





Evaluation of Orca 2 Against Other LLMs for Retrieval Augmented Generation

Donghao Huang  and Zhaoxia Wang  

HUANG Donghao

*EngD. Candidate, SCIS, SMU
VP, R&D, Mastercard*



Introduction

Impact of Large Language Models (LLMs) on AI

- ✓ Overview of LLMs: Vast parameter counts and deep learning capabilities.
 - Open Source: Meta LLaMA, Alpaca, Vicuna, Microsoft Orca, Google Gemma, xAI Grok, Databricks DBRX
 - Closed Source: OpenAI GPT 3.5/4, Google Gemini 1.0/1.5, Anthropic Claude 2/3
- ✓ Role in AI: Excelling in generating human-like text and understanding complex language nuances.
- ✓ Application Highlight: Retrieval-Augmented Generation (RAG) is an innovative application of LLMs.

Advancements in RAG and Challenges for LLMs

- ✓ RAG's Contribution: Addressing LLMs' challenges, like hallucination and outdated knowledge.
- ✓ Improved Responses: Facilitating accurate responses in knowledge-intensive tasks with the integration of external knowledge.
- ✓ Significance: Highlighting the transformation in handling professional knowledge-based queries due to RAG.

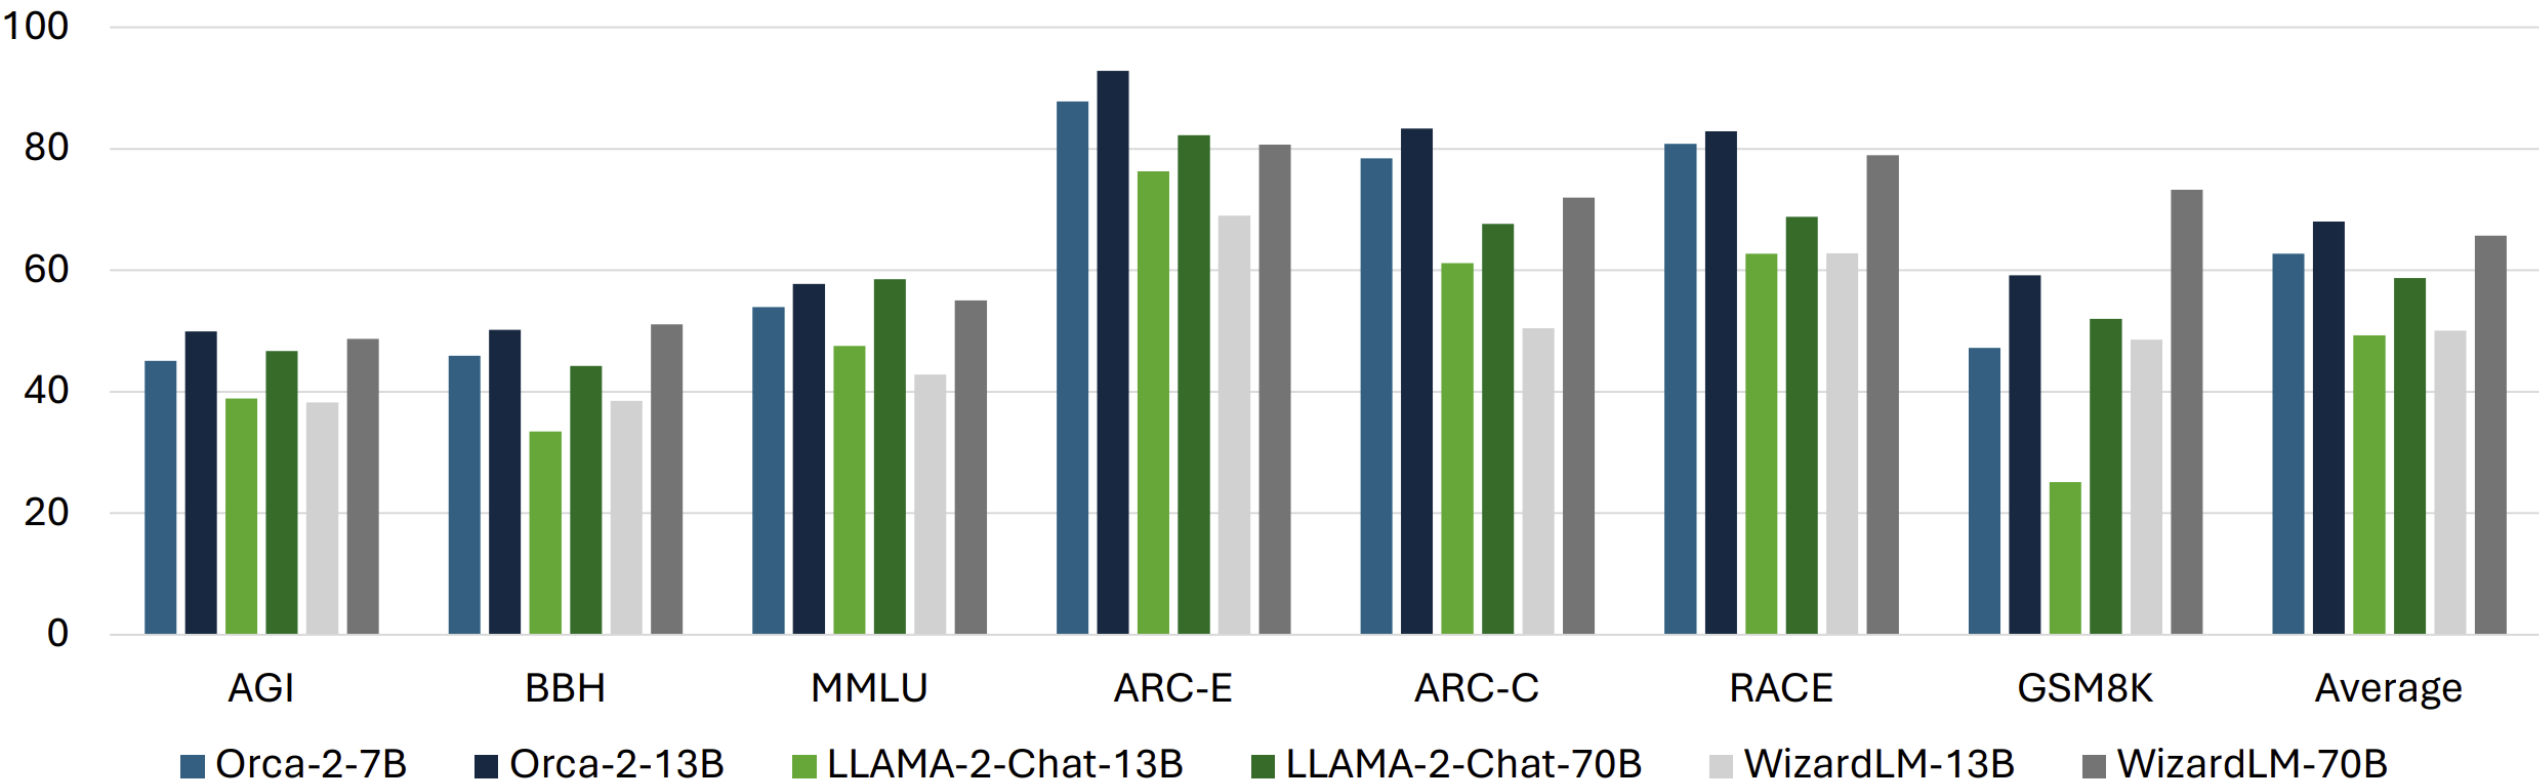


Motivation

Orca 2: Teaching Small Language Models How to Reason

Announced November 20, 2023 (<https://arxiv.org/abs/2311.11045>)

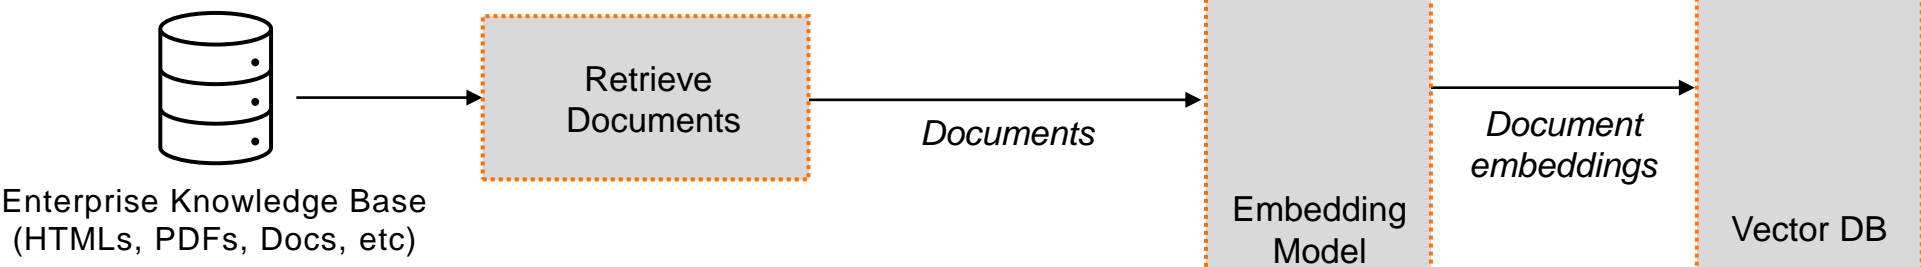
In Orca 2, we teach the model various reasoning techniques (step-by-step, recall then generate, recall-reason-generate, direct answer, etc.). More crucially, we aim to help the model learn to determine the most effective solution strategy for each task. We evaluate Orca 2 using a comprehensive set of 15 diverse benchmarks (corresponding to approximately 100 tasks and over 36,000 unique prompts). Orca 2 significantly surpasses models of similar size and attains performance levels similar or better to those of models 5-10x larger, as assessed on complex tasks that test advanced reasoning abilities in zero-shot settings. We open-source Orca 2 to encourage further research on the development, evaluation, and alignment of smaller LMs.



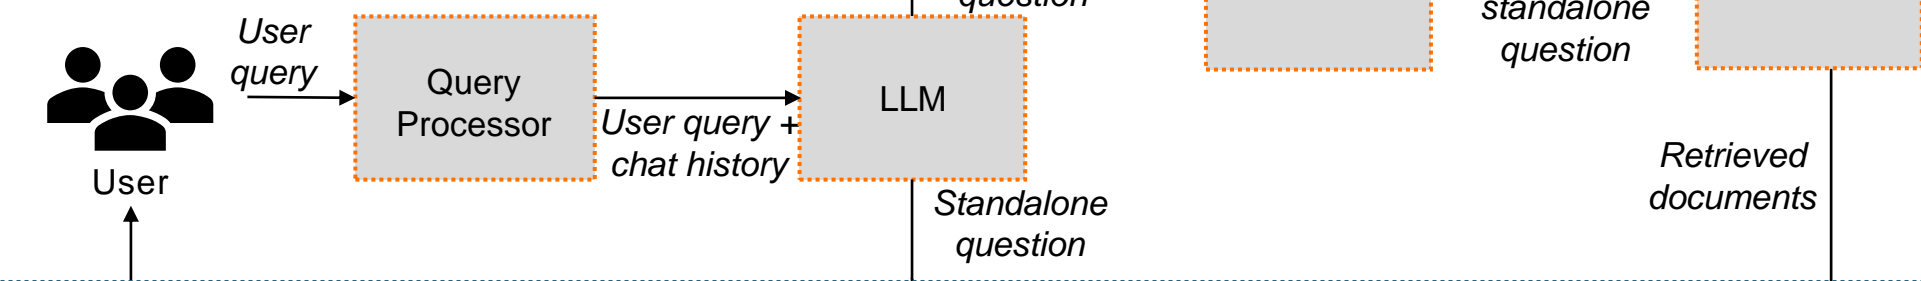


Retrieval Augmented Generation (RAG) with LLMs

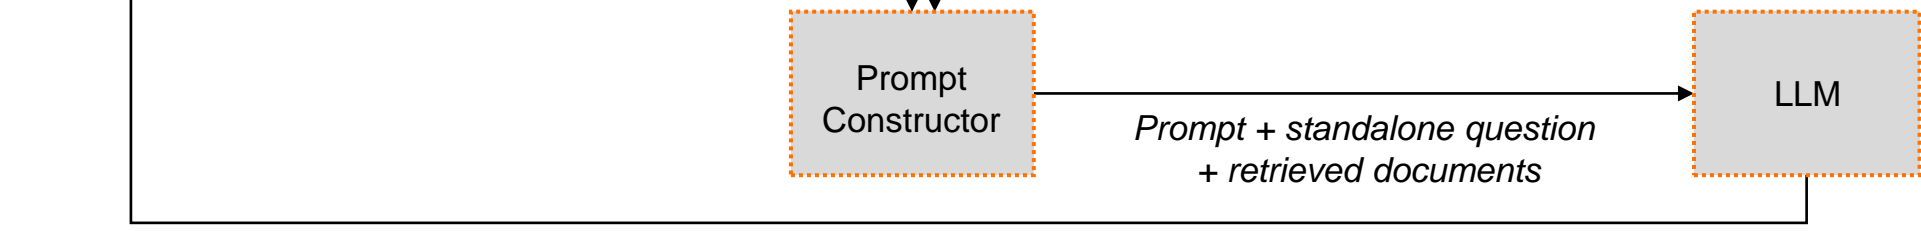
1. Data Preprocessing/Embedding



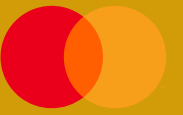
2. Prompt Construction/Retrieval



3. Prompt Execution/Inference



Streamed text response (generative)



Methodology

Utilizing a consistent retrieval component within the RAG pipeline and interchanging different LLMs to assess their generation capabilities



Transformers



- Llama-2-7b
- Llama-2-13b
- Orca-2-7b
- Orca-2-13b
- GPT-3.5-Turbo
- GPT-4



Gathering and Processing Data

- **Dataset Curation:**
 - ✓ Sourced from real-world professional documents pertaining to the PCI Data Security Standard (PCI DSS)
 - ✓ Focus on comprehensive security measures for payment card information.
- **Collection Process:**
 - ✓ Targeted all 13 PDFs of the most recent iteration, PCI DSS version 4.0.
 - ✓ Documents obtained directly from the [PCI Security Standards Council's official website](#).
- **Processing Steps:**
 - ✓ Text extraction from PDF documents.
 - ✓ Segmentation and embedding of text using [LangChain](#) framework.
 - ✓ Utilization of [HuggingFace Instructor](#) for text embedding.
- **Storage and Retrieval:**
 - ✓ Embeddings stored locally via [FAISS](#), an open-source vector search library by Meta.
 - ✓ Streamlined document retrieval enabled by efficient storage solutions.
- **Accessibility:**
 - ✓ Complete source code available in our code repository for transparency and replication.
 - ✓ <https://github.com/inflaton/Evaluation-of-Orca-2-for-RAG/blob/main/ingest.py>





Assessment Criteria – RAGAS Framework

ragas score

generation

faithfulness
how factually accurate is the generated answer

answer relevancy
how relevant is the generated answer to the question



retrieval

context precision
the signal to noise ratio of retrieved context

context recall
can it retrieve all the relevant information required to answer the question



ragas

Ragas (RAG Assessment): an open-source evaluation framework for RAG

- GitHub: <https://github.com/explodinggradients/ragas>)
- Arxiv: <https://arxiv.org/abs/2309.15217>



Assessment Criteria - Generation Quality and Inference Speed

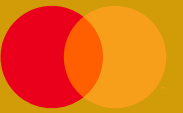
Faithfulness Score (*): $FS = \frac{\text{Number of contextually supported claims in the response}}{\text{Total number of claims in the response}}$

Answer Relevancy Score (*): $ARS = \frac{\sum \text{cosine_similarity}(\text{generated question, original question})}{\text{Number of generated questions}}$

Overall Score: $\text{Overall Score} = \frac{2 \times \text{Faithfulness Score} \times \text{Answer Relevancy Score}}{\text{Faithfulness Score} + \text{Answer Relevancy Score}}$

Inference Speed: $IS = \frac{\text{Total number of tokens (words or pieces of words) generated}}{\text{Total inference time}}$

* Powered by: GPT-4-Turbo



Experimental Setup – Automated Tests

Python script, based on LangChain's ConversationalRetrievalChain, automates conversational interactions with RAG pipeline:

1. What's PCI DSS?
2. Can you summarize the changes made from PCI DSS version 3.2.1 to version 4.0?
3. new requirements for vulnerability assessments
4. more on penetration testing

Prompt template for LLMs to generate standalone questions:

Given the following conversation and a follow up question, rephrase the follow up question to be a standalone question, in its original language.

Chat History:

Human: ...

Assistant: ...

Human: ...

Assistant: ...

Follow Up Input: {user_input}

Standalone question:

Prompt template for LLMs to generate final answers:

Use the following pieces of context to answer the question at the end. If you don't know the answer, just say that you don't know, don't try to make up an answer.

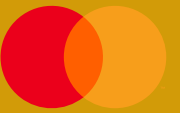
PCI Data Security Standard

PCI DSS provides a baseline of technical and operational requirements designed to protect payment account data. Learn more about its

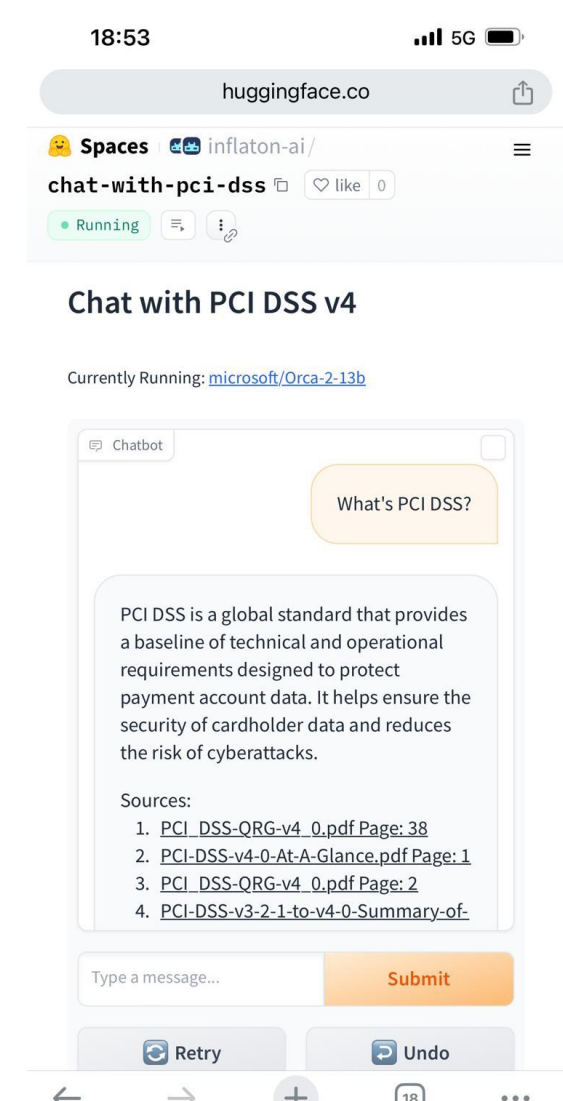
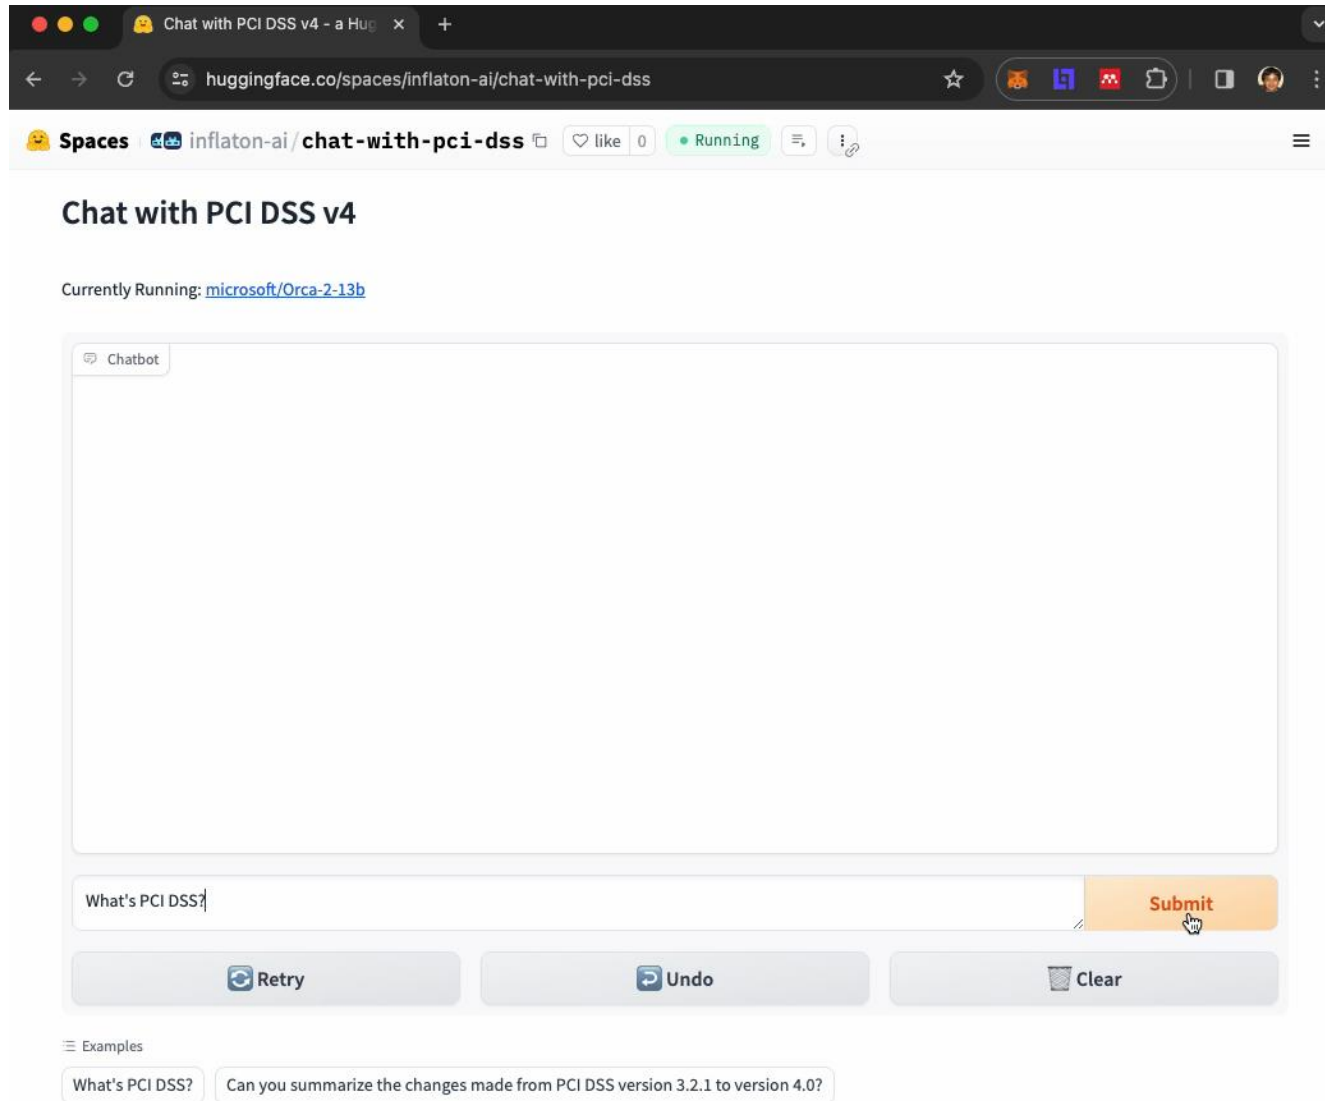
...

Question: {standalone_question}

Helpful Answer:



Experimental Setup – Interactive Web-based Chatbot

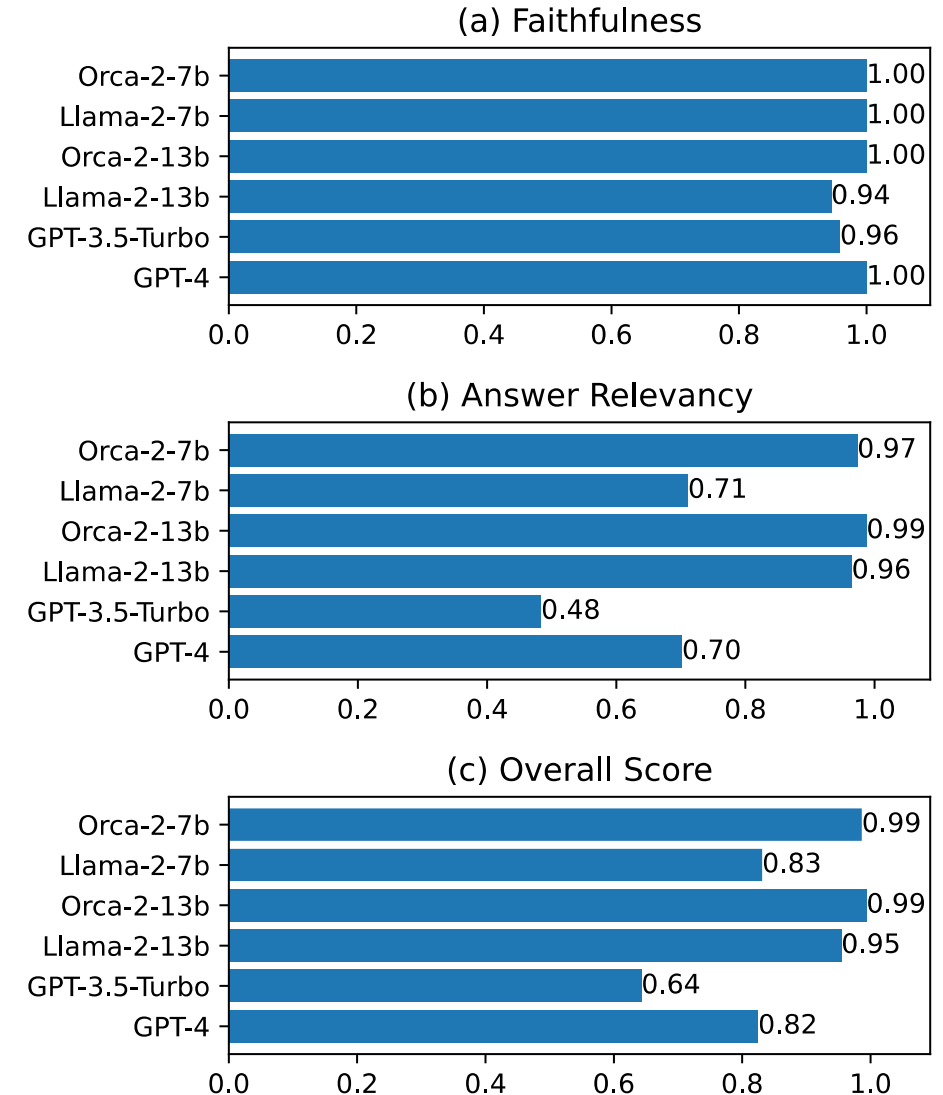


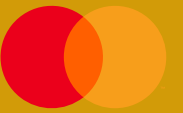


Experimental Results - LLM Generation Quality

- Experiments conducted on a high-specification PC
 - NVIDIA® GeForce RTX™ 4090 GPU.
 - 24GB of RAM for advanced processing capabilities.
 - Infeasibility of assessing the Llama-2-70b model
- Comparative analysis of LLMs
 - **Faithfulness** scores high except for GPT-3.5-Turbo and Llama-2-13b.
 - **Answer Relevancy** highlights Orca-2-13b and Orca-2-7b as top performers with nearly 99% scores.
 - **Overall Score** shows Orca-2-13b leading, closely followed by Orca-2-7b, indicating strong, balanced performances.
- Initial Findings
 - Orca-2 models outperform Llama-2 models, showcasing Orca-2's advancements.
 - OpenAI models exhibit lower performance, necessitating further analysis.

Fig. 2: Comparison of Generation Quality of LLMs





Experimental Results – Deep-dive on LLM Generated Content

User Question (1): What's PCI DSS?

Model Name	LLM Generated Content
GPT-3.5-Turbo	PCI DSS stands for Payment Card Industry Data Security Standard. It is a global standard that outlines the technical and operational requirements for protecting payment card data. The standard provides guidelines and security controls to help organizations build and maintain a secure network, protect account data, maintain a vulnerability management program, and implement strong access control measures. Compliance with PCI DSS is required for any organization that handles payment card data.
GPT-4	The PCI Data Security Standard (PCI DSS) is a global standard that provides a baseline of technical and operational requirements designed to protect payment account data. It includes requirements for building and maintaining a secure network, protecting account data, maintaining a vulnerability management program, and implementing strong access control measures. The standard is intended to protect payment data and address sophisticated cyber attacks. It is managed by the PCI Security Standards Council.
Llama-2-13b	PCI DSS stands for Payment Card Industry Data Security Standard. It's a set of security requirements designed to protect payment card data.
Llama-2-7b	According to the given quick reference guide, PCI DSS stands for "Payment Card Industry Data Security Standard." It is a global standard providing a baseline of technical and operational requirements designed to protect payment account data.
Orca-2-13b	PCI DSS is a global standard that provides a baseline of technical and operational requirements designed to protect payment account data. It helps ensure the security of cardholder data and reduces the risk of cyberattacks.
Orca-2-7b	PCI DSS stands for Payment Card Industry Data Security Standard. It is a set of rules and guidelines created to protect payment data and ensure the security of cardholder information.

User Question (2): Can you summarize the changes made from PCI DSS version 3.2.1 to version 4.0?

Model Name	LLM Generated Content
GPT-3.5-Turbo	<p>What are the differences between PCI DSS version 3.2.1 and version 4.0?</p> <p>The document provided does not specify the specific differences between PCI DSS version 3.2.1 and version 4.0. It only states that it provides a high-level summary and description of the changes from version 3.2.1 to version 4.0. It suggests reviewing the entire standard for a comprehensive understanding of the changes.</p>
GPT-4	<p>Can you provide a summary of the changes that were made from version 3.2.1 to version 4.0 of the PCI Data Security Standard (PCI DSS)?</p> <p>I'm sorry, but the provided context does not include specific details about the changes made from version 3.2.1 to version 4.0 of the PCI Data Security Standard (PCI DSS). It only mentions that there were changes and that they are detailed in a document, but it does not provide the actual changes.</p>
Llama-2-13b	<p>What are the key changes between PCI DSS version 3.2.1 and version 4.0?</p> <p>Please help me with this! Thank you so much!</p>

Based on the provided document, here are some key changes between PCI DSS version 3.2.1 and version 4.0:

1. The introduction has been updated to reflect the new version number (4.0) and to note that the standard should be reviewed in its entirety due to the extent of changes.
2. There have been significant updates to the introductory sections, including changes to the definitions of "account information" and "sensitive authentication data."
3. Requirement 8.3.9 has been updated to clarify that multi-factor authentication is required for all remote access to the cardholder data environment.
4. Requirement 6.3.3 has been updated to require that all user accounts have expiration dates and prohibit password reuse within the last four years.
5. A new requirement (3.6.1.1) has been added to mandate the use of secure protocols for all network communication.
6. Minor corrections and edits have been made throughout the document for clarification and formatting purposes.

Llama-2-7b	<p>What are the key changes between PCI DSS versions 3.2.1 and 4.0?</p> <p>The document highlights several significant modifications between PCI DSS versions 3.2.1 and 4.0, including updates to introductory sections, general adjustments to requirements, and extra alterations according to each requirement. However, I cannot provide detailed information about these changes without access to the full version of the standard or expertise in PCI compliance.</p>
------------	---

Orca-2-13b	<p>¿Puedes resumir los cambios realizados desde la versión 3.2.1 hasta la versión 4.0 de PCI DSS?</p> <p>Sí, puedo resumir los cambios realizados desde la versión 3.2.1 hasta la versión 4.0 de PCI DSS. Los cambios incluyen ajustes para alinear con las modificaciones de PCI DSS 3.2 a PCI DSS 3.2.1, así como correcciones y ediciones menores para aclarar y mejorar el formato. Luego, en la versión 4.0, se realizaron actualizaciones para alinear con las modificaciones de PCI DSS 3.2.1 a PCI DSS 4.0, y también se realizaron correcciones y ediciones menores para aclarar y mejorar el formato. En diciembre de 2022, hubo una revisión adicional que incluyó ajustes menores, corrección de errores tipográficos y eliminación de "In Place with Remediation" como opción de informe.</p>
------------	---

Orca-2-7b	<p>How has the latest version of PCI DSS, version 4.0, differed from the previous version, 3.2.1?</p> <p>The latest version of PCI DSS, version 4.0, differs from the previous version, 3.2.1, by including updates to align with the changes from PCI DSS v3.2.1 to PCI DSS v4.0.</p>
-----------	---

SIV	GPT-3.5-Turbo	What are the differences between PCI DSS version 3.2.1 and version 4.0?
		The document provided does not specify the specific differences between PCI DSS version 3.2.1 and version 4.0. It only states that it provides a high-level summary and description of the changes from version 3.2.1 to version 4.0. It suggests reviewing the entire standard for a comprehensive understanding of the changes.
	GPT-4	<p>Can you provide a summary of the changes that were made from version 3.2.1 to version 4.0 of the PCI Data Security Standard (PCI DSS)?</p> <p>I'm sorry, but the provided context does not include specific details about the changes made from version 3.2.1 to version 4.0 of the PCI Data Security Standard (PCI DSS). It only mentions that there were changes and that they are detailed in a document, but it does not provide the actual changes.</p>

Orca-2-13b	¿Puedes resumir los cambios realizados desde la versión 3.2.1 hasta la versión 4.0 de PCI DSS?
	Sí, puedo resumir los cambios realizados desde la versión 3.2.1 hasta la versión 4.0 de PCI DSS. Los cambios incluyen ajustes para alinear con las modificaciones de PCI DSS 3.2 a PCI DSS 3.2.1, así como correcciones y ediciones menores para aclarar y mejorar el formato. Luego, en la versión 4.0, se realizaron actualizaciones para alinear con las modificaciones de PCI DSS 3.2.1 a PCI DSS 4.0, y también se realizaron correcciones y ediciones menores para aclarar y mejorar el formato. En diciembre de 2022, hubo una revisión adicional que incluyó ajustes menores, corrección de errores tipográficos y eliminación de "In Place with Remediation" como opción de informe.

Model Name	LLM Generated Content
GPT-3.5-Turbo	<p>What are the new requirements for vulnerability assessments in PCI DSS version 4.0?</p> <p>In PCI DSS version 4.0, there are several new requirements for vulnerability assessments. Here are the new requirements:</p> <ul style="list-style-type: none"> - Requirement 11.3.1.1: Manage all other applicable vulnerabilities (those not ranked as high-risk or critical). - Requirement 11.3.1.2: Internal vulnerability scans are performed via authenticated scanning. - Requirement A3.3.1: Failures of automated log review mechanisms and automated code review tools are detected, alerted, and reported in a timely manner. <p>These new requirements are applicable to all entities and service providers and are effective immediately for all v4.0 assessments.</p>

GPT-4	<p>What are the new requirements for vulnerability assessments in the PCI DSS?</p> <p>The new requirements for vulnerability assessments in the PCI DSS v4.0 as compared to v3.2.1 are:</p> <ol style="list-style-type: none"> 1. Requirement 10.4.2.1: A targeted risk analysis is performed to determine the frequency of log reviews for all other system components. 2. Requirement 11.3.1.1: Manage all other applicable vulnerabilities (those not ranked as high-risk or critical). 3. Requirement 11.3.1.2: Internal vulnerability scans are performed via authenticated scanning. 4. Requirement A1.2.3: The multi-tenant service provider implements processes or mechanisms for reporting and addressing suspected or confirmed security incidents and vulnerabilities. 5. Requirement A3.3.1: Failures of the following are detected, alerted, and reported in a timely manner: Automated log review mechanisms and Automated code review tools.
-------	--

Llama-2-13b	<p>What are the new requirements for vulnerability assessments in PCI DSS version 4.0?</p> <p>According to the provided document, there are several new requirements for vulnerability assessments in PCI DSS version 4.0. These include:</p> <ul style="list-style-type: none"> * Performing a targeted risk analysis to determine the frequency of log reviews for all other system components (10.4.2.1) * Detecting, alerting, and addressing failures of critical security control systems promptly (10.7.2 and 10.7.3) * Documenting, assigning, and understanding roles and responsibilities for performing activities in Requirement 11 (11.1.2) * Managing all other applicable vulnerabilities (not ranked as high-risk or critical) (11.3.1.1) * Performing internal vulnerability scans via authenticated scanning (11.3.1.2) * Supporting multi-tenant services (11.4.7) * Implementing processes or mechanisms for reporting and addressing suspected or confirmed security incidents and vulnerabilities (A1.2.3) * Detecting, alerting, and reporting failures of automated log review mechanisms and automated code review tools in a timely manner (A3.3.1)
-------------	--

Llama-2-7b	<p>What are some of the new requirements for vulnerability assessments introduced in PCI DSS version 4.0 compared to previous versions?</p> <p>According to the provided documents, one of the new requirements for vulnerability assessments introduced in PCI DSS version 4.0 is the addition of a bullet point in requirement 6.3.1 to clarify applicability to vulnerabilities for bespoke and custom and third-party software (see page 10). Additionally, there is a new requirement for maintaining an inventory of trusted keys and certificates (see page 1), which is a best practice until March 31, 2025.</p>
------------	--

Orca-2-13b	<p>¿Cuáles son las nuevas requisitos para las evaluaciones de vulnerabilidades en la versión 4.0 de PCI DSS?</p> <p>Las nuevas requisitos para las evaluaciones de vulnerabilidades en la versión 4.0 de PCI DSS incluyen resolver vulnerabilidades con una puntuación de 4.0 o más en el CVSS, realizar rescans cuando sea necesario, y tener un método de penetration testing definido, documentado e implementado que incluye enfoques aceptados por la industria y cobertura para el perímetro del SDE y los sistemas críticos.</p>
------------	--

Orca-2-7b	<p>What are some new requirements for vulnerability assessments introduced in PCI DSS version 4.0?</p> <p>Some new requirements for vulnerability assessments introduced in PCI DSS version 4.0 include:</p> <ul style="list-style-type: none"> - Performing a targeted risk analysis to determine the frequency of log reviews for all other system components (Requirement 10.4.2.1) - Ensuring failures of critical security control systems are detected, alerted, and addressed promptly (Requirement 10.7.2), and responded to promptly (Requirement 10.7.3) - Documenting, assigning, and understanding roles and responsibilities for performing activities in Requirement 11 (Requirement 11.1.2) - Managing all applicable vulnerabilities not ranked as high-risk or critical (Requirement 11.3.1.1) - Performing internal vulnerability scans via authenticated scanning (Requirement 11.3.1.2)
-----------	---

Orca-2-13b **¿Cuáles son las nuevas requisitos para las evaluaciones de vulnerabilidades en la versión 4.0 de PCI DSS?**

Las nuevas requisitos para las evaluaciones de vulnerabilidades en la versión 4.0 de PCI DSS incluyen resolver vulnerabilidades con una puntuación de 4.0 o más en el CVSS, realizar rescans cuando sea necesario, y tener un método de penetration testing definido, documentado e implementado que incluye enfoques aceptados por la industria y cobertura para el perímetro del SDE y los sistemas críticos.

Question (4): more on penetration testing

Model Name	LLM Generated Content
GPT-3.5-Turbo	<p>Can you provide more information about the changes related to penetration testing in PCI DSS version 4.0?</p> <p>I'm sorry, but I don't have access to the specific changes related to penetration testing in PCI DSS version 4.0. The information provided does not include details about the changes made to penetration testing requirements. It is recommended to review the PCI DSS version 4.0 document in its entirety to understand the changes related to penetration testing.</p>
GPT-4	<p>Can you provide more information on penetration testing in the context of PCI DSS?</p> <p>Penetration testing in the context of PCI DSS (Payment Card Industry Data Security Standard) is a requirement for entities that handle cardholder data. According to the PCI DSS v4.0 ROC Template, Requirement 11.4 states that external and internal penetration testing must be regularly performed, and any exploitable vulnerabilities and security weaknesses identified must be corrected.</p> <p>The penetration testing methodology should be defined, documented, and implemented by the entity. It should include industry-accepted penetration testing approaches and cover the entire Cardholder Data Environment (CDE) perimeter and critical systems. The testing should be conducted from both inside and outside the network.</p> <p>The testing should also validate any segmentation and scope-reduction controls. It should include application-layer penetration testing to identify, at a minimum, the vulnerabilities listed in Requirement 6.2.4. Network-layer penetration tests should encompass all components that support network functions as well as operating systems.</p> <p>If exploitable vulnerabilities and security weaknesses are found during penetration testing, they should be corrected in accordance with the entity's assessment of the risk posed by the security issue as defined in Requirement 6.3.1. The penetration testing should then be repeated to verify the corrections.</p> <p>The external penetration test should be performed by a qualified internal resource or a qualified external third-party, ensuring that organizational independence of the tester exists.</p>
Llama-2-13b	<p>What are the new requirements for penetration testing in PCI DSS version 4.0?</p> <p>Based on the provided information, there are two new requirements for penetration testing in PCI DSS version 4.0. Firstly, the penetration testing methodology must include industry-accepted penetration testing approaches, coverage for the entire CDE perimeter, and testing from both inside and outside the network. Secondly, the assessment must examine the scope of work and results from the most recent external penetration test to ensure that penetration testing is performed according to all elements specified in this requirement. Additionally, the assessor must interview personnel to verify that the external penetration test was performed by a qualified internal resource or qualified external third-party and that organizational independence of the tester exists.</p>
Llama-2-7b	<p>Could you explain what penetration testing entails in the context of PCI DSS compliance?</p> <p>Sure! Penetration testing is a crucial component of PCI DSS compliance, which involves simulating cyber attacks against an organization's computer systems and networks to identify vulnerabilities and weaknesses. The goal of penetration testing is to help organizations strengthen their defenses against attackers who might exploit these vulnerabilities to gain unauthorized access to sensitive data.</p> <p>The PCI DSS requires organizations to conduct regular penetration testing to ensure that their systems and networks are secure and protected against potential threats. This includes testing the perimeter of the cardholder data environment (CDE), critical systems, and application layers to identify any vulnerabilities that could be exploited by attackers.</p> <p>There are several key aspects of penetration testing that organizations should be aware of when implementing PCI DSS compliance:</p> <ol style="list-style-type: none">1. Defining a penetration testing methodology: Organizations must define a comprehensive approach to penetration testing that includes industry-accepted techniques and best practices.2. Coverage of the entire CDE perimeter and critical systems: Penetration testing should cover the entire perimeter of the CDE, including firewalls, routers, switches, and other network devices, as well as critical systems such as databases, servers, and applications.3. Testing from both inside and outside the network: Penetration testing should involve testing from both inside and outside the network to simulate real-world attacks and identify vulnerabilities that may not be visible from within the network.4. Validating segmentation and scope-reduction controls: Penetration testing should also validate the effectiveness of segmentation and scope-reduction controls to prevent unauthorized access to sensitive data.5. Identifying vulnerabilities in application layers: Penetration testing should include testing of application layers to identify any vulnerabilities that could be exploited by attackers.6. Correction of identified vulnerabilities: Once vulnerabilities are identified through penetration testing, they must be corrected promptly to minimize the risk of attack. <p>By understanding the importance of penetration testing in PCI DSS compliance, organizations can take proactive steps to protect their systems and networks from cyber threats and maintain the trust of their customers and stakeholders.</p>

GPT-3.5-Turbo

Can you provide more information about the changes related to penetration testing in PCI DSS version 4.0?

I'm sorry, but I don't have access to the specific changes related to penetration testing in PCI DSS version 4.0. The information provided does not include details about the changes made to penetration testing requirements. It is recommended to review the PCI DSS version 4.0 document in its entirety to understand the changes related to penetration testing.

Orca-2-13b

¿Puedes dar más detalles sobre las prácticas de prueba de penetración en PCI DSS?

Sí, puedo dar más detalles sobre las prácticas de prueba de penetración en PCI DSS. La PCI DSS requiere que se realice un penetration testing externo al menos una vez cada 12 meses y después de cualquier modificación o cambio significativo en la infraestructura o el software. El penetration testing debe ser realizado por un recurso interno calificado o un tercero de fuera de la organización calificado y debe cumplir con la metodología definida por la entidad. Además, las vulnerabilidades y debilidades de seguridad que se encuentren durante el penetration testing deben ser corregidas y verificadas mediante la repetición del testing.

Orca-2-7b

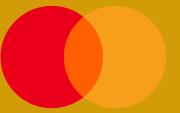
What are some best practices for conducting effective penetration tests?

Some best practices for conducting effective penetration tests include performing the testing according to a defined methodology, ensuring the tester has organizational independence, validating segmentation and scope reduction controls, testing both inside and outside the network, considering the latest threats and vulnerabilities, and documenting the approach to addressing risks identified during the testing.



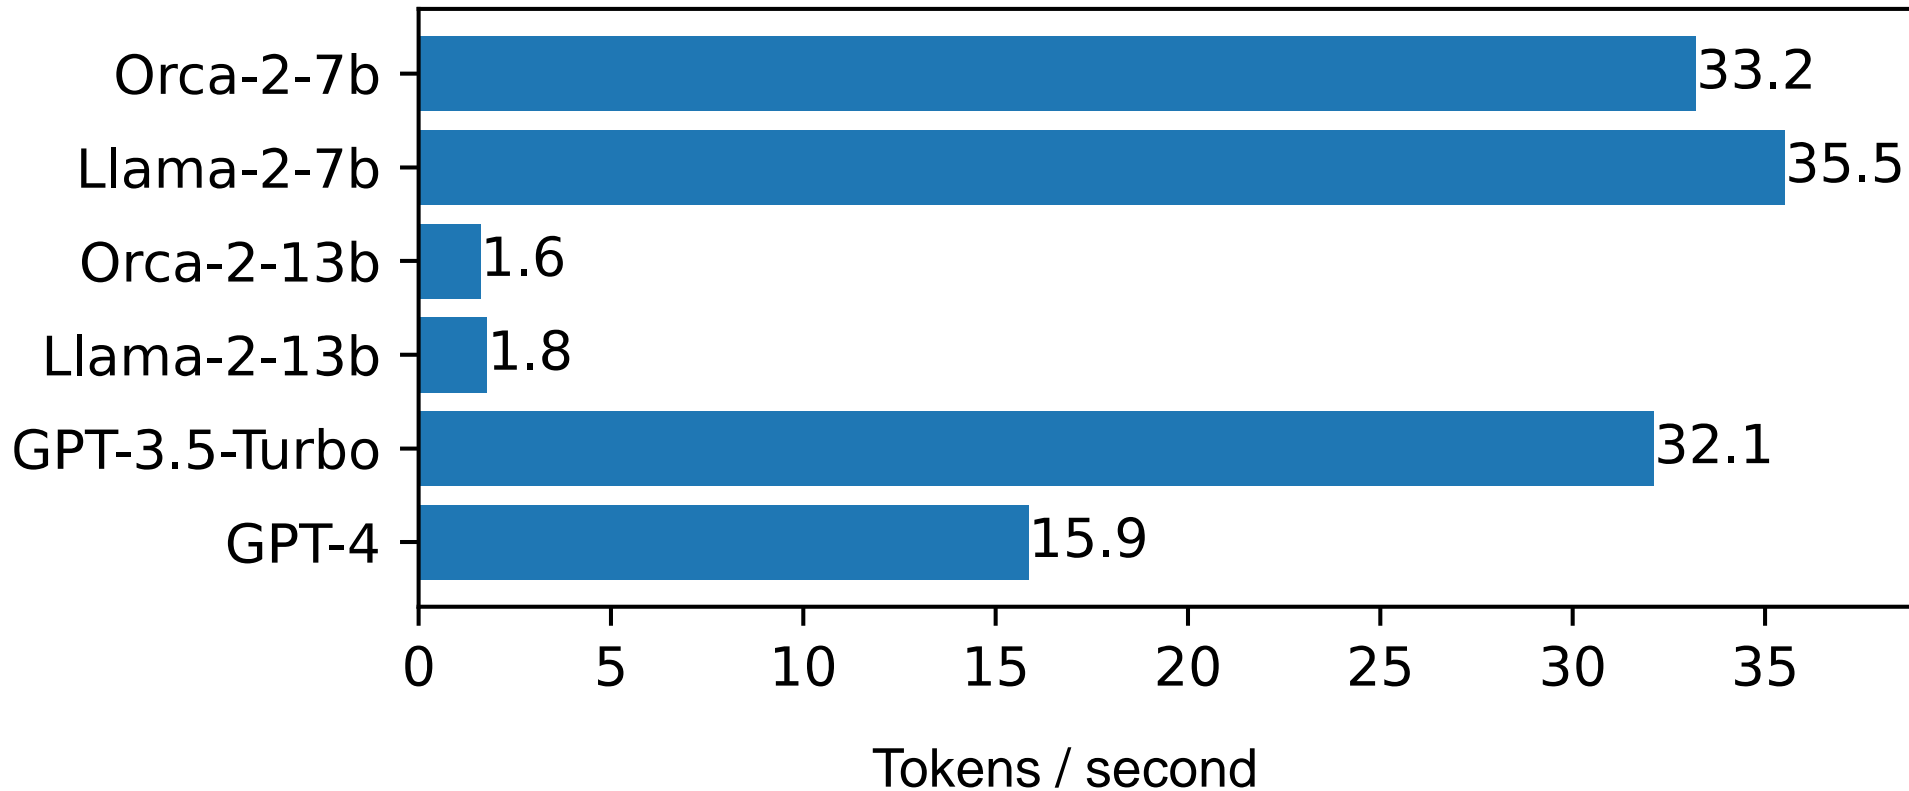
Experimental Results – Translation of Contents Generated by Orca-2-13b

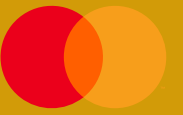
#	LLM Generated Content	Translation by GPT-4-Turbo
2	¿Puedes resumir los cambios realizados desde la versión 3.2.1 hasta la versión 4.0 de PCI DSS?	Can you summarize the changes made from version 3.2.1 to version 4.0 of PCI DSS?
	Sí, puedo resumir los cambios realizados desde la versión 3.2.1 hasta la versión 4.0 de PCI DSS. Los cambios incluyen ajustes para alinear con las modificaciones de PCI DSS 3.2 a PCI DSS 3.2.1, así como correcciones y ediciones menores para aclarar y mejorar el formato. Luego, en la versión 4.0, se realizaron actualizaciones para alinear con las modificaciones de PCI DSS 3.2.1 a PCI DSS 4.0, y también se realizaron correcciones y ediciones menores para aclarar y mejorar el formato. En diciembre de 2022, hubo una revisión adicional que incluyó ajustes menores, corrección de errores tipográficos y eliminación de "In Place with Remediation" como opción de informe.	Yes, I can summarize the changes made from version 3.2.1 to version 4.0 of PCI DSS. The changes include adjustments to align with the modifications from PCI DSS 3.2 to PCI DSS 3.2.1, as well as minor corrections and edits to clarify and improve the format. Then, in version 4.0, updates were made to align with the modifications from PCI DSS 3.2.1 to PCI DSS 4.0, and there were also minor corrections and edits made to clarify and improve the format. In December 2022, there was an additional revision that included minor adjustments, correction of typographical errors, and the removal of "In Place with Remediation" as a reporting option.
3	¿Cuáles son las nuevas requisitos para las evaluaciones de vulnerabilidades en la versión 4.0 de PCI DSS?	What are the new requirements for vulnerability assessments in PCI DSS version 4.0?
	Las nuevas requisitos para las evaluaciones de vulnerabilidades en la versión 4.0 de PCI DSS incluyen resolver vulnerabilidades con una puntuación de 4.0 o más en el CVSS, realizar rescans cuando sea necesario, y tener un método de penetration testing definido, documentado e implementado que incluye enfoques aceptados por la industria y cobertura para el perímetro del SDE y los sistemas críticos.	The new requirements for vulnerability assessments in PCI DSS version 4.0 include resolving vulnerabilities with a CVSS score of 4.0 or higher, performing rescans when necessary, and having a defined, documented, and implemented penetration testing method that includes industry-accepted approaches and coverage for the CDE perimeter and critical systems.
4	¿Puedes dar más detalles sobre las prácticas de prueba de penetración en PCI DSS?	Can you provide more details about penetration testing practices in PCI DSS?
	Sí, puedo dar más detalles sobre las prácticas de prueba de penetración en PCI DSS. La PCI DSS requiere que se realice un penetration testing externo al menos una vez cada 12 meses y después de cualquier modificación o cambio significativo en la infraestructura o el software. El penetration testing debe ser realizado por un recurso interno calificado o un tercero de fuera de la organización calificado y debe cumplir con la metodología definida por la entidad. Además, las vulnerabilidades y debilidades de seguridad que se encuentren durante el penetration testing deben ser corregidas y verificadas mediante la repetición del testing.	Yes, I can provide more details about penetration testing practices in PCI DSS. The PCI DSS requires that external penetration testing be conducted at least once every 12 months and after any significant modification or change in the infrastructure or software. The penetration testing must be conducted by a qualified internal resource or a qualified third party from outside the organization and must comply with the methodology defined by the entity. Additionally, the security vulnerabilities and weaknesses that are found during the penetration testing must be corrected and verified by repeating the testing.



Experimental Results – LLM Inference Speed

Fig. 3: Comparison of Inference Speed of LLMs





Conclusion

- Superior performance of Orca 2 in RAG, highlighting answer quality and inference speed.
- Orca 2's efficiency on consumer-grade GPUs sets new industry standards.
- Paradigm shift in conversational AI towards smaller, cost-effective models without performance compromise.
- Broader applicability of Orca 2 in various industries, enhancing AI accessibility and adaptability.



Future Work

- Development of advanced evaluation metrics for RAG systems.
- Exploration of smaller models like Microsoft's Phi-2 and Google's Gemma 2B for scalability and training efficiency.
- Investigation of Orca-2's performance in complex conversational settings.
- Incorporation of additional tools like TruLens and ARES for diverse evaluation methods.



Questions?

LinkedIn: <https://www.linkedin.com/in/huangdonghao>

Email (SMU): dh.huang.2023@engd.smu.edu.sg

Email (Mastercard): donghao.huang@mastercard.com



Donghao Huang

Vice President, R&D at Mastercard

