REPORT ORCHESTRATOR: GENERATIVE AI FRAMEWORK FOR AUTOMATED SUMMARIZATION AND PRESENTATION DECK GENERATION

Wilhelm Henri R. Alegrado
Jerald M. Constantino
Alberto Y. Zaldivar Jr.

Factory Analytics and Operations Excellence – Global Operations
Western Digital Corp., 109 Technology Ave., SEPZ, Laguna Technopark, Biñan, Laguna, Philippines 4024
wilhelmhenri.alegrado@wdc.com, jerald.constantino@wdc.com, Alberto.Zaldivar@wdc.com

ABSTRACT

In manufacturing, engineers take on the time-consuming task of reviewing, summarizing and creating operational and technical reports. To address this issue, this study presents Report Orchestrator, a Generative AI framework that automates the summarization of documents and generates ready-to-use presentation decks.

The system accepts documents, converts them into images, and leverages a Large Language Model to extract concise, relevant summaries based on pre-defined guidelines. The summarized text is structured into an HTML report and transformed into a slide deck aligned with organizational reporting standards.

Evaluation on a sample document showed that the system accurately captured technical content and organized slides coherently. Outputs included summaries of system architectures and performance metrics, formatted for technical and executive audiences significantly reducing manual workload and improving turnaround time for report preparation.

The framework shows potential for integration into real-world documentation workflows, with future directions focused on visual content generation, prompt personalization, and multi-document summarization.

1. 0 INTRODUCTION

In Western Digital, large volumes of operational, quality, maintenance, and compliance reports are generated daily. These documents often contain critical information necessary for real-time decision-making, auditing, and strategic planning. However, the process of manual reviewing, summarizing, and transforming these lengthy, technical documents into concise insights or presentation materials is time-consuming, labor-intensive, and prone to inconsistencies. As organizations seek greater agility and automation in their workflows, there is a growing need for intelligent systems that can streamline document understanding and reporting tasks.

1.1 Motivations of the Study

One of the major challenges faced by managers, engineers, and analysts in manufacturing environments is the efficient extraction of key information from complex reports and the preparation of presentation materials for meetings, reviews, or audits. This gap in automation hampers productivity and slows down decision cycles, especially in large-scale operations where the volume and diversity of documents are high.

Early approaches to document summarization predominantly utilized rule-based systems and template-driven tools, which rely on predefined patterns or manually crafted rules to extract salient information from text. While these methods offer a degree of control and interpretability, they often lack scalability and adaptability, especially when applied to domain-specific documents such as technical or manufacturing reports[1].

Traditional summarization methods, such as those based on Term Frequency–Inverse Document Frequency (TF-IDF), graph-based ranking, or clustering algorithms, have been employed to highlight keywords or sections, but these often fail to capture the contextual relevance or narrative flow required for executive-level understanding[2,3]. Meanwhile, tools for presentation deck creation are largely manual or semi-automated, requiring human intervention to interpret content and structure slides[4].

Recent advances in Natural Language Processing (NLP) and Generative AI (GenAI), particularly the advent of large language models (LLMs) like Generative Pre-trained Transformer (GPT), have shown promise in understanding and generating coherent summaries and even structured content. However, their application in domain-specific contexts such as manufacturing has been limited due to challenges in domain adaptation, content reliability, and integration with enterprise workflows.

1.2 Research Objectives

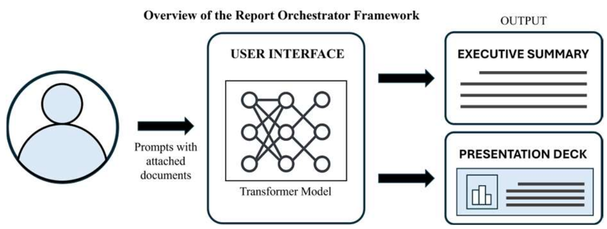

Fig. 1. Overview of the Report Orchestrator Framework. The user uploads a document, which is processed by a transformer model to generate an executive summary and a presentation deck automatically.

As shown in Fig. 1, this paper proposes the development of Report Orchestrator, a Generative AI-based framework that automates the summarization of manufacturing documents and the generation of presentation decks upon user request. The system leverages advanced GPT-based models fine-tuned for industrial reporting language, and features a user-friendly interface for document upload, intelligent summarization, and real-time presentation generation. This framework aims to bridge the gap between raw technical reports and executive-level communication—minimizing manual effort while enhancing consistency, accuracy, and efficiency in manufacturing reporting workflows.

The following section is organized into five chapters. Section 2 presents related work, discussing prior efforts in document summarization and automated presentation generation. Section 3 describes the overall design of the proposed framework, including its system architecture, components, and processing pipeline. Section 4 provides the results of the system evaluation. Finally, Sections 5 and 6 concluded the paper and discussed potential improvements and future work.

2.0 REVIEW OF RELATED WORK

This section reviews existing literature relevant to the development of automated summarization and presentation generation systems. It covers core techniques in document summarization, recent advances in AI-driven slide generation, and related methods that inform the proposed framework.

2.1 Document Summarization Techniques

Automatic text summarization is a pivotal task in NLP, aiming to condense large volumes of text into concise summaries while preserving essential information. Techniques are broadly categorized into extractive and abstractive methods[5]. Table 1 illustrates the key differences between the two.

Extractive summarization involves selecting key sentences or phrases directly from a source text. Among the prominent techniques is TextRank, an unsupervised graph-based algorithm that constructs a graph of sentences based on similarity and ranks them by importance. Its effectiveness has been demonstrated in multiple contexts. For instance, Barman et al.[6] combined TextRank with GloVe embeddings to summarize English news articles, achieving strong ROUGE scores across five domains—demonstrating high alignment with human-written summaries. Similarly, Manjari[7] applied TextRank to Telugu documents and reported an average F1-score of 0.621, highlighting the method's language-independent adaptability.

Table 1. Comparison of Extractive and Abstractive Summarization

Aspect	Extractive	Abstractive
Approach	Selecting existing sentences	Generates new sentences
Output	Exact phrases from source	Rephrased, concise content
Complexity	Simpler, rule-based or graph-based	Complex, uses deep learning models
Fluency	May be disjointed	More natural and human-like
Accuracy	Factually reliable	May introduce minor inaccuracies
Use Case	Structured, factual texts	General articles, reports with narrative flow

In contrast, abstractive summarization generates new sentences that paraphrase the source content. Ahmed Zeyad and Biradar[8] showed that Flan-T5 achieved high performance on the Gigaword dataset, with a ROUGE-L of 0.5021 and BERTScore of 0.91, indicating strong semantic similarity with reference summaries. Liaqat et al.[9] proposed a hybrid model that first extracts keywords and then generates summaries using BERT-based architecture, enhancing both accuracy and readability.

In the context of manufacturing reports, which frequently include technical jargon and structured data, hybrid summarization approaches have demonstrated promise. By first employing extractive methods to isolate critical information and then using abstractive models to rephrase it, these techniques produce summaries that are both accurate and accessible to a broader audience.

2.2 Presentation Deck Generation

The automation of presentation slide creation has gained attention with advances in AI, yet many current tools remain manual or semi-automated, requiring users to interpret and organize content manually. Recent work addresses this gap through systems that integrate NLP and machine learning to automate slide generation directly from source documents.

Fu et al.[10] introduced DOC2PPT, a sequence-to-sequence model that generates structured slide decks by integrating summarization, layout prediction, and text-image alignment. Trained in 5,800 document-slide pairs, it outperformed baselines on metrics like ROUGE-SL and TFR, ensuring both content accuracy and visual coherence.

Similarly, Wang et al.[11] developed a phrase-based method that builds bullet-structured slides by extracting key phrases and predicting hierarchical relationships. Using greedy algorithms for phrase alignment, their approach outperformed models like LexRank and PPSGen in preserving semantic relevance and structural clarity.

These studies underscore the growing potential of AI-driven frameworks to automate the creation of presentation-ready materials, especially valuable for translating complex manufacturing reports into structured, executive-level briefings.

3.0 METHODOLOGY

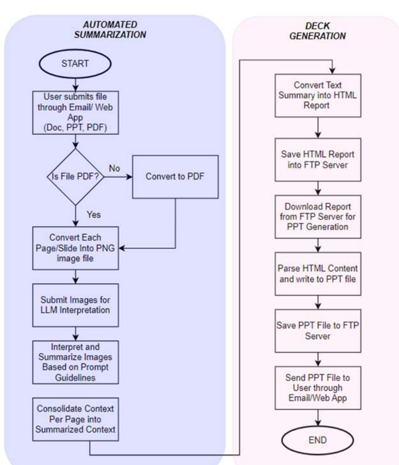

Fig. 2. Report Orchestrator Flowchart. The report orchestrator framework is divided into two stages: automated summarization and slide deck generation.

Fig. 2 shows the two core processing modules of the Report Orchestrator framework: (1) Automated Summarization, and (2) Slide Deck Generation. The process starts with document submission through web channel/email and ends with the generated presentation deck being sent back through the same web channel/email.

3.1 Automated Summarization

The summarization stage starts with the user submitting a document for summarization as an attachment through an email channel or a dedicated web application. Supported input formats include Microsoft Word, PowerPoint, and PDF files.

You are an advanced AI assistant capable of analyzing images.

Given the images, carefully examine all aspects of it, including objects, colors, and any text or symbols present.

Provide thorough and descriptive responses of what you observe.

Instructions:

1. Analyze the image in detail, noting all relevant elements.

2. Include specific details from the image in your response.

3. If applicable, make informed inferences or interpretations based on what you see.

4. Use clear and descriptive language to paint a vivid picture for the reader.

5. Organize your response into logical paragraphs or sections.

Please provide your detailed response:

Fig. 3. Image Interpretation Prompt. The prompt outlines the role and instructions needed to generate relevant and insightful summaries for each image.

In cases where the uploaded file is not in PDF format, it is converted to PDF to meet the processing requirements in the succeeding steps. Each page or slide of the PDF document is then converted into an image (PNG format). The resulting set of images is then passed through a Large Language Model (LLM)-driven interpretation module. Through this conversion step, context in differing formats (plain text, charts, tables, flowcharts, screenshots, etc.) can be fully interpreted by the LLM.

Based on pre-configured prompt guidelines shown in Fig. 3, the LLM processes each image to extract concise, relevant summaries aligned with the preferences and communication standards set for the system.

These per-page summaries are then aggregated and consolidated into a unified, coherent executive summary that captures the overall narrative and technical highlights of the original document.

3.2 Slide Deck Generation

Following the generation of the document summary, the system transitions to the slide deck generation stage.

The text summary is first converted into a structured HTML report, which acts as an intermediate representation suitable for slide generation. This HTML file is saved to an FTP server for subsequent access.

The system then retrieves the HTML report and parses its contents to generate a PowerPoint (PPT) file. The parsing process maps structured sections of the HTML (e.g., titles, bullet points, tables) into appropriate slide components, such as headings, bullet lists, and content placeholders.

After the slide content is generated, the final presentation deck (PPT file) is saved back to the FTP server. Finally, the completed PPT file is sent to the user via their original submission channel, either through email or via the dedicated web application.

For a full overview of the system architecture supporting this process, see Appendix A.

4.0 RESULTS AND DISCUSSION

This section presents the testing results of the Report Orchestrator framework, which automates the summarization of technical documents and the generation of presentation decks. To evaluate the system's effectiveness, the proponents used a published technical paper[12] containing domain-specific content on manufacturing processes as the test document.

Fig. 4. System Prompt for Summarization and Presentation Generation. The figure shows the input prompt instructing the system to summarize a technical document and generate a structured presentation deck.

A sample system prompt used for testing is shown in Fig. 4, which instructs the framework to generate a concise summary and a structured presentation deck from the provided document. The outputs generated by the system, including the executive summary and the structured presentation deck, are analyzed in terms of content accuracy, coherence, and alignment with standard reporting formats. Observations regarding the system's capabilities and limitations are also discussed.

4.1 Automated Summarization Result

The Report Orchestrator framework successfully generated an automated executive summary from the input technical document. As shown in the generated output, the summary provided a concise yet informative overview of the document's key points, including system objectives, achieved performance metrics, and areas of focus for subsequent sections.

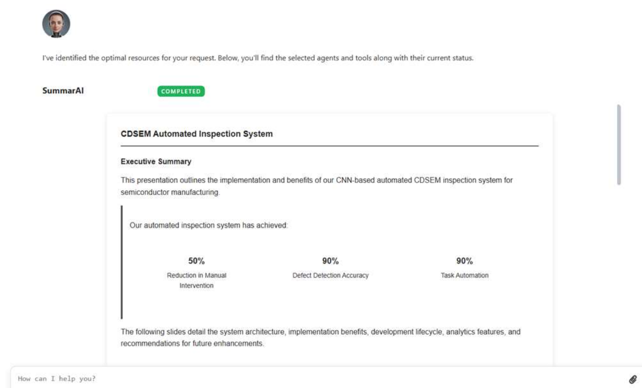

Fig. 5. Sample of Generated Executive Summary Output. The output shows the automatically generated executive summary highlighting system performance metrics and key implementation details for the manufacturing report.

The generated executive summary shown in Fig. 5 highlights the implementation and benefits of a CNN-based automated CDSEM inspection system designed to improve manufacturing efficiency and quality control. The output also emphasized quantifiable results achieved by the system, including a 50% reduction in manual intervention, 90% defect detection accuracy, and 90% task automation. These details align with typical expectations for high-level executive summaries in technical reports, providing stakeholders with relevant performance indicators immediately.

The system also generated a set of key findings summarizing critical outcomes from the source document. As shown in Appendix B, these key findings were presented in a clear, bullet-point format that emphasized essential insights, including improved manufacturing efficiency, integration with existing systems, and enhanced decision-making through real-time dashboards. The concluding statement within the generated summary effectively reinforced the system's overall impact on streamlining quality control processes across manufacturing operations.

Finally, the automated summarization output demonstrated coherence, readability, and relevance suitable for executive-level consumption. Manual inspection confirmed that the generated summary retained the core information necessary for high-level understanding, with minimal omission of essential details. This result highlights the system's capability to condense lengthy technical content into accessible, actionable insights, reducing the manual effort typically required for report summarization.

4.2 Slide Deck Generation Result

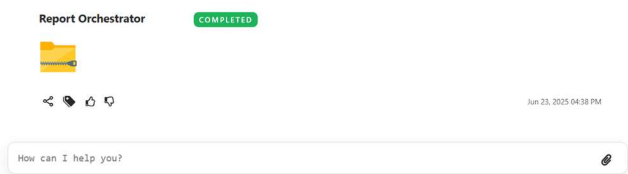

Fig. 6. Slide Deck Generation Output. The output response shows the system interface providing the downloadable ZIP file containing the automatically generated presentation slides.

Following the automated summarization process, the Report Orchestrator framework generates a compressed file containing the structured presentation deck. As shown in Fig. 6, once the summarization task is completed, the system interface provides a downloadable ZIP file that contains the generated PowerPoint slides summarizing the source document.

Upon extracting the generated file, the system produces a complete presentation deck designed to highlight the key findings, technical content, and recommendations from the original document. The generated slide deck follows standard presentation structures, including an executive summary slide, content slides organized by technical topic, and a conclusion slide summarizing overall insights and next steps.

A sample of the generated slides is shown in Fig. 7, illustrating the system architecture section. The slide presents a concise summary of the technical implementation, detailing key aspects such as the VGG-based CNN architecture, deployment strategy using containerization, and data handling processes. The content is organized in a clear, bullet-point format consistent with standard technical presentation practices, making the information accessible to both technical and non-technical stakeholders.

Additionally, the slide includes a system-generated disclaimer indicating that the content is AI-generated and requires human validation. This design choice emphasizes the intended role of the framework as an assistive tool to accelerate report processing, rather than a fully autonomous solution for technical content generation.

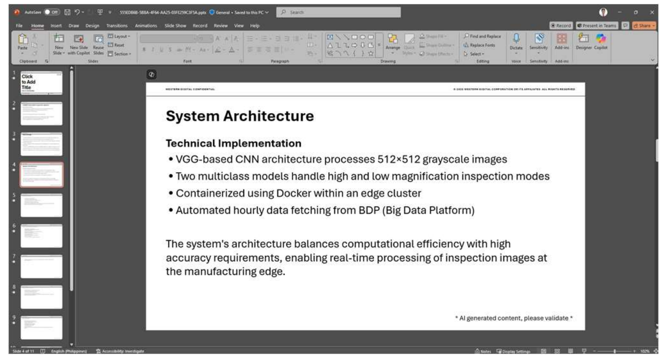

Fig. 7. Sample Generated Presentation Slide. The result shows an automatically generated slide summarizing the system architecture and technical implementation details extracted from the source document.

Lastly, the generated slide deck effectively translates complex technical documents into concise, visually structured presentations. The manual review confirmed that the generated slides accurately reflected the main points of the source document while maintaining readability and structural coherence. However, limitations were observed, particularly in the absence of graphical content such as charts or diagrams, which are commonly included in technical presentations to complement textual information.

Despite these limitations, the results demonstrate that the Report Orchestrator framework can significantly reduce the manual effort required to prepare presentation materials from technical reports, contributing to improved efficiency and consistency in reporting workflows within manufacturing environments.

5.0 CONCLUSION

This paper presented the development and deployment of the Report Orchestrator framework, a GenAI-based system designed to automate the summarization of technical documents and the generation of presentation decks for manufacturing reports. The system integrates domain-adapted LLMs to convert lengthy, complex technical documents into concise executive summaries and structured slide decks with minimal human intervention. Results demonstrated that the framework effectively condensed key information from technical reports, producing coherent, readable summaries and presentation materials aligned with standard reporting formats. Finally, the system exhibited the potential to significantly reduce manual effort and improve consistency in preparing reporting materials, addressing key challenges in manufacturing environments where large volumes of technical documentation are routinely generated.

6.0 RECOMMENDATIONS

Within the current framework, image interpretation robustness and prompt personalization are two target areas for improvement. First, the image interpretation prompt needs to be continuously tuned based on the nature of the documents across different domains in the organization. Documents can include any combination of text, tables and domain-specific charts or illustrations which may require additional processing to fully extract the context contained inside them. Second, each department has specific preferences regarding the format and structure of their presentation decks. Incorporating these user preferences into the system prompt for slide deck generation helps the system align with the user's needs.

Beyond the current framework, several areas remain open for future exploration. One key direction is the extension of the framework to support multi-document summarization, enabling the system to consolidate information from multiple related reports, which is particularly valuable for comprehensive operational reviews. Another area for future work is the integration of AI models capable of generating relevant visual content, such as charts, tables, and diagrams, to complement textual summaries and presentation materials. Finally, investigating more advanced domain adaptation techniques, such as prompt tuning or knowledge injection, could further enhance the system's ability to handle complex, technical, and highly specialized manufacturing documents.

7.0 ACKNOWLEDGMENT

The authors would like to express their gratitude to Abhinay Modala Sriranga for his valuable contributions to the project. Appreciation is also extended to the Operation Excellence team for their continuous support and collaboration. Finally, the team acknowledges the assistance provided by IT, particularly in supplying the necessary resources to accomplish this project.

8.0 REFERENCES

1. N. Moratanch and S. Chitrakala, "A survey on abstractive text summarization," in **2016 International Conference on Circuit, Power and Computing Technologies (ICCPCT)**, Nagercoil, India, **2016**, pp. 1-7, doi: 10.1109/ICCPCT.2016.7530193.

2. S. Saha and C. K. Behera, "Real-Time Text Classification and Summarization Model by Enhancing TF-IDF Extractive Summarization Method," in **2024 International Conference on Intelligent Computing and Sustainable Innovations in Technology (IC-SIT)**, Bhubaneswar, India, **2024**, pp. 1-6, doi: 10.1109/IC-SIT63503.2024.10861972.

3. A. Srikanth, A. S. Umasankar, S. Thanu and S. J. Nirmala, "Extractive Text Summarization using Dynamic Clustering and Co-Reference on BERT," in **2020 5th International Conference on Computing, Communication and Security (ICCCS)**, Patna, India, **2020**, pp. 1-5, doi: 10.1109/ICCCS49678.2020.9277220.

4. M. J. Costa, R. Matos, H. Amaro, B. Caceiro, A. Marques and H. G. Oliveira, "SmartEDU: A Platform for Generating Education-Support Materials," in **2023 6th Experiment@ International Conference (exp.at'23)**, Évora, Portugal, **2023**, pp. 55-56, doi: 10.1109/exp.at2358782.2023.10545665.

5. S. Modi and R. Oza, "Review on Abstractive Text Summarization Techniques (ATST) for single and multi documents," in **2018 International Conference on Computing, Power and Communication Technologies (GUCON)**, Greater Noida, India, **2018**, pp. 1173-1176, doi: 10.1109/GUCON.2018.8674894.

6. U. Barman, V. Barman, M. Rahman and N. K. Choudhury, "Graph Based Extractive News Articles Summarization Approach leveraging Static Word Embeddings," in **2021 International Conference on Computational Performance Evaluation (ComPE)**, Shillong, India, **2021**, pp. 008-011, doi: 10.1109/ComPE53109.2021.9752056.

7. K. U. Manjari, "Extractive Summarization of Telugu Documents using TextRank Algorithm," in **2020 Fourth International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC)**, Palladam, India, **2020**, pp. 678-683, doi: 10.1109/I-SMAC49090.2020.9243568.

8. A. M. Ahmed Zeyad and A. Biradar, "Advancements in the Efficacy of Flan-T5 for Abstractive Text Summarization: A Multi-Dataset Evaluation Using ROUGE and BERTScore," in **2024 International Conference on Advancements in Power, Communication and Intelligent Systems (APCI)**, KANNUR, India, **2024**, pp. 1-5, doi: 10.1109/APCI61480.2024.10616418.

9. M. I. Liaqat, I. Hamid, Q. Nawaz and N. Shafique, "Abstractive Text Summarization Using Hybrid Technique of Summarization," in **2022 14th International Conference on Communication Software and Networks (ICCSN)**, Chongqing, China, **2022**, pp. 141-144, doi: 10.1109/ICCSN55126.2022.9817599.

10. T.-J. Fu, W. Y. Wang, D. McDuff and Y. Song, "DOC2PPT: Automatic presentation slides generation from scientific documents," in **Proc. AAAI Conf. Artif. Intell.**, vol. 36, no. 1, **2022**, pp. 634–642, doi: 10.1609/aaai.v36i1.19943.

11. S. Wang, X. Wan and S. Du, "Phrase-based presentation slides generation for academic papers," in **Proc. AAAI Conf. Artif. Intell.**, vol. 31, no. 1, **2017**, doi: 10.1609/aaai.v31i1.10481.

12. J. Constantino and J. Gonzales, "CDSEM Automated Inspection and Defects Classification using a Multiclass Deep

Learning Approach," in **33rd ASEMEP National Technical Symposium (ANTS)**, Manila, Philippines, **2024**.

9.0 ABOUT THE AUTHORS

Wilhelm Henri Alegrado is a Sr. Data Scientist at Western Digital Storage Technologies for Operation Excellence. He holds a master's degree in data science at Asian Institute of Management. He has been with the company for 13 years with experience from previous departments with an Engineering and Development background.

Jerald Constantino is an Associate Data Scientist at Western Digital Corporation Philippines, specializing in machine vision, machine learning, and generative AI projects. He holds a Bachelor of Science in Computer Engineering from Batangas State University - TNEU and is currently studying for a Master of Science in Computer Engineering at Mapúa University.

Alberto Zaldivar is a Technology Lead at Western Digital, with a total of 20 years of experience. He holds a Bachelor of Science in Computer Engineering from Adamson University.

10.0 APPENDIX

Appendix A – Report Orchestrator Architecture.

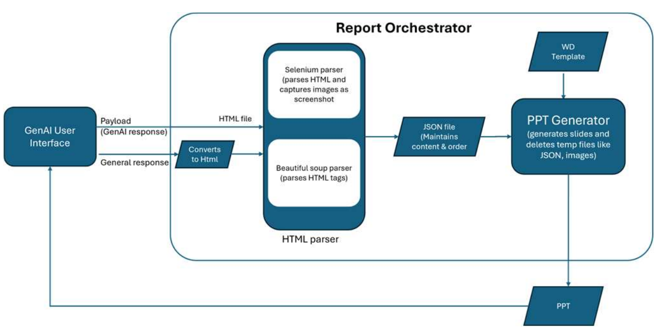

Appendix B – Generated Automated Summarization Output.