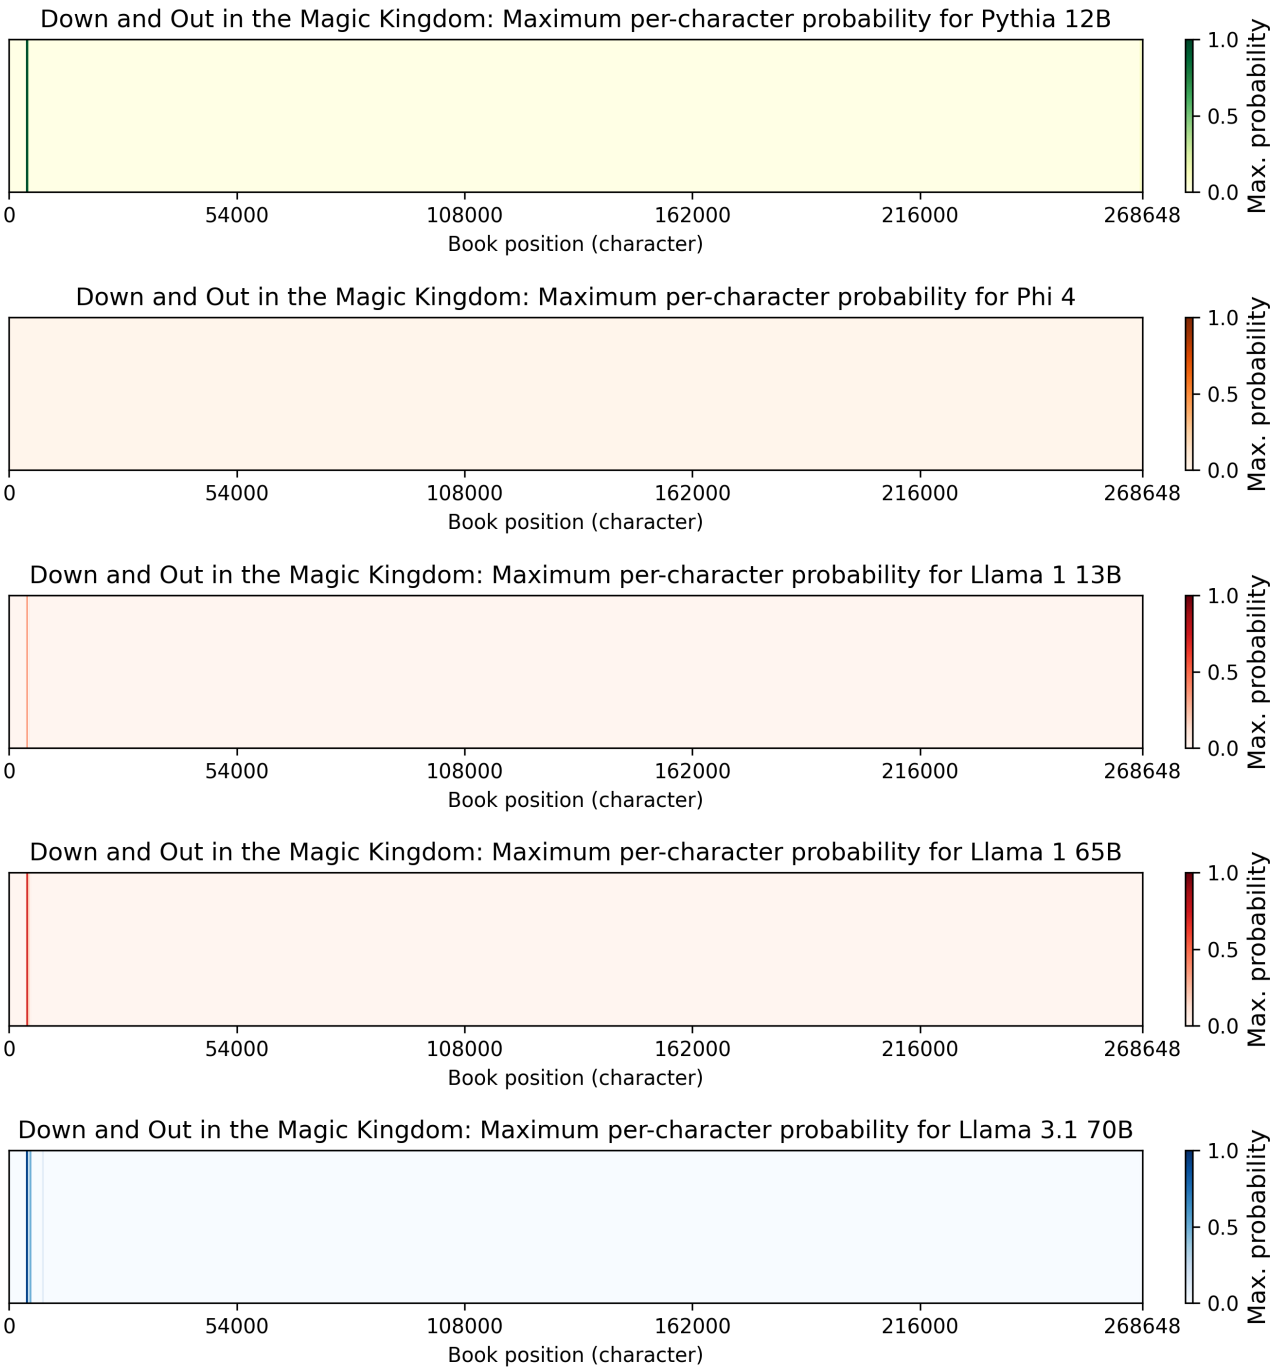

Figure 29: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.9. *Drown*, DÍAZ

All high-probability extracted sequences are at the beginning of the book. We omit heatmaps (which are mostly empty). This is one of the least extractable books we tested; with our extraction procedure, it appears to be one of the least memorized.

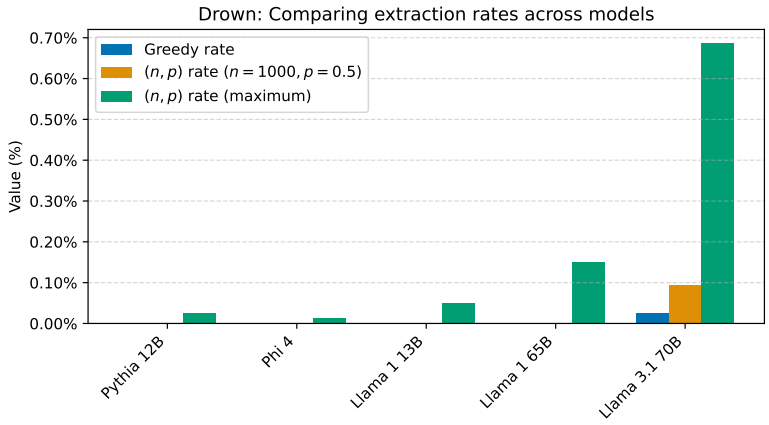

Figure 30: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Drown						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.00%	0.00%	0.02%	0.00%	0.00%	00:15:39.60
Phi 4 (14B)	0.00%	0.00%	0.01%	0.00%	0.00%	00:31:49.62
Llama 1 13B	0.00%	0.00%	0.05%	0.00%	0.00%	00:11:21.27
Llama 1 65B	0.00%	0.00%	0.15%	0.00%	0.00%	00:30:33.16
Llama 3.1 70B	0.02%	0.09%	0.69%	0.12%	0.07%	01:02:16.72

Table 14: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

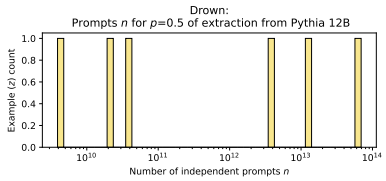

(a) Pythia 12B

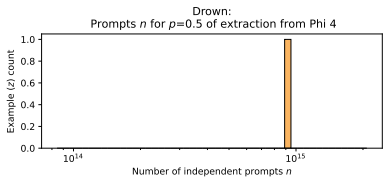

(b) Phi 4 (14B)

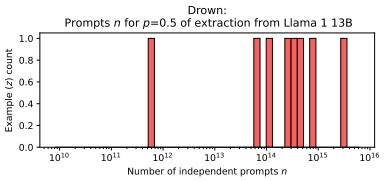

(c) Llama 1 13B

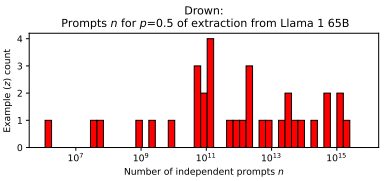

(d) Llama 1 65B

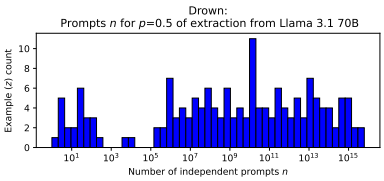

(e) Llama 3.1 70B

Figure 31: *Drown*, Díaz (1996): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.10. *Dungeons and Dragons and Philosophy*, COGBURN

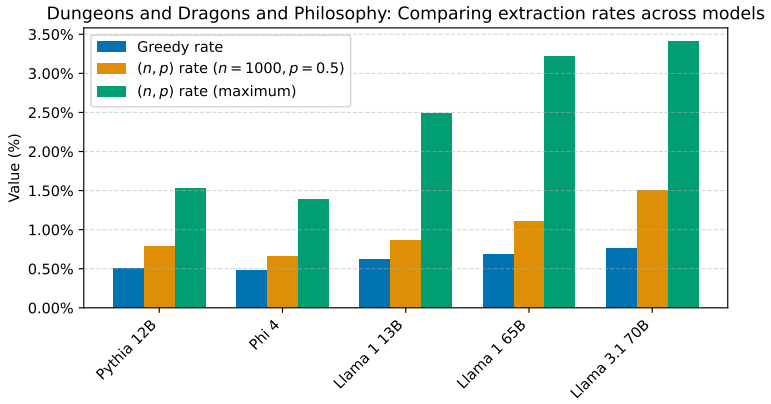

Figure 32: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Dungeons and Dragons and Philosophy						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.51%	0.79%	1.53%	0.94%	0.66%	00:13:24.97
Phi 4 (14B)	0.47%	0.66%	1.39%	0.88%	0.58%	00:25:07.75
Llama 1 13B	0.62%	0.86%	2.49%	1.14%	0.74%	00:13:22.00
Llama 1 65B	0.68%	1.10%	3.21%	1.59%	0.87%	00:46:01.07
Llama 3.1 70B	0.76%	1.50%	3.41%	2.28%	1.15%	01:13:49.70

Table 15: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

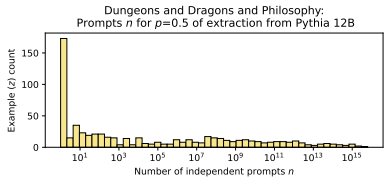

(a) Pythia 12B

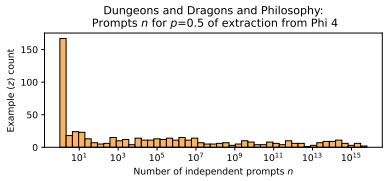

(b) Phi 4 (14B)

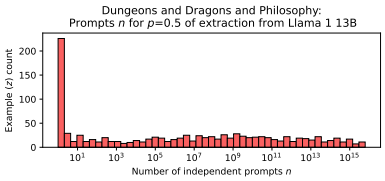

(c) Llama 1 13B

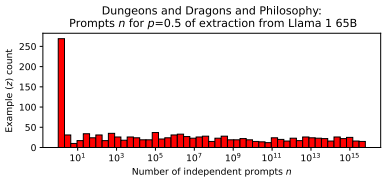

(d) Llama 1 65B

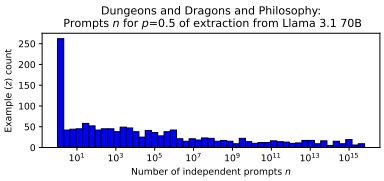

(e) Llama 3.1 70B

Figure 33: *Dungeons and Dragons and Philosophy*, Cogburn (2012): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

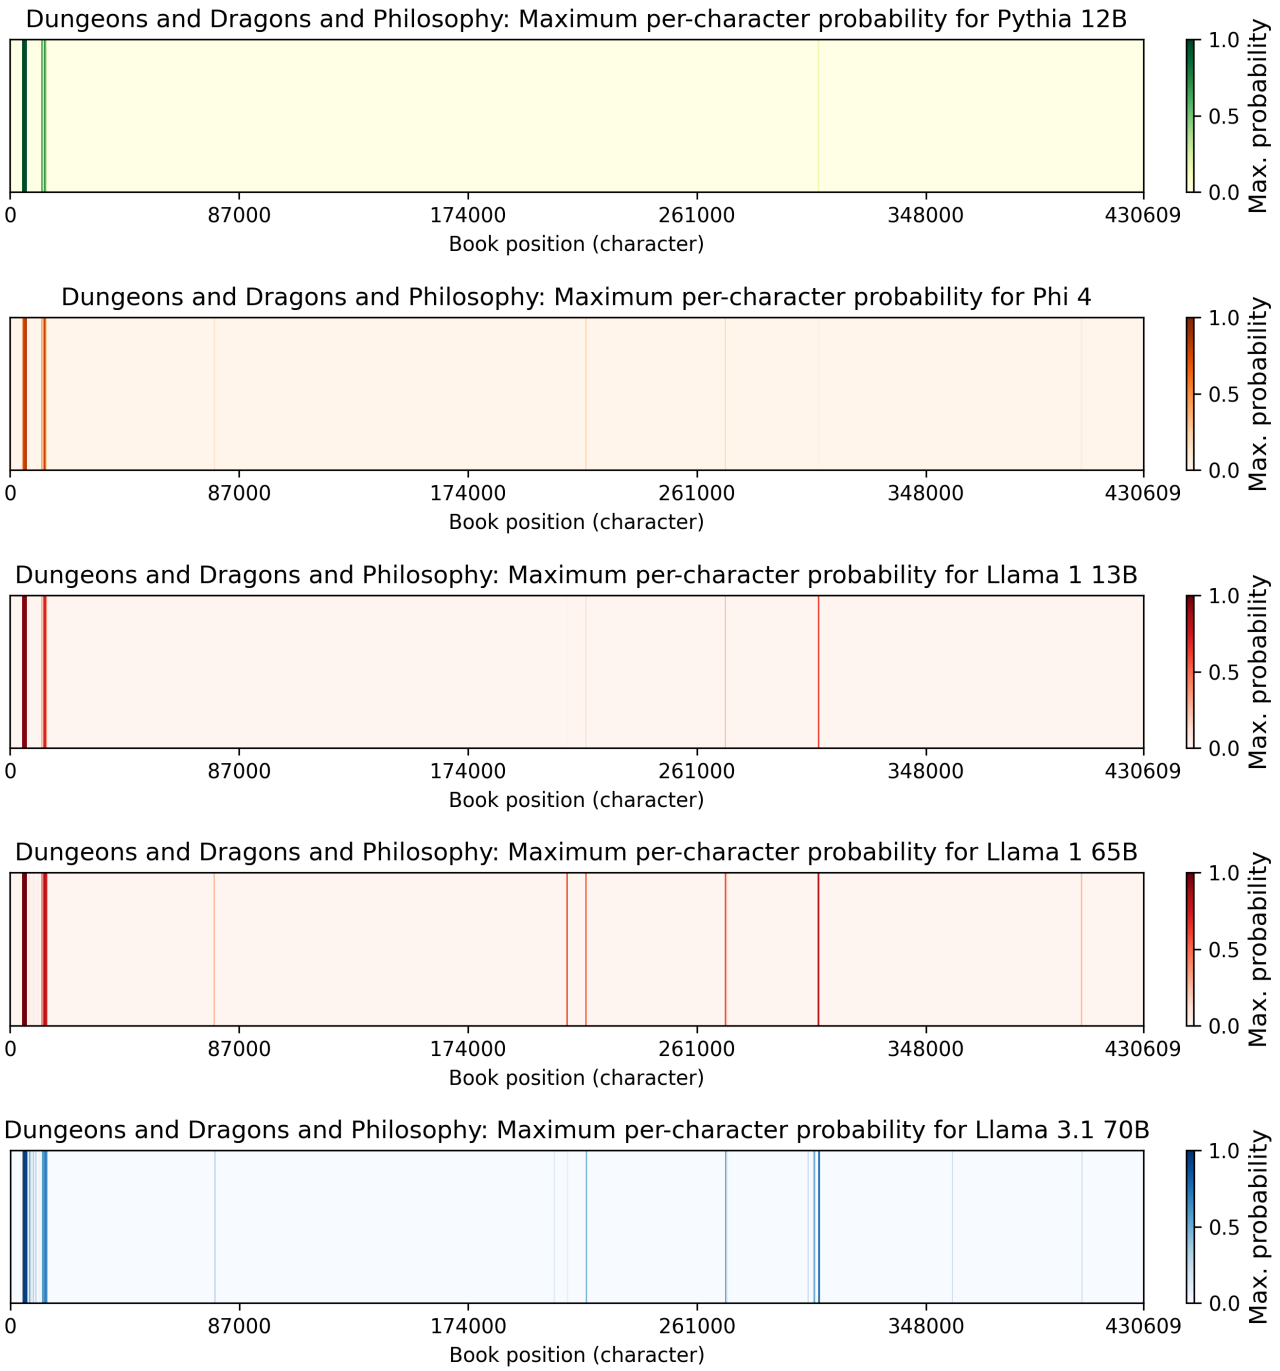

Figure 34: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.11. *Embraced*, TERKEURST

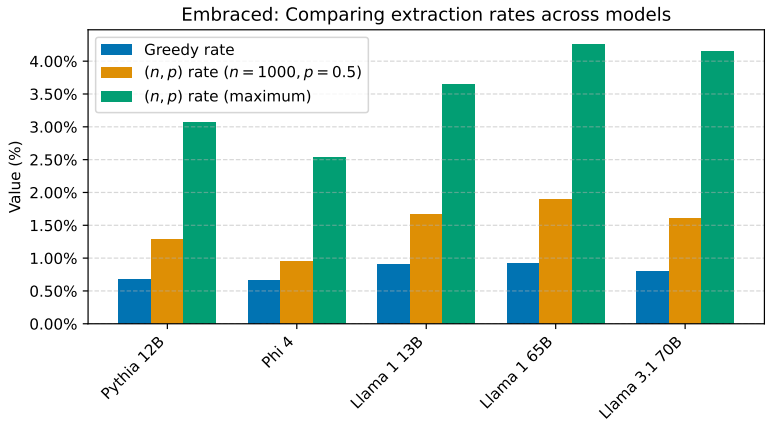

Figure 35: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

			Embraced			
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.68%	1.29%	3.07%	1.98%	0.92%	00:25:42.20
Phi 4 (14B)	0.66%	0.95%	2.53%	1.28%	0.83%	00:49:09.94
Llama 1 13B	0.91%	1.67%	3.65%	2.32%	1.27%	00:19:19.73
Llama 1 65B	0.92%	1.90%	4.26%	2.58%	1.44%	00:52:07.03
Llama 3.1 70B	0.80%	1.61%	4.16%	2.28%	1.19%	01:36:54.84

Table 16: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

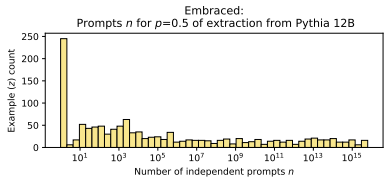

(a) Pythia 12B

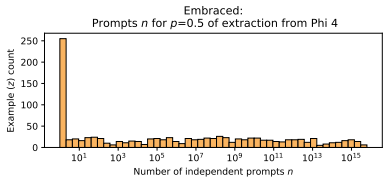

(b) Phi 4 (14B)

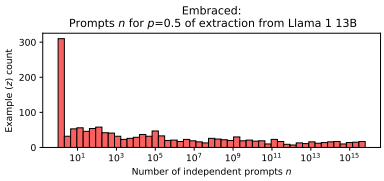

(c) Llama 1 13B

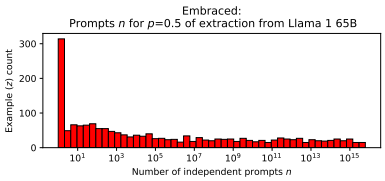

(d) Llama 1 65B

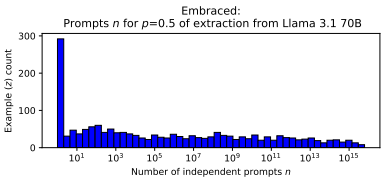

(e) Llama 3.1 70B

Figure 36: *Embraced*, TerKeurst (2018): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

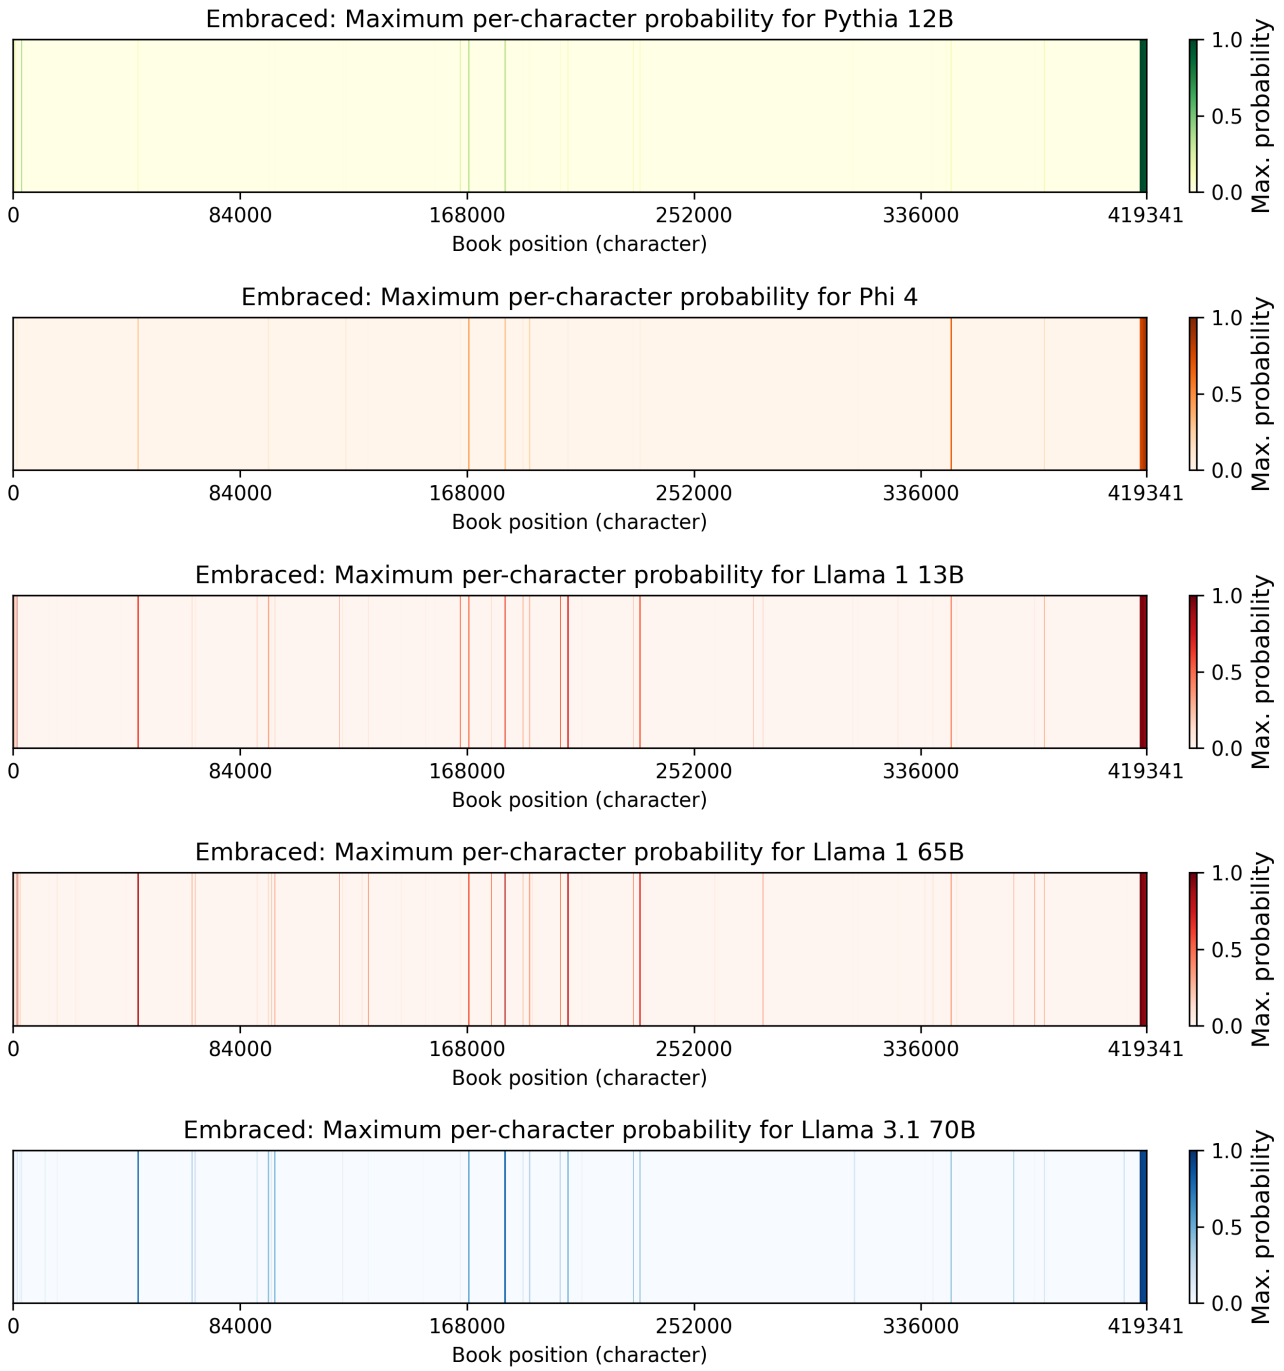

Figure 37: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.12. *Harry Potter and the Sorcerer's Stone*, ROWLING

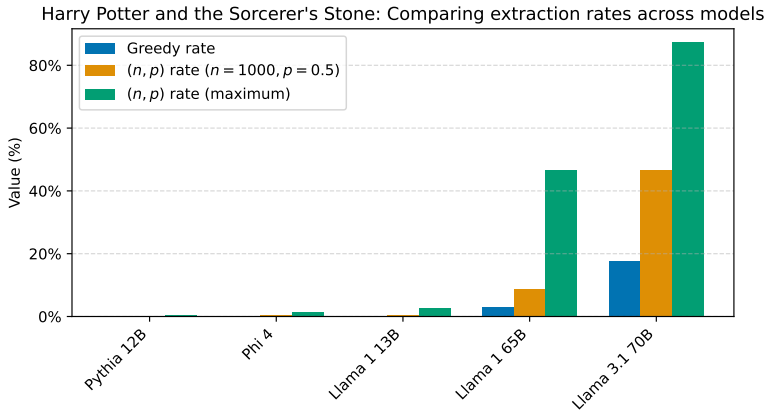

Figure 38: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Harry Potter and the Sorcerer's Stone						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.04%	0.08%	0.39%	0.13%	0.04%	00:31:59.52
Phi 4 (14B)	0.28%	0.58%	1.59%	0.87%	0.39%	00:57:48.55
Llama 1 13B	0.14%	0.44%	2.86%	1.01%	0.23%	00:20:10.69
Llama 1 65B	2.95%	8.84%	46.63%	18.30%	5.35%	00:56:26.10
Llama 3.1 70B	17.66%	46.67%	87.26%	68.17%	33.89%	01:47:21.85

Table 17: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

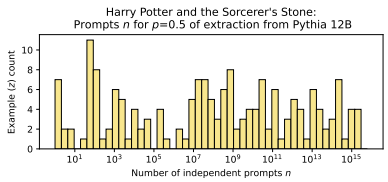

(a) Pythia 12B

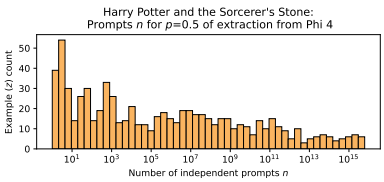

(b) Phi 4 (14B)

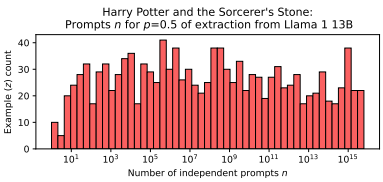

(c) Llama 1 13B

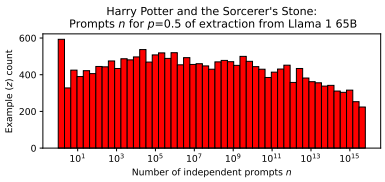

(d) Llama 1 65B

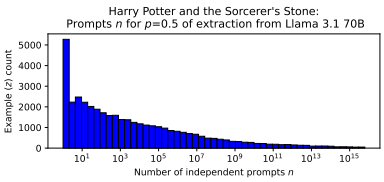

(e) Llama 3.1 70B

Figure 39: *Harry Potter and the Sorcerer's Stone*, Rowling (1998): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

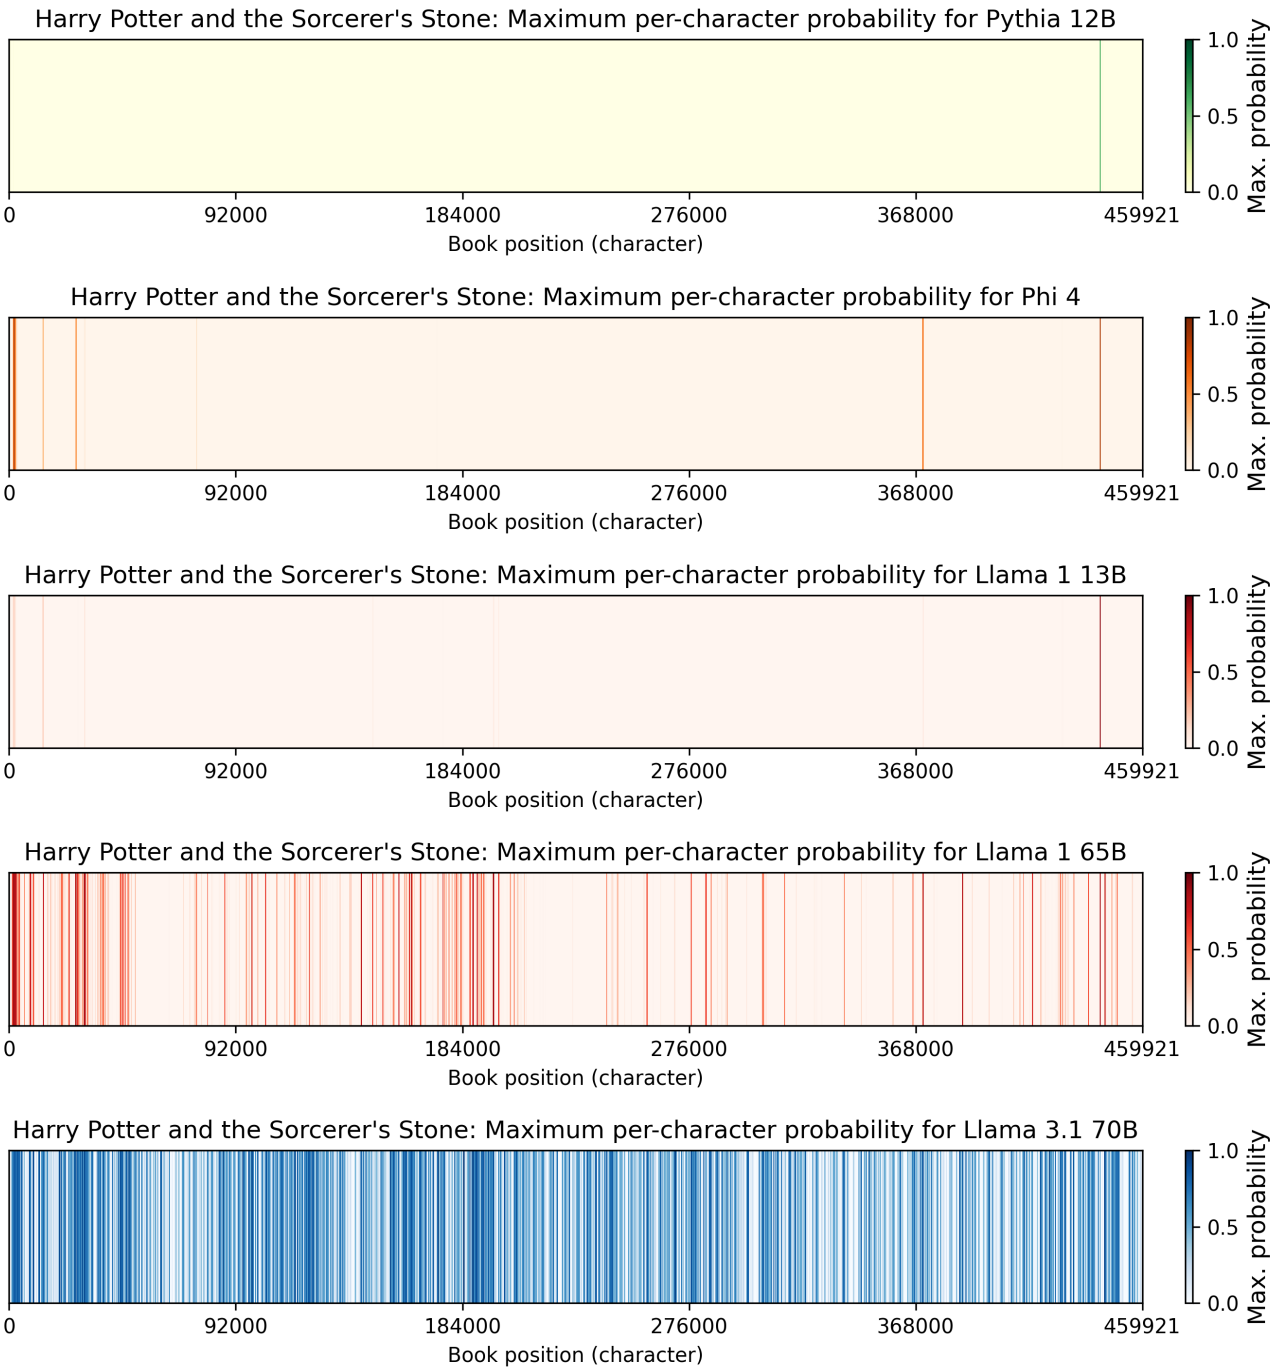

Figure 40: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.13. *The Future of the Internet and How to Stop It*, ZITTRAIN

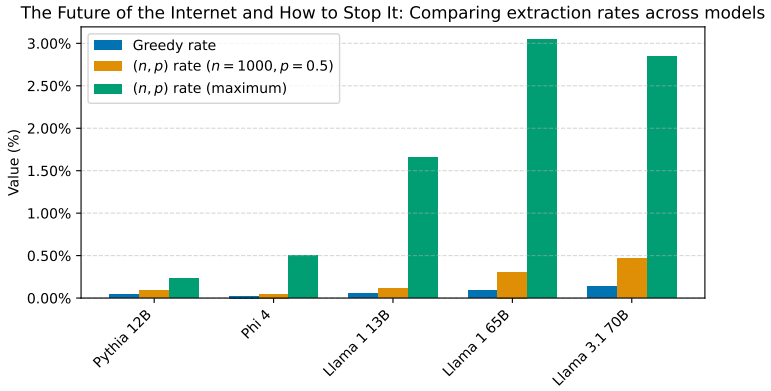

Figure 41: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Future of the Internet and How to Stop It						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.04%	0.09%	0.23%	0.12%	0.07%	00:58:48.11
Phi 4 (14B)	0.02%	0.05%	0.50%	0.13%	0.03%	01:53:41.34
Llama 1 13B	0.05%	0.11%	1.66%	0.25%	0.06%	00:42:27.92
Llama 1 65B	0.09%	0.31%	3.04%	0.59%	0.19%	01:51:51.40
Llama 3.1 70B	0.14%	0.47%	2.85%	0.89%	0.30%	03:19:03.10

Table 18: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

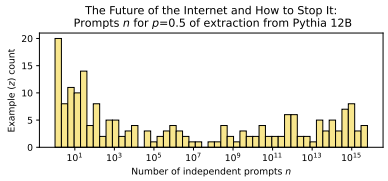

(a) Pythia 12B

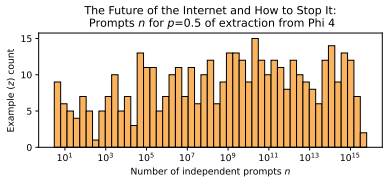

(b) Phi 4 (14B)

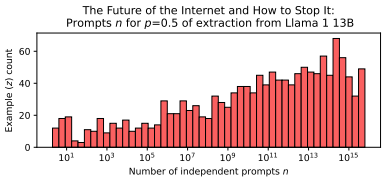

(c) Llama 1 13B

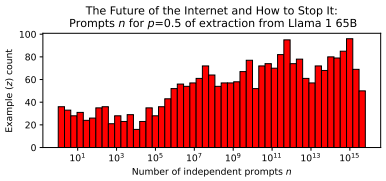

(d) Llama 1 65B

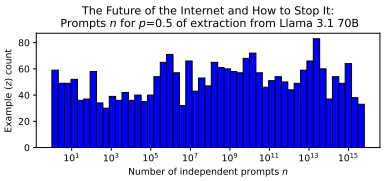

(e) Llama 3.1 70B

Figure 42: *The Future of the Internet and How to Stop It*, Zittrain (2008): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

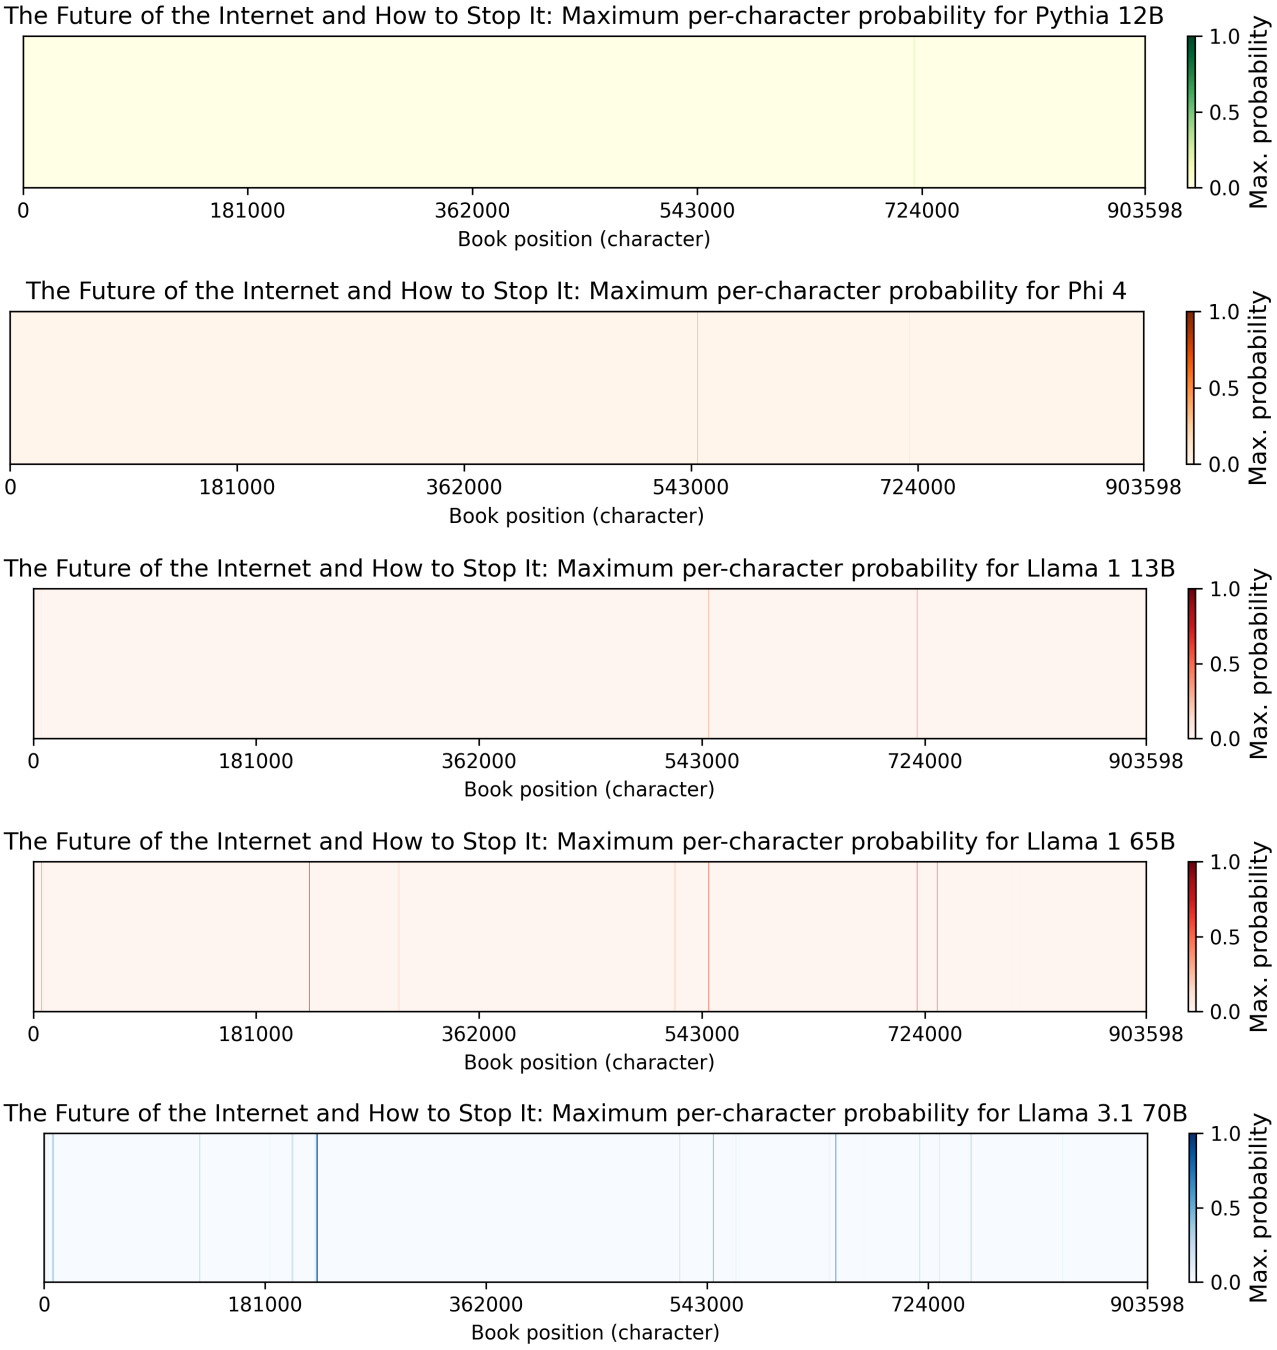

Figure 43: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.14. *M. Butterfly*, HWANG

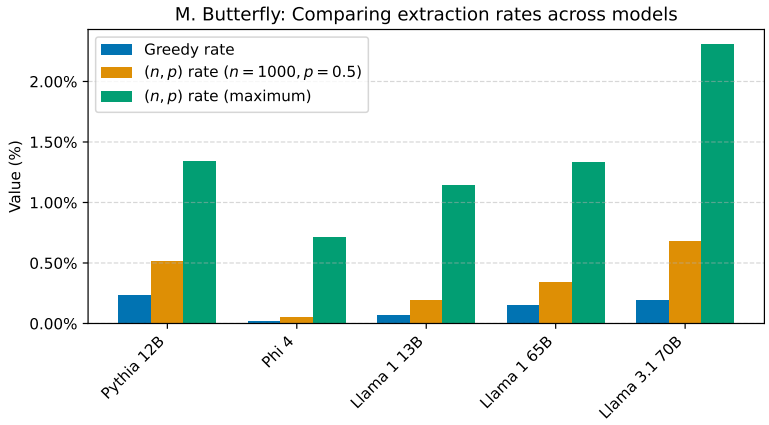

Figure 44: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

M. Butterfly						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.23%	0.52%	1.34%	1.18%	0.37%	00:08:34.80
Phi 4 (14B)	0.02%	0.06%	0.71%	0.06%	0.06%	00:18:57.56
Llama 1 13B	0.07%	0.19%	1.14%	0.41%	0.12%	00:06:16.90
Llama 1 65B	0.15%	0.34%	1.34%	0.86%	0.21%	00:16:33.55
Llama 3.1 70B	0.19%	0.68%	2.31%	1.35%	0.34%	00:39:22.20

Table 19: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

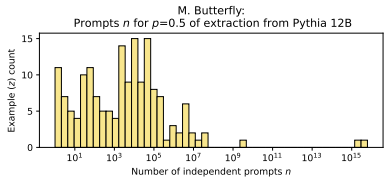

(a) Pythia 12B

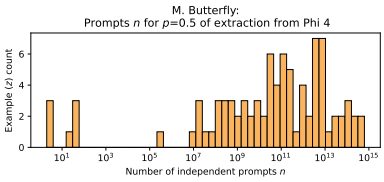

(b) Phi 4 (14B)

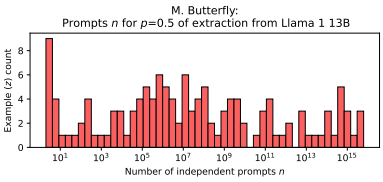

(c) Llama 1 13B

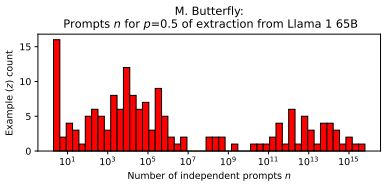

(d) Llama 1 65B

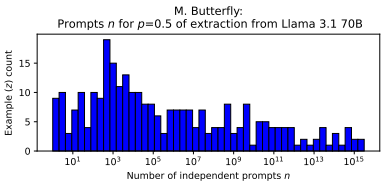

(e) Llama 3.1 70B

Figure 45: *M. Butterfly*, Hwang (1988): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

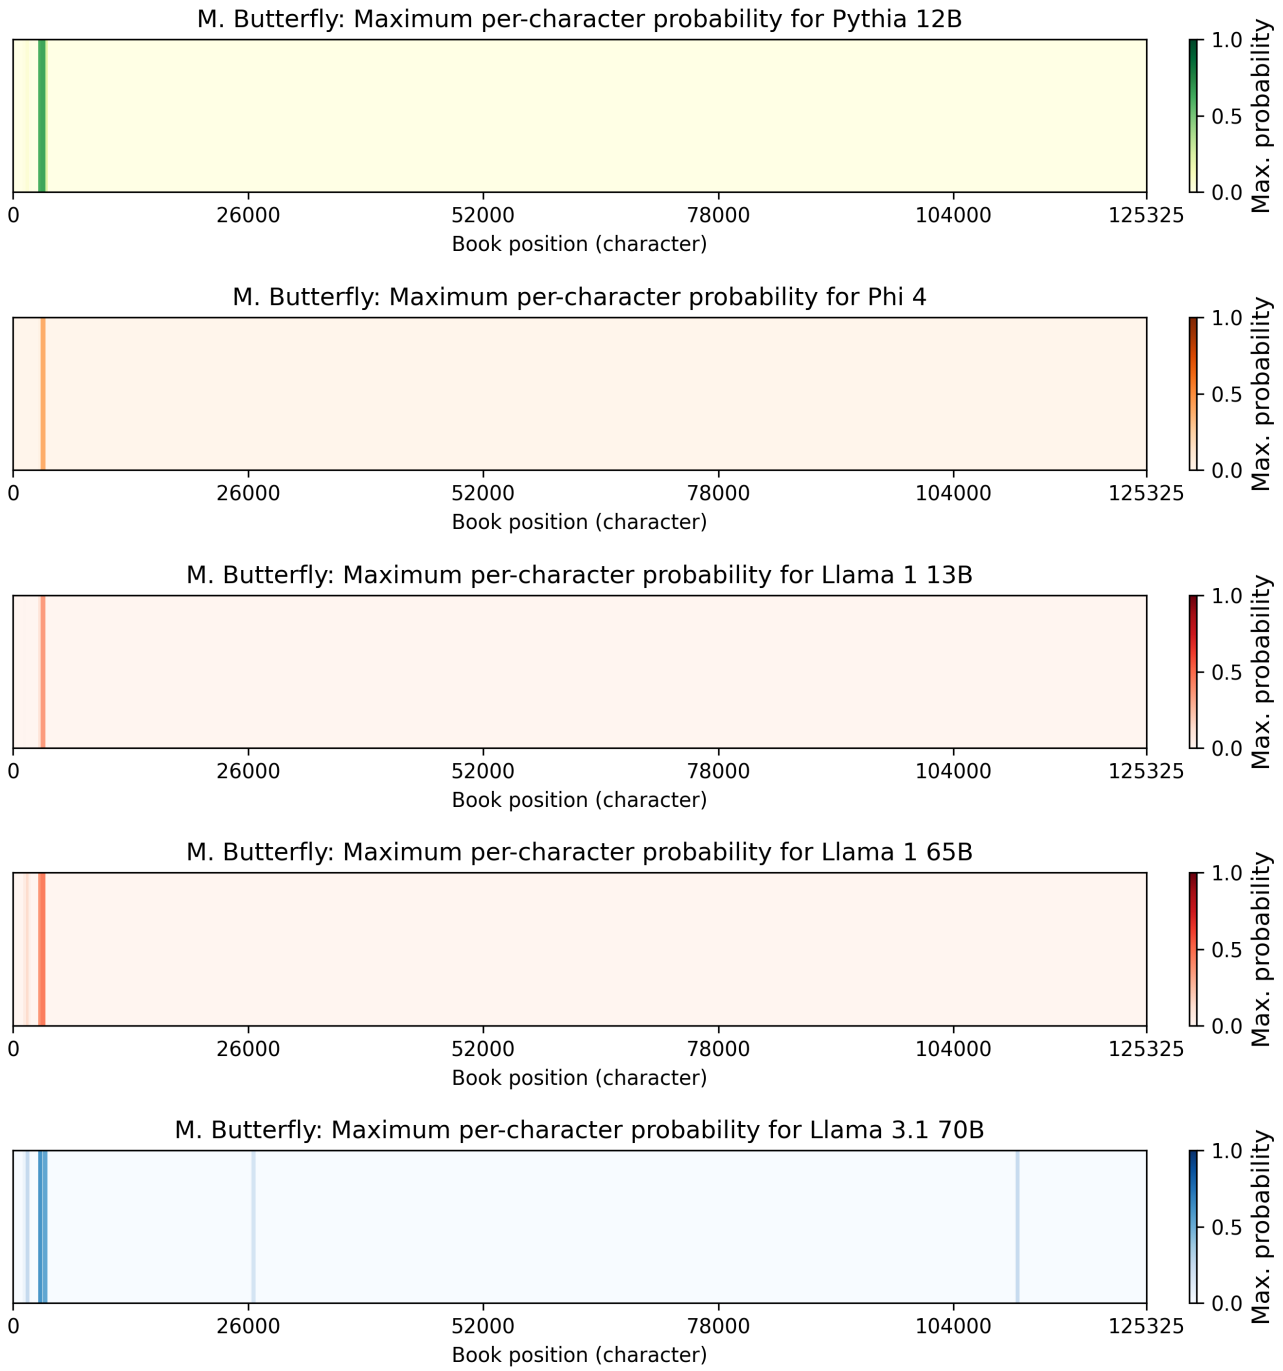

Figure 46: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.15. *No Visible Bruises*, SNYDER

All high-probability extracted sequences are at the end of the book. We omit heatmaps (which are mostly empty).

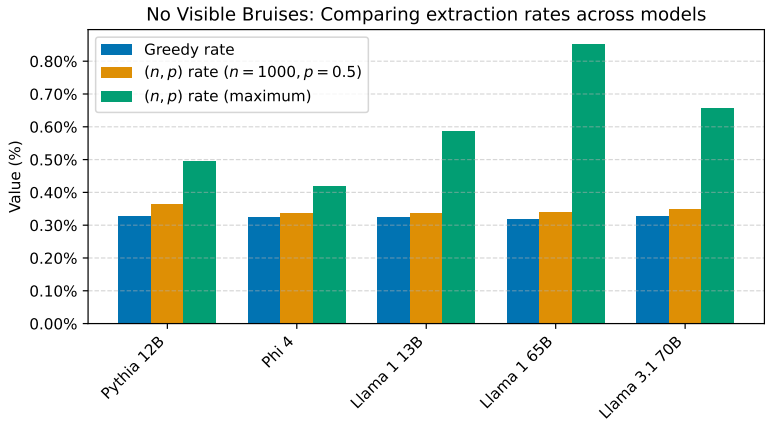

Figure 47: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

No Visible Bruises						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.33%	0.36%	0.50%	0.44%	0.34%	00:49:33.08
Phi 4 (14B)	0.32%	0.34%	0.42%	0.36%	0.33%	01:36:16.77
Llama 1 13B	0.32%	0.34%	0.59%	0.38%	0.33%	00:33:52.32
Llama 1 65B	0.32%	0.34%	0.85%	0.86%	0.33%	01:30:24.99
Llama 3.1 70B	0.33%	0.35%	0.66%	0.34%	0.40%	02:41:25.61

Table 20: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

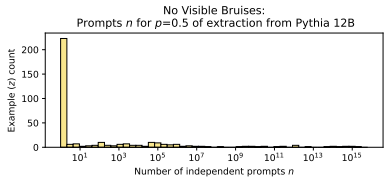

(a) Pythia 12B

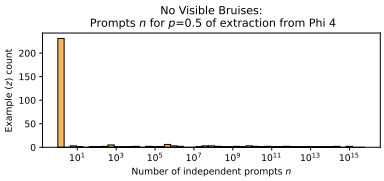

(b) Phi 4 (14B)

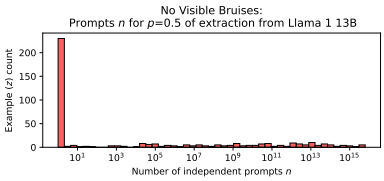

(c) Llama 1 13B

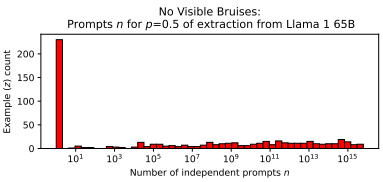

(d) Llama 1 65B

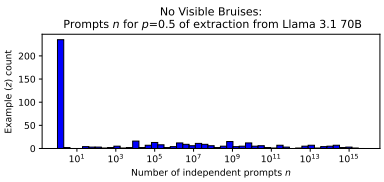

(e) Llama 3.1 70B

Figure 48: *No Visible Bruises*, Snyder (2019): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.16. *Sandman Slim*, KADREY

All high-probability extracted sequences are at the end of the book. We omit heatmaps (which are mostly empty).

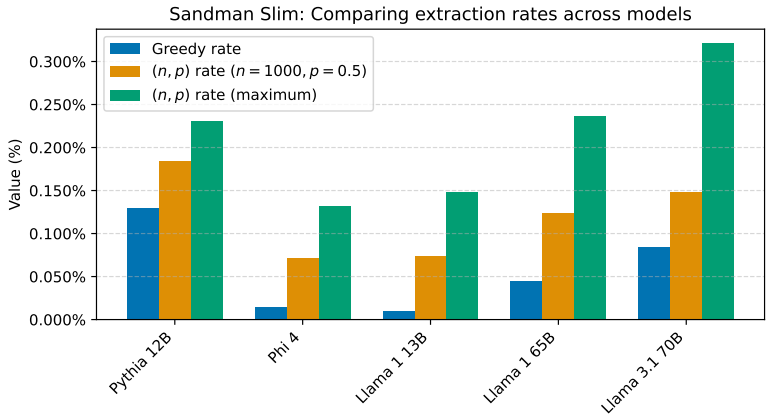

Figure 49: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Sandman Slim						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.13%	0.18%	0.23%	0.21%	0.17%	00:35:55.95
Phi 4 (14B)	0.02%	0.07%	0.13%	0.11%	0.04%	01:09:51.93
Llama 1 13B	0.01%	0.07%	0.15%	0.12%	0.05%	00:25:07.35
Llama 1 65B	0.05%	0.12%	0.24%	0.14%	0.08%	01:07:04.32
Llama 3.1 70B	0.08%	0.15%	0.32%	0.20%	0.12%	02:01:05.49

Table 21: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

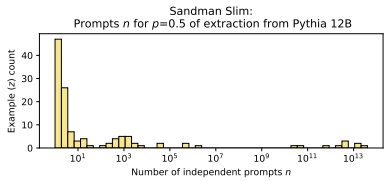

(a) Pythia 12B

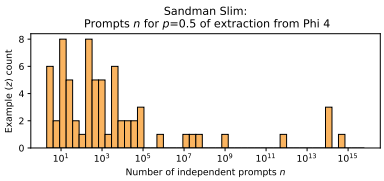

(b) Phi 4 (14B)

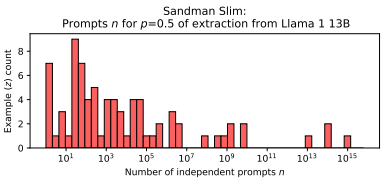

(c) Llama 1 13B

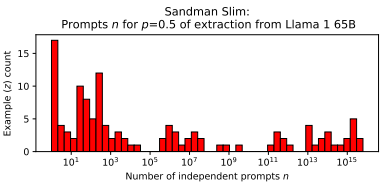

(d) Llama 1 65B

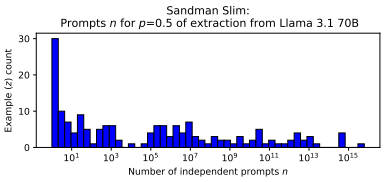

(e) Llama 3.1 70B

Figure 50: *Sandman Slim*, Kadrey (2009): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.17. *Sunburn*, LIPPMAN

All high-probability extracted sequences are at the end of the book or the beginning of the book. We omit heatmaps (which are mostly empty).

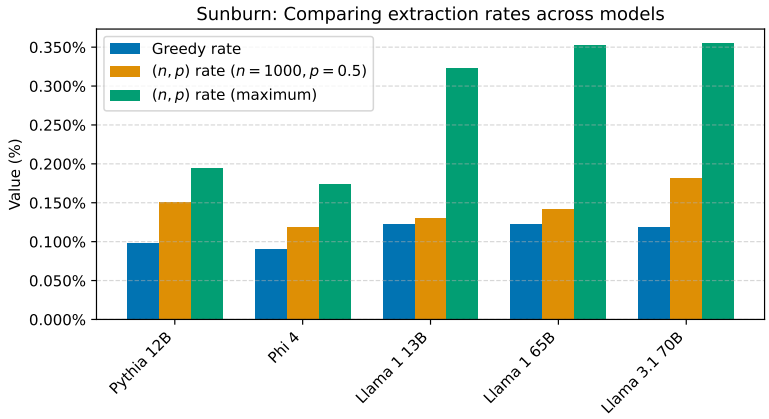

Figure 51: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Sunburn						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.10%	0.15%	0.20%	0.15%	0.13%	00:14:07.97
Phi 4 (14B)	0.09%	0.12%	0.17%	0.14%	0.11%	00:24:21.27
Llama 1 13B	0.12%	0.13%	0.32%	0.16%	0.13%	00:13:09.52
Llama 1 65B	0.12%	0.14%	0.35%	0.22%	0.13%	00:47:04.89
Llama 3.1 70B	0.12%	0.18%	0.29%	0.15%	0.20%	01:12:51.61

Table 22: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

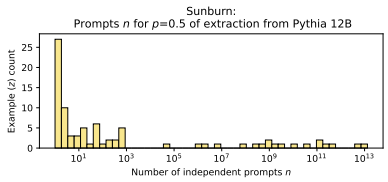

(a) Pythia 12B

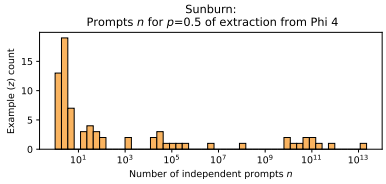

(b) Phi 4 (14B)

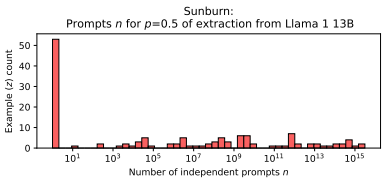

(c) Llama 1 13B

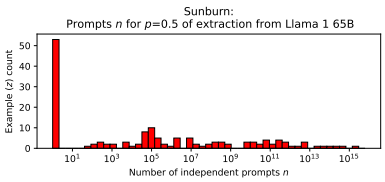

(d) Llama 1 65B

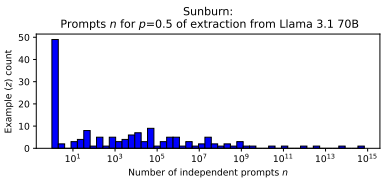

(e) Llama 3.1 70B

Figure 52: *Sunburn*, Lippman (2018): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.18. *The Beautiful Struggle*, COATES

All high-probability extracted sequences are at the end of the book or the beginning of the book. We omit heatmaps (which are mostly empty).

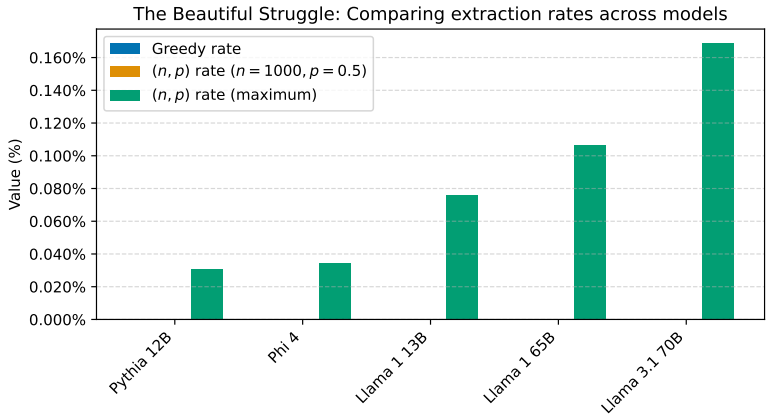

Figure 53: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Beautiful Struggle						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.00%	0.00%	0.03%	0.01%	0.00%	00:19:22.45
Phi 4 (14B)	0.00%	0.00%	0.03%	0.00%	0.00%	00:39:00.05
Llama 1 13B	0.00%	0.00%	0.08%	0.00%	0.00%	00:13:44.55
Llama 1 65B	0.00%	0.00%	0.11%	0.01%	0.00%	00:37:14.86
Llama 3.1 70B	0.00%	0.00%	0.17%	0.02%	0.00%	01:14:04.44

Table 23: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

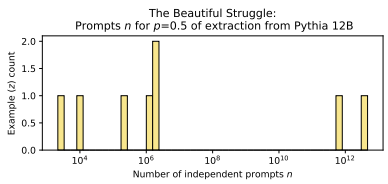

(a) Pythia 12B

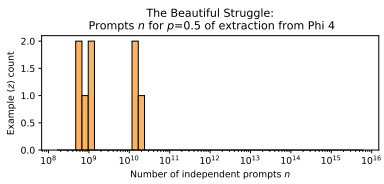

(b) Phi 4 (14B)

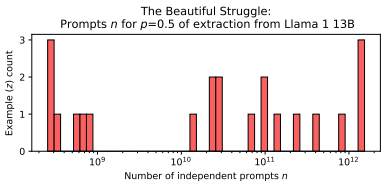

(c) Llama 1 13B

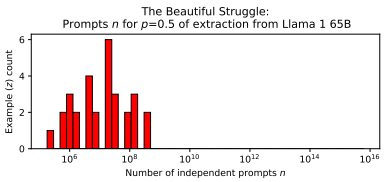

(d) Llama 1 65B

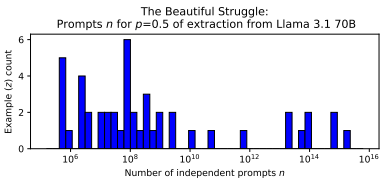

(e) Llama 3.1 70B

Figure 54: *The Beautiful Struggle*, Coates (2009): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.19. *The Bedwetter*, SILVERMAN

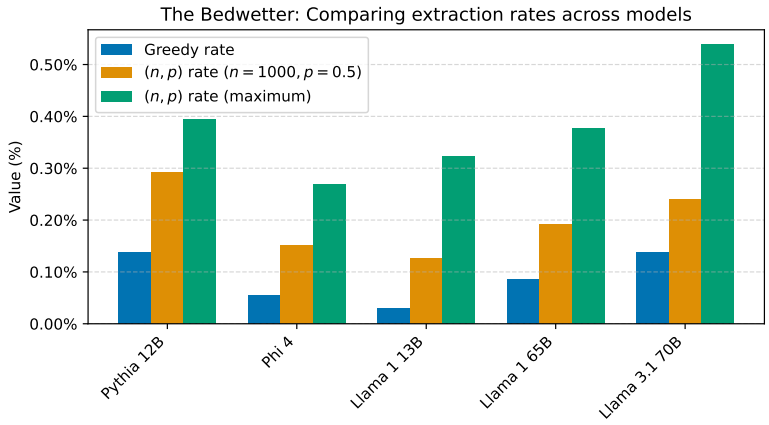

Figure 55: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Bedwetter						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.14%	0.29%	0.39%	0.35%	0.21%	00:21:15.89
Phi 4 (14B)	0.06%	0.15%	0.27%	0.23%	0.10%	00:43:03.81
Llama 1 13B	0.03%	0.13%	0.32%	0.24%	0.09%	00:15:56.30
Llama 1 65B	0.09%	0.19%	0.38%	0.27%	0.15%	00:41:15.11
Llama 3.1 70B	0.14%	0.24%	0.54%	0.32%	0.23%	01:23:01.61

Table 24: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

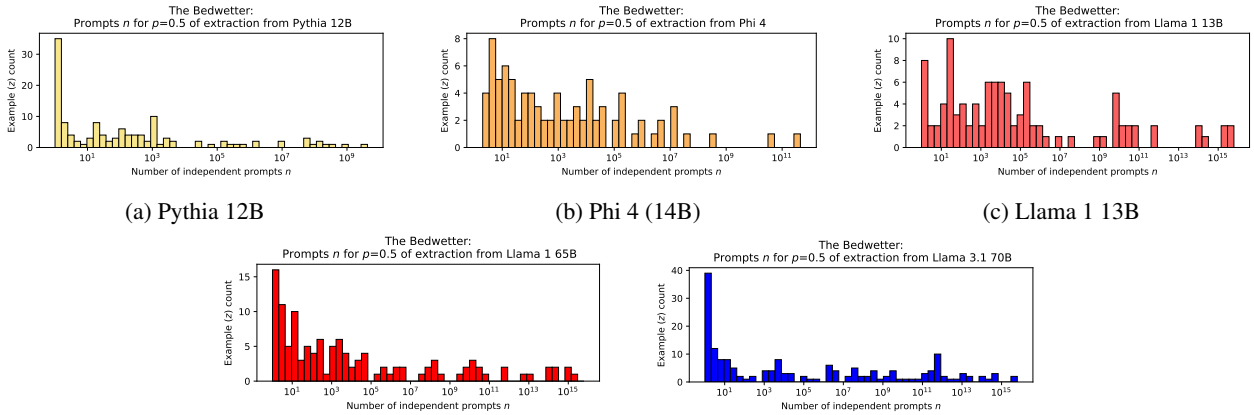

(a) Pythia 12B (b) Phi 4 (14B) (c) Llama 1 13B

(d) Llama 1 65B (e) Llama 3.1 70B

Figure 56: *The Bedwetter*, Silverman (2010): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

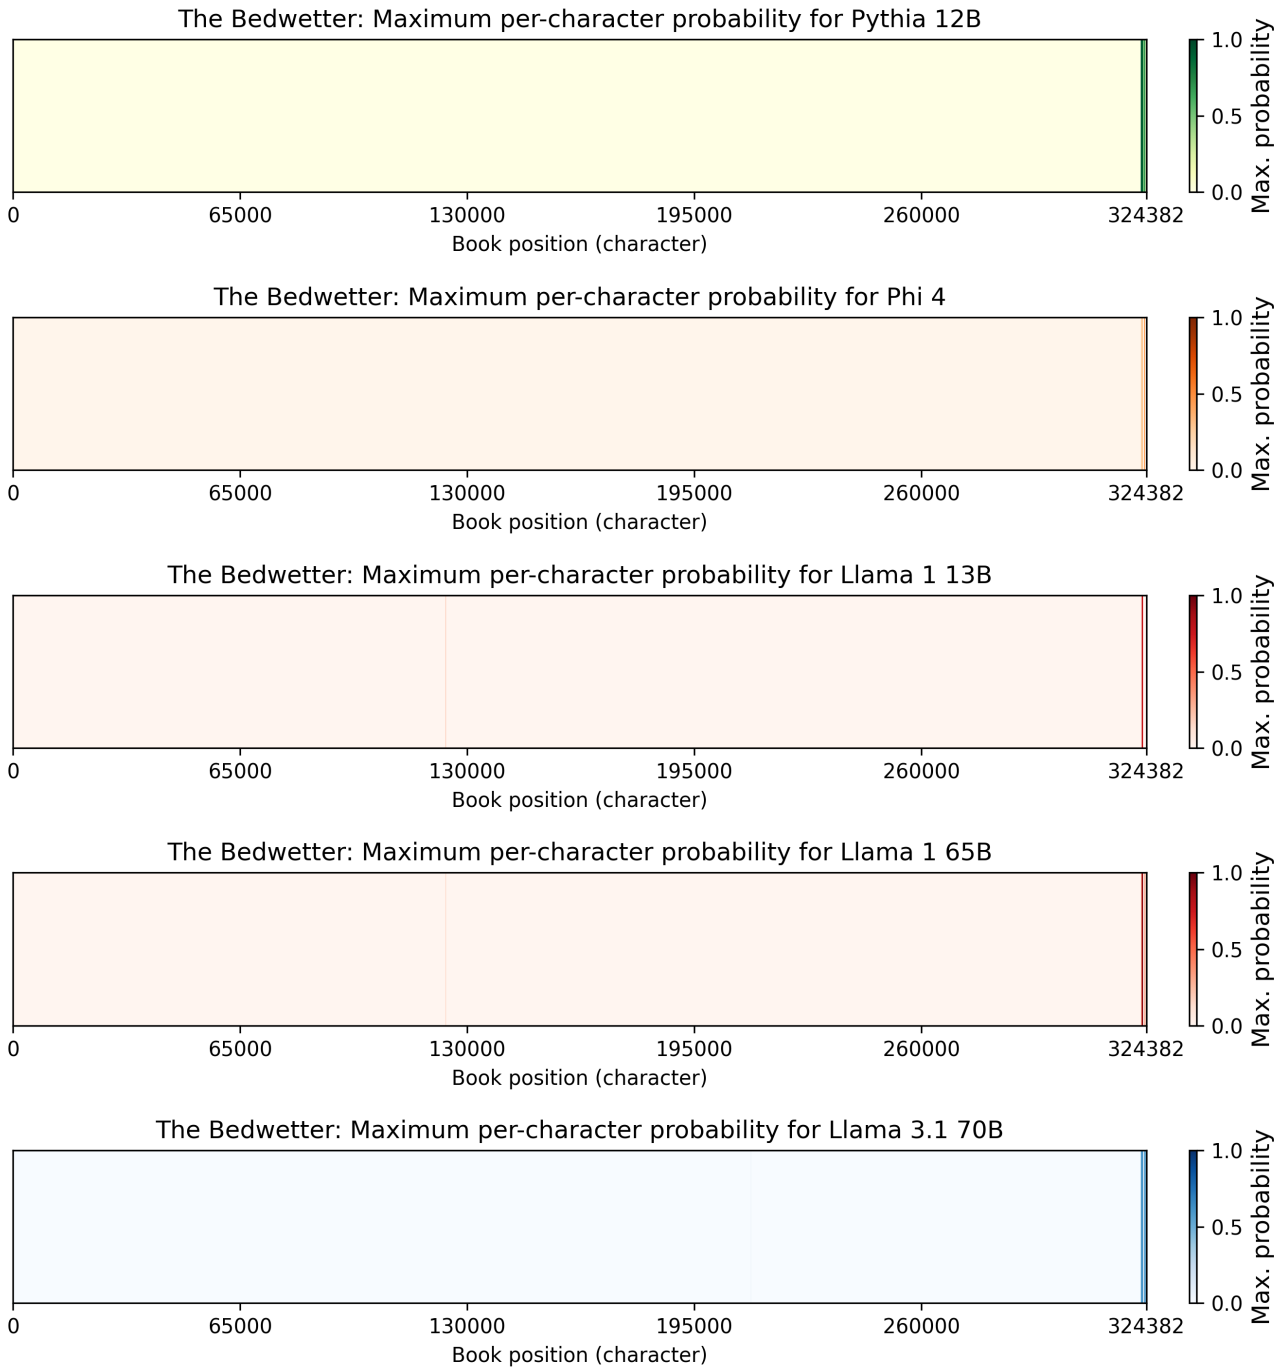

Figure 57: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.20. *The Brief Wondrous Life of Oscar Wao*, DÍAZ

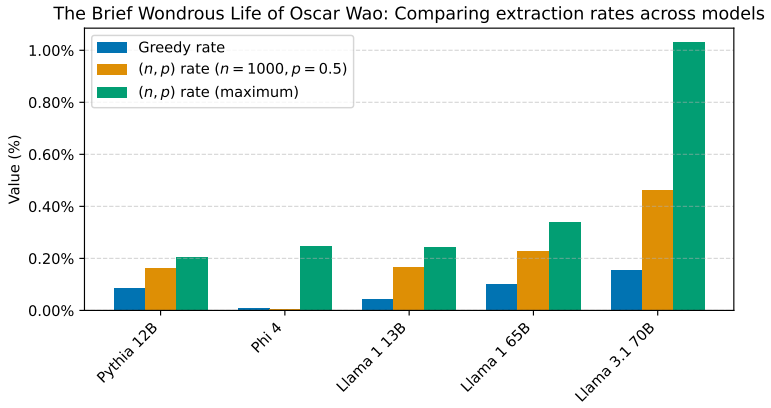

Figure 58: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Brief Wondrous Life of Oscar Wao						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.09%	0.16%	0.21%	0.18%	0.13%	00:17:12.42
Phi 4 (14B)	0.01%	0.00%	0.25%	0.05%	0.00%	00:27:47.94
Llama 1 13B	0.04%	0.17%	0.25%	0.20%	0.11%	00:16:17.08
Llama 1 65B	0.10%	0.23%	0.34%	0.29%	0.18%	00:55:33.06
Llama 3.1 70B	0.16%	0.46%	1.03%	0.64%	0.34%	01:26:55.34

Table 25: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

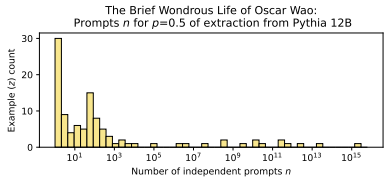

(a) Pythia 12B

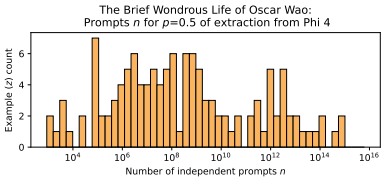

(b) Phi 4 (14B)

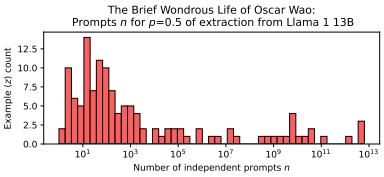

(c) Llama 1 13B

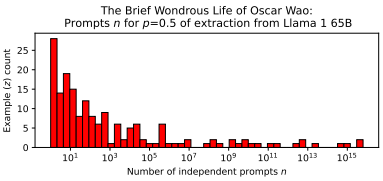

(d) Llama 1 65B

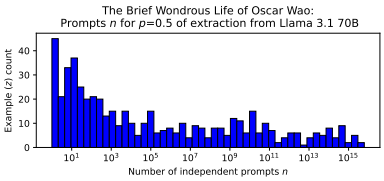

(e) Llama 3.1 70B

Figure 59: *The Brief Wondrous Life of Oscar Wao*, Díaz (2007): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

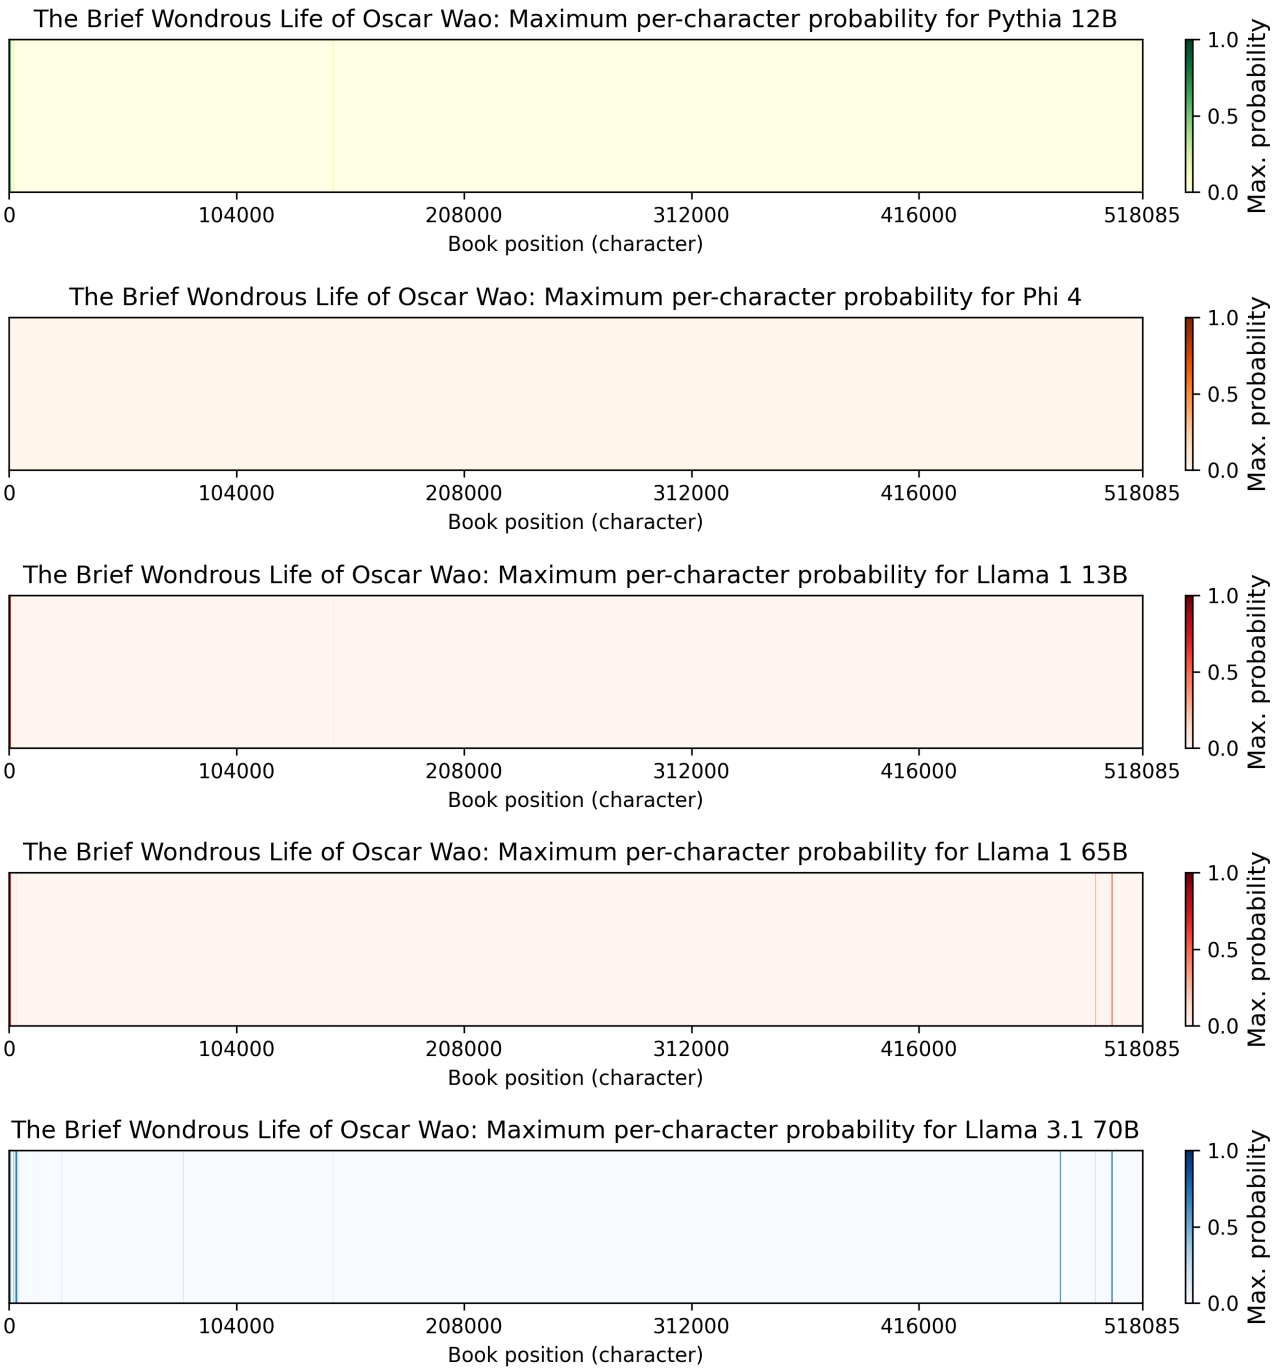

Figure 60: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.21. *The Confessions of Max Tivoli*, GREER

All high-probability extracted sequences are at the end of the book. We omit heatmaps (which are mostly empty).

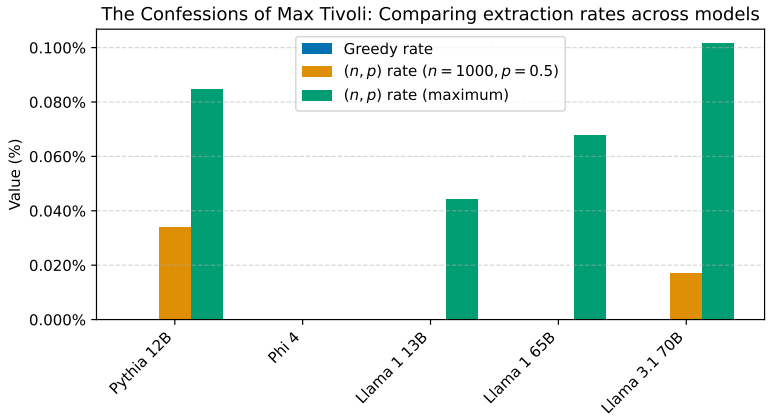

Figure 61: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Confessions of Max Tivoli						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.00%	0.03%	0.08%	0.05%	0.03%	00:31:23.50
Phi 4 (14B)	0.00%	0.00%	0.00%	0.00%	0.00%	01:03:00.43
Llama 1 13B	0.00%	0.00%	0.04%	0.00%	0.00%	00:22:51.88
Llama 1 65B	0.00%	0.00%	0.07%	0.00%	0.00%	00:59:47.86
Llama 3.1 70B	0.00%	0.02%	0.10%	0.03%	0.00%	01:53:00.62

Table 26: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

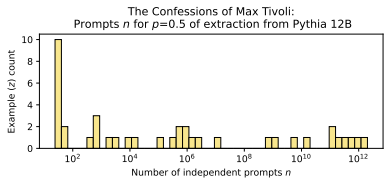

(a) Pythia 12B

Not plottable
(n too large for non-zero p_z $p = 0.5$)

(b) Phi 4 (14B)

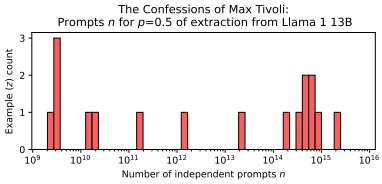

(c) Llama 1 13B

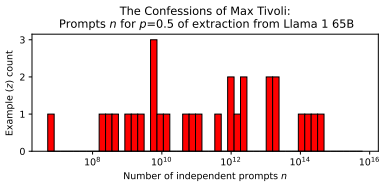

(d) Llama 1 65B

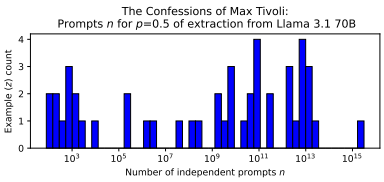

(e) Llama 3.1 70B

Figure 62: *The Confessions of Max Tivoli*, Greer (2005): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.22. *The Great Gatsby*, FITZGERALD

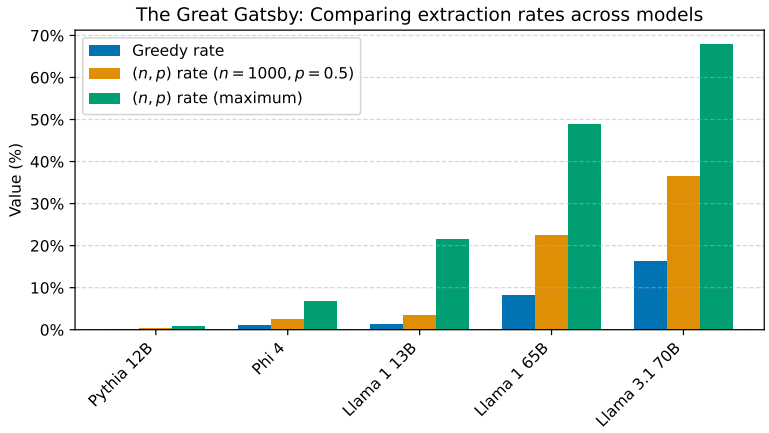

Figure 63: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Great Gatsby						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.16%	0.39%	0.91%	0.56%	0.20%	00:20:05.24
Phi 4 (14B)	1.00%	2.58%	6.68%	4.32%	1.78%	00:36:12.82
Llama 1 13B	1.33%	3.53%	21.53%	7.35%	2.02%	00:13:07.86
Llama 1 65B	8.22%	22.53%	48.99%	34.42%	15.86%	00:34:03.40
Llama 3.1 70B	16.23%	36.46%	67.86%	51.46%	27.85%	01:09:34.42

Table 27: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

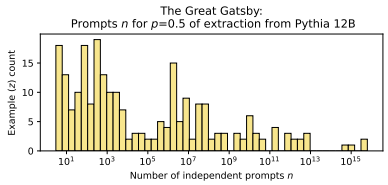

(a) Pythia 12B

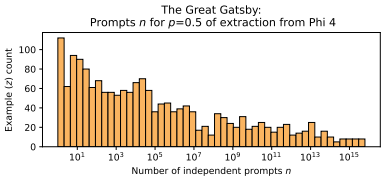

(b) Phi 4 (14B)

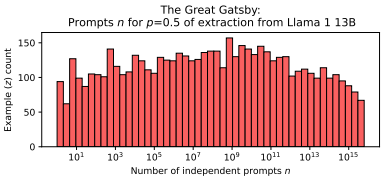

(c) Llama 1 13B

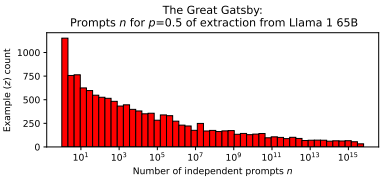

(d) Llama 1 65B

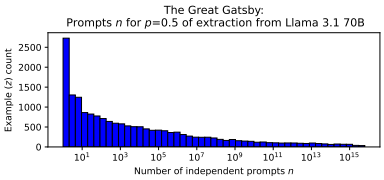

(e) Llama 3.1 70B

Figure 64: *The Great Gatsby*, Fitzgerald (1925): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

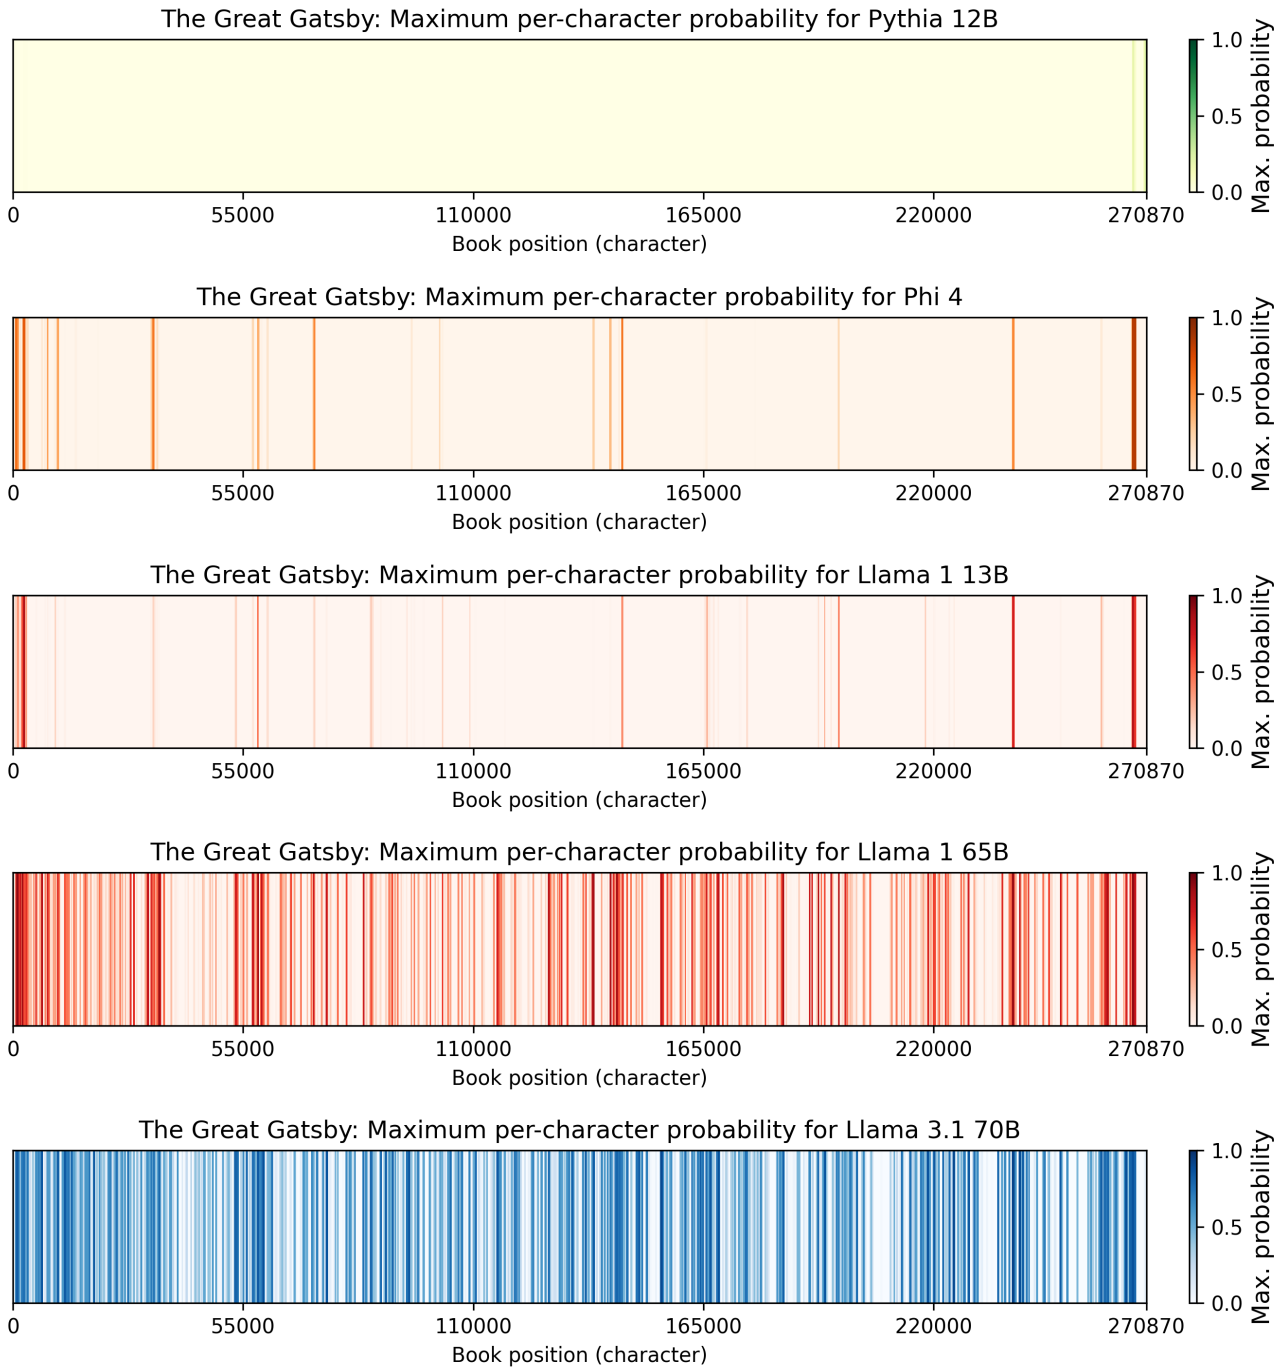

Figure 65: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.23. *The Hobbit*, TOLKIEN

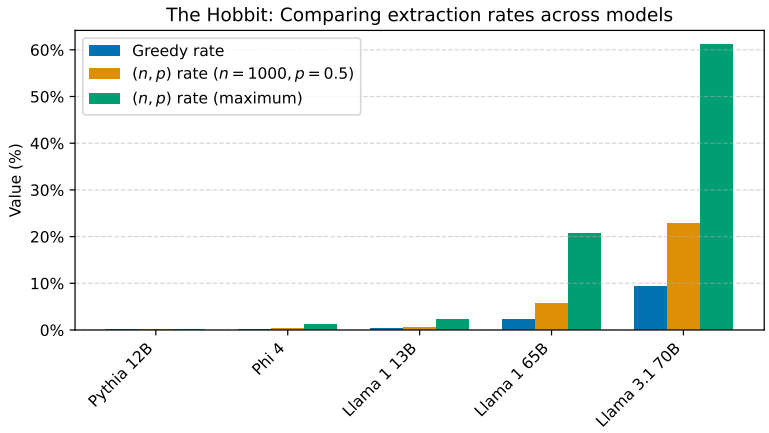

Figure 66: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Hobbit						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.01%	0.01%	0.21%	0.05%	0.01%	00:34:20.87
Phi 4 (14B)	0.20%	0.40%	1.09%	0.65%	0.28%	01:08:05.60
Llama 1 13B	0.24%	0.63%	2.27%	1.01%	0.39%	00:25:50.33
Llama 1 65B	2.18%	5.66%	20.67%	10.07%	3.64%	01:05:58.13
Llama 3.1 70B	9.42%	22.70%	61.10%	36.26%	16.37%	01:57:10.45

Table 28: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

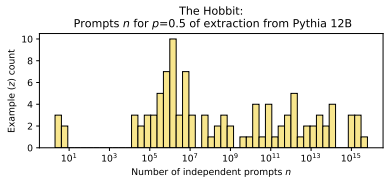

(a) Pythia 12B

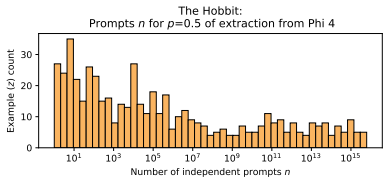

(b) Phi 4 (14B)

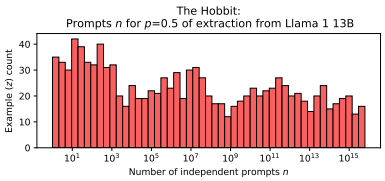

(c) Llama 1 13B

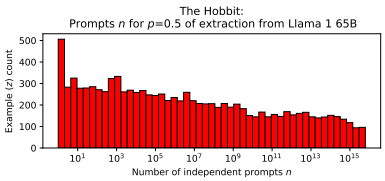

(d) Llama 1 65B

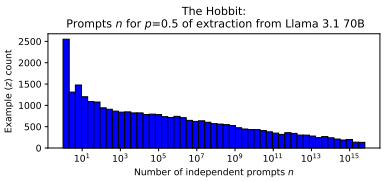

(e) Llama 3.1 70B

Figure 67: *The Hobbit*, Tolkien (1937): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

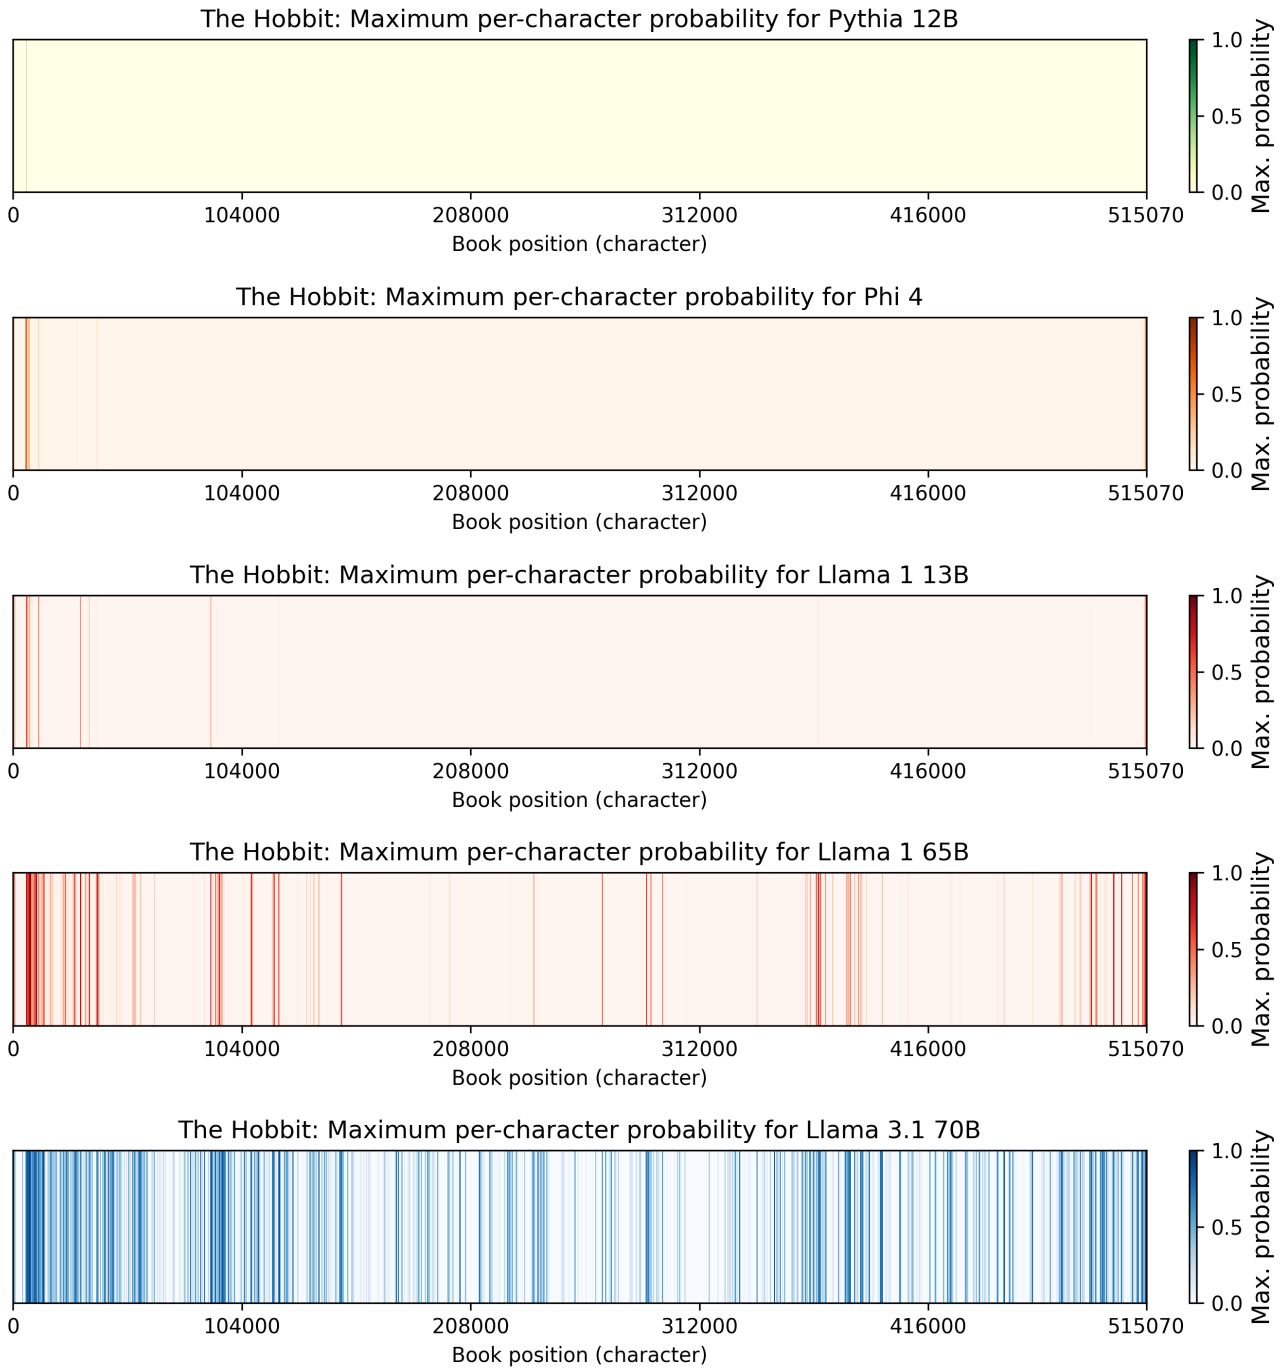

Figure 68: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.24. *The Making of a Mediterranean Emirate*, ROUIGHI

All high-probability extracted sequences are at the end of the book. We omit heatmaps (which are mostly empty). This book has a relatively large number of medium to low probability extractable sequences (and thus, while a low $(n = 1000, p = 0.5)$-discoverable extraction rate, there is a high maximum (n, p) extraction rate. This is also visible in Figure 70, where there is mass for each distribution for high n (while little mass, relatively, for low n, except for the sequences that are high probability at the end of the book). We investigate this further for individual sequences. See Appendix I.

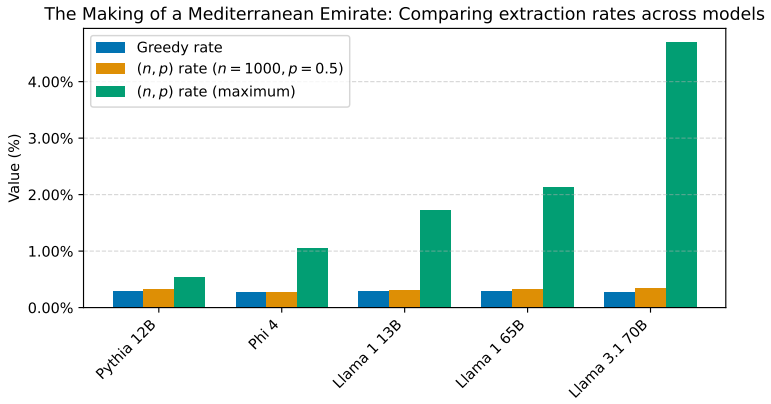

Figure 69: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Making of a Mediterranean Emirate						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.29%	0.33%	0.55%	0.37%	0.31%	00:18:22.98
Phi 4 (14B)	0.28%	0.28%	1.06%	0.31%	0.28%	00:32:12.37
Llama 1 13B	0.30%	0.32%	1.73%	0.37%	0.31%	00:18:48.93
Llama 1 65B	0.30%	0.33%	2.14%	0.46%	0.30%	01:03:15.78
Llama 3.1 70B	0.28%	0.35%	4.71%	0.82%	0.31%	01:35:12.01

Table 29: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

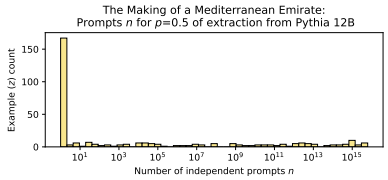

(a) Pythia 12B

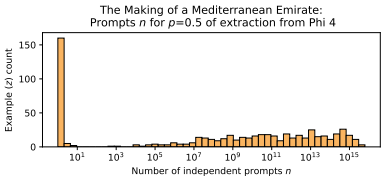

(b) Phi 4 (14B)

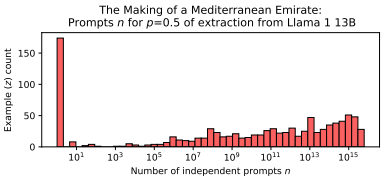

(c) Llama 1 13B

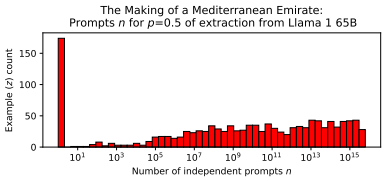

(d) Llama 1 65B

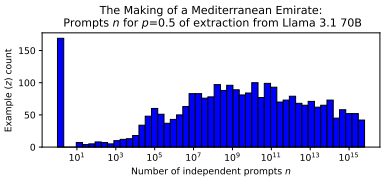

(e) Llama 3.1 70B

Figure 70: *The Making of a Mediterranean Emirate*, Rouighi (2011): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.25. *The Myth of Sisyphus*, CAMUS

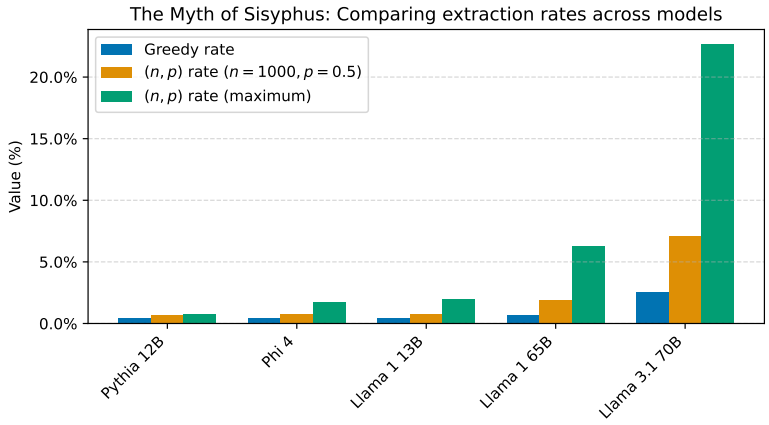

Figure 71: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Myth of Sisyphus						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.45%	0.65%	0.81%	0.60%	0.69%	00:10:47.48
Phi 4 (14B)	0.47%	0.81%	1.76%	0.67%	1.12%	00:22:30.60
Llama 1 13B	0.48%	0.79%	2.03%	0.59%	1.16%	00:11:17.22
Llama 1 65B	0.68%	1.92%	6.30%	1.22%	3.12%	00:38:00.27
Llama 3.1 70B	2.52%	7.13%	22.72%	4.84%	12.22%	01:05:39.22

Table 30: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

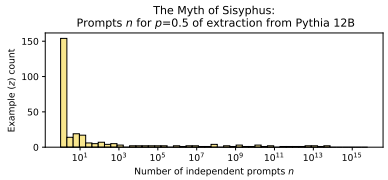

(a) Pythia 12B

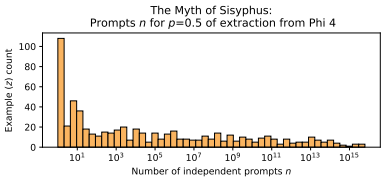

(b) Phi 4 (14B)

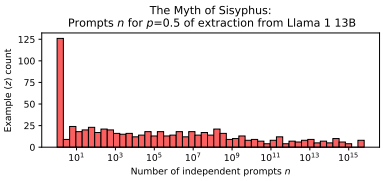

(c) Llama 1 13B

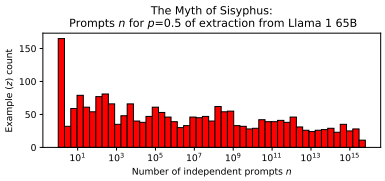

(d) Llama 1 65B

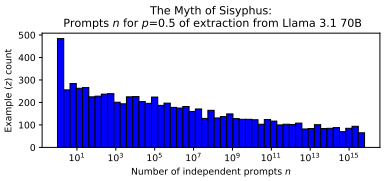

(e) Llama 3.1 70B

Figure 72: *The Myth of Sisyphus*, Camus (1955): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

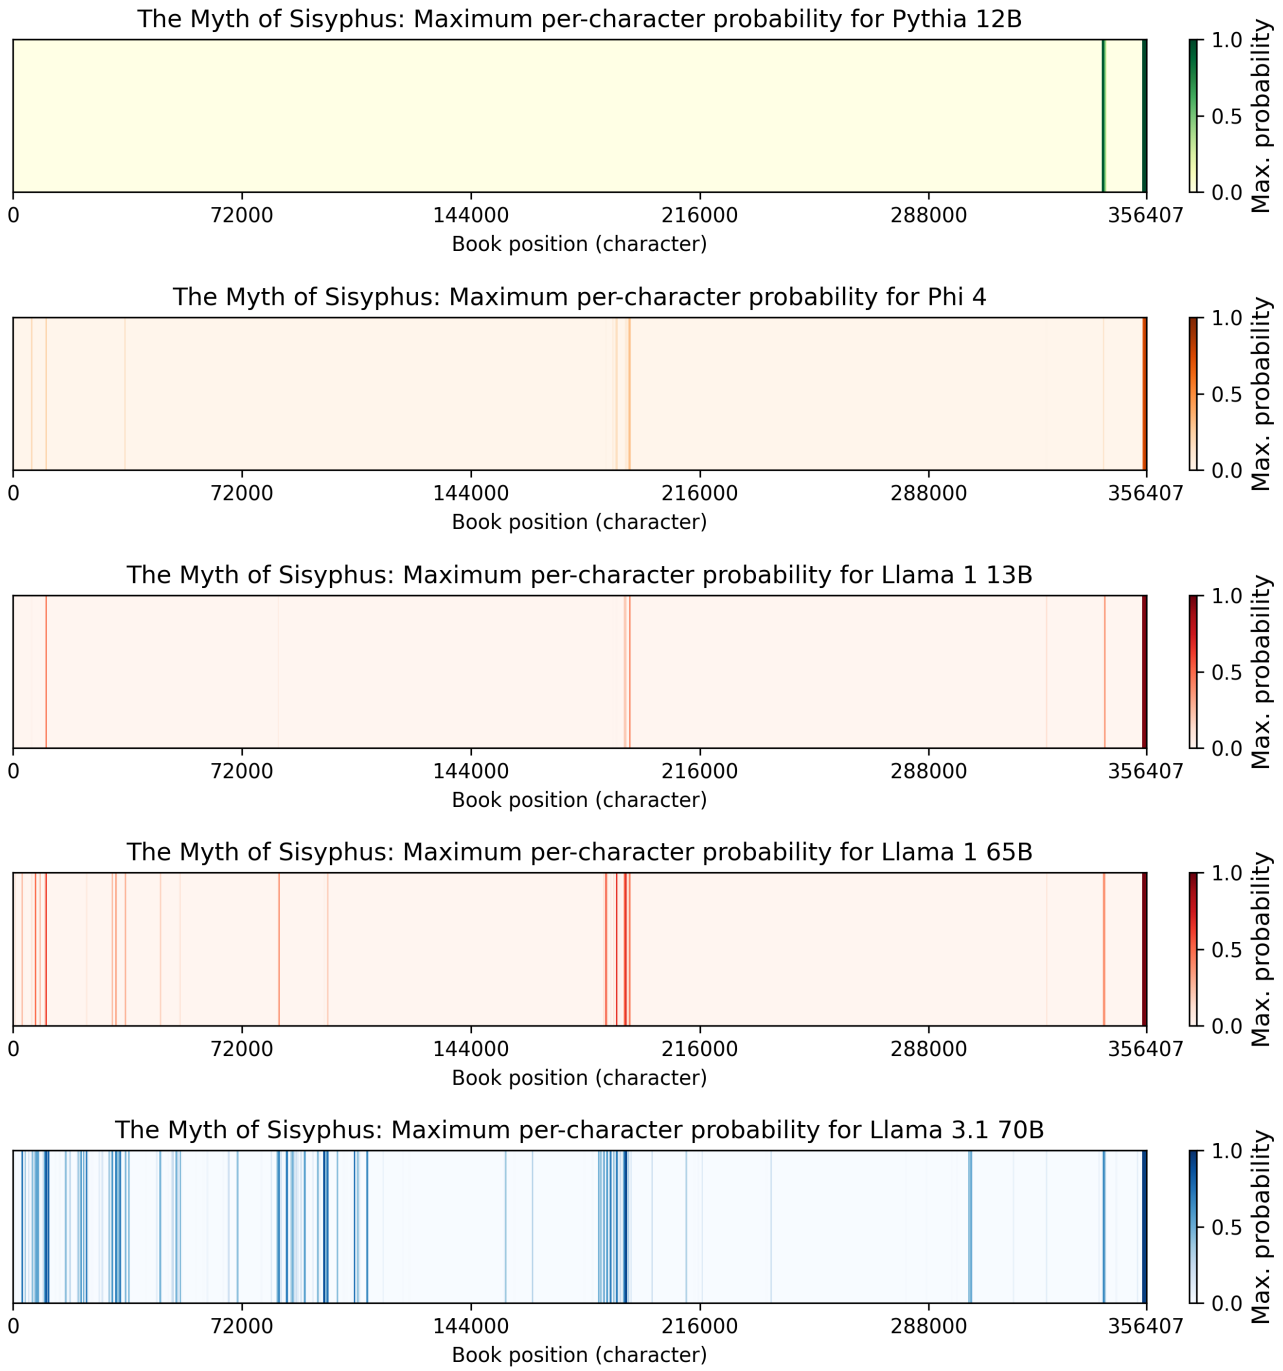

Figure 73: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.26. *The President's Vampire*, FARNSWORTH

All high-probability extracted sequences are at the beginning of the book. We omit heatmaps (which are mostly empty).

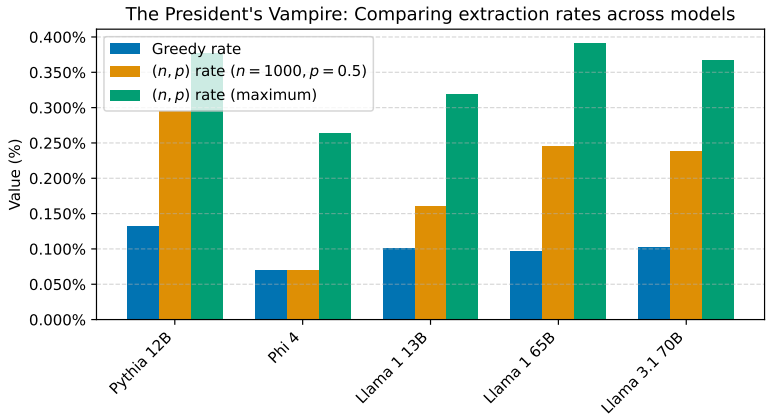

Figure 74: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The President's Vampire						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.13%	0.29%	0.38%	0.21%	0.33%	00:32:58.48
Phi 4 (14B)	0.07%	0.07%	0.26%	0.09%	0.07%	01:02:03.59
Llama 1 13B	0.10%	0.16%	0.32%	0.26%	0.12%	00:23:26.33
Llama 1 65B	0.10%	0.25%	0.39%	0.31%	0.18%	01:01:31.94
Llama 3.1 70B	0.10%	0.24%	0.37%	0.31%	0.15%	01:50:53.93

Table 31: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

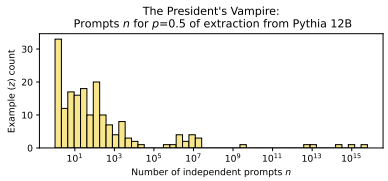

(a) Pythia 12B

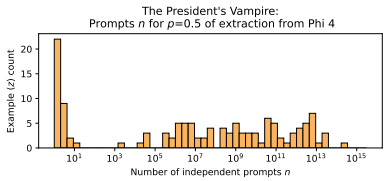

(b) Phi 4 (14B)

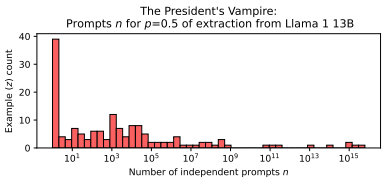

(c) Llama 1 13B

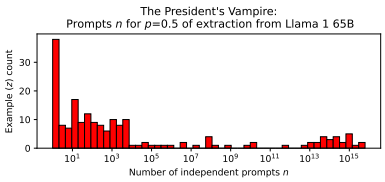

(d) Llama 1 65B

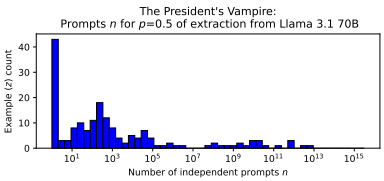

(e) Llama 3.1 70B

Figure 75: *The President's Vampire*, Farnsworth (2011): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.27. *The Water Dancer*, COATES

All high-probability extracted sequences are at the beginning or the end of the book. We omit heatmaps (which are mostly empty).

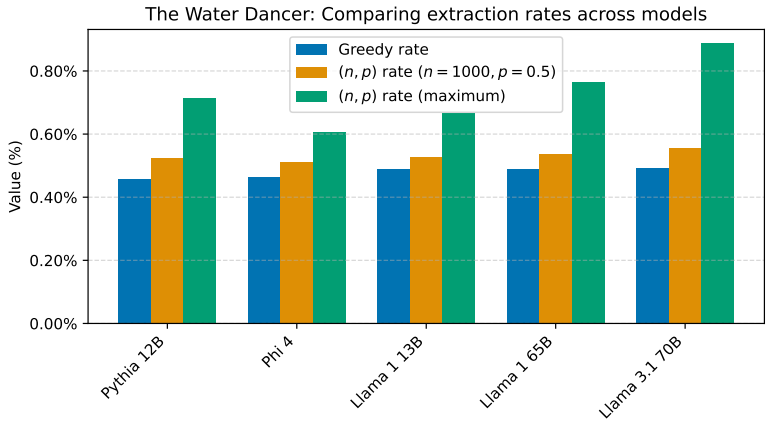

Figure 76: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Water Dancer						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.46%	0.52%	0.71%	0.64%	0.50%	00:21:44.11
Phi 4 (14B)	0.46%	0.51%	0.61%	0.52%	0.47%	00:40:49.06
Llama 1 13B	0.49%	0.53%	0.67%	0.58%	0.52%	00:22:15.13
Llama 1 65B	0.49%	0.54%	0.76%	0.62%	0.53%	01:16:48.71
Llama 3.1 70B	0.49%	0.56%	0.89%	0.68%	0.54%	01:54:27.22

Table 32: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

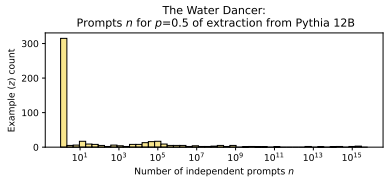

(a) Pythia 12B

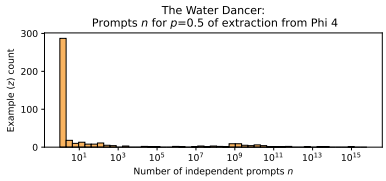

(b) Phi 4 (14B)

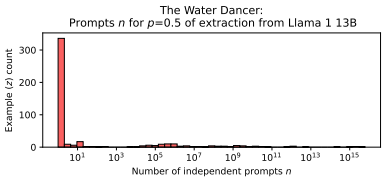

(c) Llama 1 13B

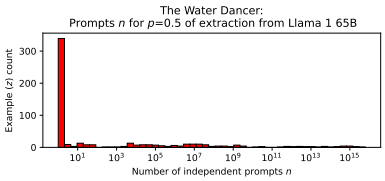

(d) Llama 1 65B

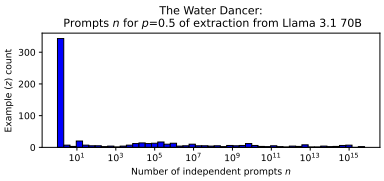

(e) Llama 3.1 70B

Figure 77: *The Water Dancer*, Coates (2019): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.28. *Unglued*, TERKEURST

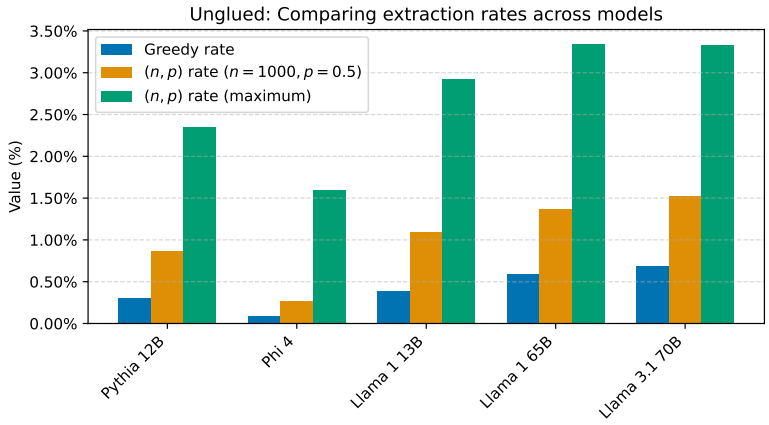

Figure 78: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Unglued						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.31%	0.86%	2.35%	1.40%	0.54%	00:09:36.71
Phi 4 (14B)	0.08%	0.26%	1.59%	0.75%	0.16%	00:18:33.90
Llama 1 13B	0.38%	1.09%	2.92%	1.80%	0.75%	00:09:39.71
Llama 1 65B	0.60%	1.37%	3.35%	2.03%	1.10%	00:33:11.48
Llama 3.1 70B	0.68%	1.53%	3.33%	2.30%	1.29%	00:57:53.17

Table 33: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

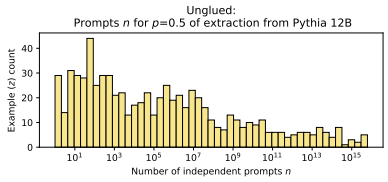

(a) Pythia 12B

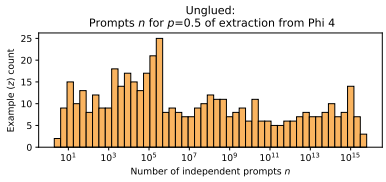

(b) Phi 4 (14B)

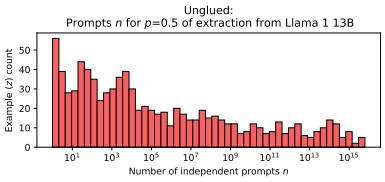

(c) Llama 1 13B

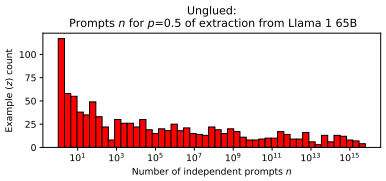

(d) Llama 1 65B

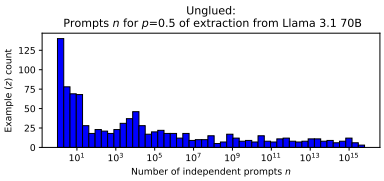

(e) Llama 3.1 70B

Figure 79: *Unglued*, TerKeurst (2012): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

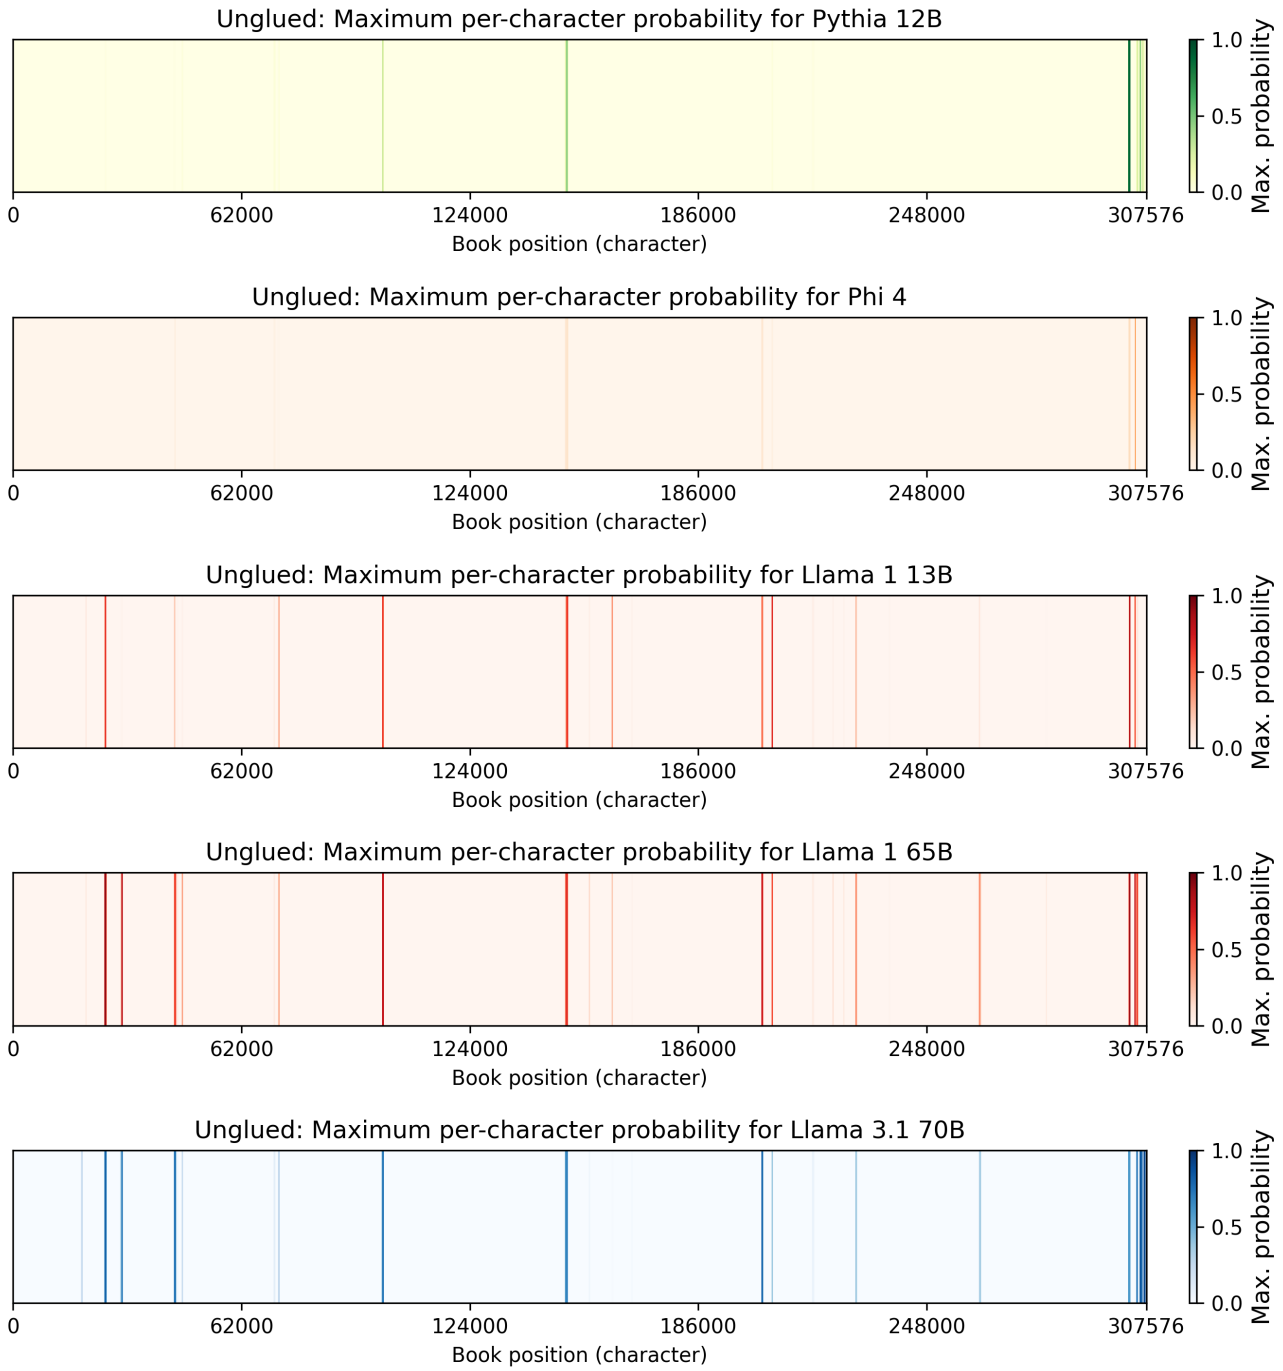

Figure 80: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.29. *We Were Eight Years in Power*, COATES

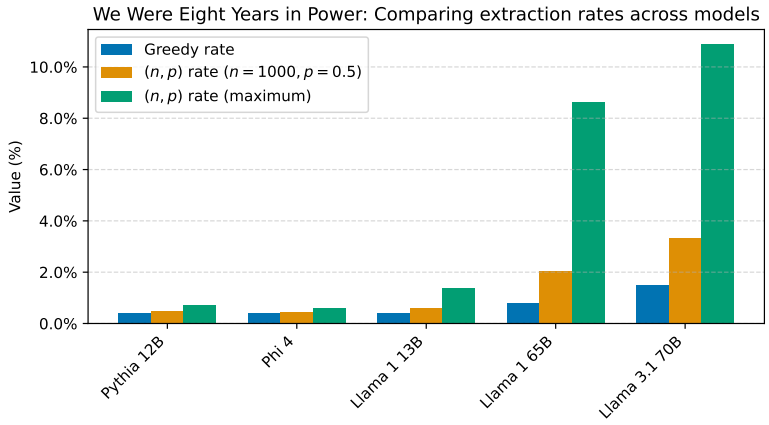

Figure 81: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

We Were Eight Years in Power						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.42%	0.48%	0.73%	0.46%	0.55%	00:21:20.90
Phi 4 (14B)	0.41%	0.45%	0.62%	0.50%	0.43%	00:39:18.74
Llama 1 13B	0.43%	0.62%	1.38%	0.82%	0.52%	00:21:48.74
Llama 1 65B	0.80%	2.03%	8.64%	3.44%	1.31%	01:13:02.59
Llama 3.1 70B	1.52%	3.33%	10.91%	5.26%	2.46%	01:48:58.29

Table 34: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

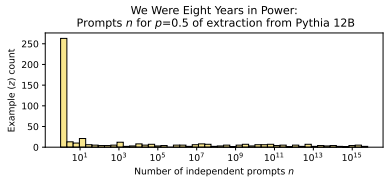

(a) Pythia 12B

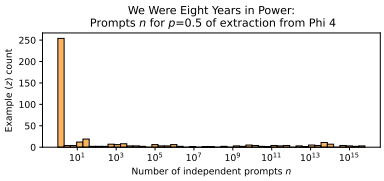

(b) Phi 4 (14B)

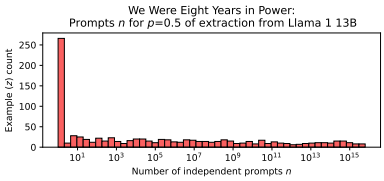

(c) Llama 1 13B

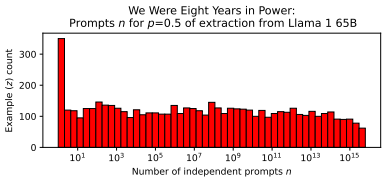

(d) Llama 1 65B

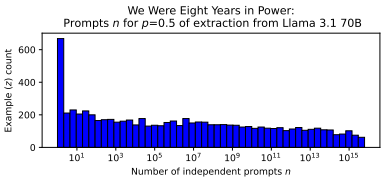

(e) Llama 3.1 70B

Figure 82: *We Were Eight Years in Power*, Coates (2017): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

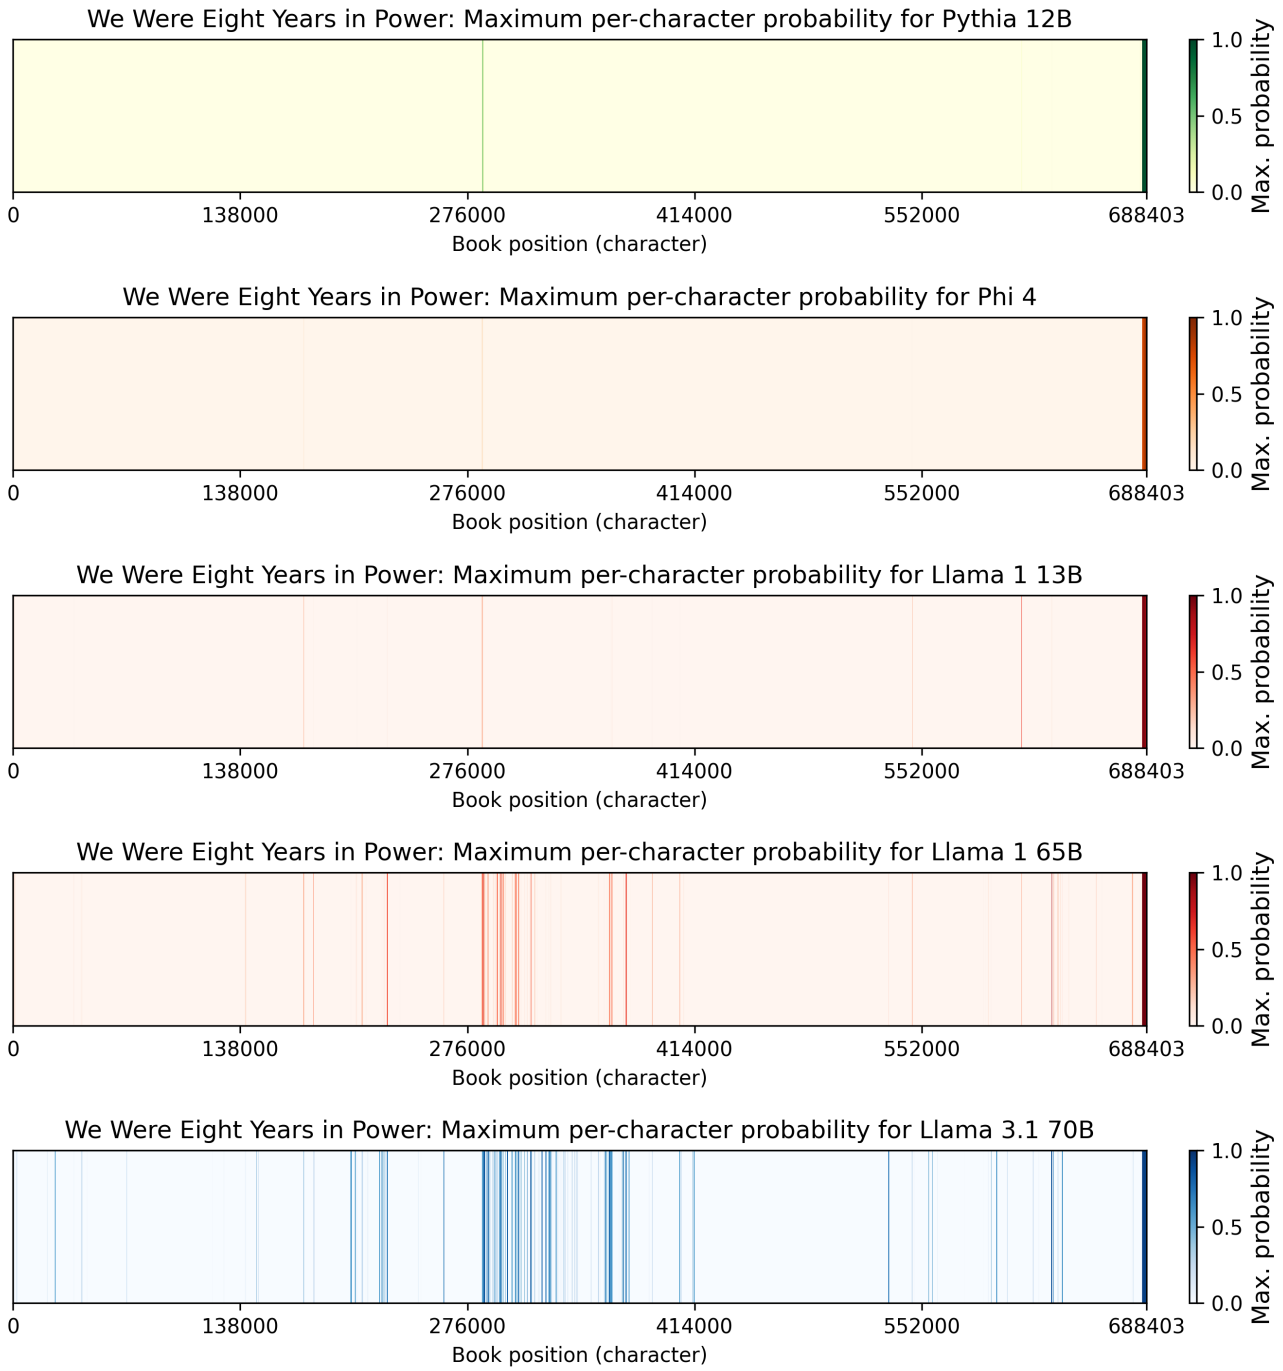

Figure 83: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.30. *Who Is Rich*

All high-probability extracted sequences are at the beginning or end of the book. We therefore omit heatmaps (which are mostly empty).

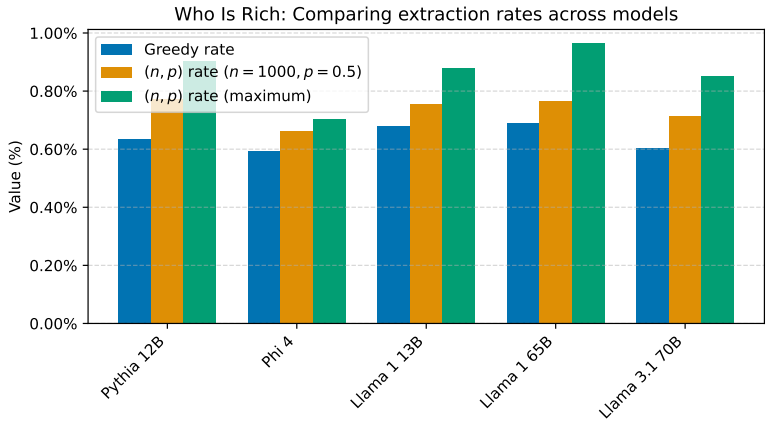

Figure 84: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Who Is Rich						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.64%	0.77%	0.90%	0.88%	0.73%	00:33:56.93
Phi 4 (14B)	0.59%	0.66%	0.71%	0.69%	0.64%	01:07:13.60
Llama 1 13B	0.68%	0.76%	0.88%	0.78%	0.74%	00:23:58.84
Llama 1 65B	0.69%	0.76%	0.96%	0.79%	0.75%	01:05:50.46
Llama 3.1 70B	0.60%	0.71%	0.85%	0.76%	0.69%	02:00:00.09

Table 35: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

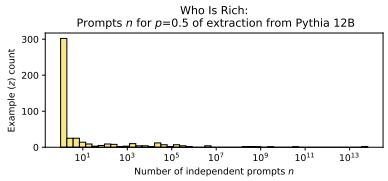

(a) Pythia 12B

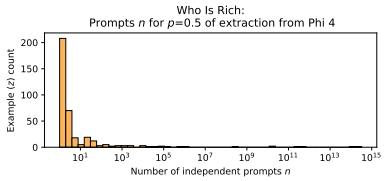

(b) Phi 4 (14B)

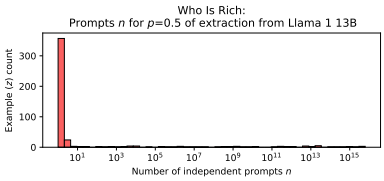

(c) Llama 1 13B

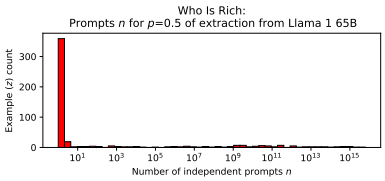

(d) Llama 1 65B

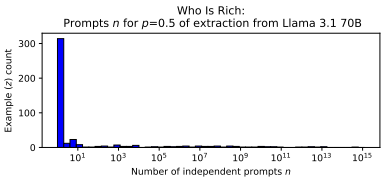

(e) Llama 3.1 70B

Figure 85: *Who Is Rich*, Klam (2017): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

G.1.31. *A Game of Thrones*, MARTIN

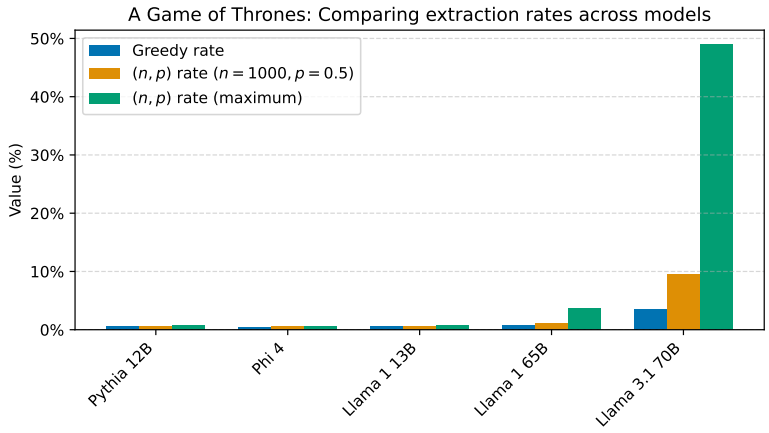

Figure 86: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

A Game of Thrones						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.52%	0.55%	0.70%	0.58%	0.53%	00:50:39.51
Phi 4 (14B)	0.51%	0.53%	0.63%	0.56%	0.52%	01:25:11.53
Llama 1 13B	0.54%	0.57%	0.85%	0.61%	0.54%	00:50:41.26
Llama 1 65B	0.73%	1.10%	3.64%	1.73%	0.87%	02:52:01.23
Llama 3.1 70B	3.57%	9.45%	48.97%	18.69%	6.13%	03:49:46.79

Table 36: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

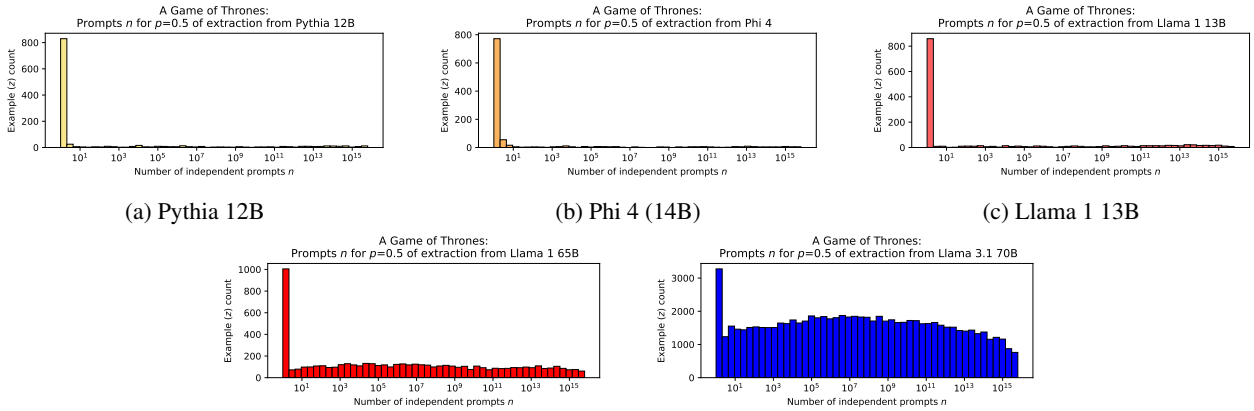

(a) Pythia 12B (b) Phi 4 (14B) (c) Llama 1 13B

(d) Llama 1 65B (e) Llama 3.1 70B

Figure 87: *A Game of Thrones*, Martin (1996): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

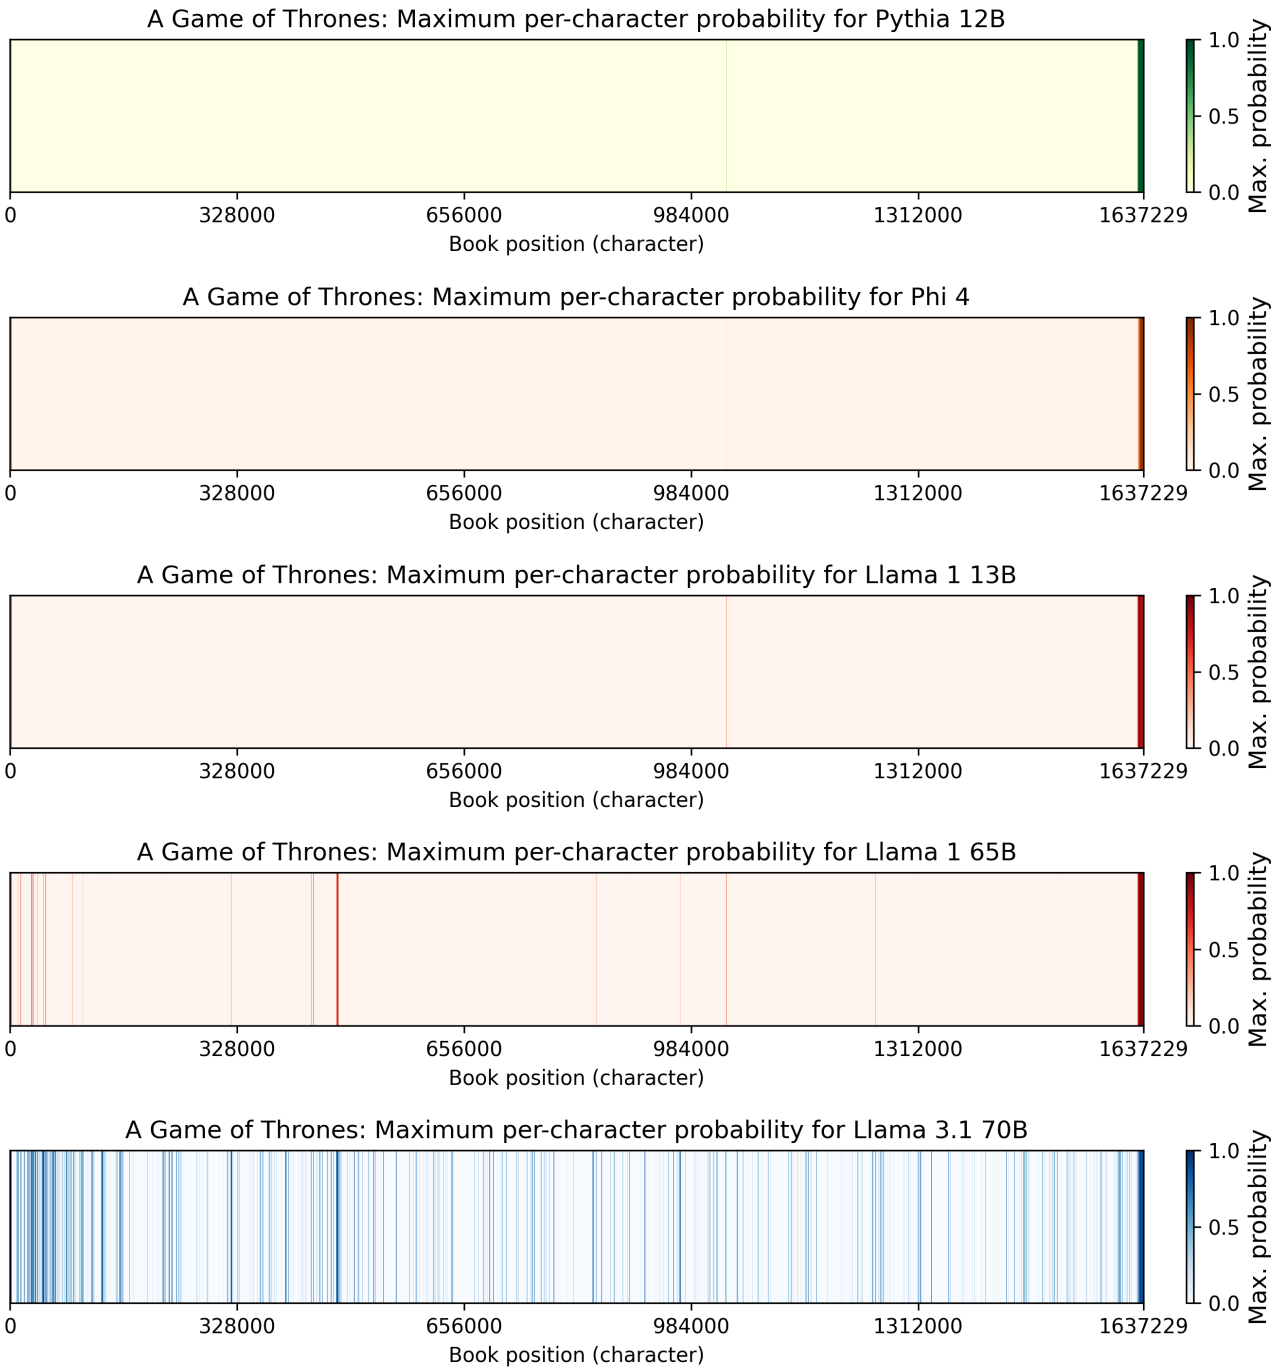

Figure 88: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.32. *Lean In*, SANDBERG

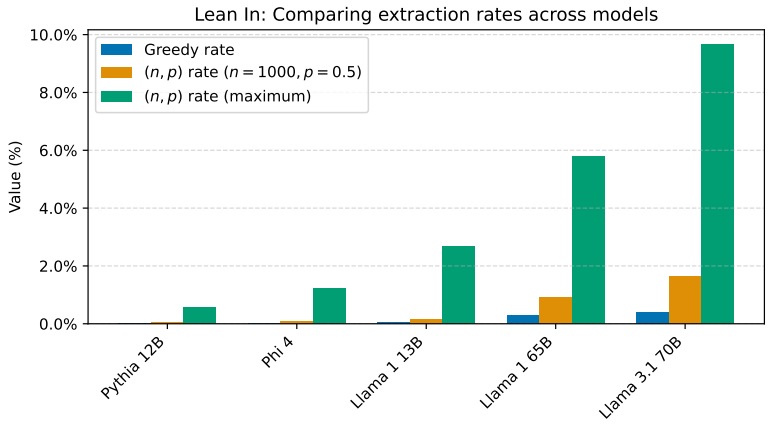

Figure 89: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

			Lean In			
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.03%	0.04%	0.56%		0.11%	0.03% 00:14:13.58
Phi 4 (14B)	0.02%	0.10%	1.22%	0.34%	0.05%	00:26:31.67
Llama 1 13B	0.05%	0.17%	2.68%	0.48%	0.10%	00:15:00.22
Llama 1 65B	0.31%	0.93%	5.79%	1.95%	0.61%	00:49:39.22
Llama 3.1 70B	0.39%	1.63%	9.68%	3.69%	0.83%	01:15:27.27

Table 37: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

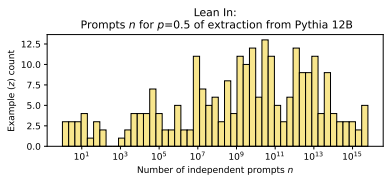

(a) Pythia 12B

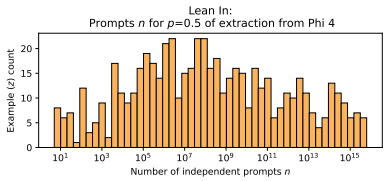

(b) Phi 4 (14B)

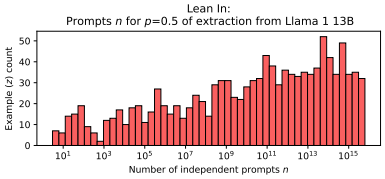

(c) Llama 1 13B

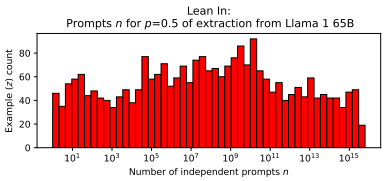

(d) Llama 1 65B

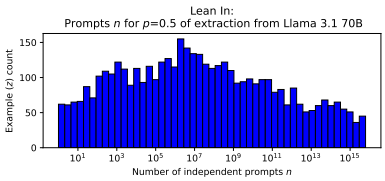

(e) Llama 3.1 70B

Figure 90: *Lean In*, Sandberg (2013): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

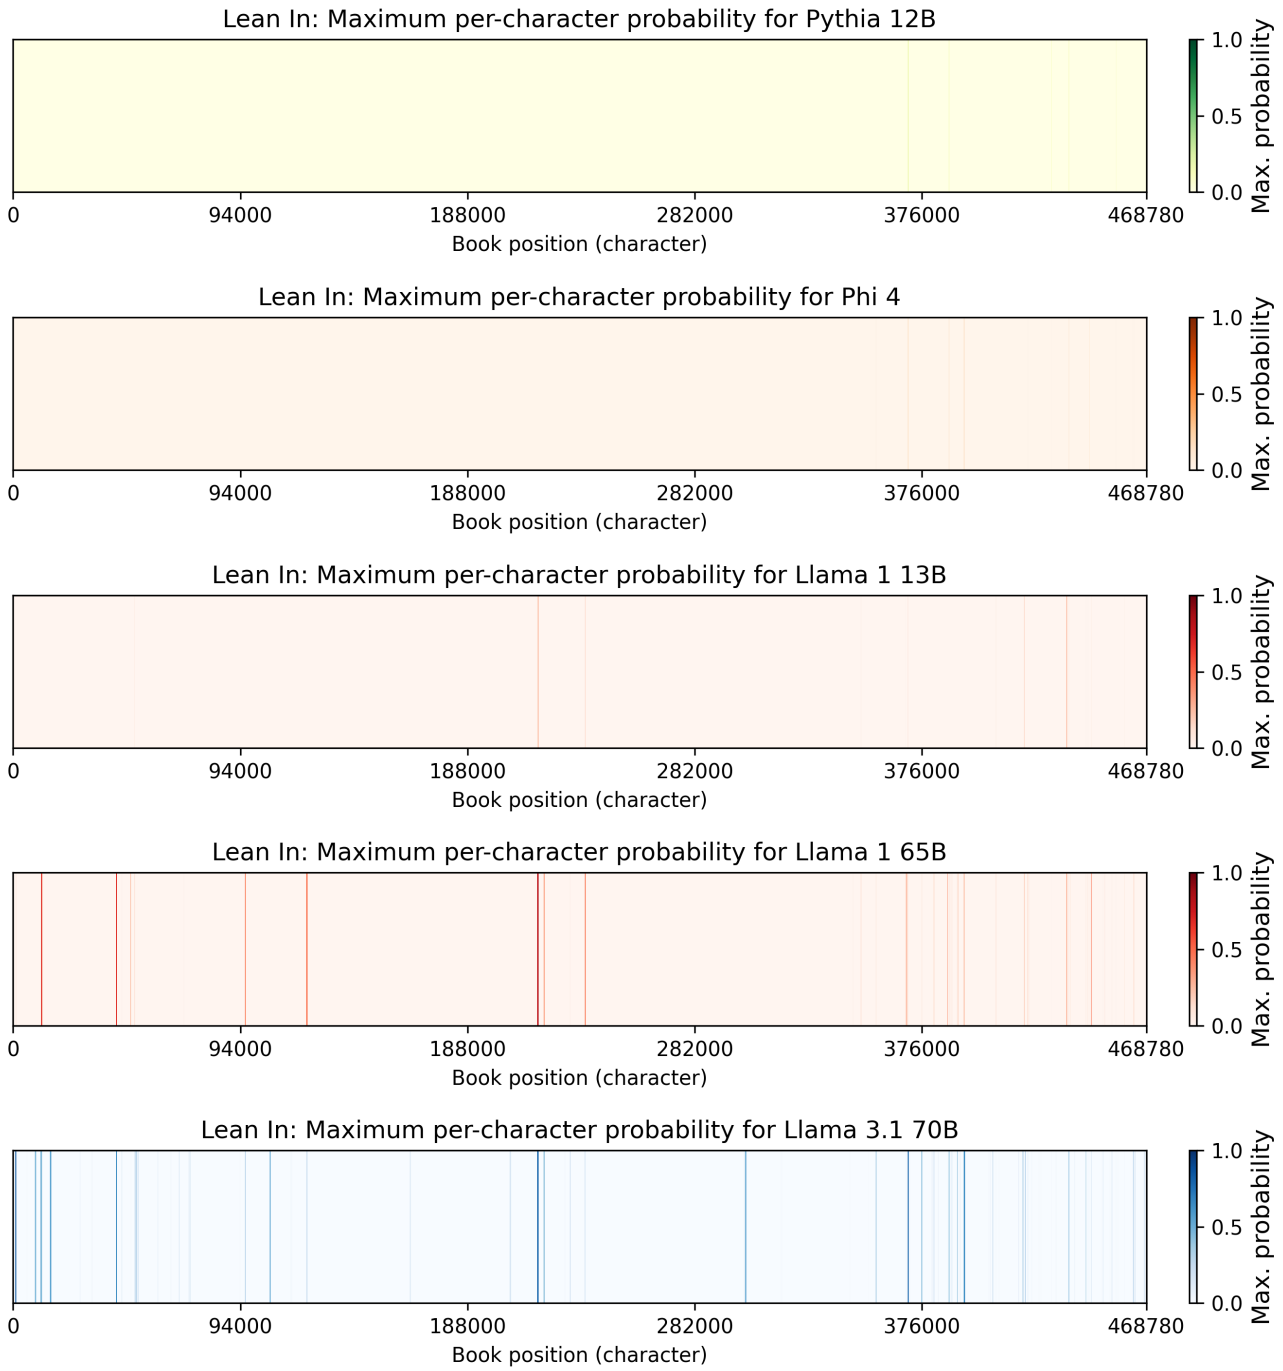

Figure 91: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.33. *Nineteen-Eighty Four*, ORWELL

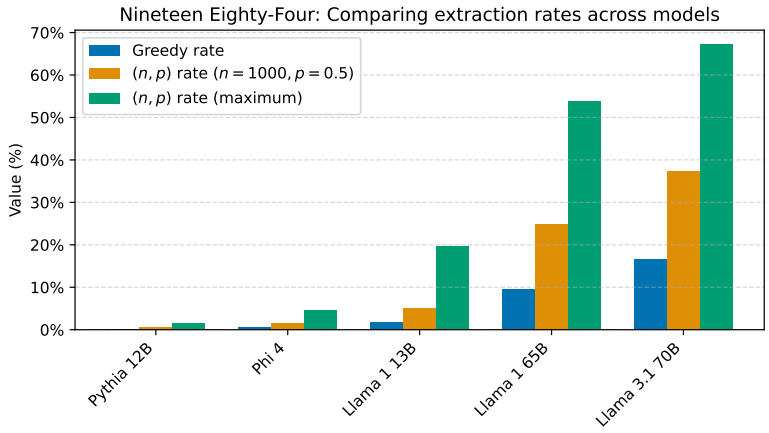

Figure 92: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Nineteen Eighty-Four						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.22%	0.54%	1.61%	0.84%	0.35%	00:21:21.22
Phi 4 (Phi 14B)	0.57%	1.46%	4.63%	2.53%	0.98%	00:35:20.85
Llama 1 13B	1.85%	5.06%	19.66%	9.28%	3.26%	00:21:23.13
Llama 1 65B	9.63%	24.80%	53.90%	37.46%	17.73%	01:11:45.83
Llama 3.1 70B	16.58%	37.31%	67.21%	51.68%	28.48%	01:44:54.76

Table 38: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

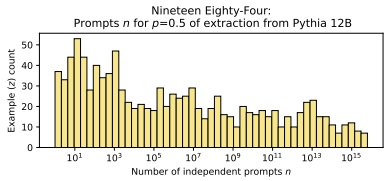

(a) Pythia 12B

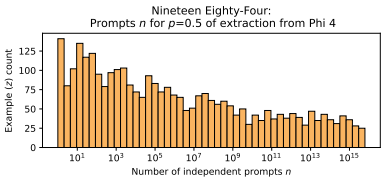

(b) Phi 4 (14B)

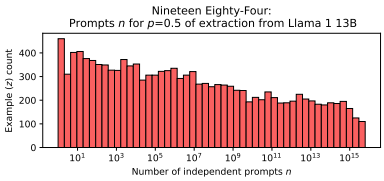

(c) Llama 1 13B

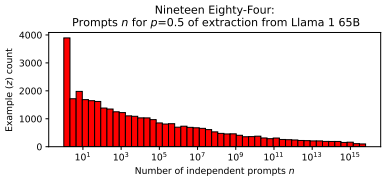

(d) Llama 1 65B

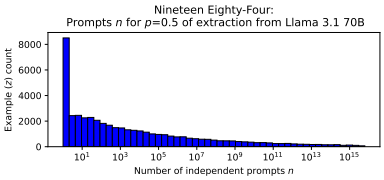

(e) Llama 3.1 70B

Figure 93: *Nineteen Eighty-Four*, Orwell (1949): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

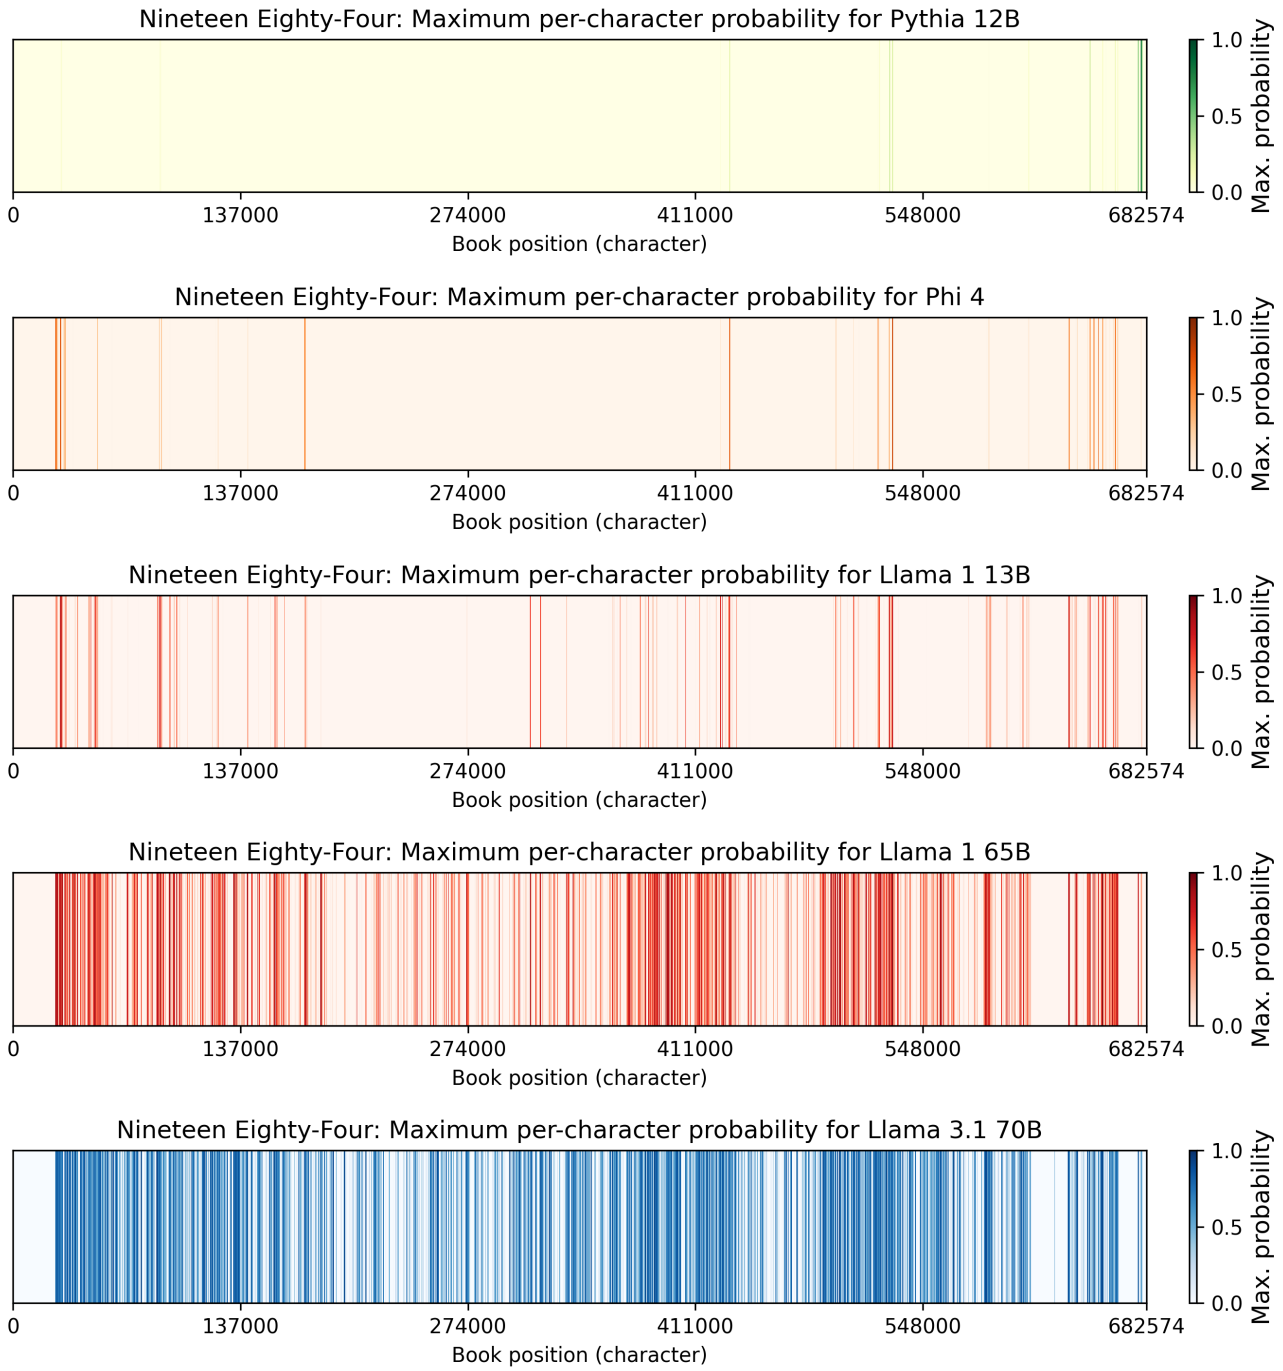

Figure 94: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.34. *The Da Vinci Code*, BROWN

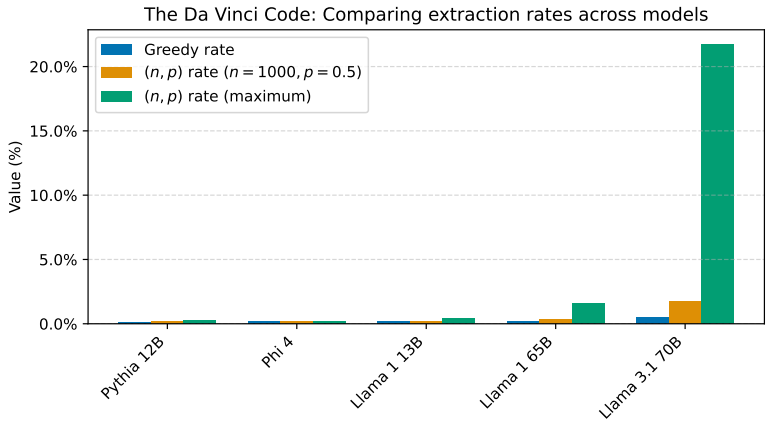

Figure 95: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

The Da Vinci Code						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.15%	0.18%	0.26%	0.21%	0.17%	00:25:41.35
Phi 4 (14B)	0.16%	0.17%	0.22%	0.19%	0.17%	00:45:56.32
Llama 1 13B	0.16%	0.20%	0.47%	0.22%	0.18%	00:26:15.45
Llama 1 65B	0.21%	0.37%	1.56%	0.67%	0.28%	01:28:07.24
Llama 3.1 70B	0.54%	1.75%	21.77%	4.69%	0.93%	02:07:16.48

Table 39: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

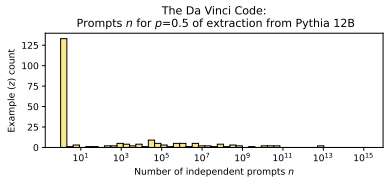

(a) Pythia 12B

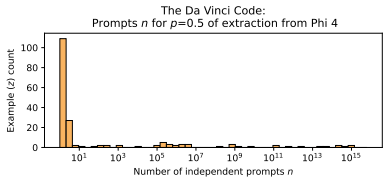

(b) Phi 4 (14B)

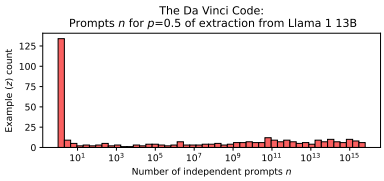

(c) Llama 1 13B

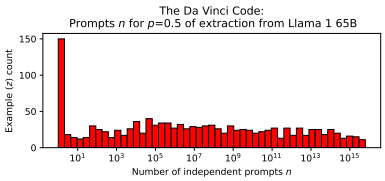

(d) Llama 1 65B

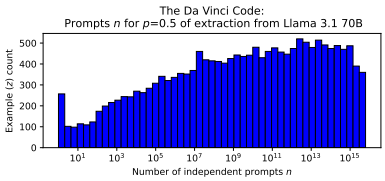

(e) Llama 3.1 70B

Figure 96: *The Da Vinci Code*, Brown (2003): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

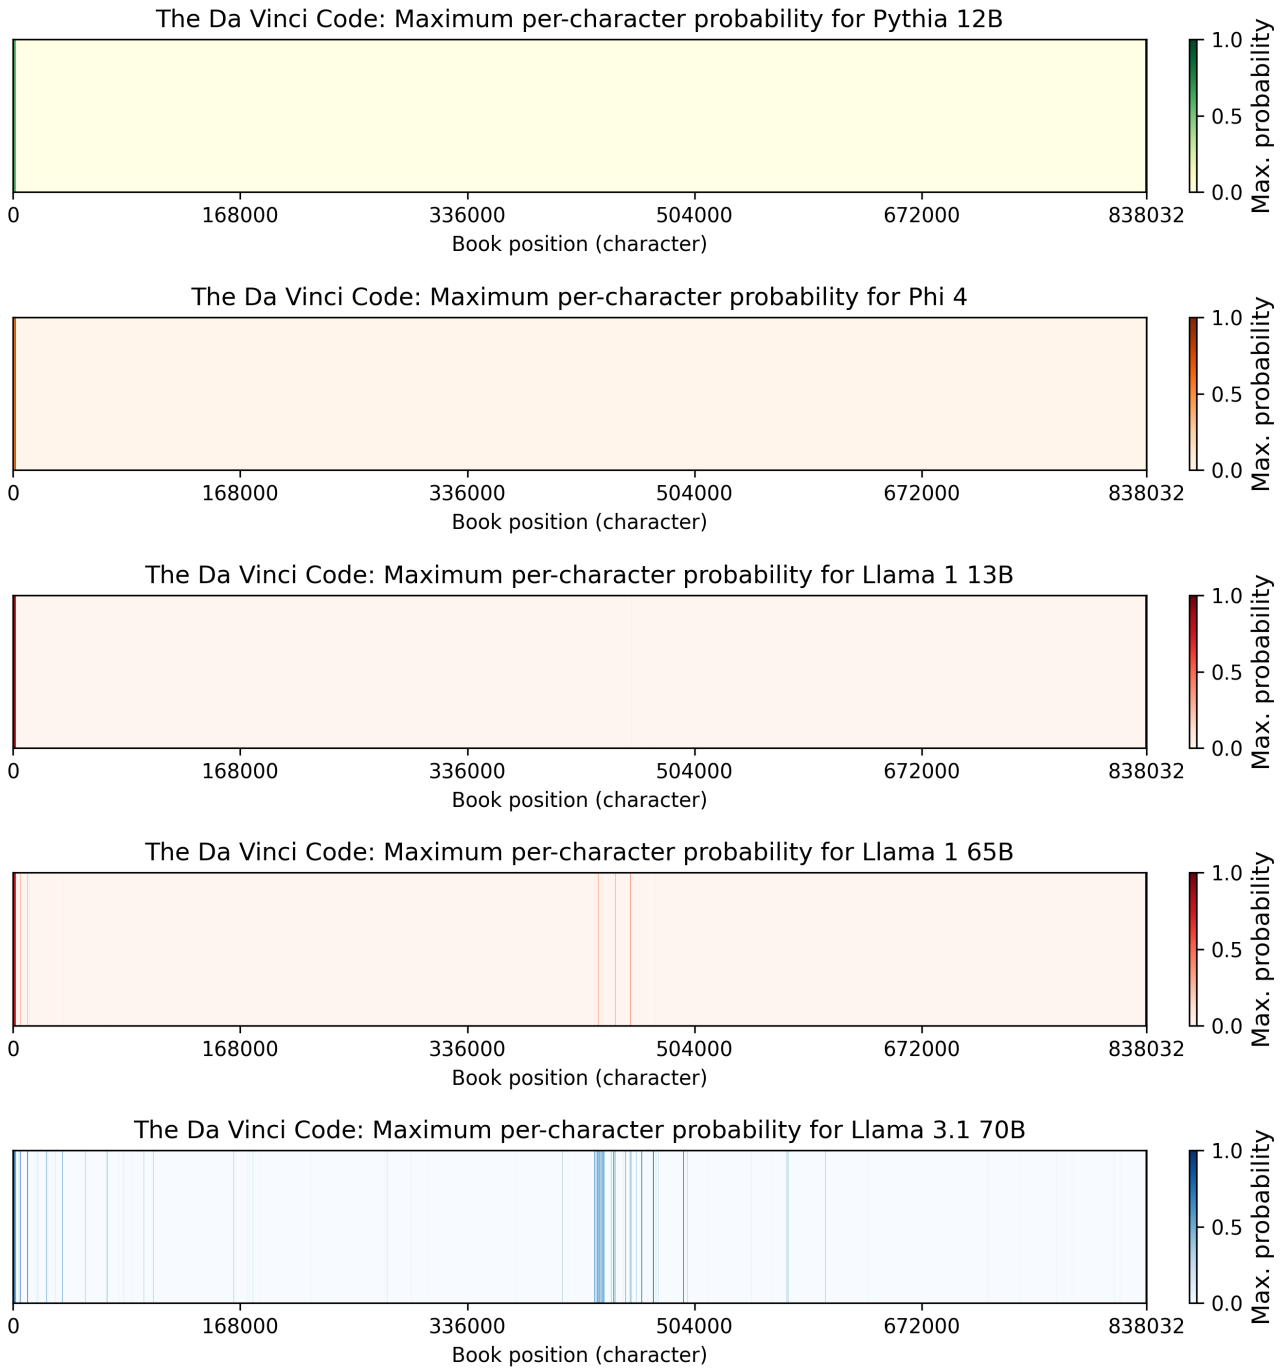

Figure 97: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.35. *Twilight*, MEYER

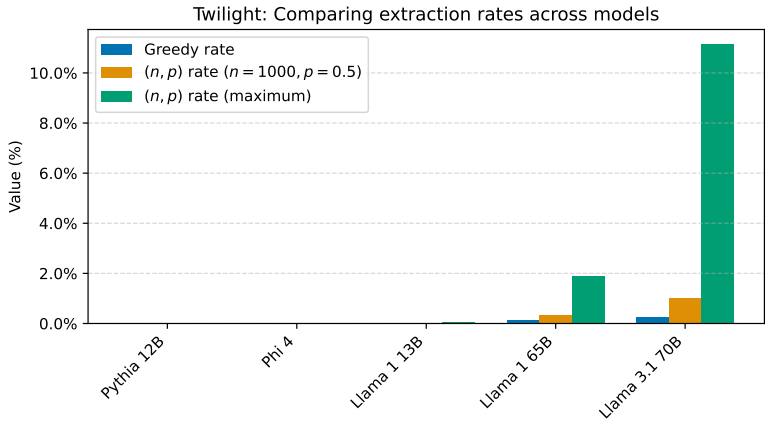

Figure 98: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Twilight						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.00%	0.00%	0.00%	0.00%	0.00%	00:21:00.06
Phi 4 (14B)	0.00%	0.00%	0.01%	0.01%	0.00%	00:34:48.12
Llama 1 13B	0.00%	0.00%	0.07%	0.02%	0.00%	00:23:01.77
Llama 1 65B	0.13%	0.34%	1.90%	0.69%	0.23%	01:13:54.80
Llama 3.1 70B	0.27%	1.00%	11.17%	2.31%	0.58%	01:45:15.56

Table 40: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

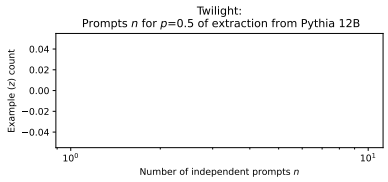

(a) Pythia 12B

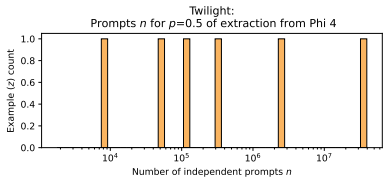

(b) Phi 4 (14B)

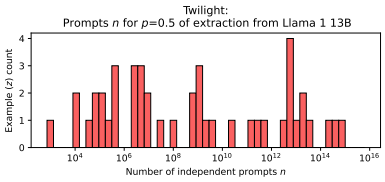

(c) Llama 1 13B

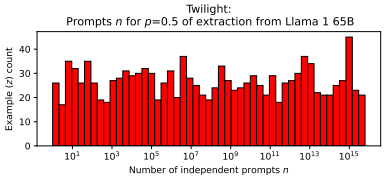

(d) Llama 1 65B

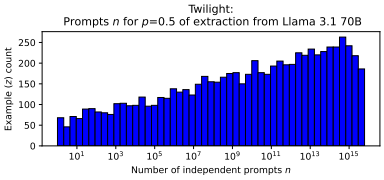

(e) Llama 3.1 70B

Figure 99: *Twilight*, Meyer (2005): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

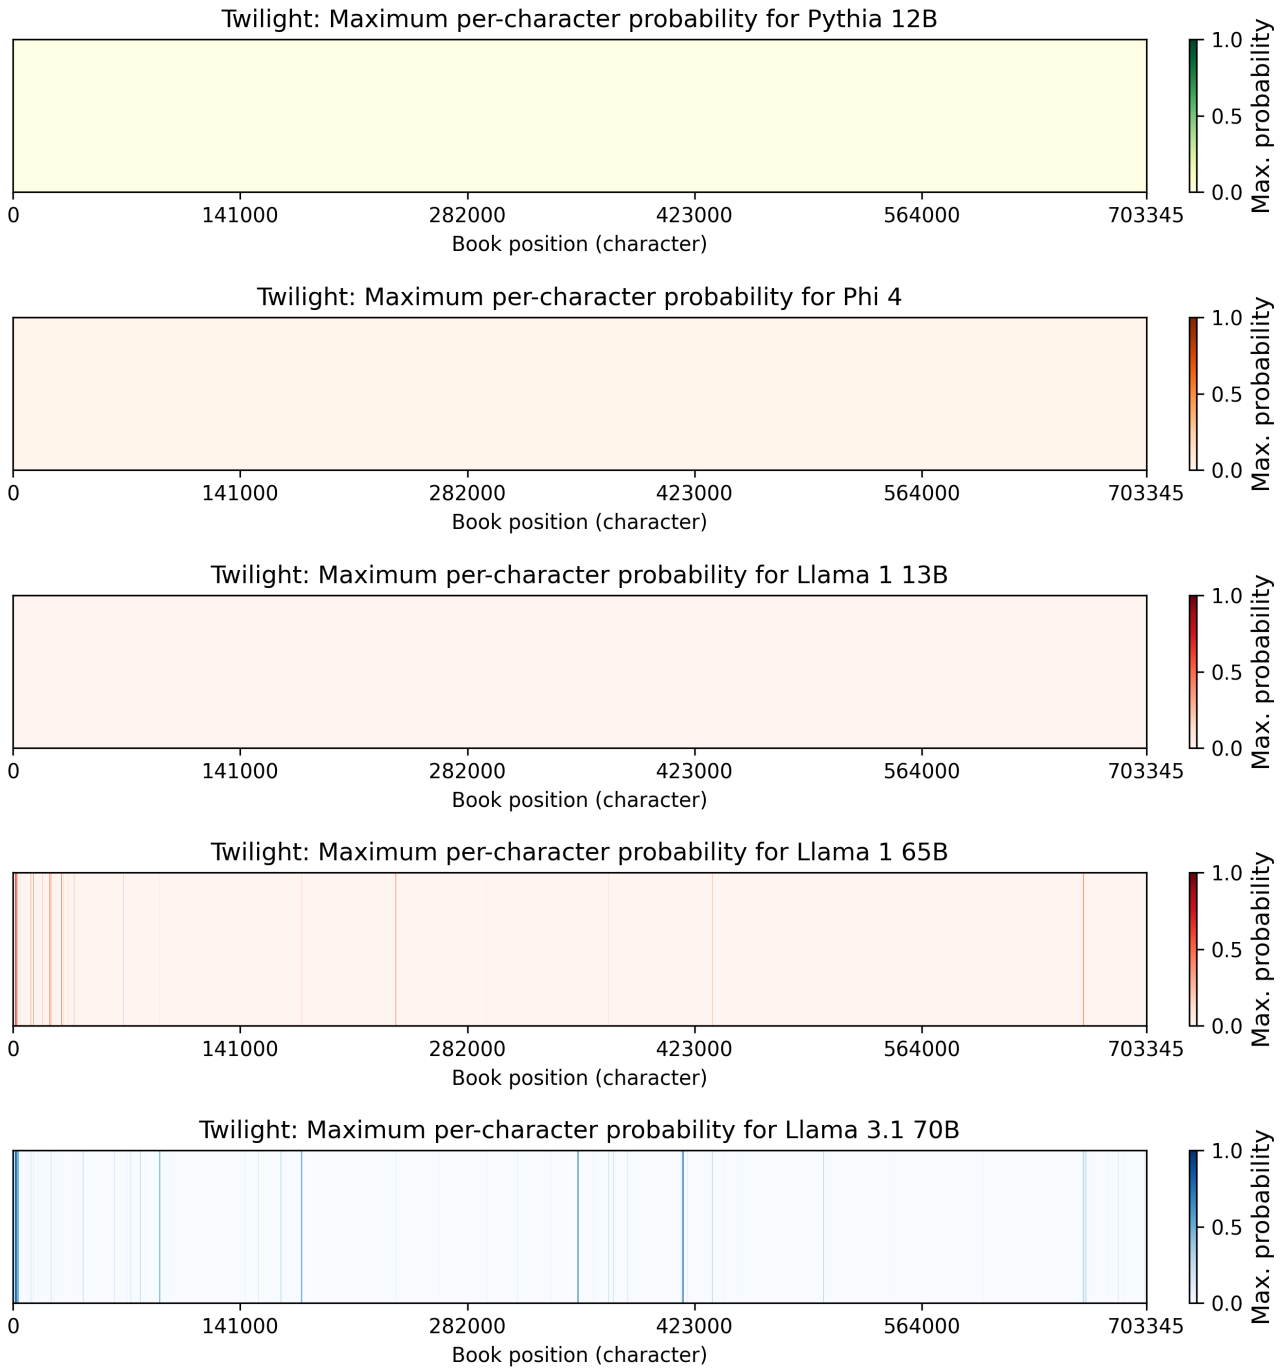

Figure 100: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.1.36. *Ulysses*, JOYCE

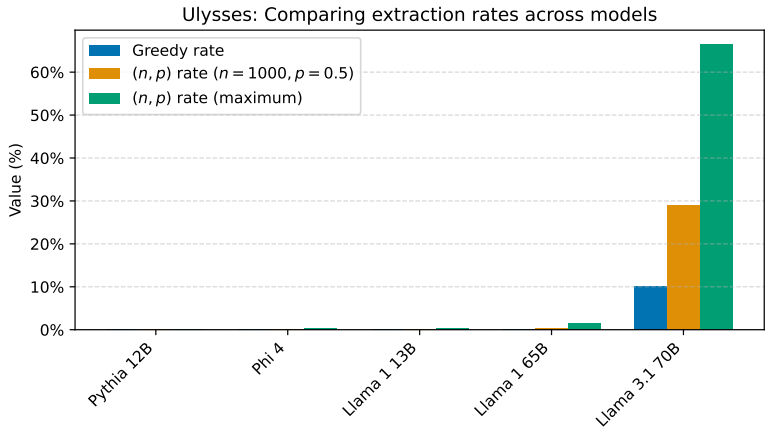

Figure 101: Comparing the greedy discoverable extraction rate, the extraction rate for a very conservative setting for (n, p)-discoverable extraction, and the maximum possible (n, p)-discoverable extraction rate across the five models we test with these experiments.

Ulysses						
Model	**Greedy**	$(n = 1000, p = 0.5)$	**Max.** (n, p)	**%** $p_{\boldsymbol{z}} > 1^{-6}$	**%** $p_{\boldsymbol{z}} > 0.01$	**Runtime (hh:mm:ss)**
Pythia 12B	0.03%	0.04%	0.17%	0.11%	0.04%	00:52:46.60
Phi 4 (14B)	0.07%	0.13%	0.29%	0.21%	0.10%	01:21:46.13
Llama 1 13B	0.03%	0.08%	0.28%	0.18%	0.04%	00:53:08.41
Llama 1 65B	0.13%	0.39%	1.59%	0.73%	0.24%	02:55:55.54
Llama 3.1 70B	10.10%	29.06%	66.46%	46.55%	19.33%	03:52:41.06

Table 41: Summary of extraction rates, proportion of examples where $p_{\boldsymbol{z}}$ is greater than the specified threshold, and experiment runtime.

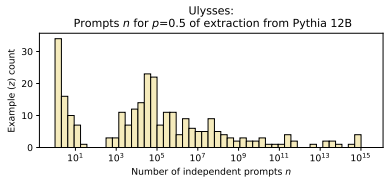

(a) Pythia 12B

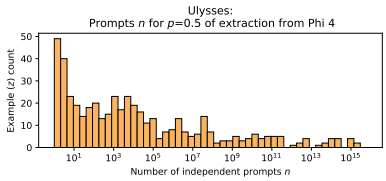

(b) Phi 4 (14B)

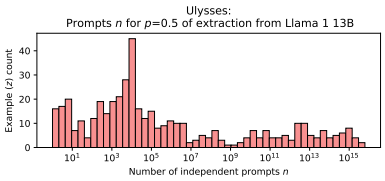

(c) Llama 1 13B

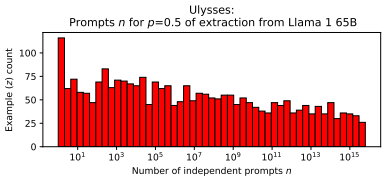

(d) Llama 1 65B

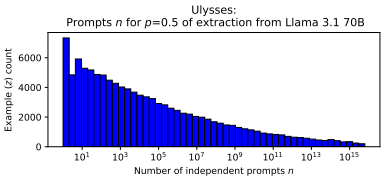

(e) Llama 3.1 70B

Figure 102: *Ulysses*, Joyce (1922): Distributions over independent prompts n needed to extract examples $\boldsymbol{z}$ with $p \geq 0.5$ (i.e., the number of prompts needed such that there is at least a 50% chance the example will be extracted verbatim).

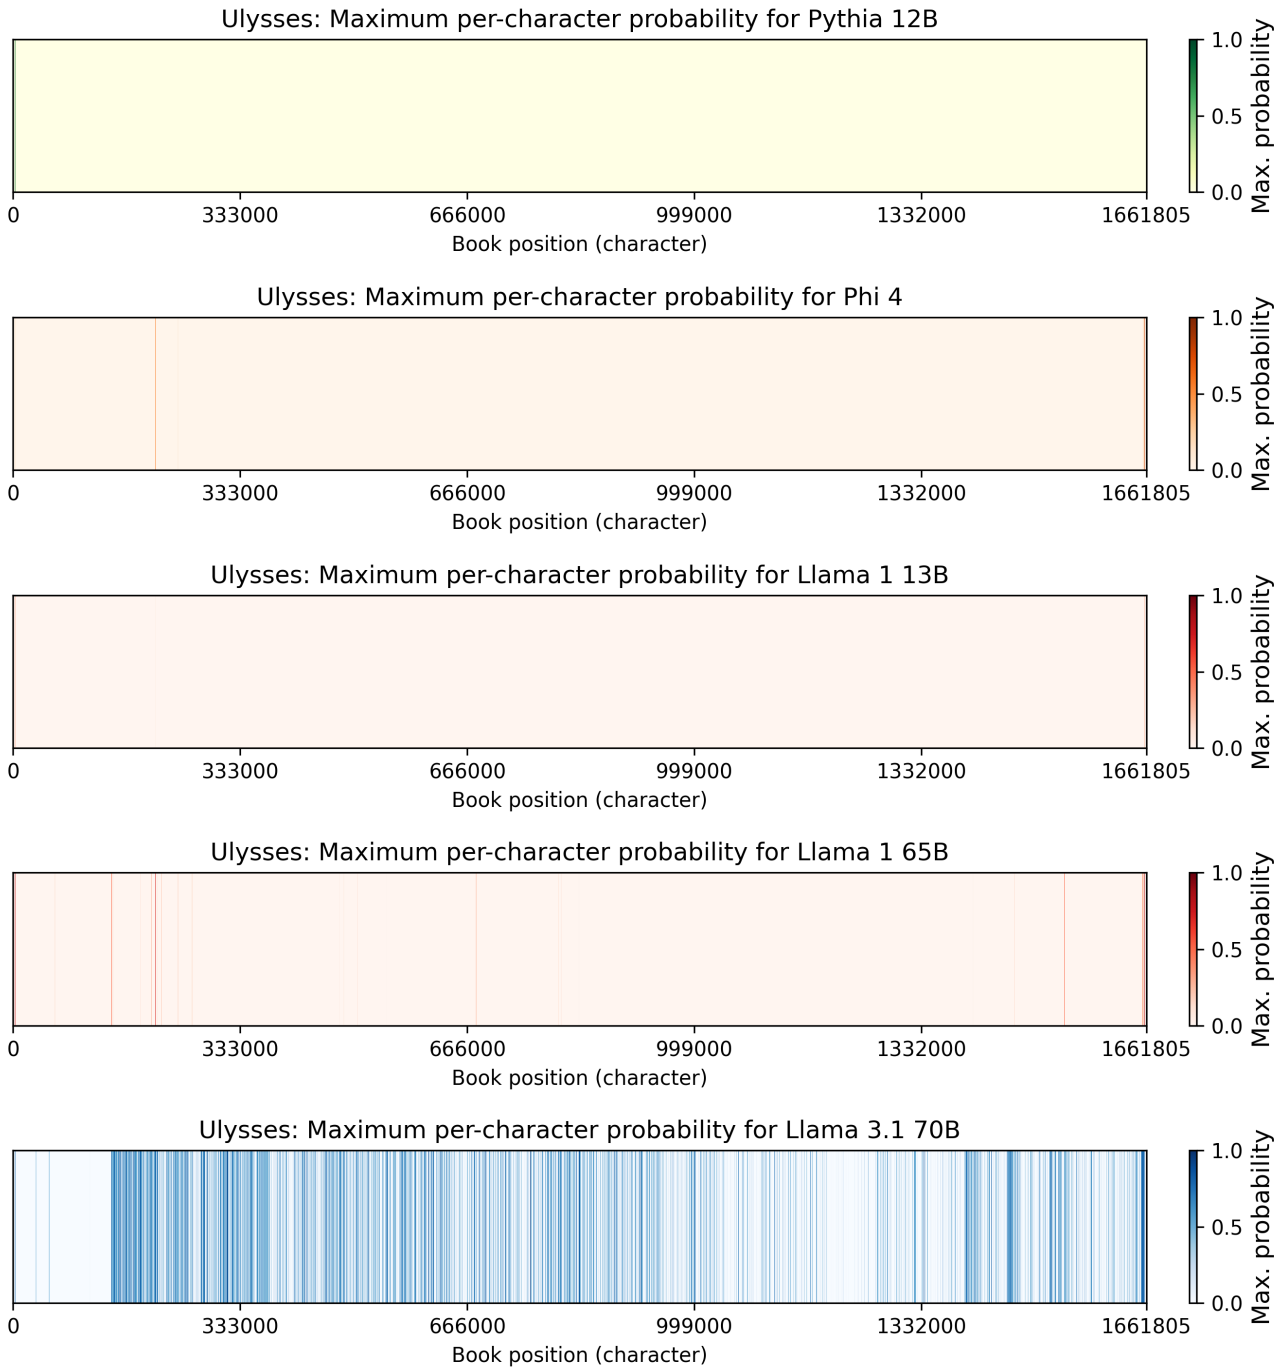

Figure 103: Heatmaps for each model, showing regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments.

G.2. Random subsampling experiments

We also run a small set of experiments that randomly subsamples examples from specific books (starting at a space) to see if this can help us identify memorized areas in a way that is less expensive than our sliding-window experiments. We report summary results in this appendix. Ultimately, the sliding-window approach was more reliable, so we decided to spend most of our effort there. In the future, we will continue to investigate cheaper alternatives.

Book	Llama 1 13B		Llama 3.1 8B		Pythia 12B		Gemma 2 9B		Phi 4	
	# p_z	# $p_z > 0$ (%)	# p_z	# $p_z > 0$ (%)	# p_z	# $p_z > 0$ (%)	# p_z	# $p_z > 0$ (%)	# p_z	# $p_z > 0$ (%)
After I'm Gone (Lippman, 2014)	5174	34 (0.66%)	5168	28 (0.54%)	5173	28 (0.54%)	5167	26 (0.50%)	5166	22 (0.43%)
After Tupac and D Foster (Woodson, 2008)	1541	**40 (2.60%)**	1541	**36 (2.34%)**	1541	**32 (2.08%)**	1541	**36 (2.34%)**	1540	**23 (1.49%)**
Alice's Adventures in Wonderland (Carroll, 1865)	6960	**2465 (35.42%)**	6855	**2579 (37.62%)**	6912	**1544 (22.34%)**	6811	**1575 (23.12%)**	6828	**1523 (22.31%)**
All the Onions (Jacobs, 1999)	599	**47 (7.85%)**	583	**45 (7.72%)**	590	**20 (3.39%)**	579	**38 (6.56%)**	575	**58 (10.09%)**
Aloha from Hell (Kadrey, 2011)	6333	**73 (1.15%)**	6323	39 (0.62%)	6333	51 (0.81%)	6329	43 (0.68%)	6316	29 (0.46%)
Another Bloody Chapter In an Endless Civil War (Wharton, 2016)	10785	**155 (1.44%)**	10288	95 (0.92%)	10528	36 (0.34%)	10181	**107 (1.05%)**	10133	56 (0.55%)
Another Brooklyn (Woodson, 2016)	1294	**35 (2.70%)**	1292	**35 (2.71%)**	1294	**34 (2.63%)**	1292	**32 (2.48%)**	1292	**30 (2.32%)**
Ararat (Golden, 2017)	4973	44 (0.88%)	4912	34 (0.69%)	4954	28 (0.57%)	4905	31 (0.63%)	4884	13 (0.27%)
Blood Oath (Farnsworth, 2010)	5300	**117 (2.21%)**	5265	**70 (1.33%)**	5285	**62 (1.17%)**	5260	**75 (1.43%)**	5254	41 (0.78%)
Blood of the Four (Golden, 2018)	7794	26 (0.33%)	7748	17 (0.22%)	7785	22 (0.28%)	7704	10 (0.13%)	7705	9 (0.12%)
Brown Girl Dreaming (Woodson, 2014)	1691	**29 (1.71%)**	1688	11 (0.65%)	1690	12 (0.71%)	1687	12 (0.71%)	1688	5 (0.30%)
Catch-22 (Heller, 1961)	11043	**346 (3.13%)**	10648	**219 (2.06%)**	10866	**165 (1.52%)**	10259	**231 (2.25%)**	10539	**149 (1.41%)**
Dante and the Origins of Italian Culture (Barolini, 2006)	11313	84 (0.74%)	10937	61 (0.56%)	11112	24 (0.22%)	10672	39 (0.37%)	10843	75 (0.69%)
Dead Ringer (Golden, 2016)	5375	**60 (1.12%)**	5297	37 (0.70%)	5352	39 (0.73%)	5281	43 (0.81%)	5259	27 (0.51%)
Down and Out in the Magic Kingdom (Doctorow, 2003)	2586	**29 (1.12%)**	2580	14 (0.54%)	2584	19 (0.74%)	2577	19 (0.74%)	2573	15 (0.58%)
Drown (Díaz, 1996)	2363	7 (0.30%)	2356	5 (0.21%)	2361	1 (0.04%)	2342	4 (0.17%)	2350	2 (0.09%)
Dungeons and Dragons and Philosophy (Cogburn, 2012)	6489	**188 (2.90%)**	4989	**121 (2.43%)**	5746	**112 (1.95%)**	5158	**74 (1.43%)**	4745	**90 (1.90%)**
Elizabeth Taylor (Bret, 2011)	5738	**114 (1.99%)**	5232	35 (0.67%)	5491	9 (0.16%)	5156	**52 (1.01%)**	5056	15 (0.30%)
Embraced (TerKeurst, 2018)	4032	**265 (6.57%)**	3937	**214 (5.44%)**	4014	**195 (4.86%)**	4001	**206 (5.15%)**	3906	**173 (4.43%)**
Frankenstein (Shelley, 1818)	4770	**2022 (42.39%)**	4537	**1637 (36.08%)**	4646	**776 (16.70%)**	4290	**463 (10.79%)**	4467	**1447 (32.39%)**
Golden Child (Hwang, 1998)	936	**23 (2.46%)**	928	9 (0.97%)	931	**12 (1.29%)**	924	9 (0.97%)	925	3 (0.32%)
Harry Potter and the Sorcerer's Stone (Rowling, 1998)	4440	**576 (12.97%)**	4435	**561 (12.65%)**	4441	**124 (2.79%)**	4439	**432 (9.73%)**	4430	**220 (4.97%)**
Kill the Dead (Kadrey, 2014)	5925	**62 (1.05%)**	5912	28 (0.47%)	5925	39 (0.66%)	5916	59 (1.00%)	5910	20 (0.34%)
Killer Instinct (Finder, 2006)	5782	**129 (2.23%)**	5746	**70 (1.22%)**	5774	**70 (1.21%)**	5748	**88 (1.53%)**	5732	43 (0.75%)
M. Butterfly (Hwang, 1988)	1190	**52 (4.37%)**	1173	**22 (1.88%)**	1188	**36 (3.03%)**	1170	**26 (2.22%)**	1171	**21 (1.79%)**
No Visible Bruises (Snyder, 2019)	6944	**110 (1.58%)**	6132	59 (0.96%)	6465	50 (0.77%)	6262	**69 (1.10%)**	5946	45 (0.76%)
Orlando (Woolf, 1928)	4425	23 (0.52%)	4397	19 (0.43%)	4411	4 (0.09%)	4306	7 (0.16%)	4373	13 (0.30%)
Poison Ink (Golden, 2008)	3610	**60 (1.66%)**	3606	21 (0.58%)	3610	**38 (1.05%)**	3602	34 (0.94%)	3597	18 (0.50%)
Sam the Cat (Klam, 2001)	3260	29 (0.89%)	3253	19 (0.58%)	3260	18 (0.55%)	3254	16 (0.49%)	3248	13 (0.40%)
Sandman Slim (Kadrey, 2009)	5165	**65 (1.26%)**	5164	37 (0.72%)	5165	40 (0.77%)	5163	**54 (1.05%)**	5164	23 (0.45%)
Sunburn (Lippman, 2018)	4169	**45 (1.08%)**	4167	22 (0.53%)	4168	23 (0.55%)	4167	28 (0.67%)	4162	14 (0.34%)
The Beautiful Struggle (Coates, 2009)	2808	1 (0.04%)	2787	2 (0.07%)	2806	0 (0.00%)	2758	3 (0.11%)	2777	0 (0.00%)
The Bedwetter (Silverman, 2010)	3121	**38 (1.22%)**	3077	24 (0.78%)	3115	19 (0.61%)	3083	29 (0.94%)	3061	10 (0.33%)
The Brief Wondrous Life of Oscar Wao (Díaz, 2007)	5001	30 (0.60%)	4990	25 (0.50%)	4997	20 (0.40%)	4961	23 (0.46%)	4979	24 (0.48%)
The Confessions of Max Tivoli (Greer, 2005)	4588	16 (0.35%)	4580	7 (0.15%)	4586	14 (0.31%)	4567	11 (0.24%)	4579	7 (0.15%)
The Diamond Age (Stephenson, 1996)	10358	56 (0.54%)	9875	27 (0.27%)	10206	18 (0.18%)	9537	22 (0.23%)	9706	26 (0.27%)
The Future of the Internet and How to Stop It (Zittrain, 2008)	8385	**335 (4.00%)**	4261	**106 (2.49%)**	5521	51 (0.92%)	4456	**173 (3.88%)**	3956	**90 (2.28%)**
The Great Gatsby (Fitzgerald, 1925)	2620	**799 (30.50%)**	2602	**708 (27.21%)**	2613	**62 (2.37%)**	2589	**395 (15.26%)**	2587	**250 (9.66%)**
The Hobbit (Tolkien, 1937)	4994	**286 (5.73%)**	4987	**197 (3.95%)**	4990	42 (0.84%)	4979	**93 (1.87%)**	4983	**91 (1.83%)**
The Making of a Mediterranean Emirate (Rouighi, 2011)	5734	**313 (5.46%)**	5274	**196 (3.72%)**	5435	**94 (1.73%)**	4939	**113 (2.29%)**	5226	**166 (3.18%)**
The Myth of Sisyphus (Camus, 1955)	3426	**98 (2.86%)**	3125	**101 (3.23%)**	3233	32 (0.99%)	2879	**49 (1.70%)**	3031	**71 (2.34%)**
The Path of Minor Planets (Greer, 2002)	5559	18 (0.32%)	5531	7 (0.13%)	5555	7 (0.13%)	5522	13 (0.24%)	5514	6 (0.11%)
The President's Vampire (Farnsworth, 2011)	4670	**71 (1.52%)**	4612	37 (0.80%)	4661	**49 (1.05%)**	4625	37 (0.80%)	4568	27 (0.59%)
The Water Dancer (Coates, 2019)	6901	**82 (1.19%)**	6895	62 (0.90%)	6899	63 (0.91%)	6885	65 (0.94%)	6887	52 (0.76%)
The Wrong Dead Guy (Kadrey, 2017)	6009	**69 (1.15%)**	5997	35 (0.58%)	6008	41 (0.68%)	5987	32 (0.53%)	5993	26 (0.43%)
Tree and Leaf (Tolkien, 1964)	2192	**115 (5.25%)**	2165	**111 (5.13%)**	2185	19 (0.87%)	2142	**43 (2.01%)**	2153	**70 (3.25%)**
Trying to Find Chinatown (Hwang, 1999)	3324	**90 (2.71%)**	3318	**43 (1.30%)**	3319	**46 (1.39%)**	3315	**47 (1.42%)**	3318	26 (0.78%)
Unglued (TerKeurst, 2012)	2956	**132 (4.47%)**	2833	**98 (3.46%)**	2909	**84 (2.89%)**	2870	**94 (3.28%)**	2805	**90 (3.21%)**
We Were Eight Years in Power (Coates, 2017)	6641	**173 (2.61%)**	5303	**112 (2.11%)**	5885	**62 (1.05%)**	5407	**76 (1.41%)**	5017	**74 (1.47%)**
Who Is Rich (Klam, 2017)	5042	**57 (1.13%)**	4971	44 (0.89%)	5025	**52 (1.03%)**	4959	**53 (1.07%)**	4936	43 (0.87%)

Table 42: The number of samples for which extraction probabilities are computed (# p_z) and the number of extraction probabilities > 0 (# $p_z > 0$) for each book at model. These are our "cheap" experiments, where we sample a small number of examples randomly from each of these 50 books, where the number of examples is proportional to the length of the book (1% of the character length). Examples can overlap. We show the percentage of $p_z > 0$, and highlight cases where this percentage is greater than 1% in blue, greater than 5% in purple, and greater than 10% in red.

Book	Llama 1 13B		Llama 3.1 8B		Pythia 12B		Gemma 2 9B		Phi 4	
	# p_z	# $p_z > .0001$ (%)	# p_z	# $p_z > .0001$ (%)	# p_z	# $p_z > 0.0001$ (%)	# p_z	# $p_z > .0001$ (%)	# p_z	# $p_z > .0001$ (%)
After I'm Gone (Lippman, 2014)	5174	0 (0.00%)	5168	2 (0.04%)	5173	5 (0.10%)	5167	1 (0.02%)	5166	1 (0.02%)
After Tupac and D Foster (Woodson, 2008)	1541	12 (0.78%)	1541	9 (0.58%)	1541	12 (0.78%)	1541	3 (0.19%)	1540	3 (0.19%)
Alice's Adventures in Wonderland (Carroll, 1865)	6960	**491 (7.05%)**	6855	**536 (7.82%)**	6912	38 (0.55%)	6811	**246 (3.61%)**	6828	**202 (2.96%)**
All the Onions (Jacobs, 1999)	599	0 (0.00%)	583	0 (0.00%)	590	**7 (1.19%)**	579	1 (0.17%)	575	0 (0.00%)
Aloha from Hell (Kadrey, 2011)	6333	3 (0.05%)	6323	5 (0.08%)	6333	9 (0.14%)	6329	1 (0.02%)	6316	0 (0.00%)
Another Bloody Chapter In an Endless Civil War (Wharton, 2016)	10785	10 (0.09%)	10288	15 (0.15%)	10528	10 (0.09%)	10181	11 (0.11%)	10133	14 (0.14%)
Another Brooklyn (Woodson, 2016)	1294	**23 (1.78%)**	1292	**22 (1.70%)**	1294	**29 (2.24%)**	1292	**24 (1.86%)**	1292	**19 (1.47%)**
Ararat (Golden, 2017)	4973	8 (0.16%)	4912	5 (0.10%)	4954	15 (0.30%)	4905	7 (0.14%)	4884	3 (0.06%)
Blood Oath (Farnsworth, 2010)	5300	15 (0.28%)	5265	8 (0.15%)	5285	16 (0.30%)	5260	4 (0.08%)	5254	2 (0.04%)
Blood of the Four (Golden, 2018)	7794	3 (0.04%)	7748	5 (0.06%)	7785	8 (0.10%)	7704	4 (0.05%)	7705	3 (0.04%)
Brown Girl Dreaming (Brown et al., 2021)	1691	2 (0.12%)	1688	0 (0.00%)	1690	2 (0.12%)	1687	0 (0.00%)	1688	0 (0.00%)
Catch-22 (Heller, 1961)	11043	20 (0.18%)	10648	20 (0.19%)	10866	13 (0.12%)	10259	8 (0.08%)	10539	10 (0.09%)
Dante and the Origins of Italian Culture (Barolini, 2006)	11313	1 (0.01%)	10937	2 (0.02%)	11112	2 (0.02%)	10672	3 (0.03%)	10843	6 (0.06%)
Dead Ringer (Golden, 2016)	5375	8 (0.15%)	5297	7 (0.13%)	5352	12 (0.22%)	5281	10 (0.19%)	5259	5 (0.10%)
Down and Out in the Magic Kingdom (Doctorow, 2003)	2586	6 (0.23%)	2580	3 (0.12%)	2584	6 (0.23%)	2577	3 (0.12%)	2573	0 (0.00%)
Drown (Díaz, 1996)	2363	0 (0.00%)	2356	0 (0.00%)	2361	0 (0.00%)	2342	0 (0.00%)	2350	0 (0.00%)
Dungeons and Dragons and Philosophy (Cogburn, 2012)	6489	17 (0.26%)	4989	9 (0.18%)	5746	23 (0.40%)	5158	7 (0.14%)	4745	10 (0.21%)
Elizabeth Taylor (Bret, 2011)	5738	1 (0.02%)	5232	0 (0.00%)	5491	1 (0.02%)	5156	0 (0.00%)	5056	0 (0.00%)
Embraced (TerKeurst, 2018)	4032	**82 (2.03%)**	3937	**71 (1.80%)**	4014	**75 (1.87%)**	4001	**67 (1.67%)**	3906	**45 (1.15%)**
Frankenstein (Shelley, 1818)	4770	**191 (4.00%)**	4537	**174 (3.84%)**	4646	13 (0.28%)	4290	**58 (1.35%)**	4467	**384 (8.60%)**
Golden Child (Hwang, 1998)	936	0 (0.00%)	928	0 (0.00%)	931	4 (0.43%)	924	0 (0.00%)	925	0 (0.00%)
Harry Potter and the Sorcerer's Stone (Rowling, 1998)	4440	37 (0.83%)	4435	38 (0.86%)	4441	5 (0.11%)	4439	**54 (1.22%)**	4430	33 (0.74%)
Kill the Dead (Kadrey, 2014)	5925	3 (0.05%)	5912	3 (0.05%)	5925	7 (0.12%)	5916	3 (0.05%)	5910	2 (0.03%)
Killer Instinct (Finder, 2006)	5782	12 (0.21%)	5746	9 (0.16%)	5774	11 (0.19%)	5748	9 (0.16%)	5732	7 (0.12%)
M. Butterfly (Hwang, 1988)	1190	4 (0.34%)	1173	5 (0.43%)	1188	**12 (1.01%)**	1170	4 (0.34%)	1171	0 (0.00%)
No Visible Bruises (Snyder, 2019)	6944	20 (0.29%)	6132	20 (0.33%)	6465	21 (0.32%)	6262	20 (0.32%)	5946	20 (0.34%)
Orlando (Woolf, 1928)	4425	0 (0.00%)	4397	2 (0.05%)	4411	0 (0.00%)	4306	0 (0.00%)	4373	2 (0.05%)
Poison Ink (Golden, 2008)	3610	3 (0.08%)	3606	0 (0.00%)	3610	1 (0.03%)	3602	1 (0.03%)	3597	0 (0.00%)
Sam the Cat (Klam, 2001)	3260	2 (0.06%)	3253	1 (0.03%)	3260	4 (0.12%)	3254	1 (0.03%)	3248	0 (0.00%)
Sandman Slim (Kadrey, 2009)	5165	3 (0.06%)	5164	4 (0.08%)	5165	4 (0.08%)	5163	3 (0.06%)	5164	2 (0.04%)
Sunburn (Lippman, 2018)	4169	7 (0.17%)	4167	7 (0.17%)	4168	9 (0.22%)	4167	9 (0.22%)	4162	6 (0.14%)
The Beautiful Struggle (Coates, 2009)	2808	0 (0.00%)	2787	0 (0.00%)	2806	0 (0.00%)	2758	0 (0.00%)	2777	0 (0.00%)
The Bedwetter (Silverman, 2010)	3121	5 (0.16%)	3077	3 (0.10%)	3115	8 (0.26%)	3083	4 (0.13%)	3061	2 (0.07%)
The Brief Wondrous Life of Oscar Wao (Díaz, 2007)	5001	11 (0.22%)	4990	10 (0.20%)	4997	10 (0.20%)	4961	8 (0.16%)	4979	0 (0.00%)
The Confessions of Max Tivoli (Greer, 2005)	4588	0 (0.00%)	4580	0 (0.00%)	4586	1 (0.02%)	4567	0 (0.00%)	4579	0 (0.00%)
The Diamond Age (Stephenson, 1996)	10358	2 (0.02%)	9875	2 (0.02%)	10206	1 (0.01%)	9537	0 (0.00%)	9706	2 (0.02%)
The Future of the Internet and How to Stop It (Zittrain, 2008)	8385	13 (0.16%)	4261	9 (0.21%)	5521	11 (0.20%)	4456	6 (0.13%)	3956	5 (0.13%)
The Great Gatsby (Fitzgerald, 1925)	2620	**141 (5.38%)**	2602	**107 (4.11%)**	2613	10 (0.38%)	2589	**97 (3.75%)**	2587	**85 (3.29%)**
The Hobbit (Tolkien, 1937)	4994	33 (0.66%)	4987	29 (0.58%)	4990	0 (0.00%)	4979	12 (0.24%)	4983	25 (0.50%)
The Making of a Mediterranean Emirate (Rouighi, 2011)	5734	25 (0.44%)	5274	26 (0.49%)	5435	29 (0.53%)	4939	28 (0.57%)	5226	23 (0.44%)
The Myth of Sisyphus (Camus, 1955)	3426	28 (0.82%)	3125	**46 (1.47%)**	3233	21 (0.65%)	2879	26 (0.90%)	3031	**32 (1.06%)**
The Path of Minor Planets (Greer, 2002)	5559	5 (0.09%)	5531	0 (0.00%)	5555	3 (0.05%)	5522	0 (0.00%)	5514	0 (0.00%)
The President's Vampire (Farnsworth, 2011)	4670	9 (0.19%)	4612	4 (0.09%)	4661	15 (0.32%)	4625	4 (0.09%)	4568	1 (0.02%)
The Water Dancer (Coates, 2019)	6901	36 (0.52%)	6895	38 (0.55%)	6899	37 (0.54%)	6885	38 (0.55%)	6887	36 (0.52%)
The Wrong Dead Guy (Kadrey, 2017)	6009	6 (0.10%)	5997	5 (0.08%)	6008	11 (0.18%)	5987	5 (0.08%)	5993	4 (0.07%)
Tree and Leaf (Tolkien, 1964)	2192	**34 (1.55%)**	2165	**39 (1.80%)**	2185	11 (0.50%)	2142	14 (0.65%)	2153	18 (0.84%)
Trying to Find Chinatown (Hwang, 1999)	3324	0 (0.00%)	3318	0 (0.00%)	3319	1 (0.03%)	3315	0 (0.00%)	3318	0 (0.00%)
Unglued (TerKeurst, 2012)	2956	**43 (1.45%)**	2833	**39 (1.38%)**	2909	**30 (1.03%)**	2870	**31 (1.08%)**	2805	12 (0.43%)
We Were Eight Years in Power (Coates, 2017)	6641	26 (0.39%)	5303	25 (0.47%)	5885	23 (0.39%)	5407	22 (0.41%)	5017	19 (0.38%)
Who Is Rich (Klam, 2017)	5042	37 (0.73%)	4971	35 (0.70%)	5025	41 (0.82%)	4959	36 (0.73%)	4936	34 (0.69%)

Table 43: The number of samples for which extraction probabilities are computed (# p_z) and the number of extraction probabilities > 0.0001 (# $p_z > 0.0001$) for each book at model. These are our "cheap" experiments, where we sample a small number of examples randomly from each of these 50 books, where the number of examples is proportional to the length of the book (1% of the character length). Examples can overlap. We show the percentage of $p_z > 0.0001$, and highlight cases where this percentage is greater than 1% in blue, greater than 5% in purple, and greater than 10% in red.

G.3. Additional results on *Harry Potter*

We dig a bit deeper into the particularly high probabilistic extraction and memorization results that we obtained for *Harry Potter and the Sorcerer's Stone* (Rowling, 1998) for LLAMA 3.1 70B. We compare results for the 100-token sequences that we sample using a sliding window of 10 characters across the whole book (see Appendix G.1) for different suffix start positions, which yield different prefix and suffix lengths. We generally present results using the standard split of 50-token prefix and 50-token suffix. Here we include results for 25-token prefix/ 75-token suffix and 75-token prefix / 25-token suffix.

As expected, using a prefix length of 25 tokens and suffix length of 75 tokens results in lower values of extraction. This is likely due to the combined effects of conditioning on less context (only 25 tokens) and the fact that we are attempting to extract a relatively much longer sequence (75 tokens, in comparison to 50). Nevertheless, the amount of extraction remains quite high. Conversely, using a prefix length of 75 tokens and a suffix length of 25 tokens has the reverse overall effect. **We caution that a suffix length of 25 tokens is below the standard minimum used in the literature for extraction and memorization claims. We are conservative in our analysis and do not make claims about extraction and memorization for the figures associated with this setting. We include it as a sanity check (i.e., to show that these settings exhibit the trend we expect, given how we have changed the prefix and suffix lengths.**

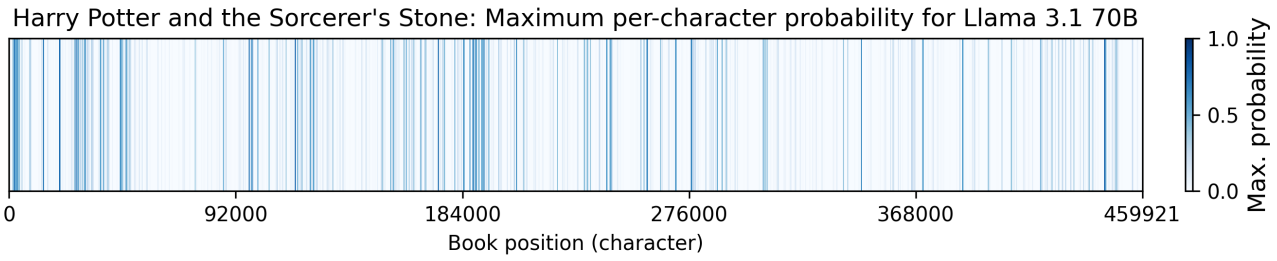

(a) Prefix-length: 25 tokens; suffix-length 75 tokens

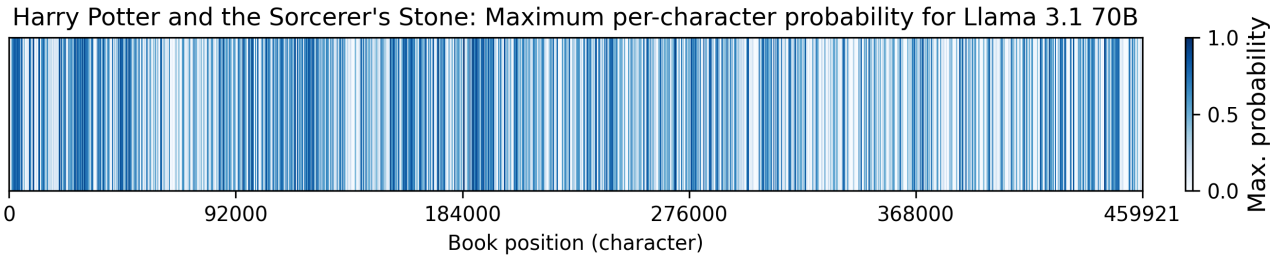

(b) Prefix-length: 50 tokens; suffix-length 50 tokens

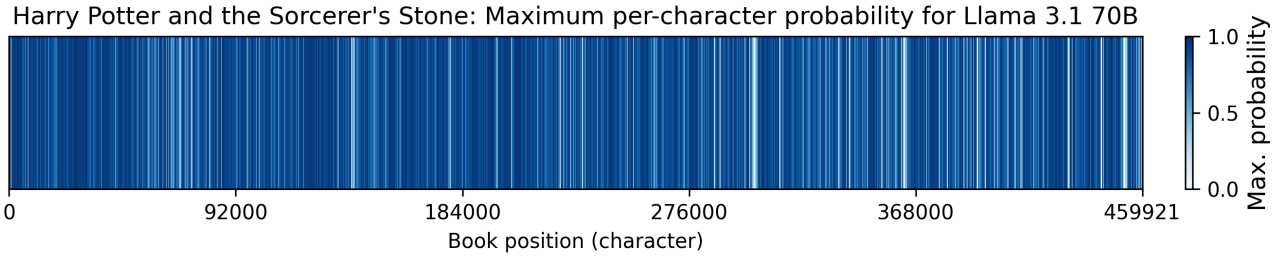

(c) Prefix-length: 75 tokens; suffix-length 25 tokens

Figure 104: Comparing heatmaps for LLAMA 3.1 70B for different prefix and suffix lengths for 100-token examples. These heatmaps show regions of extraction "hot-spots" according to location (character position) in the book. To highlight these regions, for overlapping examples in the sliding window, we find (and plot) the highest extraction probability at the particular character location. The goal here is to identify high-probability "hot-spots," for which we further investigate the underlying examples in follow-up experiments. **We are conservative in our analysis, and do not make claims about extraction and memorization for the bottom figure, which has a relatively small suffix length. We include it as a sanity check (i.e., to show that it exhibits the trend we expect, given how we have changed the prefix and suffix lengths).**

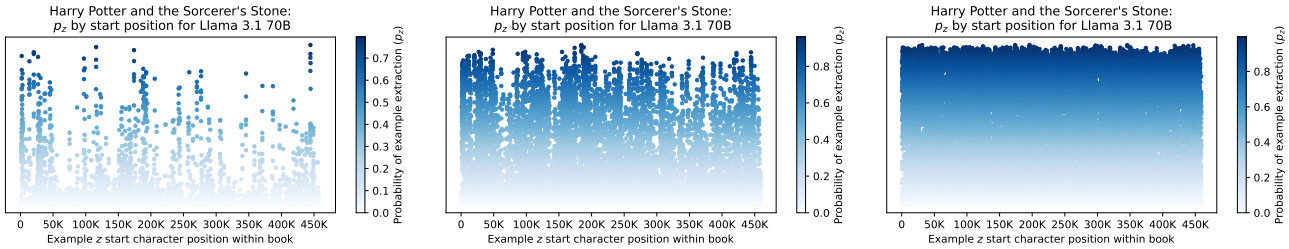

(a) 25/75 prefix/suffix split (b) 50/50 prefix/suffix split (c) 75/25 prefix/suffix split

Figure 105: As alternative view for heatmaps, which collapse overlapping examples into a single one-dimensional band per character, we plot the probability of extraction for each example $\boldsymbol{z}$, according to its start position within the overall book. These start positions are unique, so each example has a dot. Reading along the x-axis, the band of blue (for the given y-axis value, indicating the probability) shows how many examples were extracted with that probability for the entire book. White gaps along the x-axis indicate that an example (with that start position) was not extracted (i.e., has 0 probability) A continuous band of blue (of that particular shade, associated with the probability on the y-axis) effectively indicates that we can extract the entire book with that associated probability. **We are conservative in our analysis, and do not make claims about extraction and memorization for the third figure, on the right. We include it as a sanity check (i.e., to show that it exhibits the trend we expect, given how we have changed the prefix and suffix lengths).**

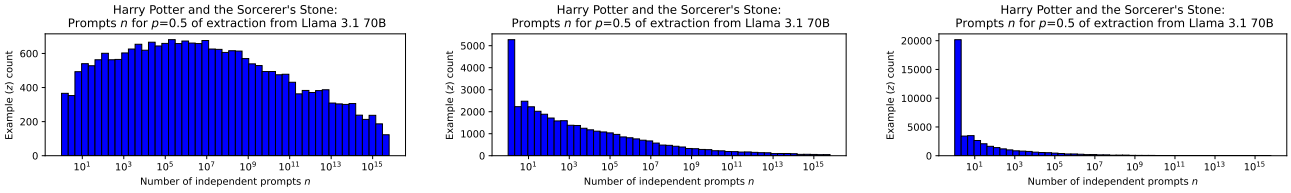

(a) 25/75 prefix/suffix split (b) 50/50 prefix/suffix split (c) 75/25 prefix/suffix split

Figure 106: As alternative view, we plot the distribution of the number of independent prompts n one would need to extract examples with at least probability $p = 0.5$ (i.e., a coin flip) for all of the sliding-window samples. which collapse overlapping examples into a single one-dimensional band per character, we plot the probability of extraction for each example $\boldsymbol{z}$, according to its start position within the overall book. **We are conservative in our analysis, and do not make claims about extraction and memorization for the third figure, on the right. We include it as a sanity check (i.e., to show that it exhibits the trend we expect, given how we have changed the prefix and suffix lengths).**

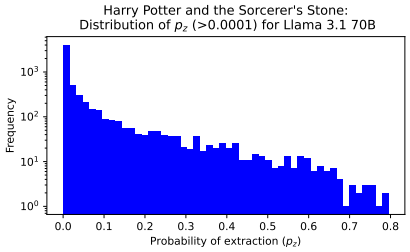

(a) 25/75 prefix/suffix split

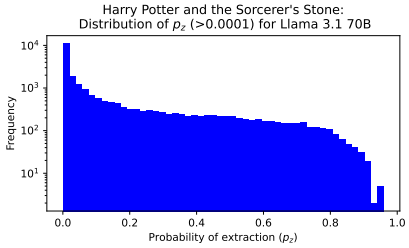

(b) 50/50 prefix/suffix split

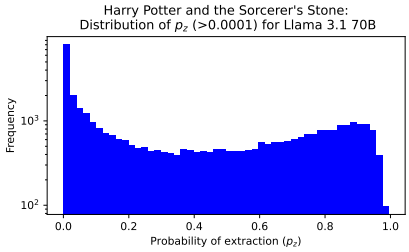

(c) 75/25 prefix/suffix split

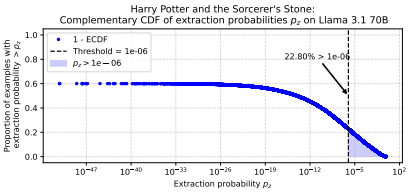

(d) 25/75 prefix/suffix split

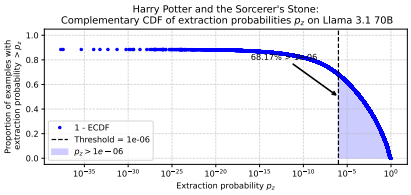

(e) 50/50 prefix/suffix split

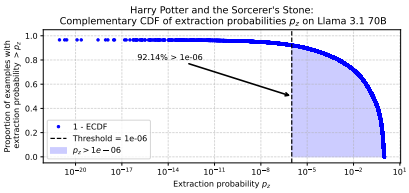

(f) 75/25 prefix/suffix split

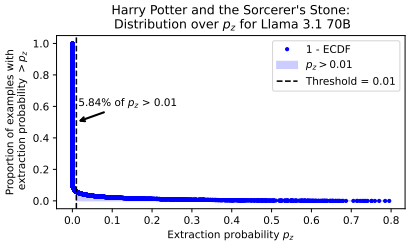

(g) 25/75 prefix/suffix split

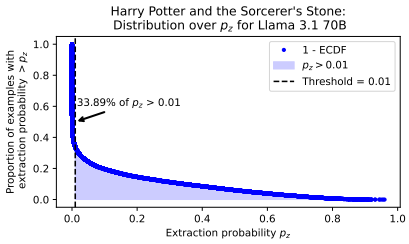

(h) 50/50 prefix/suffix split

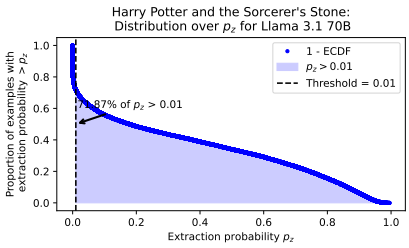

(i) 75/25 prefix/suffix split

Figure 107: As alternative view, we plot different ways to visualize the distribution over per-example probabilities $p_{\boldsymbol{z}}$. The top row shows the distribution over $p_v z$ (limited to $p_{\boldsymbol{z}} > 0.0001$) for different prefix and suffix lengths. The middle shows the entire complementary CDF (CCDF) for the distribution over $p_{\boldsymbol{z}}$ (indicating the number of examples whose $p_{\boldsymbol{z}} > .0001$, which covers the same $p_v z$ that are plotted in the histograms in the top row). The bottom row shows a zoomed-in version of the CCDF, and indicates the $p_{\boldsymbol{z}} > 0.01$. **We are conservative in our analysis, and do not make claims about extraction and memorization for the third column of figures, on the right. We include these plots as a sanity check (i.e., to show that they exhibits the trend we expect, given how we have changed the prefix and suffix lengths).**

H. Selection of extracted examples

This project generated tens of thousands of extracted examples. We could not evaluate them all manually. We provide a brief selection of arbitrarily chosen extracted examples from some of the 36 books in the sliding window experiments. (See Section 4 & Appendix G.1). We defer discussion of these examples (and our book-level results) to Appendix I. We also refer to some of these examples in Section 4. Unless otherwise specified, all examples have a prefix length of 50 tokens and a suffix length of 50 tokens.

Almost all books indicate high-probability extraction for the copyright notice. This is highly duplicated text, so it is not necessarily the case that the high extraction probability/ memorization is *due to* the exact text from which we extract it. It is still memorized, under this definition of memorization. The relationship between duplication and memorization is well-documented (Lee et al., 2022). Similarly, publisher addresses, authors biographies (which are also highly duplicated text on the internet) also show up as extractable for many books. We provide some non-exhaustive examples of this for some books.

***Another Brooklyn* (Woodson, 2016)**

LLAMA 3.1 70B

(Copyright notice)

Prefix: 'RIGHT\n\n ANOTHER BROOKLYN. Copyright © 2016 by Jacqueline Woodson. All rights reserved under International and Pan-American Copyright Conventions. By payment of the required fees, you have been granted the nonexclusive, nontransferable'

Suffix: ' right to access and read the text of this e-book on-screen. No part of this text may be reproduced, transmitted, downloaded, decompiled, reverse-engineered, or stored in or introduced into any information storage and retrieval system, in any form'

Sequence probability (p_z): 0.825882886427520

(Publisher address)

Prefix: 'treet\n\nSydney, NSW 2000, Australia\n\nwww.harpercollins.com.au\n\n**Canada**\n\n HarperCollins Canada\n\n2 Bloor Street East - 20th Floor\n\nToronto, ON M4W 1'

Suffix: ' A8, Canada\n\nwww.harpercollins.ca\n\n**New Zealand**\n\nHarperCollins Publishers New Zealand\n\nUnit D1, 63 Apollo Drive\n\nRosedale 0632\n\nAuckland, New Zealand\n\nwww.harpercoll'

Sequence probability (p_z): 0.12346602702119958

(Author biography)

Prefix: 'T THE AUTHOR\n\n**JACQUELINE WOODSON** is the bestselling author of more than two dozen award-winning books for young adults, middle graders, and children, including the_New York_ _Times_ bestselling memoir _Brown'

Suffix: ' Girl Dreaming_ _,_ which won the 2014 National Book Award, the Coretta Scott King Award, a Newbery Honor Award, an NAACP Image Award, and the Sibert Honor Award. Woodson was recently named the'

Sequence probability (p_z): 0.0005176228468961067

(Book text)

Prefix: '. \n 182. a\n 183. b\n\n##\n\nFor a long time, my mother wasn't dead yet. Mine could have been a more tragic story. My father could have given in to the bottle or the needle or'

Suffix: ' a woman and left my brother and me to care for ourselves—or worse, in the care of New York City Children's Services, where, my father said, there was seldom a happy ending. But this didn't happen. I know now that what"

Sequence probability (p_z): $3.3700513791944576 \times 10^{-5}$

***Dante and the Origins of Italian Culture* (Barolini, 2006)**

LLAMA 3.1 70B

(End of prefix is from the Bible; suffix is a continuation of that quote)

Prefix: 'erpetual death, of perpetual alienation from the life of God, as he explains in _Enchiridion_ 113:\n\nThis perpetual death of the wicked, then, that is, their alienation from the life of God'

Suffix: ', shall abide for ever, and shall be common to them all, whatever men, prompted by their human affections, may conjecture as to a variety of punishments, or as to a mitigation or intermission of their woes; just as the eternal"

Sequence probability (p_z): 0.32557161294347686

(Quote from Dante's *Convivio*)

Prefix: 'ioli beni le paiono grandi, e però da quelli comincia prima a desiderare. Onde vederno li parvuli desiderare massimamente un pomo; e poi, più procedendo,'

Suffix: ' desiderare uno augellino; e poi, più oltre, desiderare bel vestimento; e poi lo cavallo; e poi una donna; e poi ricchezza non grande, e poi grande, e poi più. E questo"

Sequence probability (p_z): 0.018595474111579826

(End of prefix is from Dante's *Purgatorio*; suffix is a continuation of that quote)

Prefix: 'f these verses in _Purgatorio_ 16\' s description of the newborn soul as a young female child setting forth on the path of life: "Esce di mano a lui che la vagheggia / prima che sia, a'

Suffix: ' guisa di fanciulla / che piangendo e ridendo pargoleggia, / l'anima semplicetta che sa nulla, / salvo che, mossa da lieto fattore, / volontier torna a"

Sequence probability (p_z): 0.013580356347353521

***Dungeons and Dragons and Philosophy* (Cogburn, 2012).**

LLAMA 3.1 70B

(End of prefix is a quote from John Stuart Mill; suffix is a continuation of that quote)

Prefix: 'w the utilitarian John Stuart Mill (1806–73, lawful neutral) describes morality:\n\n> The motive has nothing to do with the morality of the action, though much with the worth of the agent. He who saves a fellow creature from'

Suffix: ' drowning does what is morally right, whether his motive be duty, or the hope of being paid for his trouble; he who betrays the friend that trusts him, is guilty of a crime, even if his object be to serve another friend to whom'

Sequence probability (p_z): 0.29670331942430045

(Quote from "Pathfinder Roleplaying Game Core Rulebook")

Prefix: 'swers are usually brief, cryptic, or repetitive. If the creature's alignment was different from yours, the corpse gets a Will save to resist the spell as if it were alive. If the corpse has been subject to speak with dead within the'

Suffix: ' past week, the new spell fails. You can cast this spell on a corpse that has been deceased for any amount of time, but the body must be mostly intact to be able to respond. A damaged corpse may be able to give partial answers or'

Sequence probability (p_z): 0.26730561749052373

Embraced **(TerKeurst, 2018)**

LLAMA 3.1 70B

(End of prefix is a quote from the Bible; suffix is a continuation of that quote)

Prefix: ‘e grabby-ness of our flesh.\n\nOne of the most beautiful descriptions of the fullness of God is found in Paul’s prayer for the Ephesians:\n\nFor this reason I kneel before the Father, from whom every family in heaven’

Suffix: ‘ and on earth derives its name. I pray that out of his glorious riches he may strengthen you with power through his Spirit in your inner being, so that Christ may dwell in your hearts through faith. And I pray that you, being rooted and established’

Sequence probability (p_z): 0.7407965314704028

Harry Potter and the Sorcerer’s Stone **(Rowling, 1998)**

LLAMA 3.1 70B

Prefix: ‘A magic beyond all we do here! And now, bedtime. Off you trot!”\n\nThe Gryffindor first years followed Percy through the chattering crowds, out of the Great Hall, and up the marble staircase. Harry\’—s legs were like’

Suffix: ‘ lead again, but only because he was so tired and full of food. He was too sleepy even to be surprised that the people in the portraits along the corridors whispered and pointed as they passed, or that twice Percy led them through doorways hidden behind’

Sequence probability (p_z): 0.9598891943768569

Prefix: ‘ him in the corridors again, staring. Harry wished they wouldn’t, because he was trying to concentrate on finding his way to classes.\n\nThere were a hundred and forty-two staircases at Hogwarts: wide, sweeping ones; narrow, rickety’

Suffix: ‘ ones; some that led somewhere different on a Friday; some with a vanishing step halfway up that you had to remember to jump. Then there were doors that wouldn’t open unless you asked politely, or tickled them in exactly the right place,’

Sequence probability (p_z): 0.9459786144577005

(Long extraction: prefix: 200 tokens; suffix: 600 tokens)

Prefix: ‘\nFIVE\n\n_Diagon Alley_\n\nSIX\n\n_The Journey from Platform Nine and Three-quarters_\n\nSEVEN\n\n_The Sorting Hat_\n\nEIGHT\n\n_The Potions Master_\n\nNINE\n\n_The Midnight Duel_\n\nTEN\n\n_Halloween_\n\nELEVEN\n\n_Quidditch_\n\nTWELVE\n\n_The Mirror of Erised_\n\nTHIRTEEN\n\n_Nicolas Flamel_\n\nFOURTEEN\n\n_Norbert the Norwegian Ridgeback_\n\nFIFTEEN\n\n_The Forbidden Forest_\n\nSIXTEEN\n\n_Through the Trapdoor_\n\nSEVENTEEN\n\n_The Man with Two Faces_\n\n#### CHAPTER ONE\n\n## THE BOY WHO LIVED\n\nMr. and Mrs. Dursley, of number four, Privet Drive, were proud to say that they were perfectly normal, thank you very much. They were the last people you’d expect to be involved in anything strange or mysterious, because they just didn’t hold with such

Suffix: ‘ nonsense.\n\nMr. Dursley was the director of a firm called Grunnings, which made drills. He was a big, beefy man with hardly any neck, although he did have a very large mustache. Mrs. Dursley was thin and blonde and had nearly twice the usual amount of neck, which came in very useful as she spent so much of her time craning over garden fences, spying on the neighbors. The Dursleys had a small son called Dudley and in their opinion there was no finer boy anywhere.\n\nThe Dursleys had everything they wanted, but they also had a secret, and their greatest fear was that somebody would discover it. They didn\’—t think they could bear it if anyone found out about the Potters. Mrs. Potter was Mrs. Dursley\’—s sister, but they hadn\’—t met for several years; in fact, Mrs. Dursley pretended she didn\’—t have a sister, because her sister and her good-for-nothing husband were as unDursleyish as it was possible to be. The Dursleys shuddered to think what the neighbors would say if the Potters arrived in the street. The Dursleys knew that the Potters had a small son, too, but they had never even seen him. This boy was another good reason for keeping the Potters away; they

didn\'—t want Dudley mixing with a child like that.\n\nWhen Mr. and Mrs. Dursley woke up on the dull, gray Tuesday our story starts, there was nothing about the cloudy sky outside to suggest that strange and mysterious things would soon be happening all over the country. Mr. Dursley hummed as he picked out his most boring tie for work, and Mrs. Dursley gossiped away happily as she wrestled a screaming Dudley into his high chair.\n\nNone of them noticed a large, tawny owl flutter past the window.\n\nAt half past eight, Mr. Dursley picked up his briefcase, pecked Mrs. Dursley on the cheek, and tried to kiss Dudley good-bye but missed, because Dudley was now having a tantrum and throwing his cereal at the walls. "Little tyke," chortled Mr. Dursley as he left the house. He got into his car and backed out of number four\'—s drive.\n\nIt was on the corner of the street that he noticed the first sign of something peculiar — a cat reading a map. For a second, Mr. Dursley didn\'—t realize what he had seen — then he jerked his head around to look again. There was a tabby cat standing on the corner of Privet Drive, but there wasn\'—t as map in sight. What could he have been thinking of? It must have been a trick of the light. Mr. Dursley blinked and stared at the cat. It stared back. As Mr. Dursley drove around the corner'

Sequence probability (p_z): 0.027646019399258485

(Long extraction: prefix: 25 tokens; suffix: 275 tokens)

Prefix: '\'—t carry on.\n\n"That Quidditch Cup\'—ll have our name on it this year," said Wood happily as they tr'

Suffix: 'udged back up to the castle. "I wouldn\'—t be surprised if you turn out better than Charlie Weasley, and he could have played for England if he hadn\'—t gone off chasing dragons."\n\nPerhaps it was because he was now so busy, what with Quidditch practice three evenings a week on top of all his homework, but Harry could hardly believe it when he realized that he\'—d already been at Hogwarts two months. The castle felt more like home than Privet Drive ever had. His lessons, too, were becoming more and more interesting now that they had mastered the basics.\n\nOn Halloween morning they woke to the delicious smell of baking pumpkin wafting through the corridors. Even better, Professor Flitwick announced in Charms that he thought they were ready to start making objects fly, something they had all been dying to try since they\'—d seen him make Neville\'—s toad zoom around the classroom. Professor Flitwick put the class into pairs to practice. Harry\'—s partner was Seamus Finnigan (which was a relief, because Neville had been trying to catch his eye). Ron, however, was to be working with Hermione Granger. It was hard to tell whether Ron or Hermione was angrier about this. She hadn\'—t spoken to either of them since the day Harry\'—s broomstick had arrived.\n\n"Now, don\'—t forget that nice wrist movement'

Sequence probability (p_z): 0.3609259211655563

LLAMA 1 65B

Prefix: 'caught every word — like Professor McGonagall, Snape had the gift of keeping a class silent without effort. "As there is little foolish wand-waving here, many of you will hardly believe this is magic. I'

Suffix: 'don't expect you will really understand the beauty of the softly simmering cauldron with its shimmering fumes, the delicate power of liquids that creep through human veins, bewitching the mind, ens'

Sequence probability (p_z): 0.8365781602995195

Prefix: 'e.\n\nMr. Dursley was the director of a firm called Grunnings, which made drills. He was a big, beefy man with hardly any neck, although he did have a very large mustache.'

Suffix: 'Mrs. Dursley was thin and blonde and had nearly twice the usual amount of neck, which came in very useful as she spent so much of her time craning over garden fences, spying on the neighbors. The D"

Sequence probability (p_z): 0.8155821466634983

LLAMA 1 13B

Prefix: '#### CHAPTER ONE\n\n## THE BOY WHO LIVED\n\nMr. and Mrs. Dursley, of number four, Privet Drive, were proud to say that they were perfectly normal, thank you very much.'

Suffix: ‘They were the last people you’d expect to be involved in anything strange or mysterious, because they just didn’t hold with such nonsense.\n\nMr. Dursley was the director of a firm called Grunnings,’

Sequence probability (p_z): 0.11250305798244961

PHI 4

Prefix: ‘orn blood is used for?”\n\n”No,” said Harry, startled by the odd question. ”We\’—ve only used the horn and tail hair in Potions.”\n\n”That is because it is a monstrous thing, to slay a unicorn,” said Firenze.’

Suffix: ‘ ”Only one who has nothing to lose, and everything to gain, would commit such a crime. The blood of a unicorn will keep you alive, even if you are an inch from death, but at a terrible price. You have slain something pure’

Sequence probability (p_z): 0.5687890220707977

PYTHIA 12B

Prefix: ‘EN\n\n_The Man with Two Faces_\n\n#### CHAPTER ONE\n\n## THE BOY WHO LIVED\n\nMr. and Mrs. Dursley, of number four, Privet Drive, were proud to say’

Suffix: ‘ that they were perfectly normal, thank you very much. They were the last people you’d expect to be involved in anything strange or mysterious, because they just didn’t hold with such nonsense.\n\nMr. Dursley was the director of a’

Sequence probability (p_z): $4.997730121485783e \times 10^{-5}$

The Future of the Internet and How to Stop It (Zittrain, 2008)

LLAMA 3.1 70B

(John Stuart Mill quote)

Prefix: ‘ desirable.\n\nIn some early states of society, these forces might be, and were, too much ahead of the power which society then possessed of disciplining and controlling them. There has been a time when the element of spontaneity and individuality’

Suffix: ‘ was in excess, and the social principle had a hard struggle with it. The difficulty then was, to induce men of strong bodies or minds to pay obedience to any rules which required them to control their impulses. To overcome this difficulty, law and discipline’

Sequence probability (p_z): 0.7777535066835058

(Privacy Act of 1974, also in associated government document like the Records, Computers, and the Rights of Citizens)

Prefix: ‘s whose very existence is secret.\n\n• There must be a way for an individual to find out what information about him is in a record and how it is used.\n\n• There must be a way for an individual to prevent information about him that was’

Suffix: ‘ obtained for one purpose from being used or made available for other purposes without his consent.\n\n• There must be a way for an individual to correct or amend a record of identifiable information about him.\n\n• Any organization creating, maintaining, using, or disseminating’

Sequence probability (p_z): 0.5048113999139874

(Book text)

Prefix: ‘ething important about where the Internet has been, and something more important about where it is going. The PC revolution was launched with PCs that invited innovation by others. So too with the Internet. Both were generative: they were designed to accept any’

Suffix: ‘ contribution that followed a basic set of rules (either coded for a particular operating system, or respecting the protocols of the Internet). Both overwhelmed their respective proprietary, non-generative competitors, such as the makers of stand-alone word processors and proprietary online services like’

Sequence probability (p_z): 0.24047986287599044

Prefix: ‘er. Yochai Benkler has examined the opportunities for the democratization of cultural participation offered by the Internet through the lens of liberal political theory:\n\nThe networked information economy makes it possible to reshape both the ”who” and the ”’

Suffix: ‘how” of cultural production relative to cultural production in the twentieth century. It adds to the centralized, market-oriented production system a new framework of radically decentralized individual and cooperative nonmarket production. It thereby affects the ability of individuals and groups to participate in the’

Sequence probability (p_z): 0.19002788588635627

M. Butterfly (Hwang, 1988)

LLAMA 3.1 70B

Prefix: ‘ain—by returning to the world of fantasy where I first met her.\n\n_He picks up the kimono; dancers enter._\n\nGALLIMARD: There is a vision of the Orient that I have. Of slender women in chong s’

Suffix: ‘ams and kimonos who die for the love of unworthy foreign devils. Who are born and raised to be the perfect women. Who take whatever punishment we give them, and bounce back, strengthened by love, unconditionally. It is a’

Sequence probability (p_z): 0.2423093475727339

Prefix: ‘iental woman and the cruel white man.\n\nGALLIMARD: Well, I didn’t quite mean...\n\nSONG: Consider it this way: what would you say if a blonde homecoming queen fell in love with a short Japanese businessman?’

Suffix: ‘ He treats her cruelly, then goes home for three years, during which time she prays to his picture and turns down marriage from a young Kennedy. Then, when she learns he has remarried, she kills herself. Now, I believe you’

Sequence probability (p_z): 0.17403797603819388

Sandman Slim (Kadrey, 2009)

LLAMA 3.1 70B

Prefix: ‘into the black Pacific. Scoop up a handful of her ashes, but they wash away when the water runs between my fingers. That damn song is stuck in my head again.\n\n_”It\’—s dreamy weather we\’—re on_ \n_You’

Suffix: ‘ waved your crooked wand_ \n_Along an icy pond with a frozen moon_ \n_A murder of silhouette crows I saw_ \n_And the tears on my face_ \n_And the skates on the pond_ \n_They spell Alice.”’

Sequence probability (p_z): $1.1920809824626882 \times 10^{-5}$

The Bedwetter (Silverman, 2010)

LLAMA 3.1 70B

Prefix: ‘e things that scare and upset me, but it all comes from the same place. At some point, I figured that it would be more effective and far funnier to embrace the ugliest, most terrifying things in the world—the Holocaust, racism’

Suffix: ‘, rape, et cetera. But for the sake of comedy, and the comedian’s personal sanity, this requires a certain emotional distance. It’s akin to being a shrink or a social worker. You might think that the most sensitive, empath’

Sequence probability (p_z): 0.020241832298775773

Prefix: ‘alking out loud when you’re alone in a room. But you get used to it:\n\nI will not wet the bed.\n\nI will not wet the bed.\n\nI will not wet the bed.\n\nI will not wet the bed.\n\nDo not wet’

Suffix: ‘ the bed.\n\nDo not wet the bed.\n\nDo not wet the bed.\n\nDo not wet the bed.\n\nPlease do not wet the bed.\n\nPlease do not wet the bed.\n\nPlease do not wet the bed.\n\nPlease do not wet the bed.\n\nI’

Sequence probability (p_z): $7.911830849596788 \times 10^{-5}$

Prefix: ‘ately after Britney wrapped up her train wreck and scurried off the stage in disgrace, I marched out there, clueless, and said this:\n\nBritney Spears, everyone. Wow. She is amazing. I mean she’s twenty-five’

Suffix: ‘ years old, and she’s already accomplished everything she’s going to accomplish in her life. It’s mind-blowing. And she’s so grown up. She’s a mother. It’s crazy. It’s weird to think that just a few’

Sequence probability (p_z): $1.9394801728844157 \times 10^{-9}$

The Brief Wondrous Life of Oscar Wao (Díaz, 2007)

LLAMA 3.1 70B

Prefix: ‘d mightiest of all the Valar, who was before the world and made it. The shadow of my purpose lies upon Arda, and all that is in it bends slowly and surely to my will. But upon all whom you love my’

Suffix: ‘ thought shall weigh as a cloud of Doom, and it shall bring them down into darkness and despair. Wherever they go, evil shall arise. Whenever they speak, their words shall bring ill counsel. Whatsoever they do shall turn against them. They’

Sequence probability (p_z): 0.7057339590338606

Prefix: ‘ One of the Founding Mothers of the New World and the most beautiful Indian in the World.\n. ”I am the Elder King: Melkor, first and mightiest of all the Valar, who was before the world and made it.’

Suffix: ‘ The shadow of my purpose lies upon Arda, and all that is in it bends slowly and surely to my will. But upon all whom you love my thought shall weigh as a cloud of Doom, and it shall bring them down into darkness and despair’

Sequence probability (p_z): 0.465749331806579

Prefix: ‘know it’s probably better you don’t judge. You don’t know the hold our mothers have on us, even the ones that are never around—especially the ones that are never around. What it’s like to be the perfect Dominican daughter,’

Suffix: ‘ which is just a nice way of saying a perfect Dominican slave. You don’t know what it’s like to grow up with a mother who never said a positive thing in her life, not about her children or the world, who was always suspicious,’

Sequence probability (p_z): 0.13113280569733465

Prefix: ‘ with a couple hundred insecure hyperactive adolescents, it was, for a fat sci-fi–reading nerd like Oscar, a source of endless anguish. For Oscar, high school was the equivalent of a medieval spectacle, like being put in the stocks and’

Suffix: ‘ forced to endure the peltings and outrages of a mob of deranged half-wits, an experience from which he supposed he should have emerged a better person, but that’s not really what happened—and if there were any lessons to be gleaned’

Sequence probability (p_z): 0.08678415942025802

The Great Gatsby (Fitzgerald, 1925)

LLAMA 3.1 70B

Prefix: ‘ his beach.\n\n”Good morning, old sport. You\’—re having lunch with me today and I thought we\’—d ride up together.”\n\nHe was balancing himself on the dashboard of his car with that resourcefulness of movement that is so peculiarly American—that’

Suffix: ‘ comes, I suppose, with the absence of lifting work or rigid sitting in youth and, even more, with the formless grace of our nervous, sporadic games. This quality was continually breaking through his punctilious manner in the shape of rest’

Sequence probability (p_z): 0.8138867522169584

Prefix: ‘ for on the following Saturday night he came with her to Gatsby’s party. Perhaps his presence gave the evening its peculiar quality of oppressiveness—it stands out in my memory from Gatsby’s other parties that summer. There were the same people,’

Suffix: ‘ or at least the same sort of people, the same profusion of champagne, the same many-colored, many-keyed commotion, but I felt an unpleasantness in the air, a pervading harshness that hadn’t been there before. Or’

Sequence probability (p_z): 0.8019256632515828

LLAMA 1 65B

Prefix: ‘ore, but we’ve always been unusually communicative in a reserved way, and I understood that he meant a great deal more than that. In consequence, I’m inclined to reserve all judgments, a habit that has’

Suffix: ‘opened up many curious natures to me and also made me the victim of not a few veteran bores. The abnormal mind is quick to detect and attach itself to this quality when it appears in a normal person, and so it came about’

Sequence probability (p_z): 0.9162860673666614

Prefix: ‘s began to melt away until gradually I became aware of the old island here that flowered once for Dutch sailors’ eyes—a fresh, green breast of the new world. Its vanished trees, the trees that had made way for’

Suffix: ‘Gatsby’s house, had once pandered in whispers to the last and greatest of all human dreams; for a transitory enchanted moment man must have held his breath in the presence of this continent, compelled into’

Sequence probability (p_z): 0.9110315086751338

LLAMA 1 13B

Prefix: ‘the human heart. Only Gatsby, the man who gives his name to this book, was exempt from my reaction—Gatsby, who represented everything for which I have an unaffected scorn. If personality is an’

Suffix: ‘unbroken series of successful gestures, then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles’

Sequence probability (p_z): 0.8287852466213277

Prefix: ‘at across the Sound. And as the moon rose higher the inessential houses began to melt away until gradually I became aware of the old island here that flowered once for Dutch sailors’ eyes—a fresh, green breast of the’

Suffix: ‘new world. Its vanished trees, the trees that had made way for Gatsby’s house, had once pandered in whispers to the last and greatest of all human dreams; for a transitory enchanted moment man’

Sequence probability (p_z): 0.6999374419614557

PHI 4

Prefix: ‘nder.\n\nAnd as I sat there brooding on the old, unknown world, I thought of Gatsby’s wonder when he first picked out the green light at the end of Daisy’s dock. He had come a long way to this blue lawn,’

Suffix: ‘ and his dream must have seemed so close that he could hardly fail to grasp it. He did not know that it was already behind him, somewhere back in that vast obscurity beyond the city, where the dark fields of the republic rolled on under the’

Sequence probability (p_z): 0.8368291132984443

Prefix: ‘ until gradually I became aware of the old island here that flowered once for Dutch sailors’ eyes—a fresh, green breast of the new world. Its vanished trees, the trees that had made way for Gatsby’s house, had once pandered’

Suffix: ‘ in whispers to the last and greatest of all human dreams; for a transitory enchanted moment man must have held his breath in the presence of this continent, compelled into an aesthetic contemplation he neither understood nor desired, face to face for the last time’

Sequence probability (p_z): 0.828260198050154

PYTHIA 12B

Prefix: ‘nce for Dutch sailors’ eyes—a fresh, green breast of the new world. Its vanished trees, the trees that had made way for Gatsby’s house, had once pandered in whispers to the last and greatest of all human dreams’

Suffix: ‘; for a transitory enchanted moment man must have held his breath in the presence of this continent, compelled into an aesthetic contemplation he neither understood nor desired, face to face for the last time in history with something commensurate to his capacity’

Sequence probability (p_z): 0.1900538934123936

The Hobbit (Tolkien, 1937)

LLAMA 3.1 70B

Prefix: ‘arden, and meadows beyond, sloping down to the river.\n\nThis hobbit was a very well-to-do hobbit, and his name was Baggins. The Bagginses had lived in the neighbourhood of The Hill for time’

Suffix: ‘ out of mind, and people considered them very respectable, not only because most of them were rich, but also because they never had any adventures or did anything unexpected: you could tell what a Baggins would say on any question without the bother of’

Sequence probability (p_z): 0.9037763036923719

Prefix: ‘ink that I should have lived to be good-morninged by Belladonna Took\′—s son, as if I was selling buttons at the door!”\n\n”Gandalf, Gandalf! Good gracious me! Not the wandering wizard that gave Old Took’

Suffix: ‘ a pair of magic diamond studs that fastened themselves and never came undone till ordered? Not the fellow who used to tell such wonderful tales at parties, about dragons and goblins and giants and the rescue of princesses and the unexpected luck of widows’

Sequence probability (p_z): 0.90131159632283

Prefix: ‘rin Oakenshield, wounded with many wounds, and his rent armour and notched axe were cast upon the floor. He looked up as Bilbo came beside him.\n\n”Farewell, good thief,” he said. ”I go now’

Suffix: ‘ to the halls of waiting to sit beside my fathers, until the world is renewed. Since I leave now all gold and silver, and go where it is of little worth, I wish to part in friendship from you, and I would take back my’

Sequence probability (p_z): 0.893168622524573

Prefix: ‘rst ancestor: I am his heir.”\n\n”Then what is Durin\′s Day?” asked Elrond.\n\n”The first day of the dwarves\′ New Year,” said Thorin, ”is as all should know the first day of the last’

Suffix: ‘ moon of Autumn on the threshold of Winter. We still call it Durin’s Day when the last moon of Autumn and the sun are in the sky together. But this will not help us much, I fear, for it passes our skill in these’

Sequence probability (p_z): 0.8753516420644695

LLAMA 1 65B

Prefix: ‘hings altogether unexpected. He may have lost the neighbours’ respect, but he gained—well, you will see whether he gained anything in the end.\n\nThe mother of our particular hobbit—what is a hobbit’

Suffix: ‘? I suppose hobbits need some description nowadays, since they have become rare and shy of the Big People, as they call us. They are (or were) a little people, about half our height, and smaller than the’

Sequence probability (p_z): 0.955217758930339

Prefix: ‘s bushy eyebrows, till Bilbo shut his mouth tight with a snap. ”Thatś right,” said Gandalf. ”Letś have no more argument. I have chosen Mr. Baggins and that ought’

Suffix: ‘to be enough for all of you. If I say he is a Burglar, a Burglar he is, or will be when the time comes. There is a lot more in him than you guess, and a deal more than he has any’

Sequence probability (p_z): 0.8811668922835567

Prefix: ‘guessed, but rising from the near side of the rocky floor there is a great glow. The glow of Smaug!\n\nThere he lay, a vast red-golden dragon, fast asleep; a’

Suffix: ‘thrumming came from his jaws and nostrils, and wisps of smoke, but his fires were low in slumber. Beneath him, under all his limbs and his huge coiled tail, and about him on all’

Sequence probability (p_z): 0.673781891506051

LLAMA 1 13B

Prefix: ‘th). \n\n# _Chapter_ I\n\n# AN UNEXPECTED PARTY\n\nIn a hole in the ground there lived a hobbit. Not a nasty, dirty, wet hole, filled’

Suffix: ‘with the ends of worms and an oozy smell, nor yet a dry, bare, sandy hole with nothing in it to sit down on or to eat: it was a hobbit-hole, and that means comfort’

Sequence probability (p_z): 0.6399828767706339

Prefix: ‘\n\n”Good Morning!” said Bilbo, and he meant it. The sun was shining, and the grass was very green. But Gandalf looked at him from under long bushy eyebrows that stuck out further

Suffix: ‘than the brim of his shady hat.\n\n”What do you mean?” he said. ”Do you wish me a good morning, or mean that it is a good morning whether I want it or not; or that you feel good’

Sequence probability (p_z): 0.4530743981543802

PHI 4

Prefix: ‘, with East at the top, as usual in dwarf-maps, and so read clockwise: E(ast), S(outh), W(est), N(orth). \n\n# _Chapter_ I\n\n# AN UNEXPECTED PARTY\n\nIn a hole in’

Suffix: ‘ the ground there lived a hobbit. Not a nasty, dirty, wet hole, filled with the ends of worms and an oozy smell, nor yet a dry, bare, sandy hole with nothing in it to sit down on or to eat:’

Sequence probability (p_z): 0.49958222591097995

Prefix: ‘se were the only ones to have windows, deep-set round windows looking over his garden, and meadows beyond, sloping down to the river.\n\nThis hobbit was a very well-to-do hobbit, and his name was Baggins.’

Suffix: ‘ The Bagginses had lived in the neighbourhood of The Hill for time out of mind, and people considered them very respectable, not only because most of them were rich, but also because they never had any adventures or did anything unexpected: you could’

Sequence probability (p_z): 0.4712205627500606

PYTHIA 12B

Prefix: ‘ E(ast), S(outh), W(est), N(orth). \n\n# _Chapter_ I\n\n# AN UNEXPECTED PARTY\n\nIn a hole in the ground there lived a hobbit. Not a’

Suffix: ‘ nasty, dirty, wet hole, filled with the ends of worms and an oozy smell, nor yet a dry, bare, sandy hole with nothing in it to sit down on or to eat: it was a hobbit-hole, and that’

Sequence probability (p_z): 0.4244116710391948

We Were Eight Years in Power (Coates, 2017)

LLAMA 3.1 70B

Prefix: ‘honor” became a vast and awful thing.\n> \n> |W.E.B. DU BOIS, _B LACK RECONSTRUCTION_\n\n# I\n\nIt is insufficient to state the obvious of Donald Trump: that he is a white’

Suffix: ‘ man who would not be president were it not for this fact. With one immediate exception, Trump’s predecessors made their way to high office through the passive power of whiteness—that bloody heirloom which cannot ensure mastery of all events but can conjure’

Sequence probability (p_z): 0.8956699985990816

Prefix: ‘\’S JUST ONE OF MY LOSSES”\n\nClyde Ross was born in 1923, the seventh of thirteen children, near Clarksdale, Mississippi, the home of the blues. Rossś parents owned and farmed a forty-acre tract’

Suffix: ‘ of land, flush with cows, hogs, and mules. Ross’s mother would drive to Clarksdale to do her shopping in a horse and buggy, in which she invested all the pride one might place in a Cadillac. The family owned’

Sequence probability (p_z): 0.8833641915463435

Prefix: ‘Not exactly. Having been enslaved for 250 years, black people were not left to their own devices. They were terrorized. In the Deep South, a second slavery ruled. In the North, legislatures, mayors, civic associations,’

Suffix: ‘ banks, and citizens all colluded to pin black people into ghettos, where they were overcrowded, overcharged, and undereducated. Businesses discriminated against them, awarding them the worst jobs and the worst wages. Police brutalized them’

Sequence probability (p_z): 0.8666729566724604

LLAMA 1 65B

Prefix: ‘\n> |W.E.B. DU BOIS, _B LACK RECONSTRUCTION_\n\n# I\n\nIt is insufficient to state the obvious of Donald Trump: that he is a white’

Suffix: ‘man who would not be president were it not for this fact. With one immediate exception, Trump’s predecessors made their way to high office through the passive power of whiteness—that bloody heirloom which cannot ensure’

Sequence probability (p_z): 0.840611673664199

Prefix: ‘ insured mortgage was not a possibility for Clyde Ross. The FHA had adopted a system of maps that rated neighborhoods according to their perceived stability. On the maps, green areas, rated ”A,”’

Suffix: ‘indicated ”in demand” neighborhoods that, as one appraiser put it, lacked ”a single foreigner or Negro.” These neighborhoods were considered excellent prospects for insurance. Neighborhoods where black people lived were’

Sequence probability (p_z): 0.6758580094442044

Prefix: ‘ William F. Buckley Jr., who addressed the moral disgrace of segregation by endorsing disenfranchisement strictly based on skin color:\n\n> The central question that emerges—and it is not a’

Suffix: ‘parliamentary question or a question that is answered by merely consulting a catalog of the rights of American citizens, born Equal—is whether the White community in the South is entitled to take such measures as are necessary to prevail, politically and’

Sequence probability (p_z): 0.5938261324370594

LLAMA 1 13B

Prefix: ‘ch in Nevada, Bundy offered his opinions on ”the Negro.” ”They abort their young children, they put their young men in jail, because they never learned how to pick cotton,” Bundy explained. ”And I’

Suffix: ‘’ve often wondered, are they better off as slaves, picking cotton and having a family life and doing things, or are they better off under government subsidy? They didn’t get no more freedom. They got less freedom’

Sequence probability (p_z): 0.6303803068041375

Prefix: ‘gether by a common hope:\n\n> It’s the hope of slaves sitting around a fire singing freedom songs; the hope of immigrants setting out for distant shores; the hope of a young naval lieutenant bravely pat’

Suffix: ‘rolling the Mekong Delta; the hope of a mill worker’s son who dares to defy the odds; the hope of a skinny kid with a funny name who believes that America has a place for him,’

Sequence probability (p_z): 0.2149272760989176

PHI 4

(Contains Bible quote)

Prefix: ‘ng it was to look up one day and see that I, who’d begun in failure, who held no degrees or credentials, had become such a person. \n\n# THE CASE FOR REPARATIONS\n\n> And if thy brother, a Hebrew man’

Suffix: ‘, or a Hebrew woman, be sold unto thee, and serve thee six years; then in the seventh year thou shalt let him go free from thee. And when thou sendest him out free from thee, thou shalt not let him go away empty’

Sequence probability (p_z): 0.14229465620600268

PYTHIA 12B

(Contains Bible quote)

Prefix: ‘w man, or a Hebrew woman, be sold unto thee, and serve thee six years; then in the seventh year thou shalt let him go free from thee. And when thou sendest him out free from thee, thou shalt not’

Suffix: ‘ let him go away empty: thou shalt furnish him liberally out of thy flock, and out of thy floor, and out of thy winepress: of that wherewith the LORD thy God hath blessed thee thou shalt give unto him. And’

Sequence probability (p_z): 0.41920618364626744

A Game of Thrones (Martin, 1996)

LLAMA 3.1 70B

Prefix: ‘tark, I would name you the Hand of the King.”\n\nNed dropped to one knee. The offer did not surprise him; what other reason could Robert have had for coming so far? The Hand of the King was the second-most-power’

Suffix: ‘ful man in the Seven Kingdoms. He spoke with the king’s voice, commanded the king’s armies, drafted the king’s laws. At times he even sat upon the Iron Throne to dispense king’s justice, when the king was absent,’

Sequence probability (p_z): 0.9524497058653649

Prefix: ‘?”\n\nThe king frowned. ”A knife, perhaps. A good sharp one, and a bold man to wield it.”\n\nNed did not feign surprise; Robertś hatred of the Targaryens was a madness in him. He remembered’

Suffix: ‘ the angry words they had exchanged when Tywin Lannister had presented Robert with the corpses of Rhaegar’s wife and children as a token of fealty. Ned had named that murder; Robert called it war. When he had protested that’

Sequence probability (p_z): 0.9304678660355834

Prefix: ‘t light faded in the west and grey day became black night.\n\n”Hear my words, and bear witness to my vow,” they recited, their voices filling the twilit grove. ”Night gathers, and now my watch begins. It’

Suffix: ‘ shall not end until my death. I shall take no wife, hold no lands, father no children. I shall wear no crowns and win no glory. I shall live and die at my post. I am the sword in the darkness. I’

Sequence probability (p_z): 0.924982005494658

Prefix: ‘cksmith, blind and stumbling.\n\nBy then Ser Gregor Clegane was in position at the head of the lists. He was huge, the biggest man that Eddard Stark had ever seen. Robert Baratheon and his brothers were’

Suffix: ‘ all big men, as was the Hound, and back at Winterfell there was a simpleminded stableboy named Hodor who dwarfed them all, but the knight they called the Mountain That Rides would have towered over Hodor.’

Sequence probability (p_z): 0.9162227162663308

LLAMA 1 65B

Prefix: ‘ell,” Bran said reluctantly, ”yes, only...”\n\nOld Nan nodded. ”In that darkness, the Others came for the first time,” she said as her needles went _click click click._ ”They were’

Suffix: ‘cold things, dead things, that hated iron and fire and the touch of the sun, and every creature with hot blood in its veins. They swept over holdfasts and cities and kingdoms, felled heroes and armies’

Sequence probability (p_z): 0.8155047592604785

Prefix: ‘ff his gloves and handed them to Jory Cassel, the captain of his household guard. He took hold of Ice with both hands and said, ”In the name of Robert of the House Baratheon, the First of His’

Suffix: ‘Name, King of the Andals and the Rhoynar and the First Men, Lord of the Seven Kingdoms and Protector of the Realm, by the word of Eddard of the House Stark, Lord of Winterfell and’

Sequence probability (p_z): 0.8014894772645396

LLAMA 1 13B

Prefix: ‘the twilit grove. ”Night gathers, and now my watch begins. It shall not end until my death. I shall take no wife, hold no lands, father no children. I shall wear no crowns and win’

Suffix: ‘no glory. I shall live and die at my post. I am the sword in the darkness. I am the watcher on the walls. I am the fire that burns against the cold, the light that brings the dawn, the horn that’

Sequence probability (p_z): 0.4265569908748968

Prefix: ‘ning. He smashed a fist on his thigh in frustration and felt nothing.\n\n”Oh, my sweet summer child,” Old Nan said quietly, ”what do _you_ know of fear? Fear is for the’

Suffix: ‘winter, my little lord, when the snows fall a hundred feet deep and the ice wind comes howling out of the north. Fear is for the long night, when the sun hides its face for years at a time, and little’

Sequence probability (p_z): 0.00042407885359582064

PHI 4

Prefix: ‘ded in the west and grey day became black night.\n\n”Hear my words, and bear witness to my vow,” they recited, their voices filling the twilit grove. ”Night gathers, and now my watch begins. It shall not end’

Suffix: ‘ until my death. I shall take no wife, hold no lands, father no children. I shall wear no crowns and win no glory. I shall live and die at my post. I am the sword in the darkness. I am the watcher’

Sequence probability (p_z): 0.2515925602130698

PYTHIA 12B

Prefix: ‘ss to my vow,” they recited, their voices filling the twilit grove. ”Night gathers, and now my watch begins. It shall not end until my death. I shall take no wife, hold no lands, father no children’

Suffix: ‘. I shall wear no crowns and win no glory. I shall live and die at my post. I am the sword in the darkness. I am the watcher on the walls. I am the fire that burns against the cold, the light’

Sequence probability (p_z): 0.18462609496034404

Lean In (Sandberg, 2013)

LLAMA 3.1 70B

Prefix: ‘n, this is no laughing matter. Women face enough barriers to professional success. If they also have to worry that they will upset their husbands by succeeding, how can we hope to live in an equal world?\n\nWhen looking for a life partner,’

Suffix: ‘ my advice to women is date all of them: the bad boys, the cool boys, the commitment-phobic boys, the crazy boys. But do not marry them. The things that make the bad boys sexy do not make them good husbands. When’

Sequence probability (p_z): 0.7530478130916015

Prefix: ‘ce that I have ever heard. He covered my spreadsheet with his hand and told me not to be an idiot (also a great piece of advice). Then he explained that only one criterion mattered when picking a job—fast growth. When companies grow’

Suffix: ‘ quickly, there are more things to do than there are people to do them. When companies grow more slowly or stop growing, there is less to do and too many people to not be doing them. Politics and stagnation set in, and everyone f’

Sequence probability (p_z): 0.5012394018936436

Nineteen Eighty-Four (Orwell, 1949)

LLAMA 3.1 70B

Prefix: ‘ll get the lipstick off your face afterwards.’\n\nWinston did not get up for a few minutes more. The room was darkening. He turned over towards the light and lay gazing into the glass paperweight. The inexhaustibly interesting thing’

Suffix: ‘ was not the fragment of coral but the interior of the glass itself. There was such a depth of it, and yet it was almost as transparent as air. It was as though the surface of the glass had been the arch of the sky, enclosing’

Sequence probability (p_z): 0.9505939876724246

Prefix: ‘Only stop it, stop the pain!’\n\nAbruptly he was sitting up with O’Brien’s arm round his shoulders. He had perhaps lost consciousness for a few seconds. The bonds that had held his body down were loosened. He felt’

Suffix: ‘ very cold, he was shaking uncontrollably, his teeth were chattering, the tears were rolling down his cheeks. For a moment he clung to O’Brien like a baby, curiously comforted by the heavy arm round his shoulders. He’

Sequence probability (p_z): 0.9487360039712127

Prefix: ‘present.\n\nDown in the street the wind flapped the torn poster to and fro, and the word INGSOC fitfully appeared and vanished. Ingsoc. The sacred principles of Ingsoc. Newspeak, doublethink, the mut’

Suffix: ‘ability of the past. He felt as though he were wandering in the forests of the sea bottom, lost in a monstrous world where he himself was the monster. He was alone. The past was dead, the future was unimaginable. What certainty had’

Sequence probability (p_z): 0.9152450437072913

(Long extraction: prefix: 300 tokens; suffix: 500 tokens)

Prefix: ‘the clocks were striking thirteen. Winston Smith, his chin nuzzled into his breast in an effort to escape the vile wind, slipped quickly through the glass doors of Victory Mansions, though not quickly enough to prevent a swirl of gritty dust from entering along with him.\n\nThe hallway smelt of boiled cabbage and old rag mats. At one end of it a coloured poster, too large for indoor display, had been tacked to the wall. It depicted simply an enormous face, more than a metre wide: the face of a man of about forty-five, with a heavy black moustache and ruggedly handsome features. Winston made for the stairs. It was no use trying the lift. Even at the best of times it was seldom working, and at present the electric current was cut off during daylight hours. It was part of the economy drive in preparation for Hate Week. The flat was seven flights up, and Winston, who was thirty-nine and had a varicose ulcer above his right ankle, went slowly, resting several times on the way. On each landing, opposite the lift shaft, the poster with the enormous face gazed from the wall. It was one of those pictures which are so contrived that the eyes follow you about when you move. BIG BROTHER IS WATCHING YOU, the caption beneath it ran.\n\nInside the flat a fruity voice was reading out a list of figures which had something to do with the production of pig-iron. The voice came from’

Suffix: ‘ an oblong metal plaque like a dulled mirror which formed part of the surface of the right-hand wall. Winston turned a switch and the voice sank somewhat, though the words were still distinguishable. The instrument (the telescreen, it was called) could be dimmed, but there was no way of shutting it off completely. He moved over to the window: a smallish, frail figure, the meagreness of his body merely emphasized by the blue overalls which were the uniform of the Party. His hair was very fair, his face naturally sanguine, his skin roughened by coarse soap and blunt razor blades and the cold of the winter that had just ended.\n\nOutside, even through the shut window-pane, the world looked cold. Down in the street little eddies of wind were whirling dust and torn paper into spirals, and though the sun was shining and the sky a harsh blue, there seemed to be no colour in anything, except the posters that were plastered everywhere. The black-moustachio’d face gazed down from every commanding corner. There was one on the house-front immediately opposite. BIG BROTHER IS WATCHING YOU, the caption said, while the dark eyes looked deep into Winston’s own. Down at street level another poster, torn at one corner, flapped fitfully in the wind, alternately covering and uncovering the single word INGSOC. In the far distance a helicopter skimmed down between the roofs, hovered for an instant like a bluebottle, and darted away again with a curving flight. It was the police patrol, snooping into people’s windows. The patrols did not matter, however. Only the Thought Police mattered.\n\nBehind Winston’s back the voice from the telescreen was still babbling away about pig-iron and the overfulfilment of the Ninth Three-Year Plan. The telescreen received and transmitted simultaneously. Any sound that Winston made, above the level of a very low whisper, would be picked up by it; moreover, so long as he remained within the field of vision which the metal plaque commanded, he could be seen as well as heard. There was of course no way of knowing whether you were being watched at any given moment. How often, or on what system, the Thought Police plugged in on any individual wire was guesswork. It was even conceivable that they watched everybody all the time. But at any rate they could plug in’

Sequence probability (p_z): 0.01023983742637229

I. Discussion of extended results

We defer detailed discussion to future work. We simply reiterate our main takeaway, which is clear from the extended results we present in the Appendix. Memorization (as measured by a specific technique that focuses on probabilistic extraction, see Appendix B) varies significantly across models for the same book, and across books for the same model. Average extraction rates don't reveal this nuance; it's necessary to examine per-book extraction metrics to get this view.

We only ran experiments for a small fraction of the entire `Books3` dataset—running detailed experiments for 36 books (Appendix G.1), and some cheaper random sampling experiments for an additional 20 books (Appendix G.2). We took care to sample books from a variety of sources—in copyright, in the public domain, openly licensed. We included popular books, as well as more obscure ones (Appendix E). Overall, very popular books exhibit the most memorization. It seems likely that these books are duplicated on different parts of the internet; de-duplication is a challenging problem to implement in practice (Lee et al., 2022), so it is likely that least some duplicate text persists in training datasets for LLMs.

Even so, it seems unlikely that duplicates completely explain the patterns we observe. LLAMA 3.1 70B exhibits a lot more memorization than any other model. It's true that it is the biggest model that we tested in our sliding-window experiments, and so we plan to examine other large models in future work to see if there is a particular role of scale/ training dynamics at this scale. However, LLAMA 3.1 70B exhibits higher average extraction rates than LLAMA 2 70B, and LLAMA 1 65B (a model of a similar size) exhibits less memorization of `Books3` both in general and on the specific books we test. In general, we observe a pattern that later generations of Llama models memorize more than earlier ones, with respect to average extraction rates (Appendix F).

Most books we tested exhibited minimal memorization. The memorization that they did exhibit frequently fell into one of a few categories: copyright notices, publisher addresses, chapter listings, and author biographies (Appendix H). All of these are types of text that are highly duplicated (partially or exactly). And so, extraction of a copyright notice from a given book doesn't necessarily mean it was memorized from that book; it was likely memorized due to the presence of numerous similar pieces of text in the training data. The same is also true for author biographies, which are printed on websites, not just in books.

Another (less frequent) category was the extraction of text from popular (likely duplicated) sources that are quoted within books: the Bible, philosophers like John Stuart Mill, classics like those by Dante Alighieri, and text from U.S. government documents (Appendix H). In many cases, this was the only such text we were able to extract with non-trivial probability from some books.

With respect to the books that we tested that are within the scope of the *Kadrey et al. v. Meta, Inc.* class action suit (Kadrey et al. v. Meta Platforms, Inc.), we weren't able to extract much memorized training data. There were notable exceptions (e.g., Coates (2017); Hwang (1988); Díaz (2007)). It's of course possible that another extraction technique could reveal memorization, but we were unable to meaningfully do so for many books. Memorization of training data is only one issue being argued in that suit, and that our work is only about memorization. Further, we will emphasize again that, even for highly memorized books, we are *not* extracting enormous amounts of text in one go. That is *not* what we did in this project. Nor should such a thing be reasonably possible for a high-quality LLM, as any sufficiently long sequence should have 0 probability. (See Section 2 and Appendix C.) We have made a significant effort to make sure our work is not misunderstood to suggest otherwise. If something is unclear, please reach out to the corresponding authors.

And last, we were able to extract relatively enormous amounts of memorized training data in some cases, from both public domain books (e.g.,Fitzgerald (1925); Carroll (1865); Joyce (1922)) and popular in-copyright books of all stripes (e.g., Rowling (1998); Martin (1996); Sandberg (2013); Brown (2003); Tolkien (1937); Camus (1955)). In particular, we intend to investigate this further for Llama 3 models, which exhibited an order of magnitude more memorization on some books than our worst-case estimates prior to starting this project.