

Can LLMs get help from other LLMs without revealing private information?

Florian Hartmann*
Google Research
fhartmann@google.com

Duc-Hieu Tran
Google Research
hieuz@google.com

Peter Kairouz
Google Research
kairouz@google.com

Victor Cărbune
Google Research
vcarbune@google.com

Blaise Aguera y Arcas
Google Research
blaisea@google.com

Abstract

Cascades are a common type of machine learning systems in which a large, remote model can be queried if a local model is not able to accurately label a user’s data by itself. Serving stacks for large language models (LLMs) increasingly use cascades due to their ability to preserve task performance while dramatically reducing inference costs. However, applying cascade systems in situations where the local model has access to sensitive data constitutes a significant privacy risk for users since such data could be forwarded to the remote model. In this work, we show the feasibility of applying cascade systems in such setups by equipping the local model with privacy-preserving techniques that reduce the risk of leaking private information when querying the remote model. To quantify information leakage in such setups, we introduce two privacy measures. We then propose a system that leverages the recently introduced social learning paradigm in which LLMs collaboratively learn from each other by exchanging natural language. Using this paradigm, we demonstrate on several datasets that our methods minimize the privacy loss while at the same time improving task performance compared to a non-cascade baseline.

1 Introduction

Large language models (LLMs) such as Gemini Ultra by Google (2023) and GPT-4 by OpenAI (2023) are reporting remarkable performance on many tasks. These models, however, not only come with high inference costs, but they also have to run in data centers far from the local contexts where private data is available. Conversely, models that can run in private contexts, such as Gemini Nano, have more limited capabilities since they run on the user’s device.

*Corresponding author.

To unlock state-of-the-art performance in private contexts, local models with access to sensitive data need to be equipped with a privacy-preserving mechanism that enables querying a remote model without sharing any sensitive data. Although standard cascade systems in which a smaller, less capable model, queries a larger, much more capable one in order to solve a task have previously been studied (Yue et al., 2024; Chen et al., 2023), privacy-preserving ones have not yet been explored. In today’s cascade systems, the decision of whether a larger model should be leveraged or not is usually done through an additional mechanism that determines whether the query can be handled by the smaller model independently (Li et al., 2021). If determined to be handled by the larger model, the query is simply forwarded without consideration for the private data it may contain. This poses privacy threats for users, ranging from leaking sensitive data to the forwarded sample even being ingested in training datasets of the remote system.

We introduce the first privacy-preserving approach to cascade systems. Contrasting to standard cascade systems, our local model always assumes its data is private. As such, the local model should not share anything private with the remote model. Going one step further, even if the local model does not verbatim share private information, we aim to prevent a curious remote model operator from reconstructing private data by utilizing auxiliary information it might have. To focus on these challenges, we assume there are no efficiency constraints and that the local model can always ask for help from the remote model, as shown in Figure 1. Therefore our optimal cascade setup consists of minimizing the privacy loss while maximizing task performance, where an upper bound is given by querying the teacher model with the actual data, although private.

To succeed at this task, the local model, typically smaller and less capable, needs to find the right




Figure 1: The local model, the *student*, wants to label its private data. It can query a larger, remote model, the *teacher*, to get help. The student may not reveal private data to the teacher.

balance between revealing sufficient information about the problem to receive useful signals from the more capable, remote model while keeping details private. To enable learning from the remote model, the local model makes use of gradient-free learning capabilities through natural language that in-context learning (ICL) capabilities of LLMs enable (Brown et al., 2020). Throughout, we leverage the recently introduced social learning paradigm by (Mohtashami et al., 2023; Bandura and Walters, 1977) in which LLMs learn through natural language from other LLMs.

Contributions We summarize our contributions as follows: **(i)** we enable cascade systems to be used where access to private data is necessary to solve a task, but cannot be revealed **(ii)** we design and evaluate algorithms that sanitize private data while still leveraging in-context learning capabilities of private models and **(iii)** whereas previous work to the best of our knowledge analyzes settings without auxiliary information, we go one step further by considering auxiliary information and proposing a novel metric to this end **(iv)** we perform extensive experiments on a diverse range of tasks, quantifying task performance and impact on privacy using standardized measures.

2 Problem Setting

Our paper considers a variant of social learning (Mohtashami et al., 2023) where neither of the participants has any labeled data. A local model, called the *student*, has private data that it cannot label well by itself. A larger, remote model, called the *teacher*, can do a better job at labeling the data. These two models form a cascade, in which the student can improve its performance by communicating with the teacher. We call what the student sends to the teacher a *query*. Figure 1 shows a visualization of this setup.

Constraints There are two constraints on the queries from student to teacher. **(i)** The communication must be privacy-preserving, i.e. the student may not copy over its data and must not reveal anything private. **(ii)** There is only a single round of communication between student and teacher, meaning neither of them can maintain any state or update a model of the other’s capabilities.

To this end, all algorithms follow the same structure. Given (0) that the student needs help, it then (1) uses its private data to generate a query to the teacher. In turn, (2) the teacher uses the query to generate ICL examples for the student. Finally, (3) the student uses the ICL examples to go back to solving its original problem.

Simplifying assumptions To better focus on the challenges we aim to address in this paper, we furthermore make two simplifying assumptions. **(i)** We assume that communication with the remote teacher is always helpful. This assumption is reasonable because existing techniques for determining delegation in cascades, discussed in Section 6, could be combined with our methods. **(ii)** We also assume that both student and teacher are aware of the format of the examples, as shown in Table 4 in the appendix. Such an assumption is useful because we want the student to learn more complex things about the data from the teacher instead of simply learning a format or chain of thought prompt.

Given this problem setting, the goal of the student is to maximize its performance in correctly labeling its data while not revealing anything private that is part of said data.

3 Privacy Measures

The student’s data may often contain sensitive personal information that should be kept hidden from an untrusted, or partially trusted, teacher. For example, consider a query that tries to figure out what disease could best explain a set of health symptoms

that are experienced by a user after they have engaged in a specific sequence of activities. Here, being able to associate the set of activities and/or symptoms with a specific user is a privacy violation that we would like to eliminate.

To address such privacy violations, one might be tempted to resort to data anonymization techniques, such as *differential privacy* (DP) (Dwork, 2006). However, these techniques are most useful when computing aggregates across many users (e.g. average of gradient vectors computed on a batch of sensitive training examples). Using the local model of DP (Warner, 1965; Evfimievski et al., 2003; Kasiviswanathan et al., 2011) as a mechanism to mask private information in the query will end up masking both private and non-private information in the query, rendering the masked query useless for the task at hand. Alternatively, using the ICL model of DP (Liu et al., 2024; Wu et al., 2023) to privatize the sensitive parts of the student’s data suffers a major hurdle: it assumes the student has many private examples it can jointly consider when creating a query to the teacher. While we do look into grouping examples when generating a query in Section 4.4, we expect the student to have very few private examples, and want it to be able to generate privacy-preserving queries even when only having a single, private example. DP-ICL cannot work in such a setting.

Instead, we leverage data minimization privacy techniques, specifically *contextual integrity* (Nissenbaum, 2004) which describes privacy as an appropriate flow of information. Under this technique, the student would keep information that is useful for the task (e.g. the activities & symptoms in the above-mentioned example) but remove any personally identifying information that is not relevant to the query context. We note that even under perfect masking, this approach could still leak sensitive information, should the teacher model have access to auxiliary information that can be used to identify certain unique features that are strongly correlated with the “perfectly masked” prompts (Narayanan and Shmatikov, 2008; Sweeney, 2002). Thus, an important contribution of our work is a methodology for measuring and assessing leakage under auxiliary information.

The success of our approach hinges on correctly identifying and masking the sensitive parts of the query without tampering with the description of the task. To this end, we propose various techniques that can analyze information in queries to

produce safe queries that can be shared with the teacher model. To assess the privacy of the queries, we consider two concrete metrics, the *entity leak* metric that counts entities that exist in both original examples and the student’s queries, and the *mapping leak* metric that considers a setting with auxiliary information.

Entity leak metric Contextual integrity states that privacy is the appropriate flow of information. For most production applications, it is hard to say what is appropriate to share. As a proxy for this, we consider the interpretable metric of leaked entities. All entities, such as names, locations, email addresses, or numbers, in the dataset, are considered to be private. We measure how many of the entities in the original example are still part of the student’s query upon masking.

Mapping leak metric Even if all entities are removed from the student’s query, it is still possible for a curious teacher to reconstruct private information by carefully analyzing the query. Indeed, auxiliary information that the teacher may have access to can help it be more effective at this. We measure how well the teacher could do this through a worst-case analysis. More precisely, we assume the teacher is presented with 1 original example and 100 masked queries out of which exactly one was generated from the original example. We measure how often the teacher is able to correctly map the original example to this particular (masked) query out of the 100 options. Providing the teacher with a complete original example represents an upper bound on the auxiliary information the teacher could have. To do a better job at this mapping, we allow the teacher to query the student model, which is useful since it was used to generate the masked query. To conduct the mapping, we then score continuations of the original example and the 100 generated queries, and measure how often the correct query scores the highest. We show that access to such (worst-case) auxiliary information could lead to non-trivial privacy leakage even when the entities are properly masked.

4 Methods

We consider three algorithms for how the student could privately learn from the teacher, as shown in Figure 2. The first of these methods is based on the student describing the problem it is facing while the latter two methods generate similar, non-

private examples that the teacher can label. As a hyperparameter for all these methods, we consider the *expansion size* to denote how many labeled ICL examples the student will receive from the teacher.

4.1 Method 1: Creating a problem description

As an initial approach, we consider a method in which the student analyzes the problem it is given and generates a high-level description from this problem. Even if the student cannot solve the problem, it might be able to describe the type of problem it is facing. This description is the query to the teacher.

The teacher in turn wants to create few-shot examples that the student can use to solve the problem it is facing. Since the teacher has access to a template about the example structure, it knows what format to follow. To create such examples, it then uses this template as well as the student’s description to create *expansion size* many new examples.

4.2 Method 2: Generating new unlabeled examples

Instead of providing the teacher with an abstract description of the problem it is facing, the student can generate a similar, but novel problem itself. As a motivating example for why this is a sensible choice, consider GSM8k (Cobbe et al., 2021), a math dataset with problems of US middle school difficulty. Given such a math problem, it is possible to create a similar math problem that is just as educational but contains none of the same details, i.e. both problems follow a similar structure and are of similar difficulty.

Previous work has shown that LLMs are able to generate new examples from original examples that they see in-context (Shao et al., 2023; Mohtashami et al., 2023). We additionally observe that for many tasks it is easier to generate new examples than it is to solve them, meaning it is possible for the student model to synthesize similar, unlabeled examples, even if it does not do a good job at labeling them.

To this end, our second method works as follows: We (1) prompt the student LLM to generate *expansion size* new unlabeled examples. Then, (2) the teacher receives these examples and labels them. Finally, (3) the student learns in-context from that and tries to solve the original problem. Throughout, both teacher and student models utilize the task template to understand what format the labeled and

unlabeled examples follow and for where step-by-step explanations make sense.

4.3 Method 3: Replacing entities in original examples

Instead of instructing the student to generate completely novel examples, we can also ask it to keep the same example while replacing names, locations, numbers, and other entities. The student then generates a new unlabeled example that is very similar to the original but that contains none of the private information. Since there are many ways to replace the entities, we can again generate *expansion size* examples using this technique.

While this could be done using a specialized entity detection model and rule-based systems, we observe that LLMs do a fairly good job at this themselves. Thus, we decide to simply prompt the student model to find and replace private entities. The full flow of this method is the same as the method in Section 4.2, except that in step (1), we now simply replace entities instead of generating completely new examples.

4.4 Grouping unlabeled examples to reduce teacher calls

Each call to the teacher implies some chance of leaking private information. This chance needs to be traded off with how much the student can be improved through this process. Like in active learning, the teacher in our setting can thus be considered an expensive resource that needs to be used economically.

To utilize this resource well, we introduce an additional hyperparameter *group size* that denotes how many private examples the student groups together in order to create *expansion size* many ICL examples through the teacher. The student considers the entire group jointly when synthesizing descriptions and new unlabeled examples, and is thus able to combine information from the grouped, private examples. By *labeling budget = expansion size / group size*, we denote the budget of how many teacher labeled examples may be created for each original example. Note that the student does not get to choose which examples to group together.

5 Experiments

In order to evaluate the effectiveness of the methods introduced in Section 4, we evaluate them in terms of accuracy and privacy on a diverse group of datasets and compare them against two baselines.

0) <code>student</code> needs help to solve a problem		
Question: Raul had \$87 to spare so he decided to go to the bookshop. Raul bought 8 comics, each of which cost \$4. How much money does Raul have left?		
Method 1: problem description	Method 2: new problem	Method 3: masked problem
1) <code>student</code> generates a high-level description of the problem	1) <code>student</code> generates a new but similar kind of problem	1) <code>student</code> copies the example but replaces entities with different ones
This is a basic arithmetic problem involving multiplication and subtraction.	Question: A store sells a shirt for \$20. The store offers a 20% discount on the shirt. How much does the shirt cost after the discount?	Question: Emily had \$92 to spend at the ice cream shop. She bought 4 ice cream cones, each of which cost \$3. How much money does Emily have left?
2) <code>teacher</code> creates examples based on it	2) <code>teacher</code> solves the similar problem	2) <code>teacher</code> solves the similar problem
Question: Calculate $9 \times 8 - 25$ Answer: Multiply 9 by 8 to get 72. Subtract 25 from 72 to get the answer. #### 47	Answer: Step 1: Calculate the discount amount: 20% of \$20 = \$4 Step 2: Subtract the discount amount from the original price: \$20 - \$4 = \$16 #### 16	Answer: 1. Calculate the total cost of the ice cream cones: 4 cones * \$3/cone = \$12. 2. Subtract the cost of the cones from the initial amount: \$92 - \$12 = \$80 #### 80
3) <code>student</code> learns in-context from that labeled example and attempts again to solve the original problem		

Figure 2: The three methods we consider. Steps 1 and 2 show actual student queries and teacher responses as generated in our experiments when using Gemini 1.0 Nano-2 as the student and Gemini 1.0 Ultra as the teacher. Note that each method generates increasingly specific queries about the student’s problem.

Models We use the *Gemini 1.0* family of models (Google, 2023) for all of our experiments. As the teacher, we utilize *Ultra*, the most powerful model of the family. In most of our experiments, *Nano-2*, a 3.5B parameter model that can be deployed on mobile phones, is the student. The student model capabilities naturally influence the performance of our method and hence we also run experiments when *Pro* is the student. In line with previous reports on Nano’s performance (Google, 2023), we normalize task success in all our experiments by the teacher’s performance since it is an upper bound for what we can hope to achieve.

Datasets We consider a variety of datasets in our experiments to demonstrate that our methods generalize across a suite of tasks: GSM8k math problems (Cobbe et al., 2021), assistant intent detection (Srivastava et al., 2022), classifying whether statements are subjective or objective (Conneau and Kiela, 2018) and mid-resource machine translation (Tiedemann, 2020). See Appendix B for a more detailed description of the datasets.

Baselines We compare against a *weak* and *strong* baseline. For the *weak baseline*, we consider a student that does not communicate with the teacher at all. Since the student does not have any labeled data on its own, it thus falls back to the 0-shot setting while still being able to use the task’s template. As the *strong baseline*, we evaluate a student that

has access to 8 arbitrary, golden examples. We consider this to be a strong baseline since these examples are perfectly labeled and for the same task that the student is trying to solve. In practice, such data does often not exist and cannot be easily matched to the student’s problem.

5.1 Task Performance

To evaluate our methods, we run experiments for all above mentioned datasets. For ease of comparison, we consider the 8-shot performance of each method. Table 1 shows these results. Across all datasets, we outperform both the weak and strong baseline. However, we note that for GSM8k, getting close to 100% task success, as normalized by the teacher’s performance, requires Pro as a strong student model.

We observe that method 3 performs very well across all datasets. Likely this is because the queries generated by this method are the closest to the problem that the student aims to solve. Method 1 performs the worst. We find this method to be the hardest to get to work well since for some tasks, e.g. intent recognition, the student model is only able to explicitly describe the unlabeled example if it is also able to label it, rendering it a less competitive method.

Furthermore, to investigate the best use of the *labeling budget = expansion size / group size*, we run full grid searches for the different methods.

Dataset	Student	Weak Baseline: 0-shot	Strong Baseline: Golden Data 8-shot	Method 1: Descriptions 8-shot	Method 2: New Problems 8-shot	Method 3: Replacing 8-shot
GSM8k	Nano-2	11.3%	34.9%	36.7%	45.6%	55.9%
	Pro	85.4%	91.1%	78.6%	91.6%	98.3%
Intent Recognition	Nano-2	70.9%	92.4%	82.7%	92.3%	94.6%
Subj	Nano-2	55.6%	74.2%	74.2%	71.0%	79.7%
Translation en → eu	Nano-2	70.8%	72.9%	72.8%	74.8%	91.0%

Table 1: Task performance with Gemini 1.0 Nano-2 and Pro as students, and Gemini 1.0 Ultra as the teacher. All values are normalized by the teacher’s performance as reported in Table 5. For easier comparison, we only consider setups with *expansion size* = 8, *group size* = 1 here. Note that we report BLEURT (Sellam et al., 2020) for machine translation and accuracy for all other tasks. Appendix D.1 shows similar machine translation results for 6 more languages.

For each *labeling budget*, we then obtain the best performance that can be reached. As shown in Figure 3, the choice of these hyperparameters allows budgets below 1, which is not possible without grouping.

5.2 Privacy Results

To analyze how our methods fare in terms of privacy, we compute the two metrics mentioned in Section 3. We find entities in both the original example and in the query by asking Gemini 1.0 Ultra to play an entity detector that finds entities such as names, locations, numbers, and anything else one might consider private. We manually verify on a subset of examples that this does indeed find the desired entities. The results in Table 2 show the results of this analysis.

Unexpectedly, we observe that method 1 often leaks the most entities. While this method should generate the most high-level queries in theory, it is hard to get to work well in practice. On a subset of original examples, the student is not able to synthesize a high-level description and instead defaults to detailedly describing the problem it is facing. While the queries of method 3 are the closest to the original messages, they also leak the fewest entities. We hypothesize that this is because the student does not need to understand the problem well in order to find and replace entities. It can simply consider the individual tokens and replace them without needing to understand what kind of problem it is trying to get help on.

However, when analyzing the mapping metric,

which describes a worst case of how well an attacker with auxiliary information can identify original examples, the results paint a different picture. Here, method 3 performs significantly worse. While few entities leak in this method, the structure and writing style are maintained, making it especially easy to map between original and generated example. This is in particular the case for the GSM8k and Subj datasets in which examples have a distinct structure that makes them easy to identify.

We find grouping examples to work particularly well with method 2. We observe that with *group size* = 2 leaks in both metrics significantly reduce, and in the case of GSM8k even without a drop in performance.

Finally, we note that the choice of the right method depends on the concrete threat model considered. While method 1 is neither convincing in terms of quality and privacy, method 3 works remarkably well in situations where the threat model does not involve auxiliary information. Conversely, if one does consider auxiliary information, method 2, potentially with grouping, is the most appropriate to use.

5.3 Qualitative Analysis

In order to better understand where our methods help and where they fall short, we run detailed analyses on the predictions that the student model is able to make after it got help from the teacher. To do this at scale, we ask Gemini 1.0 Ultra to look at the golden label and the student’s prediction, and




Figure 3: For a given *labeling budget = expansion size / group size*, we show the accuracy reached. Grouping allows us to improve the 29.0% accuracy reached through *expansion size = group size = 1* to an accuracy of 32.5% all while just using $\frac{1}{8}$ of the budget. Furthermore, even for budgets above 1, we can well outperform the approach without grouping.

Dataset	Metric	Method 1:	Method 2:	Method 2:	Method 2:	Method 3:
		Descriptions 8-shot <i>group size = 1</i>	New Problems 8-shot <i>group size = 1</i>	New Problems 8-shot <i>group size = 2</i>	New Problems 8-shot <i>group size = 4</i>	Replacing 8-shot <i>group size = 1</i>
GSM8k	accuracy	36.7%	45.6%	45.7%	40.2%	55.9%
	entity leaks	2.7%	1.5%	0.7%	0.2%	1.2%
	mapping leak	16.4%	5.4%	2.9%	2.5%	53.8%
Intent Recognition	accuracy	82.7%	92.3%	89.9%	86.7%	94.6%
	entity leaks	5.7%	3.9%	0.7%	0.1%	0.5%
	mapping leak	1.8%	2.6%	1.0%	0.9%	3.2%
Subj	accuracy	74.2%	71.0%	64.2%	61.4%	79.7%
	entity leaks	4.3%	3.8%	1.3%	0.6%	1.4%
	mapping leak	16.6%	6.2%	3.2%	2.8%	43.3%
Translation en → eu	BLEURT	72.8%	74.8%	68.2%	69.0%	91.0%
	entity leaks	2.5%	1.3%	1.3%	0.0%	1.3%
	mapping leak	4.5%	3.5%	1.3%	1.2%	2.9%

Table 2: For Nano-2 as the student, and each dataset and method, we present our two privacy metrics. Method 3 generally achieves the best quality results while leaking few entities. Method 2 with grouping offers the strongest privacy metrics.

classify the errors into certain classes. We confirm manually for a subset of cases that these classifications make sense. Table 3 shows the results of this analysis for GSM8k based on 500 examples, for the strong baseline and the best setup for each of our methods. We show similar analyses for machine translation in Appendix D.2, as well as example student queries for all datasets in Appendix F.

6 Related Work

LLM Cascades Cascades were mostly studied for improving overall inference costs, particularly given ever-increasing LLM sizes (Hoffmann et al., 2022). Task performance steadily increases with parameter count (Schaeffer et al., 2023). Various approaches to cascade inference are compared in

(Miao et al., 2023). Some methods (Li et al., 2021; Chen et al., 2023) use a classifier to determine whether to forward a query or not, while more recent work (Yue et al., 2024) leverages a voting and consistency measure of the first model in the cascade as proxy for the inability to provide an answer. We replaced inference cost with a privacy measure optimization and quantified to what degree task performance can be preserved.

Differential Privacy (DP) DP formalizes privacy guarantees in a probabilistic framework (Dwork, 2006). This can be implemented in various ways, e.g. via the local model of DP (Warner, 1965; Evfimievski et al., 2003; Kasiviswanathan et al., 2011) or as part of in-context learning (Liu et al., 2024; Wu et al.,

Class	Strong Baseline: Golden Data 8-shot	Method 1: Descriptions 8-shot	Method 2: New Problems 8-shot	Method 3: Replacing 8-shot
Correct prediction	29.0%	32.2 %	40.0%	49.1 %
Calculation error	42.8%	35.7%	32.2%	26.8%
Flaw in reasoning	13.8%	10.8%	11.4%	8.1%
Using incorrect information	5.3%	9.1%	6.4%	6.5%
Incorrectly applying formulas	4.9%	5.1%	5.6%	4.3%
Not understanding the problem	4.3%	7.1%	4.3%	5.2%

Table 3: An analysis of the student’s predictions shows that calculation and reasoning errors of the students reduce through the ICL examples our method provides. Errors caused by using incorrect information slightly increase, likely because the student model can get confused by the similar examples it is seeing. We bold the best cell of each row to emphasize that method 3 shows the most impressive reduction in mistakes. Note that we do not normalize by teacher task success here as opposed to the other tables.

2023). While these techniques are useful when computing aggregates across many users, we want our system to work even when a user only has a single, private example, as explain in Section 3.

Data Minimization As an alternative to DP, we follow data minimization principles in the form of contextual integrity (Nissenbaum, 2004). Data minimization techniques are particularly important for removing sensitive information from LLM training datasets. (Lison et al., 2021) present an overview of many techniques relevant to enabling cascade systems in private/public setups. In this work, we investigated the effectiveness of masking operations, and instead of using a separate sequence tagging model we relied on the student LLM capability to perform such transformations. Recent studies, such as (Vats et al., 2023), have found that pre-training LLMs on datasets processed with privacy-preserving masking does not limit capabilities of models, while privacy benefits are strong.

Social Learning for LLMs (Mohtashami et al., 2023) propose the original framework that we expand here. Notable differences from that are (i) our student model can ask for help from the teacher model, (ii) additional teaching algorithms leveraging in-context learning with improved privacy metrics and (iii) showcasing how social learning can enable cascade systems in setups where they would otherwise not be usable.

Synthetic Datasets LLMs are effective at creating bootstrapping datasets, e.g. by creating task instructions through their own conditional gener-

ation (Wang et al., 2023). Similarly, (Lee et al., 2023) have shown how alignment data can be synthesized. The student model needs to have such bootstrapping capabilities and the richer this ability is, the better it produces diverse task transformations that the teacher can better use to explain it back.

7 Conclusion

In this paper, we investigated whether LLMs can privately query external LLMs to improve their performance. Indeed, we find that our methods comfortably beat strong baselines that have privacy constraints in place, even with Gemini 1.0 Nano-2 as the student, a 3.5B model that fits on phones.

To evaluate the privacy performance of our methods, we look at two metrics, a simple to interpret count of entities leaked, and another, novel, metric that measures an upper bound of what a curious teacher with auxiliary information could hope to recover from the student’s queries. For the first metric, we find masking problems (method 3) to work well, while generating new problems (method 2) with grouping does well in cases where the teacher can be expected to have auxiliary information.

Ultimately, we note that the choice of methods depends on the concrete threat model considered. For either threat model, we present a compelling system and analysis, which show that leakage can be low while beating strong quality baselines. Additionally, we show how grouping examples improves the privacy metrics, and can, under a given labeling budget constraint, even improve model quality.

Future work in this space could consider more complex forms of student-teacher interactions, further improve the privacy metrics established, and look into modalities other than text.

Limitations

While our work provides a compelling privacy analysis, consisting of an interpretable metrics based on entities and a worst-case, upper bound metric, we do not include methods with privacy guarantees. As discussed in Section 3, we do not find differential privacy to be the right notion here. However, one could consider other ways of potentially adding guarantees in the future.

A further limitation of our work is that we only study a single modality, text. Other modalities could be investigated going forward.

Finally, our work only studies the Gemini family of models. The combination of Nano, Pro and Ultra models provides interesting signal to how well LLMs can get help from other LLMs without revealing private information. However, with more budget to run experiments for different models, the experiments could be repeated for other model families.

Ethics Statement

Our work supports data minimization principles. It paves the way towards more data staying on users' devices while still offering them intelligent features based on machine learning.

Acknowledgments

We would like to express our gratitude towards Matt Sharifi, Tautvydas Misiunas, Hassan Mansoor, Dominik Roblek, Lukas Zilka, Yun Zhu, Jindong (JD) Chen, and Rif A. Saurous for providing crucial feedback on this paper. All your suggestions greatly improved this paper. Furthermore, we would also like to thank Amirkeivan Mohtashami whose contributions to the social learning code base remain useful long after his internship.

References

Albert Bandura and Richard H Walters. 1977. *Social learning theory*, volume 1. Englewood cliffs Prentice Hall.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda

Askeel, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.

Lingjiao Chen, Matei Zaharia, and James Zou. 2023. [FrugalGPT: How to use large language models while reducing cost and improving performance](#). *Preprint*, arXiv:2305.05176.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.

Alexis Conneau and Douwe Kiela. 2018. Senteval: An evaluation toolkit for universal sentence representations. *arXiv preprint arXiv:1803.05449*.

Cynthia Dwork. 2006. Differential privacy. In *International colloquium on automata, languages, and programming*, pages 1–12. Springer.

Alexandre Evfimievski, Johannes Gehrke, and Ramakrishnan Srikant. 2003. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 211–222.

Google. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. 2022. [Training compute-optimal large language models](#). *Preprint*, arXiv:2203.15556.

Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. 2011. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826.

Harrison Lee, Samrat Phatale, Hassan Mansoor, Thomas Mesnard, Johan Ferret, Kellie Lu, Colton Bishop, Ethan Hall, Victor Carbune, Abhinav Rastogi, and Sushant Prakash. 2023. [Rlaif: Scaling reinforcement learning from human feedback with ai feedback](#). *Preprint*, arXiv:2309.00267.

Lei Li, Yankai Lin, Deli Chen, Shuhuai Ren, Peng Li, Jie Zhou, and Xu Sun. 2021. [Cascadebert: Accelerating inference of pre-trained language models via calibrated complete models cascade](#). *Preprint*, arXiv:2012.14682.

- Pierre Lison, Ildikó Pilán, David Sanchez, Montserrat Batet, and Lilja Øvrelid. 2021. [Anonymisation models for text data: State of the art, challenges and future directions](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4188–4203, Online. Association for Computational Linguistics.
- Yi Liu, Gelei Deng, Yuekang Li, Kailong Wang, Zihao Wang, Xiaofeng Wang, Tianwei Zhang, Yepang Liu, Haoyu Wang, Yan Zheng, and Yang Liu. 2024. [Prompt injection attack against llm-integrated applications](#). *Preprint*, arXiv:2306.05499.
- Xupeng Miao, Gabriele Oliaro, Zhihao Zhang, Xinhao Cheng, Hongyi Jin, Tianqi Chen, and Zhihao Jia. 2023. [Towards efficient generative large language model serving: A survey from algorithms to systems](#). *Preprint*, arXiv:2312.15234.
- Amirkeivan Mohtashami, Florian Hartmann, Sian Gooding, Lukas Zilka, Matt Sharifi, et al. 2023. [Social learning: Towards collaborative learning with large language models](#). *arXiv preprint arXiv:2312.11441*.
- Arvind Narayanan and Vitaly Shmatikov. 2008. [Robust de-anonymization of large sparse datasets](#). In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pages 111–125. IEEE.
- Helen Nissenbaum. 2004. [Privacy as contextual integrity](#). *Wash. L. Rev.*, 79:119.
- OpenAI. 2023. [GPT-4 Technical Report](#). *Preprint*, arXiv:2303.08774.
- Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. 2023. [Are emergent abilities of large language models a mirage?](#) *Preprint*, arXiv:2304.15004.
- Thibault Sellam, Dipanjan Das, and Ankur P Parikh. 2020. [Bleurt: Learning robust metrics for text generation](#). *arXiv preprint arXiv:2004.04696*.
- Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. [Synthetic prompting: Generating chain-of-thought demonstrations for large language models](#). In *International Conference on Machine Learning*, pages 30706–30775. PMLR.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. 2022. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *arXiv preprint arXiv:2206.04615*.
- Latanya Sweeney. 2002. [k-anonymity: A model for protecting privacy](#). *International journal of uncertainty, fuzziness and knowledge-based systems*, 10(05):557–570.
- Jörg Tiedemann. 2020. [The tatoeba translation challenge—realistic data sets for low resource and multilingual mt](#). *arXiv preprint arXiv:2010.06354*.
- Arpita Vats, Zhe Liu, Peng Su, Debjyoti Paul, Yingyi Ma, Yutong Pang, Zeeshan Ahmed, and Ozlem Kalinli. 2023. [Recovering from privacy-preserving masking with large language models](#). *Preprint*, arXiv:2309.08628.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hananeh Hajishirzi. 2023. [Self-instruct: Aligning language models with self-generated instructions](#). *Preprint*, arXiv:2212.10560.
- Stanley L Warner. 1965. [Randomized response: A survey technique for eliminating evasive answer bias](#). *Journal of the American Statistical Association*, 60(309):63–69.
- Tong Wu, Ashwinee Panda, Jiachen T Wang, and Prateek Mittal. 2023. [Privacy-preserving in-context learning for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Murong Yue, Jie Zhao, Min Zhang, Liang Du, and Ziyu Yao. 2024. [Large language model cascades with mixture of thoughts representations for cost-efficient reasoning](#). *Preprint*, arXiv:2310.03094.

Appendix

A What Makes for a Good Student and Teacher?

To decide on promising experiments with our methods, we pose the question: what makes for a good student and teacher combination? We found the following criteria to be useful in deciding on student models and datasets.

Good student A student model is promising on a dataset, if **(i)** it can initially not solve the task well (0-shot), but **(ii)** is able to improve with in-context examples (golden 8-shot). Furthermore, **(iii)** the student model needs to be able to ask for help via a useful query to the teacher, e.g. it needs to be able to synthesize similar, unlabeled examples.

Good teacher A teacher model is a good fit for the student model if **(iv)** it can solve the task much better than the student, meaning even its 0-shot performance is significantly higher than the student’s golden 8-shot task success. Furthermore, **(v)** a good teacher needs to be able to respond to the student’s queries, e.g. by providing useful labels for them.

When evaluating whether a new dataset is promising to try with our method, we first check these five criteria.

B More Details on the Datasets

In this section we provide additional details on the four datasets we use. Table 4 shows the templates we use for each dataset.

Grade School Math GSM8k (Cobbe et al., 2021) is a dataset containing grade school math questions, annotated answers as well as step-by-step reasoning on how to reach the answer. Typical GSM8k examples are written in the form of a story with many entities that we do not want the student to reveal to the teacher.

Intent Recognition Cascade systems are especially useful for questions that users pose their personal assistant. Intent Recognition (Srivastava et al., 2022) is a dataset in which one has to classify an utterance as one of 7 assistant tasks, as shown in Table 4.

Subj The Subj dataset (Conneau and Kiela, 2018) consists of statements that are either subjective or objective. The model has to classify the statements as one of these two categories.

Machine Translation LLMs show remarkable machine translation performance. Since performance for high-resource languages is difficult to further improve via ICL, we focus on mid-resource machine translation on the Tatoeba (Tiedemann, 2020) dataset.

C Teacher Task Performance

In Tables 1 and 2 in the main text, we normalize the student’s task success by the teacher’s performance. In Table 5, we show this teacher task performance.

D More Machine Translation Results

D.1 Task Performance

For brevity’s sake, we only show results for one language pair in Table 1 of the main text. Table 6 shows the results for all seven languages we consider. Note that each time we translate from English each time since this allows the student model to synthesize useful queries to the teacher even though it does not understand the target language well.

We find our methods to work particularly well for mid-resource languages. Gemini Nano-2 al-

ready performs very well on high-resource languages, such as German and Vietnamese, even in the 0-shot setting. Though we do see a small improvement with our methods here, much bigger improvements can be achieved for mid-resource languages.

D.2 Qualitative Analysis

To better understand in which cases our techniques improve machine translation, we perform a qualitative analysis, similar to the one in Section 5.3. Tables 7 and 8 show the results of these analyses. We find most error types to significantly decrease with our methods, while the incorrect addition or omission of information slightly increases.

E A Student That Is Copying Instead of Learning In-Context

To evaluate how important ICL is in our setting, we ran additional experiments in which the student copies the teacher’s answer instead of learning from it in-context. For the case of *expansion size* > 1, the student copies the teacher’s most common answer.

We start by noting that such an approach does not satisfy the privacy constraint on many tasks. If a student were for example to achieve high task success on machine translation by simply copying the teacher’s answer, this would imply that the teacher learned the most important parts of the student’s original data.

Based on this observation, we stick to GSM8k, intent recognition and Subj for this analysis. To enable the student to achieve a good quality by copying, we rely on the masking approach introduced in Section 4.3. However, we additionally instruct the student to replace entities in a way that does not change the result. For the case of GSM8k, this means not replacing any numbers and leaving the relationship between any numbers intact.

We find that ICL outperforms copying in our experiments, as shown in Table 9. For intent recognition and Subj, copying works fairly well since there are only a few classes to cover. While most of the time, the examples generated by the student all belong to the same class, there are cases where the original example is close to two similar classes. We find ICL to help in these cases.

For GSM8k copying works much worse. This is even the case when using Pro as a significantly larger student. Looking at experiment logs, the

Dataset	Example Format
GSM8k	Question: <question> Answer: <step-by-step reasoning> #### <final number>
Intent Recognition	Utterance: <utterance> Intent: <add_to_playlist, book_restaurant, get_weather, play_music, search_screening_event, search_creative_work, rate_book>
Subj	Text: <text> Label: <subjective, objective>
Machine Translation	English sentence: <english sentence> Basque translation: <basque translation>

Table 4: The templates for the four datasets we consider. Teacher, student, and baselines can use this information in order to understand how to format examples and where step-by-step reasoning makes sense. This information can either be used in prompts or in constrained decoding configurations.

Dataset	Metric	Teacher n -shot	Teacher Task Success
GSM8k	accuracy	0	87.8%
Intent Recognition	accuracy	0	97.4%
Subj	accuracy	8	92.3%
Translation en \rightarrow el	BLEURT	0	90.6%

Table 5: Gemini 1.0 Ultra’s task success as the teacher. Even though the teacher itself is not 100% accurate, the student manages to improve through interaction with the teacher in our experiments. We use 0-shot for the teacher in most experiments, but fall back to 8-shot for Subj since this is a difficult task to do in a 0-shot setting.

student in this setup struggles to generate queries that do not affect the result.

Based on these results, we decide to stick to ICL for all other experiments, but use the results to influence our Subj prompt.

F Example Queries Our Methods Generate

Table 10 shows example student problems and queries that work well. In all of these examples, the student is able to generate a query to the teacher that does not verbatim leak sensitive information but that nevertheless allows the teacher to respond with useful examples.

In Table 11, we show examples in which the student does not generate a good query. In most of these cases, the student leaks sensitive information. In some, the student generates a query that does not make sense.

From	To	Weak Baseline: 0-shot	Strong Baseline: Golden Data 8-shot	Method 1: Descriptions 8-shot	Method 2: New Problems 8-shot	Method 3: Replacing 8-shot
English	German (de)	95.4%	96.4%	92.0%	97.6%	97.3%
English	Greek (el)	84.0%	88.0%	88.3%	88.9%	90.7%
English	Basque (eu)	70.8%	72.9%	72.8%	74.8%	91.0%
English	Hebrew (he)	81.0%	80.9%	69.1%	80.4%	86.4%
English	Georgian (ka)	45.5%	46.6%	36.3%	49.7%	64.4%
English	Tagalog (tl)	90.8%	89.7%	87.6%	90.9%	94.0%
English	Vietnamese (vi)	95.7%	95.0%	90.5%	97.5%	97.1%

Table 6: Machine translation performance (BLEURT) with Gemini 1.0 Nano-2 as the student and Gemini 1.0 Ultra as the teacher. All values are normalized by the teacher’s performance. We note that our methods significantly improve results for mid-resource languages while achieving a small improvement for high-resource languages that the student model already understands well.

Class	Strong Baseline: Golden Data 8-shot	Method 1: Descriptions 8-shot	Method 2: New Problems 8-shot	Method 3: Replacing 8-shot
Correct translation	38.4%	41.6%	40.8%	64.8%
Lexical or Semantic error	50.8%	40.4%	49.6%	27.6%
Grammatical error	6.0%	9.2%	4.0%	4.8%
Contextual or Cultural error	4.0%	5.2%	3.2%	0.4%
Omission or Incorrect Addition	0.8%	2.8%	2.4%	2.4%
Formatting error	0.0%	0.8%	0.0%	0.0%

Table 7: A qualitative error analysis for translation from English to Basque (eu). Lexical, semantic and contextual errors significantly decrease with our methods.

Class	Strong Baseline: Golden Data 8-shot	Method 1: Descriptions 8-shot	Method 2: New Problems 8-shot	Method 3: Replacing 8-shot
Correct translation	39.0%	31.8%	41.8%	50.0%
Lexical or Semantic error	39.4%	36.1%	37.5%	33.0%
Grammatical error	15.3%	12.9%	14.4%	10.6%
Contextual or Cultural error	5.0%	12.5%	5.5%	5.0%
Omission or Incorrect Addition	1.3%	3.2%	0.9%	1.2%
Formatting error	0.0%	3.4%	0.0%	0.2%

Table 8: A qualitative error analysis for translation from English to Greek (el). Lexical, semantic and grammatical errors significantly decrease with our methods.

Dataset	Student	Method 3: Replacing with copying	Method 3: Replacing with ICL
GSM8k	Nano-2	9.4%	55.9%
	Pro	18.3%	98.3%
Intent Recognition	Nano-2	92.7%	94.6%
Subj	Nano-2	74.3%	79.7%

Table 9: The student learning in-context always outperforms it simply copying the most common label from the teacher. Both methods use 8-shot.

Dataset & Method	Original Example	Student Query
GSM8k Method 1: Description	Two thirds of Jana’s puppies are Pomeranians. One third of the Pomeranians are girls. If there are 6 Pomeranian girls, how many puppies does Jana have?	Division and ratios problem involving percentages.
GSM8k Method 2: New Problem	Two thirds of Jana’s puppies are Pomeranians. One third of the Pomeranians are girls. If there are 6 Pomeranian girls, how many puppies does Jana have	If $\frac{3}{4}$ of a bag of marbles are red and the rest are blue, and there are 21 red marbles, how many marbles are blue?
GSM8k Method 3: Masking	Raul had \$87 to spare so he decided to go to the bookshop. Raul bought 8 comics, each of which cost \$4. How much money does Raul have left?	Emily had \$92 to spend at the ice cream shop. She bought 4 ice cream cones, each of which cost \$3. How much money does Emily have left?
GSM8k Method 3: Masking	The Smith family is getting ready for summer and needs to have their swimming pool filled. The pool company instructed them to measure to find the volume of the pool, then to multiply it by 5.9 to calculate how many gallons of water they need. The cost for the pool company to come and fill the pool is \$0.10 per gallon. Mr. Smith measured and found the pool is 14 feet wide, 25 feet long, and 4 feet deep. How much will it cost to fill the pool?	The Johnson family is getting ready for summer and needs to have their hot tub filled. The hot tub company instructed them to measure to find the volume of the hot tub, then to multiply it by 6.4 to calculate how many gallons of water they need. The cost for the hot tub company to come and fill the hot tub is \$0.15 per gallon. Mr. Johnson measured and found the hot tub is 8 feet wide, 12 feet long, and 3 feet deep. How much will it cost to fill the hot tub
Intent Recognition Method 3: Masking	Show the TV show “The Last Samurai.”	Play the song ‘Blinding Lights.’
Subj Method 3: Masking	’seven years later , alan is given a second chance , a reunion with tommy and the opportunity to set things right .	five years later , jesse is given a second chance , a reunion with thomas and the opportunity to set things right .
Translation Method 2: New Problem	I bought this printer yesterday	I went to the store yesterday to buy some groceries
Translation Method 3: Masking	Tom bought three dozen eggs	Jane bought five dozen apples

Table 10: Examples where the student generates good queries.

Dataset & Method	Original Example	Student Query
GSM8k Method 1: Description	Sandra, the florist around the corner, is very unhappy with Fred’s incomplete order delivery. Sandra had ordered four times as many red roses as white carnations. She also ordered 200 pink calla lilies, which were five times the number of white carnations. Sandra has threatened to switch suppliers if the missing red roses are not delivered by 5 pm. To keep Sandra’s business, how many red roses must Fred deliver by 5 pm	This is a problem of determining the number of red roses Fred must deliver in order to appease Sandra
GSM8k Method 2: New Problem	James has to buy insurance. Since he had an accident it was 60% more than normal. The normal cost is \$120 a month. How much does he pay a year?	John had an accident that was 75% more expensive than usual. If the usual monthly cost was \$150, how much did he pay for insurance for the year?
GSM8k Method 3: Masking	Raul had \$87 to spare so he decided to go to the bookshop. Raul bought 8 comics, each of which cost \$4. How much money does Raul have left?	Olivia had \$222 to spare so she decided to go to the pet shop. Olivia bought 28 dogs, each of which cost \$15. How much money does Olivia have left?
Intent Recognition Method 3: Masking	Rate “Awaiting Strange Gods: Weird and Lovecraftian Fictions” a 1.	Rate "Awaiting Strange Gods: Weird and Lovecraftian Fictions" a 10. It was so bad I wanted to cry and rip my eyes out. Even the sound of the ending credits was painful to endure.
Subj Method 2: New Problem	’adrien , a former high-flying bank robber , recently released from prison and suffering from an incurable form of cancer . . . ’	adrien , a former high-flying bank robber , was recently released from prison after serving a long sentence for his crimes . . .
Translation Method 1: Description	I went aboard	Translate the English phrase "I went aboard" into Greek.

Table 11: Examples in which the student leaks information or does not generate a useful query.