

Selective Fine-tuning on LLM-labeled Data May Reduce Reliance on Human Annotation: A Case Study Using Schedule-of-Event Table Detection

Bhawesh Kumar

Verily Life Sciences

269 East Grand Avenue, South San Francisco

BHAWESHK@VERILY.COM

Jonathan Amar

Verily Life Sciences

269 East Grand Avenue, South San Francisco

JONATHANAMAR@VERILY.COM

Eric Yang

Verily Life Sciences

269 East Grand Avenue, South San Francisco

ERYANG@VERILY.COM

Nan Li

Verily Life Sciences

269 East Grand Avenue, South San Francisco

NOTANUMBER@VERILY.COM

Yugang Jia

Verily Life Sciences

269 East Grand Avenue, South San Francisco

YUGANG@VERILY.COM

Abstract

Large Language Models (LLMs) have demonstrated their efficacy across a broad spectrum of tasks in healthcare applications. However, often LLMs need to be fine-tuned on task-specific expert-annotated data to achieve optimal performance, which can be expensive and time consuming. In this study, we fine-tune PaLM-2 (Anil et al. (2023)) with parameter efficient fine-tuning (PEFT) using noisy labels obtained from Gemini-pro 1.0 (Google (2024)) for the detection of Schedule-of-Event (SoE) tables, which specify care plan in clinical trial protocols. We introduce a filtering mechanism to select high-confidence labels for this table classification task, thereby reducing the noise in the auto-generated labels. We find that the fine-tuned PaLM-2 with filtered labels outperforms Gemini Pro 1.0 and other LLMs on this task and achieves performance close to PaLM-2 fine-tuned on non-expert human annotations. Our results show that leveraging LLM-generated labels, coupled with strategic filtering can be a viable and cost-effective strategy for improving LLM performance on specialized tasks, especially in domains where expert annotations are scarce, expensive, or time-consuming to obtain.

1. Introduction

Large Language Models (LLMs) have been found to be useful across diverse tasks like natural language understanding and generation, question-answering, summarization, programming, and creative arts (Chen et al. (2021); Radford et al. (2018, 2019); Ramesh et al. (2021)). LLMs are particularly promising in specialized fields such as healthcare, where they can significantly enhance clinical decision-making, patient care, drug discovery, and the management and utilization of medical data (Singhal et al. (2023); Ingraham et al. (2023); Tu et al. (2024a,b); Sharma et al. (2024)). However, the successful application of LLMs in specialized domains frequently depends on their ability to process and understand complex, domain-specific structured and unstructured content, which often requires fine-tuning the models with data annotated by experts (van Aken (2023)). This necessity presents considerable challenges, primarily due to the scarcity, high cost, and substantial time required to acquire expert annotations in fields like healthcare. In response to these challenges, our work investigates the potential of LLM-generated labels for fine-tuning purposes, with a specific case-study on identifying Schedule-of-Event (SoE) tables in clinical trial protocols. The accurate identification of SoE tables, which outlines plan-of-care in clinical trials (also see Appendix A for more details on SoE tables), plays a pivotal role in the digitization of clinical trial protocols which we briefly describe below.

1.1. Brief Introduction to Clinical Trial Protocols and Digitization

Clinical trials are the backbone of medical research. However, the traditional conduct of clinical trials is fraught with inefficiencies at various stages including patient recruitment, follow-ups, data acquisition and handling (Inan et al. (2020); Marquis-Gravel et al. (2019)). Clinical trials rely heavily on manual processes, leading to time-consuming, expensive, and error-prone workflows. The inefficiencies pose challenges to all stakeholders involved in the trial and also slow down the pace of medical research (Getz and Campo (2017); Jones et al. (2016)).

Clinical trial protocols are foundational documents in the trials, outlining the detailed methodologies, objectives, and care plans that guide the conduct of studies in accordance with regulatory, ethical, and scientific standards. These protocols include critical components such as the Schedule of Events (SoE) table, which details the plan of care for participants, including visits for screening, treatment, and follow-up phases, along with the assessments, treatments, and data collection scheduled for these visits. The digitization of clinical trial protocols refers to the process of converting these detailed and often voluminous paper-based documents into accurate digital workflows (Verily Life Sciences (2023); Rosa et al. (2021); Inan et al. (2020)). This transformation is not just a matter of changing the medium but involves the systematic identification, classification and ultimately extraction of key elements within the protocols, such as SoE tables, to ensure they are accurately captured and can be effectively managed and analyzed in a digital system (Inan et al. (2020)). Correctly identifying these tables, which can vary significantly in formatting, terminology, and layout across different protocols, poses a significant challenge (refer to Appendix A for more details on SoE tables as well as examples.) However, the accurate classification of such tables are crucial for any automated protocol digitization workflow; an undetected SoE table can lead to an incomplete care plan, while a misclassified table introduces erroneous

information into the system, underscoring the paramount importance of reliability in this process.

1.2. Improving Domain-Specific LLM Performance with Synthetic Labels

1.2.1. MODELING SoE DETECTION WITH LLMs

We model the problem of SoE table classification as a binary classification problem and use an LLM (PaLM-2) for accurate classification of SoE tables. To improve PaLM-2’s ability to accurately classify SoE tables with fine-tuning, we use Gemini-pro 1.0 to auto-generate training labels for fine-tuning task. This strategy aims to address the challenges of acquiring expert annotations by leveraging the capabilities of LLMs to produce high-quality, task-specific data.

1.2.2. FINE-TUNING LLM WITH LLM-GENERATED LABELS

The labels obtained from LLMs for specialized tasks like SoE table classification can be quite noisy. Thus, for a fine-tuning task to succeed, we need to remove potentially incorrect labels from auto-generated labels. For our specific task of SoE table classification, we use the consensus in Gemini-pro 1.0 model inference across dual data representations of tables – JSON and text representations – to reduce noise in the training dataset for PaLM-2. Specifically, we fine-tune PaLM-2 models on only those LLM labels, where the JSON and text based inferences of the tables are identical for the label generating LLM (Gemini-pro 1.0 in this case.)

The JSON representation of the table, which represents each of the table columns as a dictionary with the key being the row number and the value being the cell value for that column, preserves the structural details of the table. In contrast, the text representation encompasses not only the contents within the table but also all surrounding text on the page, including footnotes, titles, and any other textual content. This comprehensive capture of page content provides a fuller context and valuable redundancy for our inference process, improving the model’s ability to accurately interpret and classify the tables.

We find that enhancing the quality of the auto-generated dataset for fine-tuning PaLM-2 leads to substantial improvement in fine-tuned model performance. The PaLM-2 model trained with these subset of LLM generated labels outperforms both the baseline PaLM-2 and Gemini-pro 1.0 on SoE detection task (see Table 3). Remarkably, the fine-tuned PaLM-2 model achieves performance levels close to those obtained with human-annotated labels, showcasing the effectiveness of using LLM-generated labels for domain-specific tasks, particularly in settings where expert annotations are sparse. Finally, we explore a hybrid approach where a limited number of low-confidence LLM-generated labels, as identified by our filtering mechanism (i.e., instances where JSON and text-based inferences disagree), are selectively replaced with human annotations. This strategy aims to leverage the efficiency of LLM-generated labels while mitigating potential concerns regarding biases that might arise from discarding labels. Although the performance of the models fine-tuned with human and Gemini annotations with the filtering mechanism is already comparable, this hybrid approach utilizing only 10% of human labeling further narrows the gap with the PaLM-2 model trained solely with human labels as seen in Table 3.

1.3. Generalizable Insights about Machine Learning in the Context of Healthcare

Our work on SoE table classification, a highly specialized healthcare task, provides valuable insights into the potential of Large Language Models (LLMs) for advancing machine learning applications in healthcare. Our work demonstrates that LLM-generated labels, when combined with filtering and selective human validation, may offer a scalable and cost-effective alternative to traditional expert annotation for fine-tuning domain-specific models. This has significant implications for healthcare, where access to expert annotations is often limited due to cost, time constraints, and the limited availability of experts.

We summarize the key insights below:

- **LLMs may be effectively applied to highly specialized healthcare tasks and have the potential to transform manual workflows:** Our results show that even for tasks that require a deep understanding of domain-specific terminology and concepts, like SoE table classification, LLMs can potentially achieve strong performance, especially after fine-tuning. This adaptability of LLMs highlights their potential to automate and streamline traditionally manual and error-prone processes in healthcare.
- **LLM-generated labels offer a promising solution for addressing annotation challenges:** Leveraging LLMs to generate labels presents a scalable and cost-effective alternative to traditional manual and expert-based annotation, potentially accelerating the development of ML models in healthcare.
- **Strategic filtering of LLM generated labels is crucial for optimizing performance:** Our filtering mechanism, based on dual data representations, significantly improves the quality of LLM-generated labels and subsequent fine-tuned model performance. Our hybrid labeling strategy shows that selective human annotation on difficult examples may further improve performance and bridge the gap with models fine-tuned entirely on human labels.

2. Related Work

2.1. Large Language Models in Healthcare

Recent advances in natural language processing (NLP) and machine learning have significantly enhanced the potential for integrating these technologies into various aspects of healthcare, including clinical decision-making, patient care, drug discovery, and medical information management. A wealth of studies have underscored the capabilities of Large Language Models (LLMs) in performing crucial NLP tasks in healthcare and medicine, such as extracting medical information, summarizing patient information, facilitating automated diagnosis, and even passing board certification exams in specialty medicines (Liu et al. (2021); Shay et al. (2024); Van Veen et al. (2023); Ingraham et al. (2023); Tu et al. (2024b,a)). These applications highlight the potentially transformative impact LLMs could have on healthcare.

In the context of clinical trials, LLMs have been utilized to parse and understand interventions and findings from randomized control trials (Wadhwa et al. (2023)), and to assist

in patient-matching for clinical trials by analyzing electronic health records (EHRs) alongside clinical trial documentation (Yuan et al. (2023)). Previous research has also studied the problem of automated identification of specific elements from Schedule-of-Event (SoE) tables, such as detailed activity information, employing a human-in-the-loop approach to ensure accuracy and relevance (Dhuliawala et al. (2018)).

These emerging applications not only underscore the versatility of LLMs in managing diverse and complex healthcare datasets but also illustrate a pivotal challenge: the dependency on extensive, expert-annotated datasets for fine-tuning and evaluating LLMs in specialized tasks. This has led to a growing interest in automated label generation techniques and the exploration of fine-tuning and testing of LLMs with these synthetic labels.

2.2. Model Training on Synthetic Dataset

Successes of generative models in various tasks have spurred research into leveraging these models to augment real data for model fine-tuning and validation. Besnier et al. (2019), for example, use class-conditional GAN generated image for training model for image classification tasks. He et al. (2023) study the potential of synthetic data in zero-shot and few-shot classification using CLIP model (Radford et al. (2021)). Recently, researchers have also used LLMs to augment data for various classification tasks. Meng et al. (2022) use a pre-trained language model to generate samples by prompting it with real data and using the generated data for fine-tuning a BERT model. To control the quality of samples, they use log-probability of generated samples for filtering poor quality auto-generated samples. Yoo et al. (2021) use randomly sampled existing data samples to condition models to generate new samples, while using token probability corresponding to the label-classes to obtain soft probability for these generated samples. These soft probabilities for synthetic samples are used to train BERT-style models for classification. One of the recent studies by Li et al. (2023) has tried to understand when synthetic data can be helpful in successful model training. They find that synthetic data is less effective when a classification task is subjective or when a specific instance of data to be classified is subjective as measured by agreement amongst annotators. In the healthcare domain, Feder et al. (2024) have explored counterfactual data augmentation for improved LLM generalization showing promising result on text classification tasks.

Building on these previous research, our study employs LLM generated labels for fine-tuning another LLM for table classification task in a highly specialized context, specifically, Schedule-of-Event table classification in clinical trial protocols. Distinct from previous research, which often relies on standard benchmarks or datasets for generating synthetic data, our work showcases a novel application of synthetic labels for fine-tuning in domains where expert annotation is expensive and challenging to obtain. Furthermore, we offer a detailed comparison between models fine-tuned on LLM-generated labels versus those fine-tuned on labels annotated by human experts. Our approach also introduces an innovative label filtering mechanism that utilizes dual data representations of tables for removing potentially noisy synthetic labels and allowing for selective annotation by human annotators. Finally, our approach doesn't require access to logits for tokens and can be applied even when working with LLMs through black-box API access.

3. Methods

3.1. Problem Set-up

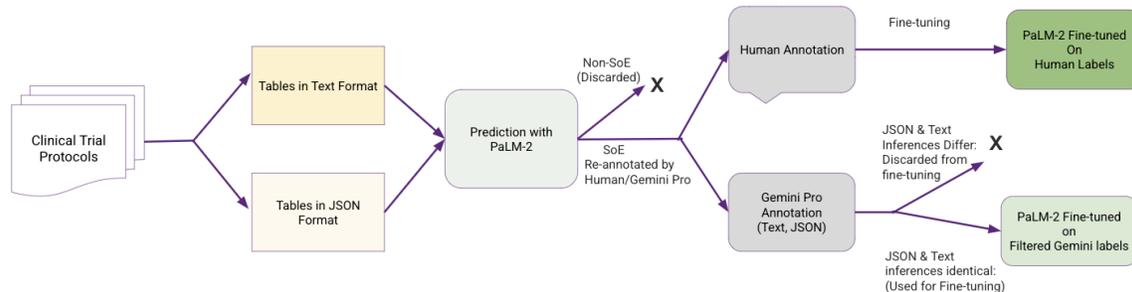


Figure 1: **Schematic of the fine-tuning process.** This figure illustrates our approach to fine-tuning PaLM-2 for SoE table classification using LLM-generated labels. A base PaLM-2 model is first applied to identify potential SoE tables in 408 unlabelled protocols. Only predicted SoE tables are then annotated by both Gemini-pro (using JSON and text representations) and human annotators. Our results (Table 3) show that our filtering approach (where we discard Gemini-pro labels with different inferences for JSON and text view before fine-tuning) achieves performance close to that of fine-tuning on human labels, highlighting the potential of LLM-generated labels for specialized healthcare tasks.

We frame the Schedule-of-Event table detection in a clinical protocol as a binary classification task of correctly classifying a table as a SoE table or a non-SoE table. Specifically, for each clinical trial protocol, the goal is to classify all the tables present inside that protocol as SoE or non-SoE table. We define a table as a SoE table when our in-house protocol digitization specialists label it as a SoE table. The goal is to achieve a very high level of precision and recall on table classification task. Since we don’t expect the model to be perfect in classification of the table, the classification algorithm is supposed to be used for reducing the annotator’s burden of going through every page in a long protocol. All protocols digitized through this semi-automated approach with human-in-the-loop goes through stringent review and quality checks to ensure accuracy.

3.2. Dataset & Models

3.2.1. TRAINING AND TEST SET

Our data set consists of a total of 499 clinical trial protocols of which 91 are expert-labeled by a team of five protocol digitization specialists and are used as the test set in our experiments. These 91 test protocols have a total of 3019 tables with 411 SoE tables (13.6%) and 2608 non-SoE tables (86.4%.) These expert-digitizers are specifically trained to manually label and digitize the clinical protocol for our in-house clinical trial management system (CTMS) software and the labeling process and digitization requires significant domain knowledge, time and effort. We take the expert annotations as ground truth for all experiments. The subset of 408 protocols that don’t have any expert-labels are used for the fine-tuning tasks (see Figure 1). Of the 408 protocols, we randomly select 300 as training set, 18 as

validation set and 90 as test set for model fine-tuning task. Note that we use a separate expert-labeled test data set of 91 protocols for all model evaluations distinct from the one used during fine-tuning. These 499 protocols in our experiments span a diverse set of clinical trials across pharmaceutical companies, academic organizations, hospitals, and government organizations.

3.2.2. MODELS

We use GPT-4 API (gpt-4-0613) (OpenAI (2024)), PaLM-2 (text-bison@001 on GCP) (Anil et al. (2023)), and Gemini-pro 1.0 (Google (2024)) for our inference tasks. The base models (without any fine-tuning) serve as the baselines. We use the PaLM-2 model for all the fine-tuning experiments. We note that Gemini-pro 1.0 and GPT-4 models have been reported as having substantially better performance on LLM benchmarks than PaLM-2 (Google (2024)). The Gemini-pro 1.0 model is not available for fine-tuning as of this writing.

3.3. Selective Human and LLM Annotation

For our fine-tuning task, we selectively collect human and Gemini-pro 1.0 annotations on previously mentioned 408 protocols. The annotation is done by a team of six non-experts annotators and they can mark complex cases for review by the expert annotators as well as directly ask about any specific annotation from an expert.

We show the fine-tuning process workflow in Figure 1. We first do inference of all 408 unlabelled protocols with the PaLM-2 model. On a subset of 60 protocols, we manually go through PaLM-2 model prediction with the help of experts to find specific patterns in incorrect model prediction. We find that the base PaLM-2 has a very high recall but also a very high false positive rate and often predicts trivial cases of non-SoE tables as SoE. Thus, we only obtain human and Gemini-pro 1.0 annotations on tables identified as SoE by the base PaLM-2 model (around 25% of all tables.) This selective annotation allows us to keep the size of annotation tasks manageable (by reducing the task to one-fourth), while also helps us over-sample the SoE table examples for fine-tuning. Additionally, this approach allows annotators to concentrate their efforts on more ambiguous cases potentially leading to higher quality annotations since annotators can spend more time on each item.

We summarize the results of non-expert and Gemini-pro 1.0 based annotations in Table 1. The train, validation, and test set for fine-tuning consist of 300, 18, and 90 protocols respectively. The number of SoE and non-SoE table annotation counts for non-expert and Gemini-pro 1.0 annotations differ (since neither the non-expert human annotators nor Gemini-pro 1.0 are perfect at identifying SoE tables and they may annotate a specific table differently), but total table counts are the same across various data splits.

We emphasize that we do not use expert annotators directly for labeling tasks. However, the annotators do have some previous experience with annotation for SoE tables and they also have access to expert annotators for any annotation they need help with. Additionally, they can choose to not annotate a table and leave it as “Do not know”. These are later annotated by an expert. Despite access to experts, non-expert annotations can be noisy due to variation in skills among the non-expert annotators. On random overlapping sets of 50 annotation, the average inter-rater agreement among non-expert annotators is 81.2%. We note that all annotations are collected only on tables predicted as SoE by base PaLM-

Table 1: Summary of Annotations

Annotation Type	Train Set	Validation Set	Test Set
Non-Expert	1536 SoE, 1264 Non-SoE	53 SoE, 74 Non-SoE	383 SoE, 413 Non-SoE
Gemini-pro	1748 SoE, 1052 Non-SoE	64 SoE, 63 Non-SoE	490 SoE, 306 Non-SoE

2 models as described previously. Thus, the inter-rate agreement is only on a portion of all tables present in the protocols that are annotated by human experts. For tables with multiple annotations, we choose the first annotation collected for that table.

Finally, to get a sense of the alignment between between Gemini-pro and the non-expert labels, we looked at the overlap between these two sets of annotations. We found that 89.6% of the annotations between human labelers and Gemini-pro are identical. We show the agreement matrix for human and Gemini-pro labels in Table 2. We can see that majority of the disagreement are cases where Gemini-pro classifies a table as SoE, but human raters don’t. This is in-line with a high recall and low precision observed for Gemini-pro model as seen in Table 3 on the test set.

Table 2: Agreement between Gemini-pro and Non-Expert Labelers

		Non-Expert Labelers	
		SoE	Non-SoE
Gemini-pro	SoE	1497	251
	Non-SoE	39	1013
Overall Agreement		89.6%	

3.4. Experiments

We use PaLM-2, Gemini-pro 1.0 and GPT-4 (gpt-4-0613) in our experiments. As described in Section 1.2.2, we use JSON and text representations of a table for inference. We use camelot (0.11.0) (Camelot Developers (2023)) for extracting JSON representations and pdfminer.six (20231228) (pdfminer.six Developers (2023)) for text extraction. The model is asked to respond with “YES” or “NO” corresponding to the prompts (see Appendix B for the prompts) for these JSON and text representations of the table. If either of the JSON or text inference output corresponding to a table is “YES”, we classify the table as SoE. This conservative approach for classification of SoE leads to higher false positives, but those are more easier to rectify in our digitization workflow than missed SoE tables, which can lead to missed plan-of-care and cause expensive manual corrections at later steps in the protocol digitization process.

For fine-tuning the PaLM-2 model, we use the 408 protocols as previously described in Section 3.2.1. We fine-tune all PaLM-2 models on table annotations obtained from 300 protocols and use 18 protocols for validation and the rest 90 protocols as test set for

our fine-tuning process. We fine-tune three sets of models with Gemini-pro 1.0 generated annotations and one model with human annotations. The first Gemini-pro based fine-tuned model uses all 2800 table annotations from 300 protocols, while the second removes all “noisy labels” from fine-tuning. Specifically, for the second fine-tuning experiment with Gemini-pro annotations, we remove all samples from training where the JSON and text annotations of Gemini-pro are not identical. This reduces the set of table annotations to 2518 tables in the training set. Finally, for the third fine-tuned model, instead of discarding the 282 samples (having different text and JSON inferences), we replace them with human annotations. All models are fine-tuned for 300 epochs with learning rate multiplier of 1, early stopping set to True, and an evaluation interval of 10 epochs with Google Cloud Vertex AI fine-tuning pipeline which uses Parameter Efficient Fine-tuning (PEFT). We track the model training through a tensorboard instance.

4. Results

We evaluate models on a comprehensive set of metrics to assess the effectiveness of each model in the context of clinical trial protocol digitization. The models are benchmarked based on recall, precision, F-1 score, and accuracy. We additionally measure model performance at various precision threshold as well as on the percentage of protocols achieving 100% recall and precision (refer Appendix D), which are critical for the practical deployment of the automated digitization pipeline. To estimate the uncertainty in our performance metrics, we use bootstrapping with 10,000 replications. This involves repeatedly resampling the test dataset with replacement and calculating the metrics for each resampled dataset. The 95% confidence intervals (CIs) are then derived from the distribution of these bootstrapped estimates.

4.1. Baselines

We start with a very simple baseline of non-finetuned models—Gemini-pro 1.0, GPT-4, and PaLM-2. We have summarized the results in Table 3. We notice that all baseline models achieve a very high recall. Among the models that are not fine-tuned, the inference with GPT-4 results in best performance with a precision of 78.2% (73.1, 83.2), f1-score of 0.845 (0.807, 0.881) and an accuracy of 94.0% (91.2, 96.2). The inference with the PaLM-2 base model achieves 59.8% (54.7, 65.0) precision, 0.710 (0.667, 0.752) f1-score and 87.6% (84.1, 90.7) accuracy. The inference with Gemini-pro 1.0 results in a performance between PaLM-2 and GPT-4 with 65.7% (60.6, 70.9) precision, 0.761 (0.719, 0.802) f1-score and 90.0% (86.8, 92.8) accuracy. In addition to these baselines, we also use baselines with naive combinations of Gemini-pro 1.0 and PaLM-2 prediction (see Appendix C for details.)

4.2. Fine-tuned Models

We run four sets of fine-tuning experiments using PaLM-2 models. We first conduct fine-tuning on PaLM-2 model using the non-expert human annotation obtained as described in Section 3.3. This serves as a strong benchmark for our remaining fine-tuning experiments that use Gemini-pro 1.0 based annotations for fine-tuning PaLM-2. The second and third fine-tuning experiments differ in the sense that while one of the experiments use entirety

Table 3: Average Recall, Precision, F-1 Score, and Accuracy on the Test Set (Values reported as mean (95% CI))

Model Type	Model/Training Data	Recall %	Precision %	F-1 Score	Accuracy %
Non-Fine-Tuned	PaLM-2	97.3 (94.9, 99.3)	59.8 (54.7, 65.0)	0.710 (0.667, 0.752)	87.6 (84.1, 90.7)
	GPT-4 (gpt-4-0613)	98.6 (97.1, 99.8)	78.2 (73.1, 83.2)	0.845 (0.807, 0.881)	94.0 (91.2, 96.2)
	Gemini-pro 1.0	99.4 (98.4, 100.0)	65.7 (60.6, 70.9)	0.761 (0.719, 0.802)	90.0 (86.8, 92.8)
Fine-Tuned PaLM-2	Human Labels	98.9 (97.3, 100.0)	87.3 (82.9, 91.4)	0.908 (0.875, 0.938)	96.0 (93.5, 98.0)
	All Gemini labels	100 (100.0, 100.0)	63.8 (58.1, 69.2)	0.744 (0.698, 0.788)	88.1 (84.2, 91.5)
	Filtered Gemini labels	97.7 (95.5, 99.5)	85.9 (81.2, 90.3)	0.894 (0.858, 0.927)	95.7 (93.1, 97.7)
	Human + Gemini labels	98.9 (97.3, 100.0)	86.4 (81.6, 90.7)	0.903 (0.867, 0.934)	95.8 (93.2, 97.8)

of Gemini-pro’s labels during fine-tuning, the other experiment is fine-tuned on only those Gemini-pro labels where there is a consensus between Gemini-pro’s inference for JSON and text-based representation of a given table. Finally, we explore a hybrid approach to further refine the fine-tuning process. Instead of discarding the samples lacking consensus between JSON and text-based inferences, we replace the Gemini-pro labels for those samples with human annotations. This results in a fine-tuning dataset consisting of 2518 Gemini-annotated samples and 282 human-annotated ones.

As seen in Table 3, fine-tuning PaLM-2 with labels generated by Gemini-pro 1.0 leads to improvements over the base PaLM-2 model’s performance. However, this improvement is nuanced. When the model is fine-tuned using the entirety of Gemini-pro’s labels, it improves over the baseline PaLM-2 model. However, the precision, f1-score, and accuracy compared to the standalone Gemini-pro 1.0 model remains inferior. Optimal fine-tuning is achieved through only incorporating labels for which there is a consensus between Gemini-pro’s JSON and text-based inferences for a table. This fine-tuned variant (Table 3 second last row), not only exceeds the performance of the base PaLM-2 and Gemini-pro 1.0 models, but also narrows the gap to the precision and f1-score of PaLM-2 fine-tuned with human labels. Our hybrid labeling strategy (Table 3 last row), incorporating human annotations for cases lacking consensus, further bridges the gap with the model fine-tuned on human annotations.

Finally, our fine-tuned models surpass the performance of naive ensemble approaches combining Gemini-pro 1.0 and PaLM-2, as detailed in Appendix C. Yet, it is worth noting that the naive ensembles still offer better results than either standalone base model. These findings underscore the value of naive ensembles for preliminary analysis and for scenarios where fine-tuning isn’t feasible, while also highlighting the potentially superior results attainable with fine-tuned models.

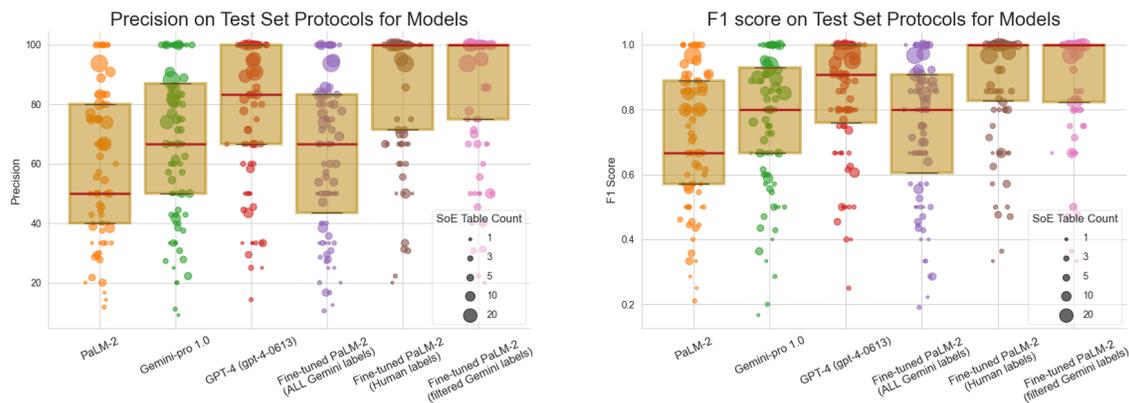


Figure 2: **Precision and F1 Score across models for 91 protocols in test set** (hybrid approach not shown for brevity.) Bubble sizes represent the number of SoE tables within a protocol. PaLM-2 models fine-tuned with human labels and consensus-based Gemini-pro 1.0 labels achieve a median precision of 100% and F1 score of 1.

We plot the results for precision and f1-score in Figure 2 corresponding to models in Table 3 for all 91 test protocols. The bubble size in the scatter plots corresponding to the protocols are proportional to the number of SoE tables present in that protocol. We also overlay the boxplot on the scatter plot to show the the precision and f1-scores across individual protocols in the test set. We see that PaLM-2 models fine-tuned on human labels as well as consensus-based Gemini-pro 1.0 labels achieve a median precision of 100% and median f1-score of 1. This means that for at least 50% of the protocols, the SoE table detection step in digitization workflow will be processed correctly without needing any further correction.

5. Discussion

Our study proposes a novel approach to fine-tuning Large Language Models for specialized domains where labels can be especially difficult to obtain. The approach of utilizing noisy labels from an LLM (Gemini-pro) for fine-tuning another LLM (PaLM-2) demonstrates a scalable and cost-effective alternative to conventional expert annotation processes. Notably, the introduction of a label filtering mechanism, which selects labels for fine-tuning only when there is agreement between dual data representations (JSON and text) of a table, leads to substantial improvements in model performance. This method not only outperforms the base model and the label-generating LLM, it also approaches the performance level of models fine-tuned with human annotations on our table classification task, highlighting the potential of our fine-tuning strategy to effectively leverage auto-generated labels. The filtering mechanism also provides an easy way to identify potential samples that may require expert labeling. As demonstrated by our hybrid labeling approach, this can significantly reduce the annotation workload while potentially bringing fine-tuned models on par with models fine-tuned on human labels. This targeted approach to human annotation allows for a more efficient allocation of resources, maximizing the benefits of both automated and expert labeling. While we use JSON and text based consensus approach as a proxy for

selecting high quality training labels for our specific table classification task, any multi-modal data representation can serve a similar purpose in filtering out potentially noisy labels.

Despite its strengths, our approach has several limitations. The task of SoE table detection is highly specialized and relies on data from our internal Clinical Trial Management System (CTMS) software, which may not fully capture the variety and complexity of clinical trial protocols encountered in broader applications. This specificity could limit the generalizability of our findings to other types of documents or on different table classification tasks. If LLMs perform poorly across the board on a specific task, automated generation of labels may not be feasible even with powerful models like Gemini-pro. Moreover, while our study underscores the feasibility and effectiveness of using LLM-generated labels for fine-tuning, it lacks a direct comparison with a baseline model fine-tuned on expert annotations due to the high costs and resource requirements associated with obtaining such annotations. This comparison could have provided a clearer benchmark for evaluating the relative performance of our approach. Importantly, the reliance on auto-label generation and consensus-based fine-tuning may introduce or perpetuate biases inherent in the models used for label generation. Depending on specific context and fine-tuning task, these biases may manifest as demographic, entity, or domain-specific biases, affecting the accuracy and fairness of the fine-tuned model, particularly in sensitive domains like healthcare. While our hybrid approach, incorporating human annotations for challenging or low-confidence cases, can help mitigate some of these biases, comprehensive fairness and bias evaluation specific to both the auto-labeling and fine-tuning steps are still essential for detecting and mitigating biases in auto-generated labels and fine-tuned models. This can pave the way for broader adoption of LLMs in data-rich, but expert-scarce domains like healthcare, where processes often rely heavily on manual workflows.

6. Data & Code Availability

Due to the terms of our data sharing agreement, we are unable to provide access to the dataset. Additionally, the code used in this study is part of our proprietary software and cannot be shared.

7. Author Contribution

BK and YJ conceived of utilizing LLM-based labels for fine-tuning. JA, EY, and BK conceived of the initial fine-tuning experiments using human labeled data. NL came up with the idea and initial prompt to use an LLM for table classification task. BK conducted all the experiments and wrote the draft of the paper. All authors reviewed the draft and edited the manuscript and take responsibility for all aspects of the work.

8. Ethics Declaration

8.1. Competing interests

All authors are employees and shareholders of Verily Life Sciences LLC.

References

- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. Palm 2 technical report, 2023.
- Victor Besnier, Himalaya Jain, Andrei Bursuc, Matthieu Cord, and Patrick Pérez. This dataset does not exist: training models from generated images. *CoRR*, abs/1911.02888, 2019. URL <http://arxiv.org/abs/1911.02888>.
- Camelot Developers. *Camelot: PDF Table Extraction for Humans*, 2023. URL <https://camelot-py.readthedocs.io/>.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021.

- Murtaza Dhuliawala, Nicholas Fay, Daniel Gruen, and Amar Das. What happens when? interpreting schedule of activity tables in clinical trial documents. In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, pages 301–306, 2018.
- Amir Feder, Yoav Wald, Claudia Shi, Suchi Saria, and David Blei. Causal-structure driven augmentations for text ood generalization. *Advances in Neural Information Processing Systems*, 36, 2024.
- Kenneth A Getz and Rafael A Campo. Trial watch: trends in clinical trial design complexity. *Nature Reviews Drug Discovery*, 16(5):307–308, 2017.
- Google. Gemini: A family of highly capable multimodal models, 2024.
- Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and Xiaojuan Qi. Is synthetic data from generative models ready for image recognition?, 2023.
- Omer T Inan, P Tenaerts, Sheila A Prindiville, HR Reynolds, DS Dizon, K Cooper-Arnold, M Turakhia, Mark J Pletcher, Kenzie L Preston, Harlan M Krumholz, et al. Digitizing clinical trials. *NPJ digital medicine*, 3(1):101, 2020.
- John B Ingraham, Max Baranov, Zak Costello, Karl W Barber, Wujie Wang, Ahmed Ismail, Vincent Frappier, Dana M Lord, Christopher Ng-Thow-Hing, Erik R Van Vlack, et al. Illuminating protein space with a programmable generative model. *Nature*, 623(7989): 1070–1078, 2023.
- W Schuyler Jones, Matthew T Roe, Elliott M Antman, Mark J Pletcher, Robert A Harrington, Russell L Rothman, William J Oetgen, Sunil V Rao, Mitchell W Krucoff, Lesley H Curtis, et al. The changing landscape of randomized clinical trials in cardiovascular disease. *Journal of the American College of Cardiology*, 68(17):1898–1907, 2016.
- Zhuoyan Li, Hangxiao Zhu, Zhuoran Lu, and Ming Yin. Synthetic data generation with large language models for text classification: Potential and limitations. *arXiv preprint arXiv:2310.07849*, 2023.
- Zhichao Liu, Ruth A Roberts, Madhu Lal-Nag, Xi Chen, Ruili Huang, and Weida Tong. Ai-based language models powering drug discovery and development. *Drug Discovery Today*, 26(11):2593–2607, 2021.
- Guillaume Marquis-Gravel, Matthew T Roe, Mintu P Turakhia, William Boden, Robert Temple, Abhinav Sharma, Boaz Hirshberg, Paul Slater, Noah Craft, Norman Stockbridge, et al. Technology-enabled clinical trials: transforming medical evidence generation. *Circulation*, 140(17):1426–1436, 2019.
- Yu Meng, Jiaxin Huang, Yu Zhang, and Jiawei Han. Generating training data with language models: Towards zero-shot language understanding. *Advances in Neural Information Processing Systems*, 35:462–477, 2022.
- OpenAI. Gpt-4 technical report, 2024.

- pdfminer.six Developers. *pdfminer.six: PDF parser and analyzer*, 2023. URL <https://github.com/pdfminer/pdfminer.six>.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation, 2021.
- Carmen Rosa, Lisa A Marsch, Erin L Winstanley, Meg Brunner, and Aimee NC Campbell. Using digital technologies in clinical trials: Current and future applications. *Contemporary clinical trials*, 100:106219, 2021.
- Puneet Sharma, Guangze Luo, Cindy Wang, Dara Brodsky, Camilia R Martin, Andrew Beam, and Kristyn Beam. Assessment of the clinical knowledge of chatgpt-4 in neonatal-perinatal medicine: a comparative analysis with chatgpt-3.5. *Journal of Perinatology*, pages 1–2, 2024.
- Denys Shay, Bhawesh Kumar, Simone Redaelli, Dario von Wedel, Manqing Liu, Mark Dershwitz, Maximilian S Schaefer, and Andrew Beam. Could chatgpt-4 pass an anaesthesiology board examination? follow-up assessment of a comprehensive set of board examination practice questions. *British Journal of Anaesthesia*, 132(1):172–174, 2024.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180, 2023.
- Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaekermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Charles Lau, Ryutaro Tanno, Ira Ktena, Anil Palepu, Basil Mustafa, Aakanksha Chowdhery, Yun Liu, Simon Kornblith, David Fleet, Philip Mansfield, Sushant Prakash, Renee Wong, Sunny Virmani, Christopher Semturs, S. Sara Mahdavi, Bradley Green, Ewa Dominowska, Blaise Aguerre y Arcas, Joelle Barral, Dale Webster, Greg S. Corrado, Yossi Matias, Karan Singhal, Pete Florence, Alan Karthikesalingam, and Vivek Natarajan. Towards generalist biomedical ai. *NEJM AI*, 1(3): AIoa2300138, 2024a. doi: 10.1056/AIoa2300138. URL <https://ai.nejm.org/doi/abs/10.1056/AIoa2300138>.
- Tao Tu, Anil Palepu, Mike Schaekermann, Khaled Saab, Jan Freyberg, Ryutaro Tanno, Amy Wang, Brenna Li, Mohamed Amin, Nenad Tomasev, et al. Towards conversational diagnostic ai. *arXiv preprint arXiv:2401.05654*, 2024b.

- Betty van Aken. Exploration and adaptation of large language models for specialized domains. 2023.
- Dave Van Veen, Cara Van Uden, Louis Blankemeier, Jean-Benoit Delbrouck, Asad Aali, Christian Bluethgen, Anuj Pareek, Malgorzata Polacin, Eduardo Pontes Reis, Anna Seehofnerova, et al. Clinical text summarization: Adapting large language models can outperform human experts. *Research Square*, 2023.
- Verily Life Sciences. Verily Viewpoint: Site CTMS and Protocol Digitization. Technical report, Q3 2023. URL https://assets.verily.com/m/5a2000e85ed78214/original/Verily-Viewpoint-Site-CTMS-ProtDig_FeatureArticle_Q3-2023-1.pdf.
- Somin Wadhwa, Jay DeYoung, Benjamin Nye, Silvio Amir, and Byron C Wallace. Jointly extracting interventions, outcomes, and findings from rct reports with llms. In *Machine Learning for Healthcare Conference*, pages 754–771. PMLR, 2023.
- Kang Min Yoo, Dongju Park, Jaewook Kang, Sang-Woo Lee, and Woomyoung Park. GPT3Mix: Leveraging large-scale language models for text augmentation. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2225–2239, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-emnlp.192. URL <https://aclanthology.org/2021.findings-emnlp.192>.
- Jiayi Yuan, Ruixiang Tang, Xiaoqian Jiang, and Xia Hu. Llm for patient-trial matching: Privacy-aware data augmentation towards better performance and generalizability. In *American Medical Informatics Association (AMIA) Annual Symposium*, 2023.

Appendix A. Schedule of Event Tables

In clinical trials, a schedule of events table is a table outlining the timeline and sequence of assessments, procedures, and data collection that will take place during the study. This table is an important part of the study protocol and provides a comprehensive overview of the study activities for both the researchers and participants. The schedule of events table typically includes the following information:

1. **Study Visits:** This includes different study visits or assessment time-points, such as screening, baseline, treatment periods, follow-up visits, and the end of study. Typically the timing of each visit (e.g., day, week, month) are also specified.
2. **Assessments and Procedures:** The Schedule-of-events table also describes the various assessments, tests, and procedures that will be performed at each study visit. This may include informed consent, physical examinations, vital sign measurements, laboratory tests, imaging studies, patient-reported outcomes, and any other relevant data collection.
3. **Data Collection:** The table includes the data that will be collected at each study visit, such as adverse events, concomitant medications, and any other relevant information.

We provide a sample SoE table (refer Table 4) based on [NIH template](#) and two non-SoE tables (refer Tables 5 and 6) below. The first table has clear screening, treatment and follow-up period. It specifies various visit with time information as well as window during which the visit can take place. The second table looks like an SoE table in terms of structure, but it doesn't have clearly demarcated screening, treatment and follow-up. Further, it only specifies specific lab tests at the start of the diagnosis and completion of therapy and lacks treatment period information. Often this table would require protocol digitization specialists to look at additional context (like surrounding texts on the page) in the protocol to determine whether or not it is a SoE table. The last table specifies pharmacokinetic collections and is not a SoE table (see prompts in Appendix B which we wrote in consultation with the digitizers for SoE tables)

Note that we are unable to provide identical sample tables from our own dataset due to limitations on data sharing. The first example of the SoE table is taken from the NIH template for SoE tables and the last two tables are fictitious and are not from any actual clinical trial protocols.

A.1. Example SoE and Non-SoE Tables

Table 4: Sample SoE Table

Assessment	Screening:	Treatment Visits					Follow-up
	Visit (Day -14 to -1)	Baseline, Enrollment, Visit 1 (Day 0)	Visit 2 (Day 7±2 Days)	Visit 3 (Day 14±2 Days)	Visit 4 (Day 21±2 Days)	Visit 5 (Day 28±2 Days)	Final Visit (Day 70±7 Days)
Informed Consent Form	X						
Demographics	X						X
DXA	X						X
Medical History	X						
General Physical Examination	X	X	X	X	X		X
Current Medications	X	X					
Blood Chemistries	X	X	X	X	X	X	X
Hematology	X	X	X	X	X	X	X
Urine Analysis	X	X	X	X	X	X	
Vital Signs	X	X	X	X	X	X	X
Inclusion/Exclusion Criteria	X	X					
Enrollment/Randomization		X					X
Treatment Administration Form		X	X	X	X	X	X
Concomitant Medications		X	X	X	X	X	
Adverse Events		X	X	X	X	X	X

Table 5: Non-SoE Table Example 1

Evaluation	Diagnosis	Months Following the Completion of Therapy			
		3	9	24	48
Physical measurements	X		X	X	X
IGF-1	X		X	X	
TSH	X	X	X	X	
Morning Cortisol (7AM-9AM)	X			X	
Systolic BP	X	X	X	X	X
Serum Sodium	X		X	X	
HbA1c	X		X	X	X
Serum Calcium	X			X	

Table 6: Non-SoE Table Example 2

Assessment or Procedure	Dose	Day	Time	Time Window	Pharmacokinetics	Immunogenicity
Cycle 1	Dose 1	Day 2	Pre-dose		X	X
	Dose 2	Day 9	Pre-dose	±2 hours	X	
	Dose 3	Day 16	Pre-dose			X
	Dose 4	Day 25	Pre-dose	±8 hours	X	
	Dose 5	Day 30	Pre-dose			
Final Assessment		End of Tx				
Follow-up Review		Day 30	Post 5 weeks		X	X

Appendix B. Prompts

B.1. Prompt for JSON based inference

We use the following prompt for inference with JSON representation of the table:

A schedule of events or activities (SoE or SoA) table in a clinical trial protocol specifies a plan of care for participants.

Here are some characteristics of SoE/SoA tables:

1. The header rows specify the name and timing of a series of visits to the research site where the participants receive some assessments or treatments.
2. The visits are usually arranged in three phases: screening visits, treatment visits and follow-up visits. A typical SoE or SoA table includes the visits of ALL the three phases or periods.
3. Body rows indicating the occurrence of an assessment or treatment during specific visits, often denoted by symbols like 'X', '✓', or '●'. Some cells may have additional textual specifications. Key terms often found in an SoE or SoA table include: "Informed Consent", "randomization", "treatment", "protocols", and "timing of visit". If you find these keywords (especially "Informed Consent"), this indicates an SoE or SoA table.

Following tables are NOT SoE or SoA tables:

1. An SoE or SoA table is NOT a table describing the timepoints when a specific assessment should be performed, such as a table specific to laboratory assessments, pharmacokinetic collections, or pharmacodynamic collections. These will often break down an assessment into hourly collections after an occurrence, like a pharmacokinetic collection that is performed many times on a single day in relation to treatment administration (0h post-dose, 2h post-dose, 6h post-dose, and so on). These are supplemental tables that greatly expand upon an abbreviated description in the SoE, but are NOT an SoE table.
2. An SoE or SoA table is NOT a document history table listing all previous protocol versions that have been amended and a summary of their changes.
3. An SoE or SoA table is NOT an objectives table, describing the research and statistical goals of the research study (also endpoints, outcomes, etc.)
4. An SoE or SoA table is NOT a table describing adequate organ function or laboratory values
5. An SoE or SoA table is NOT a table describing dose modifications and toxicity in regards to the research treatment

Given the input as a table in the JSON format, return YES if it is an SoE or SoA table or return NO if it is not. Do not output anything else.

Input table in JSON format:

```
{table}
```

Your answer (One of YES or NO):

B.2. Prompt for text based inference

The text based model inference for table classification is done with the following prompt:

A schedule of events or activities (SoE or SoA) table in a clinical trial protocol specifies a plan of care for participants. Here are some characteristics of SoE/SoA tables:

1. The header rows specify the name and timing of a series of visits to the research site where the participants receive some assessments or treatments.
2. The visits are usually arranged in three phases: screening visits, treatment visits and follow-up visits. A typical SoE or SoA table includes the visits of ALL the three phases or periods.
3. Body rows indicating the occurrence of an assessment or treatment during specific visits, often denoted by symbols like 'X', '✓', or '●'. Some cells may have additional textual specifications.

Key terms often found in an SoE or SoA table include: "Informed Consent", "randomization", "treatment", "protocols", and "timing of visit".

If you find these keywords (especially "Informed Consent"), this indicates an SoE or SoA table.

Following tables are NOT SoE or SoA tables:

1. An SoE or SoA table is NOT a table describing the timepoints when a specific assessment should be performed, such as a table specific to laboratory assessments, pharmacokinetic collections, or pharmacodynamic collections. These will often break down an assessment into hourly collections after an occurrence, like a pharmacokinetic collection that is performed many times on a single day in relation to treatment administration (0h post-dose, 2h post-dose, 6h post-dose, and so on). These are supplemental tables that greatly expand upon an abbreviated description in the SoE, but are NOT an SoE table.
2. An SoE or SoA table is NOT a document history table listing all previous protocol versions that have been amended and a summary of their changes.
3. An SoE or SoA table is NOT an objectives table, describing the research and statistical goals of the research study (also endpoints, outcomes, etc.)
4. An SoE or SoA table is NOT a table describing adequate organ function or laboratory values
5. An SoE or SoA table is NOT a table describing dose modifications and toxicity in regards to the research treatment

One way to identify an SoE or SoA table is to look at the text outside the table. Specifically, look for terms like Schedule of Events, Schedule of Assessment, Schedule of Activities, Study Calendar, Study Parameters, Study Schedule and related terms.

If you see any of these terms or related terms in the text data you can

conclude that it indicates an SoE or SoA table.

If you don't see any of these terms in the text data you should look at the whole text data to determine if it is an SoE or SoA table.

Your goal is to determine if the provide text data is from an SoE or SoA table or not. The text data includes all the text before, inside and after the table.

Return YES if it is an SoE or SoA table or return NO if it is not.

Do not output anything else.

Text Data (including before, inside and after the table):

{text}

Your answer (YES or NO):

Appendix C. Naive Combination of Gemini-pro and PaLM-2 models

To broaden our baseline comparisons, we experimented with a naive ensemble approach by combining the outputs of the Gemini-pro 1.0 and PaLM-2 models. This exploratory analysis aimed to assess whether a naive combination of model inferences could leverage the strengths of both individual models to improve the detection of Schedule-of-Event (SoE) tables.

Our ensemble strategy entailed aggregating predictions from both models, each producing two sets of inferences for the tables in clinical trial protocols based on JSON and text representations. We established varying thresholds—from a minimum of one to a maximum of four affirmative (“YES”) inferences—to determine when a table should be classified as a SoE. The performance metrics of the naive ensemble models, detailed in Table 7, indicate that the ensemble outperforms the individual models when a threshold of at least two affirmative inferences is applied. This specific threshold represents a balance, capturing the consensus across the models while mitigating the impact of any one model’s false positives or negatives. Nonetheless, the performance of naive ensemble approaches remained inferior to the fine-tuned models (PaLM-2 fine-tuned with human labels or Gemini annotated and consensus-filtered labels) at all thresholds underscoring the value of fine-tuning over simple ensemble methods in this context. The results of our naive ensemble models show that while aggregation techniques can yield benefits, they are outperformed by a more sophisticated method of fine-tuning models with carefully curated labels.

Table 7: Performance of models when using various thresholds for classifying as SoE Tables

Model	Recall	Precision	F-1 Score	Accuracy
PaLM-2-Gemini Naive Ensemble-1 (SoE if at least one inference is SoE)	100%	51.5%	0.65	83.6%
PaLM-2-Gemini Naive Ensemble-2 (SoE if ≥ 2 inferences are SoE)	99.5%	75.6%	0.83	92.6%
PaLM-2-Gemini Naive Ensemble-3 (SoE if ≥ 3 inferences are SoE)	94.9%	83.6%	0.86	94.9%
PaLM-2-Gemini Naive Ensemble-4 (SoE if all inferences are SoE)	87.2%	86.2%	0.85	95.1%

Appendix D. Additional Recall and Precision Metrics

Table 8: Additional Recall and Precision Metrics

Model	% of protocol with >60% precision	% of protocol with >80% precision	% of protocol with 100% precision	% of protocol with 100% recall
PaLM-2	44.0	22.0	14.3	93.4
GPT-4 (gpt-4-0613)	75.8	56.0	42.9	95.6
Gemini Pro 1.0	56.0	33.0	22.0	98.9
Fine-tuned PaLM (Using Human Labels)	85.7	71.4	68.1	97.8
Fine-tuned PaLM-2 (using ALL Gemini Labels)	53.8	28.6	19.8	100.0
Fine-tuned PaLM-2 (Using Filtered Gemini Labels)	82.4	69.2	64.8	95.0