

Information Extraction from Heterogenous Documents without Ground Truth Labels using Synthetic Label Generation and Knowledge Distillation

Aniket Bhattacharyya, Anurag Tripathi
Amazon

anikettb@amazon.com, tripaanu@amazon.com

Abstract

*Invoices and receipts submitted by employees are visually rich documents (VRDs) with textual, visual and layout information. To protect against the risk of fraud and abuse, it is crucial for organizations to efficiently extract desired information from submitted receipts. This helps in the assessment of key factors such as appropriateness of the expense claim, adherence to spending and transaction policies, the validity of the receipt, as well as downstream anomaly detection at various levels. These documents are heterogenous, with multiple formats and languages, uploaded with different image qualities, and often do not contain ground truth labels for the efficient training of models. In this paper we propose **Task Aware Instruction-based Labelling (TAIL)**, a method for synthetic label generation in VRD corpuses without labels, and fine-tune a multimodal Visually Rich Document Understanding Model (VRDU) on TAIL labels using response-based knowledge distillation without using the teacher model’s weights or training dataset to conditionally generate annotations in the appropriate format. Using a benchmark external dataset where ground truth labels are available, we demonstrate conditions under which our approach performs at par with Claude 3 Sonnet through empirical studies. We then show that the resulting model performs at par or better on the internal expense documents of a large multinational organization than state-of-the-art LMM (large multimodal model) Claude 3 Sonnet while being 85% less costly and $\sim 5X$ faster, and outperforms layout-aware baselines by more than 10% in Average Normalized Levenshtein Similarity (ANLS) scores due to its ability to reason and extract information from rare formats. Finally, we illustrate the usage of our approach in overpayment prevention.*

1. Introduction

Visually-rich documents (VRDs) such as invoices and receipts submitted by employees for reimbursement contain

a rich mixture of text, images, and layout information. For organizations dealing with such documents, it becomes necessary to properly extract key pieces of information such as vendor name, expense value etc from these documents to protect against abuse, fraud and wastage. Information extraction helps in assessing key factors such as validity of expense, duplicate detection, adherence with spending and transaction policies, and downstream usage such as anomaly/fraud detection through the analysis of aggregated behaviour. These documents are typically **heterogeneous**, in **multiple different formats and languages**, and **may not have ground truth labels** for queries of interest, which means models can not be trained to extract them.

Large multinational organizations receive expense reports in the millions annually with supporting invoices/receipts from a large number of vendors in dozens of languages each year. Due to the large number of expense claims, it is not possible to manually verify that the receipts submitted against the expenses are appropriate, i.e. are not duplicates of an older document, or from an unrelated transaction submitted maliciously or erroneously, etc. Employees may be manually required to submit some information, just as justification, merchant name, invoice amount, expense date, etc. However, the value in these fields do not always match the text present in the documents. Fig. 1 shows two receipts submitted by employees of a large multinational organization along with the information present on those documents in expense data. As can be seen, the invoice on the left says Heathrow Airport, while the one on the right says austostrade, whereas the employees have typed in Heathrow T5 and ASPI COLLEFERRO as merchant names respectively. We refer to these fields (typed in by employees that may not match the documents) as **noisy labels**. It may also be of interest to extract information about fields that are not typed in by employees. For such fields, even noisy labels are not available, making it difficult to design solutions to extract them.

Optical Character Recognition (OCR) is used to extract words from VRDs. Advanced OCR tools such as Amazon Textract and Tesseract are capable of returning words,



merchant_name: Heathrow T5,
amount: 6.67,
Transaction_date: 20191112



merchant_name: ASPI COLLEFFERRO,
amount: 2.6,
Transaction_date: 20200721

Figure 1. Two receipts where the manually typed merchant name does not match the name on the image

lines, key-value pairs, tables etc using geometry of recognized text ([19]). However, due to the large number of possible formats of VRDs along with multiple languages, standalone OCRs are not reliable mechanisms to extract information from them consistently. Reasoning is required to understand which part of the document corresponds to the query that is being considered. Contextual clues, such as understanding that a word represents the name of a vendor, or that CGST, VAT, and Umsatzsteuer all represent tax, or that the date mentioned at a certain area of the receipt represents transaction date even if its not labelled as the transaction date, can only be used efficiently by systems that possess the ability to reason.

NLP approaches capable of processing document text and structures have been used to solve these problems. Deep learning approaches such as RNN [1, 43], CNN [14, 18, 27] and transformers [34, 54] have been leveraged in recent years. These techniques focus on mining information such as structural components of documents. They require extensive labelling of these components, which may not be available in expense receipts, and additionally do not consider the layout of documents. Layout-aware methods such as [23, 45, 56] use layout information through mechanisms such as bounding box-embeddings to extract information from documents. However, they need to be pre-trained in tasks involving documents with unseen formats, which is challenging in expense processing cases where labels may not exist.

The recent advancements in the field of large multimodal general-purpose models have provided a new way to approach the problem. Commercial LMMs such as Claude [4] and ChatGPT [41] possess the ability to reason and process the query, the document, and provide responses in the desired format. However, these are closed source and are only available through expensive API calls, and even when instruction-tuned, using them in production environments is risky and expensive. Open-source LMM architectures ex-

ist which can be fine-tuned and hosted in internal environments. Models such as LLaVa [31] and CogVLM [52] etc have shown promising performance in question answering tasks such as SQA, VQAv2. etc. However, these are tasks that involve general descriptions of images or text instead of extraction of particular information and outputting in a precise format. Their zero-shot capabilities are poor, and accurate ground labels are often not available for invoices and receipts for fine tuning.

In this paper, we consider the problem of visually rich document understanding in the absence of labels, and propose a novel solution to leverage the advancements made in large general purpose multimodal models without expensive API calls. We generate synthetic soft labels using Claude 3 Sonnet, a state-of-the-art large multimodal model (LMM), and train a smaller multimodal model on the synthetic labels to learn from the larger model in a teacher-student architecture. We show that the multimodal model thus trained performs at par with or slightly better than Sonnet under certain conditions, and outperforms other recent solutions such as layout-aware methods. We then show how this model can be used to detect potential abuse, wastage or fraud, by focusing on one particular use case. To summarize, this paper makes the following **contributions**:

- We introduce **Task Aware Instruction-based Labelling (TAIL)**, a mechanism to generate synthetic ground truth labels for training models on VRDs without labels. We train a multimodal model **LlaVA-Net**, a VRDU model capable of extracting information in the desired format from expense documents through response-based knowledge distillation without using the teacher model’s weights or training dataset.
- We perform empirical analysis on a benchmark dataset [44] where ground truth labels are available to demonstrate conditions under which our approach performs near-par with Claude 3 Sonnet on VRD corpuses.
- We show that LlaVA-Net achieves performances comparable with instruction-tuned state-of-the-art models like Claude 3 Sonnet while reducing cost by 85% and being $\sim 5X$ faster, and outperforms layout-aware methods by up to 10% in ANLS scores in the internal expense dataset of a large multinational organization where only noisy labels are present.
- We demonstrate the usage of our approach to detect potential abuse/fraud/wastage through the overpayment detection using mismatch between noisy labels and model output.

While in this paper we explore the problems of VRDU in the context of expense documents, the learnings from this paper can be leveraged broadly in structured document understanding problems such as contract and legal document

analysis, medical document processing for purposes such as insurance, document processing in banks for automated KYC/underwriting, caption generation, digitization, etc.

2. Related Work

2.1. VRD understanding

Previous work in VRD understanding can be categorized into three groups - traditional techniques that use text and vision information, layout-aware models that use visual and layout information through learnable layers on top of NLP, and instruction-tuned and/or fine-tuned large multimodal models capable of reasoning and formatting output. **Early traditional methods** use rule-based models, or models that learn from handcrafted features [12, 17, 35, 38, 40, 49]. These techniques assume form templates are known a priori, or rely on loose rules, and do not generalize to new document templates. **Deep learning** is leveraged to get more efficient information extraction algorithms. [1] and [43] used RNNs to extract structural information about documents, such as tables and text fields. Other deep learning frameworks such as CNN [14, 18, 27] and transformers [34, 54] have also been explored. These techniques focus on mining information such as structural components of documents. They require extensive labelling of these components, which may not be available for expense receipts.

Layout-aware NLP approaches have been used in recent years to solve these problems. Transformers have been used by [6, 7, 16, 24]. Other approaches use language models [8, 15, 33] with additional learnable modules on top that capture layout of documents. [56] uses 2D position embeddings to incorporate bounding box information of detected words. [23, 45, 57] enhance the usage of position embeddings through a layout-aware attention mechanism. [29] uses a graph contrastive objective to maximize the agreement of text and image representations. These methods pre-train models on extensive document data such as IIT-CDIP [30], and requires further finetuning for document question answering tasks such as DocVQA [37] and FUNSD [25], which have labelled responses. This is a challenge as expense documents have noisy labels at best.

Large multimodal models (LMMs) have been applied to VRDU as well, recently. Commercial LMMs such as Claude [4] and ChatGPT [41] possess the ability to reason and process the query, the document, and provide responses in the desired format zero-shot, without further fine-tuning. The Claude 3 family of models has been shown in [3] to achieve state-of-the-art performances in DocVQA. Open-source LMM architectures have also been developed and tested on question answering tasks. Models such as LLaVa ([31]) and CogVLM [52] etc have shown promising performance in question answering tasks such as SQA, VQA_{v2}. etc., when fine-tuned. However, their zero-shot capabilities

are poor, and accurate ground labels are not available for invoices and receipts for fine tuning. On the other hand, while commercial LMMs show promising zero-shot performance, they require expensive API calls and have data privacy concerns.

2.2. Knowledge distillation

Response-based knowledge distillation was first proposed by [20] to compress the knowledge present in ensemble of large neural networks to a single model using the output of the teacher model as soft labels. There have been several variations on knowledge distillation proposed since, typically requiring the original data that was used to train the teacher model [13, 48, 53]. Data-free methods have been proposed since [22, 58]. However, any distillation method requiring information about the parameters or weights or training dataset of the teacher model can not be applied on closed-source LMMs. Additionally, while knowledge distillation has been studied for models processing images, expense documents are a special type of images where reasoning and layout-awareness plays a key role along with vision.

In this paper, we present **TAIL**, a method to synthetically generate labels for visually rich documents using a large state-of-the-art multimodal model **without requiring the training dataset or weights of the teacher model**. We then train a smaller student multimodal model, that we call **LLaVA-Net** on the labels generated, to teach it to extract fields of interest from expense documents through response-based knowledge distillation. We show that it is able to show at-par performance compared to the teacher model while being less costly and faster, and outperforms recent layout-aware models.

3. Problem Statement

Expense receipts and invoices are visually rich documents. Visually rich document understanding (VRDU) is a general task, and in Sec. 2 we discussed VRDU methods and datasets. In this section, we set up the VRDU task using expense document-specific vocabulary, and this can be easily extended to all structured document processing problems.

Given an expense document $E_i, i \in (1, 2, \dots, N)$, one may be interested in extracting different attributes E_{ij}^k (vendor name, amount, product purchased, location, etc) with values E_{ij}^v , in specific formats. Let M denote the total number of attributes of interest, and V_j denote all possible values for the attribute E_{ij}^k . Let $\mathbb{V} = (V_1 \times V_2 \dots \times V_M)$.

Let the entire corpus of expense documents be characterised by the set \mathbb{E} . Then, the goal of VRDU is to define a mapping $f : \mathbb{E} \rightarrow \mathbb{V}$ such that

$$f(E_i) = (E_{i1}^v, E_{i2}^v, \dots, E_{iM}^v), \forall i \in (1, 2, \dots, N). \quad (1)$$

In this paper, we examine two different cases of this problem.

Problem 1: Internal ExpenseQA dataset. This consists of internal receipts and expense reports from a multinational organization, with noisy labels as discussed earlier. To create a benchmark dataset on which different approaches could be assessed, we chose fields for which noisy labels were present. While noisy labels aren't exact, they are the closest fields available for ground truth. In particular, the problem statement formulated was as follows: given an expense document, extract the three following pieces of information - **merchant name**, **invoice/receipt amount**, and **transaction date**. Here, $M=1$. Additional details are provided in Sec. 5.

Problem 2: CORD dataset. The CORD dataset [44] consists of 1000 Indonesian receipts with differing image quality and formats. Crowdsourced exact labels are available for each CORD receipt for 20 + attributes. We use CORD dataset to understand 1) how LLaVA-Net performs on benchmark datasets, and conditions of optimal performance, 2) how training on TAIL labels affects performance with compared to training on actual labels, and 3) impact of image quality on annotation. We report results for $M=1$ (total amount) and $M=24$ setup (we leave some sparsely labelled attributes).

4. Proposed Solution

4.1. Task Aware Instruction-based Labelling - TAIL

Historically, the absence of labels has been tackled by crowdsourcing labels. This is a time-consuming and expensive task. Additionally, expense documents are sensitive containing confidential information such as employee name, which can't be shared with other parties. Instruction Tuned large multimodal models have recently shown promising performance in zero-shot learning on various visual question answering and document understanding tasks.

In [3], the Claude family of models, particularly Claude Sonnet, shows the most consistent performance across the three tasks DocVQA [37], A12D [28] and ChartQA [36]. Additionally, it is available on Amazon Bedrock, where Anthropic has set up an escrow account without outbound permission [2]. This takes care of the problem of data privacy as well. We chose Sonnet for synthetic label generation.

As discussed earlier, the internal expense receipts do not contain comprehensive labels. To control the output and to provide us with some benchmark for the assessment of our work, we chose formats present in the noisy labels. The three tasks chosen were - extraction of **merchant name**, extraction of **invoice/receipt amount**, extraction of **transaction date**. For CORD dataset, crowdsourced labels are already available for 20+ attributes in each receipt. We tested our approach on one attribute ($M=1$, E_{i1}^k = total amount) as

well as 24 attributes, with the same format as in the crowd-sourced labels.

To optimize zero-shot learning, prompts need to be efficient. General prompting of the type "What is the invoice amount?" led to responses such as "The invoice amount shown in the image is \$10.46", which does not match the format of the noisy labels or of the images. Anthropic guidelines [5] suggest the usage of examples, controlled output formats, and roles for Claude. Works such as [51] show that task-specific prompting with roles and examples lead to significant improvements in models such as Claude and GPT. Labels generated were investigated for **adherence to requirements**, matches with **noisy/exact labels**, and with the actual **expense documents**. Mismatches with noisy labels and images were investigated to understand how to get better responses. The final prompts chosen were **task-aware instructions**, i.e. contained specific instructions on what to return, and were customized for each specific task. The finalized prompt template, along with the function of each line for $M=1$, is shown in Tab. 1). The formats chosen for the internal dataset - were as few words as possible for merchant name, invoice local currency value without currency sign (such as 45 for \$45) for amount, and `yyyymmdd` for date. For CORD dataset with $M=1$, the format chosen was the total price as mentioned exactly in the receipt.

For CORD dataset with $M=24$, the format chosen for output was JSON, and examples were given for particular fields instead of the entire JSON. This was done to keep prompts succinct, as long prompts have been shown to lead to LLMs ignoring instructions towards the middle of prompts [32]. We changed the names of the receipt attributes to self-explanatory keys (such as name instead of `nm`, `line_item` instead of `menu`, etc), as it has been shown that this leads to better performance [46].

Actual prompts for CORD dataset are shown in the supplementary material. We do not share the exact prompts for the internal dataset to maintain confidentiality.

The labels thus generated were found to adhere to required format, and had an ANLS [9] of 50%, 70% and 76% respectively for the merchant name, amount and date queries with the noisy labels in the internal dataset. On CORD dataset the labels generated with TAIL prompts under $M=1$ setup showed an ANLS of 94% with crowd-sourced labels, and a tree-edit distance and TED of 93% under the $M=24$ setup. These metrics are explained in Sec. 5.3.

Note that additional good practices exist to improve performance in LMMs. In-context learning (ICL) [10] has been shown to lead to more relevant generation, especially when some similarity metric is used to include the most

Function	Prompt Template
Role/Task	You are a seasoned invoice transcriber. Given an invoice image, you return
Role/Task	<i>{value_placeholder}</i>
Example1	For example, if the <i>{value_placeholder}</i> is <i>{example1}</i> , you return <i>{example1_formatted}</i> .
Example2	If the image is a starbucks invoice worth \$45 and transaction was on 5th April 2023, you return
Example2	<i>{business_example_value}</i> .
MissingExample	If you are not confident about <i>{value_placeholder}</i> ,
MissingExample	return None. For example if in the same invoice you cant read the <i>{value_placeholder}</i> ,
MissingExample	return None.
Rules	Only the <i>{value_placeholder}</i> , no extra words.
Request	Return <i>{value_placeholder}</i> for this invoice
image data	<i>image</i>

Table 1. TAIL prompt templates, M=1

similar examples in the prompt. Chain-of-thought (COT) [55] can be used in final output as well, to get better conditional annotations in the presence of multiple likely candidates. OCR output can be leveraged inside prompts to augment an LMMs vision capabilities. However, all of these approaches lead to increased latency as well as cost, by increasing the number of tokens inputted and outputted, and by including an OCR layer. In this paper, we show how open-source LMMs can be trained to perform at-par with large state-of-the-art LMMs while decreasing latency (5X faster) and cost (by upto 85%), as cost and latency were both important given the huge volume of expense receipts present without labels. Additional good practices such as ICL, COT, advanced OCR output etc can be then leveraged by practitioners to further improve accuracy of annotations.

4.2. Student model for Knowledge Distillation: LLaVA-Net

To effectively learn from TAIL labels, student models need to possess vision, reasoning, and question-answering skills. For reasoning and question answering, open source foundational language model Vicuna [11] uses pre-trained Llama weights [50] and fine-tunes on GPT4 conversation to achieve better instruction following capabilities.

LLaVA ([31]) combines Vicuna with CLIP, a vision encoder capable of state-of-the-art zero-shot textual image representation by training models on (image, caption) samples [47]. The two modules are combined through a learnable projection matrix, and further trained on instruction following datasets. LLaVA also demonstrated multi-lingual capabilities, which are needed on the ExpenseQA task. Later versions of LLaVa also provided integration with Mistral instead of Vicuna, which has shown superior performance than Llama 13B in various benchmarks [26]. Since Vicuna is based on Llama, we chose the pre-trained weights with Mistral instead of Vicuna. We took the pre-trained LLaVA weights and fine-tuned the text component on TAIL labels using LoRA [21]. LoRA allows us to finetune the

7B parameter text component using 8 GPUs on SageMaker in 10 hours. The final model trained on TAIL labels using Sonnet is called **LLaVA-Net**.

5. Dataset and Experimental Setup

5.1. ExpenseQA Dataset Creation

TAIL prompts for all three tasks were used on a sample of 11000 expense documents. The final dataset created contained both noisy and TAIL labels for all the tasks along with the tail prompts. This was divided into a training sample of 10000 rows and a validation sample of 1000. This dataset was used to further train LLaVA-Net, and to evaluate LLaVA-Net against benchmarks. Tab. 2 shows the TAIL labels along with the noisy labels for the two receipts shown in Fig. 1. Notably, the TAIL labels match the text present in the images. While this example focuses on merchant name mismatches, other receipts in the dataset contained mismatches for the two other fields as well.

5.2. CORD Dataset

CORD [44] contains 1000 Indonesian receipts, divided into train, validation and test samples of size 800,100 and 100. Along with the images, CORD also contains crowd-sourced labels, and OCR output with bounding boxes. However, to simulate expense data without labels, we did not use the OCR output, relying instead on the vision capabilities of models under consideration, and TAIL labels for training. We used the actual labels only for testing, to understand the efficacy of our approach.

5.3. Experimental Setup and Results

We investigated five research questions to assess the performance and appropriateness of TAIL and LLaVA-Net in the ExpenseQA and CORD datasets.

In **research question 1**, We look at the performance of LLaVA-Net on both ExpenseQA and CORD, and compare it with Sonnet as well as layout-aware benchmarks.

Receipt	merchant name noisy	merchant name TAIL	amount noisy	amount TAIL	transaction date noisy	transaction date TAIL
1	Airport T5	Heathrow Airport	6.67	6.67	20191112	20191112
2	meal Nampa	Holy Cow!	2.6	2.60	20200721	20200721

Table 2. Noisy and TAIL labels for invoices in Fig. 1

In **research question 2**, we focus on the performance discrepancies between layout-aware methods and LLaVA-Net, and analyse why LLaVA-Net is more effective.

In **research question 3**, we look at the impact of image blur on the performance discrepancies between layout-aware methods, LLaVA-Net and Sonnet, and demonstrate conditions under which LLaVA-Net performs near-par with Sonnet.

In **research question 4**, we investigate whether using synthetic TAIL labels biases LLaVA-Net, and makes it prone to repeating the errors of Claude Sonnet.

In **research question 5**, we illustrate how LLaVA-Net output can be used to prevent potential fraud/abuse/wastage through the detection of overpayment risk.

5.3.1 RQ 1: Performance comparison.

Here, we look at the performance of LLaVA-Net, trained on TAIL labels generated using Sonnet on both ExpenseQA and CORD. We also compare it with the performance of LLaVA - zero shot with TAIL prompt, as well as LLaVA finetuned on noisy labels (exact labels for CORD) instead of TAIL labels, to understand the lift provided by TAIL labels. We also compare its performance with the performance of layout-aware methods which are state-of-the-art in information extraction. As discussed in the related works section, layout-aware methods need rigorous labelling for fine-tuning, which are not available in ExpenseQA. However, we have generated synthetic TAIL labels, and we used those to fine-tune LayoutLMV3 which has achieved SOTA in question-answering and document processing tasks amongst peers [23]. We did not use LayoutLMV3 on transaction date because TAIL labels for date were designed to match noisy labels instead of the date formats in invoices. We also look at the TAIL labels themselves, and compare those with the noisy labels.

For CORD, we look at performance under both M=1 and M=24 setups. We compare LayoutLMV3 with Sonnet and LLaVA-Net for M=1, trained on TAIL labels to simulate situations where labels are not available. We used LayoutLMV3’s feature extractor which uses Tesseract for training, and actual labels for testing the true performances to simulate real life. Note that while newer, more promising OCR solutions such as DocTr [39] are available, these do not always support non-Roman alphabet languages or handwritten text, crucial for industrial application.

Assessment metric We used **ANLS scores** (average normalized Levenshtein similarity score) proposed by [9] between the model being evaluated and the noisy labels (exact labels for CORD) as the metric of assessment on the validation dataset, under the M=1 setup. Under the M=24 setup, we use **tree-edit distance** [42] on the JSON output. We also compare the speeds of inference and costs. The cost of approaches are measured annually, assuming 100X documents (we withhold the exact numbers for confidentiality), Sonnet API call prices are taken as of April 2024, and the cost of TAIL labelling for 10K samples for LLaVA-Net and layoutLMV3 are included. The speed of inference is measured as documents processed per second for M=1. The results are shown in Tab. 3.

Sonnet zero-shot outputs (TAIL labels) showed ANLS scores of 50%, 70% and 76% with the noisy labels. LLaVA zero-shot capability was poor compared to Sonnet. LLaVA trained on noisy labels performs significantly better than LLaVA zero-shot, learning the required output format. However, as the noisy labels do not always match with the text in the images, the vision component of LLaVA doesn’t find the label in the image in many cases and this hampers fine-tuning. It performs worse than LLaVA-Net, showing the **lift due to TAIL labels**. LLaVA-Net is able to show near-par performances with Sonnet on the internal dataset, and performs slightly better on the merchant name and amount questions, while being $\sim 85\%$ cheaper and 5X faster than Sonnet. A surprising result was seen in CORD, where training on exact labels showed slightly worse performance than on TAIL labels, and in general LLaVA-Net performed poorer than Sonnet. We investigate this further in research question 3.

While LayoutLMV3 trained on TAIL labels is able to learn to detect the positions of merchant names and invoice amounts, it is not able to perform as well as LLaVA-Net, and in research question 2, we investigate this further.

5.3.2 RQ 2: Why does LLaVA-Net outperform layout-aware methods?

Layout-aware methods are pre-trained and fine-tuned using masked language modeling and masked image modeling rather than language-based reasoning. They learn best when the fine-tuning dataset consists of samples with repeated formats. However, to take an example, the detection of the text “Giordano’s Pizza” as merchant name, even if it

Model	Merchant name	Amount	Date	CORD, M=1	CORD, M=24	Inference speed	Annual cost
Sonnet Zero-Shot (TAIL)	50%	70%	76%	94%	93%	0.3	\$1.25X
LLaVA zero shot TAIL prompt	24%	28%	58%	48%	42%	1.5	\$0.13X
LLaVA-noisy labels	46%	63%	68%	-	-	1.5	\$0.13X
LLaVA-exact labels	-	-	-	82%	-	1.5	0.13X
LayoutLMV3-Tail	41%	58%	-	42%	-	0.8	0.2X
LLaVA-Net	52%	70%	83%	84%	69%	1.5	\$0.2X

Table 3. Performance Comparison: Sonnet, LLaVA, LLaVA-Net, LayoutLMV3

is placed in an odd part of the invoice whose format is rare, requires reasoning. Since ExpenseQA consists of a wide variety of vendors and invoice templates, LayoutLMV3 is not able to extract information from rare document templates.

To check this hypothesis, we looked at the performance of LayoutLMV3 on receipts from Uber and Lyft, two common vendors found in the dataset. LayoutLMV3 performed on par with Sonnet and LLaVA-Net on these vendors, as shown in Tab. 4. This indicates that layout-aware models are able to effectively extract information from documents from common templates but fail in heterogenous corpuses with rare templates. **The multimodal nature and reasoning capability** of LLaVA-Net allows it to use both templates and logic to extract information from expense documents.

LayoutLMV3 is also impacted by the quality of images, and of OCR. Since we used LayoutLMV3 assuming scenarios where labels are not available, in CORD dataset we used its tesseract-based feature extractor. This failed on poor-quality images leading to an ANLS of 42%. However, on images where OCR successfully detected text in this window, the ANLS jumped to 72%, showing the high dependence layout-aware methods have on underlying OCR used, as well as the superiority of Sonnet in vision-based tasks. Note that LayoutLMV3 performs well when provided OCR output and exact labels are used, showing micro-F1 of 96% across all attributes in CORD [46].

5.3.3 RQ 3: Image quality impact on performance.

To assess image quality, we used the variance of the Laplacian operator applied to an image. Intuitively, high-variance means sharp contrast between regions of high and low intensity, indicating low blur. For LayoutLMV3, we considered the general scenario, as well as receipts where its feature extractor successfully extracted text in its window.

As can be seen in Tab. 5, as focus improves, LLaVA-Net starts showing better results faster than LayoutLMV3, and at ranges with variances ≥ 200 , the ANLS exceeds 90%, which is almost at par with Sonnet. For LayoutLMV3, in general the accuracy is poor. However, when the relevant text is present in its window, the accuracy goes up to 100% for high-focus images.



Figure 2. LLaVA-Net correctly extracts amount from an invoice in portrait orientation, while Sonnet fails

This explains why LLaVA-Net performs so well on ExpenseQA, where $\sim 90\%$ images have Laplacian variance ≥ 200 , as opposed to CORD, where the number is $\sim 60\%$, where LayoutLMV3 fails due to heterogeneity.

5.3.4 RQ 4: Does LLaVA-Net replicate the errors made by Sonnet?

The ANLS score between LLaVA-Net and Sonnet on the validation dataset in ExpenseQA for merchant name, amount and transaction date was 87%, 74% and 82%, respectively. In CORD, under M=1 setup the ANLS between LLaVA-Net and Sonnet outputs was 86%, and tree-edit distance under M=24 was 87%. Overall, ANLS and TED are less than 90%, indicating that LLaVA-Net isn't a simple replica of Sonnet. On further analyses of mismatches in names, we found that Sonnet had refused to answer in 3% of samples, while LLaVA-Net had **zero refusals**. This is because LLaVA-Net was fine-tuned to generate text from the conditional distribution of the provided format. LLaVA-Net also correctly extracted merchant names in cases where Sonnet got confused by multiple logos/names in receipts (eg. Uber Eats receipts mentioning the name of restaurants fulfilling the order), or by the receipt being in landscape layout instead of the default portrait (Fig. 2). Sonnet also showed a 1.5% rate of wrongly formatted output, generating text such as "Based on the information provided in the image, it seems like the receipt was generated by...". No such phenomenon was observed in LLaVA-Net, which generated outputs exclusively in the desired format.

Model	Merchant name - Uber	Merchant name - Lyft
Sonnet Zero-Shot (TAIL)	88%	100%
LLaVA-Net	88%	100%
LayoutLMV3-Tail	88%	95%

Table 4. Performance Comparison on common vendors: Sonnet, LLaVA-Net, LayoutLMV3-Tail

Model	200	200 – 500	500 – 750	750 – 1000	≥ 1000
LLaVA-Net	70%	91%	90%	92%	92%
LayoutLMV3-Tail	35%	45%	50%	60%	33%
LayoutLMV3-Text in window	67%	70%	85%	100%	67%

Table 5. Performance Comparison on differing image quality: LLaVA-Net, LayoutLMV3-Tail

Noisy Amount Label: 600,
LLaVA-Net_output: 321.84

Figure 3. Receipt where employee typed in 600 euros as the amount to be reimbursed, and LLaVA-Net detected the correct amount of 321.64 euros

5.3.5 RQ 5: How can LLaVA-Net be used to detect potential fraud/abuse/wastage?

In this section, we show a way in which potential fraud/abuse/wastage can be detected from LLaVA-Net output. 89 invoices were found in the expenseQA dataset where the extracted amount was lower than the amount typed in by employees. This indicates risk of overpayment, as the reimbursement depends on the amount typed in rather than the invoice submitted. Fig. 3 shows an invoice with potential overpayment of ~ 280 euros. Overall, the risk from these invoices was around \$2 per document, which becomes significant with scale.

Additionally, information such as extracted transaction date and location can be cross-checked with available employee location data. Merchant name, product category can be used to check whether expenses align with spending policies. Aggregate spends on products/merchants, frequency of spends etc can be used for downstream anomaly detection.

6. Conclusion

In this paper, we have proposed TAIL, a mechanism to generate synthetic truth labels for VRDs without labels, and LLaVA-Net, a multimodal student model trained using response based knowledge distillation without using the

teacher model’s weights or training dataset. We applied LLaVA-Net to VRDU on internal and external datasets, and compared its performance with other methods that have shown promising results in literature.

Learnings LLaVA-Net performs at par with SOTA Claude 3 Sonnet output under the condition that most of the images in the VRD corpus have decent image quality. If the images are poor, Sonnet’s superior vision capabilities make it outperform LLaVA-Net. LLaVA-Net shows lower incidences of refusal and wrongly-formatted output, since fine-tuning leads it to generate text from the conditional distribution of corpus attributes. It outperforms layout-aware methods on heterogenous document corpus due to its reasoning ability and superior vision capabilities, while layout-aware methods perform well under conditions of high homogeneity and high image quality. We have also shown that TAIL leads to more efficient downstream model training than noisy labels found in expense data. LLaVA-Net, comprising of 7B parameters, can be productionized using a single Sagemaker instance with one 20GB GPU (note that we used 8 GPUs solely for faster training), and is about $\sim 85\%$ cheaper than Sonnet, while being 5X faster. We also illustrated the usage of LLaVA-Net output to detect potential abuse/wastage/fraud through overpayment prevention.

Future Directions As discussed in Sec. 4.1, good practices such as ICL and COT were not used to improve latency and reduce cost, as given the huge volume of unlabelled receipts, both cost and latency were important concerns. In the future, prompting practices such as ICL, COT, automated prompt generation etc can be leveraged on top to further improve performance. We also discussed that if image quality is poor, LLaVA-Net shows suboptimal performance compared to Sonnet, which has superior vision. Super-resolution imaging can be used to improve resolution of images, and OCR output can be passed to the LMM inside prompts to tackle these cases, depending on budget and latency constraints.

References

- [1] Milan Aggarwal, Hires Gupta, Mausoom Sarkar, and Balaji Krishnamurthy. Form2Seq : A framework for higher-order form structure extraction. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu, editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3830–3840, Online, Nov. 2020. Association for Computational Linguistics. 2, 3
- [2] Amazon. Data protection, 2024. 4
- [3] Anthropic. The claude 3 model family: Opus, sonnet, haiku, 2024. 3, 4
- [4] Anthropic. Introducing the next generation of claude, 2024. 2, 3
- [5] Anthropic. Prompt engineering, 2024. 4
- [6] Srikanth Appalaraju, Bhavan Jasani, Bhargava Urala Kota, Yusheng Xie, and R. Manmatha. Docformer: End-to-end transformer for document understanding, 2021. 3
- [7] Haoli Bai, Zhiguang Liu, Xiaojun Meng, Wentao Li, Shuang Liu, Nian Xie, Rongfu Zheng, Liangwei Wang, Lu Hou, Jiansheng Wei, Xin Jiang, and Qun Liu. Wukong-reader: Multi-modal pre-training for fine-grained visual document understanding, 2022. 3
- [8] Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Songhao Piao, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. Unilmv2: Pseudo-masked language models for unified language model pre-training, 2020. 3
- [9] Ali Furkan Biten, Ruben Tito, Andres Mafla, Lluís Gomez, Marçal Rusiñol, Ernest Valveny, C. V. Jawahar, and Dimosthenis Karatzas. Scene text visual question answering, 2019. 4, 6
- [10] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020. 4
- [11] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023. 5
- [12] Laura Chiticariu, Yunyao Li, and Frederick R. Reiss. Rule-based information extraction is dead! long live rule-based information extraction systems! In David Yarowsky, Timothy Baldwin, Anna Korhonen, Karen Livescu, and Steven Bethard, editors, *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 827–832, Seattle, Washington, USA, Oct. 2013. Association for Computational Linguistics. 3
- [13] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1, 2016. 3
- [14] Timo I. Denk and Christian Reisswig. Bertgrid: Contextualized embedding for 2d document representation and understanding, 2019. 2, 3
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2019. 3
- [16] Mohamed Dhoub, Ghassen Beltaieb, and Aymen Shabou. Docparser: End-to-end ocr-free information extraction from visually rich documents, 2023. 3
- [17] Jaekyu Ha, R.M. Haralick, and I.T. Phillips. Recursive x-y cut using bounding boxes of connected components. In *Proceedings of 3rd International Conference on Document Analysis and Recognition*, volume 2, pages 952–955 vol.2, 1995. 3
- [18] Leipeng Hao, Liangcai Gao, Xiaohan Yi, and Zhi Tang. A table detection method for pdf documents based on convolutional neural networks. In *2016 12th IAPR Workshop on Document Analysis Systems (DAS)*, pages 287–292, 2016. 2, 3
- [19] Thomas Hegghammer. Ocr with tesseract, amazon textract, and google document ai: a benchmarking experiment. *Journal of Computational Social Science*, 5, 05 2022. 2
- [20] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015. 3
- [21] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021. 5
- [22] Hengtong Hu, Lingxi Xie, Richang Hong, and Qi Tian. Creating something from nothing: Unsupervised knowledge distillation for cross-modal hashing, 2020. 3
- [23] Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. Layoutlmv3: Pre-training for document ai with unified text and image masking, 2022. 2, 3, 6
- [24] Wonseok Hwang, Jinyeong Yim, Seunghyun Park, Sohee Yang, and Minjoon Seo. Spatial dependency parsing for semi-structured document information extraction, 2021. 3
- [25] Guillaume Jaume, Hazim Kemal Ekenel, and Jean-Philippe Thiran. Funsd: A dataset for form understanding in noisy scanned documents, 2019. 3
- [26] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. Mistral 7b, 2023. 5
- [27] Anoop Raveendra Katti, Christian Reisswig, Cordula Guder, Sebastian Brarda, Steffen Bickel, Johannes H  hne, and Jean Baptiste Faddoul. Chargrid: Towards understanding 2d documents, 2018. 2, 3
- [28] Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. A diagram is worth a dozen images, 2016. 4

- [29] Chen-Yu Lee, Chun-Liang Li, Hao Zhang, Timothy Dozat, Vincent Perot, Guolong Su, Xiang Zhang, Kihyuk Sohn, Nikolai Glushnev, Renshen Wang, Joshua Ainslie, Shang-bang Long, Siyang Qin, Yasuhisa Fujii, Nan Hua, and Tomas Pfister. Formnetv2: Multimodal graph contrastive learning for form document information extraction, 2023. 3
- [30] D. Lewis, G. Agam, S. Argamon, O. Frieder, D. Grossman, and J. Heard. Building a test collection for complex document information processing. In *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '06, page 665–666, New York, NY, USA, 2006. Association for Computing Machinery. 3
- [31] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2023. 2, 3, 5
- [32] Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts, 2023. 4
- [33] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach, 2019. 3
- [34] Bodhisattwa Prasad Majumder, Navneet Potti, Sandeep Tata, James Bradley Wendt, Qi Zhao, and Marc Najork. Representation learning for information extraction from form-like documents. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6495–6504, Online, July 2020. Association for Computational Linguistics. 2, 3
- [35] S. Marinai, M. Gori, and G. Soda. Artificial neural networks for document analysis and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(1):23–35, 2005. 3
- [36] Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning, 2022. 4
- [37] Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. Docvqa: A dataset for vqa on document images, 2021. 3, 4
- [38] Mausam, Michael Schmitz, Stephen Soderland, Robert Bart, and Oren Etzioni. Open language learning for information extraction. In Jun'ichi Tsujii, James Henderson, and Marius Paşca, editors, *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 523–534, Jeju Island, Korea, July 2012. Association for Computational Linguistics. 3
- [39] Mindee. doctr: Document text recognition. <https://github.com/mindee/doctr>, 2021. 6
- [40] L. O’Gorman. The document spectrum for page layout analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1162–1173, 1993. 3
- [41] OpenAI. Gpt-4 is openai’s most advanced system, 2023. 2, 3
- [42] Benjamin Paaßen. Revisiting the tree edit distance and its backtracing: A tutorial, 2022. 6
- [43] Rasmus Berg Palm, Ole Winther, and Florian Laws. Cloudscan - a configuration-free invoice analysis system using recurrent neural networks, 2017. 2, 3
- [44] Seunghyun Park, Seung Shin, Bado Lee, Junyeop Lee, Jaeheung Surh, Minjoon Seo, and Hwalsuk Lee. Cord: A consolidated receipt dataset for post-ocr parsing. 2019. 2, 4, 5
- [45] Qiming Peng, Yinxu Pan, Wenjin Wang, Bin Luo, Zhenyu Zhang, Zhengjie Huang, Teng Hu, Weichong Yin, Yongfeng Chen, Yin Zhang, Shikun Feng, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. Ernie-layout: Layout knowledge enhanced pre-training for visually-rich document understanding, 2022. 2, 3
- [46] Vincent Perot, Kai Kang, Florian Luisier, Guolong Su, Xiaoyu Sun, Ramya Sree Boppana, Zilong Wang, Zifeng Wang, Jiaqi Mu, Hao Zhang, Chen-Yu Lee, and Nan Hua. Lmdx: Language model-based document information extraction and localization, 2024. 4, 7
- [47] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021. 5
- [48] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks, 2016. 3
- [49] A. Simon, J.-C. Pret, and A.P. Johnson. A fast algorithm for bottom-up document layout analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3):273–277, 1997. 3
- [50] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023. 5
- [51] Wenjin Wang, Yunhao Li, Yixin Ou, and Yin Zhang. Layout and task aware instruction prompt for zero-shot document image question answering, 2023. 4
- [52] Weihang Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan

- Song, Jiazheng Xu, Bin Xu, Juanzi Li, Yuxiao Dong, Ming Ding, and Jie Tang. Cogvlm: Visual expert for pretrained language models, 2024. [2](#), [3](#)
- [53] Yunhe Wang, Chang Xu, Shan You, Dacheng Tao, and Chao Xu. Cnnpack: packing convolutional neural networks in the frequency domain. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS'16, page 253–261, Red Hook, NY, USA, 2016. Curran Associates Inc. [3](#)
- [54] Zifeng Wang, Zizhao Zhang, Jacob Devlin, Chen-Yu Lee, Guolong Su, Hao Zhang, Jennifer Dy, Vincent Perot, and Tomas Pfister. Queryform: A simple zero-shot form entity query framework, 2023. [2](#), [3](#)
- [55] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023. [5](#)
- [56] Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. Layoutlm: Pre-training of text and layout for document image understanding. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '20. ACM, Aug. 2020. [2](#), [3](#)
- [57] Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, Min Zhang, and Lidong Zhou. Layoutlmv2: Multi-modal pre-training for visually-rich document understanding, 2022. [3](#)
- [58] Jingwen Ye, Yixin Ji, Xinchao Wang, Xin Gao, and Mingli Song. Data-free knowledge amalgamation via group-stack dual-gan, 2020. [3](#)