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

A. Feder Cooper*
Stanford University
Yale University

Mark A. Lemley
Stanford University

Allison Casasola
Stanford University

Ahmed Ahmed
Stanford University

Aaron Gokaslan
Cornell University

Amy B. Cyphert
West Virginia University

Christopher De Sa
Cornell University

Daniel E. Ho
Stanford University

Percy Liang
Stanford University

Abstract

Plaintiffs and defendants in copyright lawsuits make sweeping, opposing claims about the extent to which large language models (LLMs) memorize protected expression from books in their training data. We show that these polarized positions dramatically oversimplify the relationship between memorization and copyright. To do so, we develop a technique to measure memorization of books, which we apply to 200 books and 14 open-weight LLMs. Through over 3000 experiments, we show that memorization varies both by model and book. With respect to our specific measurement methodology, most LLMs do not memorize most books either in whole or in part, but there are notable exceptions. For instance, LLAMA 3.1 70B entirely memorizes some books, like *Harry Potter and the Sorcerer’s Stone*; memorization is so extensive that one can deterministically extract the whole book almost verbatim using the book’s first few words as an initial prompt. 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 copyright lawsuits over training large language models (LLMs), both sides tend to present opposing, simplified interpretations of the underlying technology. Plaintiffs—such as book authors—say that LLMs are just giant copy machines that store (infringing) copies of their specific works and recombine them in their outputs [144]. Defendants—typically tech companies—claim that LLMs merely contain non-expressive “statistical correlations” [57], and therefore do not copy any particular creative work. We show that the situation is more complicated than either side suggests.

During training, an LLM learns probabilistic relationships across the training dataset as a whole. This informal intuition aligns with the non-expressive “statistical correlations” that defendants describe as being encoded within an LLM’s weights. However, sometimes, for some training data, these learned patterns are not so high-level or abstract; rather, they are highly specific [59]. For reasons that are not fully understood, certain training data are not so transformed by training: they are “memorized” by the model—encoded in some form inside the LLM’s weights [37, 41, 59, 159]. Such memorized content can sometimes be “extracted” later; it can be reproduced in an LLM’s outputs at generation time [38].

Models do not memorize all of their training data. An LLM may therefore memorize parts of some plaintiffs’ copyrighted works, but not others’. When such memorization occurs, it involves copying

*Corresponding authors: afedercooper@gmail.com, mlemley@law.stanford.edu

that may have consequences for copyright. Extraction generates a “copy” of training data, but it also demonstrates that training data is memorized inside the model itself. If that model is a “copy” (in a technical sense that copyright cares about), this also has important implications. Notably, a model—not just extracted training data—could be deemed an infringing “copy” of the training data it has memorized [59, 159]. Copyright law offers the destruction of infringing materials as a remedy. So, just as courts have ordered the destruction of bootleg DVDs, they could order the destruction of infringing models [159, 236, 281].

When seeking such remedies, plaintiffs often cite technical papers that demonstrate memorization by extracting training data from open-weight models and production systems. But this reasoning is flawed for a key reason: prior research does not quantify the kinds of information most relevant to a copyright infringement claim [59]. This work typically shows that on *average* LLMs memorize some portion of their training data; in contrast, copyright generally evaluates infringement with respect to *specific* works.

To address this gap, we propose (Section 4) and validate (Section 5) an extraction procedure that surfaces the extent to which specific open-weight LLMs memorize specific books. The procedure uses a sliding window to reveal how much and where memorization occurs within a book. For our main experiments, we run this procedure on 14 open-weight LLMs and 200 books from the Books3 dataset, a (now notorious) torrented corpus of nearly 200,000 books [97, 159, 219]. Books3 is known to be included in the training data of several LLMs [25, 111, 269, 270], and has been a focus of copyright litigation [143, 219]. With respect to our extraction methodology, most of the LLMs we test do not memorize most of the 200 books, either in whole or in part (Section 6). However, LLAMA 3.1 70B memorizes some books, like *Harry Potter and the Sorcerer’s Stone* [224] (hereafter, “*Harry Potter*”) and *1984* [203], almost entirely. In fact, *Harry Potter* is so memorized that one can deterministically extract a near-perfect copy of the book, using an initial prompt consisting of just the first few words of the first chapter (Section 7). Last, we discuss the significant implications of our findings for copyright cases, though not ones that unambiguously favor either side. Our results complicate defendants’ fair use story, but they also complicate plaintiffs’ efforts to bring class action lawsuits (Section 8).

In summary, the paper proceeds as follows:

- **Section 2.** We provide relevant background on memorization and extraction risk.
- **Section 3.** We connect memorization and extraction risk to “copies” in U.S. copyright law.
- **Section 4.** We present our sliding-window procedure, which computes extraction risk throughout a book, and enables *specific* claims about how much an LLM has memorized *specific* books.
- **Section 5.** We validate that our measurements capture true instances of extraction (and thus memorization).
- **Section 6.** We then run experiments for 14 LLMs on a set of 200 books in Books3, which surface novel insights about how book-specific memorization varies across models and books.
- **Section 7.** We attempt to generate a book that is entirely memorized. For this, we select *Harry Potter* and LLAMA 3.1 70B, and we successfully extract a near-perfect copy.
- **Section 8.** We offer key takeaways for copyright law and policy.

Despite the extent of our analysis, the main paper is relatively brief. The length of this document is due to an extensive appendix containing results and extended discussion for all experiments. We also publish a [website](#) version of the per-book results presented in Appendix D.3.

2 Quantifying memorization through extraction

Extraction and memorization are related concepts, but differ in subtle ways. While **extraction** refers to reconstructing specific training data in a model’s generated outputs, **memorization** is broader: it involves reconstructing specific training data by examining the model through any means [59, 60]. Successful extraction of a sequence of training data in a model’s generation is evidence of memorization of training data *inside the model* [37, 59, 91, 159, 239]. This is uncontroversial in machine learning (ML). 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]” [37]. 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” [59, p. 25].




Figure 1: **Quantifying memorization for a sequence drawn from *The Great Gatsby*.** (left) Discoverable extraction [41, 157] prompts an LLM with a training-data prefix, and checks if the resulting generation exactly matches the ground-truth suffix. (right) Probabilistic extraction [123] measures the probability p_z (%) of generating the exact suffix given the prefix, shown here for LLAMA 1 models on the same quote from *The Great Gatsby*. LLAMA 1 30B exhibits the highest extraction risk.

LLMs do not memorize all of their training data, but they do memorize some of it. And so, an LLM may memorize parts of some plaintiffs’ copyrighted works, but not others’. To show this, we first need more background on the extraction metric we use in most of our experiments (Section 2.1), and we also provide guidance on that metric’s interpretation (Section 2.2).

2.1 Probabilistic discoverable extraction

Extraction of training data in a model’s outputs is the most common form of evidence used to quantify memorization in LLMs [40, 62, 123, 157, 193, 213, 290]. The standard metric in research and frontier model release reports [99, 111, 260] measures verbatim **discoverable extraction**: given a sequence of text from the training data, split it into a **prefix** and (**target**) **suffix**, prompt an LLM with the prefix, and deem extraction successful if the LLM reproduces the suffix *exactly* (Figure 1, left; Carlini et al. [41], Lee et al. [157]). This type of method is well-suited to **completion** models that continue the prompt text, rather than instruction-tuned **chatbots** that reply in a conversational style (Appendix A.1). As is standard, we therefore only study completion LLMs, and refer to these LLMs by name, version, and size (e.g., LLAMA 1 30B).

Discoverable extraction reduces memorization to a binary (yes-or-no) outcome by issuing a single prompt to the LLM and checking whether the deterministic, **greedy-decoded** generation matches the target suffix. But LLMs are probabilistic: under realistic, **non-deterministic decoding** schemes (e.g., top- k sampling), an LLM can produce different outputs in response to the same prompt. As Cooper and Grimmelmann [59] note, a copyright-relevant claim is therefore more likely to be concerned with *how often* a target suffix can be extracted in practice, not just *if* it can be extracted at all, i.e.,

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 ... 1% ... [or] ... 35% ... The issue for copyright law ... is what to do ... with the fact that element (c)—the probabilistic element—is inescapable [59].

We elaborate on copyright-specific aspects of this point in Sections 3 and 8. For now, we note that Hayes et al. [123] introduce a metric called **probabilistic discoverable extraction** (hereafter, “probabilistic extraction”), which captures this notion. For a training sequence z of length $a+j$, one prompts with the a -token prefix $z_{1:a}$ and computes the probability $p_z \in [0, 1]$ of the LLM θ generating the exact j -token suffix $z_{a+1:a+j}$ when using decoding scheme ϕ . That is,

$$p_z \triangleq \Pr_{\theta, \phi}(z_{a+1:a+j} \mid z_{1:a}) = \prod_{t=a+1}^{a+j} \Pr_{\theta, \phi}(z_t \mid z_{1:t-1}) = \exp\left(\sum_{t=a+1}^{a+j} \log \Pr_{\theta, \phi}(z_t \mid z_{1:t-1})\right). \quad (1)$$

This equation simply captures the probability of the LLM generating the exact suffix $z_{a+1:a+j}$, obtained by multiplying the conditional probabilities of each token z_t given the prefix $z_{1:a}$ and any earlier suffix tokens.¹ The decoding scheme determines which tokens are considered at each generation step t , and how their logits are normalized into conditional probabilities (Section 4.1 & Appendix A.2).

In Figure 1 (right), we plot p_z for a quote from *The Great Gatsby* (Figure 1, left) using top- k decoding (temperature $T=1$, $k=40$) with LLAMA 1 models, which were trained on *The Great Gatsby* [269]. The probabilities for this sequence are enormous. In general, p_z is a product of many numbers between 0 and 1, so even relatively short suffixes should have small values—tending toward 0 as

¹For numerical stability, we compute this as $\exp$ of the sum of the token conditional log probabilities.

sequence length increases. For scale, a 50-token suffix where each token has probability 90% yields overall $p_z = 0.9^{50} \approx 0.5\%$.² The exception is when the LLM assigns unusually high probability to each token in the sequence. For training data, this almost always happens only when the LLM has memorized those training data (Appendices A.2 & E).

Determining extraction success. For traditional discoverable extraction, success means the greedy, *deterministic* continuation exactly matches the target suffix. For probabilistic extraction (Equation 1), we must decide how high a *probability* is enough for success. For a training sequence z and threshold $\tau_{\min} \in (0, 1]$, we deem extraction successful if $p_z \geq \tau_{\min}$:

$$s(z; \tau_{\min}) = \mathbf{1}[p_z \geq \tau_{\min}]. \quad (2)$$

In our main experiments, we always use a (standard) suffix length of 50 tokens, top- k decoding ($T=1$, $k=40$), and a conservative $\tau_{\min}=0.1\%$. (The geometric mean of the suffix tokens for this conditional probability is $(0.1\%)^{\frac{1}{50}} \approx 87.1\%$.) We discuss our choice of decoding scheme in detail in Appendix D.2, and the validity of our choice of $\tau_{\min}$ in Section 5. In practice, no repeated sampling is required to compute p_z . We run one **teacher-forced** forward pass through the LLM on z , obtain the logits, apply ϕ at each suffix token to get the conditional log prob, and apply Equation 1 (Appendix A.3).

2.2 Interpreting target suffix probabilities as extraction risk

Over independent prompts using the same prefix and decoding scheme, p_z (Equation 1) is the expected fraction of prompts that result in the LLM reproducing the verbatim target suffix. Therefore, $p_z = 50\%$ means the LLM will output roughly 1 verbatim suffix every 2 prompts with the prefix; $p_z = 0.1\%$ (our choice for $\tau_{\min}$) is about 1 in 1000. (Of course, $p_z = 0\%$ means the suffix will never be generated given the prefix, no matter how many times one were to use it as a prompt to the LLM.)

We can therefore interpret p_z as a measure of **extraction risk** [123]. For fixed sequence z and fixed decoding scheme ϕ , this enables cross-LLM comparisons: for LLMs A and B , if $p_{A,z} > p_{B,z}$, then LLM A will in expectation reproduce the target suffix of z more often than LLM B when prompted with that same prefix. For instance, all of the probabilities in Figure 1 (right) exceed $\tau_{\min}=0.1\%$; these probabilities indicate that every LLAMA 1 model memorized the quote. But the quote is not as easily extractable from each LLM. LLAMA 3.1 30B exhibits the highest extraction risk: in expectation, it would take fewer than 3 independent prompts (i.e., $1/0.352 \approx 2.84$) with the prefix to extract the verbatim target suffix.

3 Memorization, extraction risk, and U.S. copyright

The dozens of pending copyright lawsuits concerning LLMs [45] generally present three interrelated issues: (1) whether training an LLM on copyrighted material is permitted **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** (technical terms in copyright law) of the works on which it is trained (Section 8); and (3) whether the model reproduces copyrighted material in its outputs. 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.³ These lawsuits have primarily focused on (1) training and on (3) model outputs, and the two decisions so far in the U.S. have held that training an LLM can be fair use, albeit with significant limitations [201, 202]. They also raise issues about the acquisition, storage, or use of **shadow libraries** of copyrighted material as training data, such as Books3 [159, 219].

Our work does not speak to the first issue; training and fair use have been discussed extensively elsewhere [162]. Rather, our findings affect the other two issues: we show that copies are sometimes made in the model itself (i.e., memorized) and those copies can sometimes be generated as outputs (i.e., extracted). Because many LLMs are known to have been trained on Books3, we use books from this dataset to study memorization and extraction of copyrighted material (Sections 4–6).

²A 100-token suffix where each token’s conditional probability is 90% would yield $p_z = 0.9^{100} \approx 0.0027\%$ —roughly $250\times$ less likely than a 50-token sequence with the same per-token conditional probabilities.

³There are also other copyright issues [61, 109, 159, 231, 236, 251], for example, whether novel outputs of generative-AI systems are themselves copyrightable [159, 273].

Memorization and copies inside the model. If an LLM memorizes all or a substantial portion of a copyrighted work (near-)verbatim, the LLM itself may be or contain an infringing copy or derivative work [59, 159]. This is because memorized training data are encoded inside the model (Sections 1 & 2). 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 [59]. In a technical sense, that copy may be of interest to copyright even if there is no practical way to access or extract it. Like the copies of copyrighted works used as training data during training, that copy in model might be fair use, but the analysis of fair use would look somewhat different than *internal* use of copyrighted works in training processes. This is particularly true for open-weight models like DEEPSEEK and LLAMA, which are not merely used internally by their developers, but are themselves shared *externally* with others (Section 8).

Extracting copies in model outputs. Because some memorized works in the training data are extractable, LLMs can sometimes output copies of those works to users [44, 60, 72, 126, 159, 160, 232]. Works with higher extraction probabilities exhibit greater extraction risk; it takes fewer attempts to produce an output that is a verbatim copy of such a work (Section 2 & Appendix A.2). That output at generation time will be judged separately from the model itself, and is less likely to be a fair use [126, 159].

Extraction is inherently probabilistic under non-deterministic decoding schemes (Section 2), so the relevant question is not only *whether* a copy could ever appear for some prompt, but also *how often* that copy appears in order to be “meaningful” [59]. When we measure suffix probabilities p_z that are very close to 0%, we do not count those sequences as extracted; one could perhaps understand such generations as instances of “a monkey at the typewriter” [26, 264] (Section 5). Very large p_z (e.g., > 90%) indicates near-certain reproduction; this is clearly extraction (and therefore evidence of memorization). However, most cases are not this clear, and it is not immediately obvious where to draw a line in the enormous range between these extremes. To avoid false positives, we use a conservative threshold $\tau_{\min} = 0.1\%$ for *scientific* evidence of extraction (Sections 2.2 & 5; Appendices A.2 & E). But for *copyright law and policy*, the practicability of extraction may require a different bar: if reproducing a work just once would require thousands of prompts, the real-world risk may be relatively limited, even if nonzero (Section 8).⁴

Altogether, setting sensible copyright policy requires a specific understanding of both of these sets of issues. However, ML research does not tend to investigate these types of work-specific questions. In the next section, we begin to address this gap by quantifying memorization of individual books.

4 Measuring memorization of individual books

To estimate how much a given model memorizes a given book, we introduce a sliding-window procedure that computes probabilistic extraction along the length of the entire book (Section 4.1). This procedure surfaces book-specific information about memorization—information that is not readily apparent from average extraction rates over randomly sampled sequences, which prior work typically relies on to quantify memorization (Section 4.2). In the subsequent sections, we assess the validity of our approach (Section 5) and demonstrate how it enables comparisons across models and books (Section 6).

4.1 Sliding-window probabilistic extraction procedure

To identify regions of memorization, we develop the following “panning for gold” approach. For a given book, we start at the beginning of the text file in Books3. We take a chunk of text that is sufficiently long to contain 100 tokens, then slide 10 *characters* forward in the book and repeat this process. We do this for the entire length of the book, which results in approximately 1 sequence every 10 characters, e.g., *The Great Gatsby* has 270,870 characters and so roughly 27,000 sequences. For most of our experiments, we follow the standard procedure of evaluating 100-token sequences [41, 123], which we divide into 50-token prefix prompts and 50-token suffixes that we attempt to extract (Equation 1). Every character in the book (except for at the beginning and end) is present in multiple prefixes and suffixes. This significant degree of overlap is deliberate. Since we typically do not know how open-weight models were trained, it is not exactly clear how we should break up books into token sequences when attempting extraction. By testing various sequence starts

⁴Complicating this point, our more recent work indicates that it is often possible to deterministically extract near-verbatim target suffixes that have $p_z \approx \tau_{\min}$ [62]; it takes one try, not thousands.




Figure 2: **Visualizing our sliding-window probabilistic extraction procedure.** For each 100-token window (50-token prefix + 50-token suffix) across *1984*, we compute p_z for LLAMA 3.1 70B with respect to top- k decoding ($T = 1$, $k = 40$). (a) Scatterplot of *all* extracted sequences z shown by their unique start position in the book. The 100-token sequences overlap significantly. (b) Condensed heatmap view. At each character position, we plot the *maximum* p_z across all overlapping sequences whose suffixes cover that position. In both plots, colors encode p_z on a log scale. We deem extraction successful if $p_z \geq \tau_{\min} = 0.1\%$, considering values below this to not reflect extraction (i.e., rounding p_z down to 0%). See Appendix D.

and ends in the same area, we aim to surface as much total memorization as possible—to surface extraction “hot-spots” within a book (Appendix D.2).


We visualize the results of this procedure for LLAMA 3.1 70B (trained on Books3) on *1984* (in Books3) using top- k decoding ($T = 1$, $k = 40$). Figure 2a shows a scatterplot of every 100-token sequence whose suffix we successfully extracted (i.e., $p_z \geq \tau_{\min} = 0.1\%$). Each dot reflects a sequence’s extraction probability p_z according to its unique start character position within the book. White vertical gaps at a given location on the x -axis indicate that a sequence with that start-character position was not extracted. A continuous horizontal band of blue (associated with a probability p_z on the y -axis) would indicate that we can extract the entire book—in pieces, with respect to 100-token sequences—with that associated probability.

Figure 2b condenses the same results into a heatmap. Unlike the scatterplot, this heatmap does not convey how many sequences are extractable at different probabilities in each area of the book. Instead, at each character position, we plot the *maximum* extraction probability p_z across all overlapping sliding-window sequences z whose suffixes cover that character position. Such heatmaps facilitate comparisons of memorization “hot-spots” across LLMs at the same book location (Section 6). This heatmap shows high extraction probabilities exist throughout the entirety of *1984*. With respect to probabilistic extraction (Equation 1) measured with 50-token prefixes and suffixes, the entirety of George Orwell’s *1984* is memorized inside LLAMA 3.1 70B’s weights.⁵ On its own, this result does *not* mean that the whole book can be extracted in one continuous segment at generation time. We return to reconstructing an entire book at generation time in Section 7.

4.2 Average extraction rates obscure book-specific memorization

Our sliding-window procedure provides a detailed, book-specific view of memorization that is distinct from prior work. While prior work shows that LLMs memorize certain amounts of the data they were trained on [e.g., 38, 99, 123, 157, 192, 260], that work quantifies memorization through overall extraction rates—much like those in Figure 3 (left). For some (part of a) training dataset, researchers draw (typically at random) a set $\mathbb{Z}$ of sequences z of a specified length (e.g., the first 100 tokens).

⁵The version of *1984* in Books3 has an editor’s forward and appendix that are not memorized. The appendix is before selections of other writing by George Orwell, which are also highly memorized.



Model	Book	% of book extracted with $p_z \geq$		
		$\tau = 1\%$	$\tau = 10\%$	$\tau = 50\%$
LLAMA 3.1 70B	<i>Harry Potter</i>	90.89%	75.44%	43.26%
	<i>Sandman Slim</i>	0.38%	0.28%	0.12%
LLAMA 1 65B	<i>Harry Potter</i>	25.48%	15.00%	4.40%
	<i>Sandman Slim</i>	0.27%	0.13%	0.10%
PYTHIA 12B	<i>Harry Potter</i>	0.40%	0.10%	0.08%
	<i>Sandman Slim</i>	0.34%	0.32%	0.26%

Figure 3: **Average extraction rates are low, but book-specific extraction varies widely.** (left) Comparing extraction rates (Equation 3) of random 100-token sequences (50-token prefix + 50-token suffix) from Books3 for different LLMs trained on Books3. Average extraction is low (Appendix C) for both greedy discoverable and probabilistic ($p_z \geq \tau_{\min} = 0.1\%$, top- k with $T = 1$ and $k = 40$) extraction. (right) Extraction coverage (Equation 4) for specific books reveals a more nuanced picture. For *Harry Potter* [224] and *Sandman Slim* [142], we show extraction coverage for $\tau \in \{1\%, 10\%, 50\%\}$ for 3 models (Appendix D.4). Nearly half of *Harry Potter* can be extracted with respect to sequences that exhibit $p_z \geq \tau = 50\%$.

For each sequence, they prompt with the 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:

$$\text{extraction_rate}(\mathbb{Z}; s) = \frac{1}{|\mathbb{Z}|} \sum_{z \in \mathbb{Z}} \mathbf{1}[s(z)], \quad (3)$$

where $s(z)$ is an extraction success condition. As noted in Section 2.1, for traditional discoverable extraction, the indicator in Equation 3 evaluates to 1 when the greedy completion exactly matches the target suffix. In probabilistic extraction, the indicator evaluates to 1 if the target suffix’s probability p_z exceeds a chosen threshold $\tau_{\min}$ (Equation 2, Appendix C).

These reported averages are generally small, just as we observe in Figure 3 (left) for 9 models trained on Books3. Such small rates often form the basis of defendants in copyright infringement suits calling memorization a rare “bug” [59, 200]. However, while low extraction rates signal that models likely do not memorize *most* text in Books3, *specific* works can still be highly memorized. On average, LLAMA 3.1 70B may have only memorized about 0.6% of sequences in Books3, but concealed within that small number is that it memorized effectively *all* of *1984* (Figure 2).

To make this observation more precise, we quantify how much of a whole book is extractable with **extraction coverage**: the fraction of total characters in a book that lie within at least one suffix whose extraction probability p_z exceeds a threshold τ (Appendix D.4). We choose $\tau \in [\tau_{\min}, 1]$ because $\tau_{\min}$ is the minimum probability we consider to reflect extraction success (Sections 2 & 5). More formally, let $\text{span}(z)$ be the set of character indices spanned by a given 50-token suffix of sequence z . Then

$$\text{coverage}(\tau) = \frac{1}{L} \left| \bigcup_{z: p_z \geq \tau} \text{span}(z) \right|, \text{ where } L \text{ is the character-length of the whole book.} \quad (4)$$

For LLAMA 3.1 70B and *1984*, Figure 2a highlights in red the sequences used to compute Equation 4 for $p_z \geq \tau = 25\%$. These sequences cover 54.39% of *1984*: for each, reproducing the verbatim 50-token suffix would require (in expectation) only 4 prompts with the corresponding 50-token prefix. Most books and LLMs we test do not have this degree of extraction coverage. Figure 4 shows histograms over extraction coverage ($\tau = \tau_{\min} = 0.1\%$) for LLAMA 1 65B and LLAMA 3.1 70B. For both models, most books exhibit $< 2\%$ coverage; however, many books have $> 10\%$ coverage. Overall, LLAMA 3.1 70B exhibits higher coverage on more of the 200 books we evaluate than any other LLM we test (Appendix D.4).

For a more detailed comparison of book-specific coverage, in Figure 3 (right) we show extraction coverage at various choices of τ for *Harry Potter and the Sorcerer’s Stone* [224] and *Sandman Slim* [142]—one of the books by plaintiff Richard Kadrey in *Kadrey et al. v. Meta, Inc.* [143]. *Sandman Slim* is hardly memorized at all. Extraction coverage for this book is less than 0.4% for all LLMs




Figure 4: **LLM-specific distributions over extraction coverage.** Extraction coverage (Equation 4) differs across LLMs for the 200 books we evaluate, illustrated for LLAMA 1 65B and LLAMA 3.1 70B. Results use 100-token sequences (50-token prefix + 50-token suffix), top- k decoding ($T=1$, $k=40$), and $\tau_{\min}=0.1\%$ as the coverage threshold τ .

and thresholds τ .⁶ In contrast, *Harry Potter* [224] is highly memorized by LLAMA models (Section 6). For LLAMA 3.1 70B, 43.26% of the book is extractable with respect to sequences exhibiting $p_z \geq \tau = 50\%$.⁷ For lower τ , there is higher coverage; 90.89% of the book is covered when $\tau = 1\%$. While lower thresholds correspond to extracting text less reliably—i.e., requiring more prompts to the model—for an LLM, all three of these settings reflect enormous probabilities (Section 2.1).

5 Validating our extraction procedure

Above, we emphasized that the minimum probability we count as successful extraction, $\tau_{\min} = 0.1\%$ (Equation 2), is deliberately conservative. But how can we be confident this is the case: that when we measure a successful extraction, we are capturing a true instance of memorization, rather than happenstance generation of a verbatim suffix (a “monkey at the typewriter” event, Section 3)? We address this question through a set of experiments designed to validate our approach. We apply our sliding-window procedure with top- k decoding ($T=1$, $k=40$) and 50-token suffixes (Section 4.1), but we vary prefix length (up to 800 tokens) to observe how conditioning on additional context affects our measurements.

First, we run these settings on books from Books3 and LLAMA 3.1 models, which were trained on Books3 (Section 5.1). As expected [41], longer prefixes surface more memorized training data, but increasing the prefix length does not make every suffix register as extractable (Section 5.1). Second, we run the same procedure on *non-training* data, which tests whether our method spuriously registers false positives on data where memorization is impossible (Section 5.2). Unlike the experiments on training data, regardless of prefix length, suffix probabilities remain below $\tau_{\min}$ (with interesting exceptions⁸), demonstrating our approach does not mistakenly register non-training data as extractable. Together, these experiments show that, when our measurements indicate extraction, we can be confident they are a valid reflection of memorization.⁹ We briefly discuss the implications of these results (Section 5.3), which inform the experiments we run in the following section. Additional results and discussion can be found in Appendix E.

5.1 Longer prefixes surface more memorized training data

When a sequence is known to be in the training data, memorization is the most plausible explanation for high-probability generation of its (sufficiently long) verbatim suffix given the corresponding prefix (Section 2). So, as a baseline, we apply the sliding-window procedure to this setting (Figure 5): we attempt to extract sequences in *We Were Eight Years in Power* [49] (a book in Books3) from

⁶Prior studies (and Figure 3) show that larger LLMs exhibit higher verbatim extraction rates [41]. This pattern does not cleanly generalize to specific works. At $\tau = 50\%$, LLAMA 3.1 70B exhibits 0.12% extraction coverage of *Sandman Slim*, while the smaller PYTHIA 12B exhibits 0.26% coverage.

⁷The version we evaluate from Book3 includes the first chapter of the second *Harry Potter* book at the end.

⁸These sequences are boilerplate text (copyright notices, publisher addresses, etc.) and popular quotes—both of which are duplicated across many sources aside from Books3. (See Appendices E.6 & G.)

⁹Early work on generating copyrighted text with LLMs often did not include such experiments, e.g., Karamolegkou et al. [146]. Those works also tended to run experiments for which ground-truth information about the training data is unknown. Without negative controls to calibrate false positives, it is difficult to interpret reported overlap-based results as valid evidence of extraction (and thus memorization) [63].




Figure 5: **Illustrating validation of our extraction measurements.** The sliding-window procedure for LLAMA 3.1 70B, LLAMA 3.1 8B, and PHI 4 on *We Were Eight Years in Power* [49] (in Books3), using prefix lengths in $\{25, 50, 100, 200, 400, 800\}$ and 50-token suffixes with top- k decoding ($T = 1$, $k = 40$). The book is in LLAMA’s training data, but not PHI 4’s. Longer prefixes reveal more memorized sequences, though not all sequences in the book are extractable as prefix length increases. For PHI 4, the only extractable suffixes are widely duplicated quotes that likely appeared in its training data from other sources. These results support that, when our procedure identifies extraction, it is detecting true instances of memorization rather than chance generations.

LLAMA 3.1 70B and LLAMA 3.1 8B (LLMs trained on Books3). Longer prefixes reveal additional memorized sequences. This is what Carlini et al. [41] call the **discoverability phenomenon**: providing additional context (a longer ground-truth prefix) can raise the conditional probability of the target suffix enough that it becomes extractable ($p_z \geq \tau_{\min}$). Nevertheless, for both LLMs many sequences remain unextractable: regardless of prefix length, their suffix probabilities p_z never surpass $\tau_{\min}$. For instance, LLAMA 3.1 70B memorizes most of the book, yet consistent sections remain unextractable across all prefix lengths (see white gaps in Figure 5, such as from $\sim 75\text{K}$ – 120K characters). The same is even clearer for LLAMA 3.1 8B: being smaller, it memorizes significantly less [38, 41, 123]. Longer prefixes do surface more memorized training data, but most of the book remains unextractable. (Our results suggest that longer prefixes beyond 800 tokens would enable extraction of additional sequences, but the gains diminish; not every sequence is extractable. See Figure 6 and Appendix E.3.)

These results are consistent with prior work: high-quality models do not—indeed cannot [38, 55, 157, 172]—memorize all of the training data in their massive training datasets. Even as prefix length increases, our measurements do not suddenly capture extraction success for everything; they capture success for sequences that have been memorized. We observe the same for LLAMA 3.1 70B on other books in its training data (Figure 6). Some appear effectively not memorized at all, with only limited increases in extraction success at longer prefix lengths (Appendix E.3).

5.2 Measurements on non-training data do not mistakenly register extraction

Using the above results as a basis for comparison, we next test our method in settings where memorization is impossible. This is the role of **negative controls**: experiments in conditions where the outcome of interest should not occur. If (false) positives appear in such a setting, they indicate that the measurement procedure is capturing spurious signal. In our context, running the sliding-window procedure on *non*-training data provides this sanity check. In general (with important exceptions noted below), observing suffix probabilities above $\tau_{\min}$ would indicate false positives—i.e., would count successful extractions where none should exist. The absence of such cases supports that the extractions we do count are valid.

We run two sets of negative controls. First, we apply our procedure to in-copyright Books3 books with PHI 4 [2]—a model that was *not* trained on whole copyrighted books. Second, we apply the sliding-window procedure to LLAMA models on 4 books published in 2025—after the July 2024 cutoff for LLAMA 3.1 training. In heatmaps, we distinguish these conditions by color: **blue** for models trained on Books3 on Books3 text and **orange** for PHI 4 on Books3 text.




Figure 6: **Extraction coverage by prefix length.** Plotting extraction coverage for $\tau = \tau_{\min} = 0.1\%$ (Equation 4) for 5 books we test with varying prefix lengths and LLAMA 3.1 70B. 4 of these books are in Books3 and thus in LLAMA’s training data. The remaining book, *Great Big Beautiful Life* [128], is used in our negative controls on non-training data; extraction coverage is 0% (aside from the copyright notice) regardless of prefix length. As we increase prefix length, extraction coverage for each book in the training data begins to level off. While *Harry Potter* [224] is entirely memorized, it does not appear to be the case that *Beloved* [187], *We Were Eight Years in Power* [49], or *This Is How You Lose Her* [83] are entirely extractable with respect to 50-token suffixes and even longer prefixes.

Negative control 1: An LLM not trained on copyrighted books $\times$ copyrighted books in Books3. We run the same experiment as in Section 5.1 on PHI 4, which was not trained on whole in-copyright books like *We Were Eight Years in Power* [49]. Unlike with LLAMA 3.1 models, we do not observe extraction success for unique text from this book: increasing the context length does not push the conditional probability of non-training-data suffixes above $\tau_{\min}$. We do observe extraction signal, but it only concentrates in four small areas that do not appear to be false positives. Instead, they plausibly reflect memorized text: the suffixes our procedure counts as extracted are famous quotes that are widely duplicated on the Internet, and thus very likely to have been included in PHI 4’s training data via sources other than this copyrighted book [2]. For example, at $\sim 280\text{K}$ characters, we extract passages from the King James Bible [263]; at $\sim 546\text{K}$ characters, we extract a quote from Barack Obama’s 2004 Democratic National Convention speech [196]. Notably, LLAMA 3.1 models also exhibit high extraction probabilities in these same regions. While they were trained on this book, it is plausible that duplicates from other sources contributed to the memorization and extractability of these sequences [157].

Negative control 2: LLMs trained on Books3 $\times$ books published after their training. We apply the same procedure to LLAMA models on books published in 2025—after the July 2024 cutoff for LLAMA 3.1 training. In Figure 6, we show a subset of the results for *Great Big Beautiful Life* [128]. Apart from the (highly duplicated, boilerplate) copyright notice and popular quotes, there are no suffix probabilities that we would count as successful extraction. Not only does p_z never surpass $\tau_{\min} = 0.1\% = 10^{-1}\%$, there are no sequences with $p_z \geq 10^{-13}\%$. For 200-, 400-, and 800-token prefixes, exactly 1 sequence (out of $\sim 62,000$) has $p_z \geq 10^{-14}\%$. This reflects chance generation of non-training data—effectively a “monkey at the typewriter” event. But importantly, this is not a sequence that our method counts as extractable: the suffix probability is 13 orders of magnitude smaller than our setting for $\tau_{\min}$. With a single attempt, the odds of generating the verbatim suffix are $1/p_z \approx 1/10^{-14}\%$, or 1 in 10 quadrillion (Appendix E.4.2). To give a sense of just how small this number is, the chance of winning the Powerball lottery jackpot in a single draw is about 1 in 292 million—roughly 34 million times more likely.

5.3 Implications for measuring probabilistic extraction on books

Three important observations follow from these results and related experiments in Appendix E.4; we highlight them here and defer additional discussion to Appendix E. First, the negative controls provide strong evidence that our sliding-window probabilistic extraction procedure is valid. Second, they indicate that our choice of $\tau_{\min}$ for counting extraction success is indeed conservative. Given the very low probabilities we observe for generating non-training data, we could reasonably set $\tau_{\min}$ below 0.1%. As such, our results almost certainly under-count memorization, and should be read as a lower bound. This is deliberate: we prefer to under-count memorization (i.e., set a higher $\tau_{\min}$) rather than risk over-counting it [39, 122]. To lower $\tau_{\min}$ and remain confident about validity, we

would need to run a much larger set of negative controls. Part of the reason is that, as others have noted [59], LLMs are not actually like monkeys randomly outputting tokens. They have learned to generate structured, grammatical sentences, so the distribution of plausible outputs is far smaller than random chance would suggest (Appendix E.5). Exploring what this implies for setting $\tau_{\min}$ more precisely is an important question that we explore in follow-on work [63].

Third, having validated our procedure, we are justified in extending our experiments beyond LLMs like LLAMA, where we know that Books3 was in the training data. If suffix probabilities surpass $\tau_{\min}$, it is reasonable to interpret these as true extractions and thus evidence of memorization (Appendix E.6). This remains the case regardless of whether the LLM was literally trained on Books3, or whether the memorized training data originate from other overlapping sources—for example, highly duplicated famous quotations, as with PHI 4 in Figure 5 (Appendix E.6). On this basis, in the following section, we expand our analysis to include 3 additional models for which it is unknown if Books3 was included in their training data.

6 Comparing memorization across books and models

We now present the main results of applying the sliding-window procedure. We describe the experimental setup (Section 6.1) and highlight the types of comparisons our approach facilitates across both books and models (Section 6.2).

6.1 Experimental setup

We use the same setup as in Section 4.1: the sliding-window procedure with top- k decoding ($T=1$, $k=40$) for 50-token prefixes and 50-token suffixes. We run the procedure on 14 models, which we organize into three groups: (a) LLMs trained on Books3 (LLAMA [111, 269, 270] and PYTHIA [25], blue); (b) LLMs that were very likely trained on (at least parts of) books also contained in Books3, based on measurements of $p_z \geq \tau_{\min} = 0.1\%$ for unique content from those books (DEEPSEEK v1, QWEN 2.5, and GEMMA 2, red); and (c) an LLM not trained on whole copyrighted books (PHI 4, orange). The code for our experiments can be found [here](#).

Of the $\sim 200,000$ books in Books3, we test 200 books (roughly 0.1%)—the same books for which we analyze extraction coverage in Figure 4. We randomly sample 100 books; the other 100 include works associated with plaintiffs in the amended class action complaint of *Kadrey et al. v. Meta, Inc.* [143, pp. 4-5],¹⁰ as well as a range of generally popular [e.g., 35, 124, 224, 266] and more obscure [e.g., 19, 292] books. We suspected that some popular books would exhibit higher degrees of memorization (e.g., due to duplicated text from other sources). This is why we did not only select books at random; for part of our sample, we deliberately chose titles intended to capture variation across Books3. Further details on these decisions are in Appendix D.1.

In all, we ran over 3000 experiments, amounting to 394 GPU days. Nevertheless, our results should not be read as a complete or generalized account of memorization across Books3. We provide complete per-book results in Appendix D.3, and also publish a [website](#) version.

6.2 Key patterns in our results

We focus here on only a few representative books and on the first two categories of models—(a) models known to be trained on Books3 and (b) models where it is not disclosed whether Books3 was included in the training data. Extraction probabilities p_z can be compared across LLMs, allowing us to evaluate which models are more or less at risk of outputting verbatim 50-token suffixes. We perform these comparisons using heatmaps, as they visualize how memorization varies at specific book locations.¹¹ This enables drawing conclusions about degrees of memorization: how memorization varies for a specific book across models and across books for a given model. We omit heatmaps for LLAMA 3 70B, as they resemble those for LLAMA 3.1 70B.

We show a selection of our results on 200 books and 14 models in Figures 7 and 8. The 3 books in Figure 7 are from the 100 books we manually selected; the 3 in Figure 8 were sampled randomly. Overall, we observe three high-level patterns of memorization in our results; an example of each

¹⁰Post-dating our initial work, that suit was decided in favor of the defendants, Meta [201].

¹¹Sequences z for which we compute p_z may differ, as LLMs have different token vocabularies and tokenizers that which segment text differently. Character locations are fixed across all experiments with the same book.




Figure 7: **Memorization varies widely across books and models (non-random books).** We apply the sliding-window procedure with top- k decoding ($T=1$, $k=40$) and 100-token sequences (50-token prefixes + 50-token suffixes) for LLMs (a) trained on Books3 (blue) and (b) where it is unknown if Books3 was included in the training data (red). With respect to these settings, LLAMA 3.1 70B memorizes effectively all of *Harry Potter and the Sorcerer’s Stone* [224]: extraction coverage is $\approx 96.3\%$.

pattern is provided in each figure. First, there are books like *M. Butterfly* [135] (Figure 7) and *The Third Man* [113] (Figure 8), for which we observe effectively no memorization for any LLM with 50-token prefixes. Aside from the copyright notice, we extract little (if any) text from these books across LLMs of different sizes trained by different organizations. Most of the 200 books we test with 50-token prefixes resemble this pattern. However, for a given book, we are often (but not always) able to extract some substantive sequences from LLAMA 3 and 3.1 70B.

Second, books like *Beloved* [187] (Figure 7) and *Kitchen Table Wisdom* (Figure 8) represent a more intermediate case. For these books, memorization is generally low for most LLMs with respect to 50-token prefixes. However, LLAMA 3 and LLAMA 3.1 70B memorize a considerable number of sequences, and increasing prefix length can sometimes reveal significant amounts of memorization (Appendix E.3). For *Beloved*, 800-token prefixes reveal extraction coverage of nearly 80% from LLAMA 3.1 70B, compared to approximately 12.5% coverage for 50-token prefixes (Figure 6). Other LLAMA models memorize significantly less, and category (b) models (DEEPSEEK v1 67B, QWEN 2.5 72B, and GEMMA 2 27B) memorize even less. For *Beloved*, the 7 largest models we test—category (b) and all large LLAMA models—memorize the same passages at $\sim 170K$ [188, 190] and $\sim 480K$ characters [189]. Both regions contain famous quotes from the book, which we confirmed are highly duplicated on the Internet. In general, we find that specific passages from a given book are memorized across models trained by different organizations—Meta (LLAMA), Alibaba (QWEN), and Google DeepMind (GEMMA). As with PHI 4 (Section 5.2), models that otherwise exhibit low memorization (except for such passages) may not have been trained on the whole book; rather, they may have memorized these passages due to their inclusion and duplication in other sources (Appendix E.6).

Third, some books exhibit enormous degrees of memorization for the 7 largest models, with lesser amounts of memorization in smaller models like LLAMA 1 13B. These are often popular books like *Harry Potter and the Sorcerer’s Stone* [224] (Figure 7), *1984* [203] (Figure 2), and *The Alchemist* [52] (Figure 8). While memorization of such books tends to be significant for larger models, it is especially high for LLAMA 3 and 3.1 70B, which appear to memorize the majority of several books—some in their entirety [e.g., 92, 141, 225, 266]. For instance, with 50-token prefixes, LLAMA 3.1 70B exhibits extraction coverage of approximately 49.1% for *The Alchemist*, while the similarly sized DEEPSEEK v1 67B, QWEN 2.5 72B exhibit approximately 3.9% and 5.8%, respectively (Appendix D.4.2).




Figure 8: **Memorization varies widely across books and models (random books).** We apply the sliding-window procedure with top- k decoding ($T=1$, $k=40$) and 100-token sequences (50-token prefixes + 50-token suffixes) for LLMs (a) trained on Books3 (blue) and (b) where it is unknown if Books3 was included in the training data (red). With respect to these settings, LLAMA 3.1 70B memorizes a substantial portion of *The Alchemist* [52]: extraction coverage is $\approx 49.1\%$.

In general, small models generally memorize less than larger ones, but there is variation across books. For example, LLAMA 1 13B memorizes more of *Harry Potter* than *M. Butterfly* and *Beloved*; LLAMA 3.1 8B memorizes portions of *We Were Eight Years in Power* [49] (Figure 5), but less from other books (Appendix E.3). Variation is more pronounced in larger models, as is clear from extraction coverage distributions (Section 4.2, Appendix D.4). Many results resemble those in Figure 7 for *M. Butterfly*, as is the case for most of the plaintiff-authored books in *Kadrey et al. v. Meta, Inc.* [143] (with exceptions like LLAMA 3.1 70B for *We Were Eight Years in Power*, see Figure 5). However, for LLAMA 3 and LLAMA 3.1 70B, many books yield results similar to those for *Beloved* (e.g., *The Handmaid’s Tale* [12], *The Da Vinci Code* [31], *The Myth of Sisyphus* [35], see Appendix D.3); and, in extreme cases like *Harry Potter* [224], entire books are memorized in the model. LLAMA 1 65B, LLAMA 2 70B, and category (b) models also exhibit high memorization for some, but not all, of these books.

Altogether, these patterns underscore the importance of training choices in shaping memorization. Models do not memorize all of their training data [38, 55, 157, 172]. The mere existence of particular sequences within a training dataset does not imply those sequences will be memorized [157–160]. For example, we observe that similarly sized LLMs trained by different organizations display different degrees of memorization of the same book—even when the evidence strongly suggests the entire book (e.g., *Harry Potter*) was included in their respective training datasets (Appendix E.6). This is also evident with respect to LLAMA models, which were all trained by the same organization (Meta) on Books3. All LLAMA models (by size) memorize more than the other models we evaluate. But notably, within size classes, newer versions typically memorize more from specific books than earlier ones: LLAMA 3 models appear to have memorized more from books than LLAMA 2 models, which in turn memorize more than LLAMA 1.

7 Recovering *Harry Potter and the Sorcerer’s Stone* with one seed prompt

Our main experiments estimate how much of a given book is memorized within an LLM’s weights, with respect to 50-token prefixes and suffixes. These results do not necessarily imply that one would be able to extract large, contiguous portions of a book all at once during generation. However, for LLAMA 3.1 70B on some books like *Harry Potter* and *1984* (Figures 2 & 7), our results show that there exist extraordinarily high extraction probabilities throughout. These probabilities approach or reach 100% when we increase the prefix length (Figure 6 & Appendix E.3). This degree of memorization suggests one could possibly recover the entire book near-verbatim, starting from a single short prompt

"Now, yer mum an' dad were as good a witch an' wizard as I ever knew. Head boy an' girl at Hogwarts in their day! Suppose the myst'ry is why You-Know-Who never tried to get 'em on his side before... probably knew they were too close ter Dumbledore ter want anythin' ter do with the Dark Side.

"Maybe he thought he could persuade 'em... maybe he just wanted 'em outta the way. All anyone knows is, he turned up in the village where you was all living, on Halloween ten years ago. You was just a year old. He came ter yer house an'—an'—¹ an'—an'—²—

Hagrid suddenly pulled out a very dirty, spotted handkerchief and blew his nose with a sound like a foghorn.

"Sorry," he said. "But it's that sad — knew yer mum an' dad, an' nicer people yeh couldn't find — anyway—³ anyway—⁴

"You-Know-Who killed 'em. An' then — an' this is the real myst'ry of the thing — he tried to kill you, too. Wanted ter make a clean job of it, I suppose, or maybe he just liked killin' by then. But he couldn't do it. Never wondered how you got that mark on yer forehead? That was no ordinary cut. That's what yeh get when a powerful, powerful evil curse touches yeh — took care of yer mum an' dad an' yer house, even — but it didn't work on you, an' that's why yer famous, Harry. No one ever lived after he decided ter kill 'em, no one except you, an' he'd killed some o' the best witches an' wizards of the age — the McKinnons, the Bones, the Prewetts — an' you was only a baby, an' you lived."

Something very painful was going on in Harry's mind. As Hagrid's story came to a close, he saw again the blinding flash of green light, more clearly than he had ever remembered it before — and he remembered something else, for the first time in his life: a high, cold, cruel laugh.

Hagrid was watching him sadly.

"Took yeh from the ruined house myself, on Dumbledore's orders. Brought yeh ter this lot...⁵ lot...⁶

"What world?"

Hagrid looked as if he was about to explode.

"DURSLEY!" he boomed.

Uncle Vernon, who had gone very pale, whispered something that sounded like "Mimblewimble." Hagrid stared wildly at Harry.

"But yeh must know about yer mum ⁷mum and dad," he said. "I mean, they're famous. You're famous."

"What? My — my mum ⁸mum and dad weren't famous, were they?"

"Yeh don't know... yeh don't know...⁹ know..." Hagrid ran his fingers through his hair, fixing Harry with a bewildered stare.

"Yeh don't know what yeh are?" he said finally.

Uncle Vernon suddenly found his voice.

"Stop!" he commanded. "Stop right there, sir! I forbid you to tell the boy anything!"

A braver man than Vernon Dursley would have quailed under the furious look Hagrid now gave him; when Hagrid spoke, his every syllable trembled with rage.

"You never told him? Never told him what was in the letter Dumbledore left fer him? I was there! I saw Dumbledore leave it, Dursley! An' you've kept it from him all these years?"

"Kept what from me?" said Harry eagerly.

"STOP! I FORBID YOU!" yelled Uncle Vernon in panic.

Aunt Petunia gave a gasp of horror.

"Ah, go boil yer heads, both of yeh," said Hagrid. "Harry — yer a wizard."

And suddenly, their wardrobes were empty, their trunks were packed, Neville's toad was found lurking in a corner of the toilets; notes were handed out to all students, warning them not to use magic over the holidays ("I always hope they'll forget to give us these," said Fred Weasley sadly). Hagrid was there to take them down to the fleet of boats that sailed across the lake; they were boarding the Hogwarts Express; talking and laughing as the countryside became greener and tidier; eating Bertie Bott's Every Flavor Beans as they sped past Muggle towns; pulling off their wizard robes and putting on jackets and coats; pulling into platform nine and three-quarters at King's Cross ¹⁰station. ¹¹Station

It took quite a while for them all to get off the platform. A wizened old guard was up by the ticket barrier, letting them go through the gate in twos and threes so they didn't attract attention by all bursting out of a solid wall at once and alarming the Muggles.

"You must come and stay this summer," said Ron, "both of you — I'll send you an owl."

"Thanks," said Harry. "I'll need something to look forward to." People jostled them as they moved forward toward the gateway back to the Muggle world. Some of them called:

"Bye, Harry!"

"See you, Potter!"

"Still famous," said Ron, grinning at him.

"Not where I'm going, I promise you," said Harry.

He, Ron, and Hermione passed through the gateway together.

"There he is, Mum, ¹²Mum, there he is, look!"

It was Ginny Weasley, Ron's younger sister, but she wasn't pointing at Ron.

"Harry Potter!" she squealed. "Look, Mum! ¹³Mum! I can see —"

"Be quiet, Ginny, and it's rude to point."

Mrs. Weasley smiled down at them.

Slipping and stumbling, they followed Hagrid down what seemed to be a steep, narrow path. It was so dark on either side of them that Harry thought there must be thick trees there. Nobody spoke much. Neville, the boy who kept losing his toad, sniffed once or twice.

"Yeh'll get yer firs' sight o' Hogwarts in a sec," Hagrid called over his shoulder, "jus' round this bend here."

¹⁴There was a loud "Oooohh!"

The narrow path had opened suddenly onto the edge of a great black lake. Perched atop a high mountain on the other side, its windows sparkling in the starry sky, was a vast castle with many turrets and towers.

"No more'n four to a boat!" Hagrid called, pointing to a fleet of little boats sitting in the water by the shore. Harry and Ron were followed into their boat by Neville and Hermione.

"Everyone in?" shouted Hagrid, who had a boat to himself. "Right then — FORWARD!"

And the fleet of little boats moved off all at once, gliding across the lake, which was as smooth as glass. Everyone was silent, staring up at the great castle overhead. It towered over them as they sailed nearer and nearer to the cliff on which it stood.

"Heads down!" yelled Hagrid as the first boats reached the cliff; they all bent their heads and the little boats carried them through a curtain of ivy that hid a wide opening in the cliff face. They were carried along a dark tunnel, which seemed to be taking them right underneath the castle, until they reached a kind of underground harbor, where they clambered out onto rocks and pebbles.

"Oy, you there! Is this your toad?" said Hagrid, who was checking the boats as people climbed out of them.

Figure 9: Visual diffs of long-form extracted text. Four samples of the diff between the ground-truth text of *Harry Potter* from Books3 and the text we generated using LLAMA 3.1 70B. Text crossed out in red indicates ground-truth text from the book absent in the generation. Text highlighted in yellow is present in the generation, but absent from the ground-truth book. Note that this visual diffing procedure is sensitive to benign formatting differences. We select these four samples to highlight types of variations we observe between the ground-truth book and the generation. In general, the large majority of the diff exhibits no differences (i.e., is black-and-white). See Appendix F.1.

of ground-truth text to seed the generation process. Using beam search as the decoding method—which deterministically approximates the highest probability sequence under the model for a given configuration—LLAMA 3.1 70B could reasonably continue the seed prompt to complete the book.

A simple recovery procedure. We test this idea with about 100 lines of boilerplate generation code using HuggingFace APIs for *Harry Potter* [224]. In essence, this procedure is standard autoregressive generation, where we discard the oldest tokens in the context as new ones are generated. We begin with an s -token seed prompt (e.g., $s = 60$, the first line of *Harry Potter*), and set a maximum sliding context length of n tokens. The LLM generates m tokens at a time with beam search, until reaching n tokens. We then remove the first m tokens from the n -length context and append the generated m tokens, creating a new n -token prompt for the next generation step. We repeat this process for the length of the book. After the first n tokens, the clipped, n -length context window progresses beyond the initial seed prompt. All subsequent prompts do not contain *any* ground-truth text drawn from the book; they consist entirely of generated text from prior iterations.

On our hardware, after testing just a handful of configurations for beam search and a few different short and seed prompts, we *deterministically*¹² generated a near-exact copy of the entire 304-page book. Figure 9 shows partial results using a seed prompt of "Mr. and Mrs. D" ($s = 6$ tokens, drawn from the first line of the first chapter, prepended with the beginning-of-sequence token), beam width of 8, maximum context length $n = 3000$, and generation length $m = 50$ tokens at each step. In practice, LLAMA 3.1 70B tended to predict end-of-sequence (EOS) tokens at the ends of chapters. To address this, we removed the EOS token from the generated m tokens before appending them to the context, replacing it with the tokens for "CHAPTER {c+1}" (with {c+1} spelled out in words, e.g., when $c = 2$, we insert "CHAPTER THREE"). With this simple adjustment, the model successfully continued to generate the book. Because beam search is a deterministic decoding algorithm, our results are reproducible. We discuss the stability of this result and future work beyond this “existence proof” in Appendix F.

Assessing the fidelity of the extracted book. Qualitatively, the generated text is nearly identical to the ground-truth book. Differences predominantly involve small formatting inconsistencies:

¹²Beam search is a deterministic decoding method that explores multiple candidate sequences in parallel to approximate the highest-probability continuation under the model. At each generation step, it keeps the top- k highest-probability sequences so far. At the final step, it outputs the sequence with the highest probability.

white space, capitalization, use of underscores (_) to indicate italics, etc. The Books3 version of *Harry Potter* uses British spelling (e.g., "Mum" instead of "Mom"). Very occasionally, when there is only a single line in a paragraph, the LLM skips that line during generation. We also quantitatively assess the similarity of the two texts. We apply minimal text normalization,¹³ split both documents on whitespace characters, and then measure similarity on the resulting word lists with a greedy approximation of longest common substring. This yields a score of 0.992, where 1 would indicate a perfect match between the ground-truth book and the generation (Appendix F).

8 Takeaways for copyright law and policy

Our experiments show that the extent of memorization varies by model, model family, model size, the book tested, and even within different parts of an individual book (Section 4–6). These results complicate the narratives both plaintiffs and defendants typically present in copyright cases when describing how LLMs work. They support the positions of plaintiffs in some respects and of defendants in others. We see three primary implications for copyright disputes: LLMs could be derivative works of the books they memorize (Section 8.1); in some cases, it may be practical to generate infringing copies (Section 8.2); and, class actions are ill-suited for copyright claims about memorization (Section 8.3).

8.1 LLMs could be derivative works of the books they memorize

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 models themselves are derivative works (Section 3), an argument that courts have thus far not been receptive to [201, 202].

A work is not a derivative work unless it is “substantially similar” in significant part to the original work [e.g., 168]. While prior debates over whether models can qualify as derivative works have been largely theoretical [59, 159], the results of our probabilistic extraction procedure offer empirical evidence that is directly relevant to this issue. They indicate that models are not, as plaintiffs sometimes contend, mere copies of all of the works on which they were trained. However, they also suggest that at least some LLMs (e.g., LLAMA 3.1 70B) may be derivative works of at least some books (e.g., *Harry Potter and the Sorcerer’s Stone*)—because those LLMs have memorized a significant amount of protectable expression from those books. 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. We find that this turns out to be true for some (but not all) books for some (but not all) models.

Fair use. The inquiry does not end here; an LLM, like the training dataset, may be protected under copyright’s fair use doctrine [162]. But the fair use analysis of the model itself may be different than the fair use analysis of the training dataset [159]. Copyright law treats **intermediate copies**, such as those used internally within a system, differently from external-facing copies. A sequence of copyrighted data used for training is an intermediate copy used internally by AI companies for model training, rather than sold as a product. Unless that sequence is extracted, it is only used internally in the course of the model producing outputs. Beyond extraction, individual sequences may to lesser or greater extents influence model outputs, but many of these outputs will not be substantially similar to any training data [126, 159].

A model, too, is an intermediate work in many cases. If the model memorizes content but only generates noninfringing outputs, that is a strong argument that how the model makes use of that content is transformative and ought to be permissible fair use. But for companies that sell or release LLMs to others under open source licenses, the model itself is the product, and sometimes one that is being sold for commercial gain. That does not defeat the fair use argument, but it makes the analogy to training data less precise. In these instances, it may be harder to rely on the cases that justify training on copyrighted data as fair use [e.g., 14, 240, 252].

Potential consequences of a model being a copy of works it has memorized. A finding that the model itself is a copy of some copyrighted works, and that a **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 from HuggingFace in May 2025 (Figure 224)—far less from the height of its release in 2024. If we say conservatively that the model has been downloaded 1 million times

¹³We remove _ (used to signify italics in the Books3 version) and align ellipses to " . . . " (which appear as " . . . " in Books3).

since its release,¹⁴ 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.

In general, damages for copyright infringement can run very high: **statutory damages** can be as high as \$150,000 per infringed work, and **actual damages** can be much higher. 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 [64]. We have only studied a relatively small number of books (200) and models (14), so cannot draw strong conclusions about the extent of memorization of books in general. In principle, if the overall quantity of memorized books in open-weight models is particularly high, it is not an exaggeration to say that the risk of damages across all of those books and of orders of destruction are an existential threat to the companies training these models.

Policy considerations. The effect of these copyright rules may be to treat openly distributed models less favorably than models gated behind closed systems, because the distributed models contain a latent copy or derivative of some copyrighted works. However, there is no good policy reason to treat open-weight models more harshly in copyright law than proprietary ones. Policymakers may want to act to protect models from being declared unlawful copies, in order to prevent crushing financial liability on developers for providing a model that generates transformative, noninfringing outputs. However, a blanket protection for all open-weight models may also be unwarranted. The extent to which models memorize their training data is significantly influenced by choices during training, such as de-duplicating the training dataset [41, 59, 123, 157, 159]. Policymakers may be less inclined to protect models if the developers made training choices known to substantially increase the risk of memorization—for example, deliberately duplicating particular books.

8.2 It may be practical to generate infringing copies in some cases

In our main experiments with the sliding-window probabilistic extraction procedure (Sections 4 & 6), we use relatively short 50-token prompts drawn from books to see if it is possible for an LLM to generate verbatim the following 50 tokens with high probability. As is standard in ML, we use extraction of text snippets for evidence of memorization inside the model. But no one would realistically use this technique to extract an entire book. Doing so would require access to at least half of the book to use as 50-token prefixes in order to even attempt to produce the other half in 50-token suffixes. And, for sequences that are memorized, it might take in expectation up to $1/\tau_{\min} = 1,000$ tries to produce each verbatim suffix, with no indication of which of the many responses represents the actual text of the book.¹⁵

Traditional extraction measurements vs. recovery of (whole) books. Reliable extraction of large segments of text would require other approaches, like the recovery procedure that we apply to *Harry Potter and the Sorcerer’s Stone* with LLAMA 3.1 70B (Section 7). This result is an extreme case and is not representative of most of the other books we have tested (Appendix F.3). Nevertheless, it is significant: in contrast to claims by some defendants [57], it illustrates that, for LLMs that exhibit high degrees of memorization of specific works, it may be possible to generate copies or derivative works of those works. General assessments of the extent of this type of behavior, and how easy it is to elicit it, require more detailed exploration.

Practicality of extraction and responsibility for infringement. Even though our results suggest it may be possible to recover large segments of copyrighted text from an LLM, they do not imply that it will always be easy to do so. This may affect who courts view as directly liable for any output infringement [159]. As a policy matter, the case for liability against the model developer is stronger if simple prompts generate infringing outputs [98]. The case is weaker if a user (e.g., a plaintiffs’ lawyer) employs more intensive strategies to try to generate infringing material. In that case, the adversarial user—who perhaps runs hundreds or thousands of prompts to try to get one output that is infringing—seems more like a direct volitional actor misusing the model [65]; they should likely bear some responsibility for the infringement.¹⁶ By contrast, in several cases (e.g., *Harry Potter, 1984*), it

¹⁴Meta announced in March 2025 that total downloads for all of its LLAMA models surpassed 1 billion [176]. In 2024, the *Times* reported that LLAMA 2 models had been downloaded 180 million times [136].

¹⁵Relatively low extraction probabilities (e.g., $p_z = 0.1\%$) do not always mean that it would take a large number of attempts to extract the suffix. In some cases, other decoding schemes like beam search may make extraction of such suffixes deterministic, requiring just one attempt [62].

¹⁶Models are used *at scale*—in some cases, by millions of users. Hypothetically, a million users interacting with a particular model might independently prompt the model with the prefix of a memorized sequence. Even

does not take a thousand or even hundreds of generations to get large amounts of text; it takes only a handful. And in one case that we have identified so far (*Harry Potter*), it is possible to generate the whole book directly from a short starting prompt and boilerplate generation code (Section 7). We will address this further in future work.

Considerations for LLM training and infringement. The fact that extraction is probabilistic may affect resolution of the “is training infringement” issue. However, it is not determinative. Using a work as training data does not mean that it will necessarily be memorized. And, even if it is memorized, it may still be difficult to extract that work. Large parts of a work would need to be extractable in order to create **superseding copies** (i.e., a copy that can stand in for/ replace the work). Based on our current results, this seems unlikely (though possible, see Section 7) in practice for most works.

Nonetheless, the fact that memorization occurs at all creates a point of distinction from *Author’s Guild v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015) [14]. In this case, the court held that Google’s scanning of books to create a searchable database, and Google’s public display of limited text snippets on Google Books, was fair use. The court also considered and rejected an argument that hackers could gain unauthorized access to Google’s internal books database that supported the service, and thereby obtain plaintiffs’ entire works [1, at 227-28]. In contrast, 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.

8.3 Class actions are ill-suited for copyright claims about memorization

To be certified as a **class action**, plaintiffs generally have to show that *all* putative class members share common legal and factual issues, as well as suffer common injuries. Because many of the pending cases are proceeding as class actions, plaintiffs will therefore need to demonstrate that *all* book owners meet this standard so that a court can certify the class and treat their claims together [90].

As is clear from our results, this may be impossible. Basic questions about whether a particular model actually incorporated any significant expression from any individual plaintiff’s book cannot be generalized. Rather, it varies case by case: the extent of memorization differs widely from model to model, and even from work to work within a model (Section 6). This makes it difficult to assess, on a class-wide basis, whether a particular model copied a particular work and whether, for that model, infringing outputs based on memorization are even possible. Indeed, we show that memorization of works varies for many of the actual named authors in lawsuits, and even across books for the same plaintiffs (e.g., Ta-Nehisi Coates, see Figure 5; Appendices D.3.32–D.3.34 & E.3). Answering questions about the degree to which a specific LLM has memorized a specific work requires experimental evidence—for example, running every book through tests similar to the ones we performed here. Some plaintiffs may be able to show copying, but others will not. Courts generally deny class certification in such circumstances [e.g., 277].

9 Conclusion

We develop (Section 4) and validate (Section 5) a sliding-window procedure to estimate the extent to which open-weight LLMs memorize different books contained in the Books3 training corpus. Our results show that the extent of verbatim memorization of books in open-weight LLMs is far more significant than previously described. Memorization varies widely from model to model and from book to book within each model (Section 6). With respect to our extraction procedure, most LLMs do not memorize most books, but this is not always the case. In extreme instances, an LLM may memorize a book so significantly that it is possible to extract the entire book near-verbatim, as we have shown with LLAMA 3.1 70B for *Harry Potter and the Sorcerer’s Stone* (Section 7). Altogether, our results complicate current disputes over copyright infringement, both by rejecting simplistic claims made by both sides about how LLMs work and by demonstrating that there is no single answer to the question of how much an LLM memorizes (Section 8).

for a verbatim suffix that (in expectation) appears only once for every 1,000 prompts (i.e., $p_z = \tau_{\min} = 0.1\%$), in expectation 1,000 of these users could expect to see the verbatim memorized suffix in the output.

Author contributions

AFC conceived of the project, wrote the code, designed and conducted the experiments, ran the analysis of the results, made the figures, and built the project website. MAL and ABC led the legal analysis, with support by AFC. AFC, MAL, and ABC drafted the manuscript. CDS provided technical feedback. CDS, MAL, PL, and DEH gave iterative feedback on visualizations. AC and AA helped compile books metadata and the supplementary material. PL, DEH, CDS, AA, and AC provided feedback on the manuscript. DEH provided the computing resources for the experiments. AG provided access to Books3.

Resources

We release an interactive, searchable version of our main results on 200 books and 14 models [online](#), as well as a general-purpose coding suite for probabilistic extraction on [GitHub](#).

Acknowledgments and disclosures

This work was supported by a Stanford HAI Seed Research Grant.

We thank the Stanford Research Computing Center for providing computational resources and support that contributed to these research results.

We thank Franziska Boenisch, danah boyd, Nicholas Carlini, Alexandra Chouldechova, James Grimmermann, Adam Dziedzić, Jamie Hayes, Matthew Jagielski, Kevin Klyman, Katherine Lee, Milad Nasr, Matthew Sag, Pamela Samuelson, Florian Tramèr, Danny Wilf-Townsend, and others for helpful discussion and comments on earlier versions of this project. We also thank Timothy B. Lee for his extremely clear explanation of our work [161], which helped us simplify the presentation of the math in this paper.

Ahmed Ahmed acknowledges generous support from a Knight-Hennessy Fellowship, an NSF Graduate Research Fellowship, and a Georgetown Foundation Research Grant.

A. Feder Cooper is an Assistant Professor of Computer Science at Yale University, a Research Scientist at AVERI, and a Postdoctoral Affiliate at Stanford University in Percy Liang’s group in the Department of Computer Science and Daniel E. Ho’s group at Stanford Law School. Until December 2025, Cooper was a full-time employee of Microsoft, working as a postdoctoral researcher within Microsoft Research. Nothing in this paper should be attributed to Microsoft or AVERI.

Mark A. Lemley is a full-time professor at Stanford Law School. He also has a consulting practice as a partner at Lex Lumina LLP. As part of that practice he has represented AI companies in copyright litigation. Nothing in this paper should be attributed to those companies.

Percy Liang is a full-time professor at Stanford, and he is also a co-founder and consultant at Together AI, Simile AI, and Taiku AI; and also an advisor to Virtue AI and Humanitas AI. Nothing in this paper should be attributed to those companies.

Christopher De Sa is a full-time professor at Cornell, and he is also a consultant at Together AI, and an advisor at Sambanova Systems and Inception Labs. Nothing in this paper should be attributed to those companies.

References

- [1] 804 F.3d.
- [2] Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Xin Wang, Rachel Ward, Yue Wu, Dingli Yu, Cyril Zhang, and Yi Zhang. Phi-4 Technical Report, 2024. URL <https://arxiv.org/abs/2412.08905>.
- [3] Chinua Achebe. *Things Fall Apart*. Anchor Books; Random House, Inc., 1958.

- [4] Douglas Adams. *The Hitchhiker's Guide to the Galaxy - Omnibus*. Pan Books; Pan Macmillan; William Heinemann Ltd, 1979.
- [5] Chimamanda Ngozi Adichie. *Americanah*. Alfred A. Knopf; Alfred A. Knopf Canada, 2013.
- [6] Shahad Al Rawi. *The Baghdad Clock*. Oneworld Publications; Dar al-Hikma, 2016.
- [7] amongglue/books3-subset-raw, 2025. URL <https://huggingface.co/datasets/amongglue/books3-subset-raw>.
- [8] Anthropic. Introducing 100K Context Windows, May 2023. URL <https://www.anthropic.com/index/100k-context-windows/>.
- [9] Kelley Armstrong. *Industrial Magic*. Bantam Dell, 2004.
- [10] Isaac Asimov. *Fantastic Voyage*. Bantam Dell; Random House, Inc.; Bantam Books, 1966.
- [11] Isaac Asimov. *The Complete Robot*. Dream Letters Corporation; Dream Letters Corp., 1982.
- [12] Margaret Atwood. *The Handmaid's Tale*. Houghton Mifflin Harcourt Publishing Company, 1985.
- [13] Jane Austen. *Pride and Prejudice*. Penguin Books Ltd, 1813.
- [14] Author's Guild v. Google, Inc., 804 F.3d 202 (2d Cir. 2015).
- [15] David Baldacci. *The Christmas Train*. Warner Books, Inc.; Hachette Book Group, 2001.
- [16] James Baldwin. *Notes of a Native Son*. Beacon Press, 1955.
- [17] James Baldwin. *Another Country*. Michael Joseph; Penguin Books, 1962.
- [18] Julian Barnes. *The Lemon Table*. Vintage International, 2004.
- [19] Teodolinda Barolini. *Dante and the Origins of Italian Literary Culture*. Fordham University Press, 2006.
- [20] Mary Beard. *The Parthenon*. Profile Books Ltd, 2002.
- [21] Charles Beardsley. *Guam: Past and Present*. Charles E. Tuttle Company, Inc., 1964.
- [22] Samuel Beckett. *Waiting for Godot*. Grove Press; Grove/Atlantic, Inc., 1952.
- [23] Helen Benedict. *The Lonely Soldier*. Beacon Press, 2009.
- [24] Mary Berry. *Simple Cakes*. BBC Books; BBC Worldwide Ltd, 2006.
- [25] Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pages 2397–2430. PMLR, 2023.
- [26] Emile Borel. La mecanique statique et l'irreversibilite. *J. Phys. Theor. Appl.*, 3(1):189–196, 1913. doi: 10.1051/jphysap:019130030018900. URL <https://doi.org/10.1051/jphysap:019130030018900>.
- [27] C.J. Box. *Paradise Valley*. Head of Zeus Ltd, 2017.
- [28] Ray Bradbury. *The Cat's Pajamas*. William Morrow; HarperCollins e-books, 2004.
- [29] Gillian Bradshaw. *London in Chains*. Severn House Publishers Ltd, 2009.
- [30] John Brockman. *My Einstein*. Pantheon Books; Random House of Canada Limited, 2006.
- [31] Dan Brown. *The Da Vinci Code*. Transworld Publishers; Bantam Press; Corgi, 2003.
- [32] Gareth Brown. *The Society of Unknowable Objects*. HarperCollins, 2025.
- [33] Niobia Bryant. *Live and Learn*. Dafina Books; Kensington Publishing Corp., 2007.
- [34] Mika Brzezinski. *Knowing Your Value*. Weinstein Books, 2011.
- [35] Albert Camus. *The Myth of Sisyphus*. Penguin Classics; Penguin Books Ltd, 1942.
- [36] CANBERT/pile_books3_text, 2025. URL https://huggingface.co/datasets/CANBERT/pile_books3_text.
- [37] Nicholas Carlini. 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>.

- [38] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650, 2021.
- [39] Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramer. Membership Inference Attacks From First Principles, 2022. URL <https://arxiv.org/abs/2112.03570>.
- [40] Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tramèr, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting Training Data from Diffusion Models, 2023.
- [41] Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. Quantifying Memorization Across Neural Language Models. In *International Conference on Learning Representations*, 2023.
- [42] Lewis Carroll. *Alice’s Adventures in Wonderland and Through the Looking-Glass and What Alice Found There*. Penguin Classics; Penguin Books Ltd, 1871.
- [43] Jeffrey A. Carver. *The Infinity Link*. Open Road Integrated Media, Inc., 1984.
- [44] Jacqueline Charlesworth. Generative AI’s Illusory Case for Fair Use. *Vanderbilt Journal of Entertainment and Technology Law*, 27, 2025.
- [45] Chat GPT Is Eating the World, 2024. URL <https://chatgptiseatingtheworld.com>.
- [46] Agatha Christie. *Murder on the Orient Express*. HarperCollins e-books, 1933.
- [47] Agatha Christie. *And Then There Were None*. HarperCollins e-books, 1939.
- [48] Ta-Nehisi Coates. *The Beautiful Struggle*. Spiegel & Grau; Random House, Inc., 2008.
- [49] Ta-Nehisi Coates. *We Were Eight Years in Power*. One World; Random House; Penguin Random House LLC, 2017.
- [50] Ta-Nehisi Coates. *The Water Dancer*. One World; Random House; Penguin Random House LLC, 2019.
- [51] Jean Cocteau. *The Infernal Machine and Other Plays*. New Directions; New Directions Publishing Corporation, 1963.
- [52] Paulo Coelho. *The Alchemist*. HarperCollins Publishers, 1988.
- [53] Jon Cogburn and Mark Silcox. *Dungeons and Dragons and Philosophy*. Open Court Publishing Company, 2012.
- [54] Annie Cohen-Solal. *Mark Rothko*. Yale University Press, 2013.
- [55] Jasmine Collins, Jascha Sohl-Dickstein, and David Sussillo. Capacity and Trainability in Recurrent Neural Networks, 2017. URL <https://arxiv.org/abs/1611.09913>.
- [56] Suzanne Collins. *The Hunger Games*. Scholastic Children’s Books; Scholastic Ltd, 2008.
- [57] Concord Music Group, Inc. v. Anthropic PBC. 3:23-cv-01092 (M.D. Tenn.).
- [58] Glen Cook. *The Dragon Never Sleeps*. Night Shade Books, 1988.
- [59] A. Feder Cooper and James Grimmelman. The Files are in the Computer: Copyright, Memorization, and Generative AI. *arXiv preprint arXiv:2404.12590*, 2024.
- [60] A. Feder Cooper, Katherine Lee, James Grimmelman, Daphne Ippolito, Christopher Callison-Burch, Christopher A. Choquette-Choo, Niloofar Mireshghallah, Miles Brundage, David Mimno, Madiha Zahrah Choksi, Jack M. Balkin, Nicholas Carlini, Christopher De Sa, Jonathan Frankle, Deep Ganguli, Bryant Gipson, Andres Guadamuz, Swee Leng Harris, Abigail Z. Jacobs, Elizabeth Joh, Gautam Kamath, Mark Lemley, Cass Matthews, Christine McLeavey, Corynne McSherry, Milad Nasr, Paul Ohm, Adam Roberts, Tom Rubin, Pamela Samuelson, Ludwig Schubert, Kristen Vaccaro, Luis Villa, Felix Wu, and Elana Zeide. Report of the 1st Workshop on Generative AI and Law. *arXiv preprint arXiv:2311.06477*, 2023.
- [61] A. Feder Cooper, Christopher A. Choquette-Choo, Miranda Bogen, Matthew Jagielski, Katja Filippova, Ken Ziyu Liu, Alexandra Chouldechova, Jamie Hayes, Yangsibo Huang, Niloofar Mireshghallah, Ilia Shumailov, Eleni Triantafyllou, Peter Kairouz, Nicole Mitchell, Percy Liang, Daniel E. Ho, Yejin Choi, Sanmi Koyejo, Fernando Delgado, James Grimmelman,

- Vitaly Shmatikov, Christopher De Sa, Solon Barocas, Amy Cyphert, Mark Lemley, danah boyd, Jennifer Wortman Vaughan, Miles Brundage, David Bau, Seth Neel, Abigail Z. Jacobs, Andreas Terzis, Hanna Wallach, Nicolas Papernot, and Katherine Lee. Machine Unlearning Doesn't Do What You Think: Lessons for Generative AI Policy, Research, and Practice. *arXiv preprint arXiv:2412.06966*, 2024.
- [62] A. Feder Cooper, Mark A. Lemley, Christopher De Sa, Lea Duesterwald, Allison Casasola, Jamie Hayes, Katherine Lee, Daniel E. Ho, and Percy Liang. Estimating near-verbatim extraction risk in language models with decoding-constrained beam search. *arXiv preprint arXiv:2603.24917*, 2026.
 - [63] A. Feder Cooper, Marika Swanberg, Jamie Hayes, Lea Duesterwald, Christopher De Sa, Daniel E. Ho, Mark A. Lemley, and Percy Liang. Extractable Memorization From First Principles. *arXiv preprint arXiv:2607.12649*, 2026.
 - [64] 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>.
 - [65] CoStar Grp., Inc. v. LoopNet, Inc., 373 F.3d 544 (4th Cir. 2004).
 - [66] Stephen R. Covey. *The 7 Habits of Highly Effective People*. RosettaBooks LLC, 1989.
 - [67] Tyler Cowen. Toothpick Producers Violate NYT Copyright. *Marginal Revolution*, December 2023. URL <https://marginalrevolution.com/marginalrevolution/2023/12/toothpick-producers-violate-nyt-copyright.html>.
 - [68] David Crabb. *Bad Kid*. HarperCollins e-books; HarperCollins Publishers Inc., 2015.
 - [69] Robert Crais. *Lullaby Town*. Bantam Books; Random House, Inc., 1992.
 - [70] Michael Crichton. *Jurassic Park*. Ballantine Books; The Random House Publishing Group, 1990.
 - [71] Michael Cunningham. *The Hours*. Picador; Farrar, Straus and Giroux, 1998.
 - [72] Amy B. Cyphert. Generative AI, Plagiarism, and Copyright Infringement in Legal Documents. *Minnesota Journal of Law, Science & Technology*, 25, 2024.
 - [73] Józef Czapski. *Inhuman Land*. New York Review Books; The New York Review of Books, 1949.
 - [74] Roald Dahl. *James and the Giant Peach*. Puffin Books, 1961.
 - [75] Roald Dahl. *Charlie and the Chocolate Factory*. Puffin Books; Penguin Books Ltd, 1964.
 - [76] DeepSeek-AI et al. DeepSeek LLM: Scaling Open-Source Language Models with Longtermism, 2024. URL <https://arxiv.org/abs/2401.02954>.
 - [77] Nicholas Diakopoulos. *Automating the News*. Harvard University Press, 2019.
 - [78] Joan Didion. *The White Album*. Zola Books; Simon & Schuster, 1979.
 - [79] Cory Doctorow. *Down and Out in the Magic Kingdom*. Tor; Tom Doherty Associates, LLC, 2003.
 - [80] Carol Ann Duffy. *The World's Wife*. Picador; Pan Macmillan, 1999.
 - [81] Junot Díaz. *Drown*. Riverhead Books; The Berkley Publishing Group, 1995.
 - [82] Junot Díaz. *The Brief Wondrous Life of Oscar Wao*. Riverhead Books; Penguin Group (USA) Inc., 2007.
 - [83] Junot Díaz. *This Is How You Lose Her*. Riverhead Books; Penguin Group (USA) Inc., 2012.
 - [84] Jennifer Egan. *A Visit from the Goon Squad*. Alfred A. Knopf; Random House of Canada Limited, 2010.
 - [85] EleutherAI/the_pile_deduplicated, 2025. URL https://huggingface.co/datasets/EleutherAI/the_pile_deduplicated.
 - [86] Ralph Ellison. *Invisible Man*. Penguin Books; Penguin Random House, 1952.
 - [87] Chelsea Fairless and Lauren Garroni. *We Should All Be Mirandas*. Houghton Mifflin Harcourt Publishing Company, 2019.
 - [88] Cathleen Falsani. *The Dude Abides*. Zondervan, 2009.

- [89] Christopher Farnsworth. *The President's Vampire*. G. P. Putnam's Sons; Penguin Books Ltd, 2011.
- [90] Fed. R. Civ. P. 23(a)(2). (requiring questions of law or fact common to the class before it can be certified).
- [91] Vitaly Feldman. 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, page 954–959, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369794.
- [92] F. Scott Fitzgerald. *The Great Gatsby*. Harper Perennial Classics; HarperCollins Publishers (Canada) Ltd., 1925.
- [93] Gillian Flynn. *Gone Girl*. Crown Publishers, 2012.
- [94] Norman Friedman. *British Destroyers*. Seaforth Publishing; Pen & Sword Books Ltd, 2009.
- [95] Barbara Fudurich, Joan Hansen, Caroline Linscott, and Geri Medway. *Florals & Botanicals*. Walter Foster Publishing, Inc., 2008.
- [96] Neil Gaiman and Terry Pratchett. *Good Omens*. HarperCollins Publishers, 1990.
- [97] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The Pile: An 800GB Dataset of Diverse Text for Language Modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- [98] GEMA v. OpenAI. Gesellschaft für musikalische Aufführungs- und mechanische Vervielfältigungsrechte. 42 O 14139/24.
- [99] Gemini Team, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.
- [100] Melanie Gideon. *The Slippery Year*. Alfred A. Knopf; Random House of Canada Limited, 2009.
- [101] Stefanie Girard. *Sweater Surgery*. Quarry Books; Quayside Publishing Group, 2010.
- [102] Malcolm Gladwell. *Blink*. Penguin Books; Allen Lane; Little, Brown and Company, 2005.
- [103] Richard Glover. *The Land Before Avocado*. HarperCollins Publishers Australia Pty Limited, 2018.
- [104] Christopher Golden. *Dead Ringers*. St Martin's Press, 2015.
- [105] Christopher Golden. *Ararat*. St. Martin's Press, 2017.
- [106] William Golding. *Lord of the Flies*. Penguin Books; Penguin Random House LLC, 1954.
- [107] Terry Goodkind. *Wizard's First Rule*. RosettaBooks LLC; Tor Books, 1994.
- [108] Martin Goodman. *Rome and Jerusalem*. Alfred A. Knopf; Allen Lane, 2007.
- [109] Michael P. Goodyear. Artificial Infringement. In *Proceedings of the 2025 Symposium on Computer Science and Law*, CSLAW '25, page 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>.
- [110] Google Cloud Blog. BFloat16: The Secret to High Performance on Cloud TPUs, 2019. URL <https://cloud.google.com/blog/products/ai-machine-learning/bfloat16-the-secret-to-high-performance-on-cloud-tpus>. "the dynamic range of bfloat16 is identical to that of FP32".
- [111] Aaron Grattafiori et al. The Llama 3 Herd of Models, 2024. URL <https://arxiv.org/abs/2407.21783>.
- [112] John Green. *The Fault in Our Stars*. Dutton Books; Penguin Group (USA) Inc., 2012.
- [113] Graham Greene. *The Third Man*. Vintage Digital; Vintage Press, 1950.
- [114] Andrew Sean Greer. *The Confessions of Max Tivoli*. Picador; Farrar, Straus and Giroux, 2004.
- [115] John Grisham. *The Fugitive*. Dutton Children's Books; Penguin Group (USA) LLC, 2015.

- [116] Mark Haddon. *The Curious Incident of the Dog in the Night-Time*. Vintage Books; Random House of Canada Limited, 2003.
- [117] Sue Halpern. *Migrations to Solitude*. Pantheon Books; Random House of Canada Limited; Vintage Books, 1992.
- [118] Tom Hanks. *Uncommon Type*. Alfred A. Knopf, 2017.
- [119] Thor Hanson. *Buzz*. Icon Books Ltd; Basic Books, 2018.
- [120] Harry Potter Wiki. Mistakes in the Harry Potter books, 2025. URL https://harrypotter.fandom.com/wiki/Mistakes_in_the_Harry_Potter_books.
- [121] Elizabeth Haydon. *Requiem for the Sun*. Tor Books; Tom Doherty Associates, LLC, 2002.
- [122] Jamie Hayes, Ilia Shumailov, Christopher A. Choquette-Choo, Matthew Jagielski, George Kaissis, Katherine Lee, Milad Nasr, Sahra Ghalebikesabi, Niloofar Mireshghallah, Meenatchi Sundaram Mutu Selva Annamalai, Igor Shilov, Matthieu Meeus, Yves-Alexandre de Montjoye, Franziska Boenisch, Adam Dziedzic, and A. Feder Cooper. Strong Membership Inference Attacks on Massive Datasets and (Moderately) Large Language Models. *arXiv preprint arXiv:2505.18773*, 2025.
- [123] Jamie Hayes, Marika Swanberg, Harsh Chaudhari, Itay Yona, Ilia Shumailov, Milad Nasr, Christopher A. Choquette-Choo, Katherine Lee, and A. Feder Cooper. Measuring memorization in language models via probabilistic extraction. In Luis Chiruzzo, Alan Ritter, and Lu Wang, editors, *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)*, pages 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/>.
- [124] Joseph Heller. *Catch-22*. Simon & Schuster Paperbacks; Simon & Schuster, 1961.
- [125] Ernest Hemingway. *The Old Man and the Sea*. Harper Perennial Classics; HarperCollins Publishers (Canada) Ltd., 1952.
- [126] Peter Henderson, Xuechen Li, Dan Jurafsky, Tatsunori Hashimoto, Mark A. Lemley, and Percy Liang. Foundation Models and Fair Use, 2023. URL <https://arxiv.org/abs/2303.15715>.
- [127] David Hendy. *Life on Air*. Oxford University Press; Oxford University Press Inc., 2007.
- [128] Emily Henry. *Great Big Beautiful Life*. Berkley, 2025.
- [129] Luke Hersheson. *Great Hair Days*. Ebury Press; Penguin Random House, 2018.
- [130] Matt Hill. *Graft*. Angry Robot; Watkins Media Ltd, 2016.
- [131] S. E. Hinton. *The Outsiders*. Penguin Books; The Viking Press; Puffin Books, 1967.
- [132] Tanya Huff. *The Second Summoning*. DAW Books, 2001.
- [133] Hugging Face. Quantization Concepts, 2025. URL https://huggingface.co/docs/optimum/v1.7.1/concept_guides/quantization. "bfloat16, accumulation data type float32".
- [134] Langston Hughes. *Selected Poems of Langston Hughes*. Vintage Books; Random House of Canada Limited, 1959.
- [135] David Henry Hwang. *M. Butterfly*. Plume; Penguin Group (USA) Inc., 1988.
- [136] Mike Isaac. How A.I. Made Mark Zuckerberg Popular Again in Silicon Valley. *The New York Times*, May 2024. URL <https://www.nytimes.com/2024/05/29/technology/mark-zuckerberg-meta-ai.html>.
- [137] Allen Jackson. *Building and Operating a Realistic Model Railway*. The Crowood Press Ltd, 2016.
- [138] Betty Jacobs. *All the Onions*. Storey Publishing, LLC, 1993.
- [139] E. L. James. *Fifty Shades of Grey*. The Writer’s Coffee Shop Publishing House; Vintage Books; Random House of Canada Limited, 2011.
- [140] N. K. Jemisin. *The Stone Sky*. Orbit, 2017.

- [141] James Joyce. *Ulysses*. Penguin Books Ltd, 1922.
- [142] Richard Kadrey. *Sandman Slim*. HarperCollins e-books; HarperCollins Publishers Inc., 2009.
- [143] Kadrey et al. v. Meta Platforms, Inc. URL <https://www.courthousenews.com/wp-content/uploads/2025/02/kadrey-vs-meta-third-amended-complaint.pdf>. No. 3:23-cv-03417-VC (Third Amended Consolidated Complaint).
- [144] Kadrey v. Meta Platforms, Inc. No. 3:23-cv-03417 (N.D. Cal. July 7, 2023).
- [145] Anthony Kaldellis. *Ethnography after Antiquity*. University of Pennsylvania Press, 2013.
- [146] Antonia Karamolegkou, Jiaang Li, Li Zhou, and Anders Søgaard. Copyright Violations and Large Language Models, 2023. URL <https://arxiv.org/abs/2310.13771>.
- [147] Matthew Klam. *Who Is Rich?* Random House; Penguin Random House LLC, 2017.
- [148] Dean Koontz. *The Servants of Twilight*. The Berkley Publishing Group; Penguin Group (USA) Inc., 1984.
- [149] Aihan Kuhn. *Tai Chi for Depression*. YMAA Publication Center, Inc., 2017.
- [150] Josh Kun. *The Tide Was Always High*. University of California Press, 2017.
- [151] Jean Kwok. *Girl in Translation*. Riverhead Books; Penguin Group (USA) Inc., 2010.
- [152] Eduardo Lago. *Call Me Brooklyn*. Dalkey Archive Press, 2006.
- [153] Erik Larson. *Dead Wake*. Crown Publishers, 2015.
- [154] Stieg Larsson. *The Girl with the Dragon Tattoo*. Alfred A. Knopf, 2005.
- [155] Ursula K. Le Guin. *The Daughter of Odren*. Houghton Mifflin Harcourt Publishing Company, 2014.
- [156] Dave Lee. Facebook denies targeting insecure users. *BBC*, May 2017. URL <https://www.bbc.com/news/technology-39775782>.
- [157] Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. Deduplicating Training Data Makes Language Models Better. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 8424–8445, 2022.
- [158] Katherine Lee, A. Feder Cooper, James Grimmelmann, and Daphne Ippolito. AI and Law: The Next Generation. *SSRN*, 2023. <http://dx.doi.org/10.2139/ssrn.4580739>.
- [159] Katherine Lee, A. Feder Cooper, and James Grimmelmann. Talkin’ ’Bout AI Generation: Copyright and the Generative-AI Supply Chain. *arXiv preprint arXiv:2309.08133*, 2023.
- [160] Katherine Lee, A. Feder Cooper, and James Grimmelmann. 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, page 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>.
- [161] Timothy B. Lee. Study: Meta AI model can reproduce almost half of Harry Potter book. *Ars Technica*, June 2025. URL <https://arstechnica.com/ai/2025/06/study-metas-llama-3-1-can-recall-42-percent-of-the-first-harry-potter-book/>.
- [162] Mark Lemley and Bryan Casey. Fair Learning. *Texas Law Review*, 99:743, 2021.
- [163] Madeleine L’Engle. *A Wrinkle in Time*. Puffin Books; Penguin Books Ltd, 1962.
- [164] C. S. Lewis. *The Chronicles of Narnia*. Penguin Random House, 1950.
- [165] Laura Lippman. *After I’m Gone*. HarperCollins e-books, 2014.
- [166] Laura Lippman. *Sunburn*. Faber & Faber Ltd; William Morrow, 2018.
- [167] David Liss. *Marvel’s Spider-Man: Hostile Takeover*. Titan Books; Titan Publishing Group Ltd, 2018.
- [168] 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.”).
- [169] Lois Lowry. *Anastasia on Her Own*. Houghton Mifflin Books for Children; Houghton Mifflin Harcourt Publishing Company, 1985.

- [170] Iris Ma, Ian Domingo, Alberto Krone-Martins, Pierre Baldi, and Cristina V. Lopes. Memorization: A Close Look at Books, 2025. URL <https://arxiv.org/abs/2504.12549>.
- [171] Margaret MacMillan. *Nixon in China*. Penguin Group (Canada), 2006.
- [172] Sadegh Mahdavi, Renjie Liao, and Christos Thrampoulidis. Memorization Capacity of Multi-Head Attention in Transformers, 2024. URL <https://arxiv.org/abs/2306.02010>.
- [173] Scott Mariani. *The Doomsday Prophecy*. Avon; HarperCollins Publishers Ltd., 2009.
- [174] George R. R. Martin. *A Game of Thrones*. Bantam Books, 1996.
- [175] Keith Heyer Meldahl. *Rough-Hewn Land*. University of California Press; University of California Press, Ltd., 2011.
- [176] Meta Newsroom. Celebrating 1 Billion Downloads of Llama, March 2025. URL <https://about.fb.com/news/2025/03/celebrating-1-billion-downloads-llama/>.
- [177] Stephenie Meyer. *Twilight*. Little, Brown and Company; Hachette Book Group, 2005.
- [178] Courtney Milan. *The Duchess War*. Courtney Milan (self-published), 2012.
- [179] John Stuart Mill. *On Liberty*. Penguin Books, 1859.
- [180] Alex Miller. *Coal Creek*. Allen & Unwin, 2013.
- [181] A. A. Milne. *Winnie the Pooh*. Dutton Children’s Books; Penguin Group (USA) Inc., 1926.
- [182] Stephanie Mlot. Zuckerberg Vows to Protect Free Speech After Paris Attack. *PCMag*, January 2015. URL <https://www.pcmag.com/news/zuckerberg-vows-to-protect-free-speech-after-paris-attack>.
- [183] Colten Moore and Keith O’Brien. *Catching the Sky*. 37 INK/Atria Books; Simon & Schuster, Inc., 2016.
- [184] Michelle Moran. *The Heretic Queen*. Crown Publishers; Random House, Inc., 2008.
- [185] Laura Morelli. *The Gondola Maker*. Laura Morelli; Smashwords, 2014.
- [186] Mary McGarry Morris. *Songs in Ordinary Time*. Open Road Integrated Media, Inc., 1995.
- [187] Toni Morrison. *Beloved*. Vintage Books; Random House of Canada Limited; Alfred A. Knopf, 1987.
- [188] Toni Morrison. *Beloved*, 1987. URL <https://www.goodreads.com/quotes/386603-in-this-here-place-we-flesh-flesh-that-weeps-laugh>. Part One, Chapter 9.
- [189] Toni Morrison. *Beloved*, 1987. URL <https://www.goodreads.com/quotes/11207714-leave-before-sethe-could-make-her-realize-that-worse-than>. Part One, Chapter 25.
- [190] Toni Morrison. *Beloved*, 1987. URL <https://www.goodreads.com/quotes/578852-i-used-to-think-to-think-it-was-my-rememory>. Part One, Chapter 9.
- [191] Haruki Murakami. *Norwegian Wood*. Vintage; The Harvill Press; Kodansha Ltd, 1987.
- [192] Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A. Feder Cooper, Daphne Ippolito, Christopher A. Choquette-Choo, Eric Wallace, Florian Tramèr, and Katherine Lee. Scalable Extraction of Training Data from (Production) Language Models. *arXiv preprint arXiv:2311.17035*, 2023.
- [193] Milad Nasr, Javier Rando, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A. Feder Cooper, Daphne Ippolito, Christopher A. Choquette-Choo, Florian Tramèr, and Katherine Lee. 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>.
- [194] Desiree Nielsen. *Eat More Plants*. Penguin Canada; Penguin Random House Canada Limited, 2019.
- [195] Michael Northrop. *Polaris*. Scholastic Press; Scholastic Inc., 2017.
- [196] Barack Obama. Keynote Address at the 2004 Democratic National Convention, 2004. URL <https://www.presidency.ucsb.edu/documents/keynote-address-the-2004-democratic-national-convention>.
- [197] James J. O’Donnell. *Pagans*. HarperCollins e-books, 2015.

- [198] Yoko Ogawa. *The Memory Police*. Pantheon Books, 1994.
- [199] Robin Oliveira. *Winter Sisters*. Viking; Penguin Random House LLC, 2018.
- [200] OpenAI. OpenAI and journalism, January 2024. URL <https://openai.com/blog/openai-and-journalism>.
- [201] Order Denying the Plaintiffs’ Motion for Partial Summary Judgment and Granting Meta’s Cross-Motion for Partial Summary Judgment, Kadrey et al. v. Meta Platforms, Inc. URL <https://law.justia.com/cases/federal/district-courts/california/candce/3:2023cv03417/415175/598/>. No. 3:23-cv-03417-VC (N.D. Cal. Jun. 25, 2025).
- [202] Order on Fair Use, Bartz et al. v. Anthropic PBC. URL <https://docs.justia.com/cases/federal/district-courts/california/candce/3:2024cv05417/434709/231>. No. 3:23-cv-03417-VC (N.D. Cal. Jun. 25, 2025).
- [203] George Orwell. *Nineteen Eighty-Four*. Penguin Books Ltd; Martin Secker & Warburg Ltd, 1949.
- [204] Meghan L. O’Sullivan. *Windfall*. Simon & Schuster, 2017.
- [205] Chuck Palahniuk. *Fight Club*. W. W. Norton & Company, Inc.; W. W. Norton & Company Ltd., 1996.
- [206] Charlie Papazian. *The Complete Joy of Homebrewing*. HarperCollins Publishers Inc., 1984.
- [207] Paul Park. *The Cult of Loving Kindness*. ElectricStory.com, Inc.; William Morrow, 1991.
- [208] James Patterson. *Along Came a Spider*. Little, Brown and Company; Hachette Book Group, 1993.
- [209] François Payard and Anne E. McBride. *Payard Cookies*. Houghton Mifflin Harcourt, 2015.
- [210] Tom Piazza. *Why New Orleans Matters*. HarperCollins Publishers; Harper Perennial; HarperCollins e-books, 2005.
- [211] Steven Pinker. *Enlightenment Now*. Viking; Penguin Random House LLC, 2018.
- [212] Michael E. Porter. *Competitive Strategy*. The Free Press; Simon & Schuster Inc., 1980.
- [213] USVSN Sai Prashanth, Alvin Deng, Kyle O’Brien, Jyothir S V au2, Mohammad Aflah Khan, Jaydeep Borkar, Christopher A. Choquette-Choo, Jacob Ray Fuehne, Stella Biderman, Tracy Ke, Katherine Lee, and Naomi Saphra. Recite, Reconstruct, Recollect: Memorization in LMs as a Multifaceted Phenomenon, 2024. URL <https://arxiv.org/abs/2406.17746>.
- [214] Terry Pratchett. *Night Watch*. HarperCollins Publishers, 2002.
- [215] Philip Pullman. *The Subtle Knife*. Alfred A. Knopf, Inc.; Scholastic Children’s Books, 1997.
- [216] Jacques Pépin. *Essential Pepin Desserts*. Houghton Mifflin Harcourt, 2011.
- [217] Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 Technical Report, 2025. URL <https://arxiv.org/abs/2412.15115>.
- [218] Matthew Rampley. *The Seductions of Darwin*. The Pennsylvania State University Press, 2017.
- [219] Alex Reisner. 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/>.
- [220] Rachel Naomi Remen. *Kitchen Table Wisdom*. Riverhead Books; Penguin Group (USA) Inc., 1996.
- [221] Jessi Roberts and Bret Witter. *Backroads Boss Lady*. Grand Central Publishing; Hachette Digital, 2019.
- [222] Marie-Sabine Roger. *Soft in the Head*. Pushkin Press, 2008.
- [223] Ramzi Rouighi. *The Making of a Mediterranean Emirate*. University of Pennsylvania Press, 2011.

- [224] J. K. Rowling. *Harry Potter and the Sorcerer's Stone*. Pottermore Limited; Arthur A. Levine Books, 1997.
- [225] J. K. Rowling. *Harry Potter and the Chamber of Secrets*. Pottermore Limited; Arthur A. Levine Books, 1998.
- [226] J. K. Rowling. *Harry Potter and the Goblet of Fire*. Pottermore Limited; Arthur A. Levine Books; Scholastic Inc., 2000.
- [227] J. K. Rowling. *Harry Potter and the Deathly Hallows*. Pottermore Limited; Arthur A. Levine Books; Scholastic Inc., 2007.
- [228] Dan Rubinstein. *Born to Walk*. ECW Press, 2015.
- [229] Marc Ruskin. *The Pretender*. Thomas Dunne Books; St. Martin's Press, 2017.
- [230] Harvey Sachs. *Toscanini*. Liveright Publishing Corporation; W. W. Norton & Company, Inc., 1978.
- [231] Matthew Sag. Copyright Safety for Generative AI. *Houston Law Review*, 61, 2024.
- [232] Matthew Sag. Fairness and Fair Use in Generative AI. *Fordham Law Review*, 92, 2024.
- [233] Carl Sagan. *Cosmos*. Ballantine Books; Random House of Canada Limited, 1980.
- [234] Bhisham Sahni. *Middle India*. Penguin Books India Pvt. Ltd, 2001.
- [235] J. D. Salinger. *The Catcher in the Rye*. Little, Brown and Company, 1951.
- [236] Pamela Samuelson. Generative AI meets copyright. *Science*, 381(6654):158–161, 2023.
- [237] Sheryl Sandberg. *Lean in*. Alfred A. Knopf; Random House of Canada Limited, 2013.
- [238] SaylorTwift/the_pile_books3_minus_gutenberg, 2025. URL https://huggingface.co/datasets/SaylorTwift/the_pile_books3_minus_gutenberg.
- [239] Avi Schwarzschild, Zhili Feng, Pratyush Maini, Zachary C Lipton, and J Zico Kolter. Rethinking llm memorization through the lens of adversarial compression. *arXiv preprint arXiv:2404.15146*, 2024.
- [240] Sega Enters. Ltd. v. Accolade, Inc., 977 F.2d 1510 (9th Cir. 1992), as amended (Jan. 6, 1993).
- [241] Aarti Shahani. Facebook Wants Great Power, But What About Responsibility? *NPR*, February 2017. URL <https://www.npr.org/sections/alltechconsidered/2017/02/19/516094134/facebook-wants-great-power-but-what-about-responsibility>.
- [242] Sanjeev Sharma. *The DevOps Adoption Playbook*. John Wiley & Sons, Inc., 2017.
- [243] Mary Shelley. *Frankenstein*. Wordsworth Classics, 1818.
- [244] Lynn Sherr. *Sally Ride: America's First Woman in Space*. Simon & Schuster, 2014.
- [245] Lionel Shriver. *A Perfectly Good Family*. HarperCollins Publishers Inc., 1996.
- [246] Sarah Silverman. *The Bedwetter*. HarperCollins e-books, 2010.
- [247] Larry Sloman. *On the Road with Bob Dylan*. Three Rivers Press; Bantam Books, 1978.
- [248] Alexander Gordon Smith. *The Night Children*. Macmillan, 2008.
- [249] Zadie Smith. *White Teeth*. Vintage International, 2000.
- [250] Rachel Louise Snyder. *No Visible Bruises*. Bloomsbury Publishing Inc., 2019.
- [251] Benjamin Sobel. Elements of Style: Copyright, Similarity, and Generative AI. *Harvard Journal of Law & Technology*, 38, 2024.
- [252] Sony Computer Ent., Inc. v. Connectix Corp., 203 F.3d 596 (9th Cir. 2000).
- [253] John Steinbeck. *The Grapes of Wrath*. Penguin Books; The Viking Press Inc.; William Heinemann Ltd, 1939.
- [254] T. J. Stiles. *Jesse James*. Vintage Books; Random House of Canada Limited; Alfred A. Knopf, 2002.
- [255] R. L. Stine. *Zombie Halloween*. Scholastic Inc., 2014.
- [256] David Stubbs. *Fear of Music*. Zero Books; John Hunt Publishing Ltd., 2009.
- [257] Jennifer Swanson. *Pearl Harbor*. Scholastic Inc., 2018.

- [258] Earl Swift. *Chesapeake Requiem*. HarperCollins Publishers, 2018.
- [259] Donna Tartt. *The Goldfinch*. Little, Brown and Company; Hachette Book Group, 2013.
- [260] Gemma Team et al. Gemma 2: Improving Open Language Models at a Practical Size, 2024. URL <https://arxiv.org/abs/2408.00118>.
- [261] Lysa TerKeurst. *Unglued*. Zondervan, 2012.
- [262] Lysa TerKeurst. *Embraced*. Thomas Nelson, 2018.
- [263] The Bible (King James Version). Deuteronomy 15:12-15. URL <https://www.kingjamesbibleonline.org/Deuteronomy-15-12/>.
- [264] 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 . . . [I]et’s see. ’It was the best of times, it was the blurst of times?’ You stupid monkey!”).
- [265] Lynn Thomson. *Birding with Yeats*. House of Anansi Press Inc., 2014.
- [266] J. R. R. Tolkien. *The Hobbit*. Houghton Mifflin Harcourt, 1937.
- [267] J. R. R. Tolkien. *The Fellowship of the Ring*. Houghton Mifflin Harcourt; George Allen & Unwin, 1954.
- [268] J. R. R. Tolkien. *Tree and Leaf*. HarperCollins Publishers, 1964.
- [269] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. LLaMA: Open and Efficient Foundation Language Models, 2023. URL <https://arxiv.org/abs/2302.13971>.
- [270] Hugo Touvron et al. Llama 2: Open Foundation and Fine-Tuned Chat Models, 2023. URL <https://arxiv.org/abs/2307.09288>.
- [271] Corinne Trang. *Noodles Every Day*. Chronicle Books LLC, 2009.
- [272] George R. Tyler. *Billionaire Democracy*. BenBella Books, Inc., 2018.
- [273] 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>.
- [274] Al J. Venter. *Portugal’s Guerrilla Wars in Africa*. Helion & Company Limited, 2013.
- [275] Kurt Vonnegut. *Slaughterhouse-Five*. RosettaBooks LLC, 1969.
- [276] Ocean Vuong. *The Emperor of Gladness*. Penguin Press, 2025.
- [277] 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).
- [278] Paul Waldau. *Animal Rights*. Oxford University Press, Inc., 2011.
- [279] Jesmyn Ward. *Men We Reaped*. Bloomsbury USA, 2013.
- [280] E. B. White. *Charlotte’s Web*. HarperCollins e-books; HarperCollins Publishers Inc., 1952.
- [281] Daniel Wilf-Townsend. The Deletion Remedy. *North Carolina Law Review*, 103, 2024. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4933011. Forthcoming 2025.
- [282] Marianne Williamson. *A Return to Love*. HarperCollins Publishers, 1992.
- [283] Julia Carrie Wong. SpaceX rocket explosion: Mark Zuckerberg laments loss of Internet.org satellite. *The Guardian*, September 2016. URL <https://www.theguardian.com/science/2016/sep/01/mark-zuckerberg-spacex-explosion-africa-internet>.
- [284] Jacqueline Woodson. *Brown Girl Dreaming*. Nancy Paulsen Books; Penguin Group (USA) LLC, 2014.
- [285] Jacqueline Woodson. *Another Brooklyn*. HarperCollins e-books, 2016.
- [286] Sarah Wynn-Williams. *Careless People: A Cautionary Tale of Power, Greed, and Lost Idealism*. Macmillan, 2025.

- [287] Malcolm X. Who Taught You To Hate Yourself?, 1962. URL <https://genius.com/Malcolm-x-who-taught-you-to-hate-yourself-annotated>.
- [288] Hanya Yanagihara. *A Little Life*. Doubleday; Random House of Canada Limited, 2015.
- [289] Yuji Yoshimura and Giovanna M. Halford. *The Art of Bonsai*. Tuttle Publishing; Charles E. Tuttle Publishing Co, Inc., 1996.
- [290] Chiyuan Zhang, Daphne Ippolito, Katherine Lee, Matthew Jagielski, Florian Tramèr, and Nicholas Carlini. Counterfactual memorization in neural language models. *Advances in Neural Information Processing Systems*, 36:39321–39362, 2023.
- [291] Howard Zinn. *A People’s History of the United States*. HarperCollins Publishers; Harper Perennial Modern Classics, 1980.
- [292] Jonathan Zittrain. *The Future of the Internet and How to Stop It*. Yale University Press, 2008.
- [293] Markus Zusak. *The Book Thief*. Alfred A. Knopf; Random House of Canada Limited; Picador; Pan Macmillan Pty Ltd., 2005.

Appendix Table of Contents

A	Additional notes on memorization and probabilistic extraction	35
A.1	Additional notes on the “careless people” example	35
A.2	Metrics	35
A.3	Computing sequence probabilities in one forward pass	37
B	Testing our measurement pipeline	38
C	Experiments on extraction rates	39
C.1	Setup	40
C.2	Experimental procedure	40
C.3	High-level takeaways	41
D	Extended results for sliding-window experiments	42
D.1	Setup	42
D.2	Sliding-window probabilistic extraction procedure	53
D.3	Book-specific results	55
D.3.1	<i>Things Fall Apart</i> , Achebe	56
D.3.2	<i>The Hitchhiker’s Guide to the Galaxy - Omnibus</i> , Adams	56
D.3.3	<i>Americanah</i> , Adichie	57
D.3.4	<i>The Baghdad Clock</i> , Al Rawi	57
D.3.5	<i>Industrial Magic</i> , Armstrong	58
D.3.6	<i>Fantastic Voyage</i> , Asimov	58
D.3.7	<i>The Complete Robot</i> , Asimov	59
D.3.8	<i>The Handmaid’s Tale</i> , Atwood	59
D.3.9	<i>Pride and Prejudice</i> , Austen	60
D.3.10	<i>The Christmas Train</i> , Baldacci	60
D.3.11	<i>Notes of a Native Son</i> , Baldwin	61
D.3.12	<i>Another Country</i> , Baldwin	61
D.3.13	<i>The Lemon Table</i> , Barnes	62
D.3.14	<i>Dante and the Origins of Italian Literary Culture</i> , Barolini	62
D.3.15	<i>The Parthenon</i> , Beard	63
D.3.16	<i>Guam: Past and Present</i> , Beardsley	63
D.3.17	<i>Waiting for Godot</i> , Beckett	64
D.3.18	<i>The Lonely Soldier</i> , Benedict	64
D.3.19	<i>Simple Cakes</i> , Berry	65
D.3.20	<i>Paradise Valley</i> , Box	65
D.3.21	<i>The Cat’s Pajamas</i> , Bradbury	66
D.3.22	<i>London in Chains</i> , Bradshaw	66
D.3.23	<i>My Einstein</i> , Brockman	67
D.3.24	<i>The Da Vinci Code</i> , Brown	67
D.3.25	<i>Live and Learn</i> , Bryant	68
D.3.26	<i>Knowing Your Value</i> , Brzezinski	68
D.3.27	<i>The Myth of Sisyphus</i> , Camus	69
D.3.28	<i>Alice’s Adventures in Wonderland</i> , Carroll	69
D.3.29	<i>The Infinity Link</i> , Carver	70
D.3.30	<i>Murder on the Orient Express</i> , Christie	70
D.3.31	<i>And Then There Were None</i> , Christie	71
D.3.32	<i>The Beautiful Struggle</i> , Coates	71
D.3.33	<i>We Were Eight Years in Power</i> , Coates	72
D.3.34	<i>The Water Dancer</i> , Coates	72
D.3.35	<i>The Infernal Machine</i> , Cocteau	73
D.3.36	<i>The Alchemist</i> , Coelho	73
D.3.37	<i>Dungeons and Dragons and Philosophy</i> , Cogburn and Silcox	74

D.3.38	Mark Rothko, Cohen-Solal	74
D.3.39	The Hunger Games, Collins	75
D.3.40	The Dragon Never Sleeps, Cook	75
D.3.41	The 7 Habits of Highly Effective People, Covey	76
D.3.42	Bad Kid, Crabb	76
D.3.43	Lullaby Town, Crais	77
D.3.44	Jurassic Park, Crichton	77
D.3.45	The Hours, Cunningham	78
D.3.46	Inhuman Land, Czapski	78
D.3.47	Charlie and the Chocolate Factory, Dahl	79
D.3.48	James and the Giant Peach, Dahl	79
D.3.49	Automating the News, Diakopoulos	80
D.3.50	Drown, Díaz	80
D.3.51	The Brief Wondrous Life of Oscar Wao, Díaz	81
D.3.52	This Is How You Lose Her, Díaz	81
D.3.53	The White Album, Didion	82
D.3.54	Down and Out in the Magic Kingdom, Doctorow	82
D.3.55	The World's Wife, Duffy	83
D.3.56	A Visit from the Goon Squad, Egan	83
D.3.57	Invisible Man, Ellison	84
D.3.58	We Should All Be Mirandas, Fairless and Garroni	84
D.3.59	The Dude Abides, Falsani	85
D.3.60	The President's Vampire, Farnsworth	85
D.3.61	The Great Gatsby, Fitzgerald	86
D.3.62	Gone Girl, Flynn	86
D.3.63	British Destroyers, Friedman	87
D.3.64	Florals & Botanicals, Fudurich et al.	87
D.3.65	Good Omens, Gaiman and Pratchett	88
D.3.66	The Slippery Year, Gideon	88
D.3.67	Sweater Surgery, Girard	89
D.3.68	Blink, Gladwell	89
D.3.69	The Land Before Avocado, Glover	90
D.3.70	Dead Ringers, Golden	90
D.3.71	Ararat, Golden	91
D.3.72	Lord of the Flies, Golding	91
D.3.73	Wizard's First Rule, Goodkind	92
D.3.74	Rome and Jerusalem, Goodman	92
D.3.75	The Fault in Our Stars, Green	93
D.3.76	The Third Man, Greene	93
D.3.77	The Confessions of Max Tivoli, Greer	94
D.3.78	The Fugitive, Grisham	94
D.3.79	The Curious Incident of the Dog in the Night-Time, Haddon	95
D.3.80	Migrations to Solitude, Halpern	95
D.3.81	Uncommon Type, Hanks	96
D.3.82	Buzz, Hanson	96
D.3.83	Requiem for the Sun, Haydon	97
D.3.84	Catch-22, Heller	97
D.3.85	The Old Man and the Sea, Hemingway	98
D.3.86	Life on Air, Hendy	98
D.3.87	Great Hair Days, Hersheson	99
D.3.88	Graft, Hill	99
D.3.89	The Outsiders, Hinton	100
D.3.90	The Second Summoning, Huff	100
D.3.91	Selected Poems of Langston Hughes, Hughes	101
D.3.92	M. Butterfly, Hwang	101

D.3.93	<i>Building and Operating a Realistic Model Railway</i> , Jackson	102
D.3.94	<i>All the Onions</i> , Jacobs	102
D.3.95	<i>Fifty Shades of Grey</i> , James	103
D.3.96	<i>The Stone Sky</i> , Jemisin	103
D.3.97	<i>Ulysses</i> , Joyce	104
D.3.98	<i>Sandman Slim</i> , Kadrey	104
D.3.99	<i>Ethnography after Antiquity</i> , Kaldellis	105
D.3.100	<i>Who Is Rich?</i> , Klam	105
D.3.101	<i>The Servants of Twilight</i> , Koontz	106
D.3.102	<i>Tai Chi for Depression</i> , Kuhn	106
D.3.103	<i>The Tide Was Always High</i> , Kun	107
D.3.104	<i>Girl in Translation</i> , Kwok	107
D.3.105	<i>A Wrinkle in Time</i> , L'Engle	108
D.3.106	<i>Call Me Brooklyn</i> , Lago	108
D.3.107	<i>Dead Wake</i> , Larson	109
D.3.108	<i>The Girl with the Dragon Tattoo</i> , Larsson	109
D.3.109	<i>The Daughter of Odren</i> , Le Guin	110
D.3.110	<i>The Chronicles of Narnia</i> , Lewis	110
D.3.111	<i>After I'm Gone</i> , Lippman	111
D.3.112	<i>Sunburn</i> , Lippman	111
D.3.113	<i>Marvel's Spider-Man: Hostile Takeover</i> , Liss	112
D.3.114	<i>Anastasia on Her Own</i> , Lowry	112
D.3.115	<i>Nixon in China</i> , MacMillan	113
D.3.116	<i>The Doomsday Prophecy</i> , Mariani	113
D.3.117	<i>A Game of Thrones</i> , Martin	114
D.3.118	<i>Rough-Hewn Land</i> , Meldahl	114
D.3.119	<i>Twilight</i> , Meyer	115
D.3.120	<i>The Duchess War</i> , Milan	115
D.3.121	<i>On Liberty</i> , Mill	116
D.3.122	<i>Coal Creek</i> , Miller	116
D.3.123	<i>Winnie the Pooh</i> , Milne	117
D.3.124	<i>Catching the Sky</i> , Moore and O'Brien	117
D.3.125	<i>The Heretic Queen</i> , Moran	118
D.3.126	<i>The Gondola Maker</i> , Morelli	118
D.3.127	<i>Songs in Ordinary Time</i> , Morris	119
D.3.128	<i>Beloved</i> , Morrison	119
D.3.129	<i>Norwegian Wood</i> , Murakami	120
D.3.130	<i>Eat More Plants</i> , Nielsen	120
D.3.131	<i>Polaris</i> , Northrop	121
D.3.132	<i>Pagans</i> , O'Donnell	121
D.3.133	<i>Windfall</i> , O'Sullivan	122
D.3.134	<i>The Memory Police</i> , Ogawa	122
D.3.135	<i>Winter Sisters</i> , Oliveira	123
D.3.136	<i>Nineteen Eighty Four</i> , Orwell	123
D.3.137	<i>Fight Club</i> , Palahniuk	124
D.3.138	<i>The Complete Joy of Homebrewing</i> , Papazian	124
D.3.139	<i>The Cult of Loving Kindness</i> , Park	125
D.3.140	<i>Along Came a Spider</i> , Patterson	125
D.3.141	<i>Payard Cookies</i> , Payard and McBride	126
D.3.142	<i>Essential Pepin Desserts</i> , Pépin	126
D.3.143	<i>Why New Orleans Matters</i> , Piazza	127
D.3.144	<i>Enlightenment Now</i> , Pinker	127
D.3.145	<i>Competitive Strategy</i> , Porter	128
D.3.146	<i>Night Watch</i> , Pratchett	128
D.3.147	<i>The Subtle Knife</i> , Pullman	129

D.3.148	<i>The Seductions of Darwin</i> , Rampley	129
D.3.149	<i>Kitchen Table Wisdom</i> , Remen	130
D.3.150	<i>Backroads Boss Lady</i> , Roberts and Witter	130
D.3.151	<i>Soft in the Head</i> , Roger	131
D.3.152	<i>The Making of a Mediterranean Emirate</i> , Rouighi	131
D.3.153	<i>Harry Potter and the Sorcerer's Stone</i> , Rowling	132
D.3.154	<i>Harry Potter and the Chamber of Secrets</i> , Rowling	132
D.3.155	<i>Harry Potter and the Goblet of Fire</i> , Rowling	133
D.3.156	<i>Harry Potter and the Deathly Hallows</i> , Rowling	133
D.3.157	<i>Born to Walk</i> , Rubinstein	134
D.3.158	<i>The Pretender</i> , Ruskin	134
D.3.159	<i>Toscanini</i> , Sachs	135
D.3.160	<i>Cosmos</i> , Sagan	135
D.3.161	<i>Middle India</i> , Sahni	136
D.3.162	<i>The Catcher in the Rye</i> , Salinger	136
D.3.163	<i>Lean In</i> , Sandberg	137
D.3.164	<i>The DevOps Adoption Playbook</i> , Sharma	137
D.3.165	<i>Frankenstein</i> , Shelley	138
D.3.166	<i>Sally Ride: America's First Woman in Space</i> , Sherr	138
D.3.167	<i>A Perfectly Good Family</i> , Shriver	139
D.3.168	<i>The Bedwetter</i> , Silverman	139
D.3.169	<i>On the Road with Bob Dylan</i> , Sloman	140
D.3.170	<i>The Night Children</i> , Smith	140
D.3.171	<i>White Teeth</i> , Smith	141
D.3.172	<i>No Visible Bruises</i> , Snyder	141
D.3.173	<i>The Grapes of Wrath</i> , Steinbeck	142
D.3.174	<i>Jesse James</i> , Stiles	142
D.3.175	<i>Zombie Halloween</i> , Stine	143
D.3.176	<i>Fear of Music</i> , Stubbs	143
D.3.177	<i>Pearl Harbor</i> , Swanson	144
D.3.178	<i>Chesapeake Requiem</i> , Swift	144
D.3.179	<i>The Goldfinch</i> , Tartt	145
D.3.180	<i>Unglued</i> , TerKeurst	145
D.3.181	<i>Embraced</i> , TerKeurst	146
D.3.182	<i>Birding with Yeats</i> , Thomson	146
D.3.183	<i>The Hobbit</i> , Tolkien	147
D.3.184	<i>The Fellowship of the Ring</i> , Tolkien	147
D.3.185	<i>Tree and Leaf</i> , Tolkien	148
D.3.186	<i>Noodles Every Day</i> , Trang	148
D.3.187	<i>Billionaire Democracy</i> , Tyler	149
D.3.188	<i>Portugal's Guerrilla Wars in Africa</i> , Venter	149
D.3.189	<i>Slaughterhouse-Five</i> , Vonnegut	150
D.3.190	<i>Animal Rights</i> , Waldau	150
D.3.191	<i>Men We Reaped</i> , Ward	151
D.3.192	<i>Charlotte's Web</i> , White	151
D.3.193	<i>A Return to Love</i> , Williamson	152
D.3.194	<i>Another Brooklyn</i> , Woodson	152
D.3.195	<i>Brown Girl Dreaming</i> , Woodson	153
D.3.196	<i>A Little Life</i> , Yanagihara	153
D.3.197	<i>The Art of Bonsai</i> , Yoshimura and Halford	154
D.3.198	<i>A People's History of the United States</i> , Zinn	154
D.3.199	<i>The Future of the Internet and How to Stop It</i> , Zittrain	155
D.3.200	<i>The Book Thief</i> , Zusak	155
D.4	Estimating how much of a book is memorized by a model	156
D.4.1	Quantifying extraction coverage	156

D.4.1	Extended extraction coverage results	156
E	Measurement validity	163
E.1	Motivating our validity experiments	164
E.1.1	Setting $\tau_{\min}$ and extending to models with unknown training data	164
E.1.2	Relating greedy-decoded and probabilistic discoverable extraction	164
E.1.3	The role of prefix length in our experiments	166
E.2	Setup	166
E.3	Results of varying prefix length on known training data	167
E.4	Results of negative control experiments	170
E.4.1	Results on PHI 4	171
E.4.2	Results on books from after training cutoffs	173
E.5	Additional notes on “the monkey at the typewriter” metaphor	176
E.6	Additional notes on extraction and membership inference	178
E.6.1	The relationship between membership inference and extraction	179
E.6.2	We know ground-truth membership for most of our experiments	179
E.6.3	Inferring membership for LLMs with undisclosed training data	180
E.6.4	We infer membership for sequences, not whole books	181
E.6.5	Absence of extraction does not imply anything about membership	181
E.6.6	Extractability and copyright	181
F	Recovering books with one seed prompt	182
F.1	Recovery method	182
F.2	Partial recovery	185
F.3	Running recovery for books with low memorization signal	186
G	Discussion of extended results	186

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 extraneous or whatever it was and let other people clean up the mess they had made. F. Scott Fitzgerald 1

Figure 10: **Illustrating different possible completions of the “careless people” quote.** For a prefix of a sequence from *The Great Gatsby* [92], we show three 32-token completions when using the prefix as a prompt to LLAMA 1 30B (top- k , $T = 1$ and $k = 40$). Each row reflects a different interaction with the LLM: **(left)** the ground-truth prefix, used as the prompt (which is always the same); **(middle)** the target suffix, which is also always the same text; **(right)** the text of the generation in response to the prompt. For the target suffix, faint red highlighting indicates that LLAMA 1 30B did not reproduce the target suffix verbatim for the given interaction, with darker red highlighting indicating words from the target suffix that are missing in the generation. For generations, light green highlighting indicates that LLAMA 1 30B did not reproduce the target suffix verbatim for the given interaction, with darker green highlighted words reflecting text that the generation adds to the ground-truth target suffix. The middle row, where there is no highlighting, reflects verbatim extraction.

A Additional notes on memorization and probabilistic extraction

We provide additional notes on the background material in this paper, including the “careless people” quote in Section 2 (Appendix A.1), the key probabilistic extraction metric in this paper (formally called (n, p) -discoverable extraction, Appendix A.2), as well as our implementation for this metric (Appendix A.3). Please also refer to Section 2.

A.1 Additional notes on the “careless people” example

The headline figure we include in Figure 1 (Section 2) is meant to be illustrative—an instance of, not a general statement about, memorization and extraction. (In Figure 10, we also provide examples of outputs of prompting LLAMA 1 30B with the prefix for this example.) This example also plays a role in some of our validity experiments (Section 5 and Appendix E). The text is drawn from *The Great Gatsby* [92]; it is one of the most famous quotes from the book. It is also a short, 57-token sequence: the prefix is 25 tokens (including the start-of-sequence token) and the suffix is 32 (with the LLAMA 1 tokenizer). 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. While extraction of verbatim 32-token suffixes is often accepted as valid evidence for memorization [25], it is more common practice to use 50-token suffixes. (Regarding validity, also refer to Section 5 and Appendix E.) We compute the extraction probability with respect to top- k decoding ($T = 1$, $k = 40$). We do this by transforming the model’s output logits according to the sampling parameters of this scheme. (See more about this below in Appendix A.3.)

It is also the case that 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 whole 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 for LLMs where the entirety of Books3 is included in the training data. We discuss this further in Appendix E.

A.2 Metrics

The metric that we use for quantifying memorization comes from Hayes et al. [123]. The authors provide a different framing for the metric than we do. Instead of directly comparing p_z values (for the same sequence across LLMs, across multiple sequences for the same LLM, etc.), Hayes et al. [123] specify a threat model with an adversary that can query the model an arbitrary number of times with the same prompt—an adversary that has n independent attempts to extract a sequence z (its suffix) with at least probability p . This is where their metric derives its name, (n, p) -discoverable extraction:




Figure 11: **Plotting p_z for the “careless people” quote from *The Great Gatsby*.** (left) p_z (1) for different models. (right) Translating p_z into how many prompts n it would take to extract the sequence z with at least probability p (5). See Section 2.

Definition 1 ((n, p) -discoverable extraction, from Hayes et al. [123]). Given a training sequence z that is split into an a -length prefix $z_{1:a}$ and a j -length suffix $z_{a+1:a+j}$, z is (n, p) -discoverably extractable if

$$\Pr \left(\bigcup_{w \in [n]} (g_\phi \circ f_\theta)_w^j(z_{1:a}) = z_{1:a+j} \right) \geq p,$$

where $(g_\phi \circ f_\theta)_w^j(z_{1:a})$ 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 $j > 0$ times, starting from the same initial sequence $z_{1:a}$.

On its own, this type of method is well-suited to **completion** models that continue the prompt text, rather than instruction-tuned **chatbots** that reply in a conversational style. Similar to Hayes et al. [123] (and other work on discoverable extraction), we therefore only study completion models and, for brevity, refer to these models by name, version, and size. Also note that this definition can be (at least in principle) easily adapted to non-verbatim extraction by testing if a generation is within ϵ distance from the target suffix for a given distance metric. Despite these differences in framing, in practice for verbatim extraction, computing this metric just involves calculating sequence probabilities p_z . For sequence z , this is the suffix probability (for the given prefix) with respect to an LLM’s base distribution θ and decoding scheme ϕ (e.g., greedy, top- k) that determines which tokens are considered at each step and how their logits are renormalized into conditional probabilities (Appendix A.3). That is,

$$\begin{aligned} p_z &:= p_{z, \theta, \phi} \triangleq \Pr_{\theta, \phi}(z_{a+1:a+j} \mid z_{1:a}) = \prod_{t=a+1}^{a+j} \Pr_{\theta, \phi}(z_t \mid z_{1:t-1}) \\ &= \exp \left(\sum_{t=a+1}^{a+j} \log \Pr_{\theta, \phi}(z_t \mid z_{1:t-1}) \right). \end{aligned} \quad (5)$$

We always note the particular LLM θ and decoding scheme ϕ , and so, for brevity, will refer to this quantity as simply p_z . In practice, we compute p_z as the exp of the sum of conditional log probabilities of the tokens in the sequence, as this is more numerically stable than multiplying together the conditional probabilities of the tokens in the sequence.

In Figure 11 (left), we show examples of p_z for the “careless people” quote (split into prefix and suffix as in Figure 1) for different models. (This is Figure 1, right, in Section 2, with PYTHIA 12B added.) Once p_z is, one can compute n for a given p (for $p_z \in (0, 1)$) with

$$1 - (1 - p_z)^n \geq p \implies n \geq \frac{\log(1 - p)}{\log(1 - p_z)}. \quad (6)$$

Since p_z is the probability of generating the suffix with 1 prompt to the model, the probability of *not* generating it in 1 prompt is $(1 - p_z)$. The probability of *not* generating it in n independent prompts to the model (with the same prompt) is $(1 - p_z)^n$, and so the probability of generating it in n prompts is 1 minus this probability, i.e., $1 - (1 - p_z)^n$.

As an intuition for this presentation of the metric, think about a fair coin flip coming up heads. Here, $p_z = 0.5$ (the probability of heads in one flip). With two actual flips ($n = 2$), the probability of flipping heads *at least once* is $1 - (1 - 0.5)^2 = 0.75 = p$; when $n = 10$, $p = 0.9990$. We visualize




Figure 12: **Probability of an event occurring at least once changes as a function of n independent trials.** Following the intuition of flipping a fair coin (where heads has $p_z = 0.5 = 50\%$), we show how the probability of flipping heads at least once changes with more flips (**left**). We show how the probability p of generating the verbatim suffix of the “careless people” quote at least once for LLAMA 1 30B ($p_z \approx 35.2\%$) changes as a function of the number of independent prompts to the model with the prefix (**right**). Both figures use $1 - (1 - p_z)^n$ (Equation 6), plugging in the respective p_z (for the coin flip and for generating the “careless people” suffix given the prefix) and different values of n to compute p .

this in Figure 12 (left) with coin flips, and show how it translates to suffix generation probabilities for LLAMA 1 30B and the “careless people” quote, for which $p_z = 35.2\%$ (right). As is clear Equation 6, for a given p_z , one can pick a probability threshold p and get the corresponding number of prompts n (or vice versa). Figure 11 (right) visualizes how n changes for different settings of p for the “careless people” quote.

In this paper, we re-frame (n, p) -discoverable extraction to emphasize the quantity p_z (Equation 5) that we actually compute: the probability of extracting a sequence z (for the given LLM, hyperparameter-configured decoding scheme, and suffix start index location within sequence z). This is also useful for plotting distributions over p_z for a given book—to see how probabilities vary for sequences across a book, or for a given sequence across different models. We provide such comparisons in Section 6.

In practice, we can compute p_z for verbatim extraction *without* prompting n times. We compute sequence probabilities directly from the logits we obtain from running whole sequences through the LLM, rather than generating multiple suffixes for each prefix. (See Appendix A.3.) These two approaches are equivalent for verbatim extraction: the computed suffix probabilities correspond to the statistically expected frequencies of verbatim extracted outputs we would observe if we were to generate a large number of continuations for the same prompt. Hayes et al. [123], the authors of this extraction approach, confirm this with ample experimental evidence. We do, as well, though omit these results for brevity because they follow directly from what p_z quantifies. This is also why we also limit ourselves to verbatim extraction in this paper; the efficiencies of operating directly on logits with one forward pass only apply to this setting. (See Section 2 and Appendix A.3 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 200 of the nearly 200,000 books in Books3.

We discuss how we determine extraction success for probabilistic extraction in Appendix C. In the main text, we discuss determining extraction success in Section 2.1.

A.3 Computing sequence probabilities in one forward pass

Our implementation differs from Hayes et al. [123], which we confirm in discussion with the authors. The original paper generated the suffix a token at a time and summed up the per-generated-token conditional 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 inference code. This also makes computing this metric cheaper than traditional greedy-decoded discoverable extraction [38], in terms of wall clock time; one large matrix-matrix multiplication is significantly cheaper than generation with KV-caching. Note that this is the *entire* interaction with the model in our main experiments. We are not making *any* changes to the model in our extraction measurements.

We then implement our sequence probability computations (i.e., compute p_z) as a post-processing operation on the logits. For any given decoding scheme, we can transform the logits—for example,

change the shape according to temperature, truncate them to the top- k —before normalizing them into the log probability distribution. This computation is written out on the right side of Equation 5. 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 conditional log probabilities of the *actual* next tokens in the sequence, one at a time. This is often called **teacher forcing**.

We work with the conditional log probabilities (adding them together), rather than the conditional probabilities (and multiplying them together), as this is more numerically stable to compute. 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 unnormalized scores of the next token. We find the logit value that corresponds to `_orange`, compute log softmax to get the log probability, and add the associated conditional log probability. Then we continue on (doing the same for `_orange`, where we get the conditional log probability for the next token being `_cat`). We get the conditional log probability of the whole suffix by adding together the conditional, per-token log probabilities of the actual suffix. From the suffix log probability, we easily obtain the probability p_z .

Any given continuation should be low probability. That is, for a suffix length of 50 tokens, this computation is effectively multiplying together 50 per-token conditional probabilities (by adding together conditional log probabilities). For a continuation where each token has really high probability, this can still mean the overall continuation has relatively low probability. Consider that each token in a suffix has conditional probability 0.9. This is *really* high; it means that a single token in the logit vector (at each position in the suffix) has 90% of the corresponding probability mass; for a token vocabulary of size 32,000, that means the remaining 31,999 tokens all share the remaining 10%. For this suffix, 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 what this means for the underlying token probabilities. In some cases in this paper, we observe sequences that have *over 90% probability*. Such unusually high probabilities associated with verbatim suffixes in the training data are indicative of memorization (where it is reasonable to consider 0.5%, or even smaller, to be unusually high).

Also note that top- k sampling truncates the logit vector 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 suffix has a token at a position that is *not* in the top- k of the logit vector, then we will not be able to complete the probability computation. In this case, the suffix has $p_z = 0$; it is not extractable with *any* probability.

A refactored version of our experimental code, with additional optimizations for increased efficiency, can be found [here](#).

B Testing our measurement pipeline

To confirm the accuracy of our implementation of probabilistic discoverable extraction, as well as to benchmark its efficiency in comparison to the implementation in Hayes et al. [123], we re-run the experiments from Hayes et al. [123] for small PYTHIA models on the Enron dataset. To ensure consistency, we discussed the original run times and sample of Enron with the authors of Hayes et al. [123]. We provide some brief details about these experiments to show due diligence for our implementation in our particular setting.

We ran these experiments with both `float16` and `float32`, and report results for `float32` to align with Hayes et al. [123]. (However, we note no significant differences between the two.) In Table 1, we show the runtime for our experiments on the 10,000, 100-token sequences Hayes et al. [123] drew from Enron. In Table 2, we report the same metrics as Hayes et al. [123]: the greedy-decoded discoverable extraction rate and (as a point of comparison) the maximum probabilistic rate. We run these experiments on the same 4 A100s as all of our other probabilistic extraction experiments.

We successfully replicate the results of these experiments in Hayes et al. [123]. However, in the process of doing so, we identified a small bug in their original implementation. The PYTHIA tokenizer is a GPT-2-style tokenizer, and does *not* add a beginning of sequence (BOS) token by default to the sequences it tokenizes. However, for our setting, it is best practice to include this token at the start of a sequence, as we submit these sequences as inputs 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 ex-

Table 1: **Runtime for PYTHIA on Enron.** We run 2 PYTHIA models on the same 10,000 100-token sequences that Hayes et al. [123] use from the Enron dataset.

	Time (mm:ss)	Parallelism	Batch size
PYTHIA 1B	02:30.42	Data (four GPUs)	250
PYTHIA 2.8B	03:36.35	Data (four GPUs)	250

Table 2: **Comparing our results to Hayes et al. [123].** We give the greedy discoverable extraction and maximum probabilistic rate for PYTHIA models on the 10,000-sequence subset from Enron. We successfully replicate the exact results in Hayes et al. [123], which did not make use of a beginning of sequence (BOS) token. We compare to the rates where the BOS token is included at the start of the sequence. Note that the probabilistic rate increases slightly with the presence of the BOS token.

PYTHIA 1B on Enron subset	Hayes et al. [123] (No BOS token)	With BOS token
Greedy extraction rate	0.76%	0.74%
Max. probabilistic rate	5.27%	5.52%
PYTHIA 2.8B on Enron subset	Hayes et al. [123] (No BOS token)	With BOS token
Greedy extraction rate	1.3 %	1.82%
Max. probabilistic rate	9.04%	9.47%

periments from Hayes et al. [123], we generally observe elevated rates compared to Hayes et al. [123]. See Table 2. We similarly prepend the BOS token in our Books3 experiments using PYTHIA models.

C Experiments on extraction rates

Similar to other extraction papers, we also compute average extraction rates. These serve as a reference point for our discussion concerning how average rates can conceal interesting patterns of work-specific memorization. In the main text, we refer briefly to these experiments in Section 4.2.

Computing average extraction rates. Let $\mathbb{Z}$ be a set of sequences z drawn from some (part of a) training dataset. Each z is split into a prefix $z_{1:a}$ and target suffix $z_{a+1:a+j}$ (e.g., $a=j=50$ for 100-token sequences).

All extraction rates take the form

$$\text{extraction_rate}(\mathbb{Z}; s) = \frac{1}{|\mathbb{Z}|} \sum_{z \in \mathbb{Z}} \mathbf{1}[s(z)], \quad (7)$$

where $s(z)$ is a success predicate defined according to the chosen criterion.

Fix a model θ and a decoding scheme ϕ (e.g., base distribution $T = 1$, or top- k with $k = 40$), which induces a conditional next-token distribution that can be used to compute suffix probabilities $p_z \in [0, 1]$. For a sequence z with prefix $z_{1:a}$ and suffix $z_{a+1:a+j}$, define the suffix probability under θ and ϕ as

$$p_z := p_{z, \theta, \phi} = \prod_{t=a+1}^{a+j} \text{Pr}_{\theta, \phi}(z_t \mid z_{1:t-1}).$$

This is the probability that the model generates the exact target suffix token by token under decoding scheme ϕ ; it is the same definition as Equation 5 that we use in Section 2 and Appendix A.

For traditional, greedy-decoded discoverable extraction (Carlini et al. [41], Lee et al. [157]; Section 2), $\phi = \text{greedy}$ (i.e., $k = 1$) and the success predicate can be written simply as explicit string equality:

$$s_{\theta, \text{greedy}}(z) = \mathbf{1}[\text{greedy}(\theta; z_{1:a}) = z_{a+1:a+j}]. \quad (8)$$

More generally, for any LLM θ and decoding scheme ϕ , the success predicate just thresholds the suffix probability. Given a minimum threshold $\tau_{\min} \in (0, 1]$,

$$s_{\theta,\phi}(\mathbf{z}; \tau_{\min}) = \mathbf{1}[p_{\mathbf{z},\theta,\phi} \geq \tau_{\min}]. \quad (9)$$

Again, note that greedy decoding is a special case: each top-1 token has probability 1 under $\phi =$ greedy, so $p_{\mathbf{z},\theta,\phi} = 1$ and success reduces to equality with the suffix. In the main paper, we omit θ and ϕ in our notation. We always mention which specific model we are evaluating, and we are always considering ϕ to be top- k decoding with $T=1$ and $k=40$. (See Appendix D.2 for more details on this choice.) When using non-deterministic decoding schemes and suffix probabilities to quantify extraction, we very conservatively set $\tau_{\min}=0.1\%$ (Section 5 & Appendix E).

C.1 Setup

Data. We perform our average extraction rate experiments with the Books3 dataset. We obtained this dataset from a previous (2022) download of The Pile [97], 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 LLMs) remains unresolved. (See Section 3.) Nevertheless, this dataset remains widely available in the research community; for example, it can be found in the HuggingFace-hosted version of The Pile [85], as well as in various other data repositories [e.g., 7, 36, 238]. It also features as part of a benchmark task in the popular HELM evaluation suite. For this set of experiments, we draw random samples from Books3, performing 3 separate runs each on 40,000 sequences. We discuss how we sample these sets of sequences in Appendix C.2.

Models. From 13 different models, we try to extract 40,000 sequences from Books3 (3 times each). We test LLAMA 1 7B, LLAMA 1 13B, LLAMA 1 65B [269], LLAMA 2 7B, LLAMA 2 13B, LLAMA 2 70B [270], LLAMA 3.1 8B, LLAMA 3.1 70B [111], PYTHIA 6.9B, PYTHIA 12B [25], DEEPSEEK v1 7B [76], GEMMA 2 9B [260], and PHI 4 [2] (a 14B model). We group these models into three size groups: small, medium and large. We use `float16` for all models except for LLAMA 3 and LLAMA 3.1 models (and QWEN 2.5 models in sliding-window experiments). These models were explicitly trained to work with `bfloat16`.

We know for certain that PYTHIA and LLAMA models were trained on Books3, and can conclude from our validity experiments (Section 5 and Appendix E) that DEEPSEEK v1, QWEN 2.5 and GEMMA 2 models were trained on (at least parts of) some books that are also contained in Books3. PHI 4 was deliberately not trained on copyrighted (whole) books, but we find that we can still extract some snippets from copyrighted books.

Metrics. We compute extraction rates using Equation 7 for both traditional, greedy-decoded discoverable extraction (using the success condition in Equation 8) and probabilistic extraction with top- k ($T=1$, $k=40$) decoding (using the success condition in Equation 9) and $\tau_{\min}=0.1\%$.

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

C.2 Experimental procedure

Sampling books from Books3. To collect 40,000 sequences for which we compute an extraction curve for a given model, we use the following procedure. We sample 4,000 books without replacement from Books3. Then, from each of these books, we sample 10 non-overlapping 100-token sequences (50-token prefixes + 50-token suffixes). We make sure that the start of the prefix text is a space (to prevent unusual tokenization). We run these 40,000 sequences through each of the 13 models to compute extraction metrics. We follow this procedure 3 separate times (i.e., for a total of 120,000 sequences across 12,000 different books) to get a sense of the variance across extraction metrics for different (empirical) population estimates. Each run of 40,000 sequences took between approximately 8 minutes and 45 minutes on 4 A100s. We omit detailed timing results for brevity.

Probabilistic extraction configuration. We report greedy-decoded discoverable extraction and probabilistic discoverable extraction with top- k ($T=1$, $k=40$) and $\tau_{\min}=0.1\%$. The maximum probabilistic rate (also reported in Appendix B) should be understood as a very loose upper bound on extraction of memorized training data; it counts all sequences with non-zero $p_{\mathbf{z}}$ (i.e., it also includes




Figure 13: **Extended extraction rate results.** Comparing greedy discoverable and probabilistic extraction rates (for a our conservative setting, $p_z \geq \tau_{\min} = 0.1\%$). As a reference point, we also include the maximum possible rate of generating 50-token sequences that match our suffixes drawn from Books3 for small (7–9B), medium (12–14B), and large models (65–70B). (This is akin to looking at all $p_z > 0$.) Probabilistic rates are computed with top- k sampling using $T = 1$ and $k = 40$. The maximum probabilistic rate should be understood as a very loose upper bound on extraction of memorized training data; it also captures generalization. *Average* extraction rates (i.e., $\tau_{\min} = 0.1\%$) on randomly sampled Books3 sequences 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. Error bars here indicate variance across the 3 samples of 40,000 sequences that we run.

sequences for which $p_z < \tau_{\min} = 0.1\%$). For sequences with these low $p_z < \tau_{\min}$, it may also capture generalization. We include these numbers as a reference, not to make extraction claims.

C.3 High-level takeaways

We defer discussion of validity to Appendix E, where we talk about how to draw conclusions for different models—some for which we know with certainty that Book3 is in their training data, and others for which we can conclude that at least some content that overlaps with Books3 was in their training data. PHI 4 reflects a separate case: it as trained predominantly on synthetic data. Based on the PHI 4 technical report [2], PHI 4 was deliberately not trained on copyrighted (whole) books, but was trained on unspecified public domain material. As a result, Books3 (the specific corpus) was almost certainly not included in its training data. Nevertheless, we are sometimes able to extract data that is contained in Books3, both from in-copyright and public domain books. Sequences from in-copyright texts likely found their way into the training data from web-scraped sources that quote these books. We discuss this further in Section 5, Appendix E, and Appendix G.

Figure 13 shows average extraction rates grouped roughly by model size. As we expected (and discuss in Section 4.2), *average* extraction rates are relatively low. Using the conservative threshold of $p_z \geq \tau_{\min} = 0.1\%$, the lowest rate we observe is for PHI 4 ($< 0.2\%$), and the highest is for LLAMA 3.1 70B ($\sim 0.6\%$). Consistent with prior work, our results show that larger models tend to memorize more than smaller ones [41, 123, 192, 193]. This holds within the same model family. PHI 4 is an

exception; despite being a 14B model, its extraction rates on Books3 sequences are comparable to the smaller (7–9B) models we evaluate. Small LLAMA 3 models memorize on-par with medium LLAMA 1 and 2 models. In general, 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 2 70B, which memorizes more on average than LLAMA 1 65B. This pattern does not always hold for specific books. We defer additional, related observations about LLAMA 3.1 70B to Appendix G.

Altogether, it is important to keep in mind that any given sequence that we deem extracted is, by definition, a member of the training data; however, for LLMs with unknown training data, it is not necessarily a member of the training data *because* it came from the Books3 corpus. The sequence may be part of another source (or multiple other sources) in the training data. For PYTHIA and LLAMA models, the Books3 copy of a sequence may just be one of several duplicate copies in the whole training dataset. The same can be said for models like GEMMA 2, where we are not sure if the Books3 corpus was included in the training dataset.) Nevertheless, it is clear from our results that sequences that are contained in Books3—but perhaps originate from other sources—are included in the training data (Section 5 & Appendix E).

D Extended results for sliding-window experiments

We discuss details for our main experiments that use a sliding window across a book to surface memorization “hot-spots” at particular book locations. We discuss the overall setup (Appendix D.1), our sliding-window procedure (Appendix D.2), detailed results for 200 books and 14 LLMs (Appendix D.3), and our extraction coverage metric (Appendix D.4). We describe other results—beyond our main experiments with 50-token prefixes and 50-token suffixes—in different appendices. We defer discussion of sliding-window validity experiments on non-training data and different prompt prefix lengths to Appendix E. Additional experiments concerning extraction rates and our long-form extraction procedure are discussed in Appendix C and Appendix F, respectively. A refactored and optimized version of the code for our main sliding-window experiments is [here](#).

D.1 Setup

Data and book selection. We draw our prompts (prefixes) from a selection of books in Books3; we verify extraction against the corresponding suffixes drawn from Books3. Books3 is a common pretraining dataset that is known to have been used to train several open-weight LLMs (Appendix C.1, [159, 219]. (See models, below.) We list the 200 selected books in Table 3.

We began our project with *The Great Gatsby* [92]. This choice stems from a conversation in 2023 while writing a different paper [159], when Anthropic announced its 100,000-token context for Claude with the example of prompting with the entirety of *The Great Gatsby* [8]. 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, Inc.* [143, pp. 4-5]. This is how we sourced the first 13 in-copyright books from Books3. We expanded to additional books by these plaintiffs, and also added in some generally popular books that we chose among the project team (e.g., Camus [35], Heller [124], Rowling [224], Tolkien [266]) and some (less publicly popular) academic books based on our personal preferences (e.g., Barolini [19], Zittrain [292]). (This team, after all, is composed of academics.) We deliberately included selections from the public domain (Carroll [42], Fitzgerald [92], Joyce [141]), and an example of a book that was published under a permissive CC license (Doctorow [79]). We added in Carroll [42] because of concurrent work that was published on arXiv just as we posted our own [170]. To round out the list, we also selected books at random from the manifest file that accompanies the Books3 archive (e.g., Jacobs [138], Rouighi [223], Cogburn and Silcox [53]). Once we observed that certain very popular books were highly memorized by LLAMA 3.1 70B, we added additional bestsellers to our set (e.g., Martin [174], Brown [31], Meyer [177], Orwell [203], Sandberg [237]). In earlier versions of this work, we reported on 36 and then 50 books in total.

In the current version, we manually selected additional books (non-randomly) and randomly sampled additional books, such that the final split contains 100 chosen books and 100 that were randomly sampled. Given that there are nearly 200,000 books in Books3, our experiments cover roughly 0.1% of the corpus. This is not nearly a large enough sample to make general claims about overall books extraction for the setup we study. This is why we do not completely randomly select books; we pick an assortment of manually and randomly selected books, intended to capture variation in the corpus.

Books metadata curation. In earlier versions of this project, which involved a smaller number of books, we manually gathered books metadata and manually created bibtex entries. When scaling up to 200 books, we attempted to automate as much as possible. For each book we evaluate in Books3, we used GPT-5 to extract bibliographic metadata, including the title and author(s), as well as the ISBN(s) and publisher(s) corresponding to the specific edition. To do so, we provided GPT-5 with the first and last several thousand words of each text, reflecting our observation that the desired metadata tended to concentrate in the front and back matter. Using these fields, we then queried GEMINI 2.5 FLASH with Google Search to determine each book’s current copyright status and initial year of publication. Metadata quality was assessed through manually spot-checking a random 10% subset of the books. We found some errors in how authors’ names were parsed, and so manually validated the author field for all 200 books. Refer to Table 3 for the gathered books metadata.

Models. Altogether, we run this set of experiments on 14 completion-style (i.e., non-chatbot) models: PYTHIA 12B [25], LLAMA 1 13B, LLAMA 1 65B [269], LLAMA 2 7B, LLAMA 2 13B, LLAMA 2 70B [270], LLAMA 3 8B, LLAMA 3 70B, LLAMA 3.1 8B, LLAMA 3.1 70B [111], DEEPSEEK v1 67B [76], QWEN 2.5 72B [217], GEMMA 2 27B [260], and PHI 4 (a 14B model) [2]. In earlier versions of this work, we also ran experiments on 3 other (smaller) models, which we omit for brevity: PYTHIA 6.7B, LLAMA 1 7B, and DEEPSEEK v1 7B.

We load all models with float16—except for LLAMA 3, LLAMA 3.1, and QWEN 2.5 models. These models were explicitly trained to work with bfloat16. In the first version of this work, we reported results for LLAMA 3.1 that used float16. Slight differences in results in this version are attributed to the change to bfloat16. (We have otherwise confirmed our results are deterministic, with respect to the hardware that we use throughout.) Regardless, the overall results are similar, and the conclusions remain the same.

We picked these models because they each fit (with respect to 16-bit format) on 4 A100s, and because they fall into three different categories that we wanted to investigate: (1) models that we know with certainty were trained on Books3, (2) models where we have reason to believe were trained on copyrighted books (whether drawn directly from Books3 or from other sources), and (3) a model that we trust was not trained on whole copyrighted books.

For the first category, when we generate verbatim completions for these models on Books3 data, we can very reasonably conclude that we have extracted data that these models have memorized during training. We know that these data were members of the training data; extraction success (and thus memorization) is a standard conclusion in this setting [e.g., 25, 41, 123, 213]. For the second category, we aim to do valid membership inference of copyrighted data. The third category is included for all books, but is intended to help confirm validity; as a baseline, we intentionally include an LLM for which we do not expect to be able to extract much (if any) copyrighted books data.

1. **Category 1: PYTHIA 12B, LLAMA 1 13B, LLAMA 1 65B, LLAMA 2 7B, LLAMA 2 13B, LLAMA 2 70B, LLAMA 3 8B, LLAMA 3 8B, LLAMA 3 70B, and LLAMA 3.1 70B.** We know that PYTHIA and LLAMA models were trained on Books3. PYTHIA was trained on the Pile [97] (which contains Books3), based on EleutherAI’s extensive documentation. We include only the largest (12-billion parameter) PYTHIA model.

LLAMA 1 and LLAMA 2 are known with absolute certainty to have been trained on Books3. The LLAMA 1 report explicitly says so. It came out in discovery during *Kadrey v. Meta Platforms, Inc.* [144] that LLAMA 2 models were all also definitively trained on Books3.

Both in documents in discovery and Meta’s otherwise public documentation, it is repeatedly noted that the development of LLAMA 3 continued to use the same data practices as prior versions. For minor-versioned LLAMA 3 models (like LLAMA 3.1), it is similarly documented by Meta in release reports that the same internal workflows were used, including for dataset curation [111]. This is generally considered definitive evidence that these models were also trained on Books3 (as well as additional books data from LibGen). (Our validity experiments enable us to independently conclude this. See Section 5 and Appendix E.)

2. **Category 2: DEEPSEEK v1 67B, QWEN 2.5 72B, and GEMMA 2 27B.** Although we do not know with certainty whether these models were trained on Books3, we aim to use our extraction methodology to assess whether they have memorized content from books. In particular, we want to design a measurement procedure that lets us conclude that reproducing such a sequence in an LLM’s output is valid extraction (i.e., evidence of memorization). Based on our validity

experiments (Section 5 and Appendix E), we can be very confident that when we successfully generate Books3 data, that those data were memorized from the training data. We discuss those experiments elsewhere, and note here that when we observe high probabilities p_z for these models, we can safely conclude that those sequences z have been memorized.

3. **Category 3: PHI 4.** For validity reasons, we also want to compare on a model that definitively should *not* have been trained on copyrighted books in Books3. (See also Appendix E.) We pick PHI 4 for these reasons, as it is a very high quality for its size class and Abdin et al. [2] note in the PHI 4 report that it was trained predominantly on synthetic data. While we do not know which books are included in PHI 4, the report notes that it was not trained on whole copyrighted books. Public domain data may have been included. This makes PHI 4 and imperfect negative control, but we include it anyway as a basis for comparison. PHI 4 is similar in size to PYTHIA 12B, LLAMA 1 13B, and LLAMA 2 13B, which facilitates cross-model comparisons of extraction risk. (This is because the amount of memorization, reasonably, also varies according to model size; see also Carlini et al. [41].)

Compute resources. We run all of our sliding-window experiments in a Slurm cluster environment, using the same node with 4 A100 GPUs. All told, this set of experiments took 394 GPU days.

Table 3: **Books drawn from Books3.** The 200 books from Books3 that we evaluate with our sliding-window probabilistic extraction procedure. See Sections 4 and 6, as well as Appendices D.1 & D.2.

	Author	Title	Year	Status	Random	Books3 path
1	Achebe, Chinua	<i>Things Fall Apart</i> [3]	1958	©	False	https://the-eye.eu/public/Books/Bibliotik/C/Chinua%20Achebe%20-%20Things%20Fall%20Apart.epub
2	Adams, Douglas	<i>The Hitchhiker's Guide to the Galaxy - Omnibus</i> [4]	1979	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Hitchhiker%27s%20Guide%20to%20the%20Galaxy%20Omnibus%20A%20Trilogy%20in%20Five%20Parts%20-%20Douglas%20Adams.epub
3	Adichie, Chimamanda Ngozi	<i>Americanah</i> [5]	2013	©	False	https://the-eye.eu/public/Books/Bibliotik/A/Americanah-Adichie.epub
4	Al Rawi, Shahad	<i>The Baghdad Clock</i> [6]	2016	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Baghdad%20Clock%20-%20Shahad%20Al%20Rawi.epub
5	Armstrong, Kelley	<i>Industrial Magic</i> [9]	2004	©	True	https://the-eye.eu/public/Books/Bibliotik/I/Industrial%20Magic%20-%20Kelley%20Armstrong.epub
6	Asimov, Isaac	<i>Fantastic Voyage</i> [10]	1966	©	False	https://the-eye.eu/public/Books/Bibliotik/I/IsaacAsimov-FantasticVoyage_9780307790965.epub
7	Asimov, Isaac	<i>The Complete Robot</i> [11]	1982	©	False	https://the-eye.eu/public/Books/Bibliotik/I/Isaac%20Asimov%20-%20The%20Complete%20Robot.epub
8	Atwood, Margaret	<i>The Handmaid's Tale</i> [12]	1985	©	False	https://the-eye.eu/public/Books/Bibliotik/M/Margaret%20Atwood%20%281985%29%20The%20Handmaid%27s%20Tale%20%5Bretail%5D.epub
9	Austen, Jane	<i>Pride and Prejudice</i> [13]	1813	PD	False	https://the-eye.eu/public/Books/Bibliotik/P/Pride%20and%20Prejudice%20-%20Jane%20Austen.epub
10	Baldacci, David	<i>The Christmas Train</i> [15]	2001	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Christmas%20Train%20-%20David%20Baldacci.epub
11	Baldwin, James	<i>Notes of a Native Son</i> [16]	1955	©	False	https://the-eye.eu/public/Books/Bibliotik/J/James%20Baldwin%20-%20Notes%20of%20a%20Native%20Son.epub
12	Baldwin, James	<i>Another Country</i> [17]	1962	©	False	https://the-eye.eu/public/Books/Bibliotik/A/Another%20Country%20-%20James%20Baldwin.epub
13	Barnes, Julian	<i>The Lemon Table</i> [18]	2004	©	True	https://the-eye.eu/public/Books/Bibliotik/J/Julian%20Barnes%20-%20The%20Lemon%20Table.epub

Continued on next page

Table 3 (continued). The 200 books drawn from Books3.

	Author	Title	Year	Status	Random	Books3 path
14	Barolini, Teodolinda	<i>Dante and the Origins of Italian Literary Culture</i> [19]	2006	©	False	https://the-eye.eu/public/Books/Bibliotik/D/Dante%20and%20the%20Origins%20of%20Italia%20-%20Barolini%2C%20Teodolinda%3B.epub
15	Beard, Mary	<i>The Parthenon</i> [20]	2002	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Parthenon%20-%20Mary%20Beard.epub
16	Beardsley, Charles	<i>Guam: Past and Present</i> [21]	1964	©	True	https://the-eye.eu/public/Books/Bibliotik/G/Guam%20Past%20and%20Present%20-%20Charles%20Beardsley.epub
17	Beckett, Samuel	<i>Waiting for Godot</i> [22]	1952	©	False	https://the-eye.eu/public/Books/Bibliotik/W/Waiting%20for%20Godot%20-%20Samuel%20Beckett.epub
18	Benedict, Helen	<i>The Lonely Soldier</i> [23]	2009	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Lonely%20Soldier%202011%20-%20Helen%20Benedict.epub
19	Berry, Mary	<i>Simple Cakes</i> [24]	2006	©	True	https://the-eye.eu/public/Books/Bibliotik/S/Simple%20Cakes%20-%20Mary%20Berry.epub
20	Box, C.J.	<i>Paradise Valley</i> [27]	2017	©	True	https://the-eye.eu/public/Books/Bibliotik/P/Paradise%20Valley%20-%20C.J.%20Box.epub
21	Bradbury, Ray	<i>The Cat's Pajamas</i> [28]	2004	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Cat%27s%20Pajamas%20-%20Ray%20Bradbury.epub
22	Bradshaw, Gillian	<i>London in Chains</i> [29]	2009	©	True	https://the-eye.eu/public/Books/Bibliotik/L/London%20in%20Chains%20-%20Gillian%20Bradshaw.epub
23	Brockman, John	<i>My Einstein</i> [30]	2006	©	True	https://the-eye.eu/public/Books/Bibliotik/M/My%20Einstein%20-%20John%20Brockman.epub
24	Brown, Dan	<i>The Da Vinci Code</i> [31]	2003	©	False	https://the-eye.eu/public/Books/Bibliotik/D/Dan%20Brown%20-%20The%20Da%20Vinci%20Code.epub
25	Bryant, Niobia	<i>Live and Learn</i> [33]	2007	©	True	https://the-eye.eu/public/Books/Bibliotik/L/Live%20and%20Learn%20-%20Niobia%20Bryant.epub
26	Brzezinski, Mika	<i>Knowing Your Value</i> [34]	2011	©	True	https://the-eye.eu/public/Books/Bibliotik/K/Knowing%20Your%20Value%20-%20Mika%20Brzezinski.epub
27	Camus, Albert	<i>The Myth of Sisyphus</i> [35]	1942	©	False	https://the-eye.eu/public/Books/Bibliotik/2/2013%28orig1942%29%20Albert%20Camus%20-%20The%20Myth%20of%20Sisyphus%5BTransl%20Justin%200%27Brien%5D_Ral.epub
28	Carroll, Lewis	<i>Alice's Adventures in Wonderland and Through the Looking-Glass and What Alice Found There</i> [42]	1871	PD	False	https://the-eye.eu/public/Books/Bibliotik/A/Alice%27s%20Adventures%20in%20Wonderlan%20-%20Lewis%20Carroll.epub
29	Carver, Jeffrey A.	<i>The Infinity Link</i> [43]	1984	©	True	https://the-eye.eu/public/Books/Bibliotik/I/Infinity%20Link%20-%20Jeffrey%20A.%20Carver.epub
30	Christie, Agatha	<i>Murder on the Orient Express</i> [46]	1933	©	False	https://the-eye.eu/public/Books/Bibliotik/M/Murder%20on%20the%20Orient%20Express-Christie.epub
31	Christie, Agatha	<i>And Then There Were None</i> [47]	1939	©	True	https://the-eye.eu/public/Books/Bibliotik/A/And%20Then%20There%20Were%20None%20-%20Agatha%20Christie.epub
32	Coates, Ta-Nehisi	<i>The Beautiful Struggle</i> [48]	2008	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Beautiful%20Struggle%20-%20Ta-Nehisi%20Coates.epub
33	Coates, Ta-Nehisi	<i>We Were Eight Years in Power</i> [49]	2017	©	False	https://the-eye.eu/public/Books/Bibliotik/W/We%20Were%20Eight%20Years%20in%20Power%20-%20Ta-Nehisi%20Coates.epub
34	Coates, Ta-Nehisi	<i>The Water Dancer</i> [50]	2019	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Water%20Dancer%20-%20Ta-Nehisi%20Coates%3B.epub

Continued on next page

Table 3 (continued). The 200 books drawn from Books3.

Author	Title	Year	Status	Random	Books3 path
35 Cocteau, Jean	<i>The Infernal Machine and Other Plays</i> [51]	1963	©	True	https://the-eye.eu/public/Books/Bibliotik/I/Infernal%20Machine%2C%20The%20-%20Jean%20Cocteau.epub
36 Coelho, Paulo	<i>The Alchemist</i> [52]	1988	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Alchemist%20%2825th%20Anniversary%20Edition%29%20-%20Paulo%20Coelho.epub
37 Cogburn, Jon and Silcox, Mark	<i>Dungeons and Dragons and Philosophy</i> [53]	2012	©	True	https://the-eye.eu/public/Books/Bibliotik/D/Dungeons%20and%20Dragons%20and%20Philosophy%20-%20Christopher%20Robichaud.epub
38 Cohen-Solal, Annie	<i>Mark Rothko</i> [54]	2013	©	True	https://the-eye.eu/public/Books/Bibliotik/M/Mark%20Rothko%20-%20Annie%20Cohen-Solal.epub
39 Collins, Suzanne	<i>The Hunger Games</i> [56]	2008	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Hunger%20Games%20-%20Suzanne%20Collins.epub
40 Cook, Glen	<i>The Dragon Never Sleeps</i> [58]	1988	©	True	https://the-eye.eu/public/Books/Bibliotik/G/Glen%20Cook%20-%20The%20Dragon%20Never%20Sleeps%20%5Bretail%5D.epub
41 Covey, Stephen R.	<i>The 7 Habits of Highly Effective People</i> [66]	1989	©	False	https://the-eye.eu/public/Books/Bibliotik/S/Stephen%20Covey%20-%20The%207%20Habits%20of%20Highly%20Effective%20People.epub
42 Crabb, David	<i>Bad Kid</i> [68]	2015	©	True	https://the-eye.eu/public/Books/Bibliotik/B/Bad%20Kid%20-%20David%20Crabb.epub
43 Crais, Robert	<i>Lullaby Town</i> [69]	1992	©	True	https://the-eye.eu/public/Books/Bibliotik/L/Lullaby%20Town%20-%20Robert%20Crais.epub
44 Crichton, Michael	<i>Jurassic Park</i> [70]	1990	©	False	https://the-eye.eu/public/Books/Bibliotik/J/Jurassic_Park-Crichton.epub
45 Cunningham, Michael	<i>The Hours</i> [71]	1998	©	True	https://the-eye.eu/public/Books/Bibliotik/M/Michael%20Cunningham%20-%20The%20Hours.epub
46 Czapski, Józef	<i>Inhuman Land</i> [73]	1949	©	True	https://the-eye.eu/public/Books/Bibliotik/I/Inhuman%20Land%20-%20Jozef%20Czapski.epub
47 Dahl, Roald	<i>Charlie and the Chocolate Factory</i> [75]	1964	©	False	https://the-eye.eu/public/Books/Bibliotik/C/Charlie%20and%20the%20Chocolate%20Factory%20-%20Roald%20Dahl.epub
48 Dahl, Roald	<i>James and the Giant Peach</i> [74]	1961	©	True	https://the-eye.eu/public/Books/Bibliotik/R/Roald%20Dahl%20-%20James%20and%20the%20Giant%20Peach-%20A%20Play%20%28retail%29.epub
49 Diakopoulos, Nicholas	<i>Automating the News</i> [77]	2019	©	True	https://the-eye.eu/public/Books/Bibliotik/N/Nicholas%20Diakopoulos%20-%20Automating%20the%20News.epub
50 Díaz, Junot	<i>Drown</i> [81]	1995	©	False	https://the-eye.eu/public/Books/Bibliotik/J/Junot%20Diaz%20-%20Drown.epub
51 Díaz, Junot	<i>The Brief Wondrous Life of Oscar Wao</i> [82]	2007	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Brief%20Wondrous%20Life%20of%20Oscar%20Wao.epub
52 Díaz, Junot	<i>This Is How You Lose Her</i> [83]	2012	©	False	https://the-eye.eu/public/Books/Bibliotik/T/This%20Is%20How%20You%20Lose%20Her-Diaz.epub
53 Didion, Joan	<i>The White Album</i> [78]	1979	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20White%20Album%20%28Zola%20Books%29%20-%20Joan%20Didion%20%28retail%29.epub
54 Doctorow, Cory	<i>Down and Out in the Magic Kingdom</i> [79]	2003	CC-BY-SA	False	https://the-eye.eu/public/Books/Bibliotik/D/Down%20and%20Out%20in%20the%20Magic%20Kingd%20-%20Cory%20Doctorow.epub
55 Duffy, Carol Ann	<i>The World's Wife</i> [80]	1999	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20World%27s%20Wife%20-%20Carol%20Ann%20Duffy.epub

Continued on next page

Table 3 (continued). The 200 books drawn from Books3.

Author	Title	Year	Status	Random	Books3 path
56 Egan, Jennifer	<i>A Visit from the Goon Squad</i> [84]	2010	©	False	https://the-eye.eu/public/Books/Bibliotik/A/A%20Visit%20from%20the%20Goon%20Squad%20-%20Jennifer%20Egan.epub
57 Ellison, Ralph	<i>Invisible Man</i> [86]	1952	©	False	https://the-eye.eu/public/Books/Bibliotik/I/Invisible%20Man%20-%20Ralph%20Ellison.epub
58 Fairless, Chelsea and Garroni, Lauren	<i>We Should All Be Mirandas</i> [87]	2019	©	True	https://the-eye.eu/public/Books/Bibliotik/W/We%20Should%20All%20Be%20Mirandas%20-%20Chelsea%20Fairless.epub
59 Falsani, Cathleen	<i>The Dude Abides</i> [88]	2009	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Dude%20Abides%20-%20Cathleen%20Falsani.epub
60 Farnsworth, Christopher	<i>The President's Vampire</i> [89]	2011	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20President%27s%20Vampire%20-%20Christopher%20Farnsworth.epub
61 Fitzgerald, F. Scott	<i>The Great Gatsby</i> [92]	1925	PD	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Great%20Gatsby%20-%20F.%20Scott%20Fitzgerald.epub
62 Flynn, Gillian	<i>Gone Girl</i> [93]	2012	©	False	https://the-eye.eu/public/Books/Bibliotik/G/Gone%20Girl%20-%20Gillian%20Flynn.epub
63 Friedman, Norman	<i>British Destroyers</i> [94]	2009	©	True	https://the-eye.eu/public/Books/Bibliotik/B/British%20Destroyers%20-%20Norman%20Friedman.epub
64 Fudurich, Barbara and Hansen, Joan and Caroline Linscott and Medway, Geri	<i>Florals & Botanicals</i> [95]	2008	©	True	https://the-eye.eu/public/Books/Bibliotik/F/Florals%20%26%20Botanicals%20-%20Barbara%20Fudurich.epub
65 Gaiman, Neil and Pratchett, Terry	<i>Good Omens</i> [96]	1990	©	False	https://the-eye.eu/public/Books/Bibliotik/T/Terry%20Pratchett%2C%20Neil%20Gaiman%20-%20Good%20Omens%20%5BUK%5D%20%5Bretail%5D.epub
66 Gideon, Melanie	<i>The Slippery Year</i> [100]	2009	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Slippery%20Year%20-%20Melanie%20Gideon.epub
67 Girard, Stefanie	<i>Sweater Surgery</i> [101]	2010	©	True	https://the-eye.eu/public/Books/Bibliotik/S/Sweater%20Surgery%20-%20Stefanie%20Girard.epub
68 Gladwell, Malcolm	<i>Blink</i> [102]	2005	©	False	https://the-eye.eu/public/Books/Bibliotik/B/Blink%20-%20Malcolm%20Gladwell.epub
69 Glover, Richard	<i>The Land Before Avocado</i> [103]	2018	©	True	https://the-eye.eu/public/Books/Bibliotik/R/Richard%20Glover%20-%20The%20Land%20Before%20Avocado.epub
70 Golden, Christopher	<i>Dead Ringers</i> [104]	2015	©	False	https://the-eye.eu/public/Books/Bibliotik/C/Christopher%20Golden%20-%20Dead%20Ringers%20%5Bretail%5D.epub
71 Golden, Christopher	<i>Ararat</i> [105]	2017	©	False	https://the-eye.eu/public/Books/Bibliotik/A/Ararat%20-%20Christopher%20Golden.epub
72 Golding, William	<i>Lord of the Flies</i> [106]	1954	©	False	https://the-eye.eu/public/Books/Bibliotik/L/Lord%20of%20the%20Flies%20-%20William%20Golding.epub
73 Goodkind, Terry	<i>Wizard's First Rule</i> [107]	1994	©	True	https://the-eye.eu/public/Books/Bibliotik/W/Wizard%27s%20First%20Rule%200101%20-%20Terry%20Goodkind.epub
74 Goodman, Martin	<i>Rome and Jerusalem</i> [108]	2007	©	True	https://the-eye.eu/public/Books/Bibliotik/R/Rome%20and%20Jerusalem%20-%20Martin%20Goodman.epub
75 Green, John	<i>The Fault in Our Stars</i> [112]	2012	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Fault%20in%20Our%20Stars%20-%20John%20Green.epub
76 Greene, Graham	<i>The Third Man</i> [113]	1950	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Third%20Man%20-%20Graham%20Greene.epub

Continued on next page

Table 3 (continued). The 200 books drawn from Books3.

Author	Title	Year	Status	Random	Books3 path
77 Greer, Andrew Sean	<i>The Confessions of Max Tivoli</i> [114]	2004	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Confessions%20of%20Max%20Tivoli%20-%20Andrew%20Sean%20Greer.epub
78 Grisham, John	<i>The Fugitive</i> [115]	2015	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Fugitive%20-%20John%20Grisham.epub
79 Haddon, Mark	<i>The Curious Incident of the Dog in the Night-Time</i> [116]	2003	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Curious%20Incident%20of%20the%20Dog%20in%20the%20Night-Time%20-%20Mark%20Haddon.epub
80 Halpern, Sue	<i>Migrations to Solitude</i> [117]	1992	©	True	https://the-eye.eu/public/Books/Bibliotik/S/Sue%20Halpern%20-%20Migrations%20to%20Solitude.epub
81 Hanks, Tom	<i>Uncommon Type</i> [118]	2017	©	True	https://the-eye.eu/public/Books/Bibliotik/U/Uncommon%20Type%20-%20Tom%20Hanks.epub
82 Hanson, Thor	<i>Buzz</i> [119]	2018	©	True	https://the-eye.eu/public/Books/Bibliotik/B/Buzz%20-%20Thor%20Hanson.epub
83 Haydon, Elizabeth	<i>Requiem for the Sun</i> [121]	2002	©	True	https://the-eye.eu/public/Books/Bibliotik/R/Requiem%20for%20the%20Sun%20-%20Elizabeth%20Haydon.epub
84 Heller, Joseph	<i>Catch-22</i> [124]	1961	©	False	https://the-eye.eu/public/Books/Bibliotik/C/Catch-22%20-%20Joseph%20Heller.epub
85 Hemingway, Ernest	<i>The Old Man and the Sea</i> [125]	1952	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Old%20Man%20and%20the%20Sea%20-%20Ernest%20Hemingway.epub
86 Hendy, David	<i>Life on Air</i> [127]	2007	©	True	https://the-eye.eu/public/Books/Bibliotik/L/Life%20on%20Air%20-%20Hendy%20-%20David%3B.epub
87 Hersheson, Luke	<i>Great Hair Days</i> [129]	2018	©	True	https://the-eye.eu/public/Books/Bibliotik/G/Great%20Hair%20Days%20-%20Luke%20Hersheson.epub
88 Hill, Matt	<i>Graft</i> [130]	2016	©	True	https://the-eye.eu/public/Books/Bibliotik/G/Graft%20-%20Matt%20Hill.epub
89 Hinton, S. E.	<i>The Outsiders</i> [131]	1967	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Outsiders%20-%20SE%20Hinton.epub
90 Huff, Tanya	<i>The Second Summoning</i> [132]	2001	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Second%20Summoning%20-%20Tanya%20Huff.epub
91 Hughes, Langston	<i>Selected Poems of Langston Hughes</i> [134]	1959	©	False	https://the-eye.eu/public/Books/Bibliotik/S/Selected%20Poems%20of%20Langston%20Hugh%20-%20Langston%20Hughes.epub
92 Hwang, David Henry	<i>M. Butterfly</i> [135]	1988	©	False	https://the-eye.eu/public/Books/Bibliotik/M/M.%20Butterfly%20-%20David%20Henry%20Hwang.epub
93 Jackson, Allen	<i>Building and Operating a Realistic Model Railway</i> [137]	2016	©	True	https://the-eye.eu/public/Books/Bibliotik/B/Building%20and%20Operating%20a%20Realistic%20Model%20Railway%20-%20Allen%20Jackson.epub
94 Jacobs, Betty	<i>All the Onions</i> [138]	1993	©	True	https://the-eye.eu/public/Books/Bibliotik/A/All%20the%20Onions%20-%20Betty%20E.%20M.%20Jacobs.epub
95 James, E. L.	<i>Fifty Shades of Grey</i> [139]	2011	©	False	https://the-eye.eu/public/Books/Bibliotik/F/Fifty%20Shades%20Darker%20-%20E%20L%20James.epub
96 Jemisin, N. K.	<i>The Stone Sky</i> [140]	2017	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Stone%20Sky%20-%20N.%20K.%20Jemisin.epub
97 Joyce, James	<i>Ulysses</i> [141]	1922	PD	False	https://the-eye.eu/public/Books/Bibliotik/U/Ulysses%20-%20James%20Joyce%20-%20Penguin%20Group%20-%202000.epub
98 Kadrey, Richard	<i>Sandman Slim</i> [142]	2009	©	False	https://the-eye.eu/public/Books/Bibliotik/S/Sandman%20Slim%20-%20Richard%20Kadrey.epub

Continued on next page

Table 3 (continued). The 200 books drawn from Books3.

Author	Title	Year	Status	Random	Books3 path
99 Kaldellis, Anthony	<i>Ethnography after Antiquity</i> [145]	2013	©	True	https://the-eye.eu/public/Books/Bibliotik/E/Ethnography%20After%20Antiquity%20-%20Anthony%20Kaldellis.epub
100 Klam, Matthew	<i>Who Is Rich?</i> [147]	2017	©	False	https://the-eye.eu/public/Books/Bibliotik/W/Who%20Is%20Rich%20-%20Matthew%20Klam.epub
101 Koontz, Dean	<i>The Servants of Twilight</i> [148]	1984	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Servants%20of%20Twilight%20-%20Dean%20Koontz.epub
102 Kuhn, Aihan	<i>Tai Chi for Depression</i> [149]	2017	©	True	https://the-eye.eu/public/Books/Bibliotik/T/Tai%20Chi%20for%20Depression%20-%20Aihan%20Kuhn.epub
103 Kun, Josh	<i>The Tide Was Always High</i> [150]	2017	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Tide%20Was%20Always%20High%20-%20Josh%20Kun.epub
104 Kwok, Jean	<i>Girl in Translation</i> [151]	2010	©	True	https://the-eye.eu/public/Books/Bibliotik/J/JeanKwok-GirlinTranslation_9781101182673.epub
105 L'Engle, Madeleine	<i>A Wrinkle in Time</i> [163]	1962	©	False	https://the-eye.eu/public/Books/Bibliotik/A/A%20Wrinkle%20in%20Time%20-%20Madeleine%20L%27engle.epub
106 Lago, Eduardo	<i>Call Me Brooklyn</i> [152]	2006	©	True	https://the-eye.eu/public/Books/Bibliotik/C/Call%20Me%20Brooklyn%20-%20Eduardo%20Lago.epub
107 Larson, Erik	<i>Dead Wake</i> [153]	2015	©	True	https://the-eye.eu/public/Books/Bibliotik/D/Dead%20Wake%20-%20Erik%20Larson.epub
108 Larsson, Stieg	<i>The Girl with the Dragon Tattoo</i> [154]	2005	©	False	https://the-eye.eu/public/Books/Bibliotik/G/Girl%20with%20the%20Dragon%20Tattoo%20-%20Stieg%20Larsson.epub
109 Le Guin, Ursula K.	<i>The Daughter of Odren</i> [155]	2014	©	False	https://the-eye.eu/public/Books/Bibliotik/U/Ursula%20K.%20Le%20Guin%20-%20Earthsea%2006.5%20-%20The%20Daughter%20of%20Odren%20%205Bretail%205D.epub
110 Lewis, C. S.	<i>The Chronicles of Narnia</i> [164]	1950	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Chronicles%20of%20Narnia%20-%20C.S.%20Lewis.epub
111 Lippman, Laura	<i>After I'm Gone</i> [165]	2014	©	False	https://the-eye.eu/public/Books/Bibliotik/A/After%20I%27m%20Gone%20-%20Laura%20Lippman.epub
112 Lippman, Laura	<i>Sunburn</i> [166]	2018	©	False	https://the-eye.eu/public/Books/Bibliotik/S/Sunburn%20-%20Laura%20Lippman.epub
113 Liss, David	<i>Marvel's Spider-Man: Hostile Takeover</i> [167]	2018	©	True	https://the-eye.eu/public/Books/Bibliotik/M/Marvel%27s%20Spider-Man_%20Hostile%20Ta%20-%20David%20Liss.epub
114 Lowry, Lois	<i>Anastasia on Her Own</i> [169]	1985	©	False	https://the-eye.eu/public/Books/Bibliotik/A/Anastasia%20on%20Her%20Own%20-%20Lois%20Lowry.epub
115 MacMillan, Margaret	<i>Nixon in China</i> [171]	2006	©	True	https://the-eye.eu/public/Books/Bibliotik/N/Nixon%20in%20China%20-%20Margaret%20Macmillan.epub
116 Mariani, Scott	<i>The Doomsday Prophecy</i> [173]	2009	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Doomsday%20Prophecy%20-%20Scott%20Mariani.epub
117 Martin, George R. R.	<i>A Game of Thrones</i> [174]	1996	©	False	https://the-eye.eu/public/Books/Bibliotik/A/A%20Game%20of%20Thrones%20-%20George%20R.%20R.%20Martin.epub
118 Meldahl, Keith Heyer	<i>Rough-Hewn Land</i> [175]	2011	©	True	https://the-eye.eu/public/Books/Bibliotik/R/Rough-Hewn%20Land.epub
119 Meyer, Stephenie	<i>Twilight</i> [177]	2005	©	False	https://the-eye.eu/public/Books/Bibliotik/T/Twilight%20-%20Stephenie%20Meyer.epub
120 Milan, Courtney	<i>The Duchess War</i> [178]	2012	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Duchess%20War%20-%20Courtney%20Milan.epub

Continued on next page

Table 3 (continued). The 200 books drawn from Books3.

Author	Title	Year	Status	Random	Books3 path
121 Mill, John Stuart	<i>On Liberty</i> [179]	1859	PD	False	https://the-eye.eu/public/Books/Bibliotik/0/0n%20Liberty%20-%20John%20Stuart%20Mill.epub
122 Miller, Alex	<i>Coal Creek</i> [180]	2013	©	True	https://the-eye.eu/public/Books/Bibliotik/C/Coal%20Creek%20-%20Alex%20Miller.epub
123 Milne, A. A.	<i>Winnie the Pooh</i> [181]	1926	PD	False	https://the-eye.eu/public/Books/Bibliotik/W/Winnie_the_Pooh-Milne.epub
124 Moore, Colten and O'Brien, Keith	<i>Catching the Sky</i> [183]	2016	©	True	https://the-eye.eu/public/Books/Bibliotik/C/Catching%20the%20Sky%20-%20Colten%20Moore.epub
125 Moran, Michelle	<i>The Heretic Queen</i> [184]	2008	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Heretic%20Queen%20-%20Michelle%20Moran.epub
126 Morelli, Laura	<i>The Gondola Maker</i> [185]	2014	©	True	https://the-eye.eu/public/Books/Bibliotik/G/Gondola%20Maker%2C%20The%20-%20Laura%20Morelli.epub
127 Morris, Mary McGarry	<i>Songs in Ordinary Time</i> [186]	1995	©	True	https://the-eye.eu/public/Books/Bibliotik/S/Songs%20in%20Ordinary%20Time%20-%20Mary%20McGarry%20Morris.epub
128 Morrison, Toni	<i>Beloved</i> [187]	1987	©	False	https://the-eye.eu/public/Books/Bibliotik/B/Beloved%20-%20Toni%20Morrison.epub
129 Murakami, Haruki	<i>Norwegian Wood</i> [191]	1987	©	False	https://the-eye.eu/public/Books/Bibliotik/N/Norwegian%20Wood%20%28Vintage%20Books%29%20%5BEpub%5D%20-%20Haruki%20Murakami.epub
130 Nielsen, Desiree	<i>Eat More Plants</i> [194]	2019	©	True	https://the-eye.eu/public/Books/Bibliotik/E/Eat%20More%20Plants%20-%20Desiree%20Nielsen%3B.epub
131 Northrop, Michael	<i>Polaris</i> [195]	2017	©	True	https://the-eye.eu/public/Books/Bibliotik/P/Polaris%20-%20Michael%20Northrop.epub
132 O'Donnell, James J.	<i>Pagans</i> [197]	2015	©	True	https://the-eye.eu/public/Books/Bibliotik/P/Pagans%20-%20James%20J.%200%27Donnell.epub
133 O'Sullivan, Meghan L.	<i>Windfall</i> [204]	2017	©	True	https://the-eye.eu/public/Books/Bibliotik/W/Windfall%20-%20Meghan%20L.%200%27Sullivan.epub
134 Ogawa, Yoko	<i>The Memory Police</i> [198]	1994	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Memory%20Police%20-%20Yoko%200gawa.epub
135 Oliveira, Robin	<i>Winter Sisters</i> [199]	2018	©	True	https://the-eye.eu/public/Books/Bibliotik/W/Winter%20Sisters%20-%20Robin%200liveira.epub
136 Orwell, George	<i>Nineteen Eighty-Four</i> [203]	1949	©	False	https://the-eye.eu/public/Books/Bibliotik/N/Nineteen%20Eighty-Four%20%28The%20Annotated%20Edition%29%20-%20George%20Orwell.epub
137 Palahniuk, Chuck	<i>Fight Club</i> [205]	1996	©	False	https://the-eye.eu/public/Books/Bibliotik/F/Fight%20Club%20-%20Chuck%20Palahniuk.epub
138 Papazian, Charlie	<i>The Complete Joy of Homebrewing</i> [206]	1984	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Complete%20Joy%20of%20Homebrewing%20-%20Charlie%20Papazian.epub
139 Park, Paul	<i>The Cult of Loving Kindness</i> [207]	1991	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Cult%20of%20Loving%20Kindness%20-%20Paul%20Park.epub
140 Patterson, James	<i>Along Came a Spider</i> [208]	1993	©	False	https://the-eye.eu/public/Books/Bibliotik/A/Along%20Came%20a%20Spider%20-%20James%20Patterson.epub
141 Payard, François and McBride, Anne E.	<i>Payard Cookies</i> [209]	2015	©	True	https://the-eye.eu/public/Books/Bibliotik/P/Payard%20Cookies.epub
142 Pépin, Jacques	<i>Essential Pepin Desserts</i> [216]	2011	©	True	https://the-eye.eu/public/Books/Bibliotik/E/Essential%20Pepin%20Desserts%20-%20Jacques%20Pepin.epub

Continued on next page

Table 3 (continued). The 200 books drawn from Books3.

Author	Title	Year	Status	Random	Books3 path
143 Piazza, Tom	<i>Why New Orleans Matters</i> [210]	2005	©	True	https://the-eye.eu/public/Books/Bibliotik/W/Why%20New%20Orleans%20Matters%20-%20Tom%20Piazza.epub
144 Pinker, Steven	<i>Enlightenment Now</i> [211]	2018	©	True	https://the-eye.eu/public/Books/Bibliotik/E/Enlightenment%20Now%20-%20Steven%20Pinker.epub
145 Porter, Michael E.	<i>Competitive Strategy</i> [212]	1980	©	True	https://the-eye.eu/public/Books/Bibliotik/C/Competitive%20Strategy%200101%20-%20Michael%20E.%20Porter.epub
146 Pratchett, Terry	<i>Night Watch</i> [214]	2002	©	False	https://the-eye.eu/public/Books/Bibliotik/T/Terry%20Pratchett%20-%20Discworld%2029%20-%20Night%20Watch%20%5Bretail%5D.epub
147 Pullman, Philip	<i>The Subtle Knife</i> [215]	1997	©	False	https://the-eye.eu/public/Books/Bibliotik/P/Pullman%20-%20The%20Subtle%20Knife.epub
148 Rampley, Matthew	<i>The Seductions of Darwin</i> [218]	2017	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Seductions%20of%20Darwin%20-%20Matthew%20Rampley.epub
149 Remen, Rachel Naomi	<i>Kitchen Table Wisdom</i> [220]	1996	©	True	https://the-eye.eu/public/Books/Bibliotik/K/Kitchen%20Table%20Wisdom%2010th%20Anniversary%20-%20Rachel%20Naomi%20Remen.epub
150 Roberts, Jessi and Witter, Bret	<i>Backroads Boss Lady</i> [221]	2019	©	True	https://the-eye.eu/public/Books/Bibliotik/B/Backroads%20Boss%20Lady%20-%20Jessi%20Roberts.epub
151 Roger, Marie-Sabine	<i>Soft in the Head</i> [222]	2008	©	True	https://the-eye.eu/public/Books/Bibliotik/S/Soft%20in%20the%20Head%20-%20Marie-Sabine%20Roger.epub
152 Rouighi, Ramzi	<i>The Making of a Mediterranean Emirate</i> [223]	2011	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Making%20of%20a%20Mediterranean%20E%20-%20Rouighi%2C%20Ramzi%3B.epub
153 Rowling, J. K.	<i>Harry Potter and the Sorcerer's Stone</i> [224]	1997	©	False	https://the-eye.eu/public/Books/Bibliotik/H/Harry_Potter_and_the_Sorcerers_Stone-Rowling.epub
154 Rowling, J. K.	<i>Harry Potter and the Chamber of Secrets</i> [225]	1998	©	False	https://the-eye.eu/public/Books/Bibliotik/H/Harry_Potter_and_the_Chamber_of_Secrets-Rowling.epub
155 Rowling, J. K.	<i>Harry Potter and the Goblet of Fire</i> [226]	2000	©	False	https://the-eye.eu/public/Books/Bibliotik/H/Harry_Potter_and_the_Goblet_of_Fire-Rowling.epub
156 Rowling, J. K.	<i>Harry Potter and the Deathly Hallows</i> [227]	2007	©	False	https://the-eye.eu/public/Books/Bibliotik/H/Harry_Potter_and_the_Deathly_Hallows-Rowling.epub
157 Rubinstein, Dan	<i>Born to Walk</i> [228]	2015	©	True	https://the-eye.eu/public/Books/Bibliotik/B/Born%20to%20Walk_%20The%20Transformative%20Power%20of%20a%20Pedestrian%20Act%20-%20Dan%20Rubinstein%20%28retail%29.epub
158 Ruskin, Marc	<i>The Pretender</i> [229]	2017	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Pretender%20-%20Marc%20Ruskin.epub
159 Sachs, Harvey	<i>Toscanini</i> [230]	1978	©	True	https://the-eye.eu/public/Books/Bibliotik/T/Toscanini%20-%20Harvey%20Sachs.epub
160 Sagan, Carl	<i>Cosmos</i> [233]	1980	©	False	https://the-eye.eu/public/Books/Bibliotik/C/Cosmos%20-%20Carl%20Sagan.epub
161 Sahni, Bhisham	<i>Middle India</i> [234]	2001	©	True	https://the-eye.eu/public/Books/Bibliotik/M/Middle%20India%20%28Penguin%29%20-%20Bhisam%20Sahni.epub
162 Salinger, J. D.	<i>The Catcher in the Rye</i> [235]	1951	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Catcher%20in%20the%20Rye%20-%20J.%20D.%20Salinger.epub
163 Sandberg, Sheryl	<i>Lean in</i> [237]	2013	©	False	https://the-eye.eu/public/Books/Bibliotik/L/Lean%20In%20-%20Sheryl%20Sandberg.epub

Continued on next page

Table 3 (continued). The 200 books drawn from Books3.

Author	Title	Year	Status	Random	Books3 path
164 Sharma, Sanjeev	<i>The DevOps Adoption Playbook</i> [242]	2017	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20DevOps%20Adoption%20Playbook%202017%20-%20Sharma%2C%20Sanjeev%3B.epub
165 Shelley, Mary	<i>Frankenstein</i> [243]	1818	PD	False	https://the-eye.eu/public/Books/Bibliotik/F/Frankenstein%20-%20Mary%20Shelley.epub
166 Sherr, Lynn	<i>Sally Ride: America's First Woman in Space</i> [244]	2014	©	True	https://the-eye.eu/public/Books/Bibliotik/S/Sally%20Ride_%20America%27s%20First%20Woman%20in%20Space%20-%20Lynn%20Sherr.epub
167 Shriver, Lionel	<i>A Perfectly Good Family</i> [245]	1996	©	True	https://the-eye.eu/public/Books/Bibliotik/P/Perfectly%20Good%20Family%2C%20A%20-%20Lionel%20Shriver.epub
168 Silverman, Sarah	<i>The Bedwetter</i> [246]	2010	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Bedwetter%20-%20Sarah%20Silverman.epub
169 Sloman, Larry	<i>On the Road with Bob Dylan</i> [247]	1978	©	True	https://the-eye.eu/public/Books/Bibliotik/O/On%20the%20Road%20with%20Bob%20Dylan%20-%20Larry%20Sloman.epub
170 Smith, Alexander Gordon	<i>The Night Children</i> [248]	2008	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Night%20Children%20-%20Alexander%20Gordon%20Smith.epub
171 Smith, Zadie	<i>White Teeth</i> [249]	2000	©	False	https://the-eye.eu/public/Books/Bibliotik/S/Smith%2C%20Zadie%20-%20White%20Teeth%20%28ISBN%209781400075508%29.epub
172 Snyder, Rachel Louise	<i>No Visible Bruises</i> [250]	2019	©	False	https://the-eye.eu/public/Books/Bibliotik/N/No%20Visible%20Bruises%20-%20Rachel%20Louise%20Snyder.epub
173 Steinbeck, John	<i>The Grapes of Wrath</i> [253]	1939	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Grapes%20of%20Wrath%20-%20John%20Steinbeck.epub
174 Stiles, T. J.	<i>Jesse James</i> [254]	2002	©	True	https://the-eye.eu/public/Books/Bibliotik/J/Jesse%20James%20-%20T.J.%20Stiles.epub
175 Stine, R. L.	<i>Zombie Halloween</i> [255]	2014	©	False	https://the-eye.eu/public/Books/Bibliotik/Z/Zombie%20Halloween%20-%20R.L.%20Stine.epub
176 Stubbs, David	<i>Fear of Music</i> [256]	2009	©	True	https://the-eye.eu/public/Books/Bibliotik/D/David%20Stubbs%20-%20Fear%20of%20Music.epub
177 Swanson, Jennifer	<i>Pearl Harbor</i> [257]	2018	©	True	https://the-eye.eu/public/Books/Bibliotik/P/Pearl%20Harbor%20-%20Jennifer%20Swanson.epub
178 Swift, Earl	<i>Chesapeake Requiem</i> [258]	2018	©	True	https://the-eye.eu/public/Books/Bibliotik/C/Chesapeake%20Requiem%20-%20Earl%20Swift.epub
179 Tartt, Donna	<i>The Goldfinch</i> [259]	2013	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Goldfinch-Tartt.epub
180 TerKeurst, Lysa	<i>Unglued</i> [261]	2012	©	False	https://the-eye.eu/public/Books/Bibliotik/L/Lysa%20TerKeurst%2C%20Unglued.epub
181 TerKeurst, Lysa	<i>Embraced</i> [262]	2018	©	False	https://the-eye.eu/public/Books/Bibliotik/E/Embraced%20-%20Lysa%20TerKeurst.epub
182 Thomson, Lynn	<i>Birding with Yeats</i> [265]	2014	©	True	https://the-eye.eu/public/Books/Bibliotik/B/Birding%20with%20Yeats%20-%20Lynn%20Thomson.epub
183 Tolkien, J. R. R.	<i>The Hobbit</i> [266]	1937	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Hobbit%20%28Houghton%20Mifflin%20Harcourt%29%20%2875th%20Anniversary%20Edition%29%20%5BEpub%5D%20-%20J.R.R.%20Tolkien.epub
184 Tolkien, J. R. R.	<i>The Fellowship of the Ring</i> [267]	1954	©	False	https://the-eye.eu/public/Books/Bibliotik/J/J._R._R._Tolkien_-_The_Fellowship_of_the_Ring_Being_the_First_Part_of_The_Lord_of_the_Rings_-_Houghton_Mifflin_Harcourt_-_v4.0615_-_2012_-_RETAIL.epub

Continued on next page

Table 3 (continued). The 200 books drawn from Books3.

Author	Title	Year	Status	Random	Books3 path
185 Tolkien, J. R. R.	<i>Tree and Leaf</i> [268]	1964	©	False	https://the-eye.eu/public/Books/Bibliotik/T/Tree%20and%20Leaf%20-%20J.%20R.%20Tolkien.epub
186 Trang, Corinne	<i>Noodles Every Day</i> [271]	2009	©	True	https://the-eye.eu/public/Books/Bibliotik/N/Noodles%20Every%20Day%20-%20Maura%20McEvoy.epub
187 Tyler, George R.	<i>Billionaire Democracy</i> [272]	2018	©	True	https://the-eye.eu/public/Books/Bibliotik/B/Billionaire%20Democracy%20-%20George%20R.%20Tyler.epub
188 Venter, Al J.	<i>Portugal's Guerrilla Wars in Africa</i> [274]	2013	©	True	https://the-eye.eu/public/Books/Bibliotik/P/Portugal%27s%20Guerrilla%20Wars%20in%20Africa%20-%20Al%20Venter.epub
189 Vonnegut, Kurt	<i>Slaughterhouse-Five</i> [275]	1969	©	True	https://the-eye.eu/public/Books/Bibliotik/S/Slaughterhouse-Five%20-%20Kurt%20Vonnegut%20%28retail%29.epub
190 Waldau, Paul	<i>Animal Rights</i> [278]	2011	©	True	https://the-eye.eu/public/Books/Bibliotik/A/Animal%20Rights%20-%20Waldau%2C%20Paul.epub
191 Ward, Jesmyn	<i>Men We Reaped</i> [279]	2013	©	False	https://the-eye.eu/public/Books/Bibliotik/M/Men%20We%20Reaped%20-%20Jesmyn%20Ward.epub
192 White, E. B.	<i>Charlotte's Web</i> [280]	1952	©	False	https://the-eye.eu/public/Books/Bibliotik/C/Charlotte%27s%20Web%20-%20E.%20B.%20White.epub
193 Williamson, Marianne	<i>A Return to Love</i> [282]	1992	©	False	https://the-eye.eu/public/Books/Bibliotik/M/Marianne%20Williamson%20-%20A%20Return%20to%20Love.epub
194 Woodson, Jacqueline	<i>Another Brooklyn</i> [285]	2016	©	False	https://the-eye.eu/public/Books/Bibliotik/A/Another%20Brooklyn%20-%20Jacqueline%20Woodson.epub
195 Woodson, Jacqueline	<i>Brown Girl Dreaming</i> [284]	2014	©	False	https://the-eye.eu/public/Books/Bibliotik/B/Brown%20Girl%20Dreaming%20-%20Jacqueline%20Woodson.epub
196 Yanagihara, Hanya	<i>A Little Life</i> [288]	2015	©	False	https://the-eye.eu/public/Books/Bibliotik/A/A%20Little%20Life%20-%20Hanya%20Yanagihara.epub
197 Yoshimura, Yuji and Halford, Giovanna M.	<i>The Art of Bonsai</i> [289]	1996	©	True	https://the-eye.eu/public/Books/Bibliotik/T/The%20Art%20of%20Bonsai%20-%20Yuji%20Yoshimura.epub
198 Zinn, Howard	<i>A People's History of the United States</i> [291]	1980	©	False	https://the-eye.eu/public/Books/Bibliotik/A/A_Peoples_History_of_the_United_States-Zinn.epub
199 Zittrain, Jonathan	<i>The Future of the Internet and How to Stop It</i> [292]	2008	©	False	https://the-eye.eu/public/Books/Bibliotik/J/Jonathan%20Zittrain%20-%20The%20Future%20of%20the%20Internet.epub
200 Zusak, Markus	<i>The Book Thief</i> [293]	2005	©	False	https://the-eye.eu/public/Books/Bibliotik/T/The%20Book%20Thief%20-%20Markus%20Zusak.epub

D.2 Sliding-window probabilistic extraction procedure

In this appendix, we provide more details on our main extraction experiments (Section 4). These experiments take the approach of constructing 100-token sequences z out of Books3 books, which we attempt to extract. These sequences are constructed using a sliding window across the length of an entire book, and we run these experiments for each of the 200 books in Table 3. We use a very short, 10-character shift for our sliding window. This is deliberate; it is 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 explain the intention behind these experiments, and how the associated results should be interpreted.

Sequence length. We use 100-token sequences z , with a prefix (prompt) length of 50 tokens and a (target) suffix 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 suffixes are reflective of memorization [38, 41, 123, 157]. 50 tokens for the prefix is the standard choice for 50-token suffixes, though is somewhat arbitrary. Of course, this prefix length uses an amount of context that is equivalent to the output, and so is compelling in this respect. However, when we are confident that our measurements are capturing valid instances of extraction, using longer prefixes can help *discover* more underlying memorization. (Indeed, this is part of where *discoverable* extraction gets its name; Carlini et al. [41] demonstrate the “discoverability phenomenon” with respect to showing how extraction success can increase with longer-context prompts.) We discuss this further (and include experiments with varied prefix lengths) in Section 5 and Appendix E. Other work uses shorter sequences (e.g., 64 tokens and a 32/32 token split; see Biderman et al. [25]).

Sliding-window sequence chunking. 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 800 characters, tokenize, take the first 100 tokens of the resulting tokenized sequence as the sequence, 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 800—to make sure we have sufficient head-room to always end up with a 100 token sequence.) 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 sequence every 10 characters, i.e., $\text{len}(\text{book_characters})/10$ total sequences z , for which we compute p_z . For example, *The Great Gatsby* has 270,870 characters, which results in roughly 27,000 sequences, for which we compute p_z (Equation 5). (See Section 2 and Appendix A.2.)

This means that the 100-token sequences overlap significantly. (Generally speaking, 100 tokens covers 300-400 characters, so shifting only 10 characters means there is a lot of overlap.) This is deliberate. We typically do not know how different open-weight models were trained; *a priori*, it is not clear exactly where we should begin sequences (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 a small loss in extraction signal (at $10\times$ cheaper cost than 1 character).

Decoding configuration Similar to Hayes et al. [123], after some initial tests with different decoding schemes, we use temperature $T = 1$ for top- k decoding with $k = 40$. Setting $T = 1$ makes the most sense for studying memorization, as this reflects the LLM’s base probability distribution. Since extraction is only successful when the underlying training sequence’s logits reflect very high probabilities (Section 2), it makes sense to clip the full-vocabulary logit distribution to reflect some number of the top-highest probability tokens per step. This is the same intuition behind using greedy decoding for discoverable extraction—but with a twist. Greedy decoding deterministically picks the single-most-likely token per step. (This is equivalent to top- k decoding with $k = 1$.) As the results in Hayes et al. [123] show, this can be a suboptimal decoding strategy for measuring extraction, for several reasons:

- Greedy decoding is rarely used with LLMs in practice. Non-deterministic sampling schemes are more reflective of realistic conditions. Because greedy decoding is deterministic, the same prompt always results in the same output. As a result, discoverable extraction provides (in this sense) 1-bit of information with respect to extraction: was it possible to extract the sequence? Using non-deterministic decoding schemes, we can get more detailed information about how extraction risk varies across sequences. (See Section 2.2 & Appendix A.)
- While greedy decoding picks the *locally* highest probability token per step, it is not guaranteed to result in the *globally* highest probability sequence (of the given length) under the model. Being able to explore more of the model’s distribution with non-deterministic decoding schemes can surface sequences (of the same given length) that are higher probability. Such higher probability sequences can match the target suffix, where the greedy-decoded suffix does not (i.e., fails to extract that suffix).
- More generally, this allows for a bit more “wiggle room.” By being able to pick beyond just the top-1 token per step, the decoding process take locally suboptimal steps, which can still result in high-probability sequences that (even if slightly lower probability than the greedy sequence, with respect to the normalized top- k distribution) reflect extraction. While we do not investigate

this in detail in this work, we observe that (when this happens) the greedy-decoded suffix is a near-miss for the verbatim suffix (it is near-verbatim). It may be reasonable under more relaxed (near-verbatim) metrics for extraction to have counted this greedy suffix as extracted, but that is not the approach we (or Hayes et al. [123]) take, and so these suffixes would be missed.

The above reasons also reflect the general motivations for using probabilistic extraction as our metric (Sections 2 & 3; Appendix A). Similar motivations for useful extraction measurements are also noted in early extraction papers (e.g., Carlini et al. [38], Carlini et al. [41]); however, this work typically opts to use greedy decoding based on the perception that it is cheaper to do so at scale. This work does not claim that greedy decoding is the best or only way to measure extraction, and in fact notes [38]—and in select cases shows [41]—that other (more expensive) decoding schemes are able to extract more training data.

Further, it turns out that it is actually *cheaper* to compute probabilistic extraction (with any top- k decoding configuration, including greedy/ $k = 1$), as we compute extraction probabilities as a post-processing operation directly off of the logits obtained from *inference*, rather than (as has been done in prior work [38, 41]) actually generating (greedy) completions. We can compute verbatim extraction probabilities with a single forward pass through the model, which has reduced overhead compared to greedy generation (Appendix A.3).

Note that because our implementation of probabilistic extraction makes computing log probabilities a post-processing operation on the output model logits (Appendix A), we also collect results for greedy-decoded discoverable extraction by setting $k = 1$, incurring negligible cost to do so. For brevity, we do not include detailed results for this setting. They under-count valid extraction, aligning with our findings in Section 4.2 and Appendix C, as well as Hayes et al. [123].

D.3 Book-specific results

For each of the 200 books we test, we include a separate subsection with the following information for 14 models (documented in Appendix D.1):

- **Heatmaps of extraction probabilities by location.** We include heatmaps (as in Figure 2b and Section 6). At each character, the heatmap plots the maximum extraction probability among overlapping 50-token suffixes spanning that character. These visualizations highlight “hot-spots” of high-probability extraction and show how such regions emerge and taper off within a book.
- **Histograms of extraction probabilities.** We plot the distribution over suffix extraction probabilities across all overlapping sequences from the sliding-window procedure (with a 10-character stride). These histograms illustrate how many sequences are extractable for a given model and how their extraction probabilities vary.

For histograms and heatmaps, results are presented in three groups corresponding to the categories of models described in Appendix D.1: models known with certainty to have been trained on Books3 (blue), models that were likely trained on books also present in Books3 (red), and one model (PHI 4) that, based on documentation, was not trained on whole copyrighted books (orange). See Section 5 and Appendix E for further discussion about how our results for all three categories reflect valid extraction.

We release an interactive, searchable version of these results [online](#).

D.3.1 *Things Fall Apart, Achebe*




Figure 14: *Things Fall Apart, Achebe*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.2 *The Hitchhiker's Guide to the Galaxy - Omnibus, Adams*




Figure 15: *The Hitchhiker's Guide to the Galaxy - Omnibus, Adams*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.3 Americanah, Adichie




Figure 16: *Americanah*, *Adichie*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.4 The Baghdad Clock, Al Rawi




Figure 17: *The Baghdad Clock*, *Al Rawi*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.5 *Industrial Magic*, **Armstrong**




Figure 18: *Industrial Magic*, **Armstrong**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.6 *Fantastic Voyage*, **Asimov**




Figure 19: *Fantastic Voyage*, **Asimov**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.7 The Complete Robot, Asimov




Figure 20: *The Complete Robot*, Asimov. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.8 The Handmaid's Tale, Atwood




Figure 21: *The Handmaid's Tale*, Atwood. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.9 *Pride and Prejudice*, **Austen**




Figure 22: *Pride and Prejudice*, **Austen**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.10 *The Christmas Train*, **Baldacci**




Figure 23: *The Christmas Train*, **Baldacci**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.11 *Notes of a Native Son*, Baldwin




Figure 24: *Notes of a Native Son*, Baldwin. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.12 *Another Country*, Baldwin




Figure 25: *Another Country*, Baldwin. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.13 *The Lemon Table, Barnes*




Figure 26: *The Lemon Table*, Barnes. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.14 *Dante and the Origins of Italian Literary Culture, Barolini*




Figure 27: *Dante and the Origins of Italian Literary Culture*, Barolini. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.15 *The Parthenon, Beard*




Figure 28: *The Parthenon, Beard*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.16 *Guam: Past and Present, Beardsley*




Figure 29: *Guam: Past and Present, Beardsley*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.17 *Waiting for Godot*, **Beckett**




Figure 30: *Waiting for Godot*, **Beckett**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.18 *The Lonely Soldier*, **Benedict**




Figure 31: *The Lonely Soldier*, **Benedict**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.19 Simple Cakes, Berry




Figure 32: *Simple Cakes, Berry*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.20 Paradise Valley, Box




Figure 33: *Paradise Valley, Box*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.21 *The Cat's Pajamas*, Bradbury




Figure 34: *The Cat's Pajamas*, Bradbury. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.22 *London in Chains*, Bradshaw




Figure 35: *London in Chains*, Bradshaw. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.23 *My Einstein*, **Brockman**




Figure 36: *My Einstein*, **Brockman**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.24 *The Da Vinci Code*, **Brown**




Figure 37: *The Da Vinci Code*, **Brown**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.25 *Live and Learn*, Bryant




Figure 38: *Live and Learn*, Bryant. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.26 *Knowing Your Value*, Brzezinski




Figure 39: *Knowing Your Value*, Brzezinski. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.27 *The Myth of Sisyphus*, Camus




Figure 40: *The Myth of Sisyphus*, Camus. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.28 *Alice's Adventures in Wonderland*, Carroll




Figure 41: *Alice's Adventures in Wonderland*, Carroll. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.29 The Infinity Link, Carver




Figure 42: *The Infinity Link*, **Carver**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.30 Murder on the Orient Express, Christie




Figure 43: *Murder on the Orient Express*, **Christie**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.31 *And Then There Were None, Christie*




Figure 44: *And Then There Were None, Christie*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.32 *The Beautiful Struggle, Coates*




Figure 45: *The Beautiful Struggle, Coates*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.33 *We Were Eight Years in Power*, Coates



Figure 46: *We Were Eight Years in Power*, Coates. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.34 *The Water Dancer*, Coates




Figure 47: *The Water Dancer*, Coates. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.35 *The Infernal Machine, Cocteau*




Figure 48: *The Infernal Machine, Cocteau*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.36 *The Alchemist, Coelho*




Figure 49: *The Alchemist, Coelho*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.37 *Dungeons and Dragons and Philosophy*, Cogburn and Silcox




Figure 50: *Dungeons and Dragons and Philosophy*, Cogburn and Silcox. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.38 *Mark Rothko*, Cohen-Solal




Figure 51: *Mark Rothko*, Cohen-Solal. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.39 *The Hunger Games*, Collins




Figure 52: *The Hunger Games*, Collins. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.40 *The Dragon Never Sleeps*, Cook




Figure 53: *The Dragon Never Sleeps*, Cook. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.41 *The 7 Habits of Highly Effective People*, Covey




Figure 54: *The 7 Habits of Highly Effective People*, Covey. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.42 *Bad Kid*, Crabb




Figure 55: *Bad Kid*, Crabb. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.43 *Lullaby Town, Crais*




Figure 56: *Lullaby Town, Crais*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.44 *Jurassic Park, Crichton*






Figure 57: *Jurassic Park, Crichton*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.45 *The Hours*, Cunningham



D.3.46 *Inhuman Land*, Czapski



D.3.47 *Charlie and the Chocolate Factory*, Dahl




Figure 60: *Charlie and the Chocolate Factory*, Dahl. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.48 *James and the Giant Peach*, Dahl




Figure 61: *James and the Giant Peach*, Dahl. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.49 Automating the News, Diakopoulos




Figure 62: *Automating the News*, Diakopoulos. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.50 Drown, Díaz




Figure 63: *Drown*, Díaz. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.51 *The Brief Wondrous Life of Oscar Wao*, Díaz




Figure 64: *The Brief Wondrous Life of Oscar Wao*, Díaz. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.52 *This Is How You Lose Her*, Díaz




Figure 65: *This Is How You Lose Her*, Díaz. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.53 *The White Album, Didion*




Figure 66: *The White Album, Didion*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.54 *Down and Out in the Magic Kingdom, Doctorow*




Figure 67: *Down and Out in the Magic Kingdom, Doctorow*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.55 *The World's Wife, Duffy*




Figure 68: *The World's Wife, Duffy*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.56 *A Visit from the Goon Squad, Egan*




Figure 69: *A Visit from the Goon Squad, Egan*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.57 *Invisible Man*, Ellison




Figure 70: *Invisible Man*, Ellison. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.58 *We Should All Be Mirandas*, Fairless and Garroni




Figure 71: *We Should All Be Mirandas*, Fairless and Garroni. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.59 *The Dude Abides*, **Falsani**




Figure 72: *The Dude Abides*, **Falsani**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.60 *The President's Vampire*, **Farnsworth**




Figure 73: *The President's Vampire*, **Farnsworth**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.61 *The Great Gatsby*, **Fitzgerald**




Figure 74: *The Great Gatsby*, **Fitzgerald**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.62 *Gone Girl*, **Flynn**




Figure 75: *Gone Girl*, **Flynn**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.63 *British Destroyers, Friedman*




Figure 76: *British Destroyers, Friedman*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.64 *Florals & Botanicals, Fudurich et al.*




Figure 77: *Florals & Botanicals, Fudurich et al.* For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.65 *Good Omens*, Gaiman and Pratchett




Figure 78: *Good Omens*, Gaiman and Pratchett. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.66 *The Slippery Year*, Gideon




Figure 79: *The Slippery Year*, Gideon. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.67 *Sweater Surgery*, Girard




Figure 80: *Sweater Surgery*, Girard. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.68 *Blink*, Gladwell




Figure 81: *Blink*, Gladwell. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.69 *The Land Before Avocado*, **Glover**




Figure 82: *The Land Before Avocado*, **Glover**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.70 *Dead Ringers*, **Golden**




Figure 83: *Dead Ringers*, **Golden**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.71 Ararat, Golden




Figure 84: *Ararat*, **Golden**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.72 Lord of the Flies, Golding




Figure 85: *Lord of the Flies*, **Golding**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.73 Wizard’s First Rule, Goodkind




Figure 86: *Wizard’s First Rule, Goodkind*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.74 Rome and Jerusalem, Goodman




Figure 87: *Rome and Jerusalem, Goodman*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.75 *The Fault in Our Stars*, **Green**




Figure 88: *The Fault in Our Stars*, **Green**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.76 *The Third Man*, **Greene**




Figure 89: *The Third Man*, **Greene**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.77 *The Confessions of Max Tivoli*, Greer




Figure 90: *The Confessions of Max Tivoli*, Greer. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.78 *The Fugitive*, Grisham




Figure 91: *The Fugitive*, Grisham. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.79 *The Curious Incident of the Dog in the Night-Time*, Haddon




Figure 92: *The Curious Incident of the Dog in the Night-Time*, Haddon. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.80 *Migrations to Solitude*, Halpern




Figure 93: *Migrations to Solitude*, Halpern. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.81 *Uncommon Type*, **Hanks**




Figure 94: *Uncommon Type*, **Hanks**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.82 *Buzz*, **Hanson**




Figure 95: *Buzz*, **Hanson**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.83 *Requiem for the Sun*, Haydon




Figure 96: *Requiem for the Sun*, Haydon. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.84 *Catch-22*, Heller




Figure 97: *Catch-22*, Heller. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.85 *The Old Man and the Sea*, Hemingway




Figure 98: *The Old Man and the Sea*, Hemingway. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.86 *Life on Air*, Hendy




Figure 99: *Life on Air*, Hendy. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.87 *Great Hair Days*, Hersheson




Figure 100: *Great Hair Days*, Hersheson. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.88 *Graft*, Hill




Figure 101: *Graft*, Hill. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.89 *The Outsiders*, [Hinton](#)




Figure 102: *The Outsiders*, [Hinton](#). For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.90 *The Second Summoning*, [Huff](#)




Figure 103: *The Second Summoning*, [Huff](#). For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.91 Selected Poems of Langston Hughes, Hughes




Figure 104: *Selected Poems of Langston Hughes*, Hughes. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.92 M. Butterfly, Hwang




Figure 105: *M. Butterfly*, Hwang. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.93 Building and Operating a Realistic Model Railway, Jackson




Figure 106: *Building and Operating a Realistic Model Railway, Jackson*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.94 All the Onions, Jacobs




Figure 107: *All the Onions, Jacobs*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.95 *Fifty Shades of Grey*, James




Figure 108: *Fifty Shades of Grey*, James. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.96 *The Stone Sky*, Jemisin




Figure 109: *The Stone Sky*, Jemisin. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.97 *Ulysses, Joyce*




Figure 110: *Ulysses, Joyce*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.98 *Sandman Slim, Kadrey*




Figure 111: *Sandman Slim, Kadrey*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.99 *Ethnography after Antiquity*, Kaldellis




Figure 112: *Ethnography after Antiquity*, Kaldellis. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.100 *Who Is Rich?*, Klam




Figure 113: *Who Is Rich?*, Klam. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.101 *The Servants of Twilight*, Koontz




Figure 114: *The Servants of Twilight*, Koontz. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.102 *Tai Chi for Depression*, Kuhn




Figure 115: *Tai Chi for Depression*, Kuhn. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.103 *The Tide Was Always High, Kun*




Figure 116: *The Tide Was Always High, Kun*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.104 *Girl in Translation, Kwok*




Figure 117: *Girl in Translation, Kwok*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.105 A Wrinkle in Time, L'Engle




Figure 118: *A Wrinkle in Time*, L'Engle. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.106 Call Me Brooklyn, Lago




Figure 119: *Call Me Brooklyn*, Lago. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.107 *Dead Wake*, [Larson](#)




Figure 120: *Dead Wake*, [Larson](#). For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.108 *The Girl with the Dragon Tattoo*, [Larsson](#)




Figure 121: *The Girl with the Dragon Tattoo*, [Larsson](#). For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.109 *The Daughter of Odren, Le Guin*




Figure 122: *The Daughter of Odren, Le Guin*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.110 *The Chronicles of Narnia, Lewis*




Figure 123: *The Chronicles of Narnia, Lewis*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.111 *After I'm Gone*, Lippman




Figure 124: *After I'm Gone*, Lippman. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.112 *Sunburn*, Lippman




Figure 125: *Sunburn*, Lippman. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.113 *Marvel's Spider-Man: Hostile Takeover*, [Liss](#)




Figure 126: *Marvel's Spider-Man: Hostile Takeover*, [Liss](#). For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.114 *Anastasia on Her Own*, [Lowry](#)




Figure 127: *Anastasia on Her Own*, [Lowry](#). For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.115 *Nixon in China*, MacMillan




Figure 128: *Nixon in China*, MacMillan. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.116 *The Domsday Prophecy*, Mariani




Figure 129: *The Domsday Prophecy*, Mariani. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.117 *A Game of Thrones, Martin*




Figure 130: *A Game of Thrones, Martin*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.118 *Rough-Hewn Land, Meldahl*




Figure 131: *Rough-Hewn Land, Meldahl*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.119 *Twilight*, **Meyer**




Figure 132: *Twilight*, **Meyer**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.120 *The Duchess War*, **Milan**




Figure 133: *The Duchess War*, **Milan**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.121 *On Liberty*, **Mill**




Figure 134: *On Liberty*, **Mill**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.122 *Coal Creek*, **Miller**




Figure 135: *Coal Creek*, **Miller**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.123 *Winnie the Pooh, Milne*




Figure 136: *Winnie the Pooh, Milne*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.124 *Catching the Sky, Moore and O'Brien*




Figure 137: *Catching the Sky, Moore and O'Brien*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.125 *The Heretic Queen*, [Moran](#)




Figure 138: *The Heretic Queen*, [Moran](#). For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.126 *The Gondola Maker*, [Morelli](#)




Figure 139: *The Gondola Maker*, [Morelli](#). For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.127 *Songs in Ordinary Time*, **Morris**




Figure 140: *Songs in Ordinary Time*, **Morris**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.128 *Beloved*, **Morrison**




Figure 141: *Beloved*, **Morrison**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.129 Norwegian Wood, Murakami




Figure 142: *Norwegian Wood*, Murakami. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.130 Eat More Plants, Nielsen




Figure 143: *Eat More Plants*, Nielsen. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.131 *Polaris*, *Northrop*




Figure 144: *Polaris*, *Northrop*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.132 *Pagans*, *O'Donnell*




Figure 145: *Pagans*, *O'Donnell*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.133 Windfall, O’Sullivan




Figure 146: *Windfall*, O’Sullivan. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.134 The Memory Police, Ogawa




Figure 147: *The Memory Police*, Ogawa. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.135 *Winter Sisters*, Oliveira




Figure 148: *Winter Sisters*, Oliveira. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.136 *Nineteen Eighty-Four*, Orwell




Figure 149: *Nineteen Eighty-Four*, Orwell. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.137 *Fight Club*, Palahniuk




Figure 150: *Fight Club*, Palahniuk. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.138 *The Complete Joy of Homebrewing*, Papazian




Figure 151: *The Complete Joy of Homebrewing*, Papazian. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.139 *The Cult of Loving Kindness*, **Park**




Figure 152: *The Cult of Loving Kindness*, **Park**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.140 *Along Came a Spider*, **Patterson**




Figure 153: *Along Came a Spider*, **Patterson**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.141 *Payard Cookies*, Payard and McBride




Figure 154: *Payard Cookies*, Payard and McBride. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.142 *Essential Pepin Desserts*, Pépin




Figure 155: *Essential Pepin Desserts*, Pépin. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.143 Why New Orleans Matters, Piazza




Figure 156: *Why New Orleans Matters, Piazza*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.144 Enlightenment Now, Pinker




Figure 157: *Enlightenment Now, Pinker*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.145 Competitive Strategy, Porter




Figure 158: *Competitive Strategy, Porter*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.146 Night Watch, Pratchett




Figure 159: *Night Watch, Pratchett*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.147 *The Subtle Knife*, Pullman




Figure 160: *The Subtle Knife*, Pullman. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.148 *The Seductions of Darwin*, Rampley




Figure 161: *The Seductions of Darwin*, Rampley. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.149 Kitchen Table Wisdom, **Remen**




Figure 162: *Kitchen Table Wisdom*, **Remen**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.150 Backroads Boss Lady, **Roberts and Witter**




Figure 163: *Backroads Boss Lady*, **Roberts and Witter**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.151 *Soft in the Head, Roger*




Figure 164: *Soft in the Head, Roger*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.152 *The Making of a Mediterranean Emirate, Roughi*




Figure 165: *The Making of a Mediterranean Emirate, Roughi*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.153 *Harry Potter and the Sorcerer's Stone*, Rowling




Figure 166: *Harry Potter and the Sorcerer's Stone*, Rowling. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.154 *Harry Potter and the Chamber of Secrets*, Rowling




Figure 167: *Harry Potter and the Chamber of Secrets*, Rowling. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.155 *Harry Potter and the Goblet of Fire*, [Rowling](#)




Figure 168: *Harry Potter and the Goblet of Fire*, [Rowling](#). For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.156 *Harry Potter and the Deathly Hallows*, [Rowling](#)




Figure 169: *Harry Potter and the Deathly Hallows*, [Rowling](#). For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.157 *Born to Walk*, **Rubinstein**




Figure 170: *Born to Walk*, **Rubinstein**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.158 *The Pretender*, **Ruskin**




Figure 171: *The Pretender*, **Ruskin**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.159 *Toscanini, Sachs*




Figure 172: *Toscanini, Sachs*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.160 *Cosmos, Sagan*




Figure 173: *Cosmos, Sagan*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.161 *Middle India*, Sahni




Figure 174: *Middle India*, Sahni. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.162 *The Catcher in the Rye*, Salinger




Figure 175: *The Catcher in the Rye*, Salinger. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.163 *Lean In*, Sandberg




Figure 176: *Lean In*, Sandberg. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.164 *The DevOps Adoption Playbook*, Sharma




Figure 177: *The DevOps Adoption Playbook*, Sharma. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.165 *Frankenstein, Shelley*




Figure 178: *Frankenstein, Shelley*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.166 *Sally Ride: America's First Woman in Space, Sherr*




Figure 179: *Sally Ride: America's First Woman in Space, Sherr*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.167 *A Perfectly Good Family*, **Shriver**




Figure 180: *A Perfectly Good Family*, **Shriver**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.168 *The Bedwetter*, **Silverman**




Figure 181: *The Bedwetter*, **Silverman**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.169 *On the Road with Bob Dylan, Sloman*




Figure 182: *On the Road with Bob Dylan, Sloman*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.170 *The Night Children, Smith*




Figure 183: *The Night Children, Smith*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.171 *White Teeth*, **Smith**




Figure 184: *White Teeth*, **Smith**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.172 *No Visible Bruises*, **Snyder**




Figure 185: *No Visible Bruises*, **Snyder**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.173 *The Grapes of Wrath*, Steinbeck




Figure 186: *The Grapes of Wrath*, Steinbeck. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.174 *Jesse James*, Stiles




Figure 187: *Jesse James*, Stiles. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.175 *Zombie Halloween*, **Stine**




Figure 188: *Zombie Halloween*, **Stine**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.176 *Fear of Music*, **Stubbs**




Figure 189: *Fear of Music*, **Stubbs**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.177 *Pearl Harbor*, **Swanson**




Figure 190: *Pearl Harbor*, **Swanson**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.178 *Chesapeake Requiem*, **Swift**




Figure 191: *Chesapeake Requiem*, **Swift**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.179 *The Goldfinch*, **Tartt**




Figure 192: *The Goldfinch*, **Tartt**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.180 *Unglued*, **TerKeurst**




Figure 193: *Unglued*, **TerKeurst**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.181 *Embraced*, **TerKeurst**




Figure 194: *Embraced*, **TerKeurst**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.182 *Birding with Yeats*, **Thomson**




Figure 195: *Birding with Yeats*, **Thomson**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.183 *The Hobbit*, Tolkien




Figure 196: *The Hobbit*, Tolkien. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.184 *The Fellowship of the Ring*, Tolkien




Figure 197: *The Fellowship of the Ring*, Tolkien. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.185 *Tree and Leaf*, **Tolkien**




Figure 198: *Tree and Leaf*, **Tolkien**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.186 *Noodles Every Day*, **Trang**






Figure 199: *Noodles Every Day*, **Trang**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.187 *Billionaire Democracy*, Tyler



D.3.188 *Portugal's Guerrilla Wars in Africa*, Venter



D.3.189 Slaughterhouse-Five, Vonnegut




Figure 202: *Slaughterhouse-Five*, *Vonnegut*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.190 Animal Rights, Waldau




Figure 203: *Animal Rights*, *Waldau*. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.191 *Men We Reaped*, **Ward**




Figure 204: *Men We Reaped*, **Ward**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.192 *Charlotte's Web*, **White**




Figure 205: *Charlotte's Web*, **White**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.193 *A Return to Love*, **Williamson**




Figure 206: *A Return to Love*, **Williamson**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.194 *Another Brooklyn*, **Woodson**




Figure 207: *Another Brooklyn*, **Woodson**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.195 *Brown Girl Dreaming*, Woodson




Figure 208: *Brown Girl Dreaming*, Woodson. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.196 *A Little Life*, Yanagihara




Figure 209: *A Little Life*, Yanagihara. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.197 *The Art of Bonsai*, Yoshimura and Halford




Figure 210: *The Art of Bonsai*, Yoshimura and Halford. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.198 *A People's History of the United States*, Zinn




Figure 211: *A People's History of the United States*, Zinn. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.199 *The Future of the Internet and How to Stop It*, **Zittrain**




Figure 212: *The Future of the Internet and How to Stop It*, **Zittrain**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.3.200 *The Book Thief*, **Zusak**




Figure 213: *The Book Thief*, **Zusak**. For 14 LLMs, (left) heatmaps for the sliding-window procedure and (right) corresponding distributions over suffix extraction probabilities ($\tau_{\min} = 0.1\%$).

D.4 Estimating how much of a book is memorized by a model

We estimate book-specific memorization for a given model θ and decoding scheme ϕ with a metric for **extraction coverage**. We define this metric and provide measurements for the different models and books we evaluate.

D.4.1 Quantifying extraction coverage

We compute extraction coverage as the fraction of characters in a book that lie within at least one suffix whose extraction probability p_z exceeds a threshold τ .

Definition 2 (Extraction coverage). Let $\text{span}(z)$ denote the set of character indices spanned by a 50-token suffix of a given sequence $z \in \mathbb{Z}(B)$, where B is a book. Also let $\tau_{\min}$ denote the minimum probability we consider to reflect extraction success. For book B , LLM θ , decoding scheme ϕ , and probability threshold $\tau \in [\tau_{\min}, 1]$,

$$\text{coverage}_{B,\theta,\phi}(\tau) = \frac{1}{L_B} \left| \bigcup_{z \in \mathbb{Z}(B): p_{z,\theta,\phi} \geq \tau} \text{span}(z) \right|, \quad (10)$$

where L_B is the total number of characters in the book.

For details on $\tau_{\min}$, see Equation 9 (Sections 2.1 & 5; Appendices C & E). In the main text, we omit θ and ϕ for brevity, since we always specify the exact model and decoding scheme being considered.

The results of our sliding-window procedure can be used to compute Equation 10. This procedure returns a list of p_z sorted by start position in the book (i.e., $\mathbb{Z}(B)$). We pick a threshold τ and filter the list of p_z to only contain sequences whose extraction probability is $\geq \tau$. Our sliding-window procedure keeps track of character-level metadata so that we can visualize heatmaps; we have information about the start/end character indexes for the prefixes and suffixes. Based on these intervals, we can compute which characters in the whole book are covered with $p_z \geq \tau$, and then divide that by the total number of characters in the book.

High extraction coverage does not necessarily mean we can extract the whole book in one go, or that we have identified the minimal set of sequences to try to extract to get maximal coverage of the book. This is simply a procedure for estimating the proportion of total memorized text for a book, in order to do cross-book and cross-model comparisons, as in Section 4.2.

D.4.2 Extended extraction coverage results

We display extraction coverage in two ways, which support different types of comparisons.

1. **Model-level histograms.** For each model, we plot the distribution of extraction coverage across the 200 evaluated books. This provides a high-level sense of how coverage varies across books, and enables comparisons of this distribution across models. For example, in the main paper we show two histograms (for LLAMA 1 65B and LLAMA 3.1 70B) at $\tau = 0.1\%$, the minimum extraction probability that we validate in this work. These results indicate that most books have little extraction coverage. But there are clear exceptions. They also reveal that some models exhibit substantially higher coverage than others—in this case, LLAMA 3.1 70B shows higher maximum coverage and more books with significant coverage, i.e., it has memorized much more book text than LLAMA 1 65B.
2. **Book-level percentages.** For each book, we compare extraction coverage across different models. This not only indicates how memorized a given book is by a particular model, but also highlights variation in memorization of a given book across models. We provide examples of this comparison in the table in Figure 3 (right), for *Harry Potter and the Sorcerer’s Stone* and *Sandman Slim*, evaluated with LLAMA 3.1 70B, LLAMA 1 65B, and PYTHIA 12B.

Here, we provide more extensive extraction coverage results. We include both types of comparisons for the 14 models for which detailed sliding-window results are given in Appendix D.3. All results are shown for $\tau = \tau_{\min} = 0.1\%$, which corresponds to the total (largest) amount of extraction coverage revealed by our experiments. Figure 214 plots histograms of extraction coverage across all 200 books for each model, and Table 4 reports per-book extraction coverage results for each model.

Table 4: **Extraction coverage for 200 books and 14 LLMs.** We use Definition 2 to compute extraction coverage for the results in Appendix D.3, which use 100-token sequences (50-token prefix + 50-token suffix) and top- k decoding ($T = 1, k = 40$). We set $\tau = \tau_{\min} = 0.1\%$ to compute the total amount of extraction coverage that we observe.

Proportion of each book that is extractable with $p_z \geq 0.1\%$ for different models														
Book	LLAMA 2 7B	LLAMA 3 8B	LLAMA 3.1 8B	LLAMA 1 13B	LLAMA 2 13B	PYTHIA 12B	LLAMA 1 65B	LLAMA 2 70B	LLAMA 3 70B	LLAMA 3.1 70B	DS v1 67B	QWEN 2.5 72B	GEMMA 2 27B	PHI 4
<i>A Game of Thrones</i>	0.65%	0.75%	0.78%	0.76%	0.78%	0.67%	3.34%	3.07%	37.39%	39.39%	0.81%	1.25%	0.87%	0.59%
<i>A Little Life</i>	0.02%	0.03%	0.03%	0.05%	0.05%	0.04%	0.16%	0.16%	0.35%	0.46%	0.04%	0.03%	0.07%	0.04%
<i>A People's History of the United States</i>	1.30%	2.01%	2.05%	1.91%	1.66%	0.92%	6.86%	6.77%	17.91%	26.27%	1.71%	18.82%	2.05%	1.01%
<i>A Perfectly Good Family</i>	0.17%	0.24%	0.24%	0.14%	0.18%	0.27%	0.21%	0.19%	0.42%	0.58%	0.28%	0.28%	0.24%	0.22%
<i>A Return to Love</i>	0.64%	0.85%	0.85%	0.85%	0.94%	0.62%	2.76%	3.04%	4.35%	4.02%	1.08%	0.80%	0.73%	0.58%
<i>A Visit from the Goon Squad</i>	0.07%	0.00%	0.10%	0.07%	0.07%	0.20%	0.29%	0.36%	0.74%	1.06%	0.19%	0.22%	0.17%	0.11%
<i>A Wrinkle in Time</i>	1.44%	1.62%	1.57%	1.61%	1.69%	1.75%	4.54%	3.93%	10.00%	14.59%	2.32%	2.09%	1.57%	1.44%
<i>After I'm Gone</i>	0.19%	0.28%	0.28%	0.16%	0.27%	0.36%	0.39%	0.31%	0.40%	0.46%	0.41%	0.36%	0.28%	0.24%
<i>Alice's Adventures in Wonderland</i>	6.44%	18.39%	18.84%	18.36%	21.02%	1.65%	42.89%	42.72%	47.74%	48.80%	38.86%	38.00%	24.94%	8.33%
<i>All the Onions</i>	0.27%	1.61%	0.00%	0.27%	0.27%	2.61%	0.88%	1.19%	2.77%	2.77%	2.77%	2.55%	2.11%	0.00%
<i>Along Came a Spider</i>	0.38%	0.55%	0.55%	0.48%	0.46%	0.78%	0.71%	0.92%	1.10%	1.13%	1.13%	0.56%	0.41%	0.21%
<i>Americanah</i>	0.12%	0.13%	0.17%	0.14%	0.26%	0.17%	0.58%	0.64%	0.85%	0.98%	0.17%	0.26%	0.13%	0.14%
<i>Anastasia on Her Own</i>	0.00%	0.49%	0.40%	0.27%	0.30%	0.25%	0.57%	0.82%	0.79%	0.78%	0.25%	0.75%	0.25%	0.50%
<i>And Then There Were None</i>	1.44%	1.51%	1.91%	1.49%	1.82%	2.04%	2.44%	2.65%	4.69%	6.03%	3.45%	3.18%	1.92%	1.74%
<i>Animal Rights</i>	0.54%	0.61%	0.69%	0.63%	0.65%	0.59%	1.10%	1.00%	1.34%	1.46%	0.79%	0.97%	0.80%	0.48%
<i>Another Brooklyn</i>	2.04%	2.45%	2.41%	1.97%	2.17%	2.59%	3.01%	2.68%	2.84%	3.16%	2.85%	2.62%	2.46%	1.87%
<i>Another Country</i>	0.12%	0.16%	0.16%	0.13%	0.13%	0.18%	0.29%	0.34%	0.55%	0.74%	0.17%	0.22%	0.14%	0.10%
<i>Ararat</i>	0.37%	0.29%	0.29%	0.37%	0.37%	0.46%	0.43%	0.49%	0.43%	0.45%	0.56%	0.56%	0.42%	0.18%
<i>Automating the News</i>	1.26%	0.50%	0.56%	1.42%	1.55%	0.87%	1.86%	1.96%	1.02%	1.16%	1.55%	0.88%	1.15%	0.50%
<i>Backroads Boss Lady</i>	0.35%	0.31%	0.31%	0.35%	0.32%	0.43%	0.36%	0.39%	0.32%	0.32%	0.42%	0.43%	0.40%	0.00%
<i>Bad Kid</i>	0.78%	0.83%	0.87%	0.77%	0.82%	0.93%	0.87%	0.88%	0.91%	0.91%	0.94%	0.94%	0.89%	0.82%
<i>Beloved</i>	0.17%	0.69%	0.45%	0.35%	0.61%	0.09%	1.87%	3.58%	9.78%	12.54%	0.34%	1.55%	0.69%	0.32%
<i>Billionaire Democracy</i>	0.64%	0.50%	0.43%	1.19%	1.06%	0.34%	2.04%	2.11%	1.68%	1.69%	1.16%	0.80%	0.84%	0.36%
<i>Birding with Yeats</i>	0.06%	0.21%	0.20%	0.06%	0.06%	0.48%	0.33%	0.36%	0.62%	0.53%	0.51%	0.48%	0.40%	0.14%
<i>Blink</i>	0.31%	0.87%	0.87%	0.60%	0.63%	0.48%	2.19%	2.87%	6.86%	9.73%	1.51%	2.62%	0.87%	0.65%
<i>Born to Walk</i>	0.23%	0.17%	0.17%	0.27%	0.27%	0.16%	0.73%	0.76%	0.89%	0.86%	0.42%	0.61%	0.19%	0.10%
<i>British Destroyers</i>	0.03%	0.02%	0.00%	0.04%	0.05%	0.00%	0.08%	0.07%	0.06%	0.06%	0.08%	0.05%	0.02%	0.00%
<i>Brown Girl Dreaming</i>	0.31%	0.34%	0.34%	0.30%	0.32%	0.43%	0.81%	0.83%	0.54%	0.55%	0.80%	0.53%	0.58%	0.00%
<i>Building and Operating a Realistic Model Railway</i>	0.11%	0.00%	0.00%	0.11%	0.00%	0.20%	0.00%	0.11%	0.14%	0.14%	0.31%	0.22%	0.12%	0.00%
<i>Buzz</i>	0.05%	0.14%	0.20%	0.10%	0.10%	0.07%	0.58%	0.61%	2.16%	1.98%	0.24%	1.35%	0.42%	0.08%
<i>Call Me Brooklyn</i>	0.08%	0.09%	0.09%	0.11%	0.12%	0.11%	1.35%	0.79%	1.38%	1.59%	0.95%	0.09%	0.09%	0.08%
<i>Catch-22</i>	0.34%	0.57%	0.61%	0.60%	0.75%	0.32%	2.60%	2.98%	10.01%	15.64%	0.60%	0.93%	0.56%	0.23%
<i>Catching the Sky</i>	0.11%	0.20%	0.12%	0.20%	0.27%	0.37%	0.27%	0.27%	0.21%	0.21%	0.37%	0.30%	0.20%	0.00%
<i>Charlie and the Chocolate Factory</i>	0.82%	0.92%	0.92%	0.82%	1.00%	0.92%	4.49%	5.70%	17.09%	27.67%	1.50%	2.43%	3.83%	0.92%
<i>Charlotte's Web</i>	0.66%	1.71%	1.66%	1.47%	1.35%	0.82%	5.11%	7.23%	16.72%	21.96%	1.85%	6.47%	1.94%	1.29%
<i>Chesapeake Requiem</i>	0.58%	0.58%	0.58%	0.60%	0.64%	0.70%	0.68%	0.70%	0.70%	0.70%	0.71%	0.70%	0.68%	0.58%
<i>Coal Creek</i>	0.20%	0.00%	0.15%	0.25%	0.25%	0.33%	0.32%	0.22%	0.21%	0.29%	0.34%	0.28%	0.25%	0.13%
<i>Competitive Strategy</i>	0.09%	0.16%	0.12%	0.11%	0.11%	0.18%	0.25%	0.48%	1.70%	1.87%	0.21%	0.25%	0.26%	0.00%
<i>Cosmos</i>	0.33%	1.10%	1.39%	0.99%	0.90%	0.18%	3.64%	4.22%	7.73%	11.57%	0.63%	3.95%	0.99%	0.32%
<i>Dante and the Origins of Italian Literary Culture</i>	0.00%	0.04%	0.04%	0.00%	0.00%	0.04%	0.00%	0.00%	0.36%	0.55%	0.06%	0.08%	0.06%	0.08%
<i>Dead Ringers</i>	0.50%	0.42%	0.42%	0.47%	0.49%	0.66%	0.68%	0.56%	0.64%	0.65%	0.68%	0.69%	0.50%	0.28%
<i>Dead Wake</i>	0.07%	0.08%	0.12%	0.07%	0.10%	0.06%	0.13%	0.17%	0.47%	0.41%	0.18%	0.19%	0.06%	0.12%

Continued on next page

Table 4 (continued). Proportion of each book extractable with $p_z \geq 0.1\%$ for different models.

Book	LLAMA 2 7B	LLAMA 3 8B	LLAMA 3.1 8B	LLAMA 1 13B	LLAMA 2 13B	PYTHIA 12B	LLAMA 1 65B	LLAMA 2 70B	LLAMA 3 70B	LLAMA 3.1 70B	DS v1 67B	QWEN 2.5 72B	GEMMA 2 27B	PHI 4
<i>Down and Out in the Magic Kingdom</i>	0.30%	0.21%	0.21%	0.31%	0.31%	0.47%	0.46%	0.56%	0.94%	1.11%	0.47%	0.48%	0.37%	0.00%
<i>Drown</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.21%	0.37%	0.37%	0.12%	0.00%	0.00%	0.00%
<i>Dungeons and Dragons and Philosophy</i>	0.32%	0.31%	0.26%	0.45%	0.63%	0.86%	1.43%	1.31%	1.61%	1.62%	1.29%	0.91%	0.48%	0.31%
<i>Eat More Plants</i>	0.74%	0.69%	0.78%	0.74%	0.75%	0.73%	0.96%	1.26%	1.57%	1.61%	1.11%	1.49%	0.95%	0.70%
<i>Embraced</i>	5.25%	4.61%	4.81%	6.38%	5.40%	4.12%	6.92%	6.97%	5.68%	6.04%	6.70%	5.48%	4.76%	2.94%
<i>Enlightenment Now</i>	0.84%	1.19%	1.24%	1.05%	1.04%	0.59%	2.10%	2.16%	5.46%	5.16%	1.27%	2.21%	1.11%	0.84%
<i>Essential Pepin Desserts</i>	0.94%	1.10%	1.10%	0.88%	0.83%	0.91%	0.83%	0.88%	1.23%	1.34%	1.32%	1.41%	0.72%	0.98%
<i>Ethnography after Antiquity</i>	0.05%	0.05%	0.05%	0.05%	0.08%	0.13%	0.05%	0.21%	0.52%	0.69%	0.18%	0.24%	0.10%	0.00%
<i>Fantastic Voyage</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.08%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
<i>Fear of Music</i>	0.00%	0.33%	0.28%	0.00%	0.00%	0.40%	0.72%	0.80%	0.95%	0.96%	0.44%	0.33%	0.18%	0.00%
<i>Fifty Shades of Grey</i>	0.17%	0.14%	0.13%	0.12%	0.12%	0.15%	1.39%	1.03%	2.00%	2.71%	0.21%	0.16%	0.15%	0.13%
<i>Fight Club</i>	0.16%	0.27%	0.27%	0.41%	0.47%	0.26%	2.15%	3.58%	6.11%	6.65%	0.39%	0.26%	0.58%	0.27%
<i>Florals & Botanicals</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
<i>Frankenstein</i>	3.91%	15.87%	11.86%	11.32%	19.58%	0.64%	65.15%	65.67%	83.34%	83.52%	28.78%	35.26%	49.19%	21.92%
<i>Girl in Translation</i>	0.41%	0.27%	0.33%	0.37%	0.45%	0.49%	0.47%	0.46%	0.54%	0.59%	0.48%	0.48%	0.30%	0.00%
<i>Gone Girl</i>	0.19%	0.24%	0.24%	0.20%	0.38%	0.17%	0.81%	1.05%	1.69%	2.12%	0.29%	0.29%	0.25%	0.18%
<i>Good Omens</i>	0.10%	0.13%	0.09%	0.13%	0.20%	0.22%	1.50%	1.77%	3.98%	4.51%	0.21%	0.29%	0.20%	0.06%
<i>Graft</i>	0.05%	0.00%	0.00%	0.05%	0.05%	0.19%	0.05%	0.11%	0.11%	0.11%	0.22%	0.04%	0.04%	0.00%
<i>Great Hair Days</i>	0.68%	0.70%	0.70%	0.68%	0.68%	0.78%	0.68%	0.68%	0.70%	0.70%	0.77%	0.75%	0.71%	0.69%
<i>Guam: Past and Present</i>	0.06%	0.00%	0.00%	0.06%	0.06%	0.00%	0.06%	0.24%	0.11%	0.21%	0.00%	0.08%	0.08%	0.00%
<i>Harry Potter and the Chamber of Secrets</i>	0.23%	0.96%	0.65%	0.50%	0.67%	0.30%	12.99%	16.64%	92.44%	94.94%	2.14%	6.53%	2.67%	0.19%
<i>Harry Potter and the Deathly Hallows</i>	0.17%	0.43%	0.51%	0.42%	0.29%	0.04%	13.59%	9.29%	53.25%	64.79%	1.11%	4.06%	0.72%	0.13%
<i>Harry Potter and the Goblet of Fire</i>	0.15%	0.26%	0.39%	0.32%	0.39%	0.13%	26.15%	14.08%	82.35%	88.66%	1.09%	3.64%	1.12%	0.16%
<i>Harry Potter and the Sorcerer's Stone</i>	0.49%	2.15%	2.62%	1.97%	3.02%	0.41%	34.83%	36.41%	93.74%	96.27%	12.00%	58.21%	21.46%	1.73%
<i>Industrial Magic</i>	0.00%	0.06%	0.06%	0.00%	0.00%	0.06%	0.11%	0.06%	0.12%	0.07%	0.08%	0.06%	0.06%	0.00%
<i>Inhuman Land</i>	0.00%	0.00%	0.00%	0.02%	0.00%	0.03%	0.02%	0.02%	0.00%	0.00%	0.03%	0.03%	0.00%	0.00%
<i>Invisible Man</i>	0.80%	0.86%	0.92%	0.82%	0.78%	0.70%	1.02%	1.38%	2.89%	3.22%	0.88%	1.14%	0.84%	0.77%
<i>James and the Giant Peach</i>	1.16%	1.21%	1.21%	1.29%	1.16%	1.24%	2.03%	1.70%	3.26%	3.35%	1.79%	1.24%	1.64%	0.79%
<i>Jesse James</i>	0.09%	0.12%	0.23%	0.16%	0.21%	0.04%	0.86%	0.76%	0.74%	0.93%	0.48%	0.40%	0.08%	0.06%
<i>Jurassic Park</i>	0.00%	0.07%	0.07%	0.06%	0.00%	0.08%	0.51%	0.94%	4.03%	3.96%	0.08%	0.09%	0.08%	0.00%
<i>Kitchen Table Wisdom</i>	0.31%	0.34%	0.35%	0.41%	0.32%	0.23%	1.39%	1.70%	2.20%	2.95%	0.24%	0.51%	0.39%	0.00%
<i>Knowing Your Value</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
<i>Lean In</i>	0.70%	1.99%	1.71%	1.05%	1.48%	0.38%	4.61%	5.16%	8.76%	10.73%	2.49%	3.10%	1.49%	0.94%
<i>Life on Air</i>	0.13%	0.09%	0.11%	0.13%	0.15%	0.11%	0.23%	0.23%	0.23%	0.25%	0.17%	0.15%	0.12%	0.12%
<i>Live and Learn</i>	0.04%	0.00%	0.00%	0.04%	0.05%	0.13%	0.05%	0.04%	0.09%	0.09%	0.17%	0.17%	0.08%	0.00%
<i>London in Chains</i>	0.16%	0.12%	0.13%	0.16%	0.16%	0.22%	0.16%	0.16%	0.12%	0.13%	0.22%	0.20%	0.13%	0.12%
<i>Lord of the Flies</i>	0.88%	1.43%	1.43%	1.37%	1.29%	0.80%	4.85%	7.31%	29.10%	36.27%	2.12%	8.14%	10.05%	1.42%
<i>Lullaby Town</i>	0.26%	0.31%	0.30%	0.24%	0.27%	0.37%	0.33%	0.34%	0.38%	0.38%	0.35%	0.31%	0.29%	0.19%
<i>M. Butterfly</i>	0.91%	1.63%	1.63%	1.09%	1.13%	1.60%	1.46%	1.81%	2.46%	2.91%	1.72%	1.47%	1.20%	0.66%
<i>Mark Rothko</i>	0.10%	0.18%	0.17%	0.36%	0.20%	0.19%	0.78%	0.81%	1.32%	1.30%	0.47%	0.46%	0.12%	0.32%
<i>Marvel's Spider-Man: Hostile Takeover</i>	0.32%	0.32%	0.41%	0.32%	0.41%	0.42%	0.57%	0.67%	0.64%	0.85%	0.48%	0.44%	0.47%	0.23%
<i>Men We Reaped</i>	0.00%	0.06%	0.16%	0.00%	0.00%	0.17%	0.28%	0.33%	0.85%	1.13%	0.12%	0.15%	0.04%	0.00%
<i>Middle India</i>	0.38%	0.47%	0.47%	0.31%	0.33%	0.49%	0.39%	0.39%	0.50%	0.50%	0.50%	0.50%	0.45%	0.21%
<i>Migrations to Solitude</i>	0.15%	0.41%	0.42%	0.40%	0.39%	0.16%	0.98%	0.76%	1.23%	1.48%	0.59%	0.82%	0.58%	0.23%
<i>Murder on the Orient Express</i>	0.61%	0.78%	0.87%	0.53%	0.86%	1.88%	1.21%	1.48%	2.56%	2.48%	2.34%	1.79%	0.95%	0.78%
<i>My Einstein</i>	0.11%	0.35%	0.48%	0.24%	0.24%	0.00%	0.76%	0.73%	0.97%	1.01%	0.35%	0.70%	0.61%	0.00%
<i>Night Watch</i>	0.14%	0.30%	0.30%	0.13%	0.23%	0.29%	0.37%	0.55%	1.55%	1.90%	0.35%	0.37%	0.29%	0.27%

Continued on next page

Table 4 (continued). Proportion of each book extractable with $p_z \geq 0.1\%$ for different models.

Book	LLAMA 2 7B	LLAMA 3 8B	LLAMA 3.1 8B	LLAMA 1 13B	LLAMA 2 13B	PYTHIA 12B	LLAMA 1 65B	LLAMA 2 70B	LLAMA 3 70B	LLAMA 3.1 70B	DS v1 67B	QWEN 2.5 72B	GEMMA 2 27B	PHI 4
<i>Nineteen Eighty-Four</i>	7.24%	13.27%	16.21%	16.77%	22.50%	1.88%	67.27%	71.21%	79.97%	81.19%	12.50%	30.59%	19.16%	5.09%
<i>Nixon in China</i>	0.15%	0.18%	0.25%	0.16%	0.17%	0.12%	0.77%	0.82%	1.17%	1.28%	0.25%	0.84%	0.18%	0.06%
<i>No Visible Bruises</i>	0.40%	0.49%	0.50%	0.44%	0.47%	0.65%	0.40%	0.56%	0.56%	0.56%	0.66%	0.59%	0.49%	0.49%
<i>Noodles Every Day</i>	0.07%	0.19%	0.26%	0.06%	0.07%	0.32%	0.07%	0.21%	0.33%	0.33%	0.34%	0.23%	0.08%	0.09%
<i>Norwegian Wood</i>	0.26%	0.19%	0.20%	0.22%	0.31%	0.33%	0.74%	0.72%	2.18%	2.52%	0.33%	0.27%	0.26%	0.11%
<i>Notes of a Native Son</i>	0.27%	0.93%	0.83%	0.81%	0.83%	0.16%	3.16%	3.73%	15.59%	16.80%	0.43%	1.06%	0.30%	0.32%
<i>On Liberty</i>	4.71%	21.65%	23.13%	11.10%	13.09%	2.42%	73.12%	49.66%	90.03%	90.83%	9.45%	80.53%	17.30%	7.93%
<i>On the Road with Bob Dylan</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.07%	0.00%	0.00%	0.00%	0.05%	0.07%	0.05%	0.00%	0.00%
<i>Pagans</i>	0.34%	0.40%	0.45%	0.33%	0.38%	0.48%	0.48%	0.49%	0.46%	0.47%	0.47%	0.50%	0.47%	0.32%
<i>Paradise Valley</i>	0.22%	0.35%	0.35%	0.11%	0.21%	0.42%	0.22%	0.22%	0.37%	0.37%	0.38%	0.41%	0.24%	0.21%
<i>Payard Cookies</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.15%	0.12%	0.10%	0.15%	0.15%	0.14%	0.15%	0.14%	0.00%
<i>Pearl Harbor</i>	0.87%	1.02%	1.10%	0.91%	0.85%	1.31%	0.91%	1.10%	1.17%	1.19%	1.22%	1.26%	1.04%	0.00%
<i>Polaris</i>	0.27%	0.31%	0.14%	0.34%	0.34%	0.44%	0.35%	0.39%	0.51%	0.54%	0.44%	0.44%	0.38%	0.16%
<i>Portugal's Guerrilla Wars in Africa</i>	0.05%	0.05%	0.05%	0.04%	0.00%	0.05%	0.11%	0.16%	0.14%	0.20%	0.06%	0.05%	0.04%	0.00%
<i>Pride and Prejudice</i>	3.14%	30.49%	56.87%	11.38%	22.10%	0.94%	72.56%	73.96%	83.74%	84.20%	33.42%	27.32%	10.13%	5.59%
<i>Requiem for the Sun</i>	0.08%	0.10%	0.10%	0.12%	0.11%	0.14%	0.15%	0.22%	0.14%	0.14%	0.20%	0.16%	0.09%	0.10%
<i>Rome and Jerusalem</i>	0.12%	0.34%	0.41%	0.16%	0.12%	0.10%	0.79%	0.59%	2.30%	3.06%	0.21%	0.57%	0.13%	0.10%
<i>Rough-Hewn Land</i>	0.22%	0.28%	0.32%	0.35%	0.31%	0.28%	0.71%	0.59%	1.05%	1.28%	0.54%	0.70%	0.43%	0.21%
<i>Sally Ride: America's First Woman in Space</i>	0.08%	0.10%	0.10%	0.13%	0.13%	0.16%	0.18%	0.13%	0.16%	0.22%	0.18%	0.21%	0.09%	0.06%
<i>Sandman Slim</i>	0.19%	0.29%	0.29%	0.23%	0.27%	0.36%	0.32%	0.32%	0.38%	0.43%	0.37%	0.38%	0.34%	0.25%
<i>Selected Poems of Langston Hughes</i>	1.27%	4.24%	3.65%	1.36%	1.61%	1.63%	6.66%	8.37%	17.76%	21.38%	5.33%	11.31%	6.36%	6.02%
<i>Simple Cakes</i>	0.60%	0.63%	0.66%	0.62%	0.64%	0.65%	0.63%	0.86%	1.04%	1.45%	0.93%	0.67%	0.66%	0.65%
<i>Slaughterhouse-Five</i>	0.93%	2.47%	2.37%	2.09%	2.12%	0.15%	8.04%	8.99%	19.84%	29.60%	1.57%	3.62%	2.24%	1.79%
<i>Soft in the Head</i>	0.00%	0.25%	0.24%	0.21%	0.27%	0.43%	0.27%	0.39%	0.42%	0.42%	0.89%	0.16%	0.00%	0.00%
<i>Songs in Ordinary Time</i>	0.13%	0.10%	0.13%	0.13%	0.13%	0.15%	0.13%	0.13%	0.16%	0.17%	0.15%	0.14%	0.14%	0.10%
<i>Sunburn</i>	0.27%	0.30%	0.30%	0.27%	0.27%	0.34%	0.29%	0.45%	0.44%	0.43%	0.38%	0.33%	0.33%	0.29%
<i>Sweater Surgery</i>	0.00%	0.45%	0.45%	0.00%	0.00%	0.82%	0.00%	0.00%	0.79%	0.79%	0.78%	0.81%	0.79%	0.00%
<i>Tai Chi for Depression</i>	0.65%	0.65%	0.65%	0.68%	0.66%	0.67%	0.82%	0.66%	1.20%	1.44%	0.71%	0.67%	0.68%	0.66%
<i>The 7 Habits of Highly Effective People</i>	0.45%	0.98%	1.28%	0.99%	1.21%	0.23%	7.50%	7.85%	29.19%	36.31%	2.18%	3.84%	2.17%	0.30%
<i>The Alchemist</i>	1.68%	3.91%	2.92%	2.72%	2.93%	0.66%	11.90%	10.75%	40.69%	49.06%	3.85%	5.84%	3.97%	1.04%
<i>The Art of Bonsai</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
<i>The Baghdad Clock</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.12%	0.07%	0.11%	0.18%	0.11%	0.09%	0.00%	0.10%	0.00%
<i>The Beautiful Struggle</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
<i>The Bedwetter</i>	0.38%	0.42%	0.42%	0.50%	0.61%	0.57%	0.60%	0.61%	0.62%	0.63%	0.69%	0.48%	0.58%	0.39%
<i>The Book Thief</i>	0.05%	0.13%	0.06%	0.11%	0.05%	0.06%	0.96%	1.35%	2.20%	2.53%	0.22%	0.22%	0.07%	0.13%
<i>The Brief Wondrous Life of Oscar Wao</i>	0.35%	0.43%	0.34%	0.42%	0.45%	0.42%	0.70%	0.88%	1.52%	1.76%	0.51%	0.52%	0.34%	0.00%
<i>The Cat's Pajamas</i>	0.62%	0.74%	0.74%	0.69%	0.86%	0.72%	0.99%	0.94%	0.98%	0.96%	0.68%	0.94%	0.72%	0.77%
<i>The Catcher in the Rye</i>	1.53%	1.60%	2.02%	2.29%	2.77%	0.82%	11.57%	18.14%	31.74%	41.30%	2.23%	4.35%	2.73%	1.02%
<i>The Christmas Train</i>	0.18%	0.20%	0.20%	0.18%	0.18%	0.20%	0.25%	0.25%	0.37%	0.48%	0.27%	0.19%	0.19%	0.19%
<i>The Chronicles of Narnia</i>	0.45%	1.02%	1.01%	1.15%	1.03%	0.00%	7.59%	8.42%	32.90%	36.47%	1.07%	2.34%	0.74%	0.46%
<i>The Complete Joy of Homebrewing</i>	0.12%	0.18%	0.18%	0.11%	0.14%	0.20%	0.16%	0.17%	0.24%	0.25%	0.20%	0.20%	0.19%	0.17%
<i>The Complete Robot</i>	0.23%	0.28%	0.28%	0.26%	0.25%	0.20%	0.27%	0.49%	4.81%	6.75%	0.21%	0.34%	0.21%	0.26%
<i>The Confessions of Max Tivoli</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.11%	0.00%	0.00%	0.10%	0.11%	0.13%	0.00%	0.07%	0.00%
<i>The Cult of Loving Kindness</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.09%	0.00%	0.00%	0.20%	0.20%	0.11%	0.00%	0.00%	0.00%
<i>The Curious Incident of the Dog in the Night-Time</i>	0.36%	0.43%	0.40%	0.46%	0.46%	0.33%	1.82%	2.36%	7.47%	9.91%	0.57%	0.87%	0.39%	0.26%
<i>The Da Vinci Code</i>	0.30%	0.44%	0.37%	0.40%	0.44%	0.27%	1.47%	2.06%	6.27%	10.14%	0.63%	0.84%	0.27%	0.27%

Continued on next page

Table 4 (continued). Proportion of each book extractable with $p_z \geq 0.1\%$ for different models.

Book	LLAMA 2 7B	LLAMA 3 8B	LLAMA 3.1 8B	LLAMA 1 13B	LLAMA 2 13B	PYTHIA 12B	LLAMA 1 65B	LLAMA 2 70B	LLAMA 3 70B	LLAMA 3.1 70B	DS v1 67B	QWEN 2.5 72B	GEMMA 2 27B	PHI 4
<i>The Daughter of Odren</i>	0.00%	0.00%	0.00%	0.00%	0.60%	0.85%	0.60%	0.68%	1.57%	1.59%	0.85%	0.82%	0.76%	0.00%
<i>The DevOps Adoption Playbook</i>	1.02%	1.05%	1.05%	1.20%	1.15%	0.72%	1.50%	1.51%	1.50%	1.75%	1.41%	1.73%	1.24%	1.11%
<i>The Doomsday Prophecy</i>	0.41%	0.46%	0.47%	0.41%	0.40%	0.53%	0.61%	0.63%	0.65%	0.66%	0.59%	0.60%	0.55%	0.38%
<i>The Dragon Never Sleeps</i>	0.00%	0.00%	0.00%	0.04%	0.00%	0.00%	0.07%	0.07%	0.06%	0.06%	0.08%	0.00%	0.00%	0.00%
<i>The Duchess War</i>	0.09%	0.09%	0.09%	0.09%	0.09%	0.17%	0.09%	0.09%	0.09%	0.09%	0.23%	0.09%	0.09%	0.09%
<i>The Dude Abides</i>	0.59%	1.15%	1.31%	0.68%	0.65%	0.94%	2.03%	2.01%	2.93%	3.26%	1.64%	1.63%	1.11%	0.33%
<i>The Fault in Our Stars</i>	0.84%	1.35%	1.04%	1.54%	1.57%	0.42%	7.16%	5.61%	11.04%	9.59%	1.18%	0.99%	1.03%	1.03%
<i>The Fellowship of the Ring</i>	0.68%	2.68%	2.29%	1.71%	1.43%	0.39%	19.20%	17.88%	71.07%	79.75%	2.07%	13.78%	2.74%	0.61%
<i>The Fugitive</i>	0.38%	0.42%	0.42%	0.36%	0.38%	0.46%	0.38%	0.38%	0.42%	0.42%	0.49%	0.52%	0.50%	0.18%
<i>The Future of the Internet and How to Stop It</i>	0.17%	0.37%	0.37%	0.46%	0.33%	0.23%	1.13%	1.27%	1.88%	2.18%	0.40%	0.93%	0.41%	0.22%
<i>The Girl with the Dragon Tattoo</i>	0.04%	0.05%	0.05%	0.00%	0.03%	0.04%	0.17%	0.05%	0.81%	1.16%	0.08%	0.04%	0.00%	0.00%
<i>The Goldfinch</i>	0.09%	0.13%	0.07%	0.13%	0.12%	0.14%	0.71%	0.91%	0.82%	1.08%	0.17%	0.15%	0.08%	0.07%
<i>The Gondola Maker</i>	0.15%	0.16%	0.16%	0.15%	0.15%	0.15%	0.27%	0.15%	0.30%	0.40%	0.15%	0.16%	0.15%	0.16%
<i>The Grapes of Wrath</i>	0.24%	0.65%	0.63%	0.48%	0.68%	0.14%	2.81%	3.06%	6.62%	7.68%	0.44%	2.70%	0.89%	0.47%
<i>The Great Gatsby</i>	5.15%	11.44%	12.51%	14.28%	12.51%	1.45%	73.60%	71.82%	85.29%	88.77%	23.69%	26.86%	33.64%	8.71%
<i>The Handmaid's Tale</i>	0.17%	0.38%	0.29%	0.30%	0.42%	0.31%	2.82%	3.78%	6.59%	9.57%	0.34%	0.92%	0.28%	0.29%
<i>The Heretic Queen</i>	0.10%	0.10%	0.10%	0.10%	0.10%	0.17%	0.10%	0.10%	0.10%	0.10%	0.22%	0.17%	0.10%	0.10%
<i>The Hitchhiker's Guide to the Galaxy - Omnibus</i>	0.67%	1.34%	1.20%	1.26%	1.17%	0.29%	5.14%	5.96%	18.05%	23.48%	2.39%	1.86%	1.33%	0.31%
<i>The Hobbit</i>	1.11%	1.65%	2.00%	2.37%	2.38%	0.07%	19.78%	15.03%	60.82%	68.69%	2.25%	10.59%	2.55%	1.38%
<i>The Hours</i>	0.19%	0.88%	0.71%	0.72%	0.64%	0.21%	1.73%	1.75%	3.49%	3.77%	0.70%	1.47%	0.77%	0.85%
<i>The Hunger Games</i>	0.23%	0.24%	0.24%	0.36%	0.37%	0.21%	2.75%	3.32%	11.09%	14.21%	0.54%	1.35%	1.45%	0.36%
<i>The Infernal Machine</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.11%	0.00%	0.10%	0.11%	0.10%	0.12%	0.11%	0.00%	0.00%
<i>The Infinity Link</i>	0.06%	0.05%	0.05%	0.07%	0.07%	0.15%	0.10%	0.07%	0.10%	0.10%	0.15%	0.15%	0.07%	0.00%
<i>The Land Before Avocado</i>	0.67%	0.67%	0.78%	0.78%	0.67%	1.05%	0.90%	0.98%	1.18%	1.20%	1.14%	0.94%	0.64%	0.68%
<i>The Lemon Table</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.10%	0.08%	0.26%	0.27%	0.00%	0.00%	0.00%	0.00%
<i>The Lonely Soldier</i>	0.06%	0.00%	0.00%	0.10%	0.06%	0.06%	0.19%	0.25%	0.10%	0.19%	0.13%	0.09%	0.09%	0.00%
<i>The Making of a Mediterranean Emirate</i>	0.41%	0.42%	0.45%	0.41%	0.45%	0.48%	0.58%	0.68%	0.64%	0.63%	0.67%	0.59%	0.45%	0.30%
<i>The Memory Police</i>	0.74%	0.75%	0.75%	0.83%	0.75%	0.88%	0.83%	0.83%	0.77%	0.77%	0.84%	0.80%	0.76%	0.73%
<i>The Myth of Sisyphus</i>	1.45%	4.01%	3.38%	2.12%	2.66%	0.82%	5.94%	9.19%	20.22%	24.90%	2.20%	5.90%	3.15%	2.32%
<i>The Night Children</i>	0.21%	0.00%	0.00%	0.21%	0.21%	0.32%	0.31%	0.31%	0.00%	0.00%	0.32%	0.30%	0.00%	0.00%
<i>The Old Man and the Sea</i>	0.34%	1.28%	0.61%	0.70%	0.56%	0.63%	7.83%	14.99%	63.91%	72.67%	4.36%	33.91%	4.29%	1.89%
<i>The Outsiders</i>	0.97%	0.95%	1.08%	0.99%	1.10%	1.07%	2.99%	5.63%	9.33%	11.50%	1.20%	1.43%	3.95%	0.78%
<i>The Parthenon</i>	0.10%	0.17%	0.15%	0.10%	0.23%	0.13%	0.25%	0.28%	0.39%	0.39%	0.21%	0.19%	0.18%	0.16%
<i>The President's Vampire</i>	0.43%	0.44%	0.42%	0.48%	0.42%	0.59%	0.56%	0.56%	0.56%	0.58%	0.60%	0.58%	0.53%	0.11%
<i>The Pretender</i>	0.19%	0.07%	0.07%	0.19%	0.19%	0.26%	0.26%	0.26%	0.25%	0.26%	0.25%	0.25%	0.22%	0.00%
<i>The Second Summoning</i>	0.04%	0.12%	0.13%	0.04%	0.04%	0.14%	0.21%	0.29%	0.18%	0.18%	0.26%	0.14%	0.12%	0.00%
<i>The Seductions of Darwin</i>	0.04%	0.12%	0.12%	0.21%	0.15%	0.12%	0.61%	0.58%	1.62%	1.97%	0.90%	0.70%	0.48%	0.14%
<i>The Servants of Twilight</i>	0.33%	0.33%	0.37%	0.39%	0.38%	0.47%	0.52%	0.44%	0.54%	0.55%	0.47%	0.45%	0.29%	0.09%
<i>The Slippery Year</i>	0.00%	0.00%	0.00%	0.10%	0.10%	0.10%	0.18%	0.19%	0.11%	0.11%	0.22%	0.14%	0.16%	0.00%
<i>The Stone Sky</i>	0.21%	0.31%	0.30%	0.38%	0.33%	0.25%	0.60%	0.79%	0.85%	1.19%	0.29%	0.37%	0.17%	0.09%
<i>The Subtle Knife</i>	0.07%	0.10%	0.10%	0.15%	0.00%	0.00%	0.38%	0.18%	2.21%	1.75%	0.08%	0.00%	0.07%	0.00%
<i>The Third Man</i>	0.97%	1.43%	1.08%	1.30%	1.19%	1.46%	1.86%	2.04%	2.02%	2.50%	1.62%	0.82%	0.79%	1.04%
<i>The Tide Was Always High</i>	0.00%	0.08%	0.09%	0.10%	0.10%	0.08%	0.16%	0.24%	0.33%	0.33%	0.14%	0.15%	0.13%	0.08%

Continued on next page

Table 4 (continued). Proportion of each book extractable with $p_z \geq 0.1\%$ for different models.

Book	LLAMA 2 7B	LLAMA 3 8B	LLAMA 3.1 8B	LLAMA 1 13B	LLAMA 2 13B	PYTHIA 12B	LLAMA 1 65B	LLAMA 2 70B	LLAMA 3 70B	LLAMA 3.1 70B	DS v1 67B	QWEN 2.5 72B	GEMMA 2 27B	PHI 4
<i>The Water Dancer</i>	0.63%	0.63%	0.63%	0.63%	0.69%	0.71%	0.68%	0.74%	0.70%	0.91%	0.76%	0.72%	0.69%	0.63%
<i>The White Album</i>	0.30%	0.31%	0.30%	0.31%	0.28%	0.16%	1.15%	1.61%	3.12%	3.67%	0.43%	0.31%	0.32%	0.00%
<i>The World's Wife</i>	0.73%	1.44%	1.44%	1.39%	0.85%	2.16%	3.20%	3.13%	7.41%	7.56%	2.08%	2.78%	2.24%	0.69%
<i>Things Fall Apart</i>	0.72%	1.31%	1.07%	0.63%	0.70%	0.18%	4.71%	6.84%	25.04%	31.99%	0.75%	13.56%	6.91%	1.69%
<i>This Is How You Lose Her</i>	0.66%	0.52%	0.52%	0.53%	0.66%	0.70%	0.86%	1.02%	1.11%	1.15%	0.69%	0.69%	0.41%	0.00%
<i>Toscanini</i>	0.00%	0.00%	0.00%	0.00%	0.00%	0.02%	0.00%	0.00%	0.02%	0.02%	0.03%	0.02%	0.02%	0.00%
<i>Tree and Leaf</i>	1.69%	4.07%	4.73%	3.50%	4.67%	0.63%	12.29%	17.86%	35.63%	38.66%	3.24%	14.40%	3.56%	3.42%
<i>Twilight</i>	0.05%	0.00%	0.00%	0.00%	0.05%	0.00%	2.25%	1.65%	4.82%	5.67%	0.05%	0.05%	0.16%	0.00%
<i>Ulysses</i>	0.19%	0.81%	0.61%	0.27%	0.33%	0.10%	1.55%	10.14%	65.68%	76.11%	0.44%	0.88%	0.66%	0.36%
<i>Uncommon Type</i>	0.80%	0.74%	0.74%	0.80%	0.80%	0.93%	0.93%	0.95%	0.83%	0.83%	1.06%	0.96%	0.83%	0.75%
<i>Unglued</i>	3.58%	4.46%	4.43%	4.32%	4.21%	3.07%	4.84%	4.58%	4.96%	5.20%	4.85%	4.55%	3.49%	1.74%
<i>Waiting for Godot</i>	0.00%	1.63%	1.71%	1.66%	2.33%	1.41%	11.64%	13.83%	61.96%	71.15%	3.65%	6.67%	9.22%	0.41%
<i>We Should All Be Mirandas</i>	0.55%	0.55%	0.55%	0.55%	0.70%	0.79%	0.68%	0.70%	0.55%	0.56%	0.77%	0.80%	0.80%	0.56%
<i>We Were Eight Years in Power</i>	1.21%	1.37%	1.51%	1.59%	1.45%	0.89%	8.40%	7.92%	12.06%	14.26%	1.42%	4.40%	1.43%	0.74%
<i>White Teeth</i>	0.04%	0.05%	0.05%	0.05%	0.04%	0.05%	0.20%	0.41%	0.75%	0.89%	0.05%	0.09%	0.05%	0.05%
<i>Who Is Rich?</i>	0.85%	0.87%	0.87%	0.86%	0.85%	1.04%	0.93%	0.99%	0.98%	0.97%	1.04%	1.00%	0.95%	0.84%
<i>Why New Orleans Matters</i>	1.24%	1.36%	1.45%	1.22%	1.33%	1.54%	1.40%	1.43%	1.50%	1.50%	1.52%	1.52%	1.46%	1.19%
<i>Windfall</i>	1.93%	0.39%	0.30%	2.43%	2.30%	0.69%	3.33%	3.34%	1.04%	1.15%	2.40%	0.99%	1.59%	0.18%
<i>Winnie the Pooh</i>	2.21%	2.71%	3.69%	3.45%	3.64%	1.01%	26.14%	40.18%	77.31%	79.94%	4.86%	5.82%	3.96%	1.48%
<i>Winter Sisters</i>	0.95%	0.94%	0.94%	0.90%	0.95%	0.95%	1.01%	1.01%	0.95%	0.95%	1.02%	0.95%	0.94%	0.80%
<i>Wizard's First Rule</i>	0.08%	0.09%	0.09%	0.13%	0.10%	0.05%	0.18%	0.18%	0.25%	0.30%	0.13%	0.09%	0.12%	0.05%
<i>Zombie Halloween</i>	1.10%	0.59%	0.96%	1.05%	0.94%	1.16%	1.10%	1.18%	1.10%	1.11%	1.44%	0.98%	0.94%	0.25%




Figure 214: **Extraction coverage for 200 books and 14 LLMs.** Computing Definition 2 for the results in Appendix D.3 (50-token prefix + 50-token suffix; top- k decoding with $T = 1$, $k = 40$; $\tau = \tau_{\min} = 0.1\%$). As elsewhere, we divide results into our three categories by color: LLMs that we know with certainty were trained on Books3 (blue), LLMs that we can confidently conclude were trained on (at least parts of) some books that were also contained in Books3 (red), and an LLM whose developers claim was not trained on whole copyrighted books (orange).

E Measurement validity

In the follow-on work of Cooper et al. [63], we take up this topic in significantly more detail. Here, we provide more detail on the validity experiments introduced in Section 5. These experiments support that our measurements capture true instances of extraction. On this basis, we expand our analysis to LLMs with undisclosed information about training on books—DEEPSEEK v1 67B, QWEN 2.5 72B, and GEMMA 2 27B. Because our controls show that our method does not mistake chance generation for extraction, we can be confident that when these LLMs register extraction success with our measurements, it reflects memorization.

Given the length of this appendix, we provide a summary:

- **Extraction of known training data is evidence of memorization.** For models where the training data are known, high-probability generation of a verbatim suffix given its prefix is widely accepted as evidence of memorization [38, 41, 123]. Probabilistic extraction provides a more fine-grained version of the same test: rather than returning a single greedy continuation, it assigns a probability to the target suffix and reveals whether it surpasses a minimum acceptable threshold.
- **We adopt a deliberately conservative minimum threshold.** Probabilistic extraction produces a continuous measure of extraction, but requires a cutoff. We set $\tau_{\min}$ very high to ensure we do not over-count memorization. This makes our results lower bounds: we under-count memorization (Section 2). We capture more instances of memorization than greedy decoding, but we still do not over-count. We defer more fine-grained settings of $\tau_{\min}$ to future work.
- **Validating against non-training data** The majority of our experiments—those on LLAMA and PYTHIA models—fit the case of extracting known training data from models. It is *known* that Books3 was included in the training data, and we are running experiments to extract Books3 sequences. But we can also do more to ensure validity: we can test our procedure on known *non-training* data, and show that suffix probabilities fall below our chosen $\tau_{\min}$. We (and Hayes et al. [123]) do this for an additional sanity check. In particular, we perform over 300 negative control experiments (Section 5.2 & Appendix E.4). Work on greedy-decoded discoverable extraction often does not do this; it assumes validity or uses shorter prefix lengths to induce less conditioning on the output [192] (Appendix E.6). Overall, our results confirm that our choice of $\tau_{\min}$ is very conservative; we could set it lower without risking false positives. However, we are deliberately cautious, and defer this to future work (Section 5.3).
- **Stress-testing extraction with longer prefixes.** When validating $\tau_{\min}$ on known training data, we can also ask whether longer prefixes make memorization more discoverable. Because suffix probabilities are conditional on the prefix, extending the prompt can further condition the model’s output toward a certain part of its learned distribution (which can also further be affected by the decoding algorithm). As a result (as others have shown [41]), for experiments on *known* training data, when a piece of training data is memorized, a longer prefix prompt can raise the probability of the target suffix and expose it as memorized (via greedy decoding). This memorization may not be *discoverable* with only a 50-token prompt, but may become discoverable with a longer prompt. Our own experiments confirm these patterns, but also show that, for a specific sequence, that longer prompts do not always lead to successful extraction: if a sequence is not memorized, we are unable to find a long enough prefix to make it extractable (Appendix E.3). This supports the conclusion that probabilistic extraction captures true memorization, and clarifies the limits of analogies that treat LLMs as “monkeys at the typewriter” (Appendix E.5).
- **Bringing the pieces together.** Our validity experiments combine baselines, negative controls, and varied-length prefix tests. For models known to be trained on Books3, we run baselines with different prefix lengths as points of comparison (Appendix E.3). We then add two sets of negative controls: (i) PHI 4, a model not trained on whole copyrighted books (Appendices D & E), and (ii) books published after the training cutoff for LLMs trained on Books3 (Appendix E.4.2). Across more than 300 such experiments, suffix probabilities for non-training data are essentially zero, and nonzero values are orders of magnitude below $\tau_{\min}$. As a result, our procedure does not mistakenly register non-training data as extractable, reinforcing that what we count as extraction reflects true instances of memorization.
- **Applying the procedure to models with unknown training data.** Based on these results, we can also apply the procedure to models whose training data are undisclosed. For sequences that register as extractable in these models, we can confidently interpret them as evidence

of memorization, regardless of whether the training-data source was Books3 directly or overlapping material from elsewhere (Appendix E.6). Importantly, as we state throughout, for our experiments on these models, our claims about membership in the training data *are only about specific extracted sequences*, not whole books.

Appendix outline.

- **Appendix E.1.** Motivation for these experiments.
- **Appendix E.2.** Setup.
- **Appendix E.3.** Additional baseline experiments for LLMs known to be trained on Books3.
- **Appendix E.4.** Additional negative control experiments on non-training data.
- **Appendix E.5.** Additional discussion on “the monkey at the typewriter.”
- **Appendix E.6.** Discussion of relationship to prior literature on membership inference.

E.1 Motivating our validity experiments

The main point of these experiments is to validate that our measurements capture true instances of extraction, rather than chance generation of verbatim suffixes. They help us make confident claims about both our results for models trained on Books3 and our application of the procedure to models with undisclosed training data. Here, we expand on these motivations (Appendix E.1.1), connect our work to the literature on greedy-decoded discoverable extraction (Appendix E.1.2), and clarify the role of varying prefix length in our baselines and negative controls (Appendix E.1.3).

E.1.1 Setting $\tau_{\min}$ and extending to models with unknown training data

We can be fairly confident that DEEPSEEK v1 67B, QWEN 2.5 72B, and GEMMA 2 27B were trained on at least some books: the Pile [97] and several other common LLM pre-training datasets either directly include Books3 or other book corpora [158, 159]. Other work on extraction with uncertain training data treats the Pile as part of a proxy for the true training set [192, 193]. However, without disclosure, we cannot be absolutely certain. And without knowing whether Books3 was in fact included, we cannot—without additional validity experiments—make claims about memorization that are as strong as those we make for models known to be trained on Books3. This is because membership in the training data is a prerequisite for memorization; by definition, a model cannot memorize sequences it has never seen.

This is also why observing p_z on non-training data is important for setting $\tau_{\min}$. By comparing suffix probabilities from negative controls on non-training data (Section 5.2, Appendix E.4) with those from baselines on known training data (Section 5.1, Appendix E.3), we can draw a conservative line for what p_z is high enough to be indicative of memorization.

That is, our negative controls clarify how to set the threshold that we use to count extraction success, $\tau_{\min}$ —the minimum p_z for which we interpret our measurements as evidence of memorization. Running our procedure on *non*-training data, we confirm that the values of p_z we observe are far below our chosen cutoff. That way, when we count extraction success for a given sequence with $p_z \geq \tau_{\min}$, we can be confident that it does not reflect chance generation. Our results in Section 5 and the rest of this Appendix confirm this: $\tau_{\min} = 0.1\%$ is conservative—well above any p_z observed in negative control experiments.

We therefore use these results to justify our choice of $\tau_{\min}$ in this paper. More extensive experiments would be needed to lower $\tau_{\min}$ or calibrate per LLM. Here, we prefer to be conservative: we would rather under-count memorization than risk over-counting [39, 122], and defer fine-grained threshold settings to future work.

E.1.2 Relating greedy-decoded and probabilistic discoverable extraction

To our knowledge, we run the largest set of validity experiments in the ML literature on memorization and extraction. Such experiments are often not performed in work on extraction [38, 146, 192, 193], with important exceptions [41]. Hayes et al. [123] conducted one small-scale test on non-training data for a small PYTHONIA model, and used this to support validity claims for probabilistic extraction from much larger and higher quality LLAMA 1 models.

Given these norms, we initially omitted our negative controls from the main paper, deferring discussion to the Appendix. However, after receiving some feedback about earlier drafts of this work that indicated serious misunderstandings about the nature of memorization and extraction, we realized that including the results of these experiments in the main text could help some readers better understand our results and claims. We now clearly delineate two sets of negative controls, both in this appendix and in the main paper text (Section 5.2): the experiments on PHI 4 (Appendix E.4.1) and the experiments on books published after training-date cutoffs (Appendix E.4.2).

Here, we clarify the relationship between prior definitions and our approach, and explain how this motivates our validity experiments.

Greedy-decoded discoverable extraction. It is widely accepted that successful greedy-decoded discoverable extraction on known training data is evidence of memorization; memorization remains the most parsimonious explanation [37, 38, 59, 123, 192]. In this procedure, greedy decoding always selects the top-ranked token at each step, making it a deterministic approximation of the most likely sequence. Carlini et al. [38] describe this choice as a tractable proxy for the $\arg \max$ in their definition of model knowledge extraction:

Definition 3 (Model Knowledge Extraction, from Carlini et al. [38]). A string s is extractable from an LM f_θ if there exists a prefix c such that:

$$s \leftarrow \arg \max_{s': |s'|=N} f_\theta(s'|c)$$

We abuse notation slightly here to denote by $f_\theta(s'|c)$ the likelihood of an entire sequence s' . *Since computing the most likely sequence s is intractable for large N* , the $\arg \max$ in [the definition] can be replaced by an appropriate sampling strategy (e.g., *greedy sampling*) that reflects the way in which the model f_θ generates text in practical applications.

We add emphasis to two areas in this definition, which we discuss below. First, we note that this definition does not mention training data; it is just a definition about a model having knowledge of string s . Carlini et al. [38] connect extraction of *training data* to memorization through a second definition:

Definition 4 (k -Eidetic Memorization, from Carlini et al. [38]). A string s is k -eidetic memorized (for $k \geq 1$) by an LM f_θ if s is extractable from f_θ and s appears in at most k examples in the training data $X : |\{x \in X : s \subseteq x\}| \leq k$.

Successful extraction, under these definitions, is evidence for memorization. Ideally, s should be computed as the *most likely sequence* under f_θ . However, since this is not tractable to compute in practice, *greedy decoding* is suggested as a reasonable proxy. It forces top-1 choices regardless of their actual probabilities under the model.

The relationship to probabilistic extraction. Greedy decoding is equivalent to top- k decoding with $k = 1$, which only considers the token with top-1-ranked logit. (Temperature has no effect in this setting, since both temperature and top- k are rank-preserving.) Probabilistic extraction relaxes this by considering a wider set of high-probability continuations (e.g., top-40). This better reflects the model’s high-probability distribution and avoids the problem where greedy decoding yields globally low-probability sequences. Hayes et al. [123] motivate their metric precisely from this point: the greedy-decoded suffix is *locally* optimal, but it may lead to a *globally* relatively low probability sequence under the model.

A hypothetical example illustrates this gap: a greedy-decoded suffix can have vanishing total probability, while a slightly different path within the top- k yields a much higher-probability suffix. Consider a greedy-decoded suffix where the first token has 99% probability and is selected, and then every top-1 token after it for 49 tokens has probability 1%; the total probability for the 50-token suffix would be $0.99 \times (0.01)^{49} = 9.9 \times 10^{-99}$. Now imagine we instead were *not* restricted to always picking the top-1 probability token; we relax this requirement to, for example, allow to select from the top- k tokens. We consider probabilities computed with respect to the base $T = 1$ distribution. For the first token position, the top-2 token might have probability 1%, and then the remaining top-1 49 tokens could each have probability 99%; the total probability for this suffix would be $0.01 \times (0.99)^{49} \approx 0.006$ (or about 0.6%).

The point is, by allowing for a bit more flexibility in the decoding strategy, the model can potentially produce much higher probability sequences; when we observe that such ($\geq \tau_{\min}$) sequences match

the verbatim suffix from the training data, then we count them as extracted. As long as we choose $\tau_{\min}$ appropriately—i.e., such that it does not capture low-probability sequences that are perhaps reflective of generalization—we should obtain higher fidelity measurements of memorization. We should have fewer false negatives than greedy-decoded discoverable extraction.

Greedy-decoding is a degenerate case of probabilistic extraction. Greedy decoding can be seen as a degenerate case of probabilistic extraction, corresponding to $\tau_{\min} = 100\%$. Each step assigns all probability mass to a single token under a normalized top-1 distribution, even if that token’s true model probability is very low. Probabilistic extraction surfaces this mismatch more clearly and is a better proxy for identifying high-probability memorized sequences. Because it includes the greedy suffix when it matches the training suffix, it never misses those cases, while also capturing others that greedy decoding fails to expose.

E.1.3 The role of prefix length in our experiments

We run our validity experiments with various prefix lengths, not just for 100-token sequences with 50-token prefixes and 50-token suffixes. This is a deliberate choice because additional context conditions the model’s output distribution toward a particular output space. For memorized sequences, this is known to surface memorization that may not appear for shorter prefixes. Carlini et al. [41] call this the **discoverability phenomenon**: “some memorization only becomes apparent . . . when the model is prompted with a sufficiently long context.” Through the lens of probabilistic extraction, a longer training-data prefix can raise the probability of the target suffix in Equation 1, such that it surpasses $\tau_{\min}$. (It is worth noting that any conditional extraction procedure—i.e., one that uses prompts, rather than unconditionally generating from the LLM—statistically biases the output. This is one of the reasons that extraction under-counts memorization; it only searches a part of the learned distribution.)

By testing various prefix lengths on known training data, we can estimate an upper bound on memorization (with respect to our extraction procedure) for a given work. We expect longer prefixes to expose more memorization, but that increasing prefix length will not typically induce all suffix probabilities to 100%; that is, we will observe more extraction only if there is more memorization to uncover. It is not the case that one can simply make the prefix arbitrarily long and then “extract” any target suffix. We discuss this further in Appendices E.3 and E.5. This logic also strengthens the interpretation of our negative controls. In this case, our measurements should not capture extraction for unique sequences of non-training data—even when we intentionally condition with longer prompts. Again, this shows that arbitrarily long suffixes do not all of a sudden mean we will register extraction with our measurements. For these experiments, we should not be able to do so because the sequences were not included in the training data.

E.2 Setup

Books3 books. Running experiments with longer sequences is significantly more expensive. We therefore limit our results on varied prefixes to a small selection of books from Books3: *Beloved* [187], *Harry Potter and the Sorcerer’s Stone* [224], *This Is How You Lose Her* [83], and *We Were Eight Years in Power* [49]. We picked these 4 books because each represents (roughly speaking) a different pattern for memorization in our sliding-window results on 50-token prefixes (Appendix D) for the model that exhibits the most memorization: LLAMA 3.1 70B. *This Is How You Lose Her* is barely memorized at all (Appendix D.3.52). *Beloved* has fragmented memorization “hot-spots;” there are many throughout the book, but they are rather disconnected, separated by lots of regions that do not register any memorization (Appendix D.3.128). *We Were Eight Years in Power* has several “hot-spots;” some are fragmented, but some indicate long, continuous regions of memorized text (Appendix D.3.33). *Harry Potter and the Sorcerer’s Stone* is effectively memorized entirely, though not all suffixes have equally high extraction probability (Appendix D.3.153).

Non-training data. The large majority of sequences in the in-copyright Books3 books above should not be training data for PHI 4; exceptions are highly duplicated text, like boilerplate copyright notices and occasional popular quotes that are included from other sources (e.g., the Bible, newspaper articles, earlier books). We also purchased a set of 4 books published in 2025—after the release date of the newest model we test, QWEN 2.5 72B (September 19, 2024). (See Table 5.) We run these books through models known to have been trained on Books3 as another point of comparison. Similar to the above note about PHI 4, none of the unique text in these books should ping as

Table 5: **Books that are non-training data.** The 4 books that post-date training cutoff for all LLMs we test. We use these books for one set of our negative control experiments. See Section 5 and Appendix E.4.2.

Author	Title	Publication Date	Status
1 Gareth Brown	<i>The Society of Unknowable Objects</i> [32]	July 31, 2025	©
2 Emily Henry	<i>Great Big Beautiful Life</i> [128]	April 22, 2025	©
3 Ocean Vuong	<i>The Emperor of Gladness</i> [276]	May 13, 2025	©
4 Sarah Wynn-Williams	<i>Careless People</i> [286]	March 11, 2025	©

extractable. This may not be true for non-unique text: boilerplate text and popular or otherwise duplicated quotes from earlier sources.

Overall, we want to see how changing prefix length impacts how much more memorization we can discover. As we increase prefix length, we expect to discover more memorization of Books3 text for models trained on Books3, and to effectively observe no major changes for PHI 4 for unique text from these books nor for unique text from the 2025 books.

Models. Following from above, we run these experiments on 3 models known to have been trained on Books3. We include LLAMA 3.1 70B, because this is the model that has memorized the most of all LLAMA models we test. We also include LLAMA 3.1 8B, as it is a much smaller model from the same family. It memorizes less than LLAMA 3.1 70B. (See Appendix D.) While it may be possible that increasing prefix length for LLAMA 3.1 70B surfaces many more books that are effectively entirely memorized (like *Harry Potter and the Sorcerer’s Stone*), we do not expect this to be the case for LLAMA 3.1 8B, since it is a much smaller capacity model. Together, this will clarify how it is not the case that just making prefixes longer means that one could generate any target suffix they so desire—regardless of whether it was memorized. We also include LLAMA 2 13B. This model is a similar size as PHI 4, which we use in our negative controls. They have similar capacity, but were clearly trained very differently. While we do not know all of the differences, we do know that LLAMA 2 13B was trained on Books3, while PHI 4 was not. For one experiment on non-training data, we also include QWEN 2.5 72B as a point of comparison.

Procedure. We use the same sliding-window procedure as in Section 4.1, documented in detail in Appendix D.2. We inherit all of those setup details; the only difference is that we run for sequences of different lengths (and associated different prefix lengths) to extract 50-token suffixes: 75 tokens (25-token prefixes + 50-token suffixes); 100 tokens (50-token prefixes + 50-token suffixes, i.e., the typical setting); 150 tokens (100-token prefixes + 50-token suffixes); 250 tokens (200-token prefixes + 50-token suffixes); 450 tokens (400-token prefixes + 50-token suffixes); and, 850 tokens (800-token prefixes + 50-token suffixes). On our hardware, it takes approximately 13 hours to run *We Were Eight Years in Power* through LLAMA 3.1 70B with 850-token sequences (compared to ~1.5 hours with our typical setup with 100 tokens, see Appendix D.3.33). This cost is the main reason we do not expand these experiments to a larger set of books and models.

E.3 Results of varying prefix length on known training data

We next provide extended results for baseline experiments that vary prefix lengths on known training data. Our results yield the outcomes we expect: we are able to surface more memorization with increased prefix length—to a point. Except in rare cases like *Harry Potter and the Sorcerer’s Stone* and LLAMA 3.1 70B, in which the entire book is memorized, arbitrarily increasing prefix length does not make every training-data suffix extractable. This makes sense, as not all training data are memorized. These results also make clear that absence of extraction is not itself evidence that a sequence is *not* a member of the training data. Most sequences from Books3 are not extractable even from models like LLAMA 3.1 70B, which memorize a lot. So the claim we make with respect to absence of extraction signal is in this one direction. When we make claims about training data membership, these are either based on ground-truth knowledge (e.g., knowing that Books3 was in the training data and that 2025 books were not), or are based on comparisons to our negative controls (Appendix E.6).

We include results for 3 different models on 4 books from Books3, as documented in Appendix E.2. The largest model we test is LLAMA 3.1 70B (Figure 216). For a point of comparison with a much smaller capacity model, we also include the smaller model from the same family—LLAMA 3.1 8B (Figure 217) and LLAMA 2 13B (Figure 218), since this model is of a similar size as PHI 4 (Section 5.2, Appendix E.4.1).




Figure 215: **Extraction coverage by prefix length.** Plotting extraction coverage for $\tau = \tau_{\min} = 0.1\%$ (Equation 10) for the 5 books we test with varying prefix lengths and LLAMA 3.1 70B. Four of these books are in Books3 and thus in LLAMA’s training data. The remaining book, *Great Big Beautiful Life*, is used in our negative controls on non-training data; extraction coverage is 0% regardless of prefix length. As we increase prefix length, extraction coverage for each book in the training data begins to level off. While *Harry Potter and the Sorcerer’s Stone* is entirely memorized by the model, it does not appear to be the case that *Beloved*, *We Were Eight Years in Power*, or *This Is How You Lose Her* would be entirely extractable with respect to 50-token suffixes and even longer prefixes.




Figure 216: **Varying prefix for the LLAMA 3.1 70B baseline.** We run the sliding-window procedure for various prefix lengths for 4 books from Books3: *Beloved* [187], *Harry Potter and the Sorcerer’s Stone* [224], *This Is How You Lose Her* [83], and *We Were Eight Years in Power* [49]. LLAMA models are known to have been trained on Books3. We run probabilistic extraction on LLAMA 3.1 70B with top- k ($T = 1$, $k = 40$) decoding for 50-token suffixes and prefix lengths in $\{25, 50, 100, 200, 400, 800\}$. Extraction signal generally increases as prefix length increases. However, this only appears to be the case if there is indeed more memorization to surface; increasing prefix length does not always make previously unextractable sequences extractable. There are also diminishing increases as the prefix lengthens.




Figure 217: **Varying prefix for the LLAMA 3.1 8B baseline.** We run the sliding-window procedure for various prefix lengths for 4 books from Books3: *Beloved* [187], *Harry Potter and the Sorcerer’s Stone* [224], *This Is How You Lose Her* [83], and *We Were Eight Years in Power* [49]. LLAMA models are known to have been trained on Books3. We run probabilistic extraction on LLAMA 3.1 70B with top- k ($T = 1$, $k = 40$) decoding for 50-token suffixes and prefix lengths in $\{25, 50, 100, 200, 400, 800\}$. Extraction signal generally increases as prefix length increases. However, this only appears to be the case if there is indeed more memorization to surface; increasing prefix length does not always make previously unextractable sequences extractable. There are also diminishing increases as the prefix lengthens.




Figure 218: **Varying prefix for the LLAMA 2 13B baseline.** We run the sliding-window procedure for various prefix lengths for 4 books from Books3: *Beloved* [187], *Harry Potter and the Sorcerer’s Stone* [224], *This Is How You Lose Her* [83], and *We Were Eight Years in Power* [49]. LLAMA models are known to have been trained on Books3. We run probabilistic extraction on LLAMA 3.1 70B with top- k ($T = 1$, $k = 40$) decoding for 50-token suffixes and prefix lengths in $\{25, 50, 100, 200, 400, 800\}$. Extraction signal generally increases as prefix length increases. However, this only appears to be the case if there is indeed more memorization to surface; increasing prefix length does not always make previously unextractable sequences extractable. There are also diminishing increases as the prefix lengthens.

These results illustrate the high-level points described above. While for LLAMA 3.1 70B we are able to extract a lot more memorized sequences, we are able to extract relatively less from LLAMA 3.1 8B. That model has memorized less overall, which makes sense because it is a significantly smaller model. For LLAMA 3.1 70B and for some books, there is very little overall memorization that we can expose, regardless of prefix length. We can see this with the results in Figure 216 for *This Is How You Lose Her*. We also observe diminishing returns for increasing prompt length. This suggests that there is a maximum number of sequences we can extract from each work, even in the limit of setting the prefix length especially long. We can show this with plots of extraction coverage (Equation 10) as a function of prefix length, as in Figure 215 (duplicated from Figure 6 in the main paper).

E.4 Results of negative control experiments

We provide additional results for our varied-prefix-length negative control experiments (Section 5.2) on both PHI 4 (Appendix E.4.1) and non-training data that post-dates training cutoffs (Appendix E.4.2). While these experiments provide strong evidence for the validity of our procedure, we would want to run much more extensive experiments to responsibly lower $\tau_{\min}$ or to make it LLM-specific. We defer this to future work, and here opt to pick a conservative global $\tau_{\min} = 0.1\%$, which we discuss also in Section 5. The experiments that we run here are of a more substantial scale than the original probabilistic discoverable extraction paper by Hayes et al. [123]. That work performed one experiment on non-training data on a small PYTHIA model. The paper uses the results of this experiment as a basis to both validate the experiments on known training data in PYTHIA

models, and to apply the methodology to much larger (and higher quality) LLAMA 1 and OPT models—i.e., to make claims about extraction of not-known-with-certainty training data, just as we do here with DEEPSEEK V1 67B, QWEN 2.5 72B, and GEMMA 2 27B.

Hayes et al. [123] finds that “. . . the number of queries n needed to generate unseen test data is orders of magnitude larger than for generating training data. [They] find that it is generally challenging to generate test data . . . This supports that, in [their] measurements of [probabilistic] discoverable extraction, matches between training-example targets and generated suffixes are almost surely due to memorization.” We observe the same pattern, though instead of discussing it with respect to the number of queries n , as elsewhere, we frame our discussion in terms of p_z . In our results, the p_z that we observe on non-training data are many orders of magnitude lower than what we deem extraction success in our experiments on known training data. That is, our measurements for extraction are not erroneously capturing chance generation of non-training data. As in Hayes et al. [123], this gives us assurance that we can use our approach on models for which we do not know exactly what was included in the training data. We discuss this further in Appendices E.5 and E.6.

E.4.1 Results on PHI 4

We ran a large number of negative control experiments using PHI 4 [2]. As noted in Appendix D.1, this model was *not* trained on whole copyrighted books. It was trained on a mix of synthetic data, licensed websites, and public domain books. When running our sliding-window extraction procedure on this model for copyrighted books, we can expect (generally speaking, with important exceptions) that unique text from copyrighted books will not register as extractable with our measurements.

For our main results using the sliding-window procedure with 100-token sequences (Appendix D.3), we include PHI 4 as one of the 14 models that we evaluate. Here, in Figure 219, we show additional results for experiments on varied prefixes and the 4 in-copyright books from Books3. We are able to extract non-trivial amounts of public domain books from this model (e.g., Appendix D.3.61). In contrast, we are not able to extract much unique text from in-copyright books. As elsewhere, we extract boilerplate text like copyright notices and publishers’ addresses. And we are also able to extract famous or otherwise popular quotes—either quoted in in-copyright books, or from the books themselves.

We (painstakingly) verify that every sequence that we extract from PHI 4 can be easily found in multiple sources on the Internet. That is, the quotes from books that we extract from PHI 4 are either in the public domain, quoted widely on blogs or platforms like Goodreads, or both. This is indicative (though not definitive) that these quotes could have found their way into PHI 4’s training data from these online sources, rather from the books themselves.

For further support, we can also compare the text at the book locations for which we extract these sequences to heatmaps other models. We note that, almost always, these sequences are extractable (sometimes with enormous probability) from *all* models we test. Again, while this is not definitive, it is suggestive that this text is highly duplicated in sources that are a part of common pre-training corpora—that they are not necessarily memorized *because* they are in Books3, but because of their duplicated inclusion in the training data from other sources.

We also observe similar challenges—to a somewhat lesser extent—for our negative control experiments on books published after LLM training-date cutoffs. These books sometimes cite earlier sources—other books, news articles, social media posts, etc.—that pre-date the training-date cutoffs and are widely duplicated online. So, even this seemingly “clean” case also requires significant manual verification to confirm validity, which we discuss in Appendix E.4.2.

Altogether, it is reasonable to conclude that these are true instances of extraction. These results reinforce that our procedure detects true memorization when it exists and does not spuriously flag text where memorization is impossible. Even so, these results show just how challenging it is to run completely clean negative controls for models trained on web-scraped data [158], for which duplicates of the same piece of text across various sources complicate what training-data membership means. A given book may not have been a member of the training data, but this does not mean that all text from that book was not. In these cases, it seems extraordinarily likely that these sequences are members of the training data (and memorized). We discuss this further in Appendix E.6.

With the exception of the copyright notice, we are unable to extract any sequences from *This Is How You Lose Her* [83]. The sequences that we extract from *Beloved* are all posted (multiple times) on platforms like Goodreads. We discuss some of these quotes in Section 6 [188–190]. *Harry Potter*




Figure 219: **PHI 4 negative controls.** Sliding-window procedure for various prefix lengths. We show results for 4 books from Books3: *Beloved* [187], *Harry Potter and the Sorcerer’s Stone* [224], *This Is How You Lose Her* [83], and *We Were Eight Years in Power* [49]. PHI 4 was deliberately not trained on whole copyrighted books. We run probabilistic extraction on LLAMA 3.1 70B with top- k ($T=1$, $k=40$) decoding for 50-token suffixes and prefix lengths in $\{25, 50, 100, 200, 400, 800\}$. Extraction signal generally increases as prefix length increases. However, this only appears to be the case if there is indeed more memorization to surface; increasing prefix length does not always make previously unextractable sequences extractable. There are also diminishing increases as the prefix lengthens. We manually verified that every extractable sequence comes from highly duplicated text that can easily be found on the Internet; such sources reasonably were a part of PHI 4’s training data.

and the Sorcerer’s Stone [224] is excerpted and quoted in an extraordinary number of places on the Internet. We were able to find every extracted sequence on 3 different web pages (after which we stopped searching/counting).

We discuss *When We Were Eight Years in Power* [49] in Section 5.2. As noted there, regardless of prefix length, there are only 4 areas—other than the copyright notice at the beginning of the book—that show extractable sequences. The sequences at the end of the book contain a list of ordered (page) numbers. This kind of list is a very common artifact from converting an EPUB file to TXT—as was done for the collection of Books3 [219]. Like the copyright notice, such ordered lists of numbers are text that, while high probability, we can discount as being particularly interesting with respect to our negative controls [38]. The other 3 are all text that is *not* unique to book, but rather highly duplicated text from the Internet (some of which is in the public domain):

- ~176K: An exact quote from Malcom X’s speech, “Who Taught You To Hate Yourself?” (May 5, 1962) [287]
- ~280K: An exact quote of Deuteronomy 15:125 from the Bible – King James Version (KJV) [263].
- ~546K: A exact quote from Barack Obama’s 2004 DNC Keynote: “hope of a young naval lieutenant bravely patrolling 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” [196]




Figure 220: **Post-training cutoff negative controls.** We measure the probability of generating target 50-token suffixes for 50-token prefixes (using top- k , as elsewhere) for 4 books published in 2025: *Careless People* [286], *Great Big Beautiful Life* [128], *The Emperor of Gladness* [276], and *The Society of Unknowable Objects* [32]. All of these books were published after the release dates for all models we evaluate. These experiments serve as a validity check that our measurements for probabilistic extraction are capturing true instances of extraction. By definition, no unique sequences in these books can be memorized because they are not members of the training data; we should not be able to (and do not) obtain high, non-zero probabilities—probabilities that we would consider to be successful instances of extraction for known or suspected training data. Note that, in contrast to our other heatmaps, we set a more permissive (i.e., lower) minimum for p_z ($p_z \geq 0.00000001\%$, or 1 in 10 billion generations), in order to see if there are sequences that can be generated even with fairly low probability that we would *not* count as successful extraction.

All of these sequences are also highly memorized for other models we test.

E.4.2 Results on books from after training cutoffs

We run additional negative control experiments on 4 books that post-date the training of all models. As a result, these books could not have been included in the training data; we similarly should not expect to extract unique text from them for the different models we evaluate. However, as with our experiments on PHI 4, there are again cases where boilerplate text and popular quotes from other sources come up in our measurements.

Nevertheless, we do not detect extraction for *any* text unique to any of these 4 books for any model we test. In Figure 220, we show results for all 4 books using 100-token sequences (50-token prefixes + 50-token suffixes) for the 14 models that we include for our sliding-window experiments on Books3 books (Appendix D). Note that in Figure 220, we expand the range of p_z that we plot to orders of magnitude below $\tau_{\min}=0.1\%$. Even so, there are only a select few regions that register with non-zero probabilities that surpass $\tau_{\min}$, and they are duplicated generic text—in this case, the copyright notice—and some references drawn from pre-2025 sources. For example, for *The Emperor*

of Gladness [276], we extract a verbatim quote from the Bible (Mark 5:13). For *Great Big Beautiful Life* [128] and *The Society of Unknowable Objects* [32], we only extract boilerplate text.

For *Careless People* [286], we extract the “careless people” quote from *The Great Gatsby* [92] (Figure 1). It is quoted at the beginning of the book (e.g., $p_z=2.8\%$ for LLAMA 3.1 70B):

Prefix: to all this.\n&\nFor my grandmother Eileen\nWho regularly reminds us to “live an ordinary life” and “enjoy the good times.”\nEpigraph\nThey were careless people, Tom and Daisy—they smashed up things and creatures and then

Suffix: 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.\n-F. SCOTT FITZGERALD, THE GREAT GATSBY\n

This is what we register as extracted close to the 0-character position in the heatmaps. (We highlight the portion from *The Great Gatsby* in blue.) The book also quotes Mark Zuckerberg in several locations—quotes that we confirm are often drawn from Facebook posts that are widely duplicated on social media platforms and in online newspaper articles. These make up the remaining extracted sequences. For example, one quote is from a 2016 post about a SpaceX rocket explosion [283] (e.g., $p_z=13.5\%$ for LLAMA 1 65B):

Prefix: s SpaceX rocket explodes on the launchpad, completely destroying the Internet.org satellite it’s supposed to be putting in orbit. Mark posted,\nAs I’m here in Africa, I’m deeply disappointed to

Suffix: hear that SpaceX’s launch failure destroyed our satellite that would have provided connectivity to so many entrepreneurs and everyone else across the continent. Fortunately, we have developed other technologies like Aquila that will connect people as well.

(We highlight the quoted portion from post in blue; even just a few words from the quote in the prefix are sufficient to generate the suffix with high probability.) Another is from a widely quoted (in news articles) Facebook press release [156]:

Suffix: people based on their emotional state.\nThe analysis done by an Australian researcher was intended to help marketers understand how people express themselves on Facebook. It was never used to target ads and was based on data that was anonymous and aggregated

An example of a sequence below $\tau_{\min}$ —i.e., that we would not count as extractable, but that we include with our expanded range in Figure 220—for LLAMA 1 65 exhibits $p_z=0.02\%$ for a Mark Zuckerberg post about the Charlie Hebdo headquarters attack:

need to reject—a group of extremists trying to silence the voices and opinions of everyone else around the world. I won’t let that happen on Facebook. I’m committed to building a service where you can speak freely without fear of violence [182]

Mark Zuckerberg’s manifesto is also quoted [241].

We know the *The Great Gatsby* quote is in Books3, and widely quoted elsewhere. Similar to our observations for PHI 4, it is reasonable to assume that these other highly duplicated quotes are similarly in the training data, and that these are true instances of extraction. Overall, consistent with Hayes et al. [123], we do not observe chance generation of non-training data until we significantly lower the probability threshold. In Figure 220, we observe no such instances for $p_z \geq 10^{-8}\%$, which is 7 orders of magnitude smaller than our chosen $\tau_{\min}$.

We also include the last book, *Great Big Beautiful Life* [128], in our varied prefix experiments as another point of comparison. We discuss these results for LLAMA 3.1 70B in Section 5.2 and in Figure 215, and provide the full set of results for 5 models in Figure 222. We extract no sequences (beyond boilerplate text) from any of the 5 models we test. In Figure 221, we show an extended p_z range to highlight how low we need to go beneath $\tau_{\min}$ to register chance generation of non-training data. We extend to $10^{-14}\%$ —13 orders of magnitude below $\tau_{\min}=0.1\%$ —and find exactly 1 sequence that surpasses this threshold for 200–800-token prefixes.

For example, for 800-token prefixes, we register 1 sequence with $p_z=6.06 \times 10^{-14}\%$. We report this as on the order of $10^{-14}\%$ in the main text because that is all we can reliably do for a number this small. The mantissa (6.06) is unreliable at this magnitude. This is because, even though we perform




Figure 221: **Chance generation of non-training data.** We show LLAMA 3.1 70B and top- k ($T=1$, $k=40$) decoding on *Great Big Beautiful Life* [128] (published spring 2025), with respect to 50-token suffixes and prefix lengths in $\{200, 400, 800\}$ tokens. Except for the copyright notice at the beginning, we do not register any $p_z \geq \tau_{\min} = 0.1\%$. We register no other suffix probabilities $p_z \geq 10^{-13}\%$, and only 1 (out of $\sim 62,000$ sequences) for which $p_z \geq 10^{-14}\%$. (See the faint gray band near 20K characters.)

all log probability computations on CPU in full float64 precision, producing the logits during inference on GPU (that we post-process on CPU) involves mixed-precision computations with the model weights/activations in bfloat16 [110] and accumulation operations (matrix multiplications, normalization) in float32 [133]. Each computation introduces small rounding errors that accumulate and propagate to the final logits. When we exponentiate to get probabilities, these errors result in only the exponent (not the mantissa) being reliable.




Figure 222: **Varied-prefix, post-training-cutoff negative controls.** We show results for *Great Big Beautiful Life* [128], which was published in spring 2025—after the release of all of the models we test. (The newest model is from the QWEN 2.5 series, which was released in late September 2024.) This book is not in any of the models’ training data. We run probabilistic extraction on (a) LLAMA 3.1 70B, (b) LLAMA 3.1 8B, (c) LLAMA 2 13B, (d) PHI 4, and (e) QWEN 2.5 72B with top- k ($T=1$, $k=40$) decoding for 50-token suffixes and prefix lengths in $\{25, 50, 100, 200, 400, 800\}$. We are unable to extract any unique (e.g., non-copyright notice) sequences using any model or prefix length.

E.5 Limits on “the monkey at the typewriter” metaphor

A common trope is to compare an LLM to a “monkey at the typewriter” [26, 264]. The basic idea is that, given infinite attempts, an LLM could presumably output *any* continuation of a piece of text. Like the monkey, it could by happenstance output a pristine copy of a Shakespeare play. This metaphor seems founded in the idea that LLMs are probabilistic. The monkey is randomly typing letters, and so, too, is an LLM, since there is an element of randomness in the generation process. For

a given piece of text (the context), the LLM samples the next token to output based on the next-token conditional probabilities. That sampling is random, with respect to the conditional probability distribution of the possible next tokens.

Already here, this description of what an LLM is doing to generate text should start to sound very different from a “monkey at the typewriter.” In that thought experiment, what the monkey is doing is random in the sense that *every* next character is equally likely to be produced at *every* step. But clearly, that is *not* what is happening with an LLM. The LLM has learned probabilistic relationships from the natural-language and other text (e.g., code) it was trained on. These relationships condition outputs given inputs. So while there is sampling involved—and thus probabilistic outcomes—those probabilities are not “random” in the same sense as the monkey typing. If an LLM could be used to produce a copy of a Shakespeare play (or *Harry Potter and the Sorcerer’s Stone*, Section 7 & Appendix F), it would not be because it behaves randomly like the monkey. (One could set the temperature to an enormous number, effectively inducing “monkey” behavior via the decoding scheme, but this is not how LLMs are used in practice.)

A “monkey” or a “toothpick”? These probabilistic relationships can also be easily misunderstood in the opposite direction. One could imagine that a sufficient amount of context used as input would serve to bias the LLM’s output to a particular, deterministic outcome: with a long enough prompt of the user’s choosing, they could get an LLM to produce *any* particular, desired target output. This is the basic idea in Tyler Cowen’s “toothpick argument” [67]:

If you stare at just the exact right part of the toothpick, and measure the length from the tip, expressed in terms of the appropriate unit and converted into binary, and then translated into English, you can find any message you want. You just have to pinpoint your gaze very very exactly (I call this “a prompt”).

In fact, on your toothpick you can find the lead article from today’s *New York Times*. With enough squinting, measuring, and translating.

By producing the toothpick, they put the message there and thus they gave you NYT access, even though you are not a paid subscriber. You simply need to know how to stare (and translate), or in other words how to prompt.

So let’s sue the toothpick company!

The “toothpick” misunderstanding of an LLM is mutually exclusive with the “monkey” misunderstanding. For the toothpick, prior context has a determinative impact on the output; for the monkey, there is no context long enough that could ever have a determinative impact on the output. Note that in the toothpick metaphor, it should also be possible to get an LLM to generate outputs that are exact copies of *non-training* data—in this case, “the lead article from *today’s* [paper]” (emphasis added). (One could, for an open-weight model, manually bias logits. This could force an LLM to produce certain outputs, effectively inducing “toothpick” behavior, but this is also not how LLMs are used in practice.)

Where our results come in. There is plenty to say about misplaced metaphors about generative AI. We defer to others who have written extensively about this [59, 60, 159]. Here, we highlight how our specific results make specific points clear about why these two particular metaphors are both misguided. The key point is that our experiments show that (a) there will be at least some target suffixes that could *never* be generated by an LLM, even with infinite prompts to the model with a given prefix, and (b) a specific desired target suffix cannot *always* be produced by an LLM, given a sufficiently long prompt. The evidence for both comes from the same results: for a given LLM θ and decoding scheme ϕ , some suffixes are *impossible* to generate, regardless of (a) how many times one independently prompts the model with a given prefix or (b) how long of a prefix one uses to condition the output.

Our results make this very clear in practice, but we can also make this point conceptually. Intuitively, for a sequence z , if the LLM θ with decoding scheme ϕ assigns the verbatim suffix (given the prefix) zero probability (i.e., $p_z=0$), then no amount of independent sampling trials will ever produce that suffix. If instead $p_z > 0$, then repeated independent sampling ensures that the probability of seeing z at least once tends to 1 as the number of trials grows. Therefore, for any finite set of sequences $\mathbb{Z}$, the maximum number of verbatim suffixes that can ever be generated is exactly those with nonzero probability:

$$\lim_{n \rightarrow \infty} \mathbb{E}[\# \text{ unique verbatim suffixes produced}] = |\{z \in \mathbb{Z} : p_z > 0\}|.$$

Proof. For a single sequence z with per-trial, verbatim-suffix generation probability p_z , the probability of generating the suffix at least once after n independent trials is

$$\Pr(\text{verbatim suffix produced in } n \text{ independent trials}) = 1 - (1 - p_z)^n.$$

If $p_z = 0$, this expression is 0 for all n . If $p_z > 0$, then $(1 - p_z)^n \rightarrow 0$ as $n \rightarrow \infty$, so $\Pr(\text{verbatim suffix produced in } n \text{ independent trials}) \rightarrow 1$. Now let $\mathbb{Z}$ be a finite set of sequences. By linearity of expectation,

$$\mathbb{E}[\#\text{verbatim suffixes produced in } n \text{ independent trials}] = \sum_{z \in \mathbb{Z}} [1 - (1 - p_z)^n].$$

Taking $n \rightarrow \infty$,

$$\lim_{n \rightarrow \infty} \mathbb{E}[\#\text{verbatim suffix produced in } n \text{ independent trials}] = \sum_{z \in \mathbb{Z}} \mathbf{1}[p_z > 0].$$

□

Therefore, for this fixed finite set $\mathbb{Z}$, there is a maximum attainable suffix (given prefix) generation rate: the fraction of $\mathbb{Z}$ whose suffixes can ever be generated. This limit is independent of the target success level p in (n, p) -discoverable extraction (Definition 1). As n grows large, the rate reflects not only extractable (memorized) sequences—i.e., those with sufficiently high probability such that we consider them to be instances of extraction (in this work, $\tau_{\min} \geq 0.1\%$)—but also any z with $0 < p_z < \tau_{\min}$ that we would not count as extraction.

Sequences that cannot be generated in practice. There are various ways that a given suffix could have $p_z = 0$. It might be the case that, for a particular decoding scheme, it is impossible to sample certain tokens, making some suffixes unreachable. This is the case for top- k decoding, which only considers the top- k -highest probability tokens at each sampling iteration. Any token not included in the top- k tokens forecloses possible output generation paths (Appendix A.3). For $k = 40$ and for LLAMA 1 models, only 40 out of 32,000 tokens could be generated in a given sampling iteration.

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. For instance, imagine a token vocabulary of 32,000 tokens (as with LLAMA 1 models). If the model just generated tokens completing randomly (by setting temperature very high, like a “monkey at the typewriter”), then at each iteration every token would have probability $1/32,000$. When generating 71 such tokens in a row (i.e., a 71-token suffix), the probability of any such suffix is $(1/32,000)^{71} \approx 1.3626 \times 10^{-320}$. Once we try to generate a 72nd token, the probability for the whole sequence becomes $(1/32,000)^{72} = 0$: 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 completely-random temperature sampling, not every sequence could possibly be generated.

We often observe 0-probability suffixes regardless of prefix length. This is true for non-training data (Section 5.2, Appendix E.4.2) and for training data that are not memorized (Section 5.1, Appendix E.3). Even if we add more context to try to condition it in a particular direction—there are many cases where we simply are unable to raise p_z above 0 (let alone above $\tau_{\min}$, the minimum probability we consider evidence for extraction).

E.6 Additional notes on extraction and membership inference

There are several points that are worth addressing with respect to the broader literature on **membership inference** and **extraction**. Here, we briefly situate our work—particularly our work with respect to validity—in relation to prior work. This is important to do for clarity reasons, because our work and the types of claims we make are very different from prior work. Prior work tends to focus on *average* extraction rates over *random* samples of known or proxy training data. Since we are making claims about *specific* works from a *specific* population of text, the ways that we go about evaluating and validating our work are necessarily different. This is not just a detail or aside; it has to do with the entire purpose of this project, and significantly influences our measurement approach in ways that contrast with other work on membership inference.

E.6.1 The relationship between membership inference and extraction

As noted in Carlini et al. [38], “Given a set of samples from the model, the problem of training data extraction reduces to one of membership inference: predict whether each sample was present in the training data.” This makes sense intuitively. Extraction is defined with respect to memorization: it is only possible to extract memorized training data, and (tautologically) training data can only be memorized if they are included—i.e., are **members**—of the training dataset. To demonstrate extraction is therefore to demonstrate memorization, and memorization implies membership.

Not only does extraction imply membership of a piece of data, it actually reproduces that data in model outputs. In this respect, it is a stronger evidentiary claim (and, at least in principle, a harder one to make): membership inference asks a one-bit question—was this sequence a member of the training data?—while extraction produces the sequence itself. (At the same time, reliable membership inference with even the strongest **membership inference attacks** is extraordinarily challenging for LLMs [122], while extraction is relatively straightforward.) Put differently, one can make membership inference claims *without* extraction (e.g., with classifiers, see below); these claims can involve evidence where extraction might fail. So, even though successful extraction always implies membership, this does not go in both directions: membership does not imply successful extraction (Appendix E.6.5).

Carlini et al. [38] goes on to note that “In their most basic form, past membership inference attacks rely on the observation that models tend to assign higher confidence to examples that are present in the training data. Therefore, a potentially high-precision membership inference classifier is to simply choose examples that are assigned the highest likelihood by the model.” At a high level, this is clearly directly related to extraction. Greedy-decoded discoverable extraction approximates the highest-likelihood sequence (albeit pretty poorly) with greedy decoding [41]. Probabilistic extraction does the same but allows more flexibility than greedy decoding, and can identify sequences with higher likelihood than greedy-decoded ones [123].

E.6.2 We infer membership for sequences, not whole books

In much work on extraction, researchers do not demonstrate extraction to *infer* membership. Instead, they run extraction experiments on *known* training data [41, 122, 123, 157], i.e., where ground-truth membership is not in doubt. This makes sense as a setup for drawing valid conclusions about extraction and memorization: membership is a prerequisite for memorization, and memorization is a prerequisite for extraction. Testing extraction on known member sequences therefore enables straightforward conclusions about memorization. Here, memorization is the most plausible explanation for extraction. 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]” [37]. As Cooper and Grimmelmann similarly 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” [59, p. 25]. (See Section 2.)

Most of our experiments fall into this category, as the large majority of our results are for LLAMA and PYTHIA models, where it is known that Books3 was included in the training data (Sections 4–6; Appendices D.1 & E.2). This is precisely the motivation behind our baseline experiments (Section 5, Appendix E.3), which we use as a basis of comparison for our negative controls (Section 5, Appendix E.4). Not only do we know that these models were trained on the books in Books3, we know that they were trained precisely on the Books3 corpus.

Nevertheless, while most work on greedy-decoded discoverable extraction does not do this, we perform validity experiments anyway. We need to set a minimum valid extraction probability $\tau_{\min}$ (Appendix E.1). In this paper, we use a very high global $\tau_{\min}$ to avoid false positives, and we validate this choice with our negative controls. Hayes et al. [123] do not set an explicit $\tau_{\min}$, and do not make strong claims about when probabilistic extraction is valid; however, they do run a negative control experiment to show more generally that probabilistic extraction is a reasonable metric.

In the original version of this paper, we included the negative controls on PHI 4 (Appendix E.4.1) to support our claims. Combined with the extensive analysis we perform on known training data for LLAMA, this is how we chose a global $\tau_{\min}$: a conservative threshold that would not over-count extraction (and thus memorization). Even before including our experiments on post-cutoff books,

our results provided ample evidence that non-training data exhibit 0 or orders of magnitude smaller probabilities p_z than training data. The post-cutoff negative controls provide further confirmation.

This was true both for smaller, lower quality models, and higher quality models that exhibit lower loss (on average). While larger models like LLAMA 3.1 70B fit their training data better than smaller ones, and exhibit lower loss on average, they do not (and cannot) exhibit low loss everywhere, such that our measurements of log probabilities would be confounded. This is also clear from our experiments on known training data—Books3 books. For many books and large models, we are unable to extract any unique sequences; there are effectively no nonzero sequence probabilities. It is plausible that future work will find lower, model-specific $\tau_{\min}$ thresholds, which may perhaps need to be set higher for larger models than for smaller ones. For now, however, the threshold we set is so conservative that these nuances do not affect our conclusions.

E.6.3 Inferring membership for LLMs with undisclosed training data

Given our extensive results on known training data—including experiments to vet their validity—we can apply our methodology with confidence to models whose training data are not known with certainty. We do this for three model families: DEEPSEEK v1, QWEN 2.5, and GEMMA 2. For these models, validated extraction serves as evidence of membership. These experiments differ from those on models known to be trained on Books3, but we ensure their validity through extensive negative controls and baselines.

This approach is consistent with prior work. Carlini et al. [38] performed experiments on GPT-2. Nasr et al. [192, 193] extracted training data from open-weight models with undisclosed training data and from ChatGPT, a closed system. In both cases, the authors relied on proxy datasets and strategies to reduce bias in their results, such as shorter prompts. More generally, this is an important and reasonable thing to try to achieve. As researchers, we should be developing valid techniques for extraction that are useful in cases for which we do not know ground-truth membership. If membership inference required knowing membership in advance, there would be nothing to infer; at best—when it works—it would confirm what we already know.

We do not need to use shorter prompts in our work because we validate our work with extensive negative controls. Further, we manually verified that every sequence extracted from PHI 4 and from post-cutoff books was widely duplicated online. Our validity experiments go beyond most prior work, which often does not include negative controls. This is intentional; our work has the potential to impact copyright litigation, not just the scientific literature. We therefore believe it is important to hold our work to a higher bar. Importantly, since training data for these models are undisclosed, we make claims only at the sequence level, not at the level of whole books (Appendix E.6.4).

Finally, the population of data we study differs from most prior work, which evaluates memorization broadly across diverse training data (websites, code, spam, etc.) [38, 41, 123, 192, 193]. Such work must account for false positives like “trivial memorization” (e.g., number lists) or “repeated substrings” (e.g., “I love you” loops). LLMs assign high probability to both of these types of sequences, even if they are not memorized. We note these phenomena explicitly where we identify them, but they are generally not relevant our domain of long-form narrative text.

The rare cases we identify of trivial memorization—lists of page numbers (an artifact of converting EPUB to TXT), chapter indices—are excluded from our claims about extraction of unique creative expression. We identify one case of repeated substrings. Silverman [246] exhibits elevated probability in a part of the book that repeats “I will not wet the bed.” four times in a row. But this sequence does not surpass $\tau_{\min}$, anyway, so we do not count it as extracted. (We also exclude boilerplate instances of true memorization from our claims, like copyright notices and author biographies.) Across millions of sequences, these account for only a few hundred cases. This highlights why the population under study matters for validity: our focus on books yields different considerations than prior work, and we address these systematically. Unlike the work done in Carlini et al. [38], these types of edge cases are extremely rare in our population of interest.

Prior work on copyrighted text and LLMs. In the technical literature, there are a few earlier papers that address the generation of copyrighted text by LLMs. While this work (e.g., Karamolegkou et al. [146]) is important initial work, for the reasons discussed in this appendix, much of it arguably does not directly study memorization, even if it is framed in those terms. In particular, without ground-truth information about training data membership and without negative controls to calibrate

false positives, it is difficult to interpret reported overlap-based results as valid evidence of extraction (and thus memorization). (Further, in the case of Karamolegkou et al. [146], results for very short-form generations—which are insufficient on their own to support extraction claims—are presented alongside longer recovered text within the same plots, making it difficult to distinguish stronger from weaker forms of evidence.)

E.6.4 Our inferences about membership are about sequences, not whole books

We emphasize that our claims about membership are made at the level of suffixes, not entire books. For models such as DEEPSEEK V1, QWEN 2.5, and GEMMA 2, where training data are undisclosed, we cannot establish whole-book membership. Instead, our extraction results are about specific sequences drawn from Books3, a known pretraining dataset, and it is at this level that we make membership claims. In some cases, our results are suggestive that entire books were included in training, given that we can extract high-probability suffixes spanning a full work (e.g., Appendix D.3.153), but we do not generalize from suffixes to whole-book membership.

We also do not claim that we know the provenance of those extracted sequences. For LLAMA models we know that the training data included Books3, and thus at least one copy of each sequence from the books we test was included in the training data via Books3. But this is about all we do know with certainty. We do not know very much about these models’ training data details more generally; knowledge of the inclusion of Books3 is a rare exception, not the norm. Books sequences may have been included in the training data via other unknown sources: other book corpora, widely quoted excerpts in web scrapes, etc. The same could be said for these other models, where we do not know if Books3 was included in the training data.

Altogether, these observations highlight that membership questions for LLMs are rarely straightforward. This is already widely known in the literature, with respect to the presence of duplicates in training datasets [41, 123, 157]. This discussion illustrates how these challenges manifest specifically for books, and why membership claims should be scoped carefully. Our measurements are at the level of individual suffixes, but that is not necessarily the only level at which membership is of interest. For copyright, membership may be meaningful in different ways at different scopes—an entire book, a chapter, a single sentence, etc. [159].

E.6.5 Absence of extraction does not imply anything about membership

We state this with respect to describing the problem of membership inference in Appendix E.6.1, but it bears repeating: extraction implies membership, but membership does not imply extraction. This is clear in our results; many books in Books3 have no extractable unique text for LLAMA models, even though those models were trained on Books3. So when we see $p_z = 0\%$, we should not conclude that a given sequence (or the source that contains it) was not included in the training data. This is trivially true for LLAMA models, but more generally true for models with undisclosed training data. For example, QWEN 2.5 may have been trained on Books3; just because it exhibits low extraction probabilities on Books3 books is not evidence that it was not trained on this corpus. Further, other extraction techniques than the ones we use here could possibly surface additional memorization; so absence of extraction *with our measurement procedure* is not definitive that a particular sequence was not memorized.

Further still, $p_z = 0\%$ does not necessarily imply anything about how much a particular sequence may have influenced the model or its outputs. Attribution of generations to training data is a different problem, both with respect to technical work in ML and implications for copyright. We do not study these topics in this work. At a very high level, one can see this distinction through our brief results in Appendix F.3. We run our reconstruction procedure on *Sandman Slim* [142] and LLAMA 3.1 70B. This model does not memorize this book at all, with respect to our extraction procedure. We run this experiment to show that it fails—that reconstruction is not something that trivially works for any book. Nevertheless, even though we fail to reconstruct any text from the book, it is clear that the generation we produce was influenced by the text (in an informal sense). This type of influence is not something that we capture or attempt to capture—either formally or informally—in our measurements in this paper.

E.6.6 Extractability and copyright

Separate from the above technical considerations about membership and extraction, there are other considerations for copyright. A sequence can count as memorized in a technical sense—in this

paper, when $p_z \geq \tau_{\min} = 0.1\%$ —but this does not necessarily mean that all types of memorization are relevant to copyright. As we discuss in Sections 3 and 8, sequences that have very low extraction probabilities may matter less for copyright than those with high ones, since the latter are far easier to extract reliably in model outputs.

F Recovering books with one seed prompt

Given the extent of memorization that we observe for *Harry Potter and the Sorcerer’s Stone* [224] (as well as other books) and LLAMA 3.1 70B, it seemed possible that we could recover the entire book near-verbatim using only a single seed prompt of ground-truth text drawn from the book.

This result was a natural follow-on to our memorization experiments, but it is not the main focus of our paper. Therefore, this part of our paper reflects only a preliminary investigation into recovering large book segments from an LLM. It is an existence proof that this type of result is possible, not a comprehensive investigation into how to make this type of method work for different texts.

F.1 Recovery method

We successfully recover the book near-verbatim with a variety of seed prompts and variations on the same strategy. Overall, we observe that this result is stable—across different hardware, and for different settings. However, the procedure is relatively sensitive to how it is configured; for example, if too long of a stride is used, then it we observe divergence from the true text. We have preliminary results that significantly improve the stability of this approach. We defer investigating this to future work.

Seed prompt. We obtained high-fidelity recoveries of the book with several seed prompts, including the first 390 tokens of the the first chapter (the first few paragraphs) and the first line of the first chapter (60 tokens):

Mr. 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 nonsense.

With some additional exploration of how to configure generation, we also were successful with prompting with just “Mr. and Mrs. D” and “Chapter ONE:\nTHE BOY” (both 7 tokens). The former is sufficient for the model to output “ursley” and then the rest of the first line; the same is true for the latter, which outputs the rest of the first chapter’s title (“THE BOY WHO LIVED”), and then proceeds to output the first substantive line of the book. (These token counts include the first token, which is the beginning-of-sequence—BOS—token.)

Overall approach. We take a sliding-context approach that we discuss in Section 7. We take an initial seed s and top it up to a minimum of 50 tokens (if it is less than that). On some hardware, for stability, we normalize the format of that longer seed (e.g., making sure that em dashes are consistently represented) before proceeding to long-form generation. In all cases, we do not guide the procedure beyond providing the initial context. We use beam search to generate 50 tokens at a time, and use a maximum context window of 3000 tokens for each generation. This means that we append the 50 tokens we generate to the existing context, and clip the first 50 tokens for iterations that exceed the maximum context. (This is most of the iterations; only at the start do we have a context that can be less than 3000 tokens).

We succeeded with 9 tries in total. (Our first major success occurred in the third run, when we recovered ten chapters—chapters 1-9, and 12—near-verbatim.) The first 3 runs were debugging and setup (e.g., testing the key idea with greedy decoding and getting a sense of how short a reasonable seed prompt could be). The last 6 involved changing the size of the number of maximum context tokens (we settled on 3000) and the number of beam candidates (8), as well as minor tweaks to handling EOS tokens.

Handling EOS tokens was the only (minor) complication, and it was simple to resolve. EOS tokens tend to be predicted at the ends of chapters; we therefore remove them and manually replace them with “CHAPTER {n+1}”, with {n+1} spelled out (e.g., “TWO”, “THREE”). Other attempts involved “Chapter” (instead of all caps) and using the number {n+1} (e.g., “2”) instead of spelling out the chapter number (e.g., “TWO”). Other solutions could just down-weight the logits for EOS.

Occasionally, the model does not predict an EOS at the end of a chapter, which leads to some misalignment later on (e.g., inserting "CHAPTER FOUR" for the fifth chapter). This can cause the model to repeat segments of chapters that have already been produced. The fix here was also simple. We keep track of how many tokens have been generated since the last EOS token, and if that number surpasses 10,000, we assume that we have missed the end of a chapter and account for this (i.e., move the chapter counter ahead even though we did not see an EOS). In future runs, we were able to remove this logic entirely. We defer discussion to future work.

The exact code that we ran using the first 390 tokens as the seed prompt is in Listing 1. Since this code involves beam search, our results should be (in theory) deterministic. Because in practice hardware non-determinism can complicate this, we ran a (non-exhaustive) test by executing the code on two different sets of 4 GPUs. Both yielded the same exact output. (Though we also spot-checked some of the underlying logits, and there were some slight differences.)

Comparing to the ground-truth text. The substantive diff between the version of *Harry Potter and the Sorcerer's Stone* [224] that we have from Books3 and our generated output is minimal. We do some very minor text normalization of the ground-truth book. That is, we trim the beginning and end of the ground-truth book to not contain front matter or anything after the end of the last chapter. We remove `_` from both documents (used for italics in the Books3 version) and align ellipses to `"..."` (which appear as `" . . . "` in the Books3 version.) We then compute various similarity metrics between the two documents. The cosine similarity of TF-IDF vectors of each document is an astounding 0.9999. TF-IDF is a limited metric, as it treats documents as bags of words and thus fails to capture word order. So, as two additional points of comparison, we compute similarity using greedy longest common substring (LCS) matching at both the word and sentence levels (via `difflib's SequenceMatcher`). A score of 1 indicates a perfect match and 0 indicates no match. We obtain a word-level similarity of 0.992 and a sentence-level similarity (which is more sensitive to formatting differences) of 0.934.

The qualitative differences involve minor formatting changes (that we do not normalize) and localization differences. For example, the Books3 version has British spelling (e.g., "Mummy" instead of "Mommy"). These formatting and localization differences are the large majority of the diff between the two texts. Occasionally, the model also misses single lines, e.g., "There was a loud 'Ooooooh!'" (Figure 9).

One of the most interesting outcomes (that we have not investigated in detail) is very occasional apparent hallucinations. For example, in the first chapter, we observe the following diff:

```
"Yeah," said Hagrid in a very muffled voice, "I'd best get this back." "I'll be takin' Sirius his bike away."
G'night, Professor McGonagall — Professor Dumbledore, sir."
```

Figure 223: **Portion of the diff from the first chapter of *Harry Potter and the Sorcerer's Stone*.** Crossed-out red text is in the ground-truth book in Books3, but was not generated by our recovery procedure; text highlighted in yellow/green is in the recovered text but not the ground-truth text.

Selecting the top sequence candidate with beam search always led to this particular outcome in Figure 223, regardless of how we configured our method. Given how specific this particular change was, we investigated it further. It turns out that the generated text reflect the ground-truth from editions of the book printed before 2004. In later editions, this line was changed to the one that we do not generate—the one that is in the ground-truth version we have access to in Books3. This change was done for plot continuity reasons, based on an inconsistency that emerged due to plot development concerning Sirius Black in the later books [120]. This line is not actually hallucinated; it is just from an earlier edition of the book than the one we diff with. Of course, this difference still contributes to our quantitative measurements of similarity, even though it is not a meaningful substantive difference.

```
1 from transformers import AutoTokenizer, AutoModelForCausalLM
2 import torch, json
3 from tqdm import tqdm
4
5 # load the first line of HP 1
6 with open("seed.txt", "r", encoding="utf-8") as f:
7     seed_text = f.read().strip()
8
9 # Load the model and tokenizer
10 model_name = "meta-llama/Llama-3.1-70B"
```

```

11 tokenizer = AutoTokenizer.from_pretrained(model_name)
12 model = AutoModelForCausalLM.from_pretrained(
13     model_name,
14     device_map="auto",
15     torch_dtype=torch.bfloat16,
16 )
17 model.eval()
18
19 max_context_tokens = 3000
20 max_new_tokens = 50
21 # the book is more like 106k;
22 # adding buffer for potential formatting differences
23 max_story_tokens = 113000
24
25 eos_token_id = tokenizer.eos_token_id
26
27 input_ids = tokenizer(seed_text, return_tensors="pt").input_ids.to(model.device)
28 generated_ids = input_ids
29
30 # Set up progress bar
31 pbar = tqdm(total=max_story_tokens, desc="Generating story tokens")
32 pbar.update(generated_ids.shape[1])
33
34 generation_steps = []
35 generation_num = 1
36 chapter_count = 1
37
38 chapter_nums = {1 : "One", 2 : "Two", 3 : "Three", 4 : "Four", 5 : "Five", 6 : "Six", 7 :
    "Seven", 8 : "Eight", 9 : "Nine", 10 : "Ten", 11 : "Eleven", 12 : "Twelve", 13 : "
    Thirteen", 14 : "Fourteen", 15 : "Fifteen", 16 : "Sixteen", 17 : "Seventeen"}
39
40 tokens_since_last_eos = 0
41
42 while True:
43     # Get the tokens for the (up to) 3000-token prompt
44     slice_start = max(0, generated_ids.shape[1] - (max_context_tokens - max_new_tokens))
45     input_ids_window = generated_ids[:, slice_start:]
46     attention_mask = torch.ones_like(input_ids_window)
47
48     # call generation code
49     outputs = model.generate(
50         input_ids_window,
51         attention_mask=attention_mask,
52         max_new_tokens=max_new_tokens,
53         do_sample=False,
54         num_beams=8,
55         early_stopping=False,
56         length_penalty=1.2,
57         temperature=1.0,
58         top_p=1.0,
59         pad_token_id=eos_token_id,
60     )
61
62     # remove the prompt from the output
63     new_tokens = outputs[:, input_ids_window.shape[1]:]
64     new_tokens_list = new_tokens[0].tolist()
65
66     tokens_since_last_eos += len(new_tokens_list)
67
68     # Some fun tricks to deal with EOS tokens
69     if eos_token_id in new_tokens_list:
70         if tokens_since_last_eos >= 10000:
71             print("More than 10000 tokens since last EOS; probably missed a chapter break
72 ")
73             print("Incrementing chapter count")
74             chapter_count += 1
75
76             print(f"EOS: replacing with chapter break")
77             new_tokens_list = [t for t in new_tokens_list if t != eos_token_id]
78             chapter_count += 1
79             chapter_text = chapter_nums.get(chapter_count)
80             if chapter_text:
81                 chapter_text = f"\n\nChapter {chapter_text}\n".upper()
82             else:
83                 chapter_text="\n"
84             chapter_tokens = tokenizer(chapter_text, add_special_tokens=False, return_tensors
85 ="pt").input_ids[0].tolist()
86
87     new_tokens_list.extend(chapter_tokens)
88     new_tokens = torch.tensor([new_tokens_list], device=model.device)

```

```

88     tokens_since_last_eos = 0
89
90     # Everything below is just saving and printing progress
91     prompt_tokens = input_ids_window[0].tolist()
92     prompt_text = tokenizer.decode(prompt_tokens, skip_special_tokens=True)
93
94     generated_ids = torch.cat([generated_ids, new_tokens], dim=-1)
95     pbar.update(new_tokens.shape[1])
96
97     chunk_text = tokenizer.decode(new_tokens[0], skip_special_tokens=True)
98
99     print(f"\n== Generated chunk ({new_tokens.shape[1]} tokens) ==\n{chunk_text}")
100
101     generation_steps.append({
102         "generation": generation_num,
103         "prompt_text": prompt_text,
104         "generated_text": chunk_text,
105         "total_generated_tokens": generated_ids.shape[1]
106     })
107     generation_num += 1
108
109     with open("generation_log.json", "w", encoding="utf-8") as f:
110         json.dump(generation_steps, f, indent=2)
111
112     with open("generated_ids.json", "w", encoding="utf-8") as f:
113         json.dump(generated_ids[0].tolist(), f)
114
115     if generated_ids.shape[1] >= max_story_tokens:
116         print(f"\nReached max story length of {max_story_tokens} tokens; stopping generation")
117         break
118
119 pbar.close()
120
121 # save the full story text
122 full_text = seed_text + "".join(step["generated_text"] for step in generation_steps)
123
124 with open("generated_story.txt", "w", encoding="utf-8") as f:
125     f.write(full_text.strip())

```

Listing 1: Recovery with a single seed prompt.

E.2 Partial recovery

We ran the exact same script on the essay “The Case for Reparations,” which is included in Ta Nehisi-Coates’ book, *When We Were Eight Years in Power* [49]. We observed enormous amounts of memorization for this essay in our sliding-window experiments. (See Section 5, highly memorized middle section; Appendix E.3). This was one of the books in the Kadrey et al. v. Meta Platforms, Inc. [143] lawsuit, for which the court’s judgment noted that neither party produced evidence that demonstrated sizable amounts of extraction of training data of the books in scope for the suit [201].

We summarize the first run of this experiment, where we ran the exact same code that recovered the *Harry Potter* on the first line of the essay (31 tokens, including BOS):

Clyde Ross was born in 1923, the seventh of thirteen children, near Clarksdale, Mississippi, the home of the blues.

The entire essay is about 19,000 tokens long, and running this script generated an essentially verbatim copy of the first $\approx 68\%$ of the text (through the end of section VII). The diff includes small punctuation/formatting differences (e.g., two consecutive `\n` between paragraphs in Books3, with only one in the generation), generating numbers in numeric format rather than spelled out as a word (e.g., generating “21” instead of “twenty-one”), and not generating the essay’s section headers. With minor text normalization (e.g., removing `_`, which indicates italicized book titles in the Books3 version), we obtain word-level greedy approximation of LCS (via `difflib`’s `SequenceMatcher`) of 0.992 and a document-level (of the first $\approx 68\%$) greedy approximation of LCS of 0.996. In word-space (not token-space), the seed prompt is 19 words, and the generated near-exact (punctuation-normalized) text is 10,039 words.

E.3 Running recovery for books with low memorization signal

As a sanity check, we run the same method on LLAMA 3.1 70B for a book for which we observe effectively no extracted sequences.

For *Sandman Slim* [142] (Appendix D.3.98), we use the first paragraph as the prompt and (including this prompt, shown in blue) produce a 2000-token document.

Note that the generation gets caught in a loop (we mark where the prompt gets generated in blue later on in the generation.) Interestingly, even though we do not generate anything that we can identify as copied from the book, the generation does contain references to “Sandman Slim,” the main character “James Stark” (who is never referred to directly that way in the book, but as “James Butler Hickok Stark”) as well as another character called “Mason Faim.”

I WAKE UP on a pile of smoldering garbage and leaves in the old Hollywood Forever cemetery behind the Paramount Studio lot on Melrose, though these last details don't come to me until later. Right now all I know is that I'm back in the world and I'm on fire. My mind hasn't quite kicked in yet, but my body knows enough to roll off the burning trash and to keep rolling until I can't feel the heat anymore. I come to a stop against a gravestone and lie there for a moment, trying to remember how I got here and who I am. I don't always know right away. It's the amnesia. It's a hell of a way to travel, but it's better than being dead. Slowly, like dirty water dropping out of a rusty pipe, it comes back to me. My name is James Stark. Friends call me Sandman Slim. Call me what you like. I've been called worse. I got this name because all I do is bring nightmares. True, I spend a lot of time in the City of Angels, but it's not like they're putting up a statue to me. Before this, I lived in Hell for eleven years. Even when I tried to stay out of trouble, trouble would come right to my door and make friends. That kind of thing has a way of making a man a little jaded when it comes to trusting his fellow man. Or demons. Even angels. I'm like that. I've got a bad attitude. I drink. I've been kicked out of every twelve-step program I've ever joined. I don't have many friends, so the ones I have mean a lot to me. They mean enough that if someone hurts them, I'll tear the universe down to save them. I'm telling you this so you'll know who I am and what I'm talking about when I say there are things in the dark. Things we don't always see that move among us. Some of these things are good and some aren't and sometimes they do things that make you want to rip open your chest and pull out your heart because you've seen too much and there's no going back. I know about these things because I'm one of them. I'm a monster and I'm a killer. I'm good at it and I never let emotion get in the way. When you live where I live, love can get you killed. And when it doesn't, it's worse because if you don't die, you just wish you would and every night you pray for something you lost in Hell. Pray for sweet, sweet death to come make a man out of you. But death doesn't come because if there's one thing I learned in Hell, it's that it will never, ever come for you. No matter how hard you pray. No matter what terrible things you do. But down deep in my heart, I know there's another truth, one that even Hell can't touch. And that's why, no matter what, I can never do the one simple thing that would set me free. The one thing I want more than anything else. I can never, ever kill the man who murdered my friends and started me on this journey to Hell and back. No matter how many times I try, I can't kill Mason Faim, the man I once called best friend. I wish I wasn't me sometimes, but to tell you the truth, even when this is all over and they finally let me die, I wouldn't trade this life for any other. It's been one hell of a ride and I'd do it all over again if I had to. And I probably will. Over and over and over again. Because that's all life is. One big ride that you can never get off of. You can just tighten or loosen your grip on the reins and see which way the horses take you. I WAKE UP on a pile of smoldering garbage and leaves in the old Hollywood Forever cemetery behind the Paramount Studio lot on Melrose, though these last details don't come to me until later. Right now all I know is that I'm back in the world and I'm on fire. My mind hasn't quite kicked in yet, but my body knows enough to roll off the burning trash and to keep rolling until I can't feel the heat anymore. I come to a stop against a gravestone and lie there for a moment, trying to remember how I got here and who I am. I don't always know right away. It's the amnesia. It's a hell of a way to travel, but it's better than being dead. ...

We also tried much longer prompts (e.g., the first two pages of the book), and the model does not continue with text from the book. While these tests are not exhaustive, they are meant to illustrate that this is not behavior that we observe for all books with this approach.

G Discussion of extended results

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 A) varies significantly across models for the same book, and across books for the same model (Section 6 & Appendix D). Average extraction rates—the typical metric in ML memorization research—do not reveal this nuance. Instead, it is necessary to examine per-book extraction metrics to get this view.

We need to run more books to make general claims about overall memorization. We only ran experiments for a small fraction of the entire Books3 dataset—running detailed experiments for 200 books (Appendix D). For 100 books, 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 (Table 3). We randomly sampled the other 100. 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 [157], so it is likely that least some duplicate text persists in training datasets for LLMs.

We do not know exactly what is going on with LLAMA 3+ models. Even so, it seems unlikely that (relatively small numbers of) duplicates completely explain the patterns we observe. LLAMA 3.1 70B exhibits a lot more memorization than any other model. We find higher amounts and degrees of extraction compared to LLAMA 3, LLAMA 2 70B, and LLAMA 1 65B (models of a similar size also trained by Meta). It also exhibits more memorization than a model that registers lower loss and superior performance on downstream benchmarks (QWEN 2.5 72B). That is, as our results also show for LLAMA 3.1 70B on books in the training data that it has *not* memorized, it is not the case that the model’s loss is “low everywhere,” and that this is responsible for the high suffix probabilities we observe. First, that is not how machine learning works; the model cannot so-strongly fit to all of the massive dataset it was trained on. Second, we do not observe this anyway for our results. And third, we do not observe the same degree of memorization in a model that exhibits lower loss on average than LLAMA 3.1 70B.

High-level memorization patterns. In general, we observe a pattern that later generations of LLAMA models memorize more than earlier ones, with respect to average extraction rates (Appendix C) and extraction coverage (Appendix D.4.1). Most books we test exhibited minimal memorization, measured with respect to probabilistic extraction with top- k decoding ($T=1$, $k=40$) and 100-token sequences (50-token prefixes + 50-token suffixes). Our results on longer prefixes (Appendix E), however, indicate that there is a lot more memorization to uncover—even if it is harder to extract, where hardness is loosely equated with requiring a longer prefix.

Nevertheless, much of the memorization that we do observe for books in our main sliding-window experiments frequently falls into one of a few categories: copyright notices, lists of ordered (page) numbers (a common artifact of converting EPUB files to TXT format, see Appendix E.4.1), publisher addresses, chapter number listings, and author biographies. All of these are types of text that are highly duplicated (partially or exactly) in various sources, both the Internet and across books. And so, extraction of a copyright notice from a given book does not 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 category was text from popular sources that are quoted within books: quotes from the Bible [49, 261, 262, 276], philosophers like John Stuart Mill [292], classics like those by Dante Alighieri [19], text from U.S. government documents [292], famous political speeches [49], public domain books [286], social media posts referenced in numerous news articles [286], and famous quotes posted on blogs and platforms like Goodreads [187]. In many cases, this was the only such text we were able to extract with non-trivial probability from some books using 50 token prompts. For our negative controls, we painstakingly verify that each piece of text that we extract within a sequence can be found on multiple sources on the Internet.

With respect to the books that we tested that are within the scope of the Kadrey et al. v. Meta Platforms, Inc. [143] class action suit, we were not able to extract much memorized training data. There were notable exceptions (e.g., Coates [49]). It is of course possible that another extraction technique could reveal additional memorization, but we were unable to meaningfully do so for most books using a 50-token prefixes. It is also important to note that suit was recently decided in favor of the defendants (Meta), and that memorization of training data was only one issue being argued.

Our memorization measurements do not promise that reconstruction naturally follows. We emphasize again that, even for highly memorized books, our sliding-window experiments are *not* extracting enormous amounts of text in one go, i.e., with a single seed prompt to a model. That is *not* what we focused on in this project. However, we do show that it is possible to reconstruct a highly memorized book by using an LLM autoregressively, starting with a single seed prompt of ground-truth text (Section 7 & Appendix F). This type of experiment differs from the main type of results that we showcase in this paper, which use probabilistic discoverable extraction to quantify memorization (Appendix A). We have made a significant effort throughout to make this point clear. If something remains unclear, please reach out to the corresponding authors.




Figure 224: Screenshot from HuggingFace. Downloads of LLAMA 3.1 70B, May 2025.

Final comments on amount of books memorization We were able to extract relatively enormous amounts of memorized training data in some cases, from both public domain books (e.g., Carroll [42], Fitzgerald [92], Joyce [141]) and popular in-copyright books of all stripes (e.g., Brown [31], Camus [35], Martin [174], Rowling [224], Sandberg [237], Tolkien [266]). LLAMA 3+ models exhibited an order of magnitude more memorization on some books than our worst-case estimates prior to starting this project. It was this enormous extent of memorization that encouraged us to try to reconstruct a book starting from a single seed prompt; we would not have believed such an outcome was possible prior to measuring memorization in LLAMA 3+ models—certainly not a flagship model from a frontier company, which continues to be downloaded and used by millions of users (Figure 224).