SoK: Memorization in General-Purpose Large Language Models

Valentin Hartmann
EPFL
valentin.hartmann@epfl.ch

Anshuman Suri
University of Virginia
anshuman@virginia.edu

Vincent Bindschaedler
University of Florida
vbindsch@cise.ufl.edu

David Evans
University of Virginia
evans@virginia.edu

Shruti Tople
Microsoft Research
shruti.tople@microsoft.com

Robert West
EPFL
robert.west@epfl.ch

Abstract

Large Language Models (LLMs) are advancing at a remarkable pace, with myriad applications under development. Unlike most earlier machine learning models, they are no longer built for one specific application but are designed to excel in a wide range of tasks. A major part of this success is due to their huge training datasets and the unprecedented number of model parameters, which allow them to memorize large amounts of information contained in the training data. This memorization goes beyond mere language, and encompasses information only present in a few documents. This is often desirable since it is necessary for performing tasks such as question answering, and therefore an important part of learning, but also brings a whole array of issues, from privacy and security to copyright and beyond. LLMs can memorize short secrets in the training data, but can also memorize concepts like facts or writing styles that can be expressed in text in many different ways. We propose a taxonomy for memorization in LLMs that covers verbatim text, facts, ideas and algorithms, writing styles, distributional properties, and alignment goals. We describe the implications of each type of memorization— both positive and negative—for model performance, privacy, security and confidentiality, copyright, and auditing, and ways to detect and prevent memorization. We further highlight the challenges that arise from the predominant way of defining memorization with respect to model behavior instead of model weights, due to LLM-specific phenomena such as reasoning capabilities or differences between decoding algorithms. Throughout the paper, we describe potential risks and opportunities arising from memorization in LLMs that we hope will motivate new research directions.

1 Introduction

Large language models (LLMs) such as ChatGPT and LLaMA have taken the world by storm, with a plethora of applications and widespread interest and use by the general public, researchers, and industry. Unlike previous machine learning (ML) models that were usually trained on comparatively small, curated datasets, large modern general-purpose generative models are trained on data that was not originally meant for model training, and is too large for careful manual curation. Increasing model sizes brings an increased capability of models to memorize substantial parts of those training sets. Memorization, especially in LLMs, can have unintended consequences such as leakage of sensitive information like social security numbers [80, 101], or the regurgitation of large parts of training documents [22, 129].

Preprint. Under review.

Type of Memorization	Performance, Alignment	Privacy	Security, Confidentiality	Copyright	Auditing
Verbatim text (Sec. 4)	capability to quote (+), ignoring instructions	leaking sensitive data	leaking proprietary information	regurgitating training data	watermarking (+)
Facts (Sec. 5)	answering questions (+), overriding prompt data	leaking personal information	leaking cryptographic keys	∅	detecting hallucination (+)
Ideas and algorithms (Sec. 6)	performing common algorithms (+), reproducing harmful ideas	decontextualization of irony	leaking secret ideas	violation through fan fiction	determining generalization capabilities (+)
Writing styles (Sec. 7)	style transfer (+), reproducing toxicity	author attribution	social engineering	violation in combination with general idea of a work	detecting underrepresented communities in training data (+)
Properties of training distribution (Sec. 8)	necessary for learning (+)	identifying human labelers	theft of preprocessing parameters	help in identifying potential violations (+)	detecting biases (+)
Alignment goals (Sec. 9)	necessary for alignment (+)	leaking preferences of human labelers	spotting weaknesses in safety alignment	∅	verifying safety claims (+)

Table 1: Main potential implications of the kinds of memorization covered in this paper across different domains. Positive implications are marked with (+), negative implication are not marked. We use ∅ to demarcate domains where we did not identify any obvious risks or benefits.

Research on the privacy and security of LLMs has so far mostly looked at such verbatim memorization, in line with prior work on membership inference [95, 153] and model inversion/attribute inference [49, 176]. We argue that text data contains information such as facts or writing styles that LLMs can memorize but are not captured by just considering memorized verbatim text. We aim to broaden the understanding around memorization in LLMs by providing a taxonomy of memorization. We further want to draw attention to risks and benefits of memorization for various domains.

Contributions. We systematize a wide range of works on the broad topic of memorization for large language models, discussing challenges with defining memorization and considerations that anyone aiming to measure memorization in a given setting has to take into account (Sec. 3). We identify the following disparate types of information that LLMs can be said to memorize: 1) verbatim text (Sec. 4), 2) facts (Sec. 5), 3) ideas and algorithms (Sec. 6), 4) writing styles (Sec. 7), 5) properties of the training distribution (Sec. 8), and 6) alignment goals (Sec. 9). For each type of memorization, we provide possible definitions and ways to detect and prevent it. We further discuss the implications of these types of memorization for model performance, risks (privacy, security), and governance (copyright, auditing). We summarize existing research, and highlight new research directions in the absence of prior work. Table 1 overviews the most important implications.

By bringing together research from LLMs, ML, privacy, security and law, we provide practitioners that want to train and deploy LLMs with an overview of the challenges associated with memorization. Researchers benefit from a new contextualization of existing research, and from the open problems and research ideas we identify within these contexts.

Related work. Memorization is often only considered explicitly in the context of verbatim memorization [22]. Ishihara [65] gives an overview of methods for detecting and exploiting such memorization. AlKhassimi et al. [2] provide a survey on LLMs as knowledge bases, i.e., stores for accessing and editing facts. Some papers [21, 75] evaluate the ability of LLMs to solve competitive programming questions, which often requires the knowledge of particular efficient algorithms, but we are not aware of any work particularly looking at algorithm or idea memorization. Tyo et al. [168]

describe and evaluate several methods for authorship attribution based on LMs, a task that requires (at least partially) memorizing writing styles. Distribution inference, the field concerned with the memorization of distributional properties, is small and emerging, and thus there are no surveys on the topic relating to language models. Wang et al. [177] describe existing methods for aligning LLMs and for evaluating alignment, which can be seen as the memorization of alignment goals. Overall, there are only few works that explicitly consider memorization, and none that approach this topic with the breadth of our paper.

The domains on which we consider the impact of memorization, on the other hand, were picked precisely because they are of large interest. Model performance is usually evaluated through extensive benchmarking, where BIG-bench [158] particularly stands out due to its comprehensiveness. Smith et al. [155] and Fan et al. [44] survey different types of privacy risks emerging from LLMs. Mozes et al. [119] and Ferrara et al. [48] discuss safety and other risks resulting from the malicious use of LLMs. Henderson et al. [56] outline copyright violation risks resulting form the deployment of foundation models such as pretrained LLMs. Mökander et al. [117] discuss and propose ways to audit LLMs, both in terms of process, and in terms of technical tools. Despite the large body of research on those domains, our paper is the first to consider them through the lens of memorization.

2 Background

Autoregressive models are trained to predict the next text token based on the sequence of previous tokens, as opposed to, e.g., bidirectional masked language models such as BERT [31], which are conditioned on both left and right context. In this work, we focus on large autoregressive language models (LLMs) with billions of parameters such as GPT-4 [126] or Llama 2 [166]. Sometimes we discuss work that uses models with fewer parameters or models that are not autoregressive. For these, we use the generic acronym LM. We write $p(\cdot|x)$ for the probability distribution of an autoregressive model given a context x.

2.1 Training LLMs

Regardless of the concrete model architecture, current LLM training typically encompasses three main stages: pretraining, supervised fine-tuning, and reinforcement learning from human feedback (RLHF).

Stage 1: Pretraining. The training data in this stage consists of text documents. The LLM is trained to predict the next token of a document given a prefix from the document. The pretraining dataset usually consists of public data such as web pages crawled from the Internet, books, etc., though the data sources are not always disclosed [126].

Stage 2: Supervised fine-tuning. The second stage uses a dataset of prompts and responses. The model is trained in a supervised fashion to give the response corresponding to a prompt in this dataset. The prompts and responses can come from public, task-specific datasets [179] or be written specifically for the model training. In the latter case, both the prompts and the responses can be written by human labelers, or the prompts can come from user requests to a language model [128], which allows for more closely modeling the prompt distribution upon deployment of the model.

Stage 3: Reinforcement learning from human feedback (RLHF). This stage consists of several steps. First, human labelers rate answers generated by the model in response to prompts. The ratings or labels can also be model-generated [166] or come from language model users [128]. These ratings are then used to train a reward model, which is finally used for optimizing the LLM via reinforcement learning. Sometimes automated tools are used in addition to human labelers [126].

2.2 Impact of LLM memorization

Memorization in LLMs impacts various domains. The implications of memorization for model performance and auditing are largely positive, whereas privacy, security and copyright are often negatively impacted. For each type of information in Sec. 4–9, we describe the implications of memorization on the domains described below. The domains are grouped by task performance

(model performance and alignment), risks (privacy; security and confidentiality), and governance (copyright; auditing).

Model performance and alignment. Model performance refers to the performance of the model in downstream applications, for example question answering, creative writing or code generation [158]. This is usually not perfectly correlated with purely technical metrics like perplexity on a test set. The related problem of alignment refers to removing undesirable behavior of the model such as toxicity, bias, etc., and instilling desirable behavior such as instruction following [9].

Privacy. These are typically risks associated with the release/inference of previously non-public information that can be linked to individuals. Even when models are trained only on public information, decontextualization of public information about individuals can be considered a privacy risk [19].

Security and confidentiality. We consider potential harms through information leakage (e.g., leakage of API keys [84]) or other means for institutions that train or use LLMs, or contribute training data. This includes memorization that enables abuses of LLMs such as efficient social engineering attacks [119].

Copyright. LLMs are typically trained on large datasets containing material with copyright protections, such as the BookCorpus dataset [11] (e.g., BERT [31] and GPT-3 [20]), components like Books3 of The Pile [52] (e.g., GPT-Neo [16] and Pythia [15]), or web crawls such as C4 [140] (e.g., T5 [140] and LLaMA [165]). Typically, no licenses are obtained for the use of this material in model training. In the US, the fair use doctrine (17 U.S.C. §107) allows for the use of copyrighted material without a license under certain conditions, but how this applies to training generative models is not yet clear. Rahman and Santacana [141] argue that training an LLM on copyrighted documents itself does not constitute a copyright violation. They draw the analogy between a model learning from copyrighted texts and students learning from such texts. Lemley and Casey [91] come to the same conclusion, arguing that the model weights constitute a transformation of the copyrighted training data. However, this debate is not completely settled yet [56]. In applying the fair use defense to outputs produced by an LLM, Rahman and Santacana point to the importance of taking into account the context in which these outputs are used: educational purposes might fall under fair use, but commercial purposes not. Henderson et al. [56] discuss the applicability of fair use to LLMs in depth, and likewise conclude that the fair use defense does not always apply. Importantly, the authors give examples showing that verbatim copying is not necessary for a copyright violation. Direct translations, re-using fictional characters for fan fiction, or re-telling a story in a different setting may all be considered copyright violations.

Auditing. Some forms of memorization can be used to audit LLMs, that is, to test their compliance with specifications or regulations pertaining to the composition of the training data or the model behavior [117].

3   Measuring memorization

Learning from data requires extracting information from the data and generalizing from this information. The learning abilities required to perform the generalization vary for different tasks. Regurgitating a specific fact seen directly in the training data requires no generalization; mimicking the choice of words and syntactic structures that make up a person's writing style requires some generalization; writing a new creative story requires a lot of generalization. Our focus is on memorization, which we define roughly as learning that involves only little generalization. We acknowledge that, as of now, there is no complete understanding of generalization [188] and no clear demarcation between memorization and generalization [45, 47]. We do not attempt to make progress on this demarcation problem here, but rather just use our intuitive notion to scope our paper. The definitions that we consider in later sections are all concerned with more specific objects than general memorization. In the remainder of this section, we describe general approaches for, and challenges with, identifying and measuring memorization.

4

3.1 Inference attacks

Privacy researchers often consider what information is stored in a model through inference attacks that are designed to make inferences about the training data. Two common types are *membership inference* and *attribute inference* [185]. In the membership inference task, an adversary is given access to a model and knowledge of a candidate record, and wants to determine whether or not that record was included in the model's training data. In attribute inference [182], the adversary is given incomplete information about a record from the training set and has to use its access to the model to infer some missing information about this record. A successful membership inference attack implies that the target model has memorized *something* about the candidate record, since the model reveals whether or not it was trained on that record, but not what precisely.

A successful attribute inference attack, on the other hand, enables inferences about data not previously known to the adversary.[1] This implies that the model has stored the information learned in some way. This does not necessarily mean that the model has *memorized* these attributes, though. It could be that it has merely learned the data distribution well enough to predict the attributes based on other attributes of the record [68]. If, however, the attribute inference attack on a counterfactual model trained on the same data with the record removed is unsuccessful in recovering the attributes, then we have evidence that the attack's success was due to memorization of the target attributes by the original model. When successful and done carefully, this can give a much more precise characterization than membership inference of what the model has memorized about a record. We consider memorization in the sense of both membership and attribute inference, since both types of memorization can have significant implications. For example, in the case of verbatim text, knowing that a particular text was used for training an LLM (a successful membership inference attack) might give an auditor information about whether a particular user's data was used in training. Having a model output the first chapter of a book when prompted with its first sentence (a successful attribute inference attack) might constitute a copyright violation.

While membership and attribute inference are concerned with individual data points, we consider several pieces of information that may affect multiple documents, such as writing styles or parameters of preprocessing methods. Memorization of such information is better captured by the concept of *distribution inference* (also called *property inference*) [8, 54, 161]. Distribution inference assumes an adversary that has a set of candidate training distributions and aims to determine the underlying training distribution of a given model. For example, one might be interested in knowing whether a certain percentage of the data was written in a particular style, or what fraction of the software engineers described in the training data are female.

3.2 Challenges in estimating memorization

It is difficult to precisely determine whether a piece of information was memorized by an LLM. There are several reasons for this, which we discuss in this subsection.

Memorization vs. extractability. Anything that a model knows about its training data needs to be stored in some way in the model's weights. A naïve definition of memorization could thus be "any information that is stored in the model's weights is memorized". However, evaluating this definition would be infeasible and basically amount to fully determining everything the model has learned. Researchers thus resort to studying proxies for this memorization via extractability or discoverability, which only captures memorized information that can be accessed through known methods. This inherently underestimates memorization, since it assumes there are no better ways to extract information from the model.

For example, many definitions are concerned with what can be inferred from model outputs—a subset of what can be inferred from model weights. This essentially assumes a black-box (API) adversary who has no direct access to the model. In some settings, such an assumption is valid and defenses such as output filtering, that do not prevent memorization in model weights but only the revelation of memorized information at prediction time, may be effective. If output token probabilities

---

[1]While a perfect membership inference attack can be used to perform attribute inference [147, 185] and thus implies that the model has memorized attributes of records, practical membership inference attacks are not perfect, and further usually only deployed for extracting one bit of information about membership or non-membership of a candidate record.

are provided, definitions can make use of them (e.g., the probabilities of the possible answers to a multiple choice question). Otherwise one has to resort to the model outputs in character space, which are produced by decoding algorithms. These decoding algorithms can be deterministic (e.g., greedy decoding), or non-deterministic (e.g., top-p sampling [58]). Furthermore, the choice of decoding algorithm can influence model hallucination (see next paragraph) [38, 89] and other behaviors [76]. Carlini et al. [22] find that swapping greedy decoding for beam search slightly increases the discoverability of memorized verbatim text. Another challenge with output-based definitions is their dependency on a prompt (with some exceptions [25]), the choice of which influences the amount of extracted memorized information [72].

Hallucination. LLMs can generate plausible content that cannot be inferred from their training or input data [70], known as hallucination. Causes of hallucination include training data that favors text generation that is not grounded in the data [32, 173], or a training objective that differs from the task objective [172]. Hallucination can also be linked to memorization [144]. Hallucinations can make it seem as if the model had memorized a piece of information, even though it was not present in its training data. This is referred to as *extrinsic hallucination*, as opposed to *intrinsic hallucination* that contradicts the training data or input [70].

Reasoning and generalization. Beyond factual outputs that are not grounded in the training data, there are outputs that are grounded in the training data but not explicitly contained in it. For example, training documents might contain the facts "[A] is a student of [B]" and "[B] is a professor at university [X]", but not the fact "[A] is a student at university [X]". Still, if the LLM is sufficiently powerful to perform deductive reasoning, it will give the correct answer to the question "Where does [A] study?" [60]. Similarly, LLMs often appear to perform inductive reasoning [183] such as guessing the nationality of a person based on their name [136] or generalizing from code seen during training to create novel algorithms [75].

Distinguishing memorization from hallucination and reasoning. A correct model response to a factual question does not reveal whether the model arrived at this response via generalization and reasoning, because of hallucination, or because it has memorized this response from the training data. When looking for memorized information in the model weights instead of in the model behavior [113], cases of reasoning and hallucination relating to input data [60, 70] can be avoided. In many cases of interest, such as personal identifiers, social security numbers or long passages of verbatim text, it is unlikely that a model could hallucinate the target information or gain knowledge of it through reasoning.

In the following five sections, we cover five different types of memorization, starting with verbatim text. Each of these sections is separated into an introduction, a subsection listing potential definitions, a subsection discussing the impact on various domains, a subsection on how to detect this type of memorization, and a final subsection that discusses mitigation strategies. Sec. 10 discusses potential mitigation strategies that are not specific to any type of memorization.

4 Verbatim text

Memorizing verbatim text is the most direct and low-level form of memorization. It is also quite prevalent—Carlini et al. [22] demonstrate how GPT-J [16, 171] memorized at least 1% of its training data according to their extractability definition (Sec. 4.1), which is similar to attribute inference.

There are different ways to look at verbatim memorization. Researchers have considered the memorization of entire training documents [56], parts of training documents [22] and, in the context of privacy risks from personally identifiable information (PII), the memorization of short sequences that comprise personal information such as names or email addresses [101]. The concept of verbatim memorization can be broadened by also considering paraphrases, such as those resulting from the replacement of words with synonyms [87]. Note that the verbatim memorization of a text sequence that describes a fact or algorithm implies the memorization of this fact or algorithm, though in a low-level representation. We restrict this section mostly to aspects that are unique to verbatim memorization, and discuss memorizing facts and algorithms in later sections.

4.1 Definitions

Since verbatim memorization is related to the tasks of membership and attribute inference [147], some definitions of inference attacks could also be applied to verbatim memorization. However, several formal definitions of specifically verbatim memorization in LLMs have been proposed, on which we focus here.

Exposure metric [24]. Carlini et al. explore memorization of out-of-distribution secrets by LMs. They consider strings of the form "The random number is r", where r is a number randomly sampled from some space $\mathcal{R}$. The authors sample one particular $r' \in \mathcal{R}$ and include the corresponding string in the training set of the LM. They define the *exposure metric*, which essentially measures how many guesses an adversary that tries to guess r' saves by computing the perplexity of the string "The random number is r" for all $r \in \mathcal{R}$ and guessing r in order of increasing perplexity, over guessing in a random order from $\mathcal{R}$. This metric requires many queries to the model to compute, in addition to a retraining of the model. Note that this approach avoids the difficulty of determining what should and should not be learned by a model, since the artificial random strings are introduced as an explicit way to insert content that should not be learned.

Extractability [25]. This definition by Carlini et al. defines a string y to be *extractable* from an LM p if there exists a prefix x such that: $y \leftarrow \arg\max_{y':|y'|=N} p(y'|x)$. Instead of the intractable computation of the $\arg\max$, the authors use a decoding algorithm such as greedy decoding in practice. They also point out that due to pathological cases any string could be extractable, e.g., when prompting the model to repeat a given input string. However, they instantiate the definition only for cases where x is the start-of-sequence token or a prefix from a document. The authors then specialize this definition to k-*eidetic memorization*, which only allows for strings that are repeated at most k times in the training data.

Counterfactual memorization. Following the observation that strings that occur more often in the training data are more likely to be reproduced by the model, Zhang et al. [189] aim to disentangle the plurality of a document in the training data and the degree to which its verbatim text is memorized. To this end, they define counterfactual memorization for a document d as the expected difference in token prediction accuracy when predicting d with models trained on datasets containing d and datasets not containing d. This definition is a variation on a definition by Feldman and Zhang for label memorization [47], and is closely related to the concept of algorithmic stability [18]. Counterfactual memorization does not measure memorization in any specific model, but the degree to which a given model architecture memorizes examples from a given distribution on average. It requires training multiple models, and also requires access to the training data.

4.2 Implications

Model performance and alignment. For certain types of tasks it is desirable that the model remembers verbatim text. E.g., reviews of a book can greatly benefit from verbatim quotes, and likewise news articles about political speeches. On the other hand, verbatim memorization can also hurt model performance. As shown by McKenzie et al. [110], there are situations where LLMs ignore task instructions and output memorized text when the prompt contains text from the training data.

Privacy. As discussed in Sec. 2.1, a model trainer may rely on user-submitted prompts and responses in training stages 2 and 3. These prompts can contain highly sensitive information, e.g., when a user asks for advice on medical or relationship issues. In such cases, the prompts may contain detailed information about the user, allowing a third party to identify them in case the model regurgitates the prompt.

Security and confidentiality. As with privacy, the security and confidentiality risks mostly come from the memorization of data from the training stages 2 and 3. LLMs used as coding assistants could leak code from non-public code bases of companies. This could give an advantage to competitors and make it easier for malicious actors to find vulnerabilities in applications based on this code. Leakage of sensitive information from code-based models has already been studied and established [61, 123]. In fact, a recent incident with Samsung [108] involved employees sharing internal, sensitive code and meeting notes with OpenAI's ChatGPT, where they may be used for train-

ing [150]. Similarly, LLMs integrated in office packages or explicitly prompted with content of internal documents of an organization could memorize and reveal secrets of the organization when trained on these documents.

Copyright. As discussed in Sec. 2.2, it is not entirely clear yet whether verbatim memorization of copyrighted documents from training data in just the model weights itself constitutes a copyright infringement. As further discussed there, even in case of regurgitation of copyrighted text, copyright violations can be context-dependent. At the same time, verbatim regurgitation may not always be necessary to constitute infringement. Lee et al. [87] investigate the related problem of plagiarism. They measure the degree of verbatim (copying passages character by character), paraphrase (para-phrasing sentences from a source document), and idea plagiarism (copying core ideas) performed by GPT-2 when prompting the model just with an end-of-sequence token. The authors show that all three types of plagiarism, from both pretraining and fine-tuning data, occur both in pretrained and fine-tuned models both, but not for all fine-tuning datasets. Fine-tuning also seems to reliably eliminate verbatim plagiarism from pretraining data. Henderson et al. [56] argue that preventing copyright violations can only be done on a higher semantic level that goes beyond verbatim text matching.

Auditing. Detecting memorized verbatim text can enable identification of certain datasets used for the training of an LLM, although likely requiring access to the specific candidate dataset [22]. For example, the benchmark BIG-bench [158] includes a randomly generated string that acts as a globally unique identifier (GUID), and a test that checks whether the model assigns an anomalously low/high probability to the GUID, which would indicate a contamination of the training data with data from the benchmark. Such canaries could be inserted by model trainers to detect theft or unlicensed use of their models, potentially aiding techniques like dataset inference [104] that might otherwise not work well with LLMs [163]. Memorized documents from different dates may also be used to determine a lower bound on the knowledge cutoff date (until when training data was collected) of a model, as recently demonstrated for Github's Copilot [30].

4.3 Detecting memorized verbatim text

There have been several attempts at detecting verbatim memorization that build upon membership inference tests. For instance, the current state-of-the-art membership inference attack for LLMs [107] relies on fluctuations in loss values around neighborhood, which is similar in idea to Merlin [69], a membership inference attack for general machine learning. However, it is important to keep in mind that successful membership inference does not necessarily imply verbatim memorization (see Sec. 3), or the converse.

While the cost to train LLMs makes it infeasible to utilize techniques that require training shadow models, as is done in most membership inference attacks, the open-endedness of prompts opens new avenues unavailable for other domains such as vision and tabular data. For instance, simply prompt-ing the model with the title and author name of a training document [56] is sometimes effective. In the same work, Henderson et al. [56] sample random snippets from books and show how some LLMs, when prompted with these snippets, return long sequences from the books. They show that instructions like *replace every 'a' with a '4' and 'o' with a '0'* can circumvent content filters. Yu et al. [186] use prompts with function signatures and code comments to extract program code from LLMs.

Prompting techniques not explicitly designed for detecting memorization could be repurposed, such as adding prefixes to encourage grounding in the training data [180]. Instruction-tuned models are much more amenable to this type of prompting, indicating that instruction tuning can make it easier to access information memorized in model weights. Some attacks rely on prefix-suffix completion to detect memorization. Ozdayi et al. [129] use prompt tuning on models by utilizing white-box access and knowledge of some training records. The attack generates a prefix that can be prepended to a prompt to maximize the likelihood of generating a suffix corresponding to training data. Outside of these prompting attacks, there have been some recent attempts at attributing memorization of examples to specific neurons in models [103].

4.4 Preventing memorization or extraction of verbatim text

One strategy for preventing memorization of verbatim text is to avoid repetitions of verbatim text in the training data, motivated by the observation that the likelihood of a sequence is memorized increases with the number of times that sequence occurs in the training data. Carlini et al. [22] extract memorized verbatim text from models of the GPT-Neo [16, 171] family. They sample text sequences from the training data, prompt the model with a prefix of each sequence and check whether the model generates the corresponding suffix. They find that the amount of memorized text increases with model size, repetition of the sequence in training data, and the length of the prefix prompt. Increased memorization from repeated sequences has been observed before, and consequently de-duplication of the training data has been proposed as a countermeasure [78, 88]. However, this might run counter to the effective upsampling (via training for more epochs) of trustworthy sources commonly performed in the training of LLMs [52, 165].

Mantri and Sasikumat [106] propose several potential pathways towards LLMs that do not regurgitate memorized copyrighted content: pruning or zeroing out parameters associated with such content; fine-tuning the model with non-copyrighted content; and the use of loss functions that discourage the model from generating text too similar to copyrighted training data. A more radical change to prevent verbatim memorization would be to use a substantially different training objective: instead of learning to predict individual tokens from training documents, the model could learn to predict the content of those documents at a higher semantic level. This has been tested for computer vision models, where the prediction of individual pixels in the training objective is replaced by latent representations of image patches [7].

Verbatim memorization can also be addressed in black-box settings by using post-processing to block text sequences in the training data from occurring in the generated output. Ippolito et al. [63] propose using a Bloom filter to detect n-grams of the model output that appear in training data, and re-generate tokens until there is no more match in the training data. However, the filter cannot account for small differences, such as changed whitespaces, and can be actively avoided by style-transfer prompts, e.g., making the model respond in all lowercase.

5 Facts

LLMs achieve good results on knowledge benchmarks even in closed book settings [165], implying sufficient memorization of facts about the world. These can be facts about the real world like "birds can fly" or facts about fictional worlds like "Harry Potter studies at Hogwarts" [27]. Li et al. [96] provide empirical evidence for memorization of the co-occurrence of words in different topical contexts in both embeddings and self-attention layers.

5.1 Definitions

Generic definitions of facts can capture the notion of what a fact is, but make it difficult to distinguish between desirable and undesirable learning of facts.

Tuple completion. Meng et al. [113] represent facts as tuples $t = (s, r, o)$ of a subject s, a relationship r and an object o. They define memorization of a fact (s, r, o) as the model completing the prompt '$s\ r$' with 'o'. The Knowledge Assessment Risk Ratio (KaRR) [34] considers the ratios between the probability of the correct object being generated by the model when given s and r, and when given only s or only r. This is aimed at removing the influence of the prior probability of the model for generating o.

Counterfactual memorization. The definition of counterfactual memorization [189] (see Sec. 4.1) might also apply to facts: Instead of a specific document, one would remove all occurrences of a specific fact from the training data, and measure how this influences the knowledge of the model about the fact. This could help with determining whether a model knows a given fact because of memorization (in that case it would not know the fact anymore after the removal) or because of reasoning (in that case the model would still know the fact).

Personally Identifiable Information. A particular class of facts that provide information about identifiable individuals is known as *personally identifiable information* (PII). Kim et al. [80] draw a

distinction between structured PII and unstructured PII. Information in structured PII follows a fixed pattern, such as email addresses and phone numbers. Unstructured PII can be expressed in different ways. e.g., "[PERSON 1] is the parent of [PERSON 2]" contains the same information as "[PERSON 2] is [PERSON 1]'s child".

PII extractability. Lukas et al. [101] define three variants of PII leakage across different threat models, which can be used to define PII memorization, but also the memorization of more general facts. The first definition is extractability, which is the probability of a piece of PII being contained in an output produced by an unprompted generation of the model.

PII reconstruction and inference. The second and third definitions of Lukas et al. measure the model's ability to associate PII with a context. A sentence containing at least one piece of PII is chosen from the training data and the PII is replaced by a [MASK] token, for example "The police arrested [MASK] near the White House on 8/20." In PII reconstruction, the (approximately) most likely PII replacement for [MASK] under the model likelihood is compared with the PII in the original sentence. PII inference differs from PII reconstruction in that the model only has to choose from a predefined set of candidates.

Linkability of PII. Kim et al. [80] argue that a random disclosure of some personal information without it being linked to an individual does not necessarily pose a privacy risk. For example, it might not be problematic if an LLM leaks an address without the name of the resident. They propose the definition of linkable PII leakage, which implies memorization of the connection between different pieces of personal information. The definition roughly states that if the likelihood of a piece of personal information a_1 under the model increases when conditioning the model on other personal information $a_2, \ldots, a_n$ linked to the same data subject, the model links a_1 to $a_2, \ldots, a_n$. This definition is also applicable to other facts, e.g., linking a football player to her teammates.

5.2 Implications

Model performance and alignment. In order to answer questions about the world, LLMs need to memorize facts about the world. Even when facts are retrieved from external sources [79, 120, 130, 142, 152], the LLM needs to have enough world knowledge to make sense of those facts. Likewise, there is information that is structurally indistinguishable from PII, but meant for public use and useful for the model to memorize. Examples for this are the phone numbers of emergency services [80] or the support email address of a company. On the other hand, memorization of invalid logical theorems can lead to the model committing logical fallacies [110]. In some cases, it might not be desirable that the model uses facts from its training, but rather information provided in the prompt. Memorized facts can hinder the model in incorporating this new information, e.g., when redefining mathematical constants [110].

Privacy. If only publicly-accessible data is used for the training of LLMs, they cannot memorize facts that are not already publicly disclosed. However, as argued by Brown et al. [19], LLMs have the potential to decontextualize information, that is, bring up the information in contexts which it was not intended for or without essential surrounding context. LLMs can reveal PII and other information that is not meant to be public but found its way online—e.g., through users who shared sensitive information in online forums via accounts that can be traced back to them. PII could also make its way into the model's memory if user-submitted queries are used in stage 2 or 3 of the training. For example, a user might ask the model to draft a reply to a letter, where the letter might contain the user's name and address.

Security and confidentiality. As with privacy, involuntarily published confidential information contained in the pretraining data can also pose problems for security and confidentiality. Leaking API keys and cryptographic secrets accidentally committed to GitHub, which is a popular source for pretraining data [52, 165], is a well-known problem [84, 111]. This risk extends to secrets obtained and published by cybercriminals [102, 154]. While an accidentally committed private key might be quickly deleted from GitHub, it could stay in the model's memory indefinitely. As for the memorization of verbatim text, confidential facts such as salaries and strategic decisions may also be memorized by the model if user-submitted data is used in stages 2 and 3 of the training process and an organization gives the LLM access to internal documents or code.

Copyright. Facts cannot be protected by copyright, and only their expression can be protected [56]. For example, textbooks on the same topic often have a large overlap in terms of the facts they contain without posing any copyright issues. On the other hand, preventing memorization of facts from a particular training document would make it impossible for the model to memorize the verbatim text of the document in its entirety, and might mitigate some copyright concerns. This could be relevant for works of fiction where one might want the model to learn the writing style, but not the details of the fictional world.

Auditing. As with verbatim memorization, facts memorized by the model provide a lower bound on the knowledge cutoff date of the model: if the model has learned a fact about an event, it has seen documents at least up to the point when that event occurred. If the downstream application requires truthfulness, it is important to detect untrue statements learned by the model [97]. Work on attributing model outputs to training documents [115, 131, 143, 180] can help in detecting hallucinations.

5.3   Detecting memorized facts

Commonly used benchmarks for assessing factual knowledge of LLMs consist of questions about facts in natural language [74, 85]. Some are in the form of multiple-choice questions [57], which are mostly suited to determining aggregate knowledge of model. Another variant is *cloze* tasks, where the model is asked to fill in masked-out entities [27, 125, 134]. The exact prompt formulation can influence the success of knowledge extraction for cloze tasks [72] and questions [158].

Jiang et al. [72] propose ensemble methods that leverage multiple prompts, generated through mining and paraphrasing content from Wikipedia. Li et al. [93] demonstrate how jailbreaking models can be exploited to increase PII leakage, as opposed to standard querying. Dong et al. [34] treat subject, relationship and object in the tuple completion task (see Sec. 5.1) as latent variables, and consider different textual aliases to generate different strings pertaining to the same s, r, o tuple. Jain et al. [67] present the model with a fact and its negation, and compare the perplexity of the model on those two strings. As a step towards identifying knowledge in the model's weights, Meng et al. [113] use causal interventions to identify components of LLMs that store facts. Specifically, they consider s, r, o tuples, where the model has to predict the object from the subject and relationship. Based on their findings, the authors posit that the MLP modules in the transformer act as two-layer key–value stores. Lehman et al. [90] study medical PII leakage of name–condition pairs via the release of embedding weights. Patil et al. [132] demonstrate, via parameter-analysis and prompting-based methods, how PII leakage persists even after information "deletion" (based on parameter editing).

5.4   Preventing memorization of facts

Fact memorization may be mitigated at various stages in the training pipeline, including removing text associated with particular facts from the pretraining data through using machine unlearning techniques to remove specific learned facts from a model.

A technique called scrubbing acts already on the level of training documents—identifying pieces of text, e.g., via named entity recognition [1, 59], and removing them or masking them with either a generic [MASK] token or more specific tokens such as [NAME]. Scrubbing has been used for removing PII [101], but could also be used for other types of facts (e.g., "The capital of Germany is [CITY].").

Shi et al. [151] propose a fine-grained variant of differential privacy (see Sec. 10.1) termed S-DP that works on the token level. Instead of protecting entire training documents with DP, it only protects tokens that have been identified as sensitive. Their proposed method selectively adds privacy noise to gradients to which the protected tokens have contributed.

Removing memorized facts might be particularly important in some jurisdictions like the EU [42] and California [124] that codify a right to be forgotten for individuals whose personal data has been used by a business. Meng et al. [113, 114] build on their hypothesis about key–values stores and propose a method to selectively change facts via values in those stores.

A general difficulty with preventing memorization of unstructured PII is that it requires a model with a deeper understanding of relationships within the training documents than for structured PII. Eldan and Russinovich [40] demonstrate a technique for unlearning data from a particular source. The technique is based on a reference model fine-tuned on the source data, along with generic text

auto-generated with GPT-4. The KGA framework [175] uses a process of unlearning to make the model's performance on target data similar to unseen data, while maintaining overall performance. This approach has been successful with more abstract data sources, such as specific personas in chat-based datasets. Bayazit et al. [13] develop a method to identify "knowledge-critical" subnetworks within GPT-2 models, using a joint loss function and demonstrate its effectiveness with concepts like 'fruit' and 'swimming'.

6 Ideas and algorithms

Ideas are similar to facts in that there are multiple ways to express them in language. Ideas can be about the physical world such as reinforcing concrete with steel, as well as about fiction, such as a story around a boy whose parents got killed when he was small and who learns wizardry at a secret school. Algorithms are a special type of idea that can be memorized in two ways: (1) as a description of a series of steps, similar to how the idea for the plot above can be memorized as a sequence of events; and (2) as a behavior, i.e., the LLM implements the algorithm and executes it in its forward pass. The former may occur via algorithm implementations in training data, the latter via input–output examples [121]. We have not found any research on whether this also works the other way around. In this section, we only cover simple algorithms such as arithmetic operations that do not require a high level of generalization. The distinction between facts and ideas can sometimes be difficult. 'Harry Potter is a wizard [...]' is a fact of literature, and likewise the steps of photosynthesis a biological fact.

6.1 Definitions

Tuple completion and linkability. For certain types of ideas, a definition similar to tuple completion or linkability for facts (Sec. 5.1) might be suitable. For ideas consisting of multiple elements (e.g., multiple plot elements) or algorithms consisting of multiple steps, the model could be prompted with all but one element or step. Then one could check whether the true missing element or step is predicted by the model, and could define this as a memorization of the idea or algorithm. As with facts, though, care would have to be taken to ensure that a correct prediction is not due to generalization abilities of the model.

Input–output behavior. Whether an LLM has memorized an algorithm in its behavior can be defined by whether it behaves correctly on all algorithm inputs [121]. If the input space is too large, the requirement can be reduced to a correct output on a selected sample of inputs, as in the algorithm tasks[2] of the BIG-bench benchmark [158].

6.2 Implications

Model performance and alignment. As with facts, one often wants the model to memorize ideas such as solutions to common problems or common tropes in fiction [187]. Similarly, one might want the model to be able to perform algorithms such as arithmetic manipulations [148], or return an implementation of dynamic programming [21, 75]. On the other hand, it is often undesirable if the model learns and reproduces harmful ideas. Meta considers three risk categories for its Llama 2 model [166]: illicit and criminal activities; hateful and harmful activities; and unqualified advice. Examples for ideas from all of these categories are very likely to be found in a dataset scraped from the Internet: plans for robberies in fictional stories; ideas around the superiority of one race in online comments; or medical speculation by non-professionals in health forums.

Privacy. Brown et al. [19] discuss the importance of the context in which information is shared. This can apply to ideas as well. For example, an individual might share an idea that is harmful or discriminatory on the Internet in an ironic way or as part of a fictional story. An LLM could, when asked about the stance of the person on that issue, attribute this idea to them without mentioning the context, opening them up to ostracism.

---

[2]https://github.com/google/BIG-bench/blob/main/bigbench/benchmark_tasks/keywords_to_tasks.md#algorithms

Security and confidentiality. Companies hide new products from the public until they are officially presented. Partners in political coalitions often negotiate in private before presenting their compromise to the media. Using an LLM for research on technical questions around a new product or legal questions around political plans could lead to a leak of confidential ideas if the LLM's provider uses user-submitted queries for training, similar to how patentable ideas could be leaked.

Copyright. Although abstract ideas typically do not enjoy copyright protection in the US [41], there have been cases of fan fiction or use of fictional characters that constituted copyright violations [56]. Algorithms are not protected by copyright law [17, 91], but concrete implementations might be (see Sec. 4). Users might use an LLM for exploring or concretizing ideas that they consider patenting. If such user-submitted data is used in model training and the model memorizes these ideas, it might reveal them to other users. This could create prior art, making the granting of a patent impossible in many jurisdictions (EU: Article 54(3) EPC; US: 35 U.S.C. §102).

Auditing. Determining how often an LLM generates new ideas instead of reproducing ideas from its training data would help determine the suitability of the LLM for tasks such as creative writing. When using the algorithmic capabilities of an LLM, it is important to know whether it has just memorized many input–output examples, or has learned the underlying algorithm, which only happens after long enough training [121].

6.3 Detecting memorized ideas and algorithms

As mentioned in Sec. 6.1, techniques for detecting memorized facts (Sec. 5.3) might also apply to ideas. The BIG-bench benchmark [158] measures the capability of LLMs to perform common algorithms such as removing duplicates from a list of numbers or finding the longest common subsequence of two strings, by testing correctness of model responses on instances of these problems. Saxton et al. [148] synthetically generate mathematics problems from fields such as algebra and arithmetic. They train models on smaller instances (e.g., smaller numbers) and test them on larger instances to determine whether the models merely learn algorithms for the domain of the training examples, or learn the algorithms in their full generality. McCoy et al. [109] compare the performance of LLMs on the same task, but with different parameters that do not change the task difficulty.

Stolfo et al. [160] perform causal mediation analysis on LLMs [133] to get insights into how LLMs perform simple arithmetic operations. They change activation values to identify layers and neurons most responsible for those operations. In a similar manner, Wang et al. [174] identify the circuit in GPT-2 responsible for detecting the grammatical object in a certain class of sentences. Nanda et al. [121] train a transformer to perform modular addition. Via careful inspection, they identify the exact algorithm that the model uses to perform this task. Interestingly, they find that in the beginning of training the model memorizes the training examples; in a second training phase it learns to perform the general algorithm; and in a third phase it removes the memorized components.

6.4 Preventing memorization or extraction of ideas and algorithms

Supervised fine-tuning, RLHF, and safety context distillation are often used [166] to deter a model from generating ideas from certain categories like criminal activities. The latter is a form of fine-tuning aimed to make the model behave as if prompts were preceded by a prompt instructing the model to, e.g., only generate safe responses. While not explicitly designed to prevent the model from outputting *memorized* ideas from these categories, this can be a side effect.

7 Writing styles

LLMs like GPT-4 can write in rhymes or imitate the writing style of Shakespeare [21]. Microsoft recently announced the "Sound like me" feature for their Copilot, allowing the LLM to write in the user's style when drafting emails [157]. With writing style, we mean the linguistic definition of style [73]—everything about a text that goes beyond pure semantics, including the choice of words and sentence structures, the characteristic use of stylistic devices such as alliterations and metaphors, and the level of formality. We consider not only writing in natural languages, but also programming languages. There, with different styles we mean functionally equivalent pieces of code that differ in

stylistic features like the naming of variables, their capitalization, the use of software design patterns, and formatting differences like the use of tabs or spaces for indentation.

7.1 Definitions

Mixture distribution. Several authors [5, 122] suggest that LMs learn to separate agents (e.g., separated by different beliefs) in their training data. This concept is formalized by Wolf et al. [181], who describe a language model as a mixture over different probability distributions. We can instantiate this formalism for writing styles, where we describe the writing style of the model as a mixture of the writing styles of authors seen during training, one writing style per author (or, alternatively, one writing style per demographic group, e.g., age groups or speakers of a dialect [167, Sec. 4.1]). The probability $p(y)$ that the model assigns to a given string y can then be decomposed as $p(y) = \sum_{\phi \in \Phi} w_\phi p_\phi(y)$, where each p_ϕ corresponds to the writing style of one author in the training data and w_ϕ is the prevalence of that style in the model outputs. We can define the memorization of one author's writing style as the presence of this writing style in the mixture, i.e., that there is one p_ϕ that corresponds to this author's writing style. Wolf et al. show that the model can be made to behave according to any p_ϕ in the mixture by a suitably long prompt, given some technical conditions. In practice, a prompt such as 'Answer in the style of [NAME]' might suffice to isolate p_ϕ.

Authorship verification and attribution. To determine whether a model can imitate an author's writing style sufficiently well, an *authorship verification* (AV) or *authorship attribution* (AA) task could be used [168]. In AV, an adversary is given two texts and has to determine whether they are written by the same author. In AA, an adversary is given texts from different authors and has to determine for a separate text by which of those authors it has been written. The advantage of an adversary in AV or AA over random guessing could be used a measure of the memorization of a writing style.

7.2 Implications

Model performance and alignment. Reif et al. [145] show that LLMs are capable of augmented zero-shot style transfer, which can help in creative writing [145] or tasks like using an LLM to make an email draft sound more formal [99]. It is also common to directly ask a model to generate an output in a specific style [99]. On the other hand, models that memorize undesirable writing styles from their training data are prone to replicating them. Allen-Zhu and Li [3] show that when pretraining a transformer with text that contains grammar mistakes, the model can respond to prompts that contain faulty grammar with text that contains grammar mistakes as well. A similar phenomenon might occur with LLMs used as coding assistants in a codebase with low-quality code. The Common Crawl dataset[3], a popular source for pretraining data, contains hate speech [100], which models might imitate. Larger models are more likely to produce toxic content [156], maybe due to their larger ability to memorize toxic writing styles from the tails of the training distribution.

Privacy. An LLM that has memorized the writing styles of individuals could be used for authorship attribution [118, 168]. One might, for example, try to determine the author of an anonymous text, if documents written by that author under their name were part of the training data. Authorship attribution via pretrained transformer models has been explored before [12, 43, 168], although with smaller models like BERT that require fine-tuning on target authors' documents. LLMs that have seen documents from multiple authors already during pretraining might be usable off-the-shelf for authorship attribution, making this technique much more accessible.

Many LLMs will have seen instances of documents written by an author together with additional information about that author (columns by a journalist that contain biographical information about authors, etc.). They might hence be able to perform author profiling [14, 167], i.e., given a document, determine attributes of the author such as age or gender. This has recently been demonstrated using Reddit comments [159].

Security and confidentiality. If a model has memorized and is able to replicate an individual's writing style, it could be used for impersonation for social engineering attacks [119]. If the number

---
[3]https://commoncrawl.org/

of memorized writing styles is large due to the large number of documents seen during training, this process could be automated at scale [29, 55]. For these types of attacks, it might not even be necessary to exactly copy the individual's writing style, but rather just the writing style within the individual's peer group [167]. Social engineering with voice cloning is already common enough for the FTC to issue an alert [137], and could be made much worse with convincing style copying.

Copyright. A writing style itself is not protected by copyright, as Henderson et al. [56] state. However, they show with one example that a text in the writing style of an author whose general idea is close to that of a copyrighted work of that author may violate copyright. The model would have to memorize at least the general idea of a text in addition to the writing style for this type of copyright violation.

Auditing. Assuming that, like verbatim text, less frequent writing styles are less prone to being memorized, the set of memorized writing styles would be an indicator of underrepresented communities in the training data. As an example, Dodge et al. [33] found that in the popular C4 text corpus [140], African-American English and Hispanic-aligned English are disproportionately affected by the blocklist used for data cleaning.

7.3 Detecting memorized writing styles

If a model can successfully apply the style of an author to its output, this is evidence for the memorization of the author's style. This could be measured by existing authorship verification or attribution methods [168]. For more generic styles like 'formal' or 'poetic', researchers use zero-shot prompting [99], augmented zero-shot prompting (giving the model examples of other styles than the target one) or few-shot prompting [145].

7.4 Preventing memorization or extraction of writing styles

If a writing style's frequency affects memorization (as it is the case for Wikipedia entities [105]), one solution is to limit the amount of text per author in the training data. Soleiman et al. [156] demonstrate that fine-tuning on non-toxic human-written prompts can reduce toxicity, arguably a form of writing style. Likewise, toxicity can be reduced by RLHF [128, 170] or incorporating human feedback in the pretraining objective [83]. Ilharco et al. [62] fine-tune models specifically for undesirable behavior and subtract the weight difference from the original model to remove such behavior. Li et al. [94] propose disconnecting model components to minimize loss on training data and maximize loss on unwanted behavior data. Both techniques could be applied to selectively remove writing styles. Li et al.'s technique could be applied to eliminate an individual author's writing style by dividing their documents into useful content and unimportant content, the latter serving as data for unwanted behavior. Mireshghallah and Berg-Kirkpatrick [116] propose a VAE-based method for obfuscating the writing style of a text document by turning it into a generic style. While the authors design this method to prevent discrimination and bias, it might be applied to the training corpus of an LLM to prevent the memorization of specific writing styles.

8 Distributional properties of the training data

Distribution inference [54, 161], also known as *property inference* [8, 51], is concerned with the leakage of (sensitive) properties of a model's training distribution. For LLMs, such properties of interest could be the proportion of documents authored by a given gender, the percentage of hateful content, the type of preprocessing used or data sources used in training.

8.1 Definitions

Distribution inference. The degree to which a distributional property is memorized by a model architecture can be measured through a distribution inference game [54, 161], wherein an adversary is presented with a model that was trained on data sampled from a distribution with one specific value for the property, and has to guess that value. One then computes the advantage of the adversary over random guessing. This amount of leakage can also be quantified using the n-leaked metric [162],

which captures how many records from the training distribution would need to be provided to leak a comparable amount of information about that distribution.

8.2   Implications

Model performance and alignment. The entire premise of training is to obtain models that know something useful about the underlying training distribution, and generalize to similar yet unseen data. When the downstream task is well-defined, such as in regression or classification settings, one can—at least theoretically—make a clear distinction between properties of the training distribution that the model should learn and properties it can ignore [54]. However, general-purpose LLMs are not trained with a clearly demarcated set of downstream tasks in mind; the next-token-prediction objective is a proxy for a large, vaguely described set of tasks [138]. It is thus unclear which properties of the training distribution the model can safely ignore without impacting its utility.

Privacy. Distributional membership inference could be used to identify individuals who contributed data in stages 2 and 3 of the model training [54]. Distribution inference attacks for author inference have already been demonstrated for text classification models [112], and for fine-tuning data for LLMs [77]. Distribution inference might further be used to inform attacks for detecting other kinds of memorization [192].

Security and confidentiality. Distribution inference can be used to identify exact training data sources, enabling targeted poisoning attempts [23] via Sybil attacks. It may further be possible to glean information about data preprocessing techniques. In the ROOTS corpus [86] (the foundation for BLOOM [149]), for instance, code files are filtered out based on length and the percentage of alphanumeric characters. Finding good values for these hyperparameters requires compute or manual data inspection, making them worthwhile targets. Since these hyperparameters impact the training data distribution, a method for distribution inference might be able to reveal them.

Copyright. When describing the training data distribution as a mixture over different data sources, one distributional property is whether a particular source—such as a specific website—is part of that mixture. This is described by distributional membership inference [54]. Inferring that information from the model would allow an author to determine whether the model was potentially trained on their copyrighted documents.

Auditing. The utility of distribution inference attacks, apart from inferring sensitive properties, also lies in auditing models without access to actual training data, which may not be available to the auditor. For instance, such attacks can be used to determine whether the training data is biased in any way, which is important, since tasks like hiring decisions and news generation require unbiased models. There has also been a demonstration of utilizing distribution inference (along with cryptographic primitives) for external auditing [35] on classification models, focusing on attestation of gender and race-related properties.

8.3   Detecting memorized distributional properties

Most techniques for distribution inference rely on some form of shadow model training [8, 51, 162, 191], which would be prohibitively expensive for LLMs. While there are a few techniques that do not require shadow models (like a simple loss-based estimation [161]), the only settings in which such attacks demonstrate non-trivial leakage require data poisoning [28]. Additionally, all of these attacks make strong assumptions about the adversary's prior knowledge of the underlying training distribution, which is a general limitation of the framework. Distribution inference may be a realistic approach for learning about the data preprocessing, though. In a setting where the training data is known (e.g., a public dataset such as The Pile [52]), but not the preprocessing, an adversary can preprocess the same documents in different ways and track the loss of the target model on all variants. Intuitively, the loss should be lowest with the preprocessing used by the model trainer, even if the documents themselves were not part of the training data.

8.4 Preventing memorization of distributional properties

Learning techniques aimed at enhancing the privacy of individuals, such as differential privacy (Sec. 10.1), do not reduce distribution inference risk (and can, in fact, worsen it [162]). Hartmann et al. [54] propose general techniques for mitigating distribution inference attacks based on causal learning, specifically invariant risk minimization (IRM) [6]. In a different context, techniques inspired by IRM have already been successfully used for making language models robust to different ways of preprocessing HTML pages in their training data [135].

9 Alignment goals

To instill instruction-following, helpful, truthful and harmless behavior into the model, data generated by human labelers is used in training stages 2 and 3. Labelers write prompts and responses, and rate model outputs along various axes [128, 166], following guidelines provided by the model trainer. The effectiveness of the alignment training (see, e.g., [166]), is proof that at least high-level goals from the guidelines such as harmlessness and helpfulness are memorized by the model. Despite the guidelines, there is still disagreement between human labelers [128, 166]. Memorization from alignment training might thus not be limited to those guidelines—models might also memorize political, ethical and other opinions of labelers.

9.1 Definitions

Distribution inference. The guidelines provided by the model trainer can be seen as a prior for the distribution of the data produced by the human labelers, i.e., as a property of that distribution. One can ignore the pretraining by considering the training in stages 2 and 3 as training that is performed with a model initialized with the weights after stage 1. This leaves one with the standard distribution inference setting, where the distributional properties of interest are the alignment guidelines. Concretely, one could, for example, be interested in whether harmlessness is prioritized over helpfulness or whether the political guidelines to the labelers reflect a particular ideology.

Membership and attribute inference. When one is interested in the text written or ratings given by individual labelers, membership and attribute inference are more apt frameworks. After all, the datasets consisting of prompt–completion pairs (stage 2) and of human ratings (stage 3) are simply datasets with each record contributed by one individual. As opposed to the standard setting, each labeler contributes multiple records to these datasets. Based on whether one is interested in one or all of the records from a labeler, one can use the item-level definition of membership inference or extensions to the user level [92].

9.2 Implications

Model performance and alignment. Supervised fine-tuning and RLHF or other methods based on human intervention [53, 139, 190] are currently essential for training aligned LLMs. While part of the alignment training can be automated [10, 126, 166], the alignment goals still need to be specified. It also seems unavoidable that in order to, e.g., be able to refuse to follow certain requests, the model needs to memorize at least high-level alignment guidelines. However, it seems desirable to have a model that *only* memorizes the alignment goals but not the data of individual human labelers, in order to avoid any bias resulting from the make up of the group of labelers.

Privacy. Preference ratings can reveal sensitive information about the labelers. For instance, truthfulness ratings of model outputs like "Taiwan is a sovereign country" might reveal political opinions. Similarly, the stance taken by a labeler in their response to the instruction "Write an essay about whether US Americans should have the right to bear arms" used for supervised fine-tuning. The often small number of labelers (e.g., fewer than 20 for Llama 2 [166]) could make them easier targets for privacy attacks, since their individual contributions have a larger impact. It is also not uncommon to have their names mentioned in acknowledgements [128, 166].

Security and confidentiality. Knowing the labeling guidelines or the reward function used for RLHF might help adversaries spot weaknesses in the safety alignment and eventually circumvent it.

Stealing the reward function could also be interesting for competitors who want to train models and avoid the cost of human labelers.

Copyright. Human labelers are likely so go through a formal process where they hand over the rights for their data to their employer. Copyright for this data, thus, may not be an issue. For instance, the participation agreement of Amazon Mechanical Turk [4] states that "all ownership rights, including all intellectual property rights, will vest with that Requester".

Auditing. Analyzing the reward function can reveal model behaviors and properties the model trainer prioritizes. This should be reflected in the instructions given to the labelers, and subsequently in the reward function used for RLHF. Access to the reward function would thus enable an auditor to verify claims of a model trainer, for example that they prioritized harmlessness over helpfulness during model training.

9.3 Detecting memorized alignment goals

Alignment training happens in the later training stages, making it more prone to detection and extraction attacks than pretraining data [46, 66]. Since the training mode for supervised fine-tuning is similar to that for pretraining, attacks aimed at extracting verbatim text memorized during pretraining (Sec. 4.3) might also work for the human-written responses to prompts used in stage 2. Different methods might be necessary for the prompts themselves, since the model is usually not directly trained to reproduce the prompts.

Regarding stage 3 training, it can be possible to reconstruct the reward function used for RLHF up to an additive, prompt-dependant constant when given access to the model after the stage 2 training, by using Eq. 5 of Rafailov et al. [139]. If the model after stage 2 is not publicly available, but based on a publicly available pretrained model—this is, e.g., the case for Llama 2 [166] and Mistral 7B [71] — one might try to approximate it through supervised fine-tuning by using one's own or public fine-tuning data such as the Flan Collection [98], used for LLaMA and Llama 2. Access to the reward function could be used to make inferences about labeler guidelines, for example, by checking whether harmful but helpful outputs are systematically higher ranked than harmless but non-helpful outputs. Individual labelers' data could be attacked using membership inference [184] or attribute inference [68] attacks. Reuter and Schulze [146] train a classifier to predict whether ChatGPT will refuse to answer a given prompt—behavior that is most likely predominantly learned in training stages 2 and 3.

9.4 Preventing memorization or extraction of alignment goals

The small number of human labelers might facilitate memorization of individual labelers' data, so increasing their number might reduce this risk. Not releasing the model after stage 1 or 2 will make it harder for an adversary to determine which stage model behavior originates from—this does not prevent memorization (since the model is not changed in any way), but might help against specific attacks.

10 General memorization mitigation strategies

This section briefly discusses general-purpose strategies related to preventing memorization that are not specific to any particular type of memorization. Differential privacy (DP) (Sec. 10.1) and near access-freeness (NAF) (Sec. 10.2) are two frameworks originally designed to solve privacy and copyright problems, respectively, which aim at preventing certain forms of memorization (DP) or reproduction of memorized information (NAF). They have several shortcomings for preventing different types of memorization in LLMs though, as we discuss next. Elkien-Koren et al. [41] argue in a very similar way why the frameworks are not suitable for preventing copyright violations. Sec. 10.3 discusses strategies that aim to mitigate copyright infringement with techniques that are external to the model.

10.1 Differential privacy

Differential privacy (DP) [37] is a privacy definition that can be satisfied by mechanisms where noise randomly sampled from appropriate distributions is incorporate in the training process to bound the impact of any individual record. Because of the random noise, the trained parameters will be nearly indistinguishable whether or not any one record was part of the training data. This makes it impossible for the model to remember any one training record. DP thus prevents the memorization of verbatim text, at least as long as this text is only contained in one training document.

For other types of memorization, DP mechanisms can, however, be both underexhaustive and overexhaustive. Underexhaustive because facts or writing styles may be present in many different training documents, books can be contained indirectly in the training data through quotes, reviews, etc. [27]; and overexhaustive because even if one only wants to prevent the model from memorizing specific pieces of information such as PII or facts, DP adds noise to all information contained in a training document.

While underexhaustiveness could be addressed by adjusting the unit of privacy so DP mechanisms will provide privacy with respect to multiple records [36], this would still require identifying the number of documents that contain a particular item (which might be expressed in different ways), and might significantly worsen the DP-induced performance drop due to larger amounts of added noise. Even for DP at the level of individual documents, El-Mhamdi et al. [39] argue that high-dimensional differentially private learning on heterogeneous data such as Internet text is impossible with high accuracy due to mean estimation being impossible under these conditions. The method by Shi et al. [151] (Sec. 5.4) that applies DP more selectively at a token level might help with the problem of overexhaustiveness, but requires exactly identifying the relevant tokens (e.g., those that contain a fact), but is not easily applicable to memorization which is not concentrated in a few tokens of a document (e.g., a writing style).

10.2 Near access-freeness

Vyas et al. [169] propose the notion of near access-free (NAF) generative models, aimed at preventing copyright violations. Given an subset $\mathcal{C}$ of its training documents, a model p is NAF with respect to $\mathcal{C}$ if for every $C \in \mathcal{C}$ the difference between the model's output distribution and the output distribution of a specific model that was not trained on C is bounded by a fixed constant. If C only occurs once in the training data, this ensures that p is unlikely to output substantial parts of C, unless those parts could have also been produced without access to C—which would not constitute a copyright violation. Vyas et al. also provide an extension for multiple occurrences.

Beyond verbatim text, NAF with the right distance metric and sufficiently small bound on the distance could also prevent the model from outputting PII such as social security numbers, or facts, as long as one knows how often they are contained in the training data. However, NAF has the same problems of under- and overexhaustiveness as DP. It is challenging to determine in how many documents a piece of information is contained, and prone to removing benign information that is unique to one or a few documents but that it is desirable for the model to learn. For the latter, consider the case where C is a book. Then a model without access to C would have a very low probability of correctly answering questions about the plot of C, so an NAF model would not be able to answer those questions either. Note that, unlike DP which is typically applied to the training process, NAF restricts the output distribution of the model. It can therefore be used to give guarantees for memorized information in the model outputs, but not for the memorization itself, i.e., someone with access to the model weights might still be able to detect memorized information.

10.3 Strategies for mitigating copyright violations via infrastructure

In some cases, copyright violations may be prevented by preventing particular types of memorization, as discussed in the previous sections. There have also been some proposals for mitigating copyright violations through appropriate infrastructure.

The most straightforward way to avoid copyright violations is by training only on data with permissive licenses [50, 81]. However, excluding data with non-permissive or unspecified licenses could significantly reduce the amount of available data [115] and might exclude high-quality data sources such as textbooks.

19

Min et al. [115] propose training a LLM only on permissively licensed documents, and augmenting it with a datastore of copyrighted documents. This datastore can then be used at prediction time to improve LLM performance on domains not covered by the training data. This setup allows for precisely pinpointing which copyrighted documents contributed to a particular model output, and for easily removing copyrighted documents if required. Determining copyrighted documents that were used in producing a model output via this and other methods [82] could allow for using documents with licenses that require author attribution [56]. Ippolito and Yu [64] describe a protocol similar to `robots.txt` wherein website owners could signal to model trainers via a file in the root directory which parts of their website are appropriate for model training, a variant of which has already been implemented by OpenAI [127] and Google [26]. In addition to deploying block requests [164], the New York Times has updated its terms of service to explicitly ban the use of its data for training ML models [178].

11   Conclusions and open research questions

In this paper, we have provided a taxonomy of memorization that goes beyond verbatim text. We have discussed challenges with defining memorization, and the consequences of memorization for various domains. Throughout the paper, we have identified several gaps in the research literature, and proposed ideas for future work. We highlight three research directions that we deem particularly important:

- Many existing definitions of memorization are specific to known attacks, and not the other way around. For example, the definitions that involve prompting the model with a prefix are adapted to the auto-regressive training of models, and, while useful for measuring a certain notion of memorization, do not cover adversaries that know a suffix of a document and want to extract a prefix.

- We did not find any research looking at memorization of data from the supervised fine-tuning and RLHF phase, even though these phases seem to contain the most severe memorization risks. In particular, it would be important to know whether user-submitted prompts that are used in these training phases can be extracted from the model.

- We want to minimize the amount of undesirable information memorized by LLMs due to the risks outlined in this paper. However, it is not clear which information LLMs need to memorize for good performance on downstream tasks, and the distinction between desirable and harmful learning is not well defined. This becomes an even more interesting question in light of the advent of LLMs enhanced with external knowledge, such as LLMs with Internet access.

Memorization in LLMs is an emerging area, and interactions between technical advances, demonstrated harms, and regulations and lawsuits will be necessary to achieve clarity on many of the issues raised in this paper. We expect this will be a long and evolving process, but one in which the research community should play an essential role.

References

[1] A. Akbik, T. Bergmann, D. Blythe, K. Rasul, S. Schweter, and R. Vollgraf. FLAIR: An easy-to-use framework for state-of-the-art NLP. In *North American Chapter of the Association for Computational Linguistics (demonstrations)*, 2019.

[2] B. AlKhamissi, M. Li, A. Celikyilmaz, M. Diab, and M. Ghazvininejad. A Review on Language Models as Knowledge Bases. *arXiv preprint arXiv:2204.06031*, 2022.

[3] Z. Allen-Zhu and Y. Li. Physics of language models: Part 1, context-free grammar. *arXiv preprint arXiv:2305.13673*, 2023.

[4] Amazon. Participation agreement. https://www.mturk.com/participation-agreement, 2023. Accessed: 2023-09-26.

[5] J. Andreas. Language models as agent models. In *Conference on Empirical Methods in Natural Language Processing*, 2022.

[6] M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.

[7] M. Assran, Q. Duval, I. Misra, P. Bojanowski, P. Vincent, M. Rabbat, Y. LeCun, and N. Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.

[8] G. Ateniese, L. V. Mancini, A. Spognardi, A. Villani, D. Vitali, and G. Felici. Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers. *International Journal of Security and Networks*, 2015.

[9] Y. Bai, A. Jones, K. Ndousse, A. Askell, A. Chen, N. DasSarma, D. Drain, S. Fort, D. Ganguli, T. Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.

[10] Y. Bai, S. Kadavath, S. Kundu, A. Askell, J. Kernion, A. Jones, A. Chen, A. Goldie, A. Mirhoseini, C. McKinnon, et al. Constitutional AI: Harmlessness from AI feedback. *arXiv preprint arXiv:2212.08073*, 2022.

[11] J. Bandy and N. Vincent. Addressing "documentation debt" in machine learning: A retrospective datasheet for BookCorpus. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.

[12] G. Barlas and E. Stamatatos. Cross-domain authorship attribution using pre-trained language models. In *Artificial Intelligence Applications and Innovations: 16th IFIP WG 12.5 International Conference*, 2020.

[13] D. Bayazit, N. Foroutan, Z. Chen, G. Weiss, and A. Bosselut. Discovering knowledge-critical subnetworks in pretrained language models, 2023.

[14] J. Bevendorff, B. Chulvi, E. Fersini, A. Heini, M. Kestemont, K. Kredens, M. Mayerl, R. Ortega-Bueno, P. Pęzik, M. Potthast, et al. Overview of PAN 2022: Authorship verification, profiling irony and stereotype spreaders, and style change detection. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, 2022.

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

[16] S. Black, L. Gao, P. Wang, C. Leahy, and S. Biderman. GPT-Neo: Large scale autoregressive language modeling with Mesh-Tensorflow, 2021.

[17] J. Bloch and P. Samuelson. Some misconceptions about software in the copyright literature. In *Symposium on Computer Science and Law*, 2022.

[18] O. Bousquet and A. Elisseeff. Stability and generalization. *The Journal of Machine Learning Research*, 2002.

[19] H. Brown, K. Lee, F. Mireshghallah, R. Shokri, and F. Tramèr. What does it mean for a language model to preserve privacy? In *ACM Conference on Fairness, Accountability, and Transparency*, 2022.

[20] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 2020.

[21] S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y. T. Lee, Y. Li, S. Lundberg, et al. Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv preprint arXiv:2303.12712*, 2023.

[22] N. Carlini, D. Ippolito, M. Jagielski, K. Lee, F. Tramer, and C. Zhang. Quantifying memorization across neural language models. In *International Conference on Learning Representations*, 2022.

[23] N. Carlini, M. Jagielski, C. A. Choquette-Choo, D. Paleka, W. Pearce, H. Anderson, A. Terzis, K. Thomas, and F. Tramèr. Poisoning web-scale training datasets is practical. *arXiv preprint arXiv:2302.10149*, 2023.

[24] N. Carlini, C. Liu, Ú. Erlingsson, J. Kos, and D. Song. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *USENIX Security Symposium*, 2019.

[25] N. Carlini, F. Tramer, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. Brown, D. Song, U. Erlingsson, et al. Extracting training data from large language models. In *USENIX Security Symposium*, 2021.

[26] G. S. Central. Overview of Google crawlers and fetchers (user agents), 2023. https://developers.google.com/search/docs/crawling-indexing/overview-google-crawlers.

[27] K. K. Chang, M. Cramer, S. Soni, and D. Bamman. Speak, memory: An archaeology of books known to ChatGPT/GPT-4. *arXiv preprint arXiv:2305.00118*, 2023.

[28] H. Chaudhari, J. Abascal, A. Oprea, M. Jagielski, F. Tramèr, and J. Ullman. SNAP: Efficient extraction of private properties with poisoning. In *IEEE Symposium on Security and Privacy*, 2023.

[29] Darktrace. Generative AI: Impact on email cyber-attacks. https://gwern.net/doc/cs/security/2023-darktrace.pdf, 2023.

[30] E. Debenedetti, G. Severi, N. Carlini, C. A. Choquette-Choo, M. Jagielski, M. Nasr, E. Wallace, and F. Tramèr. Privacy side channels in machine learning systems. *arXiv preprint arXiv:2309.05610*, 2023.

[31] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019.

[32] B. Dhingra, M. Faruqui, A. Parikh, M.-W. Chang, D. Das, and W. W. Cohen. Handling divergent reference texts when evaluating table-to-text generation. In *57th Annual Meeting of the Association for Computational Linguistics*, 2019.

[33] J. Dodge, A. Marasović, G. Ilharco, D. Groeneveld, M. Mitchell, and M. Gardner. Documenting large webtext corpora: A case study on the Colossal Clean Crawled Corpus. In *Conference on Empirical Methods in Natural Language Processing*, 2021.

[34] Q. Dong, J. Xu, L. Kong, Z. Sui, and L. Li. Statistical knowledge assessment for generative language models. *arXiv preprint arXiv:2305.10519*, 2023.

[35] V. Duddu, A. Das, N. Khayata, H. Yalame, T. Schneider, and N. Asokan. Attesting distributional properties of training data for machine learning. *arXiv preprint arXiv:2308.09552*, 2023.

[36] C. Dwork. Differential privacy. In *Automata, Languages and Programming*, 2006.

[37] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, 2006.

[38] N. Dziri, A. Madotto, O. Zaïane, and A. J. Bose. Neural Path Hunter: Reducing hallucination in dialogue systems via path grounding. In *Conference on Empirical Methods in Natural Language Processing*, 2021.

[39] E.-M. El-Mhamdi, S. Farhadkhani, R. Guerraoui, N. Gupta, L. N. Hoang, R. Pinot, and J. Stephan. On the impossible safety of large AI models. *arXiv preprint arXiv:2209.15259*, 2022.

[40] R. Eldan and M. Russinovich. Who's Harry Potter? Approximate unlearning in LLMs, 2023.

[41] N. Elkin-Koren, U. Hacohen, R. Livni, and S. Moran. Can copyright be reduced to privacy? *arXiv preprint arXiv:2305.14822*, 2023.

[42] European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union*, 2016.

[43] M. Fabien, E. Villatoro-Tello, P. Motlicek, and S. Parida. BertAA: BERT fine-tuning for authorship attribution. In *17th International Conference on Natural Language Processing*, 2020.

[44] M. Fan, C. Chen, C. Wang, and J. Huang. On the trustworthiness landscape of state-of-the-art generative models: A comprehensive survey. *arXiv preprint arXiv:2307.16680*, 2023.

[45] V. Feldman. Does learning require memorization? A short tale about a long tail. In *52nd Annual ACM SIGACT Symposium on Theory of Computing*, 2020.

[46] V. Feldman, I. Mironov, K. Talwar, and A. Thakurta. Privacy amplification by iteration. In *IEEE 59th Annual Symposium on Foundations of Computer Science*, 2018.

[47] V. Feldman and C. Zhang. What neural networks memorize and why: Discovering the long tail via influence estimation. In *Advances in Neural Information Processing Systems*, 2020.

[48] E. Ferrara. GenAI against humanity: Nefarious applications of generative artificial intelligence and large language models. *arXiv preprint arXiv:2310.00737*, 2023.

[49] M. Fredrikson, S. Jha, and T. Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *ACM SIGSAC Conference on Computer and Communications Security*, 2015.

[50] D. Fried, A. Aghajanyan, J. Lin, S. Wang, E. Wallace, F. Shi, R. Zhong, W.-t. Yih, L. Zettlemoyer, and M. Lewis. InCoder: A generative model for code infilling and synthesis. In *International Conference on Learning Representations*, 2023.

[51] K. Ganju, Q. Wang, W. Yang, C. A. Gunter, and N. Borisov. Property inference attacks on fully connected neural networks using permutation invariant representations. In *ACM SIGSAC Conference on Computer and Communications Security*, 2018.

[52] L. Gao, S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, et al. The Pile: An 800GB dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.

[53] C. Gulcehre, T. L. Paine, S. Srinivasan, K. Konyushkova, L. Weerts, A. Sharma, A. Siddhant, A. Ahern, M. Wang, C. Gu, et al. Reinforced self-training (ReST) for language modeling. *arXiv preprint arXiv:2308.08998*, 2023.

[54] V. Hartmann, L. Meynent, M. Peyrard, D. Dimitriadis, S. Tople, and R. West. Distribution inference risks: Identifying and mitigating sources of leakage. In *IEEE Conference on Secure and Trustworthy Machine Learning*, 2023.

[55] J. Hazell. Large language models can be used to effectively scale spear phishing campaigns. *arXiv preprint arXiv:2305.06972*, 2023.

[56] P. Henderson, X. Li, D. Jurafsky, T. Hashimoto, M. A. Lemley, and P. Liang. Foundation models and fair use. *arXiv preprint arXiv:2303.15715*, 2023.

[57] D. Hendrycks, C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2021.

[58] A. Holtzman, J. Buys, L. Du, M. Forbes, and Y. Choi. The curious case of neural text degeneration. In *8th International Conference on Learning Representations*, 2020.

[59] M. Honnibal, I. Montani, S. Van Landeghem, and A. Boyd. spaCy: Industrial-strength natural language processing in Python, 2020.

[60] J. Huang and K. C.-C. Chang. Towards reasoning in large language models: A survey. *arXiv preprint arXiv:2212.10403*, 2023.

[61] Y. Huang, Y. Li, W. Wu, J. Zhang, and M. R. Lyu. Do not give away my secrets: Uncovering the privacy issue of neural code completion tools. *arXiv preprint arXiv:2309.07639*, 2023.

[62] G. Ilharco, M. T. Ribeiro, M. Wortsman, S. Gururangan, L. Schmidt, H. Hajishirzi, and A. Farhadi. Editing models with task arithmetic. In *International Conference on Learning Representations*, 2023.

[63] D. Ippolito, F. Tramèr, M. Nasr, C. Zhang, M. Jagielski, K. Lee, C. A. Choquette-Choo, and N. Carlini. Preventing verbatim memorization in language models gives a false sense of privacy. *arXiv preprint arXiv:2210.17546*, 2022.

[64] D. Ippolito and Y. W. Yu. DoNotTrain: A metadata standard for indicating consent for machine learning. In *ICML Workshop on Generative AI and Law*, 2023.

[65] S. Ishihara. Training data extraction from pre-trained language models: A survey. *arXiv preprint arXiv:2305.16157*, 2023.

[66] M. Jagielski, O. Thakkar, F. Tramèr, D. Ippolito, K. Lee, N. Carlini, E. Wallace, S. Song, A. G. Thakurta, N. Papernot, and C. Zhang. Measuring forgetting of memorized training examples. In *International Conference on Learning Representations*, 2023.

[67] N. Jain, K. Saifullah, Y. Wen, J. Kirchenbauer, M. Shu, A. Saha, M. Goldblum, J. Geiping, and T. Goldstein. Bring your own data! Self-supervised evaluation for large language models. *arXiv preprint arXiv:2306.13651*, 2023.

[68] B. Jayaraman and D. Evans. Are attribute inference attacks just imputation? In *ACM SIGSAC Conference on Computer and Communications Security*, 2022.

[69] B. Jayaraman, L. Wang, K. Knipmeyer, Q. Gu, and D. Evans. Revisiting membership inference under realistic assumptions. *Privacy Enhancing Technologies*, 2021.

[70] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, and P. Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 2023.

[71] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. d. l. Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, et al. Mistral 7B. *arXiv preprint arXiv:2310.06825*, 2023.

[72] Z. Jiang, F. F. Xu, J. Araki, and G. Neubig. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 2020.

[73] D. Jin, Z. Jin, Z. Hu, O. Vechtomova, and R. Mihalcea. Deep learning for text style transfer: A survey. *Computational Linguistics*, 2022.

[74] M. Joshi, E. Choi, D. S. Weld, and L. Zettlemoyer. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In *Association for Computational Linguistics*, 2017.

[75] M. Josifoski, L. Klein, M. Peyrard, Y. Li, S. Geng, J. P. Schnitzler, Y. Yao, J. Wei, D. Paul, and R. West. Flows: Building blocks of reasoning and collaborating AI. *arXiv preprint arXiv:2308.01285*, 2023.

[76] M. Josifoski, M. Peyrard, F. Rajic, J. Wei, D. Paul, V. Hartmann, B. Patra, V. Chaudhary, E. Kıcıman, B. Faltings, and R. West. Language model decoding as likelihood–utility alignment. In *Conference of the European Chapter of the Association for Computational Linguistics (Findings)*, 2023.

[77] N. Kandpal, K. Pillutla, A. Oprea, P. Kairouz, C. A. Choquette-Choo, and Z. Xu. User inference attacks on large language models. *arXiv preprint arXiv:2310.09266*, 2023.

[78] N. Kandpal, E. Wallace, and C. Raffel. Deduplicating training data mitigates privacy risks in language models. In *International Conference on Machine Learning*, 2022.

[79] U. Khandelwal, O. Levy, D. Jurafsky, L. Zettlemoyer, and M. Lewis. Generalization through memorization: Nearest neighbor language models. In *International Conference on Learning Representations*, 2020.

[80] S. Kim, S. Yun, H. Lee, M. Gubri, S. Yoon, and S. J. Oh. ProPILE: Probing privacy leakage in large language models. *arXiv preprint arXiv:2307.01881*, 2023.

[81] D. Kocetkov, R. Li, L. B. allal, J. LI, C. Mou, Y. Jernite, M. Mitchell, C. M. Ferrandis, S. Hughes, T. Wolf, D. Bahdanau, L. V. Werra, and H. de Vries. The Stack: 3 TB of permissively licensed source code. *Transactions on Machine Learning Research*, 2023.

[82] P. W. Koh and P. Liang. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, 2017.

[83] T. Korbak, K. Shi, A. Chen, R. Bhalerao, C. L. Buckley, J. Phang, S. R. Bowman, and E. Perez. Pretraining language models with human preferences. In *International Conference on Machine Learning*, 2023.

[84] A. Kulkarni. GitHub Copilot AI is leaking functional API keys. https://analyticsdrift.com/github-copilot-ai-is-leaking-functional-api-keys/, 2021. Accessed: 2023-10-10.

[85] T. Kwiatkowski, J. Palomaki, O. Redfield, M. Collins, A. Parikh, C. Alberti, D. Epstein, I. Polosukhin, J. Devlin, K. Lee, K. Toutanova, L. Jones, M. Kelcey, M.-W. Chang, A. M. Dai, J. Uszkoreit, Q. Le, and S. Petrov. Natural Questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 2019.

[86] H. Laurençon, L. Saulnier, T. Wang, C. Akiki, A. Villanova del Moral, T. Le Scao, L. Von Werra, C. Mou, E. González Ponferrada, H. Nguyen, et al. The BigScience ROOTS corpus: A 1.6TB composite multilingual dataset. *Advances in Neural Information Processing Systems*, 2022.

[87] J. Lee, T. Le, J. Chen, and D. Lee. Do language models plagiarize? In *ACM Web Conference*, 2023.

[88] K. Lee, D. Ippolito, A. Nystrom, C. Zhang, D. Eck, C. Callison-Burch, and N. Carlini. Deduplicating training data makes language models better. In *Association for Computational Linguistics*, 2022.

[89] N. Lee, W. Ping, P. Xu, M. Patwary, P. N. Fung, M. Shoeybi, and B. Catanzaro. Factuality enhanced language models for open-ended text generation. *Advances in Neural Information Processing Systems*, 2022.

[90] E. Lehman, S. Jain, K. Pichotta, Y. Goldberg, and B. C. Wallace. Does BERT pretrained on clinical notes reveal sensitive data? In *North American Chapter of the Association for Computational Linguistics*, 2021.

[91] M. A. Lemley and B. Casey. Fair learning. *Texas Law Review*, 2020.

[92] D. Levy, Z. Sun, K. Amin, S. Kale, A. Kulesza, M. Mohri, and A. T. Suresh. Learning with user-level privacy. *Advances in Neural Information Processing Systems*, 2021.

[93] H. Li, D. Guo, W. Fan, M. Xu, and Y. Song. Multi-step jailbreaking privacy attacks on ChatGPT. *arXiv preprint arXiv:2304.05197*, 2023.

[94] M. Li, X. Davies, and M. Nadeau. Circuit breaking: Removing model behaviors with targeted ablation. *ICML Workshop on Deployment Challenges for Generative AI*, 2023.

[95] N. Li, W. Qardaji, D. Su, Y. Wu, and W. Yang. Membership privacy: A unifying framework for privacy definitions. In *IEEE Symposium on Security and Privacy*, 2013.

[96] Y. Li, Y. Li, and A. Risteski. How do transformers learn topic structure: Towards a mechanistic understanding. In *International Conference on Machine Learning*, 2023.

[97] S. C. Lin, J. Hilton, and O. Evans. TruthfulQA: Measuring how models mimic human falsehoods. In *Annual Meeting of the Association for Computational Linguistics*, 2021.

[98] S. Longpre, L. Hou, T. Vu, A. Webson, H. W. Chung, Y. Tay, D. Zhou, Q. V. Le, B. Zoph, J. Wei, et al. The Flan Collection: Designing data and methods for effective instruction tuning. In *International Conference on Machine Learning*, 2023.

[99] A. Lu, H. Zhang, Y. Zhang, X. Wang, and D. Yang. Bounding the capabilities of large language models in open text generation with prompt constraints. In *Findings of the Association for Computational Linguistics*, 2023.

[100] A. S. Luccioni and J. D. Viviano. What's in the box? An analysis of undesirable content in the Common Crawl corpus. In *Annual Meeting of the Association for Computational Linguistics*, 2021.

[101] N. Lukas, A. Salem, R. Sim, S. Tople, L. Wutschitz, and S. Zanella-Béguelin. Analyzing leakage of personally identifiable information in language models. In *IEEE Symposium on Security and Privacy*, 2023.

[102] K. MacDonald, K. Stuart, and A. Hern. Grand Theft Auto 6 leak: who hacked Rockstar and what was stolen?, 2022. https://www.theguardian.com/games/2022/sep/19/grand-theft-auto-6-leak-who-hacked-rockstar-and-what-was-stolen.

[103] P. Maini, M. C. Mozer, H. Sedghi, Z. C. Lipton, J. Z. Kolter, and C. Zhang. Can neural network memorization be localized? In *International Conference on Machine Learning*, 2023.

[104] P. Maini, M. Yaghini, and N. Papernot. Dataset inference: Ownership resolution in machine learning. In *International Conference on Learning Representations*, 2020.

[105] A. Mallen, A. Asai, V. Zhong, R. Das, H. Hajishirzi, and D. Khashabi. When not to trust language models: Investigating effectiveness and limitations of parametric and non-parametric memories. In *Association for Computational Linguistics*, 2022.

[106] K. S. I. Mantri and N. Sasikumatm. Developing methods for identifying and removing copyrighted content from generative AI models. In *ICML Workshop on Generative AI and Law*, 2023.

[107] J. Mattern, F. Mireshghallah, Z. Jin, B. Schoelkopf, M. Sachan, and T. Berg-Kirkpatrick. Membership inference attacks against language models via neighbourhood comparison. In *Findings of the Association for Computational Linguistics*, 2023.

[108] C. Mauran. Whoops, Samsung workers accidentally leaked trade secrets via ChatGPT. https://mashable.com/article/samsung-chatgpt-leak-details, 2023. Accessed: 2023-10-08.

[109] R. T. McCoy, S. Yao, D. Friedman, M. Hardy, and T. L. Griffiths. Embers of autoregression: Understanding large language models through the problem they are trained to solve. *arXiv preprint arXiv:2309.13638*, 2023.

[110] I. R. McKenzie, A. Lyzhov, M. Pieler, A. Parrish, A. Mueller, A. Prabhu, E. McLean, A. Kirtland, A. Ross, A. Liu, et al. Inverse scaling: When bigger isn't better. *arXiv preprint arXiv:2306.09479*, 2023.

[111] M. Meli, M. R. McNiece, and B. Reaves. How bad can it Git? Characterizing secret leakage in public GitHub repositories. In *Network and Distributed Systems Security Symposium*, 2019.

[112] L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov. Exploiting unintended feature leakage in collaborative learning. In *IEEE Symposium on Security and Privacy*, 2019.

[113] K. Meng, D. Bau, A. Andonian, and Y. Belinkov. Locating and editing factual associations in GPT. *Advances in Neural Information Processing Systems*, 2022.

[114] K. Meng, A. S. Sharma, A. Andonian, Y. Belinkov, and D. Bau. Mass-editing memory in a transformer. In *International Conference on Learning Representations*, 2023.

[115] S. Min, S. Gururangan, E. Wallace, H. Hajishirzi, N. A. Smith, and L. Zettlemoyer. SILO language models: Isolating legal risk in a nonparametric datastore. *arXiv preprint arXiv:2308.04430*, 2023.

[116] F. Mireshghallah and T. Berg-Kirkpatrick. Style pooling: Automatic text style obfuscation for improved classification fairness. In *Conference on Empirical Methods in Natural Language Processing*, 2021.

[117] J. Mökander, J. Schuett, H. R. Kirk, and L. Floridi. Auditing large language models: a three-layered approach. *AI and Ethics*, 2023.

[118] F. Mosteller and D. L. Wallace. Inference in an authorship problem: A comparative study of discrimination methods applied to the authorship of the disputed Federalist Papers. *Journal of the American Statistical Association*, 1963.

[119] M. Mozes, X. He, B. Kleinberg, and L. D. Griffin. Use of LLMs for illicit purposes: Threats, prevention measures, and vulnerabilities. *arXiv preprint arXiv:2308.12833*, 2023.

[120] R. Nakano, J. Hilton, S. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders, et al. WebGPT: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.

[121] N. Nanda, L. Chan, T. Liberum, J. Smith, and J. Steinhardt. Progress measures for grokking via mechanistic interpretability. In *International Conference on Learning Representations*, 2023.

[122] C. Nardo. The Waluigi effect. https://www.lesswrong.com/posts/D7PumeYTDPfBTp3i7/the-waluigi-effect-mega-post, 2023. Accessed: 2023-07-24.

[123] L. Niu, S. Mirza, Z. Maradni, and C. Pöpper. CodexLeaks: Privacy leaks from code generation language models in GitHub Copilot. In *32nd USENIX Security Symposium*, 2023.

[124] S. of California Department of Justice. California consumer privacy act (CCPA), 2018. https://oag.ca.gov/privacy/ccpa.

[125] T. Onishi, H. Wang, M. Bansal, K. Gimpel, and D. McAllester. Who did what: A large-scale person-centered cloze dataset. In *Conference on Empirical Methods in Natural Language Processing*, 2016.

[126] OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

[127] OpenAI. GPTBot, 2023. https://platform.openai.com/docs/gptbot.

[128] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 2022.

[129] M. Ozdayi, C. Peris, J. FitzGerald, C. Dupuy, J. Majmudar, H. Khan, R. Parikh, and R. Gupta. Controlling the extraction of memorized data from large language models via prompt-tuning. In *Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2023.

[130] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, and X. Wu. Unifying large language models and knowledge graphs: A roadmap. *arXiv preprint arXiv:2306.08302*, 2023.

[131] S. M. Park, K. Georgiev, A. Ilyas, G. Leclerc, and A. Madry. TRAK: Attributing model behavior at scale. In *International Conference on Machine Learning*, 2023.

[132] V. Patil, P. Hase, and M. Bansal. Can sensitive information be deleted from LLMs? Objectives for defending against extraction attacks, 2023.

[133] J. Pearl. Direct and indirect effects. In *Uncertainty in Artificial Intelligence*, 2001.

[134] F. Petroni, T. Rocktäschel, S. Riedel, P. S. H. Lewis, A. Bakhtin, Y. Wu, and A. H. Miller. Language models as knowledge bases? In *Conference on Empirical Methods in Natural Language Processing*, 2019.

[135] M. Peyrard, S. Ghotra, M. Josifoski, V. Agarwal, B. Patra, D. Carignan, E. Kiciman, S. Tiwary, and R. West. Invariant language modeling. In *Conference on Empirical Methods in Natural Language Processing*, 2022.

[136] N. Pörner, U. Waltinger, and H. Schütze. E-BERT: Efficient-yet-effective entity embeddings for BERT. In *Findings of the Association for Computational Linguistics*, 2020.

[137] A. Puig. Scammers use AI to enhance their family emergency schemes. https://consumer.ftc.gov/consumer-alerts/2023/03/scammers-use-ai-enhance-their-family-emergency-schemes, 2023. Accessed: 2023-09-26.

[138] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever. Improving language understanding by Generative Pre-Training, 2018.

[139] R. Rafailov, A. Sharma, E. Mitchell, S. Ermon, C. D. Manning, and C. Finn. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*, 2023.

[140] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 2020.

[141] N. Rahman and E. Santacana. Beyond fair use: Legal risk evaluation for training LLMs on copyrighted text. In *ICML Workshop on Generative AI and Law*, 2023.

[142] O. Ram, Y. Levine, I. Dalmedigos, D. Muhlgay, A. Shashua, K. Leyton-Brown, and Y. Shoham. In-context retrieval-augmented language models. *arXiv preprint arXiv:2302.00083*, 2023.

[143] H. Rashkin, V. Nikolaev, M. Lamm, L. Aroyo, M. Collins, D. Das, S. Petrov, G. S. Tomar, I. Turc, and D. Reitter. Measuring attribution in natural language generation models. *Computational Linguistics*, 2023.

[144] V. Raunak, A. Menezes, and M. Junczys-Dowmunt. The curious case of hallucinations in neural machine translation. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2021.

[145] E. Reif, D. Ippolito, A. Yuan, A. Coenen, C. Callison-Burch, and J. Wei. A recipe for arbitrary text style transfer with large language models. In *Association for Computational Linguistics*, 2022.

[146] M. Reuter and W. Schulze. I'm afraid I can't do that: Predicting prompt refusal in black-box generative language models. *arXiv preprint arXiv:2306.03423*, 2023.

[147] A. Salem, G. Cherubin, D. Evans, B. Köpf, A. Paverd, A. Suri, S. Tople, and S. Zanella-Béguelin. SoK: Let the privacy games begin! A unified treatment of data inference privacy in machine learning. In *IEEE Symposium on Security and Privacy*, 2023.

[148] D. Saxton, E. Grefenstette, F. Hill, and P. Kohli. Analysing mathematical reasoning abilities of neural models. In *International Conference on Learning Representations*, 2019.

[149] T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilić, D. Hesslow, R. Castagné, A. S. Luccioni, F. Yvon, M. Gallé, et al. BLOOM: A 176B-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022.

[150] M. Schade. How your data is used to improve model performance. https://help.openai.com/en/articles/5722486-how-your-data-is-used-to-improve-model-performance, 2023. Accessed: 2023-10-09.

[151] W. Shi, A. Cui, E. Li, R. Jia, and Z. Yu. Selective differential privacy for language modeling. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2022.

[152] W. Shi, S. Min, M. Yasunaga, M. Seo, R. James, M. Lewis, L. Zettlemoyer, and W.-t. Yih. RE-PLUG: Retrieval-augmented black-box language models. *arXiv preprint arXiv:2301.12652*, 2023.

[153] R. Shokri, M. Stronati, C. Song, and V. Shmatikov. Membership inference attacks against machine learning models. In *IEEE Symposium on Security and Privacy*, 2017.

[154] G. Siboni and D. Siman-Tov. Cyberspace extortion: North Korea versus the United States. 2014.

[155] V. Smith, A. S. Shamsabadi, C. Ashurst, and A. Weller. Identifying and mitigating privacy risks stemming from language models: A survey. *arXiv preprint arXiv:2310.01424*, 2023.

[156] I. Solaiman and C. Dennison. Process for adapting language models to society (PALMS) with values-targeted datasets. *Advances in Neural Information Processing Systems*, 2021.

[157] J. Spataro. Announcing Microsoft 365 Copilot general availability and Microsoft 365 Chat. https://www.microsoft.com/en-us/microsoft-365/blog/2023/09/21/announcing-microsoft-365-copilot-general-availability-and-microsoft-365-chat/, 2023. Accessed: 2023-10-01.

[158] A. Srivastava, A. Rastogi, A. Rao, A. A. M. Shoeb, A. Abid, A. Fisch, A. R. Brown, A. Santoro, A. Gupta, A. Garriga-Alonso, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Transactions on Machine Learning Research*, 2022.

[159] R. Staab, M. Vero, M. Balunović, and M. Vechev. Beyond memorization: Violating privacy via inference with large language models. *arXiv preprint arXiv:2310.07298*, 2023.

[160] A. Stolfo, Y. Belinkov, and M. Sachan. Understanding arithmetic reasoning in language models using causal mediation analysis. *arXiv preprint arXiv:2305.15054*, 2023.

[161] A. Suri and D. Evans. Formalizing and estimating distribution inference risks. *Privacy Enhancing Technologies Symposium*, 2022.

[162] A. Suri, Y. Lu, Y. Chen, and D. Evans. Dissecting distribution inference. In *IEEE Conference on Secure and Trustworthy Machine Learning*, 2023.

[163] S. Szyller, R. Zhang, J. Liu, and N. Asokan. On the robustness of dataset inference. *Transactions on Machine Learning Research*, 2023.

[164] N. Tiku. Newspapers want payment for articles used to power ChatGPT. *The Washington Post*, Oct 2023.

[165] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al. LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

[166] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

[167] E. Troiano, A. Velutharambath, and R. Klinger. From theories on styles to their transfer in text: Bridging the gap with a hierarchical survey. *Natural Language Engineering*, 2023.

[168] J. Tyo, B. Dhingra, and Z. C. Lipton. On the state of the art in authorship attribution and authorship verification. *arXiv preprint arXiv:2209.06869*, 2022.

[169] N. Vyas, S. Kakade, and B. Barak. Provable copyright protection for generative models. In *International Conference on Machine Learning*, 2023.

[170] B. Wang, W. Chen, H. Pei, C. Xie, M. Kang, C. Zhang, C. Xu, Z. Xiong, R. Dutta, R. Schaeffer, et al. DecodingTrust: A comprehensive assessment of trustworthiness in GPT models. *arXiv preprint arXiv:2306.11698*, 2023.

[171] B. Wang and A. Komatsuzaki. GPT-J-6B: A 6 billion parameter autoregressive language model, 2021. https://github.com/kingoflolz/mesh-transformer-jax.

[172] C. Wang and R. Sennrich. On exposure bias, hallucination and domain shift in neural machine translation. In *58th Annual Meeting of the Association for Computational Linguistics*, 2020.

[173] H. Wang. Revisiting challenges in data-to-text generation with fact grounding. In *International Conference on Natural Language Generation*, 2020.

[174] K. Wang, A. Variengien, A. Conmy, B. Shlegeris, and J. Steinhardt. Interpretability in the wild: a circuit for indirect object identification in GPT-2 small. In *International Conference on Learning Representations*, 2023.

[175] L. Wang, T. Chen, W. Yuan, X. Zeng, K.-F. Wong, and H. Yin. KGA: A general machine unlearning framework based on knowledge gap alignment. *arXiv preprint arXiv:2305.06535*, 2023.

[176] T. Wang, Y. Zhang, and R. Jia. Improving robustness to model inversion attacks via mutual information regularization. In *AAAI Conference on Artificial Intelligence*, 2021.

[177] Y. Wang, W. Zhong, L. Li, F. Mi, X. Zeng, W. Huang, L. Shang, X. Jiang, and Q. Liu. Aligning large language models with human: A survey. *arXiv preprint arXiv:2307.12966*, 2023.

[178] J. Weatherbed. The New York Times prohibits using its content to train AI models, 2023. https://www.theverge.com/2023/8/14/23831109/the-new-york-times-ai-web-scraping-rules-terms-of-service.

[179] J. Wei, M. Bosma, V. Y. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, and Q. V. Le. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*, 2022.

[180] O. Weller, M. Marone, N. Weir, D. Lawrie, D. Khashabi, and B. Van Durme. "According to..." prompting language models improves quoting from pre-training data. *arXiv preprint arXiv:2305.13252*, 2023.

[181] Y. Wolf, N. Wies, Y. Levine, and A. Shashua. Fundamental limitations of alignment in large language models. *arXiv preprint arXiv:2304.11082*, 2023.

[182] X. Wu, M. Fredrikson, S. Jha, and J. F. Naughton. A methodology for formalizing model-inversion attacks. In *IEEE 29th Computer Security Foundations Symposium*, 2016.

[183] Z. Yang, L. Dong, X. Du, H. Cheng, E. Cambria, X. Liu, J. Gao, and F. Wei. Language models as inductive reasoners. *arXiv preprint arXiv:2212.10923*, 2022.

[184] J. Ye, A. Maddi, S. K. Murakonda, V. Bindschaedler, and R. Shokri. Enhanced membership inference attacks against machine learning models. In *ACM Conference on Computer and Communications Security (CCS)*, 2022.

[185] S. Yeom, I. Giacomelli, M. Fredrikson, and S. Jha. Privacy risk in machine learning: Analyzing the connection to overfitting. In *Computer Security Foundations Symposium*, 2018.

[186] Z. Yu, Y. Wu, N. Zhang, C. Wang, Y. Vorobeychik, and C. Xiao. CodeIPPrompt: Intellectual property infringement assessment of code language models. In *International Conference on Machine Learning*, 2023.

[187] A. Yuan, A. Coenen, E. Reif, and D. Ippolito. Wordcraft: story writing with large language models. In *27th International Conference on Intelligent User Interfaces*, 2022.

[188] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 2021.

[189] C. Zhang, D. Ippolito, K. Lee, M. Jagielski, F. Tramèr, and N. Carlini. Counterfactual memorization in neural language models. *arXiv preprint arXiv:2112.12938*, 2021.

[190] T. Zhang, F. Liu, J. Wong, P. Abbeel, and J. E. Gonzalez. The wisdom of hindsight makes language models better instruction followers. In *International Conference on Machine Learning*, 2023.

[191] W. Zhang, S. Tople, and O. Ohrimenko. Leakage of dataset properties in multi-party machine learning. In *USENIX Security Symposium*, 2021.

[192] J. Zhou, Y. Chen, C. Shen, and Y. Zhang. Property inference attacks against GANs. In *Network and Distributed Systems Security Symposium*, 2022.