




# Virtual assistant using generative AI applied to iron ore concentration


Tiago Caixeta Nunes <sup>1,2\*</sup> 

Rodrigo Martins Gomes <sup>3</sup> 

Felipe Novaes Caldas <sup>1</sup> 

Ednei Rodrigues Rocha <sup>3</sup> 

Eric Guimarães Vieira <sup>4</sup> 

Marcelo Pereira de Castro Alves <sup>5</sup> 

## Abstract

The use of Artificial Intelligence (AI) in the competitive mining market has attracted increasing interest in recent years, as these technologies play a significant role in data interpretation, modernization of production processes, and efficient use of mineral reserves. Among the various AI approaches, Generative Artificial Intelligence (Gen-AI) stands out as one of the most promising and disruptive. This paper aims to present a virtual assistant application based on Gen-AI, capable of providing answers to user questions using natural language about operational aspects, based on several data sources to assist in decision-making that influences the performance of iron ore concentration. The assistant considers information security infrastructure, uses language models (LLMs) and Retrieval-Augmented Generation (RAG) techniques to access plant databases and documents, while a multi-agent flow centralizes application information with production data, historical data, and technical documentation.

**Keywords:** Generative artificial intelligence; Iron ore; Concentration.

## 1 Introduction

Like other industrial segments, the mining sector is constantly transforming, using and investing in Artificial Intelligence (AI) applied to business to drive efficiency and maximize value [1]. AI aims to create intelligent entities capable of learning and performing tasks that typically require human intelligence in devices or computerized systems capable of acquiring knowledge to perform tasks [2].

The accelerated development of AI is due to the growing need for technologies capable of processing large volumes of data, as well as learning and adapting to changing demands. AI methodologies are commonly classified into two main categories: reasoning and decision-making; and learning. The former involves decision-making, problem-solving, and optimization, while the latter addresses learning, communication, and perception [3].

Among learning techniques, Generative Artificial Intelligence (Gen-AI) stands out. It is an intelligent design method that boasts powerful data learning, efficient new generation capabilities, and has been continuously developing and advancing in recent years [4]. The rapid advancement of AI-related technologies is evident in innovations such

as ChatGPT, Gemini, and Copilot. Gen-AI models seek to understand and replicate the patterns inherent in a dataset, facilitating the creation of new content that closely mirrors the original data distribution [5].

In this context, the objective of this research is to develop a conversational chat-based web application that uses Gen-AI to answer questions related to the operational aspects of the iron ore concentration production process, initially considering the unit stages of mineral processing. Thus, the integration of queries regarding mine variables, feed quality, process variable stability, and unit stages can contribute to reducing response times for deviations that impact quality and metallurgical recovery.

Thus, a Gen-AI-based Virtual Assistant was developed, considering in its technical structure the sources of information, data collection, architecture, question mapping, and how it interacts with users and other systems within the mining company's application ecosystem. The AVM is capable of providing answers to user questions using natural language about operational aspects, based on various data sources, to aid in decision-making that influences iron ore concentration performance.

<sup>1</sup> Departamento de Processamento Mineral, Vale SA, Nova Lima, MG, Brasil.

<sup>2</sup> Universidade Federal de Minas Gerais - UFMG, Belo Horizonte, MG, Brasil.

<sup>3</sup> Departamento de Tecnologia, Vale SA, Nova Lima, MG, Brasil.

<sup>4</sup> Departamento de Operações, Vale SA, São Gonçalo do Rio Abaixo, MG, Brasil.

<sup>5</sup> Engenharia, AI & Data, Deloitte, Belo Horizonte, MG, Brasil.

\*Corresponding author: [tiago.caixeta@vale.com](mailto:tiago.caixeta@vale.com)

E-mails: [tiago.caixeta@vale.com](mailto:tiago.caixeta@vale.com), [rodrigo.martins@vale.com](mailto:rodrigo.martins@vale.com), [fncaldas@gmail.com](mailto:fncaldas@gmail.com), [ednei.rocha@vale.com](mailto:ednei.rocha@vale.com), [eric.vieira@vale.com](mailto:eric.vieira@vale.com), [marcealves@deloitte.com](mailto:marcealves@deloitte.com)



To implement the tool under study in this proposed work, an iron ore concentration plant was selected to develop the application. Located in Minas Gerais, it produces ore with an iron content between 42% and 45%. Processing at the plant involves comminution to reduce the ore to 8.0 mm, concentration of the -8.0 mm + 1.0 mm fraction by jigging, producing coarse sinter feed, magnetic concentration of the -1.0 mm + 0.15 mm fraction, producing fine sinter feed, and concentration of the -0.15 mm fraction by flotation, producing pellet feed. The plant's tailings are preferably filtered through disc filters and sent for stacking in the tailings disposal pile. A simplified flowchart of the concentration plant is illustrated in Figure 1.

Therefore, this article aims to utilize a Gen-AI-based virtual assistant in iron ore concentration, focusing on the operational practices of concentration plants. The Gen-AI concept is further detailed below.

## 2 Generative Artificial Intelligence (Gen-AI)

AI is a disruptive technology with great potential to transform various aspects of society and industry. Among the ongoing technological advances in AI, Gen-AI has emerged as a powerful tool in several domains [6]. It enables the use of algorithms and machine learning to generate innovative solutions or improve existing processes. Gen-AI is emerging as a unique subcategory of AI, gaining prominence through platforms such as ChatGPT, Gemini, Grok, DeepSeek, Copilot, and others. Thus, Gen-AI represents a striking example of unsupervised machine learning with enormous potential.

Collaboration between humans and AI-based systems proves essential to address the challenges and seize the opportunities afforded by Gen-AI [7]. With the continued advancement of generative algorithms, research on chatbots has intensified. Traditionally, these systems relied on Natural

Language Processing (NLP) to interpret user requests and match them with the most appropriate responses [8]. However, modern chatbots have evolved significantly, incorporating language models and deep learning techniques to provide immediate responses, expanding their capability to address NLP challenges during real-time interactions with users.

The evolution of natural language capabilities with Gen-AI is drastically changing the way businesses operate and customers consume their products and services. This interdisciplinary impact highlights the comprehensive nature of Gen-AI, shaping the technology landscape and fostering innovation and adaptation across a spectrum of industries, and mining is no exception.

To apply AI algorithms, obtaining a relevant database is essential. However, this prerequisite remains an obstacle and can impact the effectiveness of existing algorithms [9]. Although Gen-AI applications are evolving rapidly, much of the research still focuses on the end-user perspective, leaving significant gaps regarding the factors, dynamics, and outcomes involved in the adoption of Gen-AI by managers and decision-makers. There are also existing gaps and precautions to be taken during its implementation in the industry, with cybersecurity and ethical implications being important to consider.

Therefore, this study stands out for developing and applying a Gen-AI-based virtual assistant to iron ore concentration, as follows.

## 3 Development

This work developed a Gen-AI-based tool called the *Assistente Virtual da Mineração (AVM - Mining Virtual Assistant)*. Its technical structure encompasses information sources, data collection, architecture, question mapping, and interaction with users and other systems within the mining company's application ecosystem.

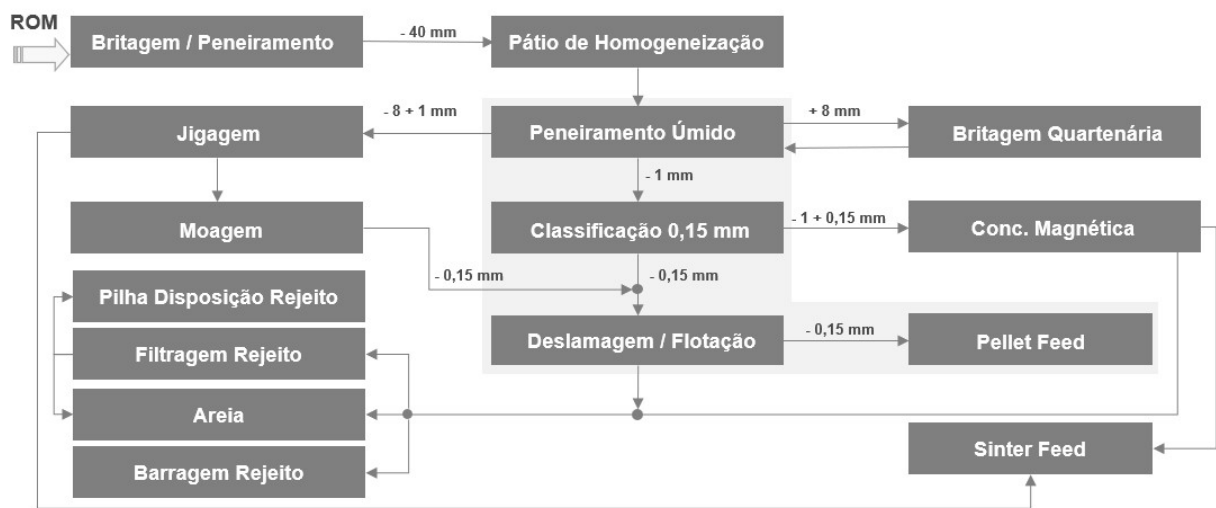


Figure 1. Summary production flow - Iron Ore Beneficiation.

The AVM is capable of providing natural language answers to user questions about operational aspects, based on various data sources, to assist in decision-making that influences iron ore concentration performance.

The AVM uses large language models (LLMs) in a multi-agent flow, which can query Vale's internal documents using the Retrieval-Augmented Generation (RAG) technique and internal data through SQL queries and access to the mining company's system data. LLMs are language models that use neural networks containing billions of parameters, trained on huge amounts of unlabeled text data using a self-supervised learning approach [10].

These combined technologies allow the application to answer user questions using natural language, drawing on information from various pre-configured data sources. This generates user-friendly responses in natural language, based on the data queried.

Initially, four data sources were defined for the AVM:

1. Process Variable Stability Index (DEP);
2. Historical Process Data (PIMS - Plant Information Management System)
3. Production Indicator Information (MES - Manufacturing Execution System);
4. Technical Documents: procedures, manuals, and technical references.

The technical structure of the AVM application is detailed below.

### 3.1 Data source

**DEP** is a system that uses Statistical Process Control (SPC) to monitor the level of agreement of **process variables** within established optimal ranges, focusing on reducing variability. The variable stability index is calculated by dividing the hours an asset operated within limits by the hours actually in operation. This tool uses data calculated by PIMS, collects data for configured variables, and presents the information in a Power BI dashboard format to the end user. The system uses criteria to classify information, allowing various data groupings, such as: operational unit, process step and sub-step, among others.

**PIMS** is a tool that stores the history of general variables across the entire plant (sensors, actuators) so they can be consulted later. Furthermore, PIMS allows you to create calculated tags of interest that use other existing measured or calculated variables or native functions of the tool. In general, hundreds of variables with hourly averages were considered for AVM using the Gen-AI knowledge base. The format of this consolidation can vary depending on the nature of each variable. To access PIMS data, a replication of this data was configured in InfoHub, which is a SQL database that stores the information for each configured tag in a consolidated 10-minute period. Thus, AVM accesses PIMS data from a SQL table with the following columns:

- Reference Date/Time: date and time of the beginning of the recording period;
- Tag: identifier of the information to be collected;
- Value: value of the collected variable consolidated in the recorded period.

Only tags that are of interest to the AVM will be replicated to InfoHub. In summary, variables can be grouped according to the following contexts, such as: percentage of time in a given condition, difference between set point and process variable, level, current, pH, grade, particle size, rate, etc.

The **MES** system contains official and up-to-date information on operational and production indicators. This system is implemented throughout the ferrous metals chain and serves as an important data source for the AVM. MES indicators are represented by codes and are configured by end users within the tool itself. There are data consolidation processes for these indicators that allow systems using MES data to have a simplified and unified interface for collecting indicator data.

Another data source used by the AVM is the unit's **Technical Reference Documentation**. Thus, it is possible to search preprocessed documents for the information needed to answer questions related to their content. This functionality allows users to obtain information more quickly and without having to read documents, such as procedures, manuals, technical references, training, and others.

The Figure 2 illustrates some data source, considering some issue questions related to each source in the AVM.

### 3.2 Architecture

The technical architecture of the implemented solution is presented in the simplified flowchart in Figure 3. The architecture demonstrates the interaction between the different agents that make up the system, from the user's question to the delivery of the final response. The flow is structured into steps that include input preprocessing, task planning, querying databases or documents, and finally, processing and formatting the response. Each agent has a specific role in this process, ensuring efficiency and accuracy in the response to the user.

The application uses a multi-agent workflow that utilizes LLM to retrieve relevant information from documents and operational data and provide contextualized answers. The infrastructure is hosted on Azure App Services, ensuring scalability and security. The system uses Azure AI Search as a vector database, where documents are stored and retrieved to provide data for answers. Additionally, queries related to operational processes are performed in Azure SQL Database, enabling integration with the organization's data. The agent concept is detailed below.

#### 3.2.1 Agents

The AVM is structured into different agents, each with a specific function within the application flow. The agents work in a complementary manner to interpret the user's question and plan the necessary actions to obtain the final answer.

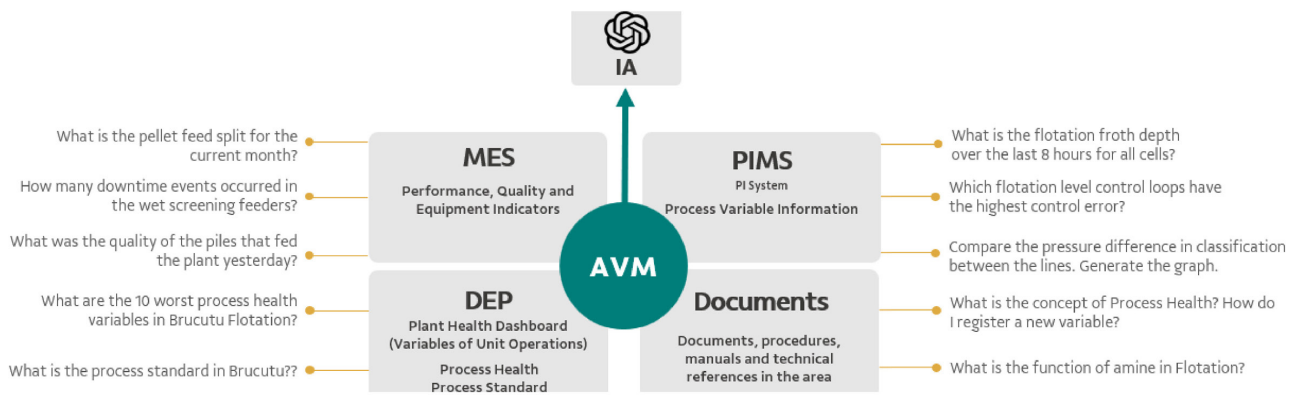


Figure 2. Data source of Virtual Assistant.

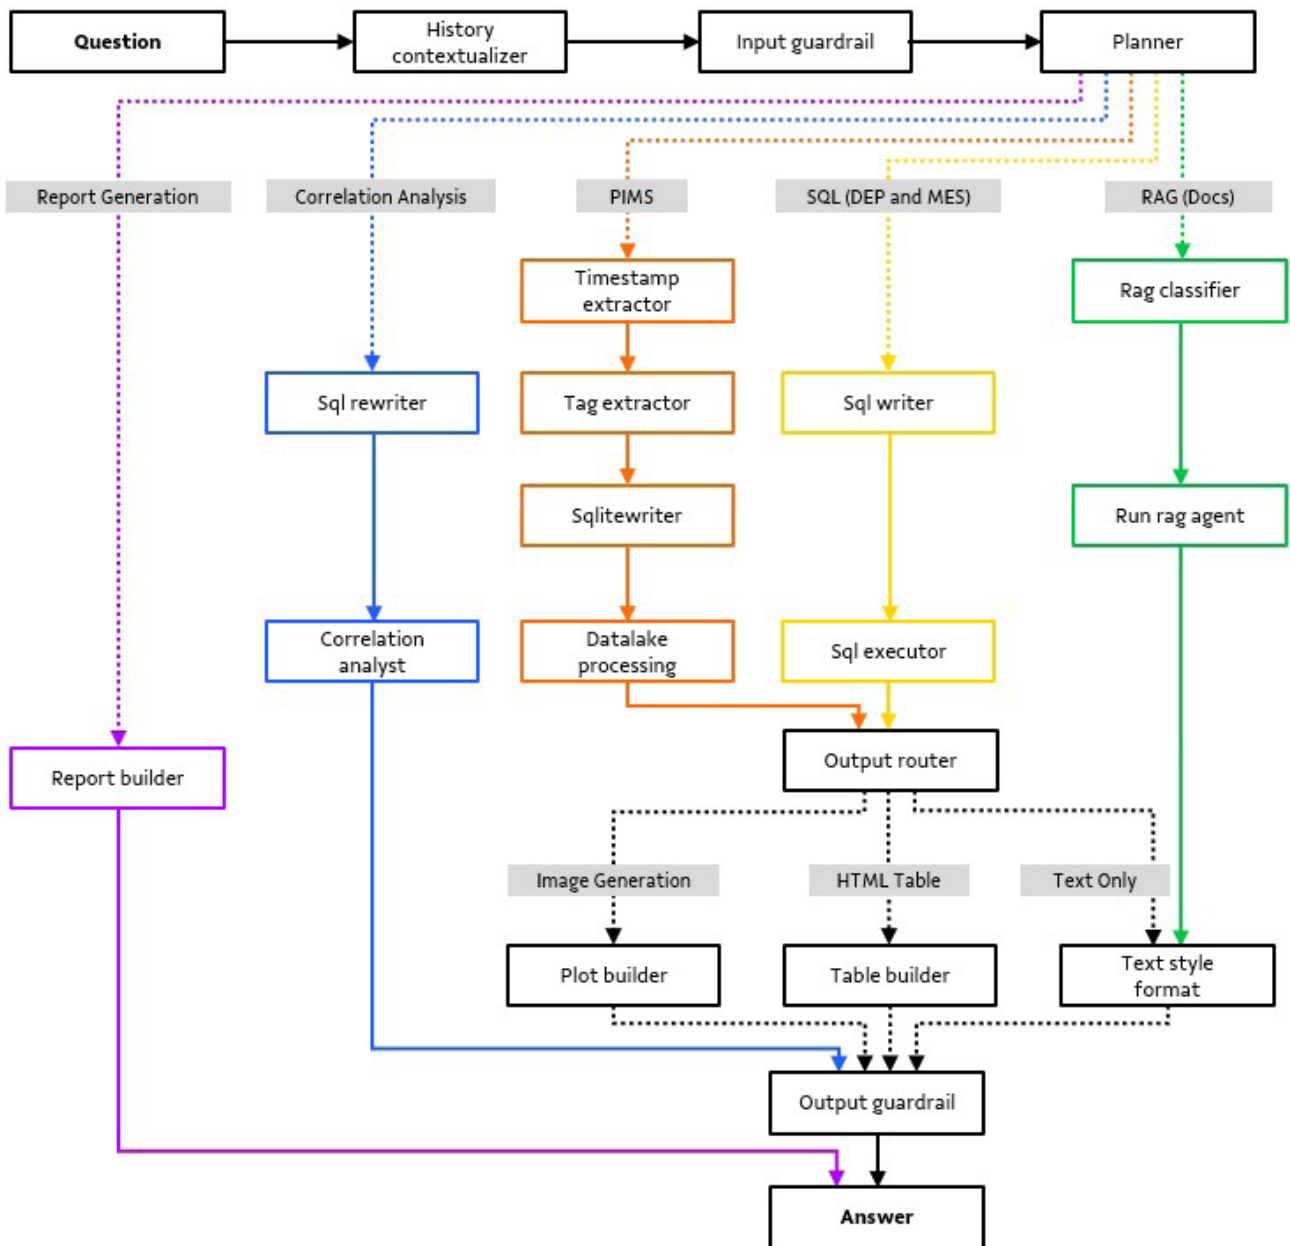


Figure 3. Simplified agent architecture and data queries.

- **Guardrails Layer:** this layer filters both the user’s questions and the answers generated by the agents;
- **Pre-Processing Agent:** this layer is where the user’s question is pre-processed so that subsequent steps can be assigned;
- **Planning and Management Agent:** after passing through the entry guardrails: through the History Contextualizer, this agent determines the actions that must be performed to fulfill the user’s request;
- **Task Execution Agents:** these are the agents that make up each path defined by the planner and are responsible for interacting with databases and documents, performing the necessary queries (RAG, SQL (DEP and GPVM), PIMS, Correlation Analysis, and Report Generation);
- **Response Agents:** this is the final layer where the searched data will be processed and formatted, transforming it into the final response sent to the user.

### 3.3 Reference question mapping

To guide the agents’ learning, approximately 300 questions related to each data source were initially mapped. To contextualize the agent within the DEP, PIMS, and MES databases, it was necessary to structure the data according to process information, such as by process area, sub-area, equipment, among others, aiming to facilitate the agent’s consultation to answer. In some cases, the answer involves consulting different sources.

In addition, several specific details must be inserted into the code for the agent’s understanding, considering tag nomenclature, synonyms, and common terms used by the operational team. For example, if an equipment tag containing the acronym “SM” is inserted, it must be specified that it is a Magnetic Separator, which can also be called a Magnetic Concentrator, Jones or Whims. Some questions used as reference in the AVM system interface are illustrated below.

- What is the Process Health (stability index) per unit stage today?
- How is the pellet feed split going last week?
- What is the quality of the batteries fed this month?
- What is the role of amine in flotation?
- Which flotation valves showed closed saturation last day?

### 3.4 User interface

The web interface is how users interact with AVM, provided they have permission to authenticate, as illustrated in Figure 4. AVM offers versatility in presenting information through multiple display formats, adapting to the data type and context of the query. Responses can be displayed in text format (explanatory responses in natural language), graphs, tables, correlation matrices, and predefined reports. Each response displayed contains the data source and the date of the last update. The window consists of:

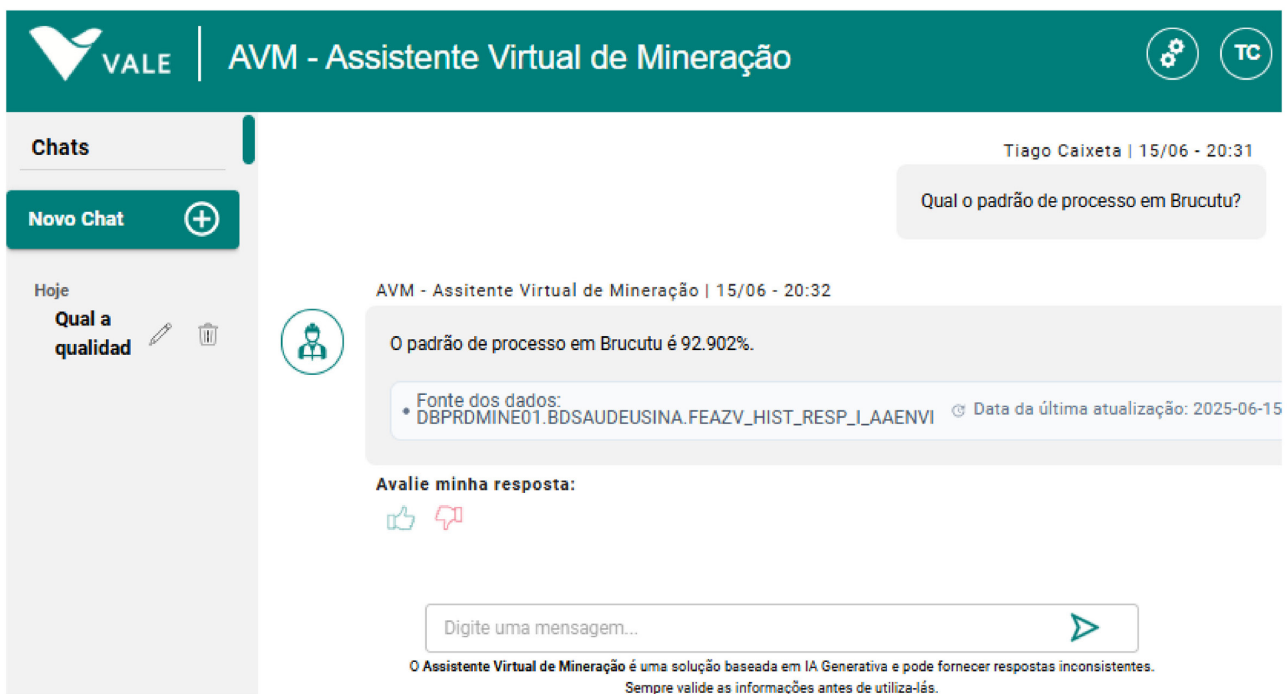


Figure 4. Virtual Assistant web interface.

- Typing bar: a lower field with a chat where the user enters a question, and the application will perform a search to display the answer;
- Session bar: a left-hand side panel located to the left of the main chat interface, responsible for managing and navigating between different user conversations with AVM, organized chronologically;
- Response display: a region of the screen for displaying the agent’s responses, maintaining a complete chronological record through timestamps that document the date and time of both the question and the generated response.

#### 4 Gen AI application results

The AVM web interface was developed for both computers and tablets, enabling consultation in the office and in the field by users from various disciplines, such as

operations, maintenance, process, engineering, and automation, among others. The answers can be used to assist in decision-making that influences plant performance. Some answers are illustrated in the figures below.

In Figure 5, operations teams can verify that the process variables of the unit stage of thickening have a worse process health index, that is, they are operating outside the specification limits for longer periods.

Figure 6 illustrates a question related to technical documents to answer what the function of starch is in flotation; in this case, operational procedures and training material on flotation and reagents were consulted.

In Figure 7, operations, maintenance, and automation teams can check which flotation control valves have the highest percentage of saturated time open or closed, which may result from some anomaly in the field that impacts process control.

The average behavior of some minerals for a %SiO<sub>2</sub> cluster can be analyzed as shown in Figure 8. In some cases, the values of certain contaminants contribute to worsening the %SiO<sub>2</sub> in the final concentrate. Figure 9 indicates a question related to the quality of the piles that fed the plant that day.

What is the process health of Brucutu's process steps

AVM - Virtual Mining Assistant | July 13 - 3:09 PM

Operating Unit	Stage	Data	Process Health	Normal Hours	Operating Hours
Brucutu	Filtering	2025-07-13 02:00:00	82.7	52.52	63.51
Brucutu	Crushing and Screening	2025-07-13 04:00:00	80.55	38.35	47.61
Brucutu	Densification	2025-07-13 05:00:00	71.95	54.74	76.09
Brucutu	Flotation	2025-07-13 08:00:00	82.7	61.2	74.0

Figure 5. AVM response for process variable stability index in DEP.

What is the role of starch in flotation?

AVM - Virtual Mining Assistant | July 13 - 4:15 PM

Starch plays an essential role in the iron ore flotation stage, acting as a **selective depressant**. Its main function is to modify the surface properties of minerals, making iron particles, such as hematite and goethite, more hydrophilic (with a greater affinity for water). This allows for the selective separation of quartz, which remains hydrophobic and can be floated.

**Functions and Mechanisms of Starch in Flotation:**

- 1. Depression of Iron Minerals :**  
Starch preferentially adsorbs to the surface of iron particles, such as hematite, preventing them from becoming hydrophobic and floating. This occurs because starch acts as a barrier, making these particles immune to the action of the collector (amine), which will be used to float the quartz.
- 2. Selective Adsorption :**  
Starch's selectivity is explained by its chemical affinity for the iron cations present in hematite. The hydroxyl (OH) groups of starch interact with these cations, while their adsorption on quartz is significantly lower. This difference is crucial to the efficiency of the process.
- 3. Starch Gelatinization :**  
Before use, starch undergoes a gelatinization process, usually carried out with sodium hydroxide (NaOH) in an oxidizing environment. This process increases its solubility and promotes the formation of COOH- salts, which have a high affinity for hematite particles, improving their adsorption capacity and effectiveness as a depressant.
- 4. Selective Separation :**  
After adding starch, the iron particles remain in the pulp, while the quartz, which is unaffected by the starch, is floated away with greater selectivity by the action of the amine.

Figure 6. AVM response to explain starch function by consulting technical documents.



Figure 7. AVM response to % saturation time of PIMS flotation control valves.

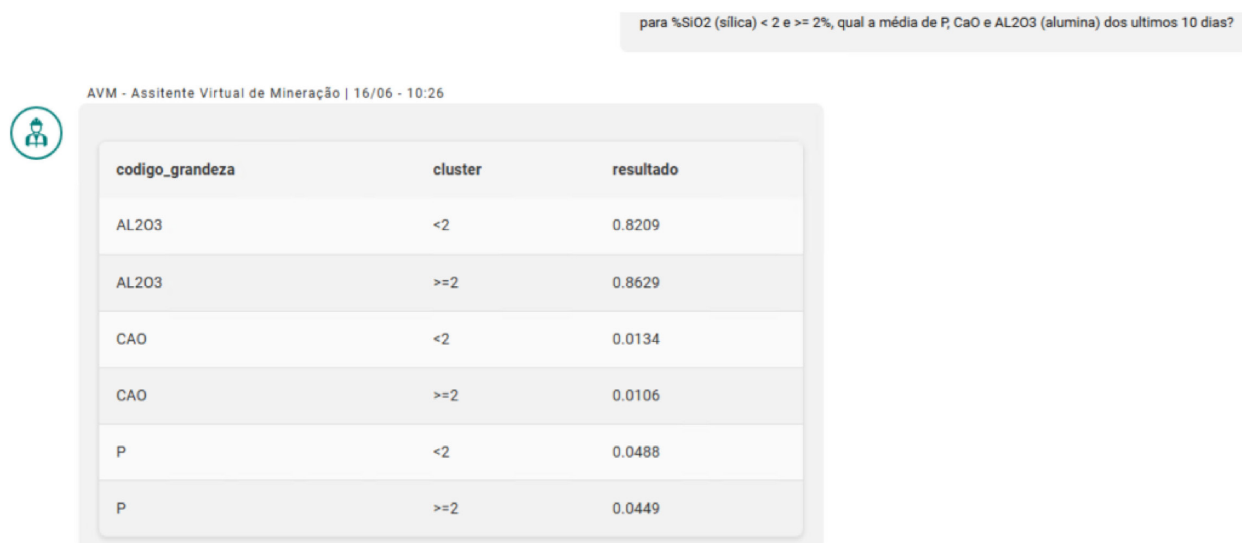


Figure 8. AVM response to average contaminants as a function of SiO<sub>2</sub> MES cluster.

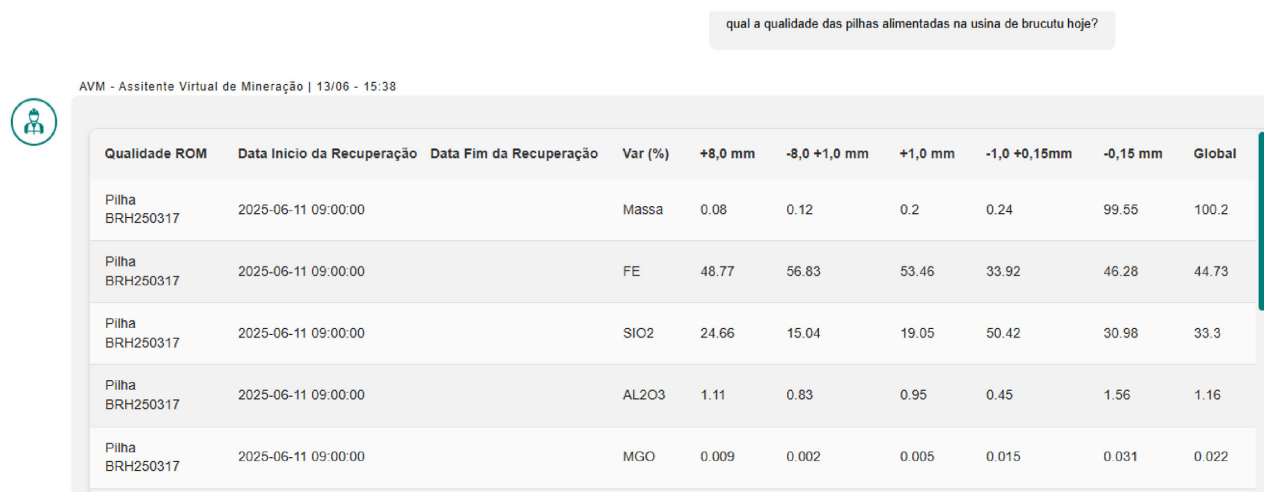


Figure 9. AVM response to the quality of the stockpile (MES) that fed the plant.

## 5 Conclusion

Over the past few years, Artificial Intelligence has evolved from a limited technological resource to a powerful agent of transformation with profound impacts on the industry. This study demonstrated that the developed Gen-AI-based virtual assistant application is capable of providing answers to user questions using natural language about operational aspects, based on various data sources, to assist in decision-making that influences the performance of iron ore concentration.

By integrating a Gen-AI-based Virtual Assistant into operations, it is expected that in the short term, it will not only explain the impact on decision-making but also act preventively on deviations to promote gains. It is also expected that explanatory questions about the causes and variables that influenced the performance of a unit step, quality loss, or volume loss, can be answered via chat or in predefined reports, and can be automatically sent passively to stakeholders in the operational areas, rather than this proposed solution, which responds based on active questions.

Based on the preliminary AVM results, several opportunities for improvement and future work were identified, aiming to integrate information from block models, yard quality, plant process variables, reliability indicators, and equipment performance. One challenge for these opportunities is to structure a database and increase the query sources for Gen-AI, which is expected to grow substantially.

Given the findings of this work, it is clear that the use of Gen-AI represents a significant advancement in the iron ore industry. Thus, this research contributes innovatively to the development of AI solutions applied to mining, paving the way for a new era of more intelligent, responsive processes that are connected to the challenges of modern mining. This study advances in a structured manner to facilitate replication across multiple plants.

## Acknowledgements

The authors are grateful to VALE SA, UFMG, CAPES, CNPq and FAPEMIG for all the support provided.

## References

- 1 Chandramohan R, Pyle M. Can AI save the mining industry? In: Proceedings of the XXXI International Mineral Processing Congress IMPC; 2024; Englewood. Englewood: SME; 2024. p. 1071-1802.
- 2 Fosso Wamba S, Bawack RE, Guthrie C, Queiroz MM, Carillo KDA. Are we preparing for a good AI society? A bibliometric review and research agenda. *Technological Forecasting and Social Change*. 2021;164:120482. <https://doi.org/10.1016/j.techfore.2020.120482>.
- 3 Fontoura L, Nascimento DLM, Neto JV, Caiado RGG. Energy Gen-AI technology framework: a perspective of energy efficiency and business ethics in operation management. *Technology in Society*. 2025;81:102847. <https://doi.org/10.1016/j.techsoc.2025.102847>.
- 4 Liao W, Lu X, Fei Y, Gu Y, Huang Y. Generative AI design for building structures. *Automation in Construction*. 2024;157:105187. <https://doi.org/10.1016/j.autcon.2023.105187>.
- 5 Jackson I, Ivanov D, Dolgui A, Namdar J. Generative artificial intelligence in supply chain and operations management: a capability-based framework for analysis and implementation. *International Journal of Production Research*. 2024;62(17):6120–6145. <https://doi.org/10.1080/00207543.2024.2309309>.
- 6 Ma X, Huo Y. Are users willing to embrace ChatGPT? Exploring the factors on the acceptance of chatbots from the perspective of AIDUA framework. *Technology in Society*. 2023;75:102362. <https://doi.org/10.1016/j.techsoc.2023.102362>.
- 7 Kumar A, Shankar A, Hollebeek LD, Behl AH, Lim WM. Generative artificial intelligence (GenAI) revolution: a deep dive into GenAI adoption. *Journal of Business Research*. 2025;189:115160. <https://doi.org/10.1016/j.jbusres.2024.115160>.
- 8 Kecht C, Egger A, Kratsch W, Roglinger M. Quantifying chatbots' ability to learn business processes. *Information Systems*. 2023;113:102176. <https://doi.org/10.1016/j.is.2023.102176>.
- 9 Patel D, Sahu CK, Rai R. Security in modern manufacturing systems: integrating blockchain in artificial intelligence-assisted manufacturing. *International Journal of Production Research*. 2024;62(3):1041-1071. <https://doi.org/10.1080/00207543.2023.2262050>.
- 10 Raiaan MAK, Mukta MDSH, Fatema K, Fahad NM, Sakib S, Mim MMJ, et al. A review on large language models: architectures, applications, taxonomies, open issues and challenges. *IEEE Access: Practical Innovations, Open Solutions*. 2024;12:26839-26874. <https://doi.org/10.1109/ACCESS.2024.3365742>.

Received: 24 Oct. 2025

Accepted: 19 Mar. 2026

Editor-in-charge:

Willy Ank de Morais 