

FINLoRA: BENCHMARKING LoRA METHODS FOR FINE-TUNING LLMs ON FINANCIAL DATASETS

Anonymous authors

Paper under double-blind review

ABSTRACT

Low-rank adaptation (LoRA) methods show great potential for scaling pre-trained general-purpose Large Language Models (LLMs) to hundreds or thousands of use scenarios. However, their efficacy in high-stakes domains like finance is rarely explored, e.g., passing CFA exams and analyzing SEC filings. In this paper, we present the open-source FinLoRA project that benchmarks LoRA methods on both general and highly professional financial tasks. First, we curated 19 datasets covering diverse financial applications; in particular, we created four novel XBRL analysis datasets based on 150 SEC filings. Second, we evaluated five LoRA methods and five base LLMs. Finally, we provide extensive experimental results in terms of accuracy, F1, and BERTScore and report computational costs in terms of time and GPU memory during fine-tuning and inference stages. We find that LoRA achieved substantial performance gains of 40.1 points on average over base models. Our FinLoRA project provides an affordable and scalable approach to democratize financial intelligence to the general public. Datasets, LoRA adapters, code, and documentation are available at <https://anonymous.4open.science/r/FinLoRA-ICLR-115E>.

1 INTRODUCTION

Large language models (LLMs) have demonstrated impressive general capabilities in various vertical domains, such as finance (Wu et al., 2023; Liu et al., 2023a; Lee et al., 2025), education (Liu et al., 2024b), and scientific discovery (Lu et al., 2022; Chen et al., 2021). In the finance sector, LLMs have been applied to tasks such as sentiment analysis (Zhang et al., 2023), question-response, and stock market prediction (Lee et al., 2025).

Cost-effective adaptation is critical for applying LLMs to vertical domains like finance, since general-purpose LLMs lack the specialized knowledge to excel in professional-level tasks. Full fine-tuning can close such performance gaps but is prohibitive for most organizations due to its computationally demanding nature. As such, parameter-efficient fine-tuning (PEFT), particularly Low-Rank Adaptation (LoRA) (Hu et al., 2022) and its variants (Dettmers et al., 2023; Mao et al., 2024; Yang et al., 2024; Huan & Shun, 2025; Liu et al., 2022; Chen et al., 2022; Meng et al., 2024), has emerged as an affordable and scalable solution. LoRA methods can enhance pre-trained general-purpose LLMs with domain-specific knowledge and improve performance on downstream tasks (Mao et al., 2024).

Recent research like FinGPT (Liu et al., 2023a; 2024e) has applied a quantized LoRA method (Dettmers et al., 2023) to general financial tasks; however, the comparative performance of various LoRA variants in complex, professional-level financial tasks remains rarely explored. Previous research shows that LLMs are struggling with professional-level financial tasks, such as analyzing SEC filings (Islam et al., 2023) and passing financial certificate exams (Callanan et al., 2024). A critical area within professional finance involves eXtensible Business Reporting Language (XBRL) data (Saeedi et al., 2007), the de facto global standard for business reporting. Despite XBRL’s importance, dedicated datasets for related analytical tasks are scarce. This deficiency, coupled with the need to evaluate different LoRA methods on highly specialized financial tasks, motivates our introduction of FinLoRA: a comprehensive benchmark designed to assess LoRA variants across diverse financial scenarios, with an emphasis on professional XBRL applications.

This paper demonstrates that fine-tuning state-of-the-art LLMs can significantly improve performance across a range of financial tasks, including specialized XBRL analysis, locally and cost-effectively

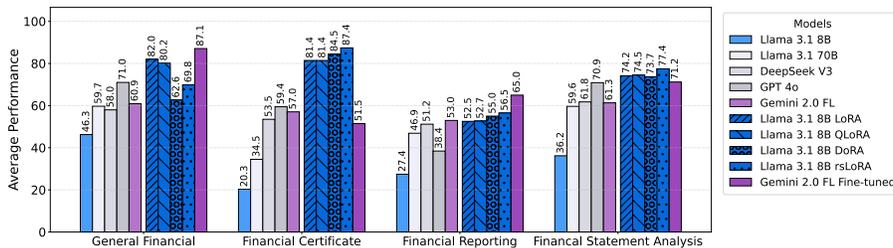


Figure 1: Average performance of base models and LoRA models.

using widely accessible GPUs. As illustrated in Fig. 1, our LoRA-adapted models achieve notable performance improvements over baseline models across four categories of financial tasks. Our main contributions are summarized as follows:

- We curated 19 financial datasets, including general financial tasks, financial analysis, and professional-level XBRL tasks. In particular, we created four novel XBRL analysis datasets. This enables future research to perform rigorous evaluation of LoRA methods in financial tasks.
- We implemented and fairly compared five LoRA methods—including LoRA (Hu et al., 2022), QLoRA (Detmeters et al., 2023), DoRA (Liu et al., 2024c), rsLoRA (Kalajdzievski, 2023), and Federated LoRA—by fine-tuning models on financial datasets. LoRA achieved an average increase of 40.1 point increase in accuracy over baseline models, which validates the effectiveness of low-rank adaptation and quantization for fine-tuning LLMs.
- We conducted an extensive analysis with 59 rounds of fine-tuning and ~ 300 rounds of evaluations for LoRA methods from four angles: (i) a comprehensive comparison across different base models and datasets, (ii) performance on various types of financial tasks, (iii) resource requirements for fine-tuning and inference, and (iv) practical considerations for LoRA deployment in finance.

2 IS FINE-TUNING OF LLMs NEEDED ON FINANCIAL TASKS?

While general-purpose LLMs demonstrate broad NLP competence, their performance often falls short on nuanced financial tasks (Islam et al., 2023; Callanan et al., 2024). This section discusses three key reasons that underscore the necessity of fine-tuning, particularly with methods like LoRA, for developing effective financial LLMs: *(i) Lack of High-Quality Financial Data in Pre-Training Datasets.* Many pre-training datasets, such as The Pile (Gao et al., 2020), primarily draw from general web crawls (e.g., GitHub, arXiv). These sources often under-represent high-quality, specialized financial data, which may be private and exist in complex formats (like XBRL). Consequently, to equip LLMs with the understanding required for complex financial analysis, targeted fine-tuning on curated, domain-specific datasets becomes essential. *(ii) General LLMs’ Failure in Specialized Financial Tasks.* General LLMs often struggle with specialized tasks that demand deep domain-specific knowledge. XBRL analysis provides a clear illustration of these difficulties: Table 1 details the typical Llama base model’s errors on XBRL questions and its improved outputs after LoRA fine-tuning. In Question 1, the base model hallucinated and created a non-existent tag that is not part of the US GAAP taxonomy. In Question 2, the base model did not know the formula of the equity multiplier, resulting not select the correct assets value. *(iii) Cost and Time.* As shown in Table 5, the training-from-scratch approach of BloombergGPT (Wu et al., 2023), which reportedly cost \$2.7 million and took 53 days to train (Table 5), is economically nonviable for most organizations. In contrast, fine-tuning existing foundational models using LoRA methods is significantly more accessible and time-efficient.

3 FINLORA BENCHMARK

3.1 BENCHMARK TASKS, DATASETS, AND METRICS

As displayed in Table 2, we consider four types of tasks: general financial tasks, financial certificate tasks, financial reporting tasks, and financial statement analysis tasks.

Table 1: Case study—XBRL tagging (Google 10-Q 2025-Q1) and XBRL formula calculation (Travelers 10-K FY-2023)—from base Llama 3.1 8B Instruct and our LoRA fine-tuned version.

| | |
|-----------------------------|--|
| Question 1 | What is the appropriate XBRL US-GAAP tag for “2.0” in 3 “...equity securities accounted for under the equity method had a carrying value of approximately \$2.0 billion” ? |
| Llama 3.1 8B | us-gaap:MajorityEquityInterest |
| Llama 3.1 8B LoRA (8bit r8) | us-gaap:EquityMethodInvestments |
| Ground truth | us-gaap:EquityMethodInvestments |
| Question 2 | What is Travelers Companies Inc’s Equity Multiplier for FY 2023? (Answer with a formula substituted with values.) {XBRL Context} |
| Llama 3.1 8B | (1,209,000,000 / 249,210,000,000) |
| Llama 3.1 8B LoRA (8bit r8) | 125,978,000,000 / 249,210,000,000 |
| Ground truth | 125,978,000,000 / 249,210,000,000 |

Public Financial Datasets FinLoRA includes 15 public financial datasets. (i) Sentiment analysis (SA): Financial Phrase Bank (FPB) (Malo et al., 2013), Financial QA Sentiment Analysis (FiQA SA) (Maia et al., 2018), Twitter Financial News Sentiment (TFNS) (Rahman, 2022), and News with GPT Instruction (NWGI) (Liu et al., 2024e), each with financial text from news or tweets and sentiment labels. (ii) Headline analysis: The Headline dataset (Sinha & Khandait, 2020) classifies financial headlines based on various questions into two classes: "yes" and "no". (iii) Named-entity recognition (NER): NER dataset (Salinas Alvarado et al., 2015) annotates one entity per sentence, categorized into one of three classes: "location", "person", and "organization". (iv) Financial certification tasks with a focus on regulations/ethics: CFA Level I, II, and III, and CPA Regulation. (v) Financial reporting: XBRL Terminology (Han et al., 2024), Financial Numeric Entity Recognition (FiNER) (Loukas et al., 2022), and Financial Numeric Extreme Labeling (FNXL) (Sharma et al., 2023). (vi): Financial statement analysis: Financial Math (Han et al., 2024) and FinanceBench (Islam et al., 2023; Han et al., 2024).

Newly-added XBRL Analysis Datasets We introduce 4 novel XBRL analysis datasets, i.e., extracting and analyzing SEC financial reports in XBRL format. These question-answering datasets, derived from the 2019-2023 annual reports of Dow Jones 30 companies, provide each example with a question, a relevant filtered XBRL text segment as source material, and a ground truth answer. The datasets cover four distinct task types: (i) **XBRL tag extraction** involves extracting a specific XBRL tag from a raw XBRL text segment given a natural language description of the tag. (ii) **XBRL value extraction** focuses on extracting a numeric value from the raw XBRL text segment given a natural language description of the value. (iii) **XBRL formula construction** tasks the LLM to first identify and select multiple relevant facts (and their corresponding XBRL tags) from the XBRL data, and then construct a standard financial formula (e.g., Net Profit Margin, Quick Ratio) using these selected tags as components. (iv) **XBRL formula calculation** builds on the previous task and requires the LLM to substitute the actual numeric values into the formula and compute the final result.

Dataset Construction Pipeline Initially, we classified financial tasks into nine categories, creating a training set for each to develop category-specific LoRA adapters per configuration. The four novel XBRL analysis datasets were constructed using XBRL-formatted 10-K annual reports from Dow Jones 30 companies (2019-2023). For these, we generated the four aforementioned types of questions by applying five distinct templates to consolidated, company-specific facts. To ensure contextual relevance, XBRL file segments were automatically filtered based on pertinent factors like year and reporting axes.

Metrics For all general financial tasks, financial analysis, XBRL tagging, financial math, and XBRL analysis tasks, we use Exact Match (EM) to evaluate the LLMs’ output and report both the accuracy and weighted F1 score (in the supplementary materials). For XBRL Term and FinanceBench, we report BERTScore F1 (Zhang et al., 2020) instead. We also report the average of scores across the tasks with BERTScore F1 multiplied by 100.

Table 2: Benchmark tasks and datasets.

| Datasets | Types | #Train/#Test | Average Prompt Length | Metrics | Sources & License |
|---|--------------------|--------------|-----------------------|--------------|---|
| General Financial Tasks (Total: 122.9k/31.7k) | | | | | |
| FPB (Malo et al., 2013) | Sentiment Analysis | 3.1k/970 | 56 | Accuracy, F1 | HF, CC BY-SA 3.0 |
| FiQA SA (Maia et al., 2018) | Sentiment Analysis | 822/234 | 48 | Accuracy, F1 | HF MIT |
| TFNS (Rahman, 2022) | Sentiment Analysis | 9.5k/2.4k | 52 | Accuracy, F1 | HF MIT |
| NWGI (Liu et al., 2023a) | Sentiment Analysis | 12.9k/4.1k | 81 | Accuracy, F1 | HF MIT |
| Headline (Sinha & Khandait, 2020) | Headline Analysis | 82.2k/20.5k | 43 | Accuracy, F1 | HF CC BY-SA 3.0 |
| NER (Salinas Alvarado et al., 2015) | NER | 13.5k/3.5k | 138 | Accuracy, F1 | HF CC BY-SA 3.0 |
| Financial Certificate Tasks (Total: 472/346) | | | | | |
| CFA Level I | Analyst Exam | 180/90 | 181 | Accuracy, F1 | Internet (Public; Not Released Due to Copyright) |
| CFA Level II | Analyst Exam | 88/77 | 1.0k | Accuracy, F1 | |
| CFA Level III | Analyst Exam | 80/78 | 961 | Accuracy, F1 | |
| CPA REG | Accountant Exam | 124/101 | 147 | Accuracy, F1 | |
| Financial Reporting Tasks (Total: 15.9k/8.3k) | | | | | |
| FiNER-139 (Loukas et al., 2022) | XBRL Tagging | 10.0k/7.4k | 1.8k | Accuracy, F1 | HF CC BY-SA 4.0 |
| FNXL (Sharma et al., 2023) | XBRL Tagging | -/247 | 7.1k | Accuracy, F1 | GitHub Public |
| XBRL Term (Han et al., 2024) | Terminology | 5.9k/651 | 25 | BERTScore | GitHub MIT |
| Financial Statement Analysis Tasks (Total: 27.9k/7.3k) | | | | | |
| Financial Math (Han et al., 2024) | Math | 800/200 | 116 | Accuracy | GitHub MIT |
| FinanceBench (Islam et al., 2023; Han et al., 2024) | Math | 86/43 | 983 | BERTScore | GitHub CC BY-NC 4.0 |
| Tags Extraction | XBRL Analysis | 10.1K/2.9k | 3.8k | Accuracy, F1 | Supplementary Materials MIT |
| Values Extraction | XBRL Analysis | 10.1k/2.5k | 3.8k | Accuracy, F1 | Supplementary Materials MIT |
| Formula Construction | XBRL Analysis | 3.4K/835 | 3.8k | Accuracy, F1 | Supplementary Materials MIT |
| Formula Calculation | XBRL Analysis | 3.4K/835 | 3.8k | Accuracy, F1 | Supplementary Materials MIT |

3.2 BASE MODELS AND LORA METHODS

Base Models We benchmark two models for both base model and LoRA fine-tuning performance—Llama 3.1 8B Instruct (Dubey et al., 2024) and Gemini 2.0 Flash Lite (Team et al., 2024). We also evaluated three additional models—Llama 3.1 70B Instruct (Dubey et al., 2024), DeepSeek V3 (Liu et al., 2024a), and GPT-4o (Hurst et al., 2024)—as base models only.

LoRA Methods We considered the following five popular LoRA methods.

- **(Vanilla) LoRA:** Low-rank adaptation (LoRA) (Hu et al., 2022) is a parameter-efficient fine-tuning method that preserves the weights of the pre-trained model and introduces a smaller set of trainable weights. The updated weights follow the low-rank decompositions $\Delta \mathbf{W} = \gamma_r \mathbf{B} \mathbf{A}$, where γ_r is a scaling factor ($\gamma_r = \frac{\alpha}{r}$ with $\alpha > 0$ and $\text{rank } r > 0$), $\mathbf{A} \in \mathbb{R}^{r \times k}$ and $\mathbf{B} \in \mathbb{R}^{d \times r}$ are trainable parameters, and $\mathbf{W}_0 \in \mathbb{R}^{d \times k}$ denote the pre-trained weights. During the fine-tuning stage, the forward pass is $\mathbf{y} = (\mathbf{W}_0 + \gamma_r \mathbf{B} \mathbf{A}) \mathbf{x} = \mathbf{W}_0 \mathbf{x} + \gamma_r \mathbf{B} \mathbf{A} \mathbf{x}$.
- **QLoRA.** Quantized LoRA (QLoRA) (Dettmers et al., 2023) further reduces memory usage by using 4-bit quantization. During fine-tuning, all weights of the pre-trained model are quantized to 4 bits. Weights will be dynamically dequantized back to 16 bits when performing computation with the input sequence \mathbf{x} and the adapter matrix \mathbf{A} and \mathbf{B} , which remain in 16-bit precision throughout the process, where $\mathbf{y} = p_{16}(\mathbf{W}_0^{\text{NF4}}) \mathbf{x} + \gamma_r \mathbf{B} \mathbf{A} \mathbf{x}$. The process is similar in the inference stage, where the merged weights \mathbf{W} are loaded in 4-bit precision.
- **DoRA.** Weight-Decomposed Low-Rank Adaptation (DoRA) (Liu et al., 2024c) decomposes $\mathbf{W}_0 \in \mathbb{R}^{d \times k}$ into a column-wise magnitude vector $\mathbf{m} \in \mathbb{R}^{1 \times k}$ and a direction matrix $\mathbf{V} \in \mathbb{R}^{d \times k}$, where $\mathbf{m} = \|\mathbf{W}_0\|_c$ (with $\|\cdot\|$ being column-wise norm) and $\mathbf{V} = \mathbf{W}_0$. Only the direction matrix receives updates through LoRA. The magnitude vector is updated separately. DoRA can achieve accuracy close to that from full fine-tuning while keeping the same parameter count as LoRA.
- **rsLoRA.** Vanilla LoRA uses a scaling factor α/r , which may cause gradients to explode or diminish as the rank r increases. Rank-Stabilized LoRA (rsLoRA) (Kalajdziewski, 2023) uses a scaling factor α/\sqrt{r} : $\mathbf{W}' = \mathbf{W}_0 + \frac{\alpha}{\sqrt{r}} \mathbf{B} \mathbf{A}$. This scaling results in gradient-scale stability at higher ranks, enabling the rank to be higher for long-context tasks like XBRL analysis.

- **LoRA with Federated Learning.** In the finance sector, multiple institutions may want to collaborate using their own proprietary datasets, but they cannot share their data due to compliance reasons and privacy concerns. Federated learning solves this issue by fine-tuning a model on local data and aggregating LoRA updates to a central node.

3.3 BENCHMARK ANGLES

Angle I: LoRA Methods’ Performance on Financial Datasets We seek to learn which LoRA method is most effective in financial tasks, in terms of both category-specific and overall performance, and how these LoRA fine-tuned models perform compared to existing state-of-the-art (SOTA) models. We fine-tuned Llama 3.1 8B Instruct using LoRA, QLoRA, rsLoRA, and DoRA, representing open-source models and fine-tuning approaches, and fine-tuned Gemini 2.0 Flash Lite using Google’s proprietary fine-tuning methods as a baseline representing closed-source counterparts.

Angle II: LoRA Suitability for Financial Tasks We wish to investigate how the benefits of LoRA fine-tuning vary across different financial tasks. This angle is motivated by the need to identify which specific applications (e.g., sentiment analysis, XBRL tagging, XBRL analysis) are most responsive to fine-tuning, and what properties of the datasets cause this.

Angle III: Resources of LoRA Fine-tuning and Inference We aim to compare which LoRA methods, out of the tested methods, are the most cost-effective in fine-tuning and compare the fine-tuning cost to closed-source fine-tuning services. We are also motivated to measure and compare the inference speeds of LoRA-fine-tuned models against their larger base model counterparts. The goal is to quantify the potential for reduced latency and increased throughput, which are critical for real-time financial applications and operational efficiency.

Angle IV: Practical Considerations for LoRA Deployment in Finance To assess the viability of deploying LoRA-fine-tuned models in real-world financial scenarios, we investigate two key concerns: (i) **Data Privacy in Collaborative Training:** While local LoRA fine-tuning enhances data protection, collaborative model training across multiple institutions often requires approaches like Federated Learning to preserve the privacy of proprietary training data. We investigate this by simulating data distribution across several nodes and evaluating LoRA fine-tuning performance against centralized training. (ii) **Catastrophic Forgetting:** Fine-tuning can risk degrading a model’s pre-existing general knowledge and capabilities. To quantify this, we evaluate our LoRA-fine-tuned models on established general-domain benchmarks, such as MMLU (Hendrycks et al., 2020), measuring any performance changes on tasks outside their financial fine-tuning scope.

4 BENCHMARK RESULTS

Setup Our experiments were conducted on four NVIDIA A5000 GPUs. For closed-source models, we employed various inference and fine-tuning APIs. (i) **Single-task fine-tuning:** we separate our datasets into 9 tasks, for each tasks we fine-tune using vanilla LoRA, QLoRA, DoRA, rsLoRA on Llama 3.1 8B Instruct, and only vanilla LoRA on Ministral 8B Instruct, and Google proprietary tuning on Gemini 2.0 Flash Lite. (ii) **Multi-task fine-tuning:** We also conducted multi-task fine-tuning that only outputs 4 LoRA adapters based on the major categories. We used a learning rate of $1e-4$ and a batch size of 2–8 based on prompt length (Refer to appendix section C for details).

4.1 ANGLE I: LoRA METHODS PERFORMANCE ON FINANCIAL DATASETS

Comparative Performance of LoRA Variants Table 4 shows the performance of base models and different LoRA fine-tuned models. Vanilla LoRA (8-bit, rank 8) achieves the highest overall average score (74.74), a 40.1 point increase over the Llama 3.1 8B base model. Fig. 1 shows the performance by category. Vanilla LoRA outperforms other LoRA variants in general financial tasks, while rsLoRA leads in financial analysis, financial reporting, and financial statement analysis.

Single-task vs. Multi-task Fine-tuning Multi-task fine-tuning produces clear gains in Financial Statement Analysis. Formula construction, formula calculation, Finance Bench, and financial

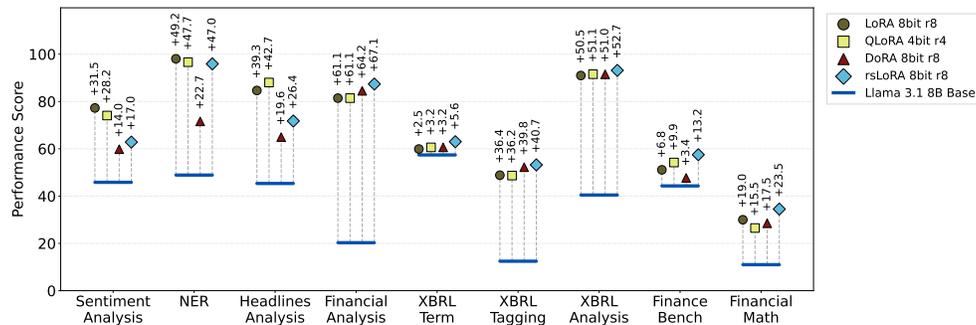


Figure 2: Task suitability.

math. These tasks share similar underlying knowledge like the structure of financial statements and numerical reasoning. Learning them together helps the model to achieve positive knowledge transfer. In contrast, we see negative transfer in General Financial and Financial Reporting tasks, where the multi-task model performs worse on TFNS, Headline, FiNER, and FNXL. The problem comes from differences in task format and objective. Even with balanced sampling, the model struggles to optimize for very different objectives at the same time. Overall, performance drops by 6 points when tasks are merged. Closely related tasks can benefit from joint training, while divergent tasks often harm each other. Refer to appendix section A.3 for more details.

rsLoRA Performs Better at High Ranks rsLoRA scales with α/\sqrt{r} instead of α/r to prevent gradient exploding or vanishing at large ranks. We set $r = 8$ for memory efficiency. rsLoRA just slightly underperforms against LoRA and QLoRA. The rsLoRA paper’s experiments (Kalajdziewski, 2023) led to lower perplexity at higher ranks (e.g., $r = 64$). This lower perplexity and the fact that higher rank LoRA captures more details suggest rsLoRA’s benefits are primarily exploited at high ranks.

DoRA Benefits from Two Learning Rates DoRA performed worse than the other three LoRA methods. We used the same learning rate for updating the magnitude vector and direction matrix. However, as shown in Table 4, this can lead to sub-optimal performance in some cases due to the gradient scales being different between the two types of updates in DoRA. This leads to DoRA sometimes under-training the magnitude vector in our experiments, which uses the same low learning rate. Thus, DoRA may achieve higher performance if the magnitude vector has its own learning rate that is higher than the low-rank update’s learning rate.

LoRA-Tuned Llama 3.1 8B vs. Baseline Models and Gemini Fine-Tuned Compared to SOTA base LLMs, the LoRA-tuned Llama 3.1 8B Instruct models generally show superior performance across most datasets, with NWGI and FNXL being the exceptions. Against another fine-tuned baseline, the Gemini 2.0 FL fine-tuned model, this Gemini model excels in general financial tasks and XBRL data reporting. However, our Llama 3.1 8B Instruct LoRA variants demonstrate stronger average performance in financial analysis and XBRL data analysis tasks.

4.2 ANGLE II: FINANCIAL TASK LORA SUITABILITY

Fig. 2 highlights LoRA’s varying effectiveness across different financial tasks. A key observation is the contrast in LoRA method improvements between XBRL Analysis tasks and FinanceBench. Although both aim to analyze financial statements, tasks based on XBRL data demonstrate substantial LoRA-induced performance improvements, whereas FinanceBench exhibits minimal gains. This disparity underscores XBRL’s superior suitability for financial statement analysis. The standardized semantics and taxonomy inherent in XBRL likely provide a more structured and consistent learning environment for LLMs, facilitating more effective adaptation compared to FinanceBench, which relies on OCR-processed PDF data lacking such rich, standardized metadata. These findings emphasize the crucial role of XBRL in enabling effective LLM integration for financial report analysis.

Table 4: Performance on financial tasks: accuracy in blue, F1 in gray, and BERTScore F1 in green.

| Datasets | Base Models | | | | | | | Fine-tuned Models | | | | | | | | | | | | | |
|--|--------------|---------------|---------------|---------------|--------------|--------------|--------------|-------------------|--------------|--------------|---------------|---------------|--------------|--------------|--------------|--------------|--------------|-------|--------------|-------|--------------|
| | 0 Shot | | | | 3 Shot | | | Single-task | | | | Multi-task | | | | | | | | | |
| | Llama 3.1 8B | Mini-stral 8B | Gemini 2.0 FL | Llama 3.1 70B | Deep-Seek V3 | GPT-4o | Llama 3.1 8B | Llama 3.1 8B | Llama 3.1 8B | Llama 3.1 8B | Mini-stral 8B | Gemini 2.0 FL | Llama 3.1 8B | LoRA | QLoRA | DoRA | rsLoRA | LoRA | N/A | LoRA | |
| General Financial Tasks | | | | | | | | | | | | | | | | | | | | | |
| FPB | 68.73 | 73.08 | 81.02 | 74.50 | 78.76 | 81.13 | 76.40 | 85.64 | 84.16 | 81.93 | 82.84 | 86.71 | 87.62 | 85.31 | 0.677 | 0.695 | 0.894 | 0.736 | 0.764 | 0.818 | 0.865 |
| FiQA | 46.55 | 52.86 | 68.09 | 47.27 | 60.43 | 72.34 | 64.68 | 81.28 | 78.30 | 78.72 | 73.19 | 80.00 | 88.09 | 82.20 | 0.922 | 0.909 | 0.901 | 0.853 | 0.868 | 0.878 | 0.873 |
| SA | 0.557 | 0.591 | 0.810 | 0.565 | 0.686 | 0.773 | 0.782 | 0.884 | 0.874 | 0.874 | 0.806 | 0.807 | 0.879 | 0.853 | 0.884 | 0.874 | 0.874 | 0.806 | 0.807 | 0.879 | 0.853 |
| TFNS | 24.16 | 22.07 | 26.38 | 68.42 | 84.38 | 73.32 | 28.81 | 88.02 | 83.84 | 59.09 | 59.51 | 45.85 | 89.49 | 34.51 | 0.339 | 0.189 | 0.385 | 0.686 | 0.846 | 0.740 | 0.415 |
| NWGI | 43.86 | 21.25 | 48.16 | 50.14 | 7.44 | 66.61 | 32.20 | 54.16 | 49.96 | 19.57 | 35.80 | 56.90 | 62.59 | 36.51 | 0.932 | 0.910 | 0.702 | 0.655 | 0.460 | 0.896 | 0.399 |
| NER | 48.89 | 58.61 | 65.13 | 46.28 | 40.82 | 52.11 | 58.85 | 98.05 | 96.63 | 71.59 | 95.92 | 98.05 | 97.29 | 76.07 | 0.690 | 0.645 | 0.281 | 0.464 | 0.538 | 0.581 | 0.475 |
| Headline | 0.569 | 0.647 | 0.769 | 0.454 | 0.360 | 0.523 | 0.707 | 0.981 | 0.966 | 0.834 | 0.963 | 0.981 | 0.973 | 0.856 | 0.583 | 0.626 | 0.614 | 0.596 | 0.097 | 0.656 | 0.399 |
| Category Avg. | 45.34 | 62.64 | 76.60 | 71.68 | 76.06 | 80.53 | 64.28 | 84.66 | 88.03 | 64.93 | 71.75 | 97.51 | 97.32 | 13.90 | 0.852 | 0.886 | 0.781 | 0.828 | 0.976 | 0.973 | 0.197 |
| | 46.25 | 48.42 | 60.90 | 59.72 | 57.98 | 71.01 | 54.20 | 81.97 | 80.15 | 62.64 | 69.83 | 77.50 | 87.07 | 54.75 | | | | | | | |
| Financial Certificate Tasks | | | | | | | | | | | | | | | | | | | | | |
| CFA Level 1 | 13.33 | 88.89 | 55.56 | 42.22 | 54.44 | 63.33 | 51.11 | 86.67 | 87.78 | 87.78 | 87.78 | 87.77 | 52.22 | 86.67 | 0.133 | 0.889 | 0.556 | 0.418 | 0.556 | 0.631 | 0.508 |
| CFA Level 2 | 19.48 | 94.80 | 56.67 | 29.87 | 46.75 | 55.84 | 37.66 | 88.31 | 83.12 | 90.91 | 92.21 | 94.80 | 51.11 | 88.31 | 0.199 | 0.948 | 0.567 | 0.303 | 0.485 | 0.563 | 0.383 |
| CFA Level 3 | 16.67 | 78.21 | 52.56 | 24.36 | 47.44 | 51.28 | 51.28 | 70.51 | 66.67 | 69.23 | 79.49 | 78.20 | 51.28 | 70.51 | 0.179 | 0.782 | 0.538 | 0.271 | 0.496 | 0.517 | 0.526 |
| CPA REG | 31.68 | 87.13 | 63.37 | 41.58 | 65.35 | 67.33 | 45.54 | 80.20 | 88.12 | 90.10 | 90.10 | 91.08 | 51.28 | 80.20 | 0.317 | 0.871 | 0.638 | 0.426 | 0.654 | 0.667 | 0.459 |
| Category Avg. | 20.29 | 87.26 | 57.04 | 34.51 | 53.49 | 59.44 | 46.40 | 81.42 | 81.42 | 84.50 | 87.40 | 87.96 | 51.47 | 81.42 | | | | | | | |
| Financial Reporting Tasks | | | | | | | | | | | | | | | | | | | | | |
| FiNER | 21.28 | 22.62 | 63.91 | 61.82 | 68.92 | 72.29 | 30.76 | 74.10 | 74.32 | 70.92 | 70.72 | 69.56 | 80.32 | 67.97 | 0.232 | 0.266 | 0.638 | 0.606 | 0.699 | 0.725 | 0.369 |
| FNXL | 3.64 | 2.55 | 37.75 | 20.14 | 27.33 | 42.41 | 13.16 | 23.57 | 23.05 | 33.50 | 35.68 | 33.67 | 47.98 | 28.23 | 0.045 | 0.026 | 0.356 | 0.210 | 0.288 | 0.398 | 0.116 |
| XBRL Term | 0.574 | 0.563 | 0.572 | 0.587 | 0.573 | 0.584 | 0.595 | 0.599 | 0.606 | 0.606 | 0.630 | 0.672 | 0.666 | 0.676 | 0.250 | 0.253 | 0.311 | 0.348 | 0.314 | 0.438 | 0.282 |
| Category Avg. | 27.44 | 27.16 | 52.95 | 46.89 | 51.18 | 57.70 | 34.47 | 52.52 | 52.66 | 55.01 | 56.47 | 56.81 | 64.97 | 54.60 | | | | | | | |
| Financial Statement Analysis Tasks | | | | | | | | | | | | | | | | | | | | | |
| Tag Extr. | 69.16 | 74.15 | 80.27 | 69.64 | 85.03 | 81.60 | 70.22 | 89.13 | 86.89 | 80.44 | 85.26 | 84.51 | 85.03 | 88.78 | 0.739 | 0.755 | 0.811 | 0.782 | 0.849 | 0.864 | 0.739 |
| Value Extr. | 52.46 | 74.21 | 98.02 | 88.19 | 98.01 | 97.01 | 72.29 | 98.49 | 97.14 | 98.57 | 99.13 | 98.80 | 99.20 | 97.62 | 0.565 | 0.743 | 0.980 | 0.904 | 0.982 | 0.974 | 0.786 |
| Formula Constr. | 12.92 | 11.91 | 61.90 | 59.28 | 22.75 | 79.76 | 17.73 | 77.61 | 89.34 | 88.02 | 89.46 | 62.39 | 67.85 | 83.33 | 0.201 | 0.187 | 0.644 | 0.665 | 0.315 | 0.820 | 0.233 |
| Formula Calc. | 27.27 | 47.62 | 53.57 | 77.49 | 85.99 | 83.59 | 33.65 | 98.68 | 92.81 | 98.92 | 98.80 | 48.50 | 54.76 | 99.04 | 0.876 | 0.898 | 0.882 | 0.893 | 0.627 | 0.786 | 0.840 |
| Finance Bench | 0.317 | 0.488 | 0.536 | 0.783 | 0.868 | 0.857 | 0.369 | 0.990 | 0.947 | 0.993 | 0.988 | 0.483 | 0.548 | 0.993 | 0.990 | 0.947 | 0.993 | 0.988 | 0.483 | 0.548 | 0.993 |
| Financial Math | 0.443 | 0.584 | 0.552 | 0.528 | 0.573 | 0.564 | 0.580 | 0.511 | 0.542 | 0.477 | 0.575 | 0.617 | 0.544 | 0.621 | 0.443 | 0.584 | 0.552 | 0.528 | 0.573 | 0.564 | 0.580 |
| Category Avg. | 11.00 | 36.00 | 19.00 | 10.50 | 21.50 | 27.00 | 32.00 | 30.00 | 26.50 | 28.50 | 34.50 | 46.00 | 66.00 | 58.00 | 0.136 | 0.378 | 0.204 | 0.134 | 0.255 | 0.296 | 0.367 |
| | 36.19 | 50.38 | 61.33 | 59.65 | 61.76 | 70.89 | 47.31 | 74.17 | 74.48 | 73.69 | 77.44 | 66.98 | 71.21 | 81.48 | | | | | | | |
| Overall Average (Using accuracy for classification and BERTScore F1 × 100 for generation) | | | | | | | | | | | | | | | | | | | | | |
| Avg. | 34.64 | 53.86 | 58.97 | 52.36 | 57.16 | 66.44 | 47.27 | 74.74 | 74.29 | 69.53 | 73.82 | 73.12 | 71.08 | 68.78 | | | | | | | |

Table 5: Comparison of fine-tuning cost. GPT-4o cost is estimated based on 4 epochs of fine-tuning at OpenAI fine-tuning pricing (OpenAI, 2025).

| Models | Time | GPUs | Est. Cost (USD) |
|---------------|---------|------------|-----------------|
| BloombergGPT | 53 days | 512 × A100 | \$2.7 M |
| LoRA | 14.9h | 4 × A5000 | \$15.50 |
| QLoRA | 14.1h | 4 × A5000 | \$14.66 |
| DoRA | 15.9h | 4 × A5000 | \$16.54 |
| rsLoRA | 14.5h | 4 × A5000 | \$15.11 |
| Gemini 2.0 FL | 8.8h | - | \$162.02 |
| GPT-4o-mini | - | - | \$312.00 |

Figure 3: Average inference time of LoRA fine-tuned Llama 3.1 8B and LoRA fine-tuned Gemini 2.0 FL across tasks

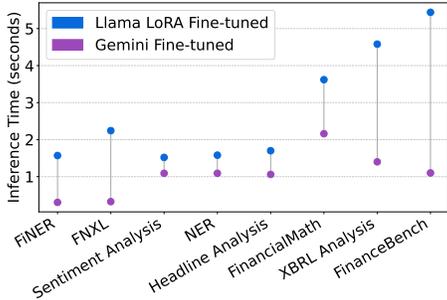


Table 6: Accuracy on MMLU, GSM8K, TriviaQA, and CoQA benchmarks for Llama 3.1 8B base and various LoRA adapters. Scores are colored relative to base: gray (same), green (higher), red (lower)

| Dataset | Llama | Llama 3.1 8B (FiNER) | | | | |
|--------------------|---------------|----------------------|---------------|----------------|--------------|----------------|
| | 3.1 8B (base) | LoRA 8bit-r8 | QLoRA 4bit-r4 | QLoRA 4bit-r32 | DoRA 8bit-r8 | rsLoRA 8bit-r8 |
| MMLU | 0.229 | 0.229 | 0.229 | 0.229 | 0.229 | 0.229 |
| GSM8K | 0.011 | 0.011 | 0.014 | 0.010 | 0.011 | 0.016 |
| TriviaQA-Open (F1) | 0.667 | 0.658 | 0.673 | 0.666 | 0.663 | 0.663 |
| CoQA (F1) | 0.711 | 0.709 | 0.697 | 0.715 | 0.707 | 0.715 |

4.3 ANGLE III: RESOURCE USAGE AND PERFORMANCE TRADE-OFFS OF LORA METHODS

Table 5 details the computational costs of LoRA fine-tuned models. Using four NVIDIA A5000 GPUs, the wall-clock time for fine-tuning ranged from 14.1 hours (QLoRA) to 15.9 hours (DoRA), corresponding to a total of approximately 56.4 to 63.6 GPU hours. At an estimated rate of \$0.26 per GPU hour, this translates to a cost of roughly \$14.66 to \$16.54. This is substantially more cost-effective than fine-tuning services from providers like Google or OpenAI. Figure 3 illustrates the inference time of fine-tuned models on various datasets. Gemini API generally exhibits lower inference latency and is less sensitive to increasing prompt lengths than local Llama 3.1 8B Instruct inference, even when accounting for network overhead for the API. However, the inference speed of locally deployed Llama models can be significantly enhanced through the use of larger batch sizes.

4.4 ANGLE IV: PRACTICABILITY OF APPLYING LORA IN REAL-WORLD FINANCIAL SCENARIOS

Federated LoRA The sensitive nature of financial data necessitates privacy-preserving techniques like Federated Learning for collaborative training. To explore this, we evaluated Federated LoRA (Sun et al., 2024), with results presented in Table 7. Our experimental setup simulated a four-node environment employing the FedAvg algorithm (McMahan et al., 2017), where the sentiment analysis dataset was partitioned across these nodes. The performance of this approach was benchmarked against both the base Llama model and standard centralized LoRA fine-tuning. While Federated LoRA did not match the performance levels of centralized LoRA, the results demonstrate a notable improvement compared to the base Llama model.

Catastrophic Forgetting A major concern with PEFT is that fine-tuning on domain-specific tasks leads to the model forgetting pre-training knowledge. To investigate this, we evaluated eight adapters—covering both sentiment and FiNER tasks and all four LoRA variants—as well as the Llama 3.1 8B Instruct base model on two out-of-domain benchmarks, MMLU (Hendrycks et al., 2020) and GSM8K (Cobbe et al., 2021). We used a zero-shot, no chain-of-thought setting to isolate stored knowledge. Table 6 shows identical MMLU accuracy across all adapters and the base model,

Table 7: Performance comparison of central LoRA and LoRA federated learning (8 bit, rank 8) on Llama 3.1 8B using four nodes: accuracy (blue) and F1 score (gray).

| | Sentiment | | | | XBRL Analysis | | | |
|-------------------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| | FPB | FiQA SA | TFNS | NWGI | Tag Extr. | Value Extr. | Formula Constr. | Formula Calc. |
| Base | 68.73 0.677 | 46.55 0.557 | 24.16 0.339 | 43.86 0.538 | 69.16 0.739 | 52.46 0.565 | 12.92 0.201 | 27.27 0.317 |
| Central (LoRA) | 85.64 0.922 | 81.28 0.884 | 88.02 0.932 | 54.16 0.690 | 89.13 0.886 | 98.49 0.986 | 77.61 0.876 | 98.68 0.990 |
| FedAvg (McMahan et al., 2017) | 82.43 0.902 | 76.17 0.860 | 73.41 0.842 | 56.02 0.698 | 69.05 0.745 | 79.76 0.814 | 13.10 0.170 | 28.57 0.305 |

and equal or higher scores on GSM8K. Hence, at the ranks r we tested (4 and 8) with $\alpha:r$ equal to 8:1 or 4:1, we observe that LoRA does not exhibit catastrophic forgetting. In fact, the slight GSM8K performance improvements hint at cross-domain knowledge transfer—fine-tuning on financial data may improve the model’s numerical reasoning skills.

5 RELATED WORKS

5.1 FINANCIAL LLMs AND BENCHMARKS

BloombergGPT (Wu et al., 2023), a 50B parameter model trained from scratch on mixed financial and general data, was the first specialized financial LLM, showing strong performance on domain-specific tasks. In contrast, FinGPT (Liu et al., 2023a; 2024f;e) applies LoRA fine-tuning to open-source LLMs, an efficient alternative that improves performance over base models, even surpassing BloombergGPT. Financial benchmarks like FinBen (Xie et al., 2024) and PIXIU (Xie et al., 2023) show that while current LLMs handle textual analysis well, they struggle with complex financial reasoning.

5.2 PARAMETER-EFFICIENT FINE-TUNING (PEFT)

To mitigate the high cost of full fine-tuning, Parameter-efficient fine-tuning (PEFT) methods update only a small subset of model parameters (Mao et al., 2024). A prominent PEFT technique is Low-rank Adaptation (LoRA) (Hu et al., 2022), which injects trainable low-rank matrices and achieves performance comparable to full fine-tuning. Quantized LoRA (QLoRA) (Dettmers et al., 2023) further enhances efficiency by applying LoRA to a 4-bit quantized model, drastically reducing memory usage.

5.3 LoRA WITH FEDERATED LEARNING

Federated learning is necessary to fine-tune LLMs on decentralized private data, which is common in finance. Consequently, methods combining LoRA with federated learning have been proposed, such as PrivateLoRA (Wang et al., 2023) and FFA-LoRA (Sun et al., 2024).

6 CONCLUSION AND FUTURE WORK

In this paper, we present FinLoRA, a benchmark that evaluates LoRA methods on both general and highly specialized financial tasks. We curated 19 diverse datasets covering a wide range of financial applications. Our study includes 59 rounds of fine-tuning and ~ 300 rounds of evaluation to thoroughly assess and analyze commonly used LoRA methods. FinLoRA offers insights into overall performance, task-specific results, resource requirements for fine-tuning and inference, and practical considerations for real-world deployment—including data privacy in collaborative training and catastrophic forgetting. Our results demonstrate that fine-tuning can significantly enhance the effectiveness of LLMs on financial tasks. Additionally, FinLoRA provides a comprehensive collection of datasets with baseline results, laying a solid foundation for future research in this field. Moving forward, we plan to expand FinLoRA by incorporating additional LoRA methods into the project.

486 REPRODUCIBILITY STATEMENT
487

488 To ensure the reproducibility of our results, all code, datasets, and pretrained LoRA adapters have been
489 made publicly and anonymously available at <https://anonymous.4open.science/r/FinLoRA-ICLR-115E/>. The repository includes detailed documentation and README files with instructions for
490 setting up the environment and running the experiments. Appendix C includes details on experimental
491 setup, such as hyperparameters. Our XBRL analysis dataset and project code are also available in
492 supplementary materials on Open Review. Furthermore, a comprehensive description of our novel
493 XBRL analysis dataset, including all data collection and processing steps, is provided in Appendix E.
494

495
496 REFERENCES
497

- 498 Ethan Callanan, Amarachi Mbakwe, Antony Papadimitriou, Yulong Pei, Mathieu Sibue, Xiaodan
499 Zhu, Zhiqiang Ma, Xiaomo Liu, and Sameena Shah. Can GPT models be financial analysts? an
500 evaluation of ChatGPT and GPT-4 on mock CFA exams. In *Proceedings of the Eighth Financial
501 Technology and Natural Language Processing and the 1st Agent AI for Scenario Planning*, pp.
502 23–32, Jeju, South Korea, 3 August 2024. -. URL <https://aclanthology.org/2024.finnlp-2.2/>.
503
- 504 Lulu Chen, Yingzhou Lu, Chiung-Ting Wu, Robert Clarke, Guoqiang Yu, Jennifer E Van Eyk,
505 David M Herrington, and Yue Wang. Data-driven detection of subtype-specific differentially
506 expressed genes. *Scientific Reports*, 11(1):332, 2021.
- 507 Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo.
508 Adaptformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural
509 Information Processing Systems*, 35:16664–16678, 2022.
- 510 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
511 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve
512 math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- 513 Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. QLoRA: Efficient finetuning
514 of quantized LLMs. In *Thirty-seventh Conference on Neural Information Processing Systems*,
515 2023. URL <https://openreview.net/forum?id=OUIFPHEgJU>.
- 516 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
517 Letman, Akhil Mathur, and et al. The Llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.
518
- 519 Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang,
520 Horace He, Anish Thite, Noa Nabeshima, et al. The Pile: An 800GB dataset of diverse text for
521 language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- 522 Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach,
523 et al. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021.
- 524 Shijie Han, Haoqiang Kang, Bo Jin, Xiao-Yang Liu, and Steve Y Yang. XBRL Agent: Leveraging
525 large language models for financial report analysis. In *Proceedings of the 5th ACM International
526 Conference on AI in Finance, ICAIF ’24*, pp. 856–864, 2024. ISBN 9798400710810.
- 527 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and
528 Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint
529 arXiv:2009.03300*, 2020.
- 530 Edward J Hu, Yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
531 and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International
532 Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.
533
- 534 Muchen Huan and Jianhong Shun. Fine-tuning transformers efficiently: A survey on LoRA and
535 its impact. *Preprints*, February 2025. doi: 10.20944/preprints202502.1637.v1. URL <https://doi.org/10.20944/preprints202502.1637.v1>.
536
537
538
539

- 540 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-
541 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. GPT-4o system card. *arXiv preprint*
542 *arXiv:2410.21276*, 2024.
- 543
544 Pranab Islam, Anand Kannappan, Douwe Kiela, Rebecca Qian, Nino Scherrer, and Bertie Vidgen.
545 FinanceBench: A new benchmark for financial question answering, 2023. URL [https://](https://arxiv.org/abs/2311.11944)
546 arxiv.org/abs/2311.11944.
- 547 Damjan Kalajdzievski. A Rank Stabilization Scaling Factor for Fine-Tuning with LoRA, 2023. URL
548 <https://arxiv.org/abs/2312.03732>.
- 549
550 Jean Lee, Nicholas Stevens, and Soyeon Caren Han. Large language models in finance (fin-
551 llms). *Neural Computing and Applications*, January 2025. ISSN 1433-3058. doi: 10.1007/
552 s00521-024-10495-6. URL <http://dx.doi.org/10.1007/s00521-024-10495-6>.
- 553 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,
554 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. DeepSeek-V3 technical report. *arXiv preprint*
555 *arXiv:2412.19437*, 2024a.
- 556
557 Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and
558 Colin A Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context
559 learning. *Advances in Neural Information Processing Systems*, 35:1950–1965, 2022.
- 560 Jiayu Liu, Zhenya Huang, Tong Xiao, Jing Sha, Jinze Wu, Qi Liu, Shijin Wang, and Enhong
561 Chen. SocraticLM: Exploring socratic personalized teaching with large language models. In
562 *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024b. URL
563 <https://openreview.net/forum?id=qkoZgJhxA>.
- 564 Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-
565 Ting Cheng, and Min-Hung Chen. DoRA: Weight-decomposed low-rank adaptation. In *Forty-first*
566 *International Conference on Machine Learning*, 2024c.
- 567
568 Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-
569 Ting Cheng, and Min-Hung Chen. DoRA: Weight-decomposed low-rank adaptation, 2024d. URL
570 <https://arxiv.org/abs/2402.09353>.
- 571 Xiao-Yang Liu, Guoxuan Wang, Hongyang Yang, and Daochen Zha. Data-centric FinGPT: Democra-
572 tizing internet-scale data for financial large language models. In *Workshop on Instruction Tuning*
573 *and Instruction Following, NeurIPS*, 2023a.
- 574
575 Xiao-Yang Liu, Guoxuan Wang, Hongyang Yang, and Daochen Zha. FinGPT: Democratizing
576 internet-scale data for financial large language models. *arXiv preprint arXiv:2307.10485*, 2023b.
- 577
578 Xiao-Yang Liu, Jie Zhang, Guoxuan Wang, Weiqin Tong, and Anwar Walid. Efficient Pretraining and
579 Finetuning of Quantized LLMs with Low-Rank Structure . In *IEEE 44th International Conference*
580 *on Distributed Computing Systems (ICDCS)*, pp. 300–311, July 2024e.
- 581
582 Xiao-Yang Liu, R. Zhu, Daochen Zha, J. Gao, S. Zhong, Matt White, and Meikang Qiu. Differentially
583 private low-rank adaptation of large language model using federated learning. *ACM Transactions*
584 *on Management Information Systems*, 2024f.
- 585
586 Xiao-Yang Liu, Rongyi Zhu, Daochen Zha, Jiechao Gao, Shan Zhong, Matt White, and Meikang Qiu.
587 Differentially private low-rank adaptation of large language model using federated learning. *ACM*
588 *Transactions on Management Information Systems*, 16(2):1–24, 2025.
- 589
590 Lefteris Loukas, Manos Fergadiotis, Ilias Chalkidis, Eirini Spyropoulou, Prodromos Malakasiotis,
591 Ion Androutopoulos, and Georgios Paliouras. FiNER: Financial numeric entity recognition for
592 XBRL tagging. In *Proceedings of the 60th Annual Meeting of the Association for Computational*
593 *Linguistics (Volume 1: Long Papers)*, Dublin, Ireland, May 2022.
- 594
595 Yingzhou Lu, Chiung-Ting Wu, Sarah J Parker, Zuolin Cheng, Georgia Saylor, Jennifer E Van Eyk,
596 Guoqiang Yu, Robert Clarke, David M Herrington, and Yue Wang. COT: an efficient and accurate
597 method for detecting marker genes among many subtypes. *Bioinformatics Advances*, 2(1), 2022.

- 594 Macedo Maia, Siegfried Handschuh, Andre Freitas, Brian Davis, Ross McDermott, Manel Zarrouk,
595 and Alexandra Balahur. Wwv'18 open challenge: Financial opinion mining and question answer-
596 ing. pp. 1941–1942, 04 2018.
- 597
- 598 Pekka Malo, Ankur Sinha, Pyry Takala, Pekka Korhonen, and Jyrki Wallenius. Good debt or bad
599 debt: Detecting semantic orientations in economic texts, 2013. URL [https://arxiv.org/
600 abs/1307.5336](https://arxiv.org/abs/1307.5336).
- 601
- 602 Yuren Mao, Yuhang Ge, Yijiang Fan, Wenyi Xu, Yu Mi, Zhonghao Hu, and Yunjun Gao. A survey
603 on lora of large language models. *Frontiers of Computer Science*, 19(7), December 2024. ISSN
604 2095-2236. doi: 10.1007/s11704-024-40663-9. URL [http://dx.doi.org/10.1007/
605 s11704-024-40663-9](http://dx.doi.org/10.1007/s11704-024-40663-9).
- 606
- 607 Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas.
608 Communication-Efficient Learning of Deep Networks from Decentralized Data. In Aarti Singh
609 and Jerry Zhu (eds.), *Proceedings of the 20th International Conference on Artificial Intelligence
610 and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 1273–1282. PMLR,
20–22 Apr 2017. URL <https://proceedings.mlr.press/v54/mcmahan17a.html>.
- 611
- 612 Fanxu Meng, Zhaohui Wang, and Muhan Zhang. Pissa: Principal singular values and singular vectors
613 adaptation of large language models. *Advances in Neural Information Processing Systems*, 37:
121038–121072, 2024.
- 614
- 615 OpenAI. OpenAI API pricing. <https://platform.openai.com/docs/pricing>, May
616 2025. Accessed: 2025-05-14.
- 617
- 618 Md. Abdur Rahman. Twitter financial news sentiment. [http://precog.iiitd.edu.in/
619 people/anupama](http://precog.iiitd.edu.in/people/anupama), 2022.
- 620
- 621 Ali Saeedi, Jim Richards, and Barry Smith. An introduction to XBRL. In *British Accounting
622 Association's Annual Conference*, 2007.
- 623
- 624 Julio Cesar Salinas Alvarado, Karin Verspoor, and Timothy Baldwin. Domain adaption of named
625 entity recognition to support credit risk assessment. In Ben Hachey and Kellie Webster (eds.),
626 *Proceedings of the Australasian Language Technology Association Workshop 2015*, pp. 84–90,
Parramatta, Australia, December 2015. URL <https://aclanthology.org/U15-1010>.
- 627
- 628 Soumya Sharma, Subhendu Khatuya, Manjunath Hegde, Afreen Shaikh, Koustuv Dasgupta, Pawan
629 Goyal, and Niloy Ganguly. Financial numeric extreme labelling: A dataset and benchmarking. In
Findings of the Association for Computational Linguistics: ACL 2023, pp. 3550–3561, July 2023.
- 630
- 631 Ankur Sinha and Tanmay Khandait. Impact of news on the commodity market: Dataset and results,
632 2020. URL <https://arxiv.org/abs/2009.04202>.
- 633
- 634 Youbang Sun, Zitao Li, Yaliang Li, and Bolin Ding. Improving loRA in privacy-preserving federated
635 learning. In *The Twelfth International Conference on Learning Representations*, 2024. URL
<https://openreview.net/forum?id=NLPzL6HWN1>.
- 636
- 637 Gemini Team, Rohan Anil, Sebastian Borgeaud, et al. Gemini: A family of highly capable multimodal
638 models, 2024. URL <https://arxiv.org/abs/2312.11805>.
- 639
- 640 Yiming Wang, Yu Lin, Xiaodong Zeng, and Guannan Zhang. PrivateLoRA for efficient privacy
641 preserving LLM, 2023. URL <https://arxiv.org/abs/2311.14030>.
- 642
- 643 Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhan-
644 jan Kambadur, David Rosenberg, and Gideon Mann. BloombergGPT: A large language model for
finance. *arXiv preprint arXiv:2303.17564*, 2023.
- 645
- 646 Qianqian Xie, Weiguang Han, Xiao Zhang, Yanzhao Lai, Min Peng, Alejandro Lopez-Lira, and Jimin
647 Huang. PIXIU: A comprehensive benchmark, instruction dataset and large language model for
finance. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and
Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=vTrRq6vCQH>.

648 Qianqian Xie, Weiguang Han, Zhengyu Chen, Ruoyu Xiang, Xiao Zhang, Yueru He, Mengxi Xiao,
649 Dong Li, Yongfu Dai, Duanyu Feng, Yijing Xu, Haoqiang Kang, Ziyang Kuang, Chenhan Yuan,
650 Kailai Yang, Zheheng Luo, Tianlin Zhang, Zhiwei Liu, Guojun Xiong, Zhiyang Deng, Yuechen
651 Jiang, Zhiyuan Yao, Haohang Li, Yangyang Yu, Gang Hu, Huang Jiajia, Xiao-Yang Liu, Alejandro
652 Lopez-Lira, Benyou Wang, Yanzhao Lai, Hao Wang, Min Peng, Sophia Ananiadou, and Jimin
653 Huang. FinBen: An holistic financial benchmark for large language models. In *The Thirty-eight
654 Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.
655 URL <https://openreview.net/forum?id=loDHZstVP6>.

656 Menglin Yang, Jialin Chen, Yifei Zhang, Jiahong Liu, Jiasheng Zhang, Qiyao Ma, Harshit Verma,
657 Qianru Zhang, Min Zhou, Irwin King, and Rex Ying. Low-rank adaptation for foundation models:
658 A comprehensive review, 2024. URL <https://arxiv.org/abs/2501.00365>.

659 Boyu Zhang, Hongyang Yang, Tianyu Zhou, Muhammad Ali Babar, and Xiao-Yang Liu. Enhancing
660 financial sentiment analysis via retrieval augmented large language models. In *ACM International
661 Conference on AI in Finance*, pp. 349–356, 2023.

662 Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating
663 text generation with bert. In *International Conference on Learning Representations*, 2020.

664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

Appendix

A DATASETS

A.1 OVERVIEW

We aim to assess models' capabilities in four key areas of financial expertise—general financial tasks, financial certificate tasks, financial reporting tasks, and financial statement analysis tasks.

A.2 TASK DETAILS

In this subsection, we describe the different tasks on which we benchmarked models.

A.2.1 GENERAL FINANCIAL TASKS

General financial tasks include tasks commonly used to benchmark financial LLMs, such as BloombergGPT Wu et al. (2023) and FinGPT Liu et al. (2023b). We aim to enhance and benchmark our model with three main general capabilities: sentiment analysis, headline analysis, and name entity recognition.

Sentiment Analysis The following datasets all consist of input texts from various sources like news or tweets and sentiment labels such as "positive", "negative", and "neutral".

- **Financial Phrase Bank (FPB)** Malo et al. (2013) contains sentiment-annotated sentences extracted from financial news and reports. We manually created the train/test split with 25% train.
- **Financial question-answering sentiment analysis (FiQA SA)** Maia et al. (2018) is another sentiment analysis dataset with the same labels as FPB from microblog headlines and financial news.
- **Twitter financial news sentiment (TFNS)** Rahman (2022) comprises annotated tweets related to financial news labeled with sentiment categories.
- **News with GPT instruction (NWGI)** Liu et al. (2024e) comprises samples with seven labels ranging from "strong negative" to "strong positive".

Headline classification The Headline dataset Sinha & Khandait (2020) classifies headlines based on various questions into two classes: "yes" and "no".

Named entity recognition (NER) The NER dataset Salinas Alvarado et al. (2015) annotates one entity per sentence, categorized into one of three classes: "location", "person", and "organization"

A.2.2 FINANCIAL CERTIFICATE TASKS

To measure LLMs' performance in financial analysis and regulation compliance, we introduced questions from three financial analyst certification exams and one accountant certification exam. The questions we test models on are from the three levels of Chartered Financial Analyst (CFA) exams as well as the Certified Public Accountant (CPA) REG (Regulation) exam. We manually created the train/test split. We combined the train sets from each exam into one question set to better fine-tune Llama 3.1 8B Instruct due to the overlapping topics of questions. The test sets remain as four separate question sets, with CFA Level I to III and CPA Regulation.

A.2.3 FINANCIAL REPORTING TASK

Financial reporting refers to the tasks of creating financial reports that are in compliance with U.S. Securities and Exchange Commission (SEC) regulations or other regulatory agencies around the world. Specifically, the SEC requires publicly traded companies to file financial reports in XBRL formats using the U.S. Generally Accepted Accounting Principles (GAAP), an accounting standard.

Terminology XBRL Term Han et al. (2024) includes over 6,000 XBRL terminology entries and their explanations from the XBRL International website. The train/test split was sampled randomly.

Tagging XBRL tagging is a key step in creating XBRL reports, where numerical entities in a financial statement such as earning calls must be tagged with US GAAP tags. We use two existing datasets to benchmark tagging performance. FiNER Loukas et al. (2022) includes train and test sets of sentences annotated with 139 types of XBRL tags. We processed the dataset using the batched approach so each input comprises tag options and four questions of highlighted entity tagging, and the answer includes the four correct tags. FNXL Sharma et al. (2023) (Financial Numeric Extreme Labeling) is similar to FiNER but with 2,794 labels. Only a test split was published by the authors and used for our benchmark, so it was tested on the models that were fine-tuned on FiNER. We processed the dataset with a QA style prompt and batched approach to improve efficiency and reduce token usage. The input includes multiple (e.g., four) numerical entities to be tagged simultaneously, and a list of potential US GAAP tag options is provided. Providing the options enables the model to understand which tags are valid choices without needing pre-existing knowledge of the entire taxonomy. The model can infer the appropriate tag from its name and context.

A.2.4 FINANCIAL STATEMENT ANALYSIS TASKS

Financial statement analysis refers to the task of extracting and analyzing information from financial statements or reports. We include three main sub-tasks: financial math, XBRL analysis, and FinanceBench.

Financial Math Financial math Han et al. (2024) is the basics of financial analysis. It tests the model’s capabilities in applying financial formulas to questions. The dataset includes 50 formulas, each with 20 questions. We created a manual train/test split with 16 of those 20 questions for each formula as the train split and the rest as the test split.

XBRL Analysis One approach for financial statement analysis is to directly examine XBRL reports in their raw form. We created the XBRL extraction dataset. The QA style dataset consists of questions and answers derived from XBRL filings from 2019 to 2023 for Dow Jones 30 companies. Each example includes a question, a text segment from an XBRL file containing the answer, and the ground truth. The prompt consists of a portion of a filtered XBRL report in XML and the question.

- **XBRL tag extraction** involves extracting a specific XBRL tag from an XBRL raw text segment given a natural language description of the tag.
- **XBRL value extraction** focuses on extracting a numeric value from the raw XBRL text segment given a natural language description of the value.
- **XBRL formula construction** retrieves multiple tags and outputs a financial formula.
- **XBRL formula calculation** is similar to formula construction but with the goal of outputting formulas substituted with numerical values.

To allow better instruction following for the base model, we use one-shot prompting by providing an example question and answer.

FinanceBench FinanceBench Islam et al. (2023) is another approach for financial analysis which reads the financial reports in PDF format. We include 139 of the 150 questions (some financial report PDFs are no longer accessible) and created a manual train/test split. During dataset pre-processing, we use Google Cloud Document OCR to convert relevant PDF pages to text.

A.3 TASK SIMILARITY

To further measure the correlation between in-category task similarity and multi-task performance, we estimate the in-category task similarity by averaging the cosine similarity between each pair of LoRA adapters within each task category. The cosine similarity is computed by first calculating the effective weight update matrix ($\Delta W = BA$) for each attention layer in a LoRA adapter. These matrices are then flattened into vectors. For any two tasks, we compute the cosine similarity between their

810 corresponding vectors for each shared attention layer and then average these layer-wise similarities.
 811 The final value in the table is the average of these pairwise similarities across all tasks within the
 812 category.

813 The final task similarity and the average single-task and multi-task performance (and their difference)
 814 are displayed below.
 815

816 Table 8: Correlation between Task Similarity and Multi-Task Performance
 817

| 818 Task Category | Avg. Cos. Similarity | Avg. Single-Task Perf. | Avg. Multi-Task Perf. | Avg. Perf. Δ (Multi - Single) |
|----------------------------------|----------------------|------------------------|-----------------------|--------------------------------------|
| 819 General Financial | 0.0043 | 81.97 | 54.75 | -27.22 |
| 820 Financial Reporting | 0.0048 | 52.52 | 54.60 | +2.08 |
| 821 Financial Statement Analysis | 0.0108 | 74.17 | 81.48 | +7.31 |

822 The results show a clear correlation between the similarity of LoRA adapters and the success of
 823 multi-task learning.
 824

- 825 • **Positive Transfer:** The Financial Statement Analysis tasks show the highest average
 826 similarity (0.0108). This corresponds to successful multi-task learning, where the combined
 827 model outperformed the average of single-task models by 7.31 points. This suggests the tasks
 828 in this category are complementary, and their learned parameter updates are constructive.
 829
- 830 • **Negative Interference:** The General Financial and Financial Reporting tasks have low
 831 similarity scores (0.0043 and 0.0048). This dissimilarity leads to interference in the multi-
 832 task setting, causing a significant performance drop of 27.22 points for General Financial
 833 Tasks and only 2.08 points increase in Financial Reporting Tasks compared to their single-
 834 task counterparts. This indicates that the parameter updates required for these tasks are
 835 conflicting or marginally constructive.
 836

837 A.4 INPUT/OUTPUT EXAMPLES 838

839 In this subsection, we provide input and output examples for the tasks on which we benchmarked
 840 models. We only provide the prompt templates without question or answer choice text for financial
 841 certificate tasks due to copyright concerns. Sample CFA and CPA REG questions can be found
 842 online.
 843

844 A.4.1 GENERAL FINANCIAL TASKS 845 846

| 847 Dataset | Financial Phrase Bank (FPB) Malo et al. (2013) |
|-------------|---|
| 848 Task | Sentiment Analysis |
| 849 Input | Instruction: What is the sentiment of this news? Please choose an answer from {negative/neutral/positive}. Input: Pharmaceuticals group Orion Corp reported a fall in its third-quarter earnings that were hit by larger expenditures on R&D and marketing. Answer: |
| 853 Output | negative |

| 858 Dataset | FiQA SA Maia et al. (2018) |
|-------------|---|
| 859 Task | Sentiment Analysis |
| 860 Input | Instruction: What is the sentiment of this news? Please choose an answer from {negative/neutral/positive}. Input: Johnson Matthey raises prospect of investor payout Answer: |
| 863 Output | positive |

| | | |
|-----|---------|---|
| 864 | Dataset | Twitter Financial News Sentiment (TFNS) Rahman (2022) |
| 865 | Task | Sentiment Analysis |
| 866 | Input | What is the sentiment of this tweet? Please choose an answer from {negative/neutral/positive}. Input: \$BYND - JPMorgan reels in expectations on Beyond Meat |
| 867 | | https://t.co/bd0xbFGjkT Answer: negative |
| 868 | Output | negative |
| 869 | | |
| 870 | | |
| 871 | | |
| 872 | | |
| 873 | | |
| 874 | | |
| 875 | | |
| 876 | | |
| 877 | Dataset | News with GPT (NWGI) |
| 878 | Task | Sentiment Analysis |
| 879 | Input | Instruction: What is the sentiment of this news? Please choose an answer from {strong negative/moderately negative/mildly negative/neutral/mildly positive/moderately positive/strong positive}. Input: Amid a soft performance for the major equity indices on Tuesday, Nvidia (NASDAQ: NVDA) posted a particularly glaring loss. Shares continued to fall in sympathy with fellow semiconductor specialist Micron Technology (NASDAQ: MU) following its disappointing earnings results last week. Answer: moderately negative |
| 880 | Output | moderately negative |
| 881 | | |
| 882 | | |
| 883 | | |
| 884 | | |
| 885 | | |
| 886 | | |
| 887 | | |
| 888 | | |
| 889 | | |
| 890 | | |
| 891 | | |
| 892 | | |
| 893 | | |
| 894 | | |
| 895 | | |
| 896 | Dataset | Financial Headline Analysis Sinha & Khandait (2020) |
| 897 | Task | Headline Analysis |
| 898 | Input | Instruction: Does the news headline talk about price? Please choose an answer from Yes/No. Input: Gold futures edge up after two-session decline Answer: No |
| 899 | Output | No |
| 900 | | |
| 901 | | |
| 902 | | |
| 903 | | |
| 904 | | |
| 905 | | |
| 906 | | |
| 907 | | |
| 908 | Dataset | Named Entity Recognition (NER) Salinas Alvarado et al. (2015) |
| 909 | Task | Named Entity Recognition |
| 910 | Input | Instruction: What is the entity type of '40 William St' in the input sentence. Options: person, location, organization Input: This LOAN AND SECURITY AGREEMENT dated January 27, 1999, between SILICON VALLEY BANK (" Bank "), a California - chartered bank with its principal place of business at 3003 Tasman Drive, Santa Clara, California 95054 with a loan production office located at 40 William St., Ste. Answer: location |
| 911 | Output | location |
| 912 | | |
| 913 | | |
| 914 | | |
| 915 | | |
| 916 | | |
| 917 | | |

918 A.4.2 FINANCIAL CERTIFICATE TASKS
919

| 920 Dataset | CFA Level I |
|-------------|--|
| 921 Task | Question Answering |
| 922 Input | Please answer the following question with the exact letter and choice text and no explanations that you choose. Nothing else like an explanation should be in your response. |
| 923 | |
| 924 | |
| 925 | |
| 926 | Question: [CFA Level I context and question omitted due to copyright] |
| 927 | |
| 928 | Choices: [Question choices omitted due to copyright] |
| 929 Output | [Letter of selected answer with answer text omitted due to copyright] |
| 930 | |

| 932 Dataset | CFA Level II |
|-------------|---|
| 933 Task | Question Answering |
| 934 Input | Look at the following context. |
| 935 | [CFA Level II question-set-level context omitted due to copyright] |
| 936 | |
| 937 | Based on this, please answer the following question with the exact letter and choice text and no explanations that you choose. Nothing else like an explanation should be in your response. |
| 938 | |
| 939 | |
| 940 | |
| 941 | Question: [CFA Level II question omitted due to copyright] |
| 942 | |
| 943 | Choices: [Question choices omitted due to copyright] |
| 944 Output | [Letter of selected answer with answer text omitted due to copyright] |
| 945 | |

| 947 Dataset | CFA Level III |
|-------------|---|
| 948 Task | Question Answering |
| 949 Input | Look at the following context. |
| 950 | [CFA Level III question-set-level context omitted due to copyright] |
| 951 | |
| 952 | Based on this, please answer the following question with the exact letter and choice text and no explanations that you choose. Nothing else like an explanation should be in your response. |
| 953 | |
| 954 | |
| 955 | |
| 956 | Question: [CFA Level III question omitted due to copyright] |
| 957 | |
| 958 | Choices: [Question choices omitted due to copyright] |
| 959 Output | [Letter of selected answer with answer text omitted due to copyright] |
| 960 | |

| 962 Dataset | CPA REG |
|-------------|--|
| 963 Task | Question Answering |
| 964 Input | Please answer the following question with the exact letter and choice text and no explanations that you choose. Nothing else like an explanation should be in your response. |
| 965 | |
| 966 | |
| 967 | |
| 968 | Question: [CPA REG question omitted due to copyright] |
| 969 | |
| 970 | Choices: [Question choices omitted due to copyright] |
| 971 Output | [Letter of selected answer with answer text omitted due to copyright] |

A.4.3 FINANCIAL REPORTING TASKS

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

| Dataset | XBRL Term Han et al. (2024) |
|---------------|---|
| Task | XBRL Terminology |
| Input | Explain this XBRL term briefly in one sentence: Frequent Flier Liability, Current. Answer: |
| Output | Value of revenue deferred or cost to provide future products or services, primarily air transportation, associated with programs used by airlines to encourage passenger loyalty by providing rewards geared to the frequency of travel on the sponsoring airline, typically in the form of frequent flyer miles, points, or segments that can be accumulated and converted into free or discounted travel or other redemption options. Used to reflect the current portion of the liability (within one year or within the normal operating cycle if longer) |

| Dataset | FiNER Loukas et al. (2022) |
|---------------|---|
| Task | XBRL Tagging |
| Input | You are XBRL expert. Here is a list of US GAAP tags options: [omitted for this table] Answer the following 4 independent questions by providing only 4 US GAAP tags answers in the order of the questions. Each answer must be separated by a comma (,). Provide nothing else. 1. What is best tag for entity "66,32" in sentence: "Unvested restricted stock outstanding as of March 31, 2020 and 2019 were 42,690 and 66,321 shares, respectively.?" 2. What is best tag for entity "625,000" in sentence: "Unvested AO LTIP Units outstanding as of March 31, 2020 and 2019 were 625,000 and 625,000, respectively.?" 3. What is best tag for entity "625,000" in sentence: "Unvested AO LTIP Units outstanding as of March 31, 2020 and 2019 were 625,000 and 625,000, respectively.?" 4. What is best tag for entity "0.20" in sentence: "Distributions declared per common unit for each of the three month periods ended March 31, 2020 and 2019 was \$ 0.20 per unit.?" Output US GAAP tags:" |
| Output | ShareBasedCompensationArrangementByShareBasedPaymentAwardEquityInstrumentsOtherThanOptionsNonvestedNumber, ShareBasedCompensationArrangementByShareBasedPaymentAwardEquityInstrumentsOtherThanOptionsNonvestedNumber, ShareBasedCompensationArrangementByShareBasedPaymentAwardEquityInstrumentsOtherThanOptionsNonvestedNumber, CommonStockDividendsPerShareDeclared |

| | | |
|------|----------------|---|
| 1026 | Dataset | FiNER Loukas et al. (2022) |
| 1027 | | |
| 1028 | Task | XBRL Tagging |
| 1029 | Input | You are XBRL expert. Here is a list of US GAAP tags options: [omitted for this table] Answer the following 4 independent questions by providing only 4 US GAAP tags answers in the order of the questions. Each answer must be separated by a comma (,). Provide nothing else. |
| 1030 | | 1. What is best tag for entity "66,32" in sentence: |
| 1031 | | "Unvested restricted stock outstanding as of March 31, 2020 and 2019 were 42,690 and 66,321 shares, respectively.?" |
| 1032 | | 2. What is best tag for entity "625,000" in sentence: |
| 1033 | | "Unvested AO LTIP Units outstanding as of March 31, 2020 and 2019 were 625,000 and 625,000, respectively.?" |
| 1034 | | 3. What is best tag for entity "625,000" in sentence: |
| 1035 | | "Unvested AO LTIP Units outstanding as of March 31, 2020 and 2019 were 625,000 and 625,000, respectively.?" |
| 1036 | | 4. What is best tag for entity "0.20" in sentence: |
| 1037 | | "Distributions declared per common unit for each of the three month periods ended March 31, 2020 and 2019 was \$ 0.20 per unit.?" Output US GAAP tags:" |
| 1038 | Output | ShareBasedCompensationArrangementByShareBasedPaymentAwardEquityInstrumentsOtherThanOptionsNonvestedNumber, ShareBasedCompensationArrangementByShareBasedPaymentAwardEquityInstrumentsOtherThanOptionsNonvestedNumber, ShareBasedCompensationArrangementByShareBasedPaymentAwardEquityInstrumentsOtherThanOptionsNonvestedNumber, CommonStockDividendsPerShareDeclared |
| 1039 | | |
| 1040 | | |
| 1041 | | |
| 1042 | | |
| 1043 | | |
| 1044 | | |
| 1045 | | |
| 1046 | | |
| 1047 | | |
| 1048 | | |
| 1049 | | |
| 1050 | | |
| 1051 | | |
| 1052 | | |
| 1053 | | |
| 1054 | | |
| 1055 | | |
| 1056 | | |
| 1057 | | |
| 1058 | | |
| 1059 | | |
| 1060 | | |
| 1061 | | |
| 1062 | | |
| 1063 | | |
| 1064 | | |
| 1065 | | |
| 1066 | | |
| 1067 | | |
| 1068 | | |
| 1069 | | |
| 1070 | | |
| 1071 | | |
| 1072 | | |
| 1073 | | |
| 1074 | | |
| 1075 | | |
| 1076 | | |
| 1077 | | |
| 1078 | | |
| 1079 | | |

| | | |
|------|---|---|
| 1080 | Dataset | FNXL Sharma et al. (2023) |
| 1081 | | |
| 1082 | Task | XBRL Tagging |
| 1083 | Input | You are XBRL expert. Choose the best XBRL US GAAP tag for each highlighted entity in the sentences below. |
| 1084 | | Provide only the US GAAP tags, comma-separated, in the order of the sentences and highlighted entity. Provide nothing else. [Tags omitted for this table] |
| 1085 | | 1. What is the best us gaap tag for entity "6.3" in sentence: "The projected benefit obligation and fair value of plan assets for U.S. pension plans with projected benefit obligations in excess of plan assets was \$6.3 billion and \$4.7 billion, respectively, as of December31, 2019 and \$5.5 billion and \$4.1 billion, respectively, as of December31, 2018."? |
| 1086 | | 2. What is the best us gaap tag for entity "124,043" in sentence: "Capitalized software, net of accumulated amortization of \$124,043 in 2020 and \$104,237 in 2019"? |
| 1087 | | 3. What is the best us gaap tag for entity "1.5" in sentence: "The Company purchased 30 million and 57 million shares under stock repurchase programs in fiscal 2020 and 2019 at a cost of \$1.5 billion and \$3.8 billion, respectively."? |
| 1088 | | 4. What is the best us gaap tag for entity "651,313" in sentence: "This multi-tenant mortgage loan is interest-only with a principal balance due on maturity, and it is secured by seven properties in six states, totaling approximately 651,313 square feet."? |
| 1089 | Output | DefinedBenefitPlanPensionPlanWithProjectedBenefitObligation InExcessOfPlanAssetsProjectedBenefitObligation, CapitalizedComputerSoftwareAccumulatedAmortization, PaymentsForRepurchaseOfCommonStock, AreaOfRealEstateProperty |
| 1090 | | |
| 1091 | | |
| 1092 | | |
| 1093 | | |
| 1094 | | |
| 1095 | | |
| 1096 | | |
| 1097 | | |
| 1098 | | |
| 1099 | | |
| 1100 | | |
| 1101 | | |
| 1102 | | |
| 1103 | | |
| 1104 | | |
| 1105 | | |
| 1106 | | |
| 1107 | | |
| 1108 | | |
| 1109 | | |
| 1110 | | |
| 1111 | | |
| 1112 | A.4.4 FINANCIAL STATEMENT ANALYSIS | |
| 1113 | | |
| 1114 | | |
| 1115 | | |
| 1116 | Dataset | XBRL Tags Extraction |
| 1117 | Task | XBRL Analysis |
| 1118 | Input | You are a knowledgeable XBRL assistant. Your task is to analyze the XBRL context and provide an accurate and very concise answer to the question, The example question can help you to learn the answer format. |
| 1119 | | DO NOT output xml, code, explanation or create new question. |
| 1120 | | Example question: What is the US GAAP XBRL tag for Cash and Cash Equivalents as reported by Example Company Inc for the Fiscal Year ending in FY 2022 |
| 1121 | | Example answer: us-gaap:AnExampleTagName |
| 1122 | | |
| 1123 | | XML File: [XBRL omitted] |
| 1124 | | |
| 1125 | | Question: What is the US GAAP XBRL tag for Total Equity as reported by Home Depot Inc for the Fiscal Year ending in FY 2023? |
| 1126 | | Answer: |
| 1127 | | us-gaap:StockholdersEquity |
| 1128 | Output | |
| 1129 | | |
| 1130 | | |
| 1131 | | |
| 1132 | | |
| 1133 | | |

| | | |
|------|----------------|--|
| 1134 | Dataset | XBRL Values Extraction |
| 1135 | | |
| 1136 | Task | XBRL Analysis |
| 1137 | Input | You are a knowledgeable XBRL assistant. Your task is to analyze the XBRL context and provide an accurate and very concise answer to the question, The example question can help you to learn the answer format. DO NOT output xml, code, explanation or create new question. |
| 1138 | | |
| 1139 | | Example question: What is the value of Example company's income for the Fiscal year ending in FY 2020? |
| 1140 | | Example answer: 80000000 |
| 1141 | | XML File: [XBRL omitted] |
| 1142 | | Question: What is the value of Unitedhealth Group Inc's Total Revenue for the Fiscal Year ending in FY 2023? |
| 1143 | | Answer: |
| 1144 | Output | 371622000000.0 |
| 1145 | | |
| 1146 | | |
| 1147 | | |
| 1148 | | |
| 1149 | | |
| 1150 | | |
| 1151 | | |
| 1152 | | |
| 1153 | | |
| 1154 | | |
| 1155 | | |
| 1156 | | |
| 1157 | | |
| 1158 | | |
| 1159 | | |
| 1160 | | |
| 1161 | | |
| 1162 | | |
| 1163 | | |
| 1164 | | |
| 1165 | | |
| 1166 | | |
| 1167 | | |
| 1168 | | |
| 1169 | Dataset | XBRL Formula Construction |
| 1170 | | |
| 1171 | Task | XBRL Analysis |
| 1172 | Input | You are a knowledgeable XBRL assistant. Your task is to analyze the XBRL context and provide an accurate and very concise answer to the question, The example question can help you to learn the answer format. DO NOT output xml, code, explanation or create new question. |
| 1173 | | |
| 1174 | | Example question: What is the formula for the Gross Profit Margin of Example Inc, formatted with the relevant US GAAP XBRL tags |
| 1175 | | Example answer: us-gaap:ExampleTag / us-gaap:AnotherExampleTag) * 100 |
| 1176 | | XML File: [XBRL omitted] |
| 1177 | | Question: Can you provide the formula of the Cash Flow Margin for Microsoft Corp, using the appropriate US GAAP XBRL tags, as it applies to the FY 2023? * 100') |
| 1178 | | Answer: |
| 1179 | Output | (us-gaap:NetCashProvidedByUsedInOperatingActivities / us-gaap:RevenueFromContractWithCustomerExcludingAssessedTax) * 100 |
| 1180 | | |
| 1181 | | |
| 1182 | | |
| 1183 | | |
| 1184 | | |
| 1185 | | |
| 1186 | | |
| 1187 | | |

| | | |
|------|----------------|---|
| 1188 | Dataset | XBRL Formula Calculation |
| 1189 | Task | XBRL Analysis |
| 1190 | Input | You are a knowledgeable XBRL assistant. Your task is to analyze the XBRL context and provide an accurate and very concise answer to the question, The example question can help you to learn the answer format. |
| 1191 | | DO NOT output xml, code, explanation or create new question. |
| 1192 | | Example question: Can you provide the formula for Operating Profit Margin from Example Corp for the Fiscal Year ending in FY 2022? |
| 1193 | | Example answer: $(50000000 / 3590000000) * 100$ |
| 1194 | | XML File: [XBRL omitted] |
| 1195 | | Question: Can you provide the value for Net Profit Margin from Apple Inc for the Fiscal Year ending in FY 2023? Answer with a formula substituted with values. |
| 1196 | | Answer: |
| 1197 | Output | $(96995000000 / 383285000000) * 100$ |

| | | |
|------|----------------|--|
| 1205 | Dataset | Financial Math Han et al. (2024) |
| 1206 | Task | Math |
| 1207 | Input | Use formula Operating Cash Flow (OCF) to answer the question. Answer with a numerical answer with 2 decimal places and with no explanations or other text. |
| 1208 | | Formula: $OCF = \text{Net Income} + \text{Depreciation} + \text{Changes in Working Capital}$, |
| 1209 | | Explanation: Net Income: Company's total earnings, |
| 1210 | | Depreciation: Depreciation expense, |
| 1211 | | Changes in Working Capital: Change in current assets minus current liabilities. |
| 1212 | | Question: "What is the Operating Cash Flow when a company's financial sheets reflect a Net Income of \$260,000, Depreciation of \$48,000, and a \$35,000 decline in Working Capital?". Answer: |
| 1213 | Output | 343000.0 |

| | | |
|------|----------------|--|
| 1222 | Dataset | FinanceBench Han et al. (2024) |
| 1223 | Task | XBRL QA |
| 1224 | Input | Answer the following question very briefly using scanned text from pages. No explanation needed. |
| 1225 | | Question: Has AMCOR's quick ratio improved or declined between FY2023 and FY2022? |
| 1226 | | If the quick ratio is not something that a financial analyst would ask about a company like this, then state that and explain why. |
| 1227 | | Document Pages Context: [XBRL omitted] Answer: |
| 1228 | Output | The quick ratio has slightly improved from 0.67 times to 0.69 times between FY 2023 and FY 2022.(3.4% jump) |

| | | |
|------|--|--|
| 1236 | | |
| 1237 | | |
| 1238 | | |
| 1239 | | |
| 1240 | | |
| 1241 | | |

B LORA FOUNDATIONS AND METHODS

B.1 WHAT IS LORA?

LoRA is a technique to efficiently update the parameters of pre-trained language models when fine-tuning on new tasks.

B.2 FOUNDATIONS OF LORA

In this subsection, we introduce two fundamental concepts needed to understand LoRA—ranks and fine-tuning.

B.2.1 RANKS

Rank is the number of linearly independent rows or columns in a matrix. Linearly independent columns, for example, are columns whose entries cannot be written as an integer-weighted sum of earlier columns.

$$W = \begin{bmatrix} 1 & 7 & 2 & 8 & 5 \\ 2 & 10 & 4 & 12 & 10 \\ 3 & 15 & 12 & 18 & 27 \\ 4 & 12 & 16 & 16 & 36 \end{bmatrix}, \quad \text{Dimensions: } 4 \times 5 \text{ (rows} \times \text{columns)}$$

In this matrix there are **two** linearly independent columns, so $\text{rank}(W) = 2$.

- Column 1 is independent (nothing precedes it).
- Column 2 cannot be written as a multiple of Column 1, so it is also independent.
- Columns 3–5 are dependent:

$$C_3 = 2C_1, \quad C_4 = C_1 + C_2, \quad C_5 = C_1 + 2C_2.$$

Re-expressing those dependencies in vector form:

$$W = \underbrace{\begin{bmatrix} 1 & 7 \\ 2 & 10 \\ 3 & 15 \\ 4 & 12 \end{bmatrix}}_{B \in \mathbb{R}^{4 \times 2}} \underbrace{\begin{bmatrix} 1 & 0 & 2 & 1 & 1 \\ 0 & 1 & 0 & 1 & 2 \end{bmatrix}}_{A \in \mathbb{R}^{2 \times 5}}.$$

$$\text{Dimensions}(W) = d \times k = 4 \times 5,$$

$$\text{Dimensions}(B) = d \times r = 4 \times 2,$$

$$\text{Dimensions}(A) = r \times k = 2 \times 5,$$

$$\text{Dimensions}(BA) = d \times k = \text{Dimensions}(W).$$

$$\text{Parameters}(W) = 4 \times 5 = 20,$$

$$\text{Parameters}(B) = 4 \times 2 = 8,$$

$$\text{Parameters}(A) = 2 \times 5 = 10,$$

$$\text{Parameters}(BA) = 8 + 10 = 18.$$

Thus storing B and A uses fewer parameters than storing W directly—a key idea behind *low-rank* adaptation (LoRA).

B.2.2 FULL FINE-TUNING

Consider a pre-trained model M with 500 million parameters. Suppose we pre-trained M with two tasks. Task 1 is Masked Language Modeling (MLM), where we mask some words in a sentence, and the task is to predict the sentence with the masked tokens filled in. Task 2 is Next Sentence Prediction

(NSP), where the task is to predict if, given 2 sentences, whether or not sentence A comes before sentence B.

If we want to fine-tune the pre-trained model M on a new task, Named Entity Recognition (NER), where the task is to annotate one entity (location/person/organization) per sentence in a financial task.

When we perform full fine-tuning on model M , all parameters are updated based on the gradients we compute during backpropagation. In backpropagation, we compute the loss (the difference between the predicted output and the target output) and propagate the loss backward through the model. As we propagate the loss backward, we compute the gradient of the loss with respect to each parameter. The optimizer uses these gradients to update the model’s parameters.

If we want to fine-tune model M on another task Financial Phrase Bank (FPB), where the task is to annotate sentences from financial news and reports with sentiment, we still need to update all 500 million parameters. This is costly and can lead to overfitting and the model forgetting pre-training tasks.

B.2.3 FINE-TUNING WITH ADAPTERS (PARAMETER-EFFICIENT FINE-TUNING—PEFT)

Parameter Efficient Fine-Tuning (PEFT) adds small adapter layers per transformer block as shown below. Let’s consider a scenario in which we use PEFT to fine-tune the pre-trained model M and add two adapter layers per transformer layer.

Now, when we fine-tune M on NER, only the adapter parameters are updated. This only consists of a tiny fraction of the original parameters. The rest of the model’s parameters are frozen. This means that, during backpropagation, the gradients of loss pass through them, but the parameters aren’t updated. While we do have to swap the adapters and store the updated parameters separately for FPB, the number of parameters required to fine-tune on FPB is much smaller than full fine-tuning.

B.3 LORA METHODS

In this subsection, we introduce the five LoRA methods we use in our paper. We choose LoRA Hu et al. (2022) and QLoRA Dettmers et al. (2023) due to their standard use in fine-tuning. We choose DoRA Liu et al. (2024d) and rsLoRA Kalajdziewski (2023) due to their performance enhancements: DoRA proposes fine-grained updates for achieving accuracy through LoRA, and rsLoRA proposes a scaling factor to achieve gradient stability. Lastly, we choose LoRA with federated learning for its practical ability to allow financial institutions to collaborate in fine-tuning models while using private, confidential data.

B.3.1 LOW-RANK ADAPTATION (LORA)

LoRA adds a scaled low-rank update $\Delta W = \gamma_r B A$ —where γ_r is a scaling factor ($\gamma_r = \frac{\alpha}{r}$ with $\alpha > 0$ and rank $r > 0$), $B \in \mathbb{R}^{d \times r}$, and $A \in \mathbb{R}^{r \times k}$ —to the frozen pre-trained weight matrix $W_0 \in \mathbb{R}^{d \times k}$.

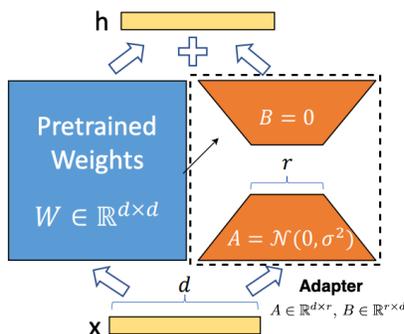


Figure 4: LoRA mechanism Hu et al. (2022).

For each multi-head attention layer, we have query, key, and value weight matrices, which we can factorize as follows:

$$W_Q^{(n)} = B_Q^{(n)} A_Q^{(n)}, \quad W_K^{(n)} = B_K^{(n)} A_K^{(n)}, \quad W_V^{(n)} = B_V^{(n)} A_V^{(n)}.$$

During fine-tuning, the weight matrices are updated as follows with the scaled low-rank update:

$$\begin{aligned} W_{Q,\text{new}}^{(n)} &= W_{Q,\text{old}}^{(n)} + \gamma_r B_Q^{(n)} A_Q^{(n)}, \\ W_{K,\text{new}}^{(n)} &= W_{K,\text{old}}^{(n)} + \gamma_r B_K^{(n)} A_K^{(n)}, \\ W_{V,\text{new}}^{(n)} &= W_{V,\text{old}}^{(n)} + \gamma_r B_V^{(n)} A_V^{(n)}. \end{aligned}$$

Because the update is in-place, no extra layers are added, and inference latency is unchanged.

B.3.2 QUANTIZED LORA (QLORA)

Table 9: Finetuning LLMs with QLoRA methods: listing the number of parameters and GPU memory.

| Models | Parameters | GPU Memory | | Model size |
|----------------------------|------------|------------|-----------|------------|
| | | Batch = 4 | Batch = 8 | |
| Llama-3.1-8B-16bit (base) | 8.03 B | - | - | 16.06 GB |
| Llama-3.1-8B-r8-16bit | 4.72 M | 30.91 GB | 30.91 GB | 16.08 GB |
| Llama-3.1-8B-r8-8bit | 4.72 M | 11.41 GB | 11.81 GB | 8.04 GB |
| Llama-3.1-8B-r8-4bit | 4.72 M | 8.26 GB | 8.65 GB | 4.02 GB |
| Llama-3.1-8B-r4-16bit | 2.36 M | 30.90 GB | 30.90 GB | 16.07 GB |
| Llama-3.1-8B-r4-8bit | 2.36 M | 11.40 GB | 11.78 GB | 8.03 GB |
| Llama-3.1-8B-r4-4bit | 2.36 M | 8.25 GB | 8.61 GB | 4.02 GB |
| Llama-3.1-70B-16bit (base) | 70.56 B | - | - | 151.53 GB |
| Llama-3.1-70B-r8-16bit | 22.28 M | - | - | 151.57 GB |
| Llama-3.1-70B-r8-8bit | 22.28 M | 173.57 GB | 258.17 GB | 75.79 GB |
| Llama-3.1-70B-r8-4bit | 22.28 M | 42.78 GB | 42.78 GB | 37.89 GB |
| Llama-3.1-70B-r4-16bit | 11.14 M | - | - | 151.55 GB |
| Llama-3.1-70B-r4-8bit | 11.14 M | 173.36 GB | 258.11 GB | 75.70 GB |
| Llama-3.1-70B-r4-4bit | 11.14 M | 42.73 GB | 42.73 GB | 37.89 GB |

When fine-tuning, LoRA requires a large amount of GPU memory. To solve this issue, we can use QLoRA. QLoRA drastically reduces memory usage and lets you fine-tune on a single GPU.

In QLoRA, we quantize the weights of the adapter layers, reducing both parameter count and memory usage. Quantization is a technique that reduces the precision of the weights to reduce the number of bits used to store them. It consists of two parts: Rounding to the nearest integer and truncating to remove the decimal portion of a floating point number. QLoRA specifically uses 4-bit NormalFloat (NF4), an optimal data type for normally distributed weights, quantization. Pre-trained weights are usually normally distributed and centered around 0, which is why NF4 is ideal for quantization.

If we quantize from Float16 to Int4, we can represent 16 different values (bins) because Int4 has 4 bits and $2^4 = 16$. Inputs are usually normalized from -1 to 1. Very close together values, however, will be mapped to the same bin. This means that the precision is lost if we want to convert back to Float16. However, we can use blockwise quantization, where we divide the input range into blocks and quantize each block separately. QLoRA uses a 64 blocksize for better precision.

Since regular quantization relies on the bins being equally probable, QLoRA uses NormalFloat where the bins are weighted by the normal distribution. The spacing between bins is therefore closer together near 0 and further apart further away from 0.

Each block in QLoRA has a quantization constant. QLoRA employs double quantization, where it quantizes the quantization constants themselves to further save space.

The last part of QLoRA is paged optimizers. Paged optimizers reduce GPU memory spikes by switching pages to CPU memory when GPU RAM becomes full when processing long sequences, and the pages are not needed for the current computation of the forward/backward pass.

The forward pass for QLoRA is $\mathbf{y} = p_{16}(\mathbf{W}_0^{\text{NF4}})\mathbf{x} + \gamma_r \mathbf{B}\mathbf{A}\mathbf{x}$.

B.3.3 WEIGHT-DECOMPOSED LOW-RANK ADAPTATION (DORA)

LoRA makes simple changes to the model weights, so it sometimes doesn't capture the full complexity of the data and its relationships. DoRA solves this issue of capturing data complexity.

DoRA decomposes the weight matrix into a *magnitude vector* and a *direction matrix*. The magnitude vector consists of the lengths of the columns in the weight matrix and is computed by taking each column's ℓ_2 norm. The direction matrix \mathbf{V} is the collection of the original columns. Its unit-column form $\hat{\mathbf{V}} = \mathbf{V}/\|\mathbf{V}\|_c$ is obtained by dividing each column by its ℓ_2 norm.

The magnitude vector \mathbf{m} is of size $1 \times k$, where k is the number of columns. The direction matrix \mathbf{V} is of size $d \times k$, where d is the number of rows in a weight matrix.

The decomposition can be written as:

$$\mathbf{W}_0 = \mathbf{m} \frac{\mathbf{V}}{\|\mathbf{V}\|_c} = \|\mathbf{W}_0\|_c \frac{\mathbf{W}_0}{\|\mathbf{W}_0\|_c},$$

where $\|\cdot\|_c$ denotes the column-wise ℓ_2 norm (i.e., the norm is taken independently for each column) and \mathbf{W}_0 is the frozen pretrained weight.

Here is an example of the decomposition:

$$\mathbf{W}_0 = \begin{bmatrix} 1 & 7 & 2 & 8 & 5 \\ 2 & 10 & 4 & 12 & 10 \\ 3 & 15 & 12 & 18 & 27 \\ 4 & 12 & 16 & 16 & 36 \end{bmatrix}, \quad \mathbf{W}_0 \in \mathbb{R}^{4 \times 5}.$$

For column j

$$\|\mathbf{w}_j\|_2 = \sqrt{\sum_{i=1}^4 W_{ij}^2}.$$

These norms form a 1×5 magnitude vector:

$$\mathbf{m} = [5.4772, 22.7596, 20.4939, 28.0713, 46.3681].$$

The unit-column direction matrix is

$$\hat{\mathbf{V}} = \begin{bmatrix} 0.182574 & 0.307562 & 0.097590 & 0.284988 & 0.107833 \\ 0.365148 & 0.439375 & 0.195180 & 0.427482 & 0.215666 \\ 0.547723 & 0.659062 & 0.585540 & 0.641223 & 0.582297 \\ 0.730297 & 0.527250 & 0.780720 & 0.569976 & 0.776396 \end{bmatrix}.$$

Every column of $\hat{\mathbf{V}}$ now has unit length:

$$\|\mathbf{v}_j\|_2 = 1, \quad \text{for all } j.$$

These are updated separately. The magnitude vector \mathbf{m} is trained directly, while the direction matrix \mathbf{V} is fine-tuned using LoRA: $\Delta\mathbf{V} = \mathbf{B}\mathbf{A}$ with $\mathbf{B} \in \mathbb{R}^{d \times r}$ and $\mathbf{A} \in \mathbb{R}^{r \times k}$.

After the updates, the new weight matrix is

$$\mathbf{W}' = \mathbf{m} \frac{\mathbf{V} + \Delta\mathbf{V}}{\|\mathbf{V} + \Delta\mathbf{V}\|_c} = \mathbf{m} \frac{\mathbf{W}_0 + \mathbf{B}\mathbf{A}}{\|\mathbf{W}_0 + \mathbf{B}\mathbf{A}\|_c}.$$

B.3.4 RANK-STABILIZED LORA (RSLORA)

LoRA scales the weight matrix update $\mathbf{B}\mathbf{A}$ by α/r , which can cause gradients to explode or diminish as the rank r increases. In contrast, rsLoRA uses a scaling factor α/\sqrt{r} : $\mathbf{W}' = \mathbf{W}_0 + \frac{\alpha}{\sqrt{r}}\mathbf{B}\mathbf{A}$. This scaling results in gradient-scale stability at higher ranks, enabling the rank to be higher to capture more details in long-context tasks like XBRL extraction. rsLoRA also results in lower perplexity—the model assigns higher probabilities to correct words—than LoRA at higher ranks.

1458 B.3.5 LORA WITH FEDERATED LEARNING

1459
1460 In the finance sector, multiple banks may want to work together on a model to predict credit risk and
1461 whether a borrower will default on a loan. Each bank may have a different dataset, but they cannot
1462 share their data due to compliance reasons and privacy concerns. Federated learning solves this issue
1463 by fine-tuning a model on local data and aggregating updates during backpropagation to a centralized
1464 model via a server.

1465 Differentially Private Low-Rank Adaptation (DP-LoRA) Liu et al. (2025) is a method to use federated
1466 learning with LoRA.

1467 DP-LoRA first uses a server to send the current global LoRA weights (the A and B matrices from
1468 earlier) to all clients.

1469 Every client does the following: 1) Gets a minibatch of its private data 2) Computes the gradient
1470 for only its local A and B weights clipped with an ℓ_2 norm (square root of the sum of the squares
1471 of elements in the vector) 3) Adds Gaussian noise to the gradients 4) Updates the A and B LoRA
1472 matrices 5) Sends the updated A and B matrices to the server.

1473
1474 By adding noise, DP-LoRA prevents the centralized model from inferring the private data later on.
1475 This would allow the banks in the credit risk example to work on a model together.

1476 As in normal federated learning, the server then aggregates the weights from all clients in a weighted
1477 average and sends the updated weights to all clients.

1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511

C EXPERIMENTAL SETUP

C.1 GPUS AND APIS USED

For Llama 3.1 8B Instruct, we use four NVIDIA RTX A5000 GPUs, each with 24GB of GPU memory, for all fine-tuning and a single A5000 for all inference. For Gemini 2.0 Flash Lite, we use Google Cloud Vertex AI for fine-tuning and inference. For the rest of the baseline models, we use the OpenAI, Together AI, DeepSeek, and Fireworks.ai API services for inference.

C.2 FINE-TUNING

C.2.1 HYPERPARAMETERS AND CONFIGURATIONS

All Llama models fine-tuning experiments were conducted using the Axolotl library, using its configuration files to manage hyperparameters. The base model for all experiments was meta-llama/Llama-3.1-8B-Instruct. A consistent learning rate of 1×10^{-4} was utilized across all fine-tuning runs. Table 10 outlines the common hyperparameters and settings derived from the Axolotl configuration template that were largely consistent across experiments, unless specified otherwise in the experiment-specific configurations. For Gemini 2.0 Flash Lite, we use Google Cloud Vertex AI for fine-tuning using the same number of epochs as Llama models and default settings for the rest of the hyperparameters. The Gemini 2.0 Flash Lite hyperparameters are mostly hidden by Google.

Table 10: Common Axolotl Hyperparameters and Settings.

| Parameter | Value / Setting |
|--------------------------------|--|
| Base Model | meta-llama/Llama-3.1-8B-Instruct |
| Model Type | LlamaForCausalLM |
| Tokenizer Type | AutoTokenizer |
| Chat Template | llama3 |
| Special Pad Token | < end_of_text > |
| Sequence Length | 4096 |
| Adapter (underlying mechanism) | lora |
| LoRA Alpha | 16 |
| LoRA Dropout | 0.05 |
| LoRA Target Modules | q_proj, v_proj, k_proj |
| Optimizer | adamw_bnb_8bit (for 8-bit quantization); adamw_torch_fused (for 4-bit quantization) |
| Learning Rate Scheduler | cosine |
| Base Learning Rate | 1×10^{-4} |
| Warmup Steps | 10 |
| Weight Decay | 0.0 |
| Gradient Checkpointing | Enabled |
| BF16 Precision | auto |
| TF32 Precision | Disabled |
| Flash Attention | Disabled |
| DeepSpeed Configuration | deepspeed_configs/zerol.json |
| Dataset Format | Instruction-following: context (instruction), target (output). |
| Format | [INST] {instruction} [/INST] |
| Validation Set Size | 2% of training data |
| W&B Project Name | finlora_models |
| Logging Steps | Every 500 steps |
| Evaluations per Epoch | 4 |
| Checkpoints per Epoch | 1 |

Table 11 provides a detailed breakdown of the hyperparameters for each specific fine-tuning run, grouped by the PEFT method employed.

1566 Table 11: Experiment-specific Fine-tuning Hyperparameters. All runs use
 1567 meta-llama/Llama-3.1-8B-Instruct and a learning rate of 1×10^{-4} .

| 1569 | Adapter Task | Train Dataset | Epochs | Batch size | Grad. Accum. | Effective batch size |
|------|---|----------------------------|--------|------------|--------------|----------------------|
| 1570 | <i>LoRA, 8-bit Quantization, Rank 8</i> | | | | | |
| 1571 | Sentiment | finlora_sentiment_train | 4 | 8 | 2 | 16 |
| 1572 | Headline | headline_train | 4 | 8 | 2 | 16 |
| 1573 | NER | ner_train | 4 | 8 | 2 | 16 |
| 1574 | Certificate | certificate_train | 4 | 8 | 2 | 16 |
| 1575 | XBRL Term | xbml_term_train | 1 | 4 | 2 | 8 |
| 1576 | XBRL Tagging | finer_train_batched | 4 | 1 | 8 | 8 |
| 1577 | XBRL Analysis | xbml_extract_train | 1 | 1 | 8 | 8 |
| 1578 | Formula | formula_train | 1 | 4 | 2 | 8 |
| 1579 | FinanceBench | financebench_train | 4 | 1 | 2 | 2 |
| 1580 | <i>QLoRA, 4-bit Quantization, Rank 4</i> | | | | | |
| 1581 | Sentiment | finlora_sentiment_train | 4 | 8 | 2 | 16 |
| 1582 | Headline | headline_train | 4 | 8 | 2 | 16 |
| 1583 | NER | ner_train | 4 | 8 | 2 | 16 |
| 1584 | Certificate | certificate_train | 4 | 8 | 2 | 16 |
| 1585 | XBRL Term | xbml_term_train | 1 | 4 | 2 | 8 |
| 1586 | XBRL Tagging | finer_train_batched | 4 | 1 | 8 | 8 |
| 1587 | XBRL Analysis | xbml_extract_train | 1 | 1 | 8 | 8 |
| 1588 | Formula | formula_train | 1 | 4 | 2 | 8 |
| 1589 | FinanceBench | financebench_train | 4 | 1 | 2 | 2 |
| 1590 | <i>DoRA, 8-bit Quantization, Rank 8</i> | | | | | |
| 1591 | Sentiment | finlora_sentiment_train | 4 | 8 | 2 | 16 |
| 1592 | Headline | headline_train | 4 | 8 | 2 | 16 |
| 1593 | NER | ner_train | 4 | 8 | 2 | 16 |
| 1594 | Certificate | certificate_train | 4 | 8 | 2 | 16 |
| 1595 | XBRL Term | xbml_term_train | 1 | 4 | 2 | 8 |
| 1596 | XBRL Tagging | finer_train_batched | 4 | 1 | 8 | 8 |
| 1597 | XBRL Analysis | xbml_extract_train | 1 | 1 | 8 | 8 |
| 1598 | Formula | formula_train | 1 | 4 | 2 | 8 |
| 1599 | FinanceBench | financebench_train | 4 | 1 | 2 | 2 |
| 1600 | <i>rsLoRA, 8-bit Quantization, Rank 8</i> | | | | | |
| 1601 | Sentiment | finlora_sentiment_train | 4 | 8 | 2 | 16 |
| 1602 | Headline | headline_train | 4 | 8 | 2 | 16 |
| 1603 | NER | ner_train | 4 | 8 | 2 | 16 |
| 1604 | Certificate | certificate_train | 4 | 8 | 2 | 16 |
| 1605 | XBRL Term | xbml_term_train | 1 | 4 | 2 | 8 |
| 1606 | XBRL Tagging | finer_train_batched | 4 | 1 | 8 | 8 |
| 1607 | XBRL Analysis | xbml_extract_train_batched | 1 | 1 | 8 | 8 |
| 1608 | Formula | formula_train | 1 | 4 | 2 | 8 |
| 1609 | FinanceBench | financebench_train | 4 | 1 | 2 | 2 |

1602

1603

1603 C.3 INFERENCE

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619

We use temperature 0.0 for all inference. For all fine-tuned Llama models, we use 8-bit quantization for inference. For all inference, the maximum new token lengths for inference for various financial datasets are: 10 tokens for FPB, FiQA SA, TFNS, NWGI, NER, and Headline; 20 tokens for Tag Extraction and Value Extraction; 30 tokens for Formula Construction and Formula Calculation; 50 tokens for XBRL Term, FinanceBench, and Financial Math; and 100 tokens for FiNER and FNXL.

D XBRL TAGGING/ANALYSIS DEMO

To help readers understand XBRL tasks and how our fine-tuned models' performance compares with base models' performance, we have created an interactive demo showcasing two tasks: XBRL tagging and XBRL analysis. The user can select one of the built-in examples as input or enter their custom prompt.

D.1 XBRL TAGGING DEMO

XBRL Enhanced LLM Interactive Demo

XBRL Tagging
XBRL Extraction

Use Case

- XBRL tagging is a key step in creating XBRL reports. Numerical entities in texts such as earning calls can be tagged with US GAAP tags.
- This demo previews the capabilities of our fine-tuned model for XBRL tagging. It can be further integrated with existing work flows.

Usage

- Input:** Provide a sentence containing financial information.
- Output:** Key entities and their corresponding US GAAP (Generally Accepted Accounting Principles) tags will be generated by the base model and our fine-tuned model.

Sentence

The effective tax rate was 47.6 percent and 13.4 percent for the three months ended March 31, 2020 and 2019, respectively.

Clear

Submit

Llama 3.1 8b (base) output

| Entites | US GAAP tags |
|---------|----------------------------------|
| 47.6 | us-gaap_IncomeTaxExpenseEfctRate |
| 13.4 | us-gaap_EffectiveTaxRate |

Llama 3.1 8b (fine-tuned for XBRL tagging) output

| Entites | US GAAP tags |
|---------|-------------------------------------|
| 47.6 | EffectiveIncomeTaxRateContinuingOpe |
| 13.4 | EffectiveIncomeTaxRateContinuingOpe |

Ground Truth Answer

| Entites | US GAAP tags |
|---------|-------------------------------------|
| 47.6 | EffectiveIncomeTaxRateContinuingOpe |
| 13.4 | EffectiveIncomeTaxRateContinuingOpe |

Figure 5: XBRL tagging demo.

The XBRL tagging demo showcases base and fine-tuned models designed to identify and tag numerical entities within financial texts, such as earning calls, with appropriate US GAAP (Generally Accepted Accounting Principles) taxonomies—a critical step in generating XBRL reports. Users can input a sentence containing financial data, and the system will output the key entities and their corresponding US GAAP tags as identified by both a base model and the enhanced fine-tuned model, illustrating the latter's advanced capabilities and potential for integration into existing workflows.

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

D.2 XBRL ANALYSIS DEMO

XBRL Enhanced LLM Interactive Demo

[XBRL Tagging](#) [XBRL Analysis](#)

Use Case

- Analyze an existing XBRL report with ease using our fine-tuned model as a chatbot.
- The model allows extraction of US GAAP tags, values, or financial formulas from the XBRL report.
- Current support for XBRL extraction is limited to small portion of a XBRL file.

Usage

- **Input:** A financial question and an XBRL file name.
- **Output:** The answer to the question will be generated by the base model and our fine-tuned model.

Feel free to explore the examples below or enter your own question.

Question

Can you provide the value for Operating Profit Margin from Salesforce Inc for the Fiscal Year ending in FY 2023? Answer with a formula substituted with values.

XBRL File Name

crm-20230131/crm-20230131_htm.xml_context_ib41f5e45110a4b88b9616fd4fdb14e1b_D20220201-20230131

Clear
Submit

[Examples](#)

Llama 3.1 8b (Base) output

$(21962000000 / 31352000000) * 100 = 70.05$

Llama 3.1 8b (fine-tuned for XBRL extraction) output

$(1030000000 / 31352000000) * 100 = 3.29$

Ground truth answer

$(1030000000 / 31352000000) * 100 = 3.29$

Flag

Figure 6: XBRL analysis demo.

The XBRL analysis demo showcases base and fine-tuned models acting as chatbots to simplify the analysis of existing XBRL reports. It enables users to extract US GAAP tags, values, or financial formulas, although current support is limited to analyzing small portions of an XBRL file. To use it, a user inputs a financial question along with an XBRL file name, and the system provides answers generated by both a base model and the fine-tuned model.

1728 E DATASHEET FOR XBRL ANALYSIS DATASETS

1729
1730 The following section answers the questions listed in Datasheets for Datasets Gebru et al. (2021) for
1731 our newly constructed XBRL datasets.
1732

1733 E.1 MOTIVATION

1734
1735 • **For what purpose were the datasets created?**

1736 XBRL is a key format in financial report, but XBRL data analysis is rarely explored in LLM
1737 benchmark.

1738
1739 • **Who funded the creation of the datasets? If there is an associated grant, please provide
1740 the name of the grantor and the grant name and number.**

1741 N/A.

1742 E.2 COMPOSITION

1743
1744 • **What do the instances that comprise the datasets represent (e.g., documents, photos,
1745 people, countries)?**

1746 The datasets represent questions and answers derived from XBRL reports.
1747

1748
1749 • **How many instances are there in total (of each type, if appropriate)?**

1750 In total, there are about 34.2k instances in the new XBRL datasets. There are around 13k
1751 instances of Tags Extraction (10.1k train; 2.9k test), 12.6k instances of Values Extraction
1752 (10.1k train; 2.5k test), 4.3k instances of Formula Construction (3.4k train; 835 test), and 4.3k
1753 instances of Formula Calculation (3.4k train; 835 test).

1754
1755 • **Do the datasets contain all possible instances or is it a sample (not necessarily random) of
1756 instances from a larger set?**

1757 The datasets contain all instances of the questions we curated.

1758
1759 • **What data does each instance consist of?**

- 1760 – instruction: The instruction for the LLM to work on this task.
- 1761 – input: The question.
- 1762 – output: The ground truth answer.
- 1763 – year: The year of the XBRL report the extract is based on.
- 1764 – company: The ticker symbol of the company the XBRL report belongs to.
- 1765 – doc_path: The relative path of the XBRL report required for this entry. Zipped XBRL
1766 reports are within the Hugging Face dataset repository in Git LFS format.
- 1767 – context_id: In XBRL, the `context_id` is the unique identifier for a `<context>`
1768 element, which defines crucial metadata for reported financial facts, such as the reporting
1769 period, and any applicable dimensional information (like business segments or scenarios).
1770 We provide the relevant `context_id` for each entry in this dataset to narrow down the
1771 scope of data the LLM needs to consider. This pre-selection of the context ensures that
1772 the LLM can focus on the question using a targeted set of facts, rather than search the
1773 entire XBRL report. Finding the appropriate `context_id` is not considered part of
1774 the LLM's task for several reasons: for questions related to consolidated, company-wide
1775 figures for a specific year (which is the focus of our dataset), the correct `context_id`
1776 can generally be identified programmatically. This typically involves filtering the context
1777 definitions by the target date or year, and then selecting the context that represents the
1778 consolidated entity (often identifiable by its lack of dimensional axis information). This
1779 process is a deterministic and straightforward data preprocessing step that can be efficiently
1780 implemented with code, making it unnecessary to leverage an LLM for this aspect of data
1781 retrieval.

1782
1783 • **Is there a label or target associated with each instance?**

Yes, each question in the dataset has an associated ground truth answer, referred to as "output".

- 1782
- 1783 • **Is any information missing from individual instances? If so, please describe why.**
- 1784 No.
- 1785
- 1786 • **Are relationships between individual instances made explicit (e.g., users' movie ratings, social-network links)?**
- 1787 N/A. Each instance is generally treated as independent.
- 1788
- 1789 • **Are there recommended data splits (e.g., training, development/validation, testing)?**
- 1790 Yes, the data is split chronologically by year. Questions based on XBRL reports from 2019-
- 1791 2022 form the training split, while questions based on 2023 XBRL reports constitute the
- 1792 test split. This approach is used to ensure no data leakage, as XBRL reports often contain
- 1793 comparative data from previous years.
- 1794
- 1795 • **Are there any errors, sources of noise, or redundancies in the dataset?**
- 1796 To the best of our knowledge, no.
- 1797
- 1798 • **Are the datasets self-contained, or do they link to or otherwise rely on external resources?**
- 1799 The dataset entries do not directly embed the full raw XBRL text, even when filtered by
- 1800 `context_id`, due to the considerable size of these texts. However, we provide code in our
- 1801 GitHub repository that allows users to fetch the source XBRL documents and reconstruct the
- 1802 complete data for each entry. Otherwise, the structured data provided is self-contained.
- 1803
- 1804 • **Do the datasets contain data that might be considered confidential?**
- 1805 No.
- 1806
- 1807 • **Do the datasets contain data that, if viewed directly, might be offensive, insulting, threaten-**
- 1808 **ing, or might otherwise cause anxiety?**
- 1809 No.
- 1810
- 1811 • **Do the datasets relate to people?**
- 1812 No.
- 1813
- 1814 • **Do the datasets identify any sub-populations (e.g., by age, gender)?**
- 1815 No.
- 1816
- 1817 • **Is it possible to identify individuals, directly or indirectly, from the datasets?**
- 1818 No.
- 1819

1820 E.3 COLLECTION PROCESS

- 1821
- 1822 • **How was the data associated with each instance acquired? Was the data directly ob-**
- 1823 **servable (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or**
- 1824 **indirectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses**
- 1825 **for age or language)? If data was reported by subjects or indirectly inferred/derived from**
- 1826 **other data, was the data validated/verified? If so, please describe how.**
- 1827 The questions were generated based on 22 financial concepts and 7 financial formulas. The
- 1828 ground truth answer were extracted from the XBRL documents using the Arelle package.
- 1829
- 1830 • **What mechanisms or procedures were used to collect the data? How were these mecha-**
- 1831 **nisms or procedures validated?**
- 1832 N/A.
- 1833
- 1834 • **If the datasets are a sample from a larger set, what was the sampling strategy?**
- 1835 No. The provided datasets are complete.
- 1836
- 1837 • **Who was involved in the data-collection process and how were they compensated?**
- 1838 N/A.

- 1836 • **Over what timeframe was the data collected? Does this timeframe match the creation**
1837 **timeframe of the data?**
1838 N/A.
1839
- 1840 • **Were any ethical-review processes conducted (e.g., by an institutional review board)?**
1841 No.
1842
- 1843 • **Do the datasets relate to people?**
1844 No.
1845
- 1846 • **Did you collect the data directly from the individuals in question, or obtain it via third**
1847 **parties or other sources?**
1848 N/A.
- 1849 • **Were the individuals in question notified about the data collection?**
1850 N/A.
1851
- 1852 • **Did the individuals in question consent to the collection and use of their data?**
1853 N/A.
1854
- 1855 • **If consent was obtained, were the consenting individuals provided with a mechanism to**
1856 **revoke their consent in the future or for certain uses?**
1857 N/A.
- 1858 • **Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a**
1859 **data-protection impact analysis) been conducted?**
1860 No.
1861

1862 E.4 PREPROCESSING / CLEANING / LABELING

- 1863
- 1864 • **Was any preprocessing, cleaning, or labeling of the data done?**
1865 The Arelle package was used to process XBRL files in order to create the ground truth answer.
1866
- 1867 • **Was the raw data saved in addition to the preprocessed/cleaned/labeled data?**
1868 Raw XBRL files will be available for download at Hugging Face.
1869

1870 E.5 USES

- 1871 • **Have the datasets been used for any tasks already?**
1872 The dataset is used for XBRL analysis tasks.
1873
- 1874 • **Is there a repository that links to any or all papers or systems that use the datasets?**
1875 N/A.
1876
- 1877 • **What other tasks could the datasets be used for?**
1878 XBRL/financial statement analysis related tasks.
1879
- 1880 • **Is there anything about the datasets' composition or collection that might impact future**
1881 **uses?**
1882 No.
1883
- 1884 • **Are there tasks for which the datasets should *not* be used?**
1885 N/A.

1886 E.6 DISTRIBUTION

- 1887
- 1888 • **Will the datasets be distributed to third parties outside the entity on behalf of which the**
1889 **dataset was created?**
Dataset will be released publicly after review period ends.

- 1890
- **When will the datasets be distributed?**
1891 The datasets are available now.
1892
 - **Will the datasets be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?**
1893
1894
1895 The datasets are released under the MIT License.
1896
 - **Have any third parties imposed IP-based or other restrictions on the data?**
1897 No.
1898
 - **Do any export controls or other regulatory restrictions apply to the datasets or to individual instances?**
1899
1900
1901 No.
1902

1903 E.7 MAINTENANCE

- 1904
- **Who will support, host, or maintain the datasets?**
1905 We, the authors of the FinLoRA paper, will.
1906
 - **Is there an erratum?**
1907 No.
1908
 - **If the datasets relate to people, are there applicable limits on the retention of the data associated with the instances (e.g., were the individuals in question told that their data would be retained for a fixed period of time and then deleted)?**
1909
1910
1911
1912 N/A.
1913
 - **If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so?**
1914
1915 Yes, using Open Review Comments.
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943