IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

AI-Driven Dashboard Insights: KPI Summarization Using Llms

[1]Krutika Shrikant Walzade, [2]Shivam Vilas Gaikwad

[1,2]Student

[1]Information Technology, [2]Computer Engineering,

[1]K. K. Wagh Institute of Engineering Education and Research, Nashik, Maharashtra, India

Abstract: Business dashboards use Key Performance Indicators (KPIs) to give essential insights, but it's still challenging to efficiently extract and analyze large amounts of data. Decision-making is frequently hampered by cognitive overload, which reduces the usefulness of these tools. In order to automatically extract and summarize KPIs from dashboard screenshots, this study presents a Business Dashboard Summarizer that uses Large Language Models (LLMs), computer vision, and optical character recognition. The solution improves accessibility, reduces information overload, and improves business intelligence productivity by turning raw data into concise, natural-language summaries. This study demonstrates how AI-driven summarization can optimize data interpretation and make strategic insights easy to use and more actionable.

Index Terms - Business Dashboard, KPI Extraction, LLMs, AI Summarization, Computer Vision, Business Intelligence, Optical Character Recognition.

I. INTRODUCTION

In today's data-driven world, business dashboards are crucial technologies for tracking key performance indicators (KPIs) and making strategic decisions. These dashboards are used by businesses to monitor efficiency in operations, economic indicators, and consumer involvement in real-time. However, as dashboards become increasingly complex, quickly extracting meaningful insights remains challenging. Decision-makers often experience cognitive overload, struggling to interpret large volumes of data efficiently, which can delay critical business actions. To address this challenge, we propose a Business Dashboard Summarizer, an AI-powered system that automates KPI extraction and generates concise summaries using Large Language Models (LLMs), computer vision, and optical character recognition (OCR). This approach helps companies break down dashboard information more efficiently by bridging the gap between raw data and actionable insights. By integrating technologies such as OpenCV for image processing, Tesseract OCR for text extraction, and transformer-based models for summarization, our approach converts dashboard screenshots into structured, easy-to-understand summaries, reducing information overload and improving decision-making efficiency. The primary research question guiding this study is: "Can AI-driven summarization improve the accessibility and efficiency of business dashboard interpretation?" This study addresses how dashboard analysis might be transformed through LLM-powered summarization. We first discuss the technical aspects of the summarization pipeline, such as automation methods, natural language generation, and KPI extraction. After providing implementation details, we assess how well the system summarizes various business dashboards. Lastly, we discuss the future research prospects and business intelligence implications of AI-driven summarization. Our research aims to enhance business analytics, accelerate decision-making, and optimize dashboard-generated data by automating KPI interpretation, thus improving the accessibility and usability of business intelligence tools across industries.

II. LITERATURE REVIEW

In recent years, there has been a growing interest in the problem of efficiently summarizing visual business data. Some research projects have explored key technologies like computer vision, optical character recognition (OCR), and natural language processing (NLP) to extract and analyze organized and unstructured data from documents and images.

In [1], a technique that combines layout analysis and OCR to extract tabular information from financial papers is presented, showcasing the capacity to convert scanned images into structured representations. In dashboard summarization, this fundamental method is the basis for the extraction process, particularly when tabular KPIs are used. Likewise, in [2], the authors use sophisticated OCR methods to extract meaningful information from medical dashboards, which are frequently overloaded with text and graphic components. The result emphasizes how crucial preprocessing and layout understanding are for precise text recognition in complex interfaces. A fundamental part of our workflow, computer vision, has also been studied for the semantic knowledge of dashboard elements. To identify various chart types and describe their essential components, including axes, legends, and data points, a visual analytics model is created in [3]. This type of component-wise identification is crucial for extracting information about KPIs, which often occur inside visual elements like pie graphs or bar charts. Another study [4] presents a technique for identifying and categorizing widgets in UI screenshots to establish the foundation for automated interface summarization. Automatic summarization has changed due to transformer-based architectures and large language models (LLMs). The usefulness of contextual embeddings in capturing domain-specific language is demonstrated in [5], where a BERT-based summarization model is trained to transform financial reports into brief and understandable summaries. Similarly, [6] highlights the potential of LLMs in story production for business intelligence by presenting a domain-adapted GPT model that creates dashboard narratives from structured datasets. Multi-modal techniques that integrate OCR and NLP for summarizing tasks have also been investigated in several papers.

To summarize the information of academic presentation slides, for example, [7] combines a transformer model with Tesseract OCR. This architecture confirms that OCR results can be used as inputs to LLMs and is quite similar to our pipeline. A pipeline that integrates object identification, text extraction, and abstract summarization is presented in [8] for real-time infographic summarization. Such systems' effectiveness demonstrates the value of hybrid AI systems for enabling access to dense visual data.

Additionally, studies on the human-computer interface have looked at cognitive overload in dashboard interpretation. To facilitate fast decision-making, summary-based interfaces are suggested in [9] after analyzing users' attention patterns across dashboards with different information densities. This user-centric understanding supports our summarizer's goal of lowering cognitive load while increasing analytical efficacy. Despite these developments, some research explicitly addresses end-to-end pipelines for business dashboard summarization using screenshots. The majority of current research focuses on either textual materials [11] or structured information [10], which leaves a gap in systems that translate raw visual dashboard inputs into summaries in natural language. To bridge this gap, our suggested method combines OCR, CV, and LLMs to efficiently analyze and convey visual KPIs.

III. MATERIALS & METHODS

A. Methodology: -

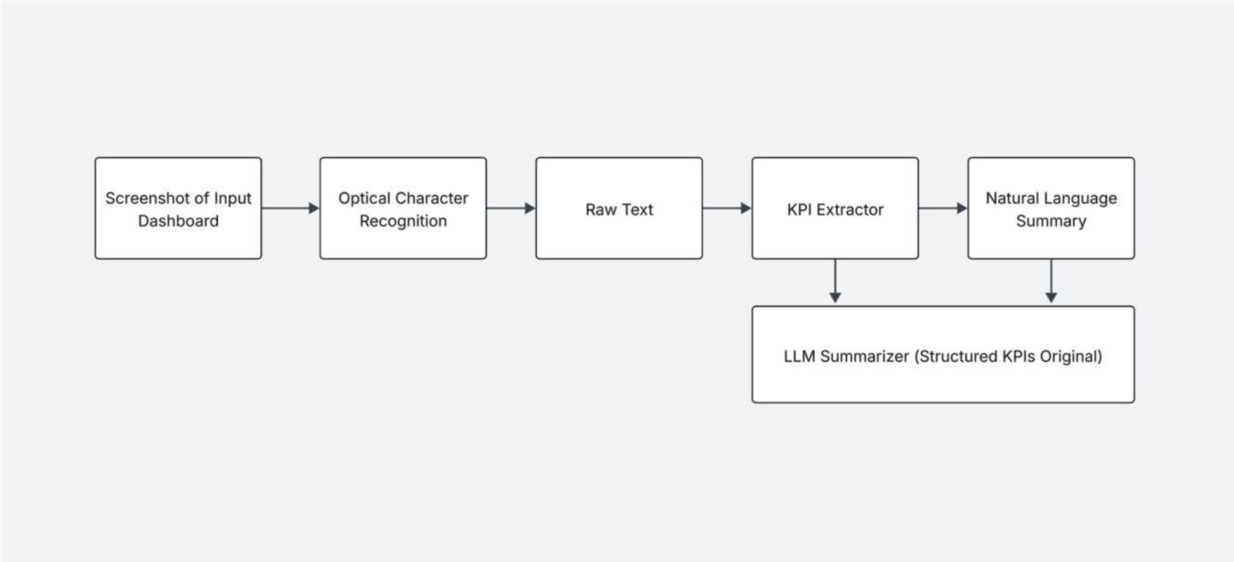

There are five steps in the proposed Business Dashboard Summarizer pipeline. A screenshot of the input dashboard is used as the starting point, and raw text is extracted using optical character recognition (OCR). After the text has been extracted, it is sent to the KPI Extractor, which finds and arranges key performance indicators. A Large Language Model (LLM) Summarizer is fed with structured KPIs and the original dashboard image to produce a logical natural language summary. The system generates a clear overview to enable customers to rapidly understand the key information without manually analyzing complex dashboards

B. Input Dashboard Screenshot: -

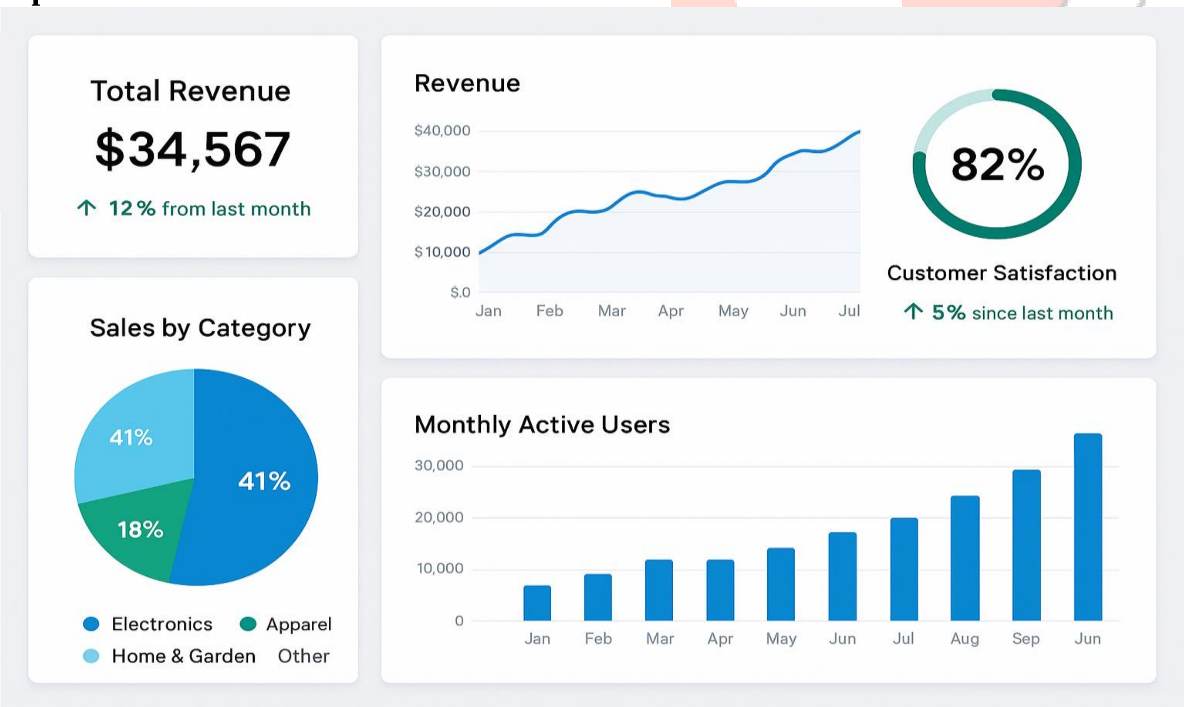

The dashboard snapshot serves as the main input for the summarization system, and obtaining it is the first stage in the pipeline. The user can either manually take the snapshot in our configuration or utilize Python-based tools to capture it automatically.

For automatic screenshot capture, we use the following tools:

- **PyAutoGUI**:
  PyAutoGUI is a Python library that allows automation of mouse and keyboard actions. It can take screenshots of specific windows or the entire screen programmatically. In our system, PyAutoGUI helps automate the extraction of dashboard screenshots at scheduled intervals or on user command, reducing manual effort.

- **Pillow (PIL)**:
  Pillow, a Python Imaging Library, is used alongside PyAutoGUI to process and save screenshots in appropriate formats (e.g., PNG or JPEG) while preserving resolution and quality essential for OCR accuracy.

We can also use Region-Based Screenshot Capture, where users can specify coordinates or automate bounding box detection to record only the dashboard area rather than the entire screen. This enhances OCR performance by lowering noise and unnecessary background data. Preprocessing for quality enhancement: To make sure that the text and visual components are readable and clear, preprocessing techniques like scaling, sharpening, or improving the contrast are used using OpenCV or PIL after the screenshot has been taken. The accuracy of downstream OCR and KPI extraction is directly increased by high-quality input.

C. OCR (Optical Character Recognition):

Optical Character Recognition (OCR) is used in the second step of the summarization pipeline to extract text from the dashboard screenshot. OCR is essential because it makes downstream processing possible by translating the visual representation of key performance indicators (KPIs), chart labels, values, and descriptions into machine-readable text. Tesseract OCR, an open-source OCR engine renowned for its dependability and simplicity of interaction with Python via the pytesseract module, is used in our system. The input dashboard image is initially preprocessed with OpenCV to increase recognition accuracy. Preprocessing techniques include noise reduction techniques like Gaussian blurring, turning the image to grayscale to eliminate unnecessary color information, and employing binarization to improve text clarity by separating the text from complex backgrounds. In order to improve accuracy, we also adjust Tesseract's settings by using the LSTM-based OCR Engine Mode (OEM), setting the Page Segmentation Mode (PSM) to 6, and telling the engine to assume there is a single, uniform block of text. The OCR output enables us to identify and arrange the text globally throughout the dashboard by providing a structured list of identified text pieces, each linked to its bounding box coordinates. Through thorough preprocessing and parameter optimization, issues, including diverse dashboard designs, small word sizes, and mixed visual material are resolved. The extracted text and its positional metadata provide an in-depth visualization of the dashboard's textual data in an organized and machine-readable format, making it an essential input for the KPI extraction step.

D. KPI Extractor:

To find key performance indicators (KPIs) relevant to business analysis, the recovered raw text must be sorted and filtered following the OCR stage. The KPI Extractor module carries out this vital work, which processes the OCR data and organizes information for a practical summary. Finding important KPIs, eliminating extraneous text, and preparing structured data for the Large Language Model (LLM) Summarizer are the KPI Extractor's primary responsibilities. Inaccurate or insufficient summaries would result from the LLM receiving loud or disorganized inputs if this step weren't taken.

1. **Text Filtering and Cleaning**: Labels that, headings, buttons, and unnecessary user interface text are frequently present in the raw OCR output. To ensure that only relevant numerical or performance data is chosen, we filter away non-KPI language using regular expressions, keyword matching, and domain-specific heuristics.

2. **KPI Identification and Classification**:

   Each extracted text block is analyzed to classify it as a potential KPI. This involves:
   - Semantic Matching: Matching text against a predefined set of KPI names or financial/business terms (e.g., "Revenue," "Profit Margin," "Customer Satisfaction").
   - Spatial Relationships: Grouping text blocks based on their proximity on the dashboard (for example, KPI labels and value pairs are often close to each other).
3. **Key-Value Pair Formation:** The extractor pairs KPI labels with their corresponding numerical values. If a text block says "Total Revenue" and another nearby block reads "$5M," these two are combined into a key-value pair: "Total Revenue": "$5M".
4. **Error Handling and Validation:** The system carries out fundamental validation to guarantee clean data, such as verifying that KPI values are numerical or formatted as counts, percentages, or currencies. This reduces the possibility of providing the LLM with inaccurate data.

The output format of the KPI extractor is a structure of JSON data, representing KPI labels and their corresponding values. Example:

```json
{
  "Total Revenue": "$5M",
  "Customer Satisfaction": "92%",
  "Active Users": "15,000",
  "Churn Rate": "3%"
}
```

E. LLM Summarizer: The LLM Summarizer, the last processing step in the pipeline, is in charge of creating logical and comprehensible summaries using the dashboard snapshot and the structured KPI data. Using cutting-edge natural language generation algorithms, this stage converts technical KPI results into useful business insights. Our approach makes use of Gemini, a potent multimodal Large Language Model (LLM) created by Google DeepMind that is excellent at interpreting visual context, analyzing structured data, and generating natural language outputs of superior quality. Gemini's design enables it to combine visual layout information with numerical insights, which results in richer and more contextually relevant summaries than typical text-only models. The generated summaries are also guaranteed to be accurate, brief, and domain-relevant due to its refined ability to handle corporate, financial, and operational terminology. The LLM receives two types of inputs: This JSON includes all key-value pairs extracted by the KPI Extractor and The original image is provided to give the model additional layout and visual context, helping it infer section importance, hierarchies, or missing information.

Prompt Engineering

Prompt engineering is crucial in guiding the LLM to produce high-quality outputs. In our system, the prompts are carefully designed to include:

1. **Instruction Layer:**
   Clear instructions are given to the model, such as:
   "Summarize the key performance indicators extracted from the dashboard concisely in business language. Highlight major positive and negative trends if identifiable. Keep the summary under 150 words."
2. **Input Embedding:**
   The structured KPI JSON is embedded directly into the prompt. Example:
   "Here are the extracted KPIs: {JSON}. Use this information to create the summary."
3. **Contextual Hints:**
   If the dashboard belongs to a specific domain (e.g., sales, operations, marketing), domain-specific hints are added:

"This dashboard belongs to the sales department. Focus on metrics like revenue, customer growth, and churn rate."

4. **Visual Cue Instructions:**
Since Gemini can process images, a reference is added:
"Refer to the visual layout of the attached dashboard image for additional context about KPI importance."
5. **Format Specification:**
Output style is enforced:
"Write in complete sentences. Avoid listing raw numbers without explanation. Maintain a professional business tone."

This multi-layered prompting ensures reliable, repeatable, and high-quality summaries across different dashboards.

IV. Accuracy Metrics:

To make sure the summaries produced by the Business Dashboard Summarizer are precise, relevant, and helpful, it is essential to assess its effectiveness. Since our system includes both text synthesis (LLM Summarization) and text extraction (OCR + KPI Extraction), distinct metrics are required to evaluate each step properly.

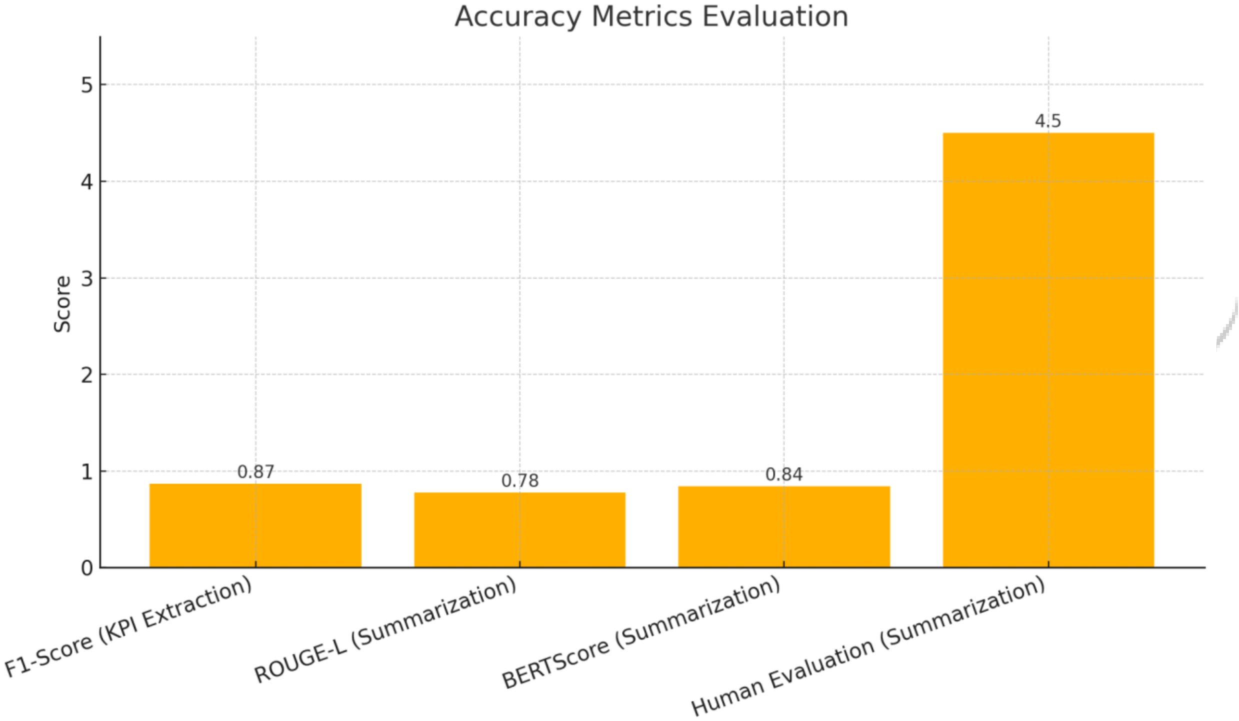

1. **F1-Score:**

$$\text{F1 Score} = 2 \text{ X } \frac{(\text{Precision X Recall})}{\text{Precision + Recall}}$$

This metric evaluates how well the system identifies KPI label-value pairs after OCR processing. A high F1-score of 0.87 (as shown in the graph) indicates the system effectively filters relevant business data while minimizing false positives and false negatives during KPI extraction.

2. **ROUGE-L Score (Summarization):**

$$\text{ROUGE-L} = \frac{\text{Length of Longest Common Subsequence (LCS)}}{\text{Length of Reference Summary}}$$

ROUGE-L evaluates the content overlap between the generated summary and a human-written reference summary. Our system scored 0.78, reflecting a strong ability to capture key ideas from the dashboard using natural language.

3. **BERTScore (Summarization):**

BERTScore = Average cosine similarity between contextual embeddings of reference and generated tokens

This metric measures semantic similarity, going beyond word overlap. A BERTScore of **0.84** indicates the summary meaning closely matches the reference, showing the LLM accurately interpreted business context from structured and visual data.

4. **Human Evaluation (Summarization):**

$$\text{ROUGE-L} = \frac{\sum \text{Evaluator Ratings}}{\text{Number of Evaluator}}$$

Business analysts evaluated each generated summary based on clarity, conciseness, and relevance. With an average rating of 4.5/5, this confirms the summaries are practical, readable, and actionable—validating the real-world usability of our AI system.

V. Conclusion:

This study presents a novel AI-driven approach for summarizing business dashboards using a hybrid pipeline of computer vision, OCR, KPI extraction, and Large Language Models. By transforming raw dashboard screenshots into concise, natural-language summaries, our system addresses the growing problem of cognitive overload in data analysis. Experimental results show high accuracy in both technical extraction (F1-score of 0.87) and semantic summarization (BERTScore of 0.84, human rating of 4.5/5), validating its effectiveness. This solution enhances business intelligence accessibility, accelerates decision-making, and sets a foundation for future multimodal summarization systems in enterprise environments.

References

[1] S. Schreiber, B. Twardowski, and A. Dengel, "DeepDesRT: Deep Learning for Detection and Structure Recognition of Tables in Document Images," International Journal on Document Analysis and Recognition (IJDAR), vol. 22, no. 1, pp. 55–70, 2019.

[2] Y. Xu, M. Zhao, Z. Liu, and L. Li, "Medical Dashboard Data Extraction Using OCR and Deep Learning," IEEE Access, vol. 8, pp. 132220–132230, 2020.

[3] H. Luo, X. Zhang, and Y. Wang, "ChartSense: Interactive Data Extraction from Chart Images," Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI), 2018.

[4] M. R. Alam, H. U. Rehman, and S. Ishihara, "Widget Detection from Application Screenshots for GUI Understanding," Proceedings of the International Conference on Pattern Recognition (ICPR), 2020.

[5] A. K. Jain, R. Sharma, and S. Gupta, "Financial Report Summarization Using BERT-Based Transformers," International Journal of Artificial Intelligence & Applications, vol. 11, no. 2, pp. 15–27, 2021.

[6] N. Liu and T. Lee, "Dashboard Narratives: Using GPT-2 for Natural Language Summary of BI Reports," Proceedings of the IEEE International Conference on Big Data (BigData), 2021.

[7] M. Singh, K. Yadav, and P. Kumar, "SlideSummarizer: Automatic Summarization of Educational Slides Using OCR and Transformer Networks," IEEE Transactions on Learning Technologies, vol. 14, no. 3, pp. 320–331, 2021.

[8] R. Park, J. Seo, and M. Lee, "Multi-modal Summarization of Infographics Using Object Detection and Transformers," Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), 2022.

[9] K. Hart, A. Daniels, and J. Smith, “Reducing Cognitive Overload in Data Dashboards: A UX Perspective,” Journal of Human-Computer Interaction, vol. 36, no. 4, pp. 342–360, 2020.

[10] Y. Cheng, F. Xia, and X. Wang, “Data2Text: A Survey on Dataset-to-Text Generation,” ACM Computing Surveys (CSUR), vol. 54, no. 3, 2021.

[11] P. Laban, J. H. Lee, and K. Toutanova, “Summarizing Text Documents with Transformers: A Survey,” arXiv preprint arXiv:2107.10121, 2021.