AIMS Geosciences

AIMS Geosciences, 10(3): 641–661.
DOI: 10.3934/geosci.2024033
Received: 29 May 2024
Revised: 26 July 2024
Accepted: 13 August 2024
Published: 02 September 2024

http://www.aimspress.com/journal/geosciences

Research article

Reservoir geophysical monitoring supported by artificial general intelligence and Q-Learning for oil production optimization

Paolo Dell'Aversana*

Eni S.p.A. Via Emilia 1, San Donato Milanese, 20097, Milan (I), Italy

* **Correspondence:** Email: paolo.dell'aversana@eni.com; Tel: +393400736394.

Abstract: Artificial general intelligence (AGI), or strong AI, aims to replicate human-like cognitive abilities across diverse tasks and domains, demonstrating adaptability and learning like human intelligence. In contrast, weak AI refers to systems designed for specific tasks, lacking the broad cognitive flexibility of AGI. This paper introduces a novel approach to optimize oil production by integrating fundamental principles of AGI with geophysical data inversion and continuous monitoring techniques. Specifically, the study explored how AGI-inspired algorithms, combined with established reinforcement learning (RL) techniques, can enhance borehole electric/electromagnetic monitoring and reservoir fluid mapping technology. This integration aims to mitigate the risk of unwanted water invasion in production wells while optimizing oil extraction. The proposed methodology leverages real-time geophysical data analysis and automated regulation of oil production. The paper begins by outlining the key features of AGI and RL, and then discusses their application in electric/electromagnetic monitoring to define optimal production policies. The effectiveness of this approach was verified through synthetic tests, showing significant improvements in production efficiency, resource recovery, and environmental impact reduction.

Keywords: artificial general intelligence; reinforcement learning; geophysical inversion; monitoring; oil production optimization

1. Introduction

The advancement of artificial intelligence (AI) technologies has opened new perspectives in optimizing industrial processes. Artificial general intelligence (AGI) stands at the forefront of the

evolving landscape of artificial intelligence, representing a paradigm shift in our pursuit of creating intelligent systems [1–5]. Unlike narrow or specialized AI systems designed for specific tasks (often indicated as weak AI), AGI aspires to replicate the broad cognitive abilities of the human mind, possessing the capacity to understand, learn, and adapt across a diverse range of domains. The concept of AGI envisions machines not only excelling in predefined tasks, but also comprehending the intricacies of novel challenges, like the flexible intelligence demonstrated by humans. This level of adaptability implies that AGI could solve complex problems and continuously improve its performance over time. The journey toward achieving AGI involves enabling machines to not only process information but also to understand and learn in a manner reminiscent of human intelligence. In this context, AGI emerges as a promising solution to address complex challenges in the oil production sector. We are aware that implementing an effective AGI system is still a theoretical goal, far to be reached with current technology. However, we can develop some basic algorithms inspired by AGI criteria to improve the performances of artificial agents.

This article explores the role of AGI in optimizing oil production, with a focus on managing the risk of water invasion while maximizing hydrocarbon recovery [6–10]. Water invasion occurs when water from the surrounding formations infiltrates the production well. This can significantly impact the efficiency and economics of hydrocarbon extraction. Unfortunately, traditional methods often fall short in providing real-time, accurate assessments of subsurface conditions, leading to suboptimal decision-making and increased operational risks.

Our approach to addressing the problem of water breakthrough is based on the combination of electric and/or electromagnetic data inversion, continuous monitoring techniques, reinforcement learning, and AGI. By installing electrodes and magnetometers within a production well, a persistent monitoring system is established, enabling the ongoing detection and analysis (through time-lapse inversion methods) of variations in electric potentials and/or electromagnetic fields associated with reservoir fluids displacement [11–14]. This approach provides insights into subsurface fluid dynamics, aiding in the understanding of reservoir behavior, fluid migration patterns, and potential geo-hazards.

The physical background of our approach can be summarized as follows. The electrodes capture electrical signals (electric potentials) resulting from fluid movements, allowing real-time data collection on resistivity variations around the wellbore. Similarly, magnetometers measure changes in magnetic fields induced by fluid displacements, enhancing the accuracy and reliability of gathered geophysical data. This dual-sensor system [14] forms a robust network that enables scientists and engineers to track and interpret subsurface processes with precision (mainly fluids displacement), facilitating informed decisions about well operations, reservoir management, and environmental impact assessments.

As an illustrative example, Figure 1 shows a synthetic resistivity model related to fluids displacement at reservoir depth. This model was obtained through the inversion of (simulated) electric potentials recorded in a horizontal production well. The black dots represent the location of electrodes permanently installed in the well (in real measurements, the electrodes must be properly insulated from the metallic casing [14]). The colors represent the different resistivity values. Low resistivity (blue) is associated with the water table. This is clearly detected, in contrast with the more resistive background (oil filled reservoir). Our approach consists of obtaining many models like this through time-lapse resistivity surveys, to map the fluids displacements over the time, during production. The final goal is to prevent undesired water invasion in the production well. To reach this target, we need to regulate the oil production in real time, considering the water displacements around the well.

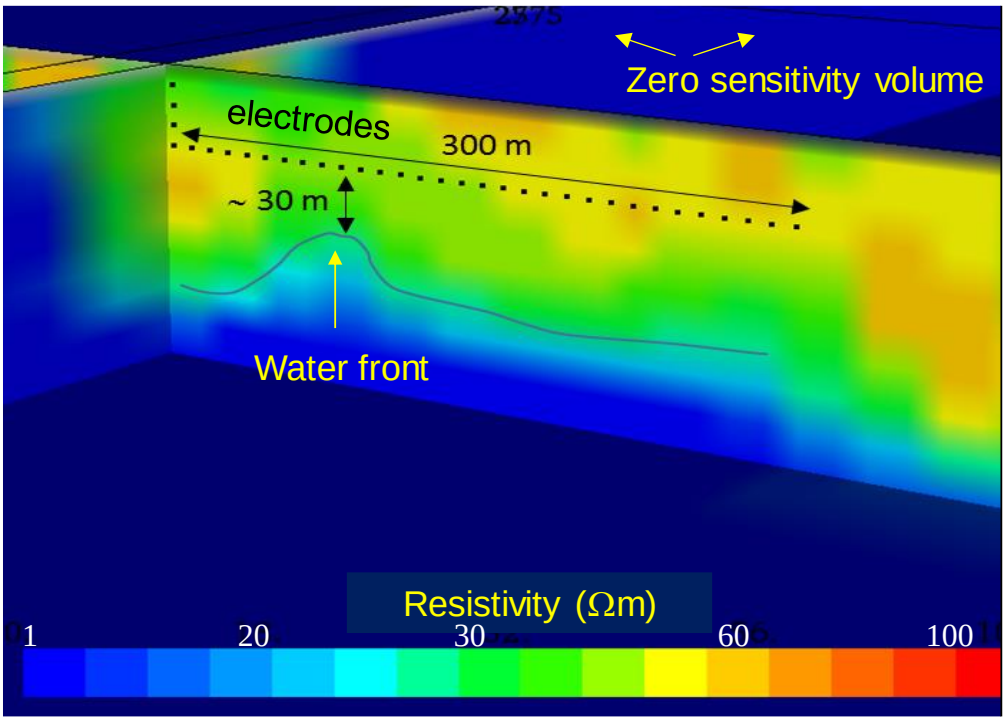

Figure 1. An illustrative modeling example: resistivity section extracted from a 3D model around the production horizontal well, obtained through inversion of synthetic electric data.

In that context, AGI plays a crucial role in interpreting and using the datasets obtained by electric and/or electromagnetic sensors. Through advanced deep learning algorithms, AGI allows analyzing complex patterns in the data, identifying subtle changes associated with fluid movements. This enables real-time decision-making by providing actionable insights into fluid dynamics, optimizing oil production, and minimizing risks associated with potential fluid invasion during production.

A key novelty of our approach lies in its holistic integration of AGI criteria and reinforcement learning (RL) techniques with real-time geophysical monitoring. In the next section, we will discuss the AGI/RL workflow applied to our specific problem. Then, we will briefly recall the key aspects of borehole electric/electromagnetic acquisition/inversion methods. We will discuss how all these technologies can be integrated into a comprehensive methodology. Finally, we will present synthetic tests in realistic scenarios, showcasing the potential of the entire methodology.

2. AGI/RL workflow

The specific structure of the AGI and RL workflow proposed in this article can be summarized in the following two main aspects.

First, we combine deep learning techniques with self-monitoring, self-learning, self-consciousness, and self-improvement abilities. We start from the assumption that these abilities, typical of biological entities, can be partially implemented (at a very basic level) in artificial agents [15,16]. In a previous work [17], we explained how all these abilities can be included in practice into the architecture of deep neural networks. We named this approach "self-aware deep learning" (or, briefly, SAL). In the specific case discussed in this paper, we apply the SAL approach to time-lapse datasets acquired through continuously monitoring geophysical methods aimed at mapping fluids displacements during production.

The SAL network can dynamically adjust its own architecture (hyper-parameters) to maximize its performance in response to changing conditions in various operational environments. We remark that this dynamic ability is a key feature of AGI systems that try to emulate dynamic problem-solving capabilities of natural intelligence. Figure 2 shows the workflow of the SAL approach.

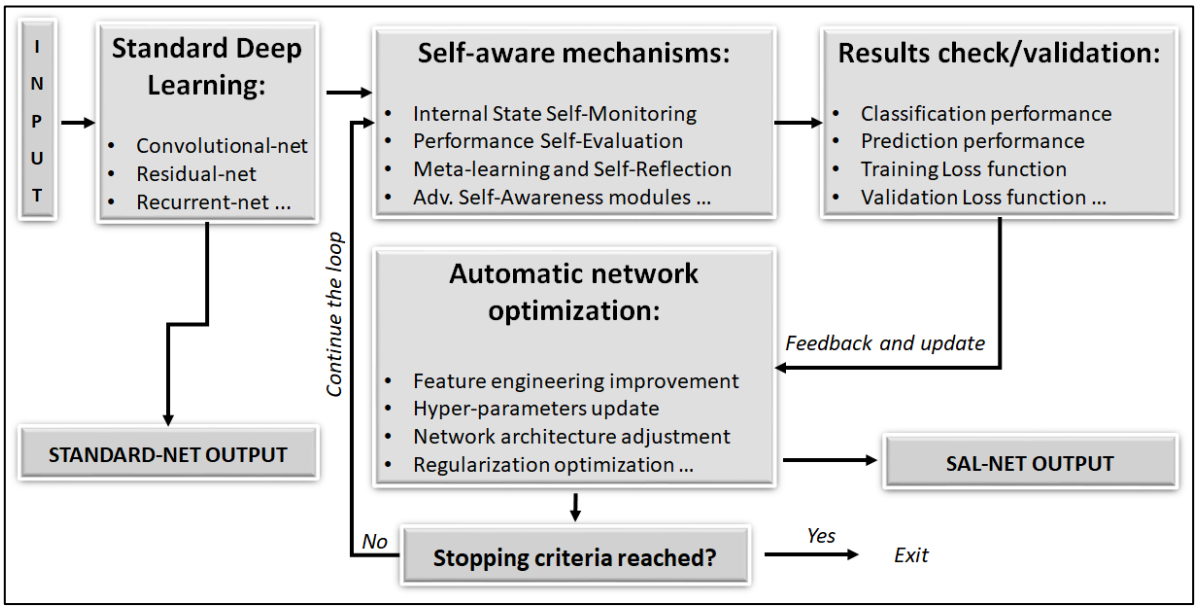

Figure 2. Block diagram of the SAL workflow addressed to self-learning and self-improvement in deep neural networks. (Obtained from [17].)

Initially, conventional neural networks process input data and produce an initial output (see left part of Figure 2). These "standard neural network outputs" are used as *reference results* and for comparison purposes. Self-aware mechanisms (such as performance self-evaluation, meta-learning algorithms, and so forth) continuously monitor the network's performance (results check/validation block), aiming to autonomously evaluate real-time outcomes and adjust the network's architecture (automatic network optimization block) to improve results. In this context, improving the results means decreasing validation loss and reducing overfitting with respect to the standard network output. This iterative process continues until certain termination conditions (stopping criteria) are met. These may include reaching a maximum iteration limit, achieving hyper-parameter convergence, enhancing performance, encountering resource constraints, and so on.

The second crucial aspect of our workflow consists of integrating self-aware deep learning (SAL) with reinforcement learning (RL) techniques. Among the many available RL methods, in this paper we selected the well-known Q-Learning algorithm, which we recall briefly. The name of this method comes from the Q-function that provides a measure of the quality (in terms of effectiveness for a given task) of an action, a, that the agent takes starting from a certain state, s. It is defined as follows:

$$Q(s, a)=\mathrm{S}\times\mathrm{A}\rightarrow\mathrm{R}. \tag{1}$$

The Bellman equation provides an operative definition of the maximum cumulative reward:

$$Q(s,a) = r + \gamma\, max_{a'}\, Q(s',a'). \tag{2}$$

It is given by the reward r that the agent received for entering the current state s and action a, plus the maximum future reward for the next state s', taking all the possible actions a' from that state. The symbol γ represents the discount factor. It has the crucial role of balancing the contribution of future rewards with respect to the immediate reward. The value of $Q(s, a)$ is found recursively, through the following procedure. First, the algorithm starts by using random values (or any type of guess value) for the *Q-function*. Next, while the agent explores its environment, the initial Q values progressively converge toward the optimal ones, based on the positive and/or negative feedback that the agent receives from its environment.

Such an integration between SAL and RL creates a comprehensive approach that enhances the adaptability and performance of artificial agents. In our combined methodology, self-awareness mechanisms are embedded within the RL framework, enabling agents to monitor their internal states, evaluate their performance, and dynamically adjust various aspects of their behavior. In summary, self-awareness is focused to the network architecture, enabling agents to modify their structure, such as the number of hidden layers or learning rates, in response to changing conditions. Meanwhile, reinforcement learning provides a framework for agents to interact with their dynamic environment, learn from experiences, and optimize their actions based on reward signals (Figure 3). In the following section, we are going to explain how such a novel deep learning workflow is applied to support time-lapse electric/electromagnetic inversion/monitoring.

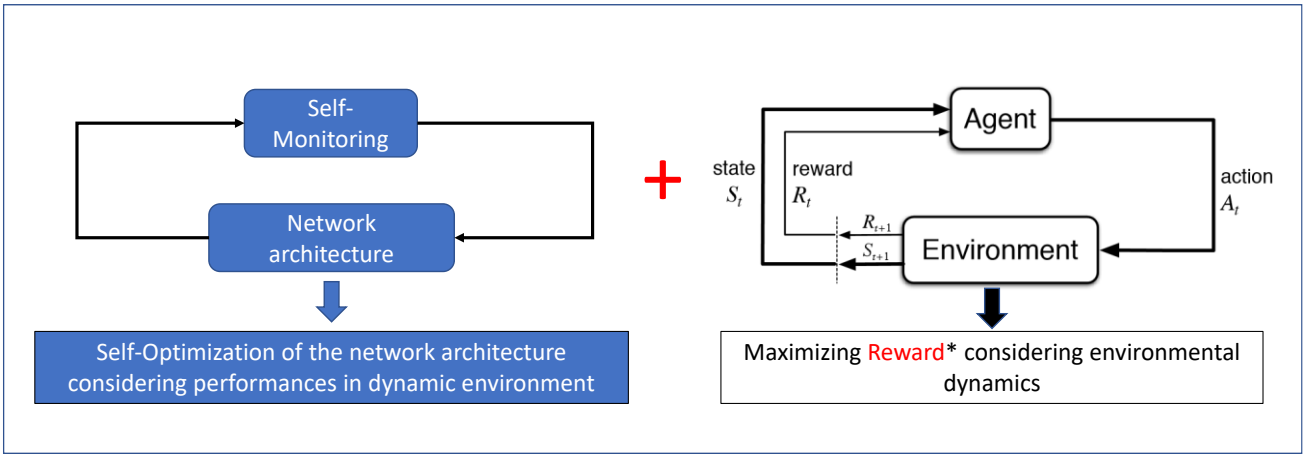

Figure 3. Synergy between self-aware deep learning (left blocks) and reinforcement learning (right blocks). In this scheme, intelligent agents exhibit self-improvement of their own architecture and, at the same time, also adapt to environmental dynamics.

3. EM monitoring supported by AGI/RL system for oil production

Our continuous geophysical monitoring approach involves the permanent installation of electrodes and magnetometers within a production well, establishing a persistent system that enables the ongoing detection and analysis of variations in electromagnetic fields associated with fluid displacement. Installing magnetometers and electrodes in a production well is not a simple task. It can work in practice only if appropriate completion technology is applied [14].

This innovative approach allows comprehending the dynamic changes occurring in the subsurface environment surrounding the well. For instance, borehole electric tomography, often combined with

surface electrode layouts, has been widely applied over the past few years for continuous monitoring of fluid displacement (water and oil, as well as CO_2) in the subsoil [18–29].

Um et al. [18] discuss a complex approach based on multiphysics network supported by deep learning techniques. The entire method is designed for real-time imaging of CO_2 saturation in geological carbon storage. In their workflow, the authors combine seismic, electromagnetic, and other data to improve imaging accuracy and reduce uncertainties, providing an advanced tool for monitoring CO_2 storage.

Bergmann et al. [19,20] focus on the application of time-lapse surface-downhole electrical resistivity tomography (ERT) for monitoring CO_2 storage at the Ketzin site in Germany. Their research highlights the effectiveness of ERT in detecting CO_2 migration and monitoring storage integrity.

Binley [21] and Binley & Kemna [22] review tools and techniques related to DC electrical and induced polarization methods, particularly focusing on their applications in geophysics and hydrogeophysics. The authors cover the principles and practices of direct current resistivity methods and their role in subsurface investigations, for applications to monitoring as well as exploration problems.

Buscheck et al. [23] discuss the strategy of pre-injection brine production to manage pressure in CO2 storage reservoirs. The approach is aimed at enhancing CO_2 storage capacity and ensuring the stability of storage sites.

Christensen et al. [24] provide an additional contribution to the discussion about the use of cross-hole electrical resistivity tomography (ERT) for monitoring CO_2 injection. Their study illustrates the capability of ERT in providing detailed images of subsurface CO_2 movement during injection.

Descloitres et al. [25] apply time-lapse resistivity mapping to investigate infiltration in a Sahelian gully erosion area. The authors show how resistivity mapping can be used to study water infiltration processes in erosion-prone regions.

Goldman & Kafri [26] explore hydrogeophysical applications in aquifers, focusing on methods like electrical resistivity for studying groundwater in vulnerable coastal environments. The research emphasizes the importance of geophysical techniques in managing coastal water resources.

Schmidt-Hattenberger et al. [27] discuss the development and application of electrical resistivity tomography (ERT) for monitoring CO_2 migration. Like Bergmann et al., these authors focus the discussion on the results obtained at the Ketzin pilot site. In this specific work, the authors cover the progression from tool development to practical reservoir surveillance.

Kazakis et al. [28] apply electrical resistivity tomography (ERT) and hydrochemical data to map seawater intrusion in the coastal area of the eastern Thermaikos Gulf, Greece. Their work demonstrates the effectiveness of ERT in detecting and monitoring saltwater intrusion in coastal aquifers.

Tarrahi & Afra [29] discuss the optimization of geological carbon sequestration in heterogeneous saline aquifers. The study focuses on managing CO_2 injection to achieve uniform distribution within the storage reservoir, enhancing the efficiency of carbon sequestration.

In the above-mentioned research, a key message emerges: the electrodes play a crucial role in capturing electrical signals that result from the movement of fluids causing resistivity variations in the subsurface and around the wellbore. As these fluids are displaced, the electric fields are perturbed, and the electrodes, strategically positioned, provide real-time data on these alterations. This comprehensive monitoring system offers useful insights into the subsurface fluid dynamics, aiding in the understanding of reservoir behavior, fluid migration patterns, and potential geo-hazards.

Similarly, magnetometers contribute to this continuous monitoring process by measuring changes in the magnetic fields induced by fluid displacements and the consequential variations in electric

resistivity distribution over space and time. This dual-sensor system, comprising both electrodes and magnetometers, forms a robust network that enhances the accuracy and reliability of the gathered geophysical data. By leveraging these technologies, scientists and engineers gain the ability to track and interpret subsurface processes with a high degree of precision, enabling them to make informed decisions about well operations, reservoir management, and environmental impact assessments.

In this geophysical monitoring frame, artificial general intelligence combined with RL techniques can play a crucial role in analyzing/interpreting/using data from electric/electromagnetic sensors. Let us better clarify these concepts with the support of illustrative synthetic tests.

4. Synthetic tests

In this section, we delve into a simulation test designed to elucidate the core principles of our approach and its significance within the realms of geosciences and the oil industry. Our aim is to implement and apply some of the key elements of an AGI system through a software platform based on the integration of self-aware mechanisms with Q-learning. We discuss a synthetic test with the specific goal of showing how we can optimize oil production from a single horizontal well while mitigating the risk of water intrusion in the borehole. The section is divided into two sub-sections. The first part is dedicated to the simulation of a DC (geoelectric) experiment performed in a horizontal well. The second part is focused on the analysis of the simulated data through our AGI/RL approach. This provides prescriptive rules for optimizing oil production and, at the same time, reducing the risk of water invasion in the well itself.

4.1. Data simulation

We generate time-dependent data pertaining to a) waterfront distances from the production well and b) the rate of oil production.

We simulate a scenario in which waterfront distances, variable over time, have been obtained through direct current (DC) borehole tomography. For that scope, we used the open-source software package R2, based on Python libraries [30,31]. Additionally, the pyres Python wrapper package for R2 was employed, providing a robust and adaptable programmatic interface for the modeling and inversion of electrical resistivity tomography (ERT) datasets. The inverse solution method is grounded in a regularized objective function, incorporating weighted least squares principles akin to an "Occam's" type solution.

Our simulation introduces oscillations in waterfront distances, contaminated by random noise, with the risk of water invasion being derived from these distances (wherein risk escalates as the distance diminishes). To aid comprehension, in this specific synthetic test, both waterfront distances and oil production data are standardized within the range of 0 to 1.

As already discussed in a previous paper [32], Figure 4 illustrates our schematic 2D scenario wherein the water table, marked by dashed lines and with a resistivity of 1 Ωm, exhibits an upward movement following an irregular dynamic trend. This scenario mirrors typical situations wherein overproduction of oil or water flooding techniques lead to undesired "cresting" or "coning" phenomena. In our simulation, we deploy an electrical resistivity tomography (ERT) permanent layout, with electrode spacing set at 10 m, within a horizontal well situated at a relative elevation of 0 m. Our objective is to monitor water movements to prevent any potential invasion within the producing well.

Through repeated monitoring surveys akin to this setup, we compile a dynamic database comprising electric potential measurements over time.

Figure 5 shows the apparent resistivity pseudo-section obtained through the monitor ERT survey at time T0, when the water table is marked by the black dashed curve in the previous figure.

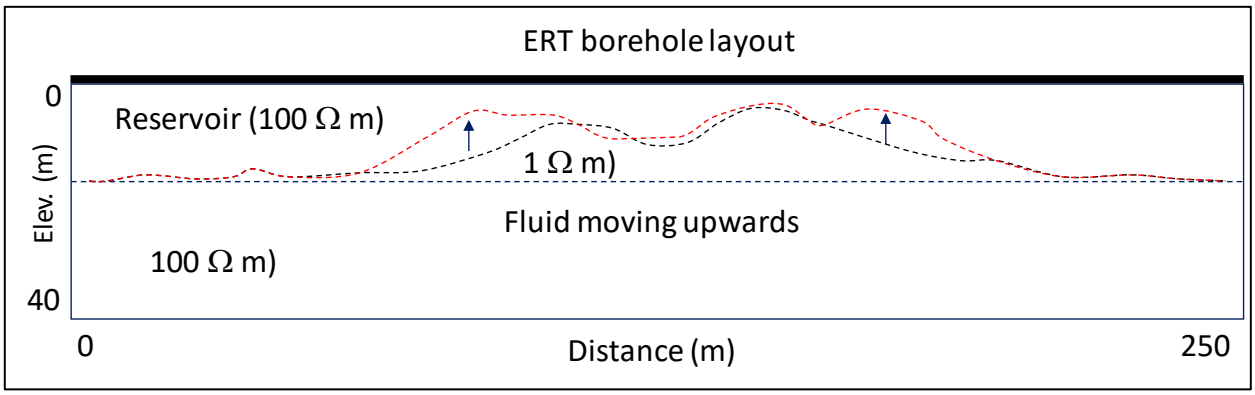

Figure 4. Simulation scenario. Black dashed curve: waterfront at time step T0. Red dashed line: waterfront at a successive time step. The vertical axis represents a relative depth with respect to the horizontal well (thick black line). (Obtained from [32].)

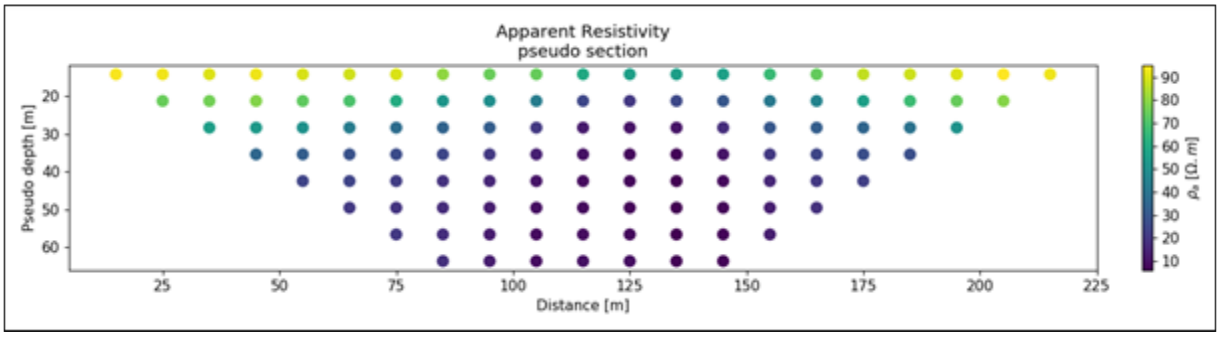

Figure 5. Apparent resistivity pseudo-section at time step T0. (Obtained from [32].)

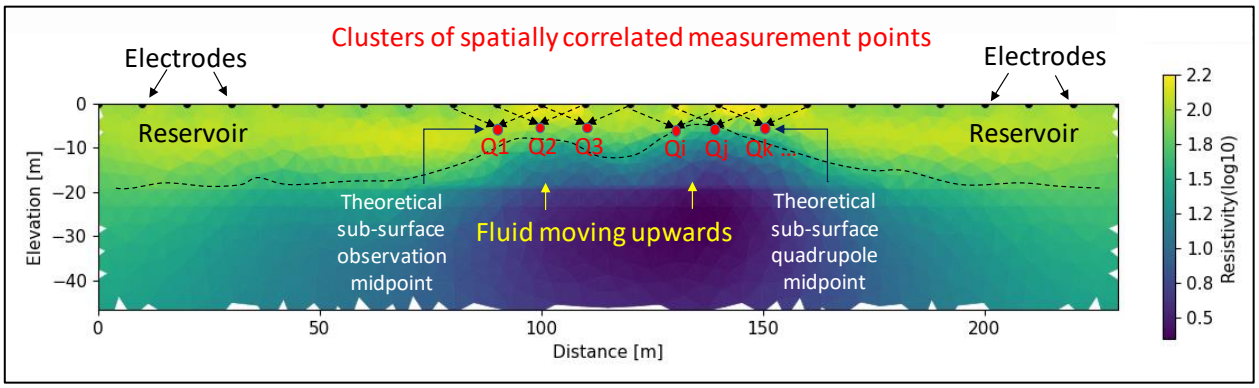

Figure 6. Example of ERT model (obtained through inversion of the pseudo-section of figure 5) in a scenario of detection of fluids below the measurement points. (Obtained from [32].)

Figure 6 shows the ERT section obtained by inversion of the apparent resistivity pseudo-section of Figure 5.

We created many scenarios like those depicted in the previous figures, simulating an oscillatory trend of the waterfront distance from the well induced by oil production over time.

4.2. Prescriptive AGI data analysis

To optimize oil production considering the variable risk of water invasion directly linked with oil production rate, we applied our AGI approach, based on a combination of RL and self-aware mechanisms, properly implemented in a dedicated deep learning architecture. Let us explain more in detail such an algorithm.

In the frame of the RL context, we selected a Q-Learning technique [33–37]. First, we set up Q-learning parameters, including the number of episodes, learning rate, discount factor, and exploration rate (epsilon). The Q-Learning algorithm is then trained over a fixed number of episodes specified by the user; at each episode, the agent begins at the initial state (time step 0) and interacts with the environment until the end of the data sequence. It utilizes an epsilon-greedy policy [37] to select actions (increasing or decreasing production) based on Q-values. The agent executes the chosen action and observes the resulting reward, which is positive for increasing production while maintaining a stable waterfront distance, and negative for decreasing production and decreasing waterfront distance. Subsequently, the agent moves to the next state (time step) based on its selected action, with Q-values being updated using the Bellman equation [37].

Figure 7 illustrates how the Q-Learning algorithm, combined with self-aware deep neural networks (that continuously optimizes the network architecture), automatically regulates oil production, considering the oscillations of the waterfront. The left graph shows the total reward obtained by the agent in each episode, offering insight into the agent's performance improvement over time. The graph on the right displays the normalized waterfront distance and oil production values across iterations, showing how these variables evolve based on the agent's actions during the training process. This example refers to the oscillations of the waterfront (and of the oil production rate) observed locally at the center of the measurement system (to clarify, in Figure 6, this point corresponds to the distance x = 125 m). Of course, similar plots are produced for all the other observation points where we take measurements of the electric potential, waterfront distance, and oil production rate.

Looking at Figure 7, a simple correlation function between water table distance and oil production might appear sufficient for our optimization purposes. That is reasonable in an idealized setting. Furthermore, we remark that this example considers the curves calculated as we had only one measurement point along the well. In real situations, the optimized adaptation of the network architecture as well as the use of Q-Learning techniques ensure that the network structure remains flexible and capable of handling more complex scenarios, where several production sectors and many observation points can operate simultaneously along the borehole. In these cases, simple correlation functions could not be sufficient for our goals. Through the Q-Learning approach, the artificial agent can learn from prolonged experience in dynamic scenarios, optimizing production in various sectors of the well according to a complex relationship. The final number of hidden neural layers after convergence provides insight into the network's learned complexity; while two layers may be sufficient for the given simple simulation, real-world scenarios may benefit from additional layers to better capture and model complex relationships.

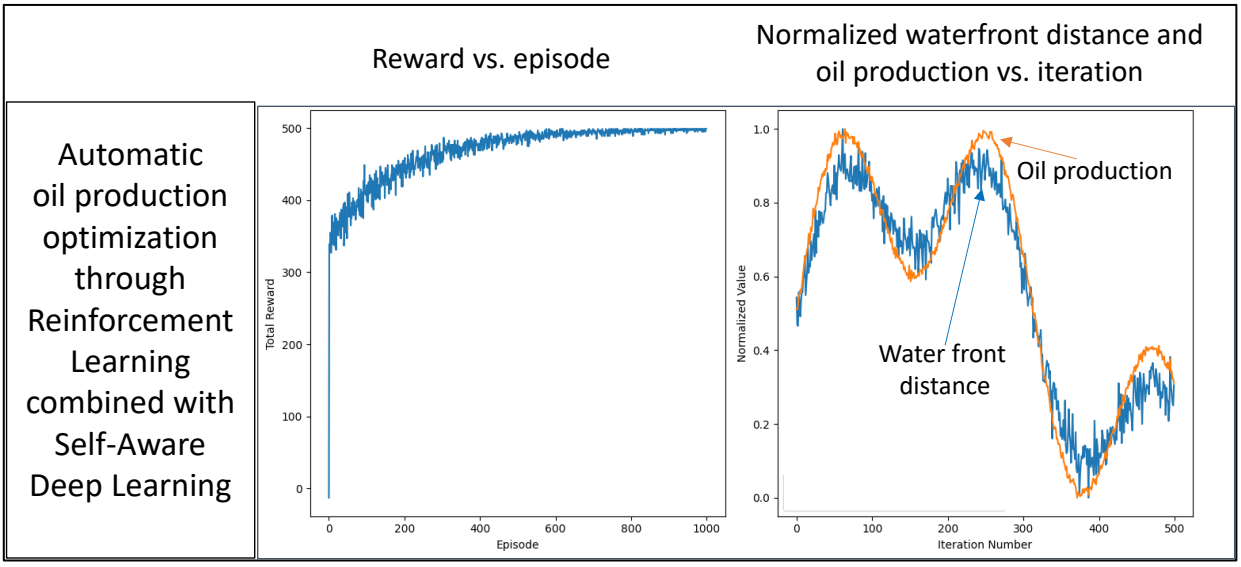

Figure 7. The Q-Learning algorithm, combined with self-aware deep neural networks, automatically regulates oil production, considering the oscillations of the waterfront estimated at a given observation point in the well.

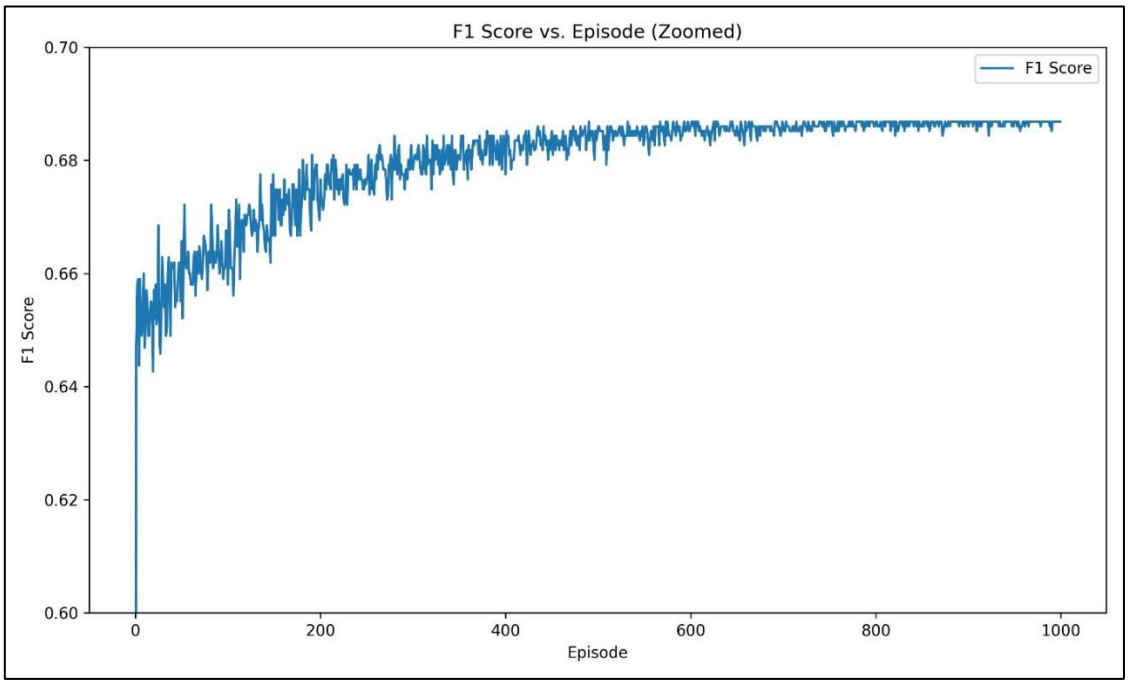

Figure 8. F1-score zoomed in to the range between 0.6 and 0.7.

Furthermore, it is worth noting that an "episode" signifies a single run or trial of the reinforcement-learning agent in the environment. Within the context of our algorithm, an episode starts with the agent at the initial state (typically, time step 0) and proceeds with a sequence of decisions (actions) until the end of the data sequence is reached, with the aim of optimizing its actions over this sequence. The ultimate objective is to increase oil production while maintaining constant or increasing the water table distance. Throughout an iteration, the agent selects an action, observes the resulting

state transition, receives a reward, and updates its knowledge (e.g., Q-values in Q-learning). In our code, an "iteration" corresponds to the progression from one time step to the next within an episode, during which the agent decides whether to increase or decrease production based on its current state (water table distance and associated risk of water invasion) and policy, repeating this decision-making process until the episode concludes.

Finally, in order to estimate the performance variations of the neural network, Figure 8 shows the trend of a well-known performance parameter, the F1-score, zoomed in to the range between 0.6 and 0.7.

5. Further improvements and tests

Next, we replicated the same type of test, but increasing the complexity of the simulated scenario. In this new test, we simulated a more complex trend of waterfront oscillations (at a single observation point). Consequently, the production trend becomes more complex, because the AGI/RL system tries to adapt it to the waterfront variations to maintain a safe distance of undesired water from the well, preventing possible phenomena of cresting or coning. Furthermore, to face such an increment of complexity in the data, we improved some aspects of the code. For instance, we refined self-reflection and self-learning mechanisms: the epsilon decay schedule was adjusted to decay over episodes, allowing for more initial exploration and gradually reducing it over time. In addition, we applied upper and lower bounds to the learning rate, ensuring it stays within reasonable limits during adjustments. Moreover, the number of hidden layers was now bounded by minimum and maximum values, preventing the network from becoming too complex or too simple. Finally, the learning rate and hidden layer adjustments were now based on recent performance relative to a moving average of rewards, leading to more informed adjustments. Figure 9 shows the reward trend and the normalized oscillations of waterfront and oil production for this new case.

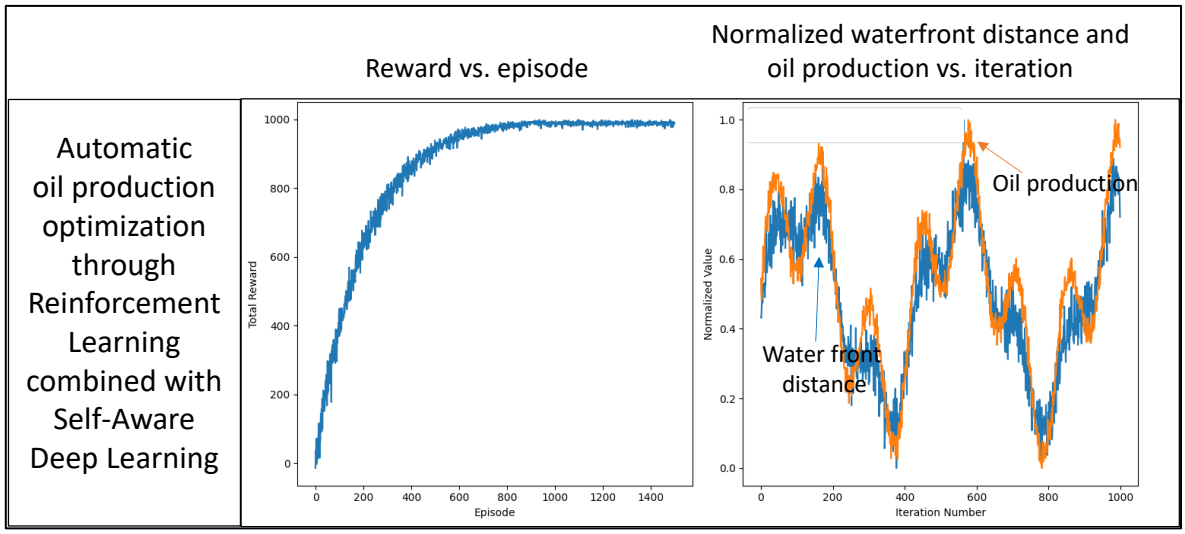

Figure 9. Reward vs. episodes (left) and oscillations of the waterfront-oil production oscillations vs. iteration number (right).

6. Why is this an AGI-inspired approach?

The approach here applied through these tests incorporates many of the crucial mechanisms of an AGI system. These mechanisms enable an artificial agent to autonomously evaluate its performance, adjust neural network architecture, and modify learning parameters based on recent experiences, simulating the key aspects of a general form of operative intelligence. By dynamically adapting to changing conditions, the system exhibits a level of self-reflection and self-improvement characteristic of a typical AGI, allowing it to continuously enhance its performance over time without human intervention. Let us highlight these aspects in more detail, with reference to our tests.

Generalization: The self-adjustment mechanism enables the system to generalize its knowledge and adaptability to a wide range of situations. By continuously evaluating its performance and adjusting parameters, the system learns to generalize effective strategies for optimizing oil production across diverse reservoir conditions and operational scenarios.

Learning ability: The system's ability to adjust autonomously neural network architecture and learning parameters reflects its capacity for continuous learning. Through experience-driven adaptation, the system improves its understanding of complex reservoir dynamics and refines its decision-making abilities over time.

Reasoning and problem-solving: By dynamically modifying its internal structure and parameters based on recent performance, the system exhibits sophisticated reasoning and problem-solving capabilities. It can analyze complex production challenges, infer relationships between different variables, and formulate effective strategies to optimize oil production while mitigating risks.

Autonomy: The self-adjustment mechanism empowers the system with a high degree of autonomy in decision-making and adaptation processes. Without requiring human intervention, the system autonomously evaluates its performance, adjusts its strategies, and enhances its effectiveness in optimizing oil production in real time.

Self-improvement: Through iterative self-adjustment and learning, the system engages in continuous self-improvement. By tracking adjustments and analyzing performance metrics, the system identifies areas for enhancement and iteratively refines its strategies, algorithms, and neural network architecture to achieve better outcomes over time.

Adaptability: The system's ability to adjust dynamically its internal parameters and strategies in response to changing conditions demonstrates its adaptability. It can adapt to fluctuations in reservoir dynamics, production demands, and environmental factors, ensuring robust performance and optimization in diverse operational environments.

Perception and sensing: The system's integration with geophysical monitoring and inversion techniques implies a capability for perception and sensing. By analyzing data from sensors and inverting/interpreting geophysical signals, the system perceives changes in reservoir conditions and senses potential risks or opportunities for optimization.

Continuous learning: The iterative nature of the self-adjustment mechanism facilitates continuous learning and adaptation. By leveraging experiences and performance feedback, the system engages in ongoing learning processes, continuously refining its understanding and strategies for optimizing oil production in dynamic operational contexts.

All these capabilities are particularly useful for optimizing oil production because they enable the system to respond to (and predict) evolving reservoir dynamics and production challenges. By considering factors such as water intrusion risk, our approach can allow informed decisions to regulate

production strategies effectively. The integration of self-adjustment mechanisms ensures that the system remains adaptive and responsive in real time, ultimately leading to improved production efficiency and risk mitigation. Thus, this approach demonstrates the potential of AGI-inspired techniques in enhancing oil production management by facilitating autonomous decision-making and adaptation in complex operational environments.

7. Discussion: benefits, limitations, and next steps

The integration of AGI, RL, and SAL with continuous geophysical monitoring represents a significant advancement in the field of oil production and reservoir management. By combining these techniques in real-time data analysis and decision-making with electromagnetic monitoring techniques, our approach offers a comprehensive solution to optimize oil production while minimizing operational risks, particularly with reference to unwanted water invasion.

Compared to traditional (non-AI) methods used for the optimization of oil-producing wells, our method offers several significant advantages. The first important benefit is that traditional methods typically rely on periodic manual data collection and analysis, which can result in delays in decision-making. Instead, our approach provides continuous, real-time monitoring and analysis of geophysical data, enabling immediate adjustments to production strategies. This ensures more responsive and timely decision-making.

A second advantage concerns adaptive learning and self-improvement. Traditional methods usually follow predefined rules and procedures that may not adapt well to changing conditions. On the other side, our method employs self-adjusting algorithms that learn from new data and experiences, continuously improving their performance and adapting to evolving reservoir dynamics without human intervention.

This implies a significant step forward in risk management. In fact, in traditional methods, risk assessments and mitigation strategies are often based on historical data and fixed models, which may not accurately predict future events. Instead, our approach utilizes advanced machine learning algorithms to predict and manage risks such as unwanted fluid invasion (e.g., water ingress) in real time. This proactive approach helps prevent potential issues before they escalate.

Additional benefits include higher accuracy and precision, increased efficiency and optimization, comprehensive monitoring system, scalability, and generalization using AGI-inspired criteria.

On the other side, there are limitations in our methodology. These include a possible open problem of complex implementation. In fact, implementing AGI systems and continuous geophysical monitoring infrastructure may require significant initial investment, new technologies of well completion, and novel expertise.

Furthermore, the effectiveness of the approach heavily relies on the quality of geophysical data collected and the accuracy of AGI algorithms in inverting/interpreting this data. That is true, especially considering the difficulties in recording borehole electric and electromagnetic data with a good signal/noise ratio at the depth of kilometers, in oil or gas reservoirs. Potential issues include signal distortion from well casings. This can interfere with data accuracy, requiring advanced signal processing for correction. Furthermore, low resistivity contrasts between reservoirs and water tables can complicate fluid differentiation, necessitating the use of enhanced inversion algorithms and additional geophysical data. In addition, fluid properties that change over time can affect measurement reliability, demanding adaptive algorithms and regular recalibration. Environmental factors and

external noise can distort measurements, which calls for robust noise-filtering techniques and high signal-to-noise ratio designs. Finally, integrating diverse data sources and interpreting them accurately is complex, necessitating advanced data fusion methods and AGI's pattern recognition capabilities. Addressing these limitations effectively is crucial for ensuring accurate, reliable, and actionable insights in oil production optimization.

An additional discussion is required about the synthetic tests presented in this paper. These serve as valuable demonstration of the proposed methodology's efficacy in optimizing oil production and mitigating the risk of water invasion. By simulating scenarios of fluid dynamics and employing a combination of reinforcement learning and self-aware mechanisms within a deep learning architecture, the tests show how our approach can autonomously regulate production strategies to maintain optimal conditions. In fact, the results illustrate the adaptability and learning capabilities of the entire system, as it successfully navigates complex scenarios and optimizes production while considering the risk of water intrusion. Moreover, the inclusion of self-reflection mechanisms further enhances the system's performance over time, highlighting the potential for continuous improvement and adaptation in real-world applications. However, these tests are based on a simple 2D simulation, and additional tests and experiments on real data are necessary for validating the entire methodology.

Moving forward, several avenues for future development and research emerge from this work. As anticipated earlier, conducting field trials to validate the effectiveness of the proposed methodology in real-world oil production environments would provide valuable insights and practical feedback. Furthermore, further refining AGI algorithms to improve their accuracy and efficiency in interpreting geophysical data could enhance the system's performance and reliability.

Of course, many additional applications can be done in different fields using a similar approach. In a multi-disciplinary scenario, other authors have successfully applied fruitful cross-disciplinary approaches that combine geophysics and advanced machine learning algorithms in various Earth disciplines, including mineral automatic classification, geothermal exploration, reservoir characterization, CO_2 storage and multiphysics monitoring, and integrated geophysical field analysis [38–41].

8. Conclusions

This article underscores the transformative potential of artificial general intelligence (AGI) and reinforcement learning (RL) in improving oil production management. By integrating geophysical data monitoring, interpretation, continuous learning, and dynamic production control, AGI and RL offer a paradigm shift in reservoir management practices.

The simple simulated examples presented in this study serve as theoretical illustrations of our approach in optimizing oil production while minimizing the risk of unwanted fluid invasion. Through advanced machine learning algorithms and self-adjustment mechanisms, the entire system autonomously regulates production strategies in real time, adapting to changing reservoir dynamics and environmental conditions.

In conclusion, this study shows the potential of geophysical monitoring and AGI/RL integrated technology to be a game-changer in oil production management. As the industry continues to embrace digital transformation and automation, our approach offers a pathway toward more sustainable, efficient, and resilient reservoir management practices.

9. Appendix: Tutorial Python code

To provide further insight in our approach, we prepared a sequence of tutorial Python codes with didactical purposes. This tutorial represents just a simplified version of our AGI software platform. It is shared exclusively for illustrative (neither commercial nor industrial) purposes. The entire Jupyter notebook can be requested directly from the author (paolo.dell'aversana@eni.com). In the following, we include an example of the code, extracted from the notebook. This can be copied by the users into their Python environment for running didactical tests.

The code simulates the data to be used for tests like the one discussed in the paper. We remark again that the rationale behind the method is to use reinforcement learning, specifically Q-learning, to optimize oil production while minimizing the risk of water invasion in the borehole. By training an agent to learn optimal actions based on simulated data, the approach aims to automate decision-making processes and improve production efficiency. The inclusion of self-reflection and self-adjustment mechanisms further enhances the agent's ability to adapt to changing conditions and optimize its actions over time.

In this tutorial code, the input data consist of a set of distances of the water table from the production well (here, we skip the step that simulates the measurements of the electric potentials as well as the oil production rate). The following are the main blocks of the code.

- Simulated data generation: Simulates time-dependent data for waterfront distances and oil production. The waterfront distances exhibit complex oscillation with added random noise, while oil production is initialized, and then it is determined based on the risk of water invasion, which is calculated from the waterfront distances and using Q-Learning technique.
- Normalization: This block normalizes the simulated data between 0 and 1 to facilitate training.
- Q-Learning setup: This block initializes Q-learning parameters such as the number of episodes, learning rate, discount factor, and exploration rate (epsilon) and sets up a Q-table to store Q-values for state-action pairs.
- Training with Q-Learning: This block executes Q-learning training with self-reflection and self-adjustment mechanisms. The agent interacts with the environment (simulated data) over multiple episodes. At each step, the agent selects an action (increase or decrease production) based on an epsilon-greedy policy. The agent receives a reward based on the action taken and updates its Q-values accordingly.
- Self-reflection and self-adjustment: Here, the code periodically evaluates the agent's performance and adjusts the neural network architecture (number of hidden layers), learning rate, and exploration rate accordingly. This adjustment enables the agent to adapt to changing conditions and improve its performance over time without manual intervention.
- Plotting: Finally, the code allows visualizing the training progress and results using matplotlib. It plots the total reward obtained during training and the normalized waterfront distance and oil production over the iterations.

Code example

```
# This code is provided just for illustrative purposes.
# The AGI-RL workflow is implemented at a very basic level, only for didactical scope.
# Additional, more complex tutorial codes are available on demand (paolo.dell'aversana@eni.com)
```

```
import numpy as np
import tensorflow as tf
import matplotlib.pyplot as plt

# Simulated data
time_steps = np.linspace(0, 10, num=500)
water_front_distances = (
    150 + 20 * np.sin(0.5 * time_steps) + 10 * np.sin(1.5 * time_steps) +
    np.random.normal(0, 2, size=500)
)
oil_production = (
    100 - 50 * np.exp(-0.02 * water_front_distances) + 20 * np.sin(0.5 * time_steps) +
    15 * np.sin(1.5 * time_steps) + np.random.normal(0, 0.5, size=500)
)

# Normalize data between 0 and 1
normalized_water_front_distances = (
    (water_front_distances - water_front_distances.min()) /
    (water_front_distances.max() - water_front_distances.min())
)
normalized_oil_production = (
    (oil_production - oil_production.min()) /
    (oil_production.max() - oil_production.min())
)

# Initialize Q-learning parameters
num_episodes = 1000
learning_rate = 0.1
discount_factor = 0.8
epsilon = 0.3

# Initialize Q-table
num_actions = 2    # Two actions: Increase production (0) or decrease production (1)
Q_table = np.zeros((len(normalized_water_front_distances), num_actions))

# Initialize self-adjustment parameters
hidden_layers = [10]    # Initial hidden layer configuration
self_adjustment_rate = 0.1

# Create lists for tracking rewards, exploration rate, self-adjustment,
# learning rate adjustments, and hidden layer adjustments
rewards = []
exploration_rates = []
```

```
self_adjustments = []
learning_rate_adjustments = []
hidden_layer_adjustments = []

# Q-learning training with self-reflection and self-adjustment
for episode in range(num_episodes):
    state = 0    # Starting state
    total_reward = 0

    while state < len(normalized_water_front_distances) - 1:
        # Epsilon-greedy policy
        if np.random.uniform(0, 1) < epsilon:
            action = np.random.choice(num_actions)
        else:
            action = np.argmax(Q_table[state])

        # Apply the chosen action
        next_state = state + 1

        # Calculate reward based on the change in water table distance and oil production
        if action == 0:    # Increase production
            if normalized_water_front_distances[next_state] >=
normalized_water_front_distances[state]:
                reward = 1
            else:
                reward = -1
        else:    # Decrease production
            if normalized_water_front_distances[next_state] <
normalized_water_front_distances[state]:
                reward = 1
            else:
                reward = -1

        # Q-learning update
        Q_table[state, action] += learning_rate * (
            reward + discount_factor * np.max(Q_table[next_state]) -
            Q_table[state, action]
        )

        total_reward += reward
        state = next_state

    rewards.append(total_reward)
    exploration_rates.append(epsilon)
```

```
        # Self-reflection and self-adjustment
        if episode > 0 and episode % 50 == 0:
            # Check recent performance and adjust the neural network architecture
            if total_reward > np.mean(rewards[-50:]):
                hidden_layers.append(10)  # Increase the number of neurons in the hidden layer
                hidden_layer_adjustments.append(1)  # 1 indicates an increase in hidden layers
            else:
                if len(hidden_layers) > 1:
                    hidden_layers.pop()  # Reduce the number of neurons in the hidden layer
                    hidden_layer_adjustments.append(-1)  # -1 indicates a decrease in hidden layers
                else:
                    hidden_layer_adjustments.append(0)  # 0 indicates no adjustment

            # Adjust the neural network learning rate based on self-reflection
            if total_reward > np.mean(rewards[-50:]):
                learning_rate *= (1 + self_adjustment_rate)  # Increase learning rate
                learning_rate_adjustments.append(1)  # 1 indicates an increase in learning rate
            else:
                learning_rate_adjustments.append(0)  # 0 indicates no adjustment

            self_adjustments.append(
                1 if total_reward > np.mean(rewards[-50:]) else 0
            )  # 1 indicates an adjustment, 0 indicates no adjustment
        else:
            self_adjustments.append(0)  # 0 indicates no adjustment

        epsilon *= 0.995  # Decay epsilon for exploration over time

# Plotting results
plt.figure(figsize=(12, 6))
plt.subplot(1, 2, 1)
plt.plot(range(num_episodes), rewards)
plt.xlabel('Episode')
plt.ylabel('Total Reward')
plt.title('Total Reward vs. Episode')

plt.subplot(1, 2, 2)
plt.plot(normalized_water_front_distances, label='Normalized Water Front Distance')
plt.plot(normalized_oil_production, label='Normalized Oil Production')
plt.xlabel('Iteration Number')
plt.ylabel('Normalized Value')
plt.title('Normalized Water Table Distance and Oil Production vs. Iteration Number')
plt.legend()
```

```
plt.tight_layout()
plt.show()
```

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Conflict of Interest

No conflict of interest.

References

1. Wang P, Goertzel B (2007) Introduction: Aspects of Artificial General Intelligence. Artificial General Intelligence: Concepts, Architectures and Algorithms. IOS Press, 1–16.
2. Goertzel B, Pennachin C (2007) *Artificial General Intelligence*, Springer Science and Business Media.
3. Goertzel B (2021) The General Theory of General Intelligence: A Pragmatic Patternist Perspective. *arXiv:2103.15100*.
4. Franz A, Gogulya V, Löffler MW (2019) William: A monolithic approach to agi. International Conference on Artificial General Intelligence, Springer International Publishing, 44–58.
5. Franz A (2015) Artificial general intelligence through recursive data compression and grounded reasoning: a position paper. *arXiv preprint arXiv:1506.04366*.
6. Al-Sikaiti SH, Regtien J (2008) Challenging Conventional Wisdom, Waterflooding Experience on Heavy Oil Fields in Southern Oman. Canada.
7. Chierici GL, Ciucci GM, Pizzi G (1964) A Systematic Study of Gas and Water Coning By Potentiometric Models. *J Pet Technol* 16: 923–929. http://dx.doi.org/10.2118/871-PA
8. Chaperon I (1986) Theoretical Study of Coning Toward Horizontal and Vertical Wells in Anisotropic Formations: Subcritical and Critical Rates. SPE Annual Technical Conference and Exhibition. https://doi.org/10.2118/15377-MS
9. Wheatley MJ (1985) An Approximate Theory of Oil/Water Coning. *SPE Annual Technical Conference and Exhibition*, USA. https://doi.org/10.2118/14210-MS
10. Karami M, Khaksar Manshad A, Ashoori S (2014) The Prediction of Water Breakthrough Time and Critical Rate with a New Equation for an Iranian Oil Field. *Pet Sci Technol* 32: 211–216. https://doi.org/10.1080/10916466.2011.586960
11. Bottazzi F, Dell'Aversana P, Molaschi C, et al. (2020) A New Downhole System for Real Time Reservoir Fluid Distribution Mapping: E-REMM, the Eni-Reservoir Electro-Magnetic Mapping System. International Petroleum Technology Conference. https://doi.org/10.2523/IPTC-19807-MS
12. Dell'Aversana P, Servodio R, Bottazzi F, et al., (2019) Asset Value Maximization through a Novel Well Completion System for 3d Time Lapse Electromagnetic Tomography Supported by Machine Learning. Abu Dhabi International Petroleum Exhibition & Conference. https://doi.org/10.2118/197573-MS

13. Dell'Aversana P, Rizzo E, Servodio R (2017) 4D borehole electric tomography for hydrocarbon reservoir monitoring. 79th EAGE Conference & Exhibition 2017. European Association of Geoscientists & Engineers. 1–5.
14. Dell'Aversana P, Bottazzi F, Molaschi C, et al. (2022) Monitoring and mapping system of the space-time distribution of formation fluids in a reservoir and a completion and production plant of a well for the extraction of formation fluids. U.S. Patent 11,268,374.
15. Chella A, Frixione M, Gaglio S (2008) A cognitive architecture for robot self-consciousness. *Artif Intell Med* 44: 147–154. https://doi.org/10.1016/j.artmed.2008.07.003
16. Dehaene S, Lau H, Kouider S (2017) What is consciousness, and could machines have it? *Science* 358: 486–492. https://doi.org/10.1126/science.aan8871
17. Dell'Aversana P (2024) Enhancing Deep Learning and Computer Image Analysis in Petrography through Artificial Self-Awareness Mechanisms. *Minerals* 14: 247. https://doi.org/10.3390/min14030247
18. Um ES, Alumbaugh D, Commer M, et al. (2024) Deep learning multiphysics network for imaging CO_2 saturation and estimating uncertainty in geological carbon storage. *Geophys Prospect* 72: 183–198. https://doi.org/10.1111/1365-2478.13257
19. Bergmann P, Schmidt-Hattenberger C, Kiessling D, et al. (2012) Surface-downhole electrical resistivity tomography applied to monitoring of CO_2 storage at Ketzin, Germany. *Geophysics* 77: B253–B267. https://doi.org/10.1190/GEO2011-0515.1
20. Bergmann P, Schmidt-Hattenberger C, Labitzke T, et al. (2017) Fluid injection monitoring using electrical resistivity tomography—five years of CO_2 injection at Ketzin, Germany. *Geophys Prospect* 65: 859–875. https://doi.org/10.1111/1365-2478.12426
21. Binley A (2015) Tools and Techniques: DC Electrical Methods. *Treatise Geophys* 11: 233–259. https://doi.org/10.1016/B978-0-444-53802-4.00192-5
22. Binley A, Kemna A (2005) DC Resistivity and Induced Polarization Methods, In: Rubin Y, Hubbard SS, Authors, *Hydrogeophysics,* Springer, 129–156.
23. Buscheck TA, White JA, Chen M, et al. (2014) Pre-injection brine production for managing pressure in compartmentalized CO_2 storage reservoirs. *Energy Procedia* 63: 5333–5340. https://doi.org/10.1016/j.egypro.2014.11.565
24. Christensen NB, Sherlock D, Dodds K (2006) Monitoring CO_2 injection with cross-hole electrical resistivity tomography. *Explor Geophys* 37: 44–49. https://doi.org/10.1071/EG06044
25. Descloitres M, Ribolzi O, Le Troquer Y (2003) Study of infiltration in a Sahelian gully erosion area using time-lapse resistivity mapping. *Catena* 53: 229–253. https://doi.org/10.1016/S0341-8162(03)00038-9
26. Goldman M, Kafri U (2006) Hydrogeophysical applications in coastal aquifers. *Applied Hydrogeophysics*, Springer: Dordrecht, 71: 233–254. https://doi.org/10.1007/978-1-4020-4912-5_8
27. Schmidt-Hattenberger C, Bergmann P, Bösing D, et al. (2013) Electrical resistivity tomography (ERT) for monitoring of CO_2 migration—from tool development to reservoir surveillance at the Ketzinpilot site. *Energy Procedia* 37: 4268–4275. https://doi.org/10.1016/j.egypro.2013.06.329
28. Kazakis N, Pavlou A, Vargemezis G, et al. (2016) Seawater intrusion mapping using electrical resistivity tomography and hydrochemical data. An application in the coastal area of eastern Thermaikos Gulf, Greece. *Sci Total Environ* 543: 373–387. https://doi.org/10.1016/j.scitotenv.2015.11.041

29. Tarrahi M, Afra S (2015) Optimization of Geological Carbon Sequestration in Heterogeneous Saline Aquifers through Managed Injection for Uniform CO_2 Distribution. Carbon Management Technology Conference. https://doi.org/10.7122/440233-MS
30. Befus KM (2018) Pyres: A Python Wrapper for Electrical Resistivity Modeling with R2. *J Geophys Eng* 15: 338–346.
31. Binley A (2016) R2 version 3.1 Manual. Lancaster, UK.
32. Dell'Aversana P (2021) Reservoir prescriptive management combining electric resistivity tomography and machine learning. *AIMS Geosci* 7: 138–161. https://doi.org/10.3934/geosci.2021009
33. Kaelbling LP, Littman ML, Moore AW (1996) Reinforcement Learning: A Survey. *J Artif Intell Res* 4: 237–285. https://doi.org/10.1613/jair.301
34. Raschka S, Mirjalili V (2017) *Python Machine Learning: Machine Learning and Deep Learning with Python, scikit-learn, and TensorFlow*, 2nd Eds., PACKT Books.
35. Ribeiro C, Szepesvári C (1996) Q-learning combined with spreading: Convergence and results, *Proceedings of the ISRF-IEE International Conference: Intelligent and Cognitive Systems (Neural Networks Symposium)*, 32–36.
36. Russell S, Norvig P (2016) *Artificial Intelligence: A Modern approach*, Global Edition, Pearson Education, Inc., publishing as Prentice Hall.
37. Ravichandiran S (2020) *Deep Reinforcement Learning with Python: Master classic RL, deep RL, distributional RL, inverse RL, and more with OpenAI Gym and TensorFlow*, Packt Publishing.
38. Ullah J, Li H, Ashraf U, et al. (2023) Knowledge-based machine learning for mineral classification in a complex tectonic regime of Yingxiu-Beichuan fault zone, Sichuan basin. *Geoenergy Sci Eng* 229: 212077. https://doi.org/10.1016/j.geoen.2023.212077
39. Ullah J, Luo M, Ashraf U, et al. (2022) Evaluation of the geothermal parameters to decipher the thermal structure of the upper crust of the Longmenshan fault zone derived from borehole data. *Geothermics* 98: 102268. https://doi.org/10.1016/j.geothermics.2021.102268
40. Ullah J, Li H, Soupios P, et al. (2024) Optimizing geothermal reservoir modeling: A unified bayesian PSO and BiGRU approach for precise history matching under uncertainty. *Geothermics* 119: 102958. https://doi.org/10.1016/j.geothermics.2024.102958
41. Eppelbaum LV, Zheludev V, Averbuch A (2014) Diffusion maps as a powerful tool for integrated geophysical field analysis to detecting hidden karst terranes. *Izv Acad Sci Azerb Rep Ser.: Earth Sciences* 2014: 36–46.

AIMS Press