YRKESHÖGSKOLAN
NOVIA

Engine Product Quality Reporting with the Help of AI

Ebba Pada

Bachelor's thesis for Wärtsilä Finland

Industrial Management

Vasa 2025

BACHELOR'S THESIS

Author:	Ebba Pada
Degree Programme:	Industrial Management and Engineering, Vaasa
Specialization:	Energy Technology
Supervisor(s):	Johan Nydahl – Wärtsilä Finland
	Mikael Ehrs – Novia University of Applied Sciences

Title: Engine Product Quality Reporting With The Help of AI

Date: 12.5.2025 Number of pages: 41 Appendices: 0

Abstract

This thesis written on behalf of Wärtsilä Finland Oy, focuses on how Artificial Intelligence could be utilized to make engine-related quality reports that would make it easier to improve engine product quality. The engine-related issues are recorded in different internal databases, which makes it hard to find recurring issues when these vast databases need to be searched by hand. This could be made more efficient by implementing an AI that could analyze the databases and, based on that, make a quality report that gives a better overview of different trends it finds. The ability to find patterns is something that AI excels in and that humans can not compete with.

The method used in this thesis is a combination of qualitative and quantitative methods. The material has been collected from internal documents and internal tools like WärtsiläGPT, Wärtsila's own generative AI. All the theory has been collected from websites, news articles, and blog posts to ensure that the information is up to date.

WärtsiläGPT did not have access to the internal documents needed, and the gathering of data from internal databases was crucial. The results will present some of the responses that WärtsiläGPT provided when asked to make quality reports for different projects. The responses will be analyzed and will be discussed in the results chapter.

Due to sensitive material, the results will not be available to the public.

Language: English
Keywords: artificial intelligence, report generating, quality report

EXAMENSARBETE

Författare:	Ebba Pada
Utbildning och ort:	Produktionsekonomi, Vasa
Inriktning:	Energiteknik
Handledare:	Johan Nydahl – Wärtsilä Finland Oy
	Mikael Ehrs – Yrkeshögskolan Novia

Titel: Motorkvalitetsrapportering med hjälp av AI

Datum: 12.5.2025 Sidantal: 41 Bilagor: 0

Abstrakt

Detta examensarbete är skrivet for Wärtsilä Finland Oy och fokuserar på hur artificiell intelligens kan användas för att producera motorrelaterade kvalitetsrapporter som skulle vara till nytta för dem som jobbar med att förbättra motorprodukt kvalitén. De motorrelaterade problemen registreras i nuläget i olika interna databaser, vilket gör hela processen att manuellt hitta återkommande problem i dessa stora databaser väldigt lång. Detta skulle kunna göras mer effektivt genom att implementera AI som kan lätt analysera databaserna och utifrån detta producera en kvalitetsrapport som skulle ge en bättre överblick över olika trender den lyckats identifiera. AI har en god förmåga att upptäcka mönster i text, något som människor har svårt att konkurrera med.

Metoden som används i detta arbete är en kombination av kvalitativ och kvantitativ. Materialen har insamlats från interna dokument och genom användning av interna verktyg som WärtsiläGPT, Wärtsiläs egen generativa AI. Allt i teoridelen har samlats in från olika websidor, nyhetsartiklar och blogginlägg för att säkerställa att all information är aktuellt.

WärtsiläGPT hade inte tillgång till de interna dokuments som behövdes och därför var insamlingen av data från de olika interna databaserna en stor del av arbetet. Resultatet kommer presentera de rapporter som har genererats av WärtsiläGPT som gjort för några olika projekt, eller rättare sagt olika kraftverk med Wärtsilä motorer. Svaren kommer att analyseras och diskuteras i resultatkapitlet, i samband med svaren till frågorna ställda i introduktionen.

På grund av känsligt material så kommer resultatkapitlet inte vara tillgängligt för allmänheten.

Språk: engelska
Nyckelord: artificiell intelligens, rapport generering, kvalitetsrapport

Table of contents

Figures

1 Introduction

The use of AI has grown exponentially the last few years, not only for individuals but also for companies worldwide. Companies are exploring ways to utilize AI to maximize their profits and make everything more efficient.

According to a survey done by McKinsey, 65 percent of the respondents use AI in their organizations regularly, and this percentage has nearly doubled since their last survey done almost a year earlier (Singla, Sukharevsky, Yee, Chui, & Hall, 2024).

National University wrote about AI usage in companies, 56 % of the respondents answered that they use AI within customer service. Other frequently mentioned implementations are cybersecurity and fraud management, digital personal assistant, and inventory management. They also mentioned that people often believe that if companies implement AI, their jobs are going to be replaced, but that is not the case. When companies start to use AI, it creates more jobs than it replaces. (National University, 2024)

It seems there is an ongoing trend for companies to implement different AI solutions to maximize their profits. Wärtsilä is a company that has already succeeded in the implementation of AI and is aiming to find even more solutions in different parts of the company. They even have a department mainly for developing different AI solutions within the company.

1.1 Background

This thesis is made for the Product Quality team that is within the Quality Department at Wärtsilä Energy. The task was to find a way to gather all recorded issues on a specific engine or engine type from different internal data systems and compile them into a short quality report, with the help of an AI. I will not be making an AI, instead, I will reflect on how this could be done and if this is even possible to implement in Wärtsilä's internal system. This task was given to me by my supervisor at Wärtsilä during the autumn of 2024.

Wärtsilä Energy is focusing on delivering power plants across the globe, driven by Wärtsilä engines and other auxiliaries. The team that presented the task works with product quality improvement and receives information directly from sites regarding the engines and auxiliaries, before handing over and sometimes together with the warranty team.

Those who work with the product quality of engines usually face the challenge of finding the root cause of the problem. There are different databases where it is possible to find known issues with the engine, but the databases have thousands of entries. Finding the root cause of the specific issue the person is looking for could take hours to figure out. Instead of searching in these databases, it is usually much easier to replace the part that has broken and hope that it will fix he issue, instead of figuring out why it broke in the first place.

This tool to gather all information about an engine in a short quality report would make problem-solving and investigation much more efficient and straightforward. The recurring issues that frequently happen could be solved much earlier and even be prevented, which would be much more time and money saving. This would not only be beneficial for this team but also for other teams that are working with engine product quality.

1.2 Purpose and Problem Area

The purpose of this bachelor's thesis is to research how an AI could generate a quality report for engine-related issues for an engine type or several different power plants. The AI would go through internal databases and concentrate the requested information into a simple quality report, either an existing AI like WärtsiläGPT or a completely new one. I will be researching how an AI functions and how to make it work the way that it is intended to work. The most important test I will be executing in this thesis is to see if WärtsiläGPT would be able to make these quality reports if the right data is fed into it, then there would not be any need for a new AI. If this is possible, then quality reporting would be much easier to implement in the future.

In this thesis, I will be focusing on answering these questions:

- How does generative AI work, and is it fit for report generation?

This will be answered based on the theory chapter and my own conclusions when all the tests have been executed.

- What information will be taken from which internal database?

I will go through the internal databases that are needed in this thesis and see how they work and how to navigate them.

- Does Wärtsilä's generative AI access to the internal documents needed for this thesis?

I will be asking the WärtsiläGPT itself and find more information on this internally.

- Is WärtsiläGPT able to make quality reports if it has access to the correct data?

This will be tested in this thesis and presented in the results chapter.

- How does better prompting contribute to a more accurate answer?

During my testing, I have utilized the theory regarding prompting to get better responses from WärtsiläGPT.

The existing internal tools, platforms, and systems are a significant part of this thesis work and will be analyzed more in-depth to figure out what information will be taken from which databases. I will be testing Wärtsilä's own ChatGPT platform to figure out how much information it has access to at this moment, and if it has access to the internal documents that relate to this thesis work or not.

Wärtsilä has different internal tools, platforms, and systems where all the information about different engines is stored. The issue with these tools and systems is that they all have vast databases that feel almost impossible to go through, at least by a human being. AI is a tool that can be used to go through countless amounts of data in a short period of time, and this is something that could be implemented to handle these databases and compile the necessities more easily.

To put this to the test, I will try to search in these databases to find issues related to a specific engine type. I will make multiple quality reports with the help of WärtsiläGPT where it lists all the key issues that could be crucial to know when trying to figure out the root cause of an issue. The quality report will consist of recurring issues that happen in multiple engines, where the issue is located in the engines, and in that way give a better overview to prevent those issues in the future. There will be multiple tests with specific engine types and projects, to really see if the AI picks up these recurring issues. I will note down all the issues that I face while searching in these databases and see how long it will take to find the crucial information needed to make the reports.

1.3 Delimitation

The thesis is limited to the engines in the Wärtsilä Energy sector, all the engines that are currently in the power plants. The internal tools, platforms, and systems used for information in this thesis are going to be connected to the Wärtsilä Energy sector and the power plant within this sector. WärtsiläGPT, Wärtsiläs own generative AI, is available for all workers in the company.

1.4 Scope

The first chapter gives an overview of the background, problem, and purpose of the thesis.

The second chapter contains more in-depth information on Wärtsilä, where specific information about the internal technical information will be handled.

The third chapter gives the reader the theory behind this thesis and all the technical information that is connected to this.

The fourth chapter explains the choice of method and shows how the information has been gathered. This chapter also states how different tests have been executed.

The fifth chapter will present the results gathered during the thesis, more specifically, the WärtsiläGPT responses and other results.

The sixth chapter gives the conclusion of this thesis work, including challenges during this thesis, future research work, and a summary.

2 Wärtsilä Energy

Wärtsilä is a global multinational company founded in the year 1834 in the municipality of Tohmajärvi, which is in eastern Finland in the county of Karelia (Wärtsilä, 2024). It started as a simple sawmill and since then has grown into a multibillion company and has around 17,800 employees in 79 different countries around the world. Their net sales in 2024 were 6,449 million, and they are almost equally split between the Energy side and the Marine side, and the Portfolio business has 10 %.

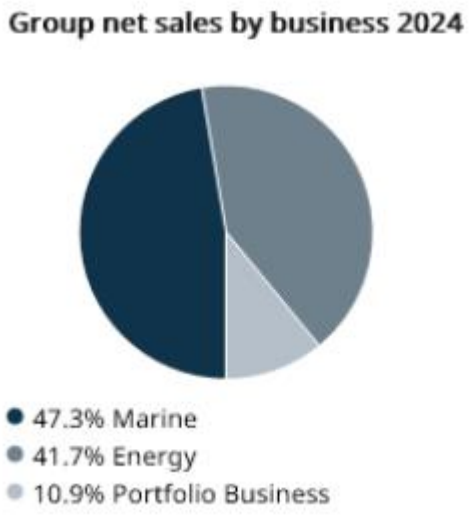

Figure 1: Wärtsilä net sales split among the different businesses (Wärtsilä, 2025)

Wärtsilä is split into two different businesses, Energy and Marine. Energy focuses on engine power plants and energy storing while also providing services for the power plants throughout the whole lifecycle. Meanwhile, the Marine business focuses on providing engines and other technologies to ships and other marine vessels and providing service during the whole lifecycle. (Wärtsilä, 2024).

At the beginning of 2024, Wärtsilä Energy had 5430 employees and accounted for 43.4% of the whole company's net sales (Wärtsilä, 2023). Wärtsilä Energy provides engine power plants around the world; they do not only provide the engine but also all the auxiliary equipment from other suppliers for the whole power plant. They have provided 79 GW of power plant capacity around the world, and each power plant consists of one or more engines.

Besides the engine power plants, they are actively working towards a transition to 100 % renewable energy. The engines used in the power plants are designed to be compatible with future fuels, which are expected to be more carbon-neutral than the fuels we have today. They also have hybrid solutions for the engines, so that they can use many different fuels and make them more flexible. (Wärtsilä, 2024).

The production of both Marine and Energy engines is located at STH, Sustainable Technology Hub, which is located in Vaskiluoto, Vaasa. Besides the production of engines, they also research sustainable fuels and test the engines. STH opened in 2022 and has become a crucial part of the company to become more sustainable. (Wärtsilä, 2022).

Figure 2: Sustainable Technology Hub, Vaasa (Wärtsilä, 2022)

Wärtsilä's Central Distribution Centre is the biggest and most important logistics center and is in Kampen, Netherlands. This logistics center covers both Marine and Energy customers and has 30 million parts in stock, and annually, they deliver more than 140,000 packages to over 180 countries. To hold all these millions of parts, the whole warehouse is 37,000 square meters, which is massive. The warehouse handles everything from order confirmation to delivering the packages to the customers. (Wärtsilä, 2024).

2.1 Non-Conformity Management

The Non-Conformity Management team is responsible for sending replacement parts to the power plants when something has broken, which could be either an engine component or an auxiliary part. Auxiliary equipment is typically bought from external suppliers, and if it breaks, it is managed through supplier contracts. On the other hand, engine-related parts are ordered from an internal warehouse, located in the Netherlands. The biggest challenge is often to figure out the root cause of the part failure, whether it is a recurring problem or a single case.

Auxiliary equipment does not have any registers where the issues are reported, as these are managed within the supplier's data systems. Our engines have many different internal registers where the issues are reported, whether they are during the manufacturing process, installation on site, or in operational use.

This team consists of six people who have permanent employment and one trainee, who manages the non-conformity cases. Most of the non-conformity cases relate to auxiliary equipment, and some of us are more focused on engine-related cases.

The whole product quality department focuses on the improvement of the engine quality, this thesis is not only for my team but also for all the other teams that deal with the engines. This team does not focus on improving the engine quality; we only report to other teams that something has broken, and we send the requested parts, or figure out which part we need to send. If the root causes of the engine issues were easier to figure out, there would not be nearly as many engine-related reports as it is today.

The reports from projects are sent to our shared mailbox, those reports tell us what the issue is and usually what part it could be. The number of reports in the mailbox varies every day; one day, it could be around 30, and the next day, less than 10. The process of handling the case varies as well; sometimes it is done in a day, and others could take up to a year. Usually, the auxiliary cases are longer when the suppliers are involved.

3 Theory

This chapter will cover the relevant theory of this thesis. The purpose of this chapter is to get a better understanding of the different aspects related to this topic, as well as background information on how Generative AI works and how it could be utilized. The chapter will provide information regarding prompts that will benefit the testing phase in this thesis.

3.1 Important Terms to Know Regarding AI

Before entering the rest of the theory chapter, different terms need to be explained. These terms are used frequently in this chapter, and it will help the reader to know these terms before reading the whole chapter.

3.1.1 AI Algorithm

To start off this chapter, the Artificial Intelligence Board of America had a well-written explanation of how algorithms work:

"Algorithms are the basic building blocks that allow machines to replicate what can be defined in part by human intelligence and perform tasks independently. These algorithms work by receiving and analyzing inputs and producing outputs based on them. Their success depends on the compatibility of input data and the models upon which their functioning is based." (Artificial Intelligence Board of America, 2024)

The algorithm needs training data to learn, and this data depends on what type of AI algorithm is being made. The algorithm uses training data to complete tasks, and in that way, it learns and grows. Some AI algorithms learn on their own, and others learn by having programmers streamline the process. There are three key categories of AI algorithms; supervised, unsupervised, and reinforcement learning. (Tableau, 2024).

Supervised learning is the most used method for learning AI algorithms, and it is based on clearly labeled data. The supervised learning part is literally what it says, a team of experts supervise the algorithm to make sure it gives the right results, and they test the models to ensure accuracy. (Tableau, 2024).

Unsupervised learning is the opposite of supervised learning; instead of labeled data, the data is unlabeled. The algorithms use the data to evaluate relationships between data points to get a better understanding of the data. (Tableau, 2024).

Reinforcement learning works by getting feedback for actions, usually called a reward, that could be either positive or negative. The website described it like this:

"*A reinforcement algorithm is usually composed of two major parts: an agent that performs an action, and the environment in which the action is performed. The cycle begins when the environment sends a "state" signal to the agent. That queues the agent to perform a specific action within the environment. Once the action is performed, the environment sends a "reward" signal to the agent, informing it on what happened, so the agent can update and evaluate its last action. Then, with that new information, it can take the action again. That cycle repeats until the environment sends a termination signal."* (Tableau, 2024).

AI algorithms are a complex topic, and there are countless applications where this could be used. This was the basics needed to have a better understanding of the theory chapter. (Tableau, 2024).

3.1.2 Machine Learning and Deep Learning

The theory in this chapter has been taken from an article released by Google Cloud named "What's the difference between deep learning, machine learning, and artificial intelligence?", which had all the relevant theory needed for this chapter.

The terms AI and machine learning are usually discussed like they are the same thing; in fact, they are correlated, but they are not the same thing. AI is a broad subject, and machine learning is the part that makes AI act like a human, like learning by doing and responding in the same way as a human would.

"Machine learning is a subset of artificial intelligence that enables a system to autonomously learn and improve without being explicitly programmed. Machine learning algorithms work by recognizing patterns and data and making predictions when new data is inputted into the system." (Google Cloud, 2024).

There are many different machine learning models, but this thesis will mainly focus on the deep learning model. Deep learning is a branch of machine learning that processes and analyzes information with the help of ANN, short for Artificial Neural Networks. These neural networks use computational nodes with different layers within the deep learning algorithms. These layers each contain an input layer, an output layer, and a hidden layer. To help the algorithm learn and improve its accuracy, the neural network needs to be fed training data. When a neural network contains three or more layers, it is considered deep, and that is where deep learning is gets its name. (Google Cloud, 2024).

Synoptek had made a perfect illustration of how Artificial Intelligence, Machine Learning, Deep Learning, and Generative AI are connected. These are key elements of the generative AI that we have today. (Purohit, 2023).

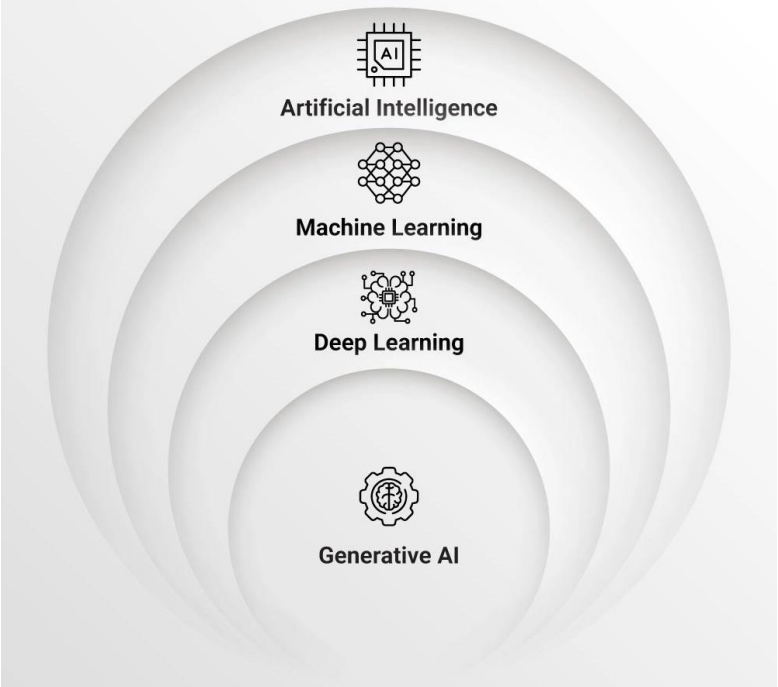

Figure 3: How AI, ML, DL, and GenAI are connected (Purohit, 2023)

3.2 Generative AI

Generative AI is a type of AI that can be used for generating text, audio, images, videos, and other codes. To get a generative AI to work, it needs to have a special machine learning model called deep learning, which makes the AI identify patterns in vast amounts of data and helps it use that information to make a response. To get a generative AI to work, it needs training, it needs the data that it is going to use, and then the algorithm. The foundation model is made with a deep learning algorithm on a vast amount of data, and this is when the training begins. The algorithm tries to predict the next word in a sentence without deviating from the actual data that has been presented. It has given feedback based on its responses, and then it will improve over time. When the training is done, the AI can be fed prompts or other inputs to get the desired response. (Stryker & Scapicchio, 2024).

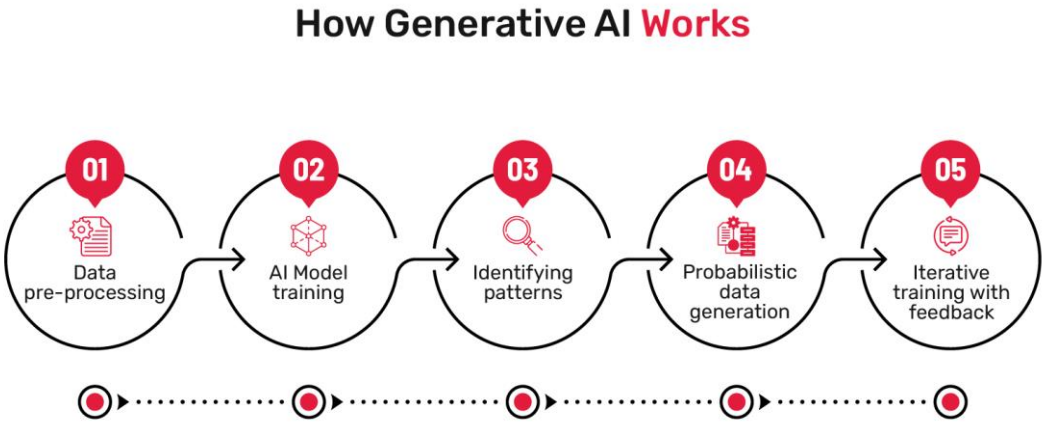

Figure 4: How generative AI is trained step by step (Jaro Education, 2024)

The most popular AI that uses this generative AI is the ChatGPT robot, which almost everyone has tried out at least once. Here I asked ChatGPT an easy question and it provided me with the correct information, which was found on the internet:

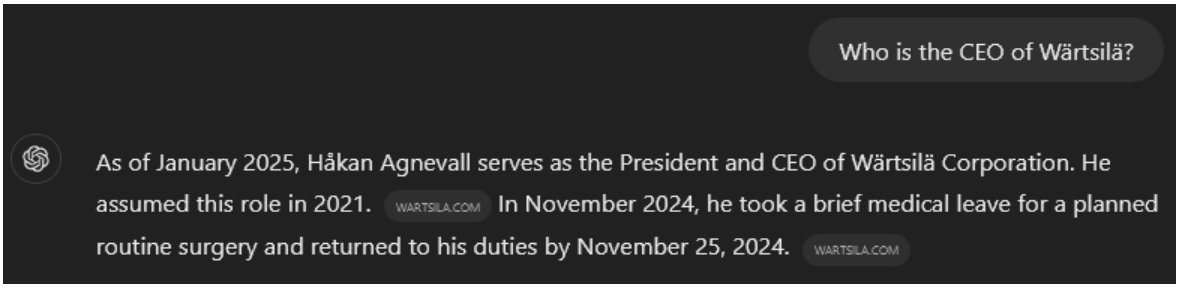

Figure 5: ChatGPT response on a basic prompt (ChatGPT, 2024)

When talking about generative AI, they often have a limit on how big of an input it can comprehend. The limit is often stated in tokens, which are like the number of words but not the same thing. OpenAI had a good list explaining the rules of thumb regarding tokens:

- 1 token ~= 4 chars in English
- 1 token ~= ¾ words
- 100 tokens ~= 75 words

Or

- 1-2 sentence ~= 30 tokens
- 1 paragraph ~= 100 tokens
- 1,500 words ~= 2048 tokens

Figure 6: How tokens work with generative AI (OpenAI, 2025)

The tokens are also language-dependent because the word splits are different in different languages. The best way to minimize the amount of tokens is to touch up the prompt or make the text shorter. ChatGPT has an input limit of 128,000 tokens, depending on the model. (OpenAI, 2025). The token limits differ from different AI models, and the biggest one right now is the Google Gemini 1.5 Pro, which has a limit of 1 million tokens (Pichai & Hassabis, 2024).

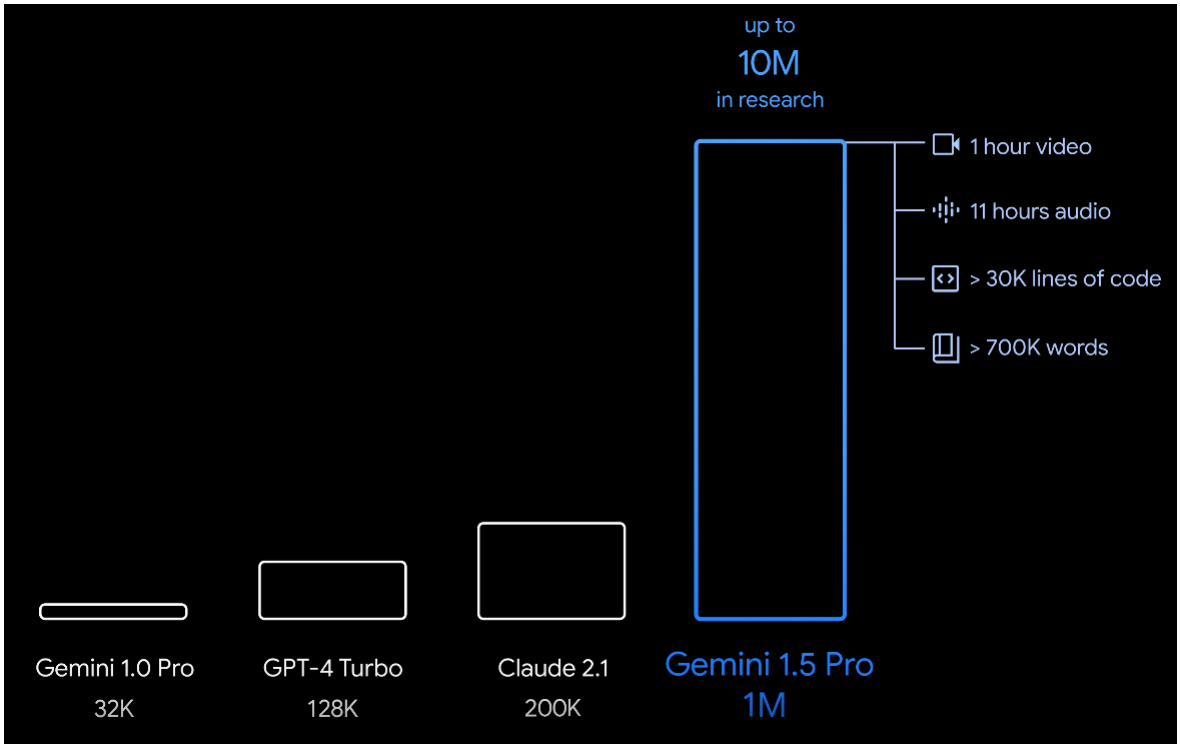

Figure 7: Google Gemini 1.5 Pro token limit compared to other AI models (Pichai & Hassabis, 2024)

3.2.1 Prompting

When a user is asking a question to an AI, it is called a prompt. Prompts could be everything from an easy question about a historical person to having it generate a complex code for a coding program. Generative AI usually learns by doing and will understand better the more it is trained, but it is still important to know how to make prompts more effective. Generative AI could sometimes be inaccurate and provide the wrong answer, and that is why it is important to be very specific about what you want as a response. Harvard University had a good article about this and had a great example of this. Instead of writing "Write a story", it should be more specific, like what genre it is, who is the audience is, whether it should be funny or serious, and how long it should be. The more specific the prompt is, the more accurate response will be generated. (Harvard University, 2023).

Figure 8: Examples of good vs bad prompts (Yudovski, 2024)

Even though the prompts should be specific and detailed, they should not be too long. Usually, the AI does not have a hard limit on the amount of characters in the input, but it could still struggle if the prompt is exceedingly long. These kinds of prompts make it hard for the AI to differentiate which parts of the prompts prioritize more. This could be easily solved by writing the prompt in a step-by-step format, which ensures that the AI understands each step of the requested job it needs to do to get the correct output. With this method, it is important to think about not exceeding the token limits, in other words, try not to make too many instructions. Generative AI is also sensitive to the choice of words, usually, it is not recommended to use any slang or metaphors. This differs between different AI models, but to be sure, it is always safer to use direct and clear words. (Bigelow S. J., 2025).

To ask the AI to behave like something, like a specific type of person, object, or role, will make the AI start acting in the role and give answers more accurately. A great example of this is to ask the AI to explain something to you, acting like an elementary school teacher, then the response will be easy to understand and not too complex. It is also possible to have the AI explain something to you, and have it written so that a 7th grader would understand it, and then the information would be much easier to take in. It is also important to tell the AI in what format you want your answer, wheather you would like a picture, story, report, code, or audio, etc. To use “do” and “don’t” is an important thing when prompting, if you want something included, use “do”, and if you think something is not needed in the prompt, use “don’t”. It is possible to give the AI examples of how the output should look like, if a report should be generated, give an example report that has the layout that is desired. When giving examples to an AI, it is important not to give it any copyrighted works. (Harvard University, 2023).

When making a prompt, it does not have to be perfect the first time, it is always possible to improve it over time by asking multiple questions and adding things to it. If the AI gave a response that was not correct, it should be corrected and told to fix it. Generative AI learns over time, and if it gets feedback on its work, it will improve and give better results. (Harvard University, 2023).

3.2.2 Companies That Use Generative AI

All theory in this chapter has been taken from a blog posted by April Miller on the website Indatalabs, the name of the post is “15 companies using generative AI for business efficiency”. This blog post had some interesting examples of how different companies already use generative AI to ensure efficiency.

There are many companies that already use generative AI solutions in different aspects of the company. In the introduction, it was mentioned that the most common AI that companies use is chatbots used in customer service. Almost every company’s home page has one of these chatbots; sometimes they are helpful, and sometimes, they do not provide any help at all.

Mastercard is a good example of a company that uses AI to help with customer service. Besides the chatbot, they also use generative AI to prevent fraud, the AI learns the customers' behaviors, and if their behavior changes in some way, they will let someone in the company know about this unusual behavior.

Microsoft uses generative AI in their search engine Bing, in that way, the search engine could answer the question that the user has searched for directly, using multiple results for a more accurate answer. They also integrated their AI Copilot into the Windows 11 operating system as a virtual assistant to help with different tasks.

Toyota has an interesting implementation of generative AI; they use it to design new car models. Their generative AI considers different factors like ergonomics, handling, and safety to generate a car design with these prompts. The engineers use prompts to request a car design with specific qualifications, making the process of coming up with new designs a lot smoother and easier. (Miller, 2024).

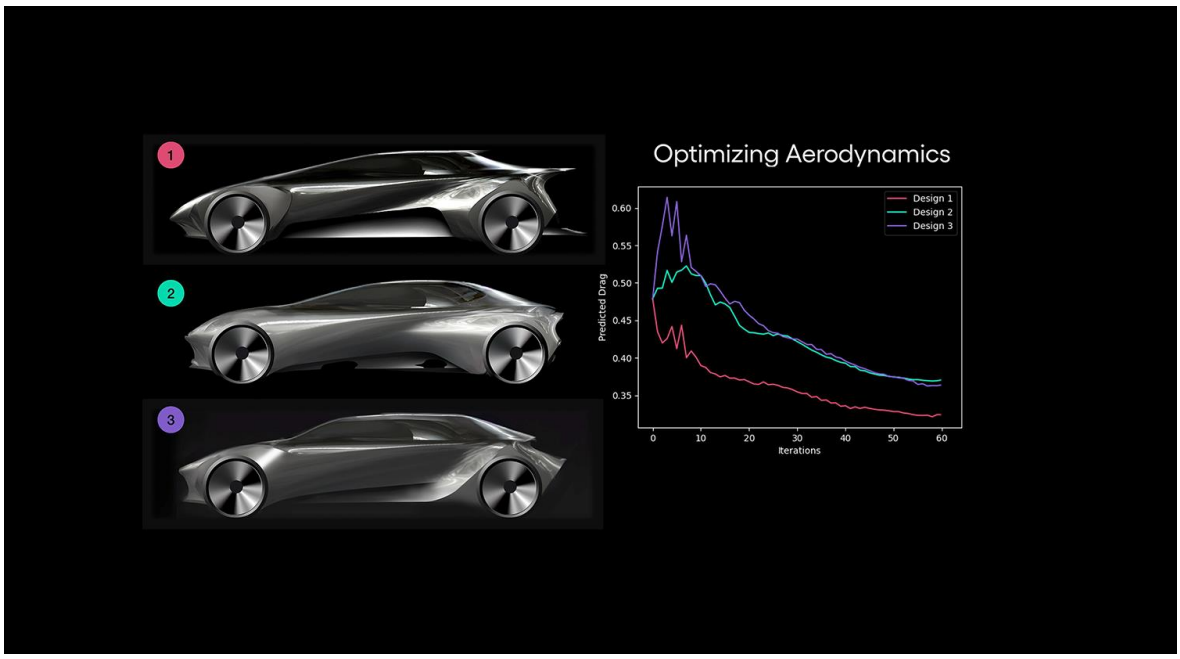

Figure 9: AI-optimized car model from Toyota website (Bigelow P. , 2023)

3.2.3 Report Generating with Generative AI

Generative AI can be used to make texts, and a type of text is reports, so the AI can make reports from the data that it has access to. To make this work as intended, it is crucial that the deep learning algorithm has access to the correct data and that it has had the correct training to provide reliable responses. When the algorithm has been fully trained, the user can provide the AI with a prompt that requests a report on a specific subject, and the AI will provide a report based on the data that it has learned. The AI will not only be able to make reports but also make improvements to existing texts and make graphs of the data provided. (Barbarossa & Filisetti, 2024).

This type of generative AI is already used in financial sectors; they use AI to generate financial reports that used to take hours upon hours to make and now are made in a matter of seconds. Also, in healthcare hospitals have implemented generative AI to make reports of patient data to make it easier for healthcare professionals to pinpoint illnesses and, in that way, get better treatment for the patients. (Kaur, 2024).

3.2.4 Issues with Generative AI in general

There are many things that are great about generative AI, but there are also issues with it. The most common issue with the generative AI is its hallucination problem, as it does not know how to stop giving an answer, even though it does not know the answer. Then it could make things up instead of admitting that it does not know, which is why it is really important to train the AI properly. (IBM, 2023).

Bias

The most popular generative AIs are trained on vast amounts of public online texts, which means that the AI has the same bias as those texts. An AI cannot have the same judgement as a human, and that leads to biased answers based on the information that it is trained on. It is hard to pinpoint exactly where bias originates from, it could be the algorithm, the data, or human reviewers who provide feedback in the testing phase. If an AI is biased, it could provide misleading information or hallucinations. It could also be demographic, ideological, cultural, or human-biased, which is not ideal for an AI that should be neutral in these kinds of topics. (The University of Kansas, 2024).

A good example of how bias in AI could end horribly is the AI recruitment tool that Amazon used which was biased against women. Amazon wanted to automate their hiring process by having an AI go through the resumes, and it would spit out the top 5 candidates most suitable for the position. The issue with this was that the data that it had been trained on was resumes that had been submitted over a 10-year period, that was mostly submitted by men. When the AI started working, it penalized the resumes that contained the word "Women's", which led to those male candidates becoming preferable. This AI recruitment tool was scrapped when they realized that it was biased against women, and they never used this kind of system again. (Dastin, 2018).

Explainability

When the generative AI provides an answer to a prompt, it sometimes does not explain how it came to that answer, this is commonly called a "black box". The engineers and data scientists cannot even explain how the AI concluded, or they do not understand what is happening inside that black box. That is why it is important to train the AI correctly to make sure that it can provide an explanation behind its answer. (IBM, 2024).

Data privacy

Generative AI has access to vast amounts of data, and if this data consists of some sensitive information, then the AI has access to that. This could result in unauthorized access, exposure, or even privacy breaches. This concerns both personal private details and internal company details. (Thejus, 2024).

This article gave a good example of this; if ChatGPT had access to Facebook and someone searched for information about an individual, then a lot of personal details would be revealed. On the other hand, the information fed into AI by users could also pose a risk of exposing private details if the AI is not designed to protect that sort of information. The worst-case scenario would be that these risks could potentially lead to identity theft. This is already happening with the help of deepfakes, which is the generative AI's ability to make realistic pictures, videos, or sounds depicting an individual. (Thejus, 2024).

In companies, it has been reported that generative AI has caused potential data breaches and privacy violations. Therefore, it is important to implement measures to protect sensitive information, for example, minimizing data, making the data anonymous, obtaining consent from users, and monitoring the generative AI. (Thejus, 2024).

3.2.5 Prompts for Minimizing Risk of Hallucinations

I found this article named “Three Prompt Engineering Methods to Reduce Hallucinations” released by Prompthub that explained well how to minimize AI hallucinations by using different methods while prompting.

Earlier in this chapter, I mentioned different methods to get more accurate answers, but despite these, there is still the risk of AI hallucinations. The hallucinations are usually issues that happen due to how they are trained, but with better prompting, it is possible to minimize the risk.

The first method that researchers have found to give 20 % more accurate responses is to include “According to…” in prompts. This means that, if the AI is forced to take information from a trusted source, it is more likely to give a more accurate response. This first method is the easiest of the methods that was found.

The second method is called Chain-of-Verification prompting, which means that verification prompts are made, and then the final output is done. Firstly, the original prompt is made, and when the answers are allocated, then verification prompts are made. When these verification prompts are done, then the original question will be answered again, and this time correctly. The website Prompthub had a great example of this:

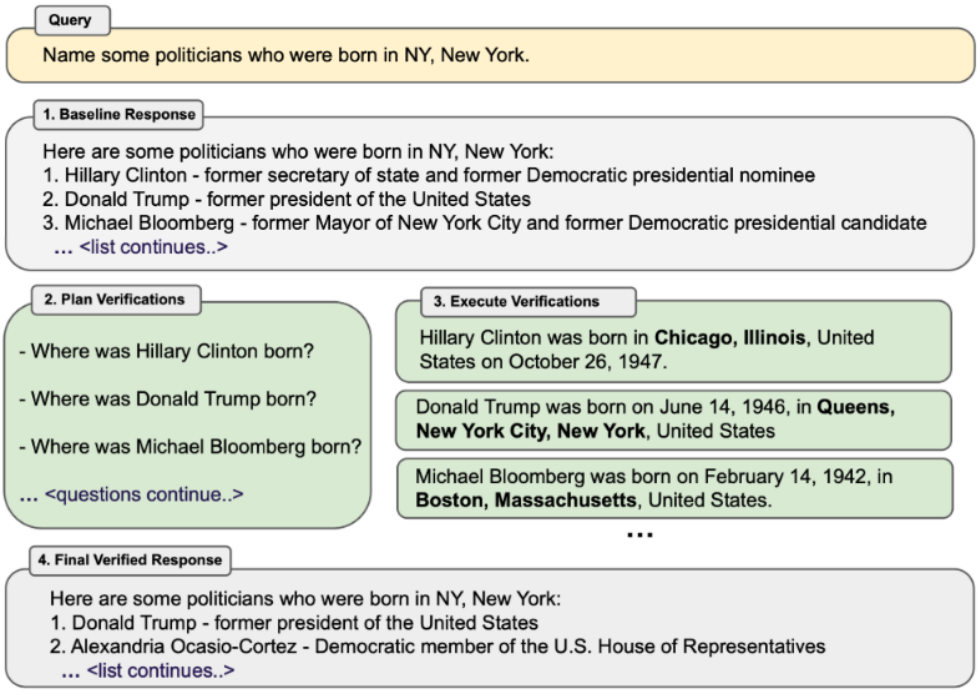

Figure 10: Example of Chain-of-Verification prompting (Cleary, 2025)

The third and last method is the Chain of Thought, and that is to make the AI think through the problem before answering. Adding “think through this task step-by-step” at the end of the prompt could help generate a more accurate response. Step-back prompting is like the Chain of Thought, but could generate a more accurate response. The website gave an example of this:

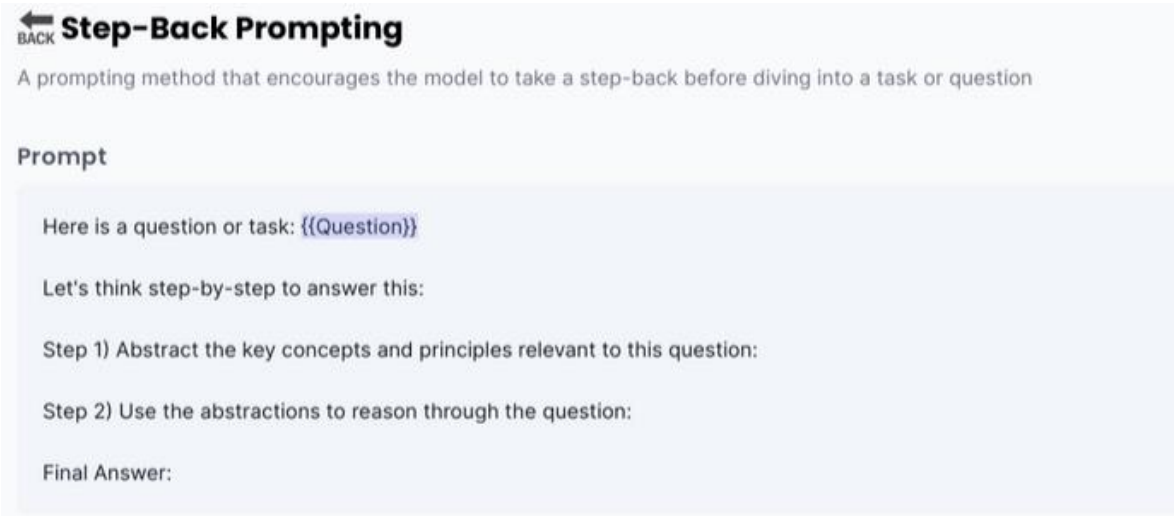

Figure 11: Example of Step-back prompting (Cleary, 2025)

These were some methods that could be useful for getting more accurate responses and minimizing the risk of AI hallucination. The biggest cause of AI hallucinations is often how it is programmed or that it is not trained enough, but with better prompts, it could still minimize the risk. (Cleary, 2025).

4 Methods

This chapter will show how the material and information have been gathered to present the results in Chapter 5.

4.1 Choice of Methods

There are many kinds of research methods, but the most used are qualitative and quantitative. Quantitative focuses mainly on numbers and graphs, and qualitative focuses on words and meaning. The quantitative method consists of experiments or surveys, using more and more physical approaches. I could have done a survey asking employees in Wärtsilä about this reporting tool, but this type of research does not necessarily need to be done to get the result that I am searching for.

So, for this thesis, I will be using a combination of qualitative and quantitative research methods. The reasoning is that the goal for this thesis is to get a more in-depth knowledge about the topic, and the information is mainly from internet articles, websites, and discussions with people working with engine quality. The use of internet articles and websites as my source of information is to ensure that I get the latest information about AI, which is a relatively new invention. Besides a lot of information gathering, I will test the WärtsiläGPT and its ability to make reports on the information fed into it. To test the WärtsiläGPT I will be searching in these databases and find the right documents that need to be included in a quality report for a specific engine or engine type, with the right prompts, the AI will generate a report.

4.2 WärtsiläGPT Testing

I will be testing the WärtsiläGPT to make these quality reports to see how much it could handle and what the reports could look like. I will be starting with only factory issues and gradually add more information along the way. Internal Wärtsilä documents have different security levels, and it is strict which of these are allowed to be used in WärtsiläGPT. DCM365 is internally available to all Employees in Wärtsilä, and that is the reason why the data is taken from there.

4.3 Databases

Wärtsilä has different databases where the information on the issues is recorded, and they are all linked in different ways. In this chapter, I will be explaining how they are connected to each other and what information is found in these databases.

4.3.1 SAP

SAP is an ERP, Enterprise Resource Planning system that helps companies manage their data and makes it easier for companies to share information between different business functions (SAP, 2024). SAP is being used in many different business areas in Wärtsilä, and there are many kinds of uses.

The non-conformity team uses SAP to record the issues that have been reported on the site, and this is also where the documents are being made and printed, which are needed in the work. All the engine and auxiliary-related issues that have been recorded in SAP may be found in DCM365, due to DCM365 consisting of the original reports that were later put into SAP. All the issues used in the WärtsiläGPT testing are taken from DCM365 due to SAP entries being already recorded there.

4.3.2 Wärtsilä DCM365

DCM365 is an internal database for the different project stakeholders to use, where they report issues that happen on-site or where they find issues with the engine and auxiliary equipment. All the stakeholders can modify reports and update information related to projects, and in this way keep good communication with each other. All the issues recorded in DCM365 are put into SAP, and that means that all things in SAP are in DCM365. (Wärtsilä Corporation, 2017).

The program can export all the data for one project into Excel, which will help with the testing of WärtsiläGPT. Every project has its own list of data in DCM365, which means that the list depends on which project has been searched for. This makes it easier to find data for one project at a time.

Non-conformity team gets the reports of the issues through DCM365, and then an SAP notification is made, and the notification number is put into DCM365. This makes it easier for the site or project team to keep up with what the issue is and how it is getting resolved.

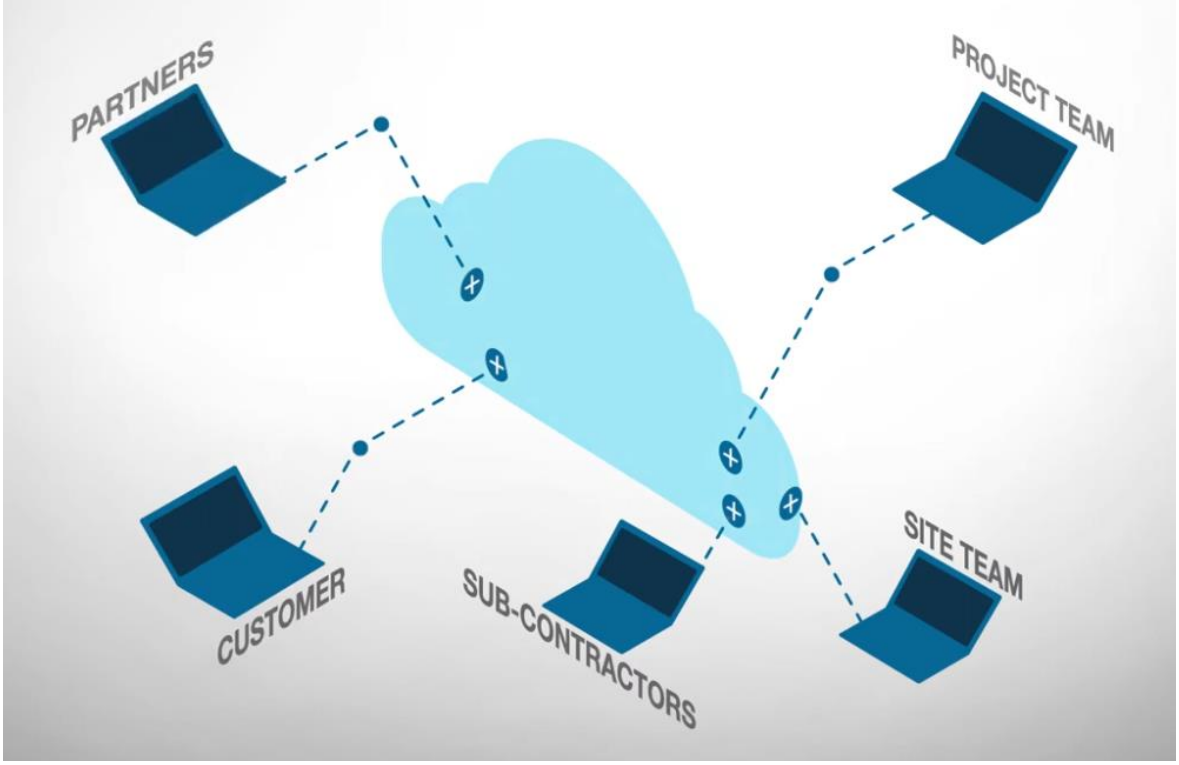

Figure 12: Screenshot describing how DCM365 works (Wärtsilä Corporation, 2017)

4.3.3 Issue register

Issue register is a database relying on Excel, with limited access. This database contains all the issues that have been recorded from when the engine was manufactured. Most of the issues in the Issue register may be found in DCM365, but only the information that seems necessary for the site personnel to know. To gain access to the Issue register, there is a need of a personal invitation from those working with the Issue register. Due to the need for a personal invitation, it seemed safer to only take the information from DCM365.

4.4 Discussions with employees

During this thesis, I have had discussions with those in my team involved with engine quality to get a better understanding of why this would benefit them in their work. This is something that they had thought about for a longer period.

When starting with this thesis, I was completely new to Issue register, due to having not worked that much with engine-related non-conformities. I got some help with navigating the system and what information was stored in this database. I also learned that most of the data here is also stored in DCM365. During this, I also got access to the Issue register after requesting the entry, to be able to research some more.

Before starting the WärtsiläGPT testing, I had to make sure that the information that would be fed into the AI was okay to use. I mentioned in the theory chapter that AI learns from the data that it is being fed, so it is important to know what can be used or not. WärtsiläGPT does have terms and conditions on how to use the AI, but to be one hundred percent sure, I asked an employee working with WärtsiläGPT.

To get started with WärtsiläGPT testing, I needed some help with completely understanding how DCM365 works, even though I used it pretty much every day at work. There are not only non-conformity cases stored there, but there are also different factory issues and other issues recorded there as well. I was taught how to filter the database and find what was needed for the engine-related quality reports.

4.4.1 Factory issues

Before starting the testing, I had to check if WärtsiläGPT would happen to have access to the internal databases that would be needed for these reports. I asked some colleagues if they knew something about it, and they said it did not have access to internal documents. I checked the roadmap to see if the team behind WärtsiläGPT had any plans on implementing the access to internal documents, but could not find anything about that. I also asked WärtsiläGPT itself, and I got this response:

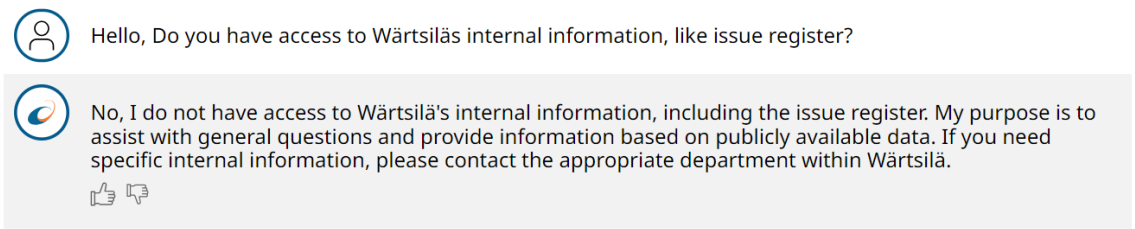

Figure 13: WärtsiläGPT internal document access

To test the WärtsiläGPT I gathered the information that should be analyzed by the AI, and I chose specific projects that used the same types of engines. To begin with, I only gathered the factory issues for different projects that could be found in DCM365. Here I searched for the project using its project number. I put the project number in the search field and found the project site.

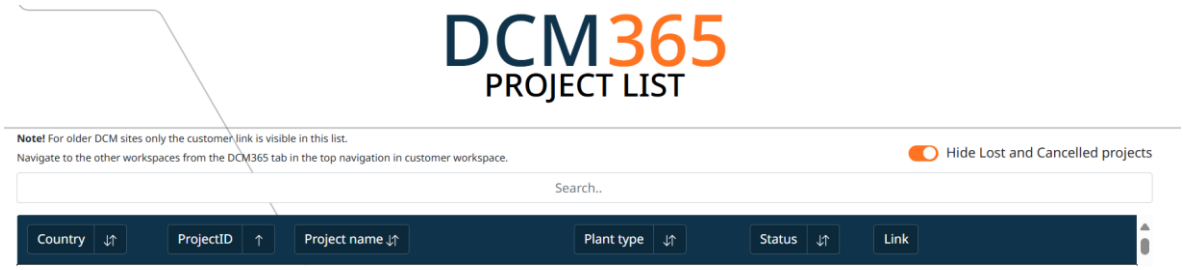

Figure 14: DCM365 search field

Then I located its scope where all the entries that I needed were recorded. The database looked like this:

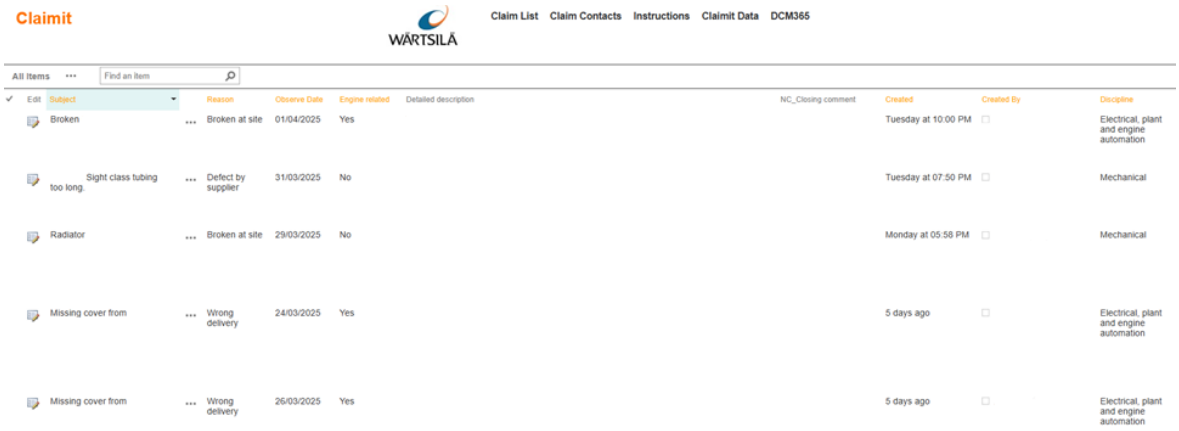

Figure 15: Example screenshot of how DCM365 looks like

In this first test, I only gathered information from four different projects, located in the same country and had the same engine types. In DCM365 I found the export to Excel button and then converted the data to Excel, where I filtered the tables so that only engine-related and factory issues were present. All the exported data from DCM365 includes all the raw data, and both engine and auxiliary related issues are exported. I filtered the information by only allowing entries that had “yes” in the engine-related field by clicking the drop-down menu in the column. Then I removed all the columns except the subject, reason, observed date, engine related, detailed description, closing comments, and serial/engine number. In the reason field, I filtered that it only showed the factory issue register, an issue recorded

during the assembly of the engine in the factory, by using the drop-down menu with the filters.

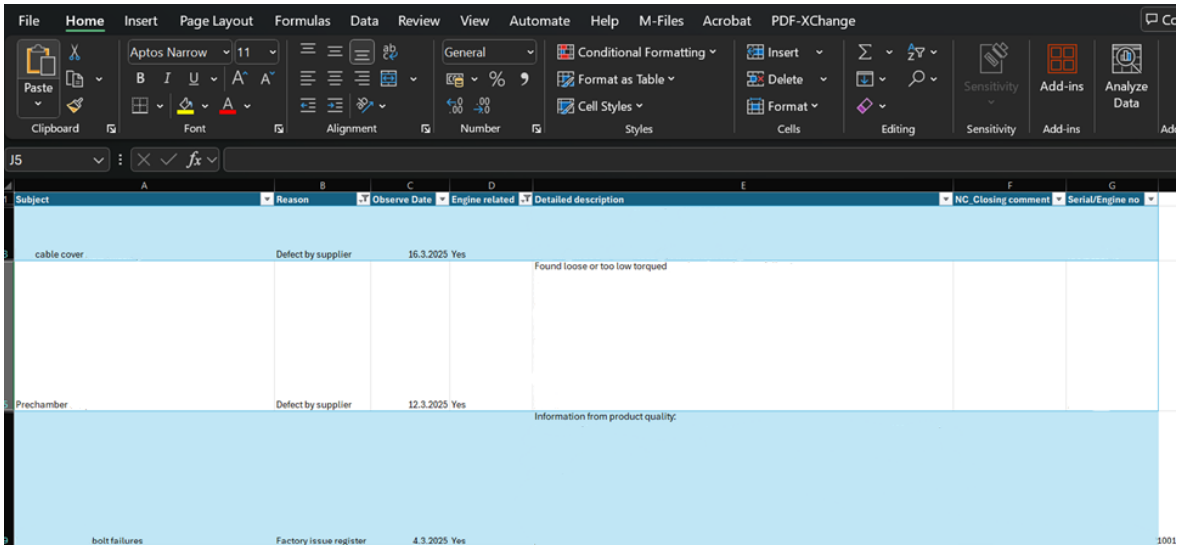

Figure 16: Screenshot of how the Excel tables looked like

WärtsiläGPT is only able to analyze Word, text, and PDF documents, and the easiest choice here was a Word document. So, I copied the whole table, and copy-pasted it to Word to make it readable by the AI, here I tried to save the Excel as a .cvs but I concluded that it took the same amount of time as putting it into Word. I am more comfortable with Word than other file types, so the choice was easy. I finalized the Word document by removing the column that was named engine related, to make the document somewhat smaller.

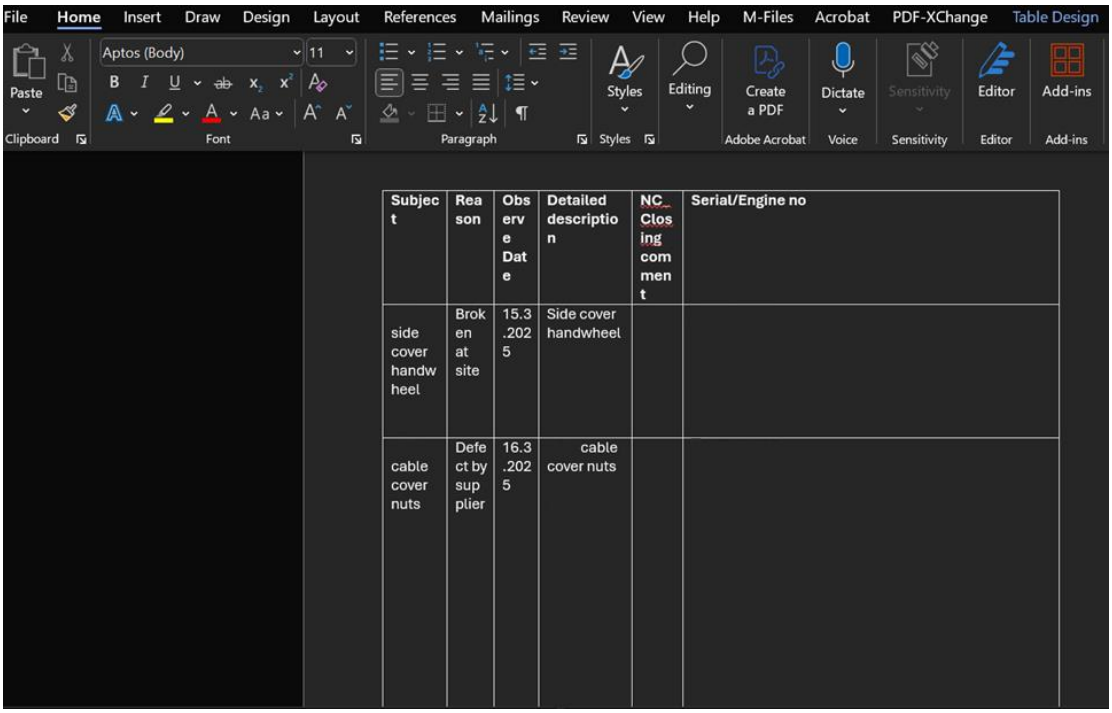

Figure 17: Screenshot of how the Word documents looked like

I also made a different Word document that included each project and how many engines the projects have. This process for all the different projects took about two hours. I was not really used to the system, and several times I had put on the wrong filter or forgotten to put on the right filter. After several tries, I got the finalized Word documents needed, where I deleted the column stating if it was an engine-related issue or not, to make the document more compact.

I started making a quality report for one project, and after that, I made a separate report for the second project. Here is the prompt that I started with:

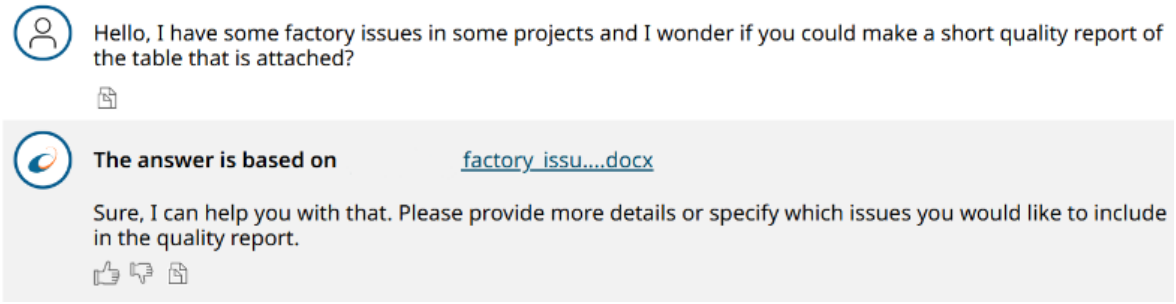

Figure 18: The starting prompt and answer from WärtsiläGPT

When I got the answer, I started to plan how I would formulate my prompt. In the theory chapter, I mentioned some ways to improve prompts, and during my testing, I have utilized the method of making the prompts more specific, specifying exactly that a report is needed and what should be included or not. Also, the prompt was not perfect the first time, and I improved it over time. I added new restrictions to the prompt when I realized that the answer did not include everything that I aimed for. I developed my prompt by trial and error. In the beginning, I forgot to mention that I needed the locations where the issues had occurred, and it gave the report in bullet point format, so I mentioned the need for locations in the engine and to write it in a running text. Then it had some issues with only listing almost all the issues and did not give any more detail, so I added that it should mention if something is a recurring issue, it should be checked. When I thought it was finished, I noticed that all the dates and engine numbers were listed in the report, and that made it too long, and that is why I restricted it from adding them to the text.

After developing the prompt several times, I finally ended up with this prompt:

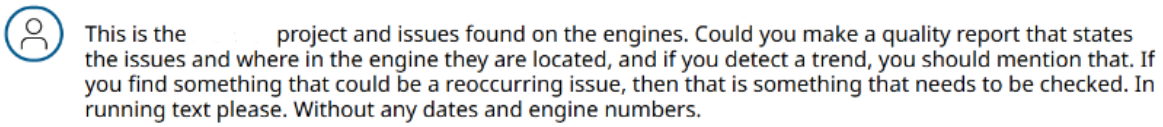

Figure 19: Prompt for quality report for one project

After the first successful report, after some back-and-forth fixing, I also made an individual report for the second project with the same prompt as the previous one. When I had the first two reports, I asked AI to summarize these reports into a combined quality report. I wanted to find similarities between the two projects, and here I also added to not mention dates or engine numbers and to write it in a running text. After this, the question looked like this:

Figure 20: Prompt for the quality report on two different projects

All the individual reports were made with the same prompt, except that I changed the document and the name of the project in the prompt. When all four projects had quality reports, I asked the AI to summarize all these reports into one final quality report, using this prompt:

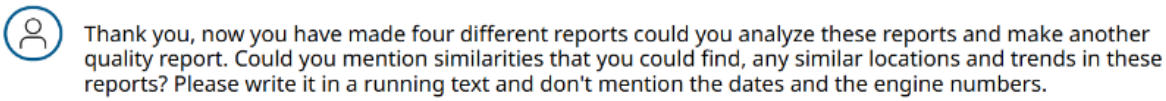

Figure 21: Prompt for summarizing multiple projects into one final quality report

When the report was finalized, I added the document including how many engines are in each project, to get the percentages of engines that had these different issues, with this prompt:

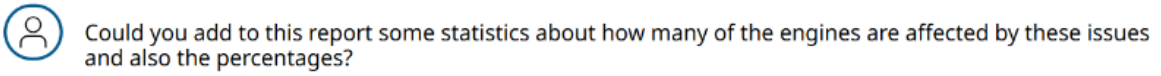

Figure 22: Prompt for adding statistics to finalized quality report

When this method was tested, I emptied the chat and made a prompt for making a report directly of the four projects. I added all the documents, including the issues from all the projects, into one document, and the prompt looked like this:

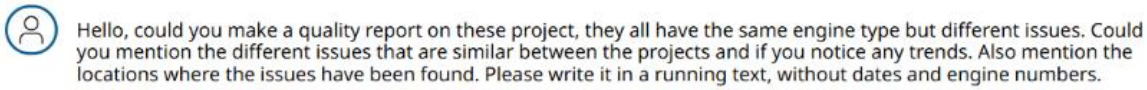

Figure 23: Prompt for making quality reports for multiple projects at the same time

When the AI was done generating the answer, it mentioned that it was not able to make quality reports, which was strange because it had already made multiple reports. This was when I realized that the WärtsiläGPT is not able to analyze documents due to technical limitations, something the team behind WärtsiläGPT is still developing. It also has a token limit of 40 000, and this big document should still be under that limit, so the limit was not the issue here. The best explanation of what tokens mean is that 100 tokens equal about 75 words, as mentioned in the theory chapter in Figure 3.

After these tests, it was concluded that WärtsiläGPT can make quality reports with the right prompts and information, but if all the information is fed into it at the same time, it gets overwhelmed. If WärtsiläGPT would have access to these internal databases, it would be much easier to get these quality reports. This process of gathering the information and then getting the right prompt to get the quality report took about 4 hours until I had all the reports for these factory issue register entries.

4.4.2 All the Issues Recorded

In this test, I used the same Excel documents that I used in the first test because they had all the data that was needed in this test. Then I used the drop-down menu in the engine-related column and checked the yes box, to get all the engine-related entries to only be present. I filtered the reason so that all the issues related to the engines were showing, except the lost at site issues. Lost at site often means that someone has accidentally lost a part or equipment, and that does not really belong in a quality report. I filtered this by using the drop-down menu and checking all the boxes except the lost at site. Then I copy pasted it into Word to make it readable by WärtsiläGPT and deleted the engine-related column to make it more compact. This process did not take more than half an hour, after getting some practice from the test before.

Firstly, I took the first project and made a prompt to get a quality report on all the issues that have occurred. It took several tries until I got the perfect prompt for this. The first time, I had forgotten to mention that I do not need any dates, and it mentioned issues that were resolved that did not really need any follow-up, so I wrote that it did not need to mention that. Then it was that issue that it only mentioned to check things at the site, so I corrected it to mention how to correct it already in the factory. Here I also mentioned writing it in a running text and skipping the dates, because I forgot to do it right away. This process of coming up with a prompt was easier after already making quality reports from the factory issue register data. After several tries, I ended up with this prompt:

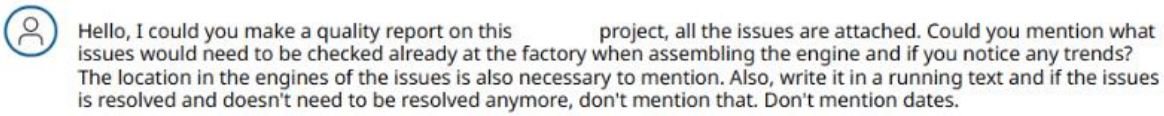

Figure 24: Prompt for the quality report on one project with all recorded issues

When the prompt was working the way it was intended to, I used it to make the reports for the three remaining projects. When all the individual reports for each project were done, I had the AI compile them into one report. To make this work, I had to change the prompt a few times, the text was long in the beginning. To make it easier for the AI, I mentioned that all these projects have the same type of engine. I had to tell the AI not to include the engine numbers because the report was too long with them included. Also, I had to write that this report could not be longer than one page long, the purpose of a quality report is that it should be short and easy to read. After some corrections, I ended up with this prompt:

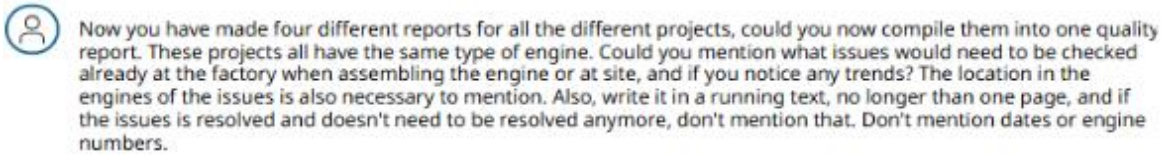

Figure 25: Prompt for summarizing all the project recorded issues into one quality report

Once the prompts were optimized, the reporting went well. After realizing that WärtsiläGPT has technical limitations, I figured that it would be impossible to try to get a quality report of all the projects with the information in one Word document. This did not work in the last test, and in this test, I had far more text in the Word documents, so the safest way was to make them individually and then summarize the projects into one quality report. This time the whole process of gathering information and getting the correct prompt took about 1,5 hours, even though I had some practice from before.

5 Results (Confidental)

The content of the results chapter has been left out from the public version of this thesis due to confidentiality.

6 Conclusion

The purpose of this thesis was to figure out how an AI could help with making quality reports for engine-related issues, with the help of an existing AI or a completely new one. WärtsiläGPT is Wärtsila's own generative AI, which was tested during this thesis to see if it had the capability to make quality reports if it had access to the correct data. In the introduction part of this work, a couple of questions needed to be answered, and the answers to those questions are in the next sub-chapter.

During this thesis work, information has been gathered from websites, internet articles, blog posts, discussions with colleagues, and from internal documents. The discussions with colleagues were necessary to get the help that was needed to figure out how these internal databases worked.

Based on the results that have been gathered during the testing of the WärtsiläGPT, it would be possible to have an AI make these reports. WärtsiläGPT had some technical limitations that made the report generating harder, but they are making improvements regularly. The solution to this would be to use the already implemented generative AI and develop it to be more prone to report generation. If there is a possibility that one day the WärtsiläGPT has access to internal documents, that would make quality reporting much smoother.

The rest of this chapter will summarize what challenges I faced and the answers to the questions from the introduction. As well as, if there are any further research ideas and how that could be achieved.

6.1 Answering Research Questions

In this subchapter, I will be discussing more about the questions that were stated in the introductory chapter.

- How does generative AI work, and is it fit for report generation?

This was explained in the theory chapter, where I explained how Generative AI works and how it is used in report generation. To answer this question more specifically, I made the WärtsiläGPT testing and concluded that Generative AI can make reports if it is trained correctly, and if it is provided the correct data and prompt.

- What information will be taken from which internal database?

In Chapter 4, I discussed the method, and all the databases are stated there. Here I mentioned how they are related to each other and from where the information was taken for these reports. I also mentioned the difficulty of finding the correct data and how much time was spent on this.

- Does Wärtsilä's generative AI access to the internal documents needed for this thesis?

In the methods chapter, it was explained that WärtsiläGPT does in fact not have access to these internal documents needed for this thesis. I asked both colleagues and the WärtsiläGPT itself, and both parts denied that it had access. Also, the roadmap did not provide any information about this implementation in the near future.

- Is WärtsiläGPT able to make quality reports if it has access to the correct data?

The results that were generated from WärtsiläGPT show that it is able to make quality reports, but some improvements would be needed to make the result even better. Some of these reports were somewhat vague and indicated the need for fine-tuning.

- How does better prompting contribute to a more accurate answer?

During the testing of WärtsiläGPT, I used some of the methods of prompting mentioned in the theory chapter to improve the responses generated from the AI. Not all of the methods were implemented due to them not really fitting the purpose of the result that I was going for. The methods that were used during the prompting did help me get the output that I was looking for.

6.2 Challenges

One of the biggest issues I had with this thesis work was to find reliable information sources for my theory part. Generative AI is relatively new, and most of the information that I found was from internet websites, news articles online, or blog posts. Most of the time, I had to research the websites to make sure that the information provided could be trusted, and this made the theory part more time-consuming than it should have been.

WärtsiläGPT is not fully developed yet, and it is not able to analyze text perfectly, and that that made the report generating a lot harder. This was something that I needed to work around, and at last, I did end up with reliable reports.

6.3 Proposal for Further Research

When researching AI, it is very important to find these trusted websites, I personally found a lot of information from university websites from all over the world. In the future, there is probably more information available about generative AI and what uses it could offer.

This thesis could have included that I would have worked with the team that has developed the WärtsiläGPT and made a prototype that would have access to the correct data and would make the quality reports from that. The interest in the company of such an AI tool could also be investigated, to see how many employees in Wärtsilä would benefit from this.

The team which is behind the WärtsiläGPT also has a roadmap of things that will be implemented in the WärtsiläGPT in the future. I did notice that at the end of 2025, they will be adding the ability to read Excel files, which would have been helpful during my testing. This is something that will make the future of data searching much easier, due to the databases having the ability to export to Excel.

6.4 Summary

To summarize this thesis work, I learned a lot about how generative AI works and how it can be used to make certain processes more effective. This engine quality reporting is something that certainly could be executed in Wärtsilä when they already have the AI that could be utilized for this. Quality reporting could be a great way for the employees working with engine quality to get an overview of which recurring issues are happening and, in that way, prevent those from happening in the future.

7 References

Artificial Intelligence Board of America. (2024, October 18). *AI Algorithms: What They Are and How They Work?* Retrieved from ARTiBA: https://www.artiba.org/blog/ai-algorithms-what-they-are-and-how-they-work

Barbarossa, S., & Filisetti, D. (2024). *How generative AI is re-inventing narrative reporting and disclosure management*. Retrieved from Wolterkluwer: https://www.wolterskluwer.com/en/expert-insights/how-genai-is-re-inventing-disclosure-management

Bigelow, P. (2023). *Toyota unleashes generative AI to meld vehicle design with engineering principles*. Retrieved from Autonews: https://www.autonews.com/mobility-report/toyota-researches-generative-ai-use-car-design/

Bigelow, S. J. (2025). *12 prompt engineering best practices and tips*. Retrieved from techtarget: https://www.techtarget.com/searchenterpriseai/tip/Prompt-engineering-tips-and-best-practices

ChatGPT. (2024). *ChatGPT*. Retrieved from ChatGPT: https://chatgpt.com/

Cleary, D. (2025, February 7). *Three Prompt Engineering Methods to Reduce Hallucinations*. Retrieved from PromptHub: https://www.prompthub.us/blog/three-prompt-engineering-methods-to-reduce-hallucinations

Dastin, J. (2018). *Amazon scraps secret AI recruiting tool that showed bias against women*. Retrieved from Reuters: https://www.reuters.com/article/world/insight-amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK0AG/

Google Cloud. (2024). *What's the difference between deep learning, machine learning, and artificial intelligence?* Retrieved from Google Cloud: https://cloud.google.com/discover/deep-learning-vs-machine-learning

Harvard University. (2023). *Getting started with prompts for text-based Generative AI tools*. Retrieved from Harvard University Information Technology: https://www.huit.harvard.edu/news/ai-prompts

IBM. (2023). *What are AI hallucinations?* . Retrieved from IBM: https://www.ibm.com/think/topics/ai-hallucinations

IBM. (2024). *What is explainable AI?* Retrieved from IBM: https://www.ibm.com/think/topics/explainable-ai

Jaro Education. (2024, March 20). *Future of Generative AI*. Retrieved from Jaro Education: https://www.jaroeducation.com/blog/future-of-generative-ai/

Kaur, I. (2024). *5 Use Cases Of GenAI Automated Report Generation*. Retrieved from crossml: https://www.crossml.com/use-cases-of-genai-automated-report-generation/

Miller, A. (2024, July 16). *15 companies using generative AI for business efficiency*. Retrieved from indatalabs: https://indatalabs.com/blog/companies-using-generative-ai

National University. (2024, March 1). *131 AI Statistics and Trends for 2025*. Retrieved from National University: https://www.nu.edu/blog/ai-statistics-trends/

OpenAI. (2025). *What are tokens and how to count them*. Retrieved from OpenAI: https://help.openai.com/en/articles/4936856-what-are-tokens-and-how-to-count-them

Pichai, S., & Hassabis, D. (2024, February 15). *Our next-generation model: Gemini 1.5*. Retrieved from Google: https://blog.google/technology/ai/google-gemini-next-generation-model-february-2024/#architecture

Purohit, A. (2023, July 25). *AI, ML, DL, and Generative AI Face Off: A Comparative Analysis*. Retrieved from Synoptek: https://synoptek.com/insights/it-blogs/data-insights/ai-ml-dl-and-generative-ai-face-off-a-comparative-analysis/

SAP. (2024). *What is SAP?* Retrieved from SAP: https://www.sap.com/about/what-is-sap.html

Singla, A., Sukharevsky, A., Yee, L., Chui, M., & Hall, B. (2024). *The state of AI in early 2024*. Retrieved from Mckinsey & Company: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Stryker, C., & Scapicchio, M. (2024). *What is generative AI?* Retrieved from IBM: https://www.ibm.com/think/topics/generative-ai

Tableau. (2024). *Artificial intelligence (AI) algorithms: a complete overview*. Retrieved from Tableau: https://www.tableau.com/data-insights/ai/algorithms

The University of Kansas. (2024). *Helping students understand the biases in generative AI* . Retrieved from The University of Kansas: https://cte.ku.edu/addressing-bias-aihttps://cte.ku.edu/addressing-bias-ai

Thejus, Z. (2024). *Generative AI and Privacy: What You Need to Know*. Retrieved from Webtoffee: https://www.webtoffee.com/generative-ai-and-privacy/

Wärtsilä. (2022). *STH*. Retrieved from Wärtsilä: https://www.wartsila.com/sth

Wärtsilä. (2023). *Factsheet*. Retrieved from Wärtsilä: https://www.wartsila.com/investors/wartsila-as-an-investment/factsheet

Wärtsilä. (2024). *About*. Retrieved from Wärtsilä: https://www.wartsila.com/about

Wärtsilä. (2024). *History*. Retrieved from Wärtsilä: https://www.wartsila.com/about/history

Wärtsilä. (2024). *Our businesses*. Retrieved from Wärtsilä: https://www.wartsila.com/about/our-businesses

Wärtsilä. (2024). *Wärtsilä's Central Distribution Centre*. Retrieved from Wärtsilä: https://www.wartsila.com/about/190/article/the-engine-behind-efficient-logistics-wartsilas-central-distribution-centre

Wärtsilä. (2025). *Key figures*. Retrieved from Wärtsilä: https://www.wartsila.com/investors/financial-information/key-figures

Wärtsilä Corporation. (2017). *What is DCM365? | Wärtsilä*. Retrieved from Youtube: https://www.youtube.com/watch?v=RwlAx_U3MxM

Yudovski, H. (2024). *50+ Best ChatGPT Prompts for Marketing + Examples*. Retrieved from Mayple: https://www.mayple.com/resources/digital-marketing/chatgpt-prompts-for-marketing