Geosciences

Research article

Reservoir prescriptive management combining electric resistivity tomography and machine learning

Paolo Dell'Aversana*

Eni SpA, San Donato Milanese, Milan, Italy

* Correspondence: Email: paolo.dell'aversana@eni.com, dellavers@tiscali.it; Tel: 390252063217.

Abstract: In this paper, I introduce a comprehensive workflow aimed at optimizing oil production and CO_2 geological storage. I show that the same methodology can be applied to different categories of problems: a) real-time reservoir fluid mapping for predicting and delaying water breakthrough time as far as possible in oil production; b) real-time CO_2 mapping for maximizing the sweep efficiency and storage capacity of CO_2 in geological formations. Despite their intrinsic differences, these types of problems share common aspects, issues and possible solutions. In both cases, various geophysical techniques can be applied, including Electric Resistivity Tomography (briefly ERT) for accurate fluid mapping and monitoring. This method is highly effective and sensitive for detecting the type of fluid and for estimating saturation in the geological formations. The robustness and the accuracy of the ERT models increase if densely spaced electrodes layouts are permanently deployed into the production and injection wells. In the first part of the paper, I discuss how in both scenarios of oil production and CO_2 storage, we can apply time-lapse borehole ERT method for mapping fluids in the reservoir. Next, I discuss how to apply various techniques of time-series analysis for predicting the evolution of the fluids distribution over time. Finally, using Q-Learning, that is a specific Reinforcement Learning method, I discuss how we can optimize the decisional workflow using our models about past, real-time and predicted fluids displacement. The result is the definition of a "best policy" addressed to both problems of optimized oil production and safe CO_2 geological storage. In the second part of the paper, I show benefits and limitations of my approach with the support of synthetic tests.

Keywords: reservoir monitoring; water breakthrough; electric tomography; machine learning

1. Introduction

Water and gas movement to production oil zone represents a relevant issue in the oil industry. In fact, coning and cresting phenomena [1–3] can cause significant decrease in well productivity and recovery factors. In extreme cases, the unpredicted arrival of undesired fluids into the production well can cause the premature shut-in of the well itself. Furthermore, the production of large quantities of water, especially in the depletion stage of the reservoirs, commonly causes additional high costs for treatment and disposal [4]. Early water breakthrough can be a serious problem when the waterflooding approach is applied in the phase of secondary recovery. In this technique, water is injected into the reservoir formation to displace residual oil. The water from injection wells sweeps the displaced oil to adjacent production wells, causing significant fluids movements in the reservoir.

In order to mitigate all the problems associated with undesired fluids displacement, it is crucial to map fluids' distribution in the reservoir zone, possibly in real time. Furthermore, it is important to make probabilistic predictions about when and where phenomena like coning or cresting can occur at the production wells [5]. Such a key information can be used for tuning production in optimum rate to avoid or to delay water breakthrough time as far as possible.

A different category of problem associated with the displacement of fluids in the reservoir, is how to optimize the storage of CO_2 into deep geological formations. CO_2 capture and storage (CCS) represents nowadays one of the viable options for reducing the emissions of anthropogenic CO_2 into the atmosphere [6,7]. An open question associated with the CCS strategy is how to maximize the sweep efficiency and storage capacity of CO_2 in geological formations. This goal can be achieved by adopting a controlled injection scenario. As in the case of the water displacement problem during oil production, a crucial objective is mapping and predicting the movements of the CO_2 plume during and after injection.

In both cases, (mapping and predicting water and CO_2 displacement), several geophysical techniques can be applied, including seismic as well as non-seismic methods. All these methods are based on the principle that the different types of fluid (oil, water and CO_2) show different physical properties and, consequently, can be detected using geophysical methods sensitive to the variations of these properties. However, detecting the fluids, estimating their saturation in the reservoir and predicting future displacements represent partial aspects of the general problem of reservoir management optimization. An additional pragmatic aspect is how to use the entire available information in order to take the best actions for avoiding premature water breakthrough as well as for optimizing the CO_2 storage.

Despite obvious intrinsic differences, both types of problems share many common features. Therefore, they can be faced through a similar, comprehensive approach. The workflow proposed in this paper includes the following main steps.

1) Fluid mapping: defining a reliable model of fluid distribution and saturation in the reservoir, before, during and after production/injection; 2) Fluid prediction: using the available historical information acquired through a sequence of monitor surveys for predicting the space-time evolution of the fluid distribution; 3) Actions prescription: using the entire information about past, present and predicted fluid distribution for decision/action prescription.

The final goal of these steps is management optimization addressed to hydrocarbon production as well as to CO_2 storage operations.

Some authors have proposed interesting solutions for managing both oil production and CO_2 storage optimization problems [8–11].

Liao et al. [12] have discussed ensemble-based optimization techniques of the water-alternating-gas (WAG) injection process.

Shamshiri and Jafarpour [13] have developed a framework for equalizing and delaying (under constant total injected volume of water) the breakthrough time of the injected fluid at production wells. Furthermore, they have shown how to apply the same approach to optimize CO_2 sequestration. In that case, well producers are not present, so the authors have introduced the concept of pseudo production wells. These represent just virtual wells used to indicate the arrival of CO_2 (and its fractional production rate) at selected locations. The final scope is to equalize production rates and to delay CO_2 breakthrough. The authors have supported their approach with synthetic tests, showing its theoretical effectiveness in both cases of water breakthrough risk mitigation and CO_2 storage optimization. However, they did not discuss any method for effectively mapping over time the fluids displacement in the reservoir zone.

In this paper, I apply an approach similar to that of Shamshiri and Jafarpour, but I introduce additional ideas and methods for real-time fluid mapping. Furthermore, I introduce a Machine Learning approach for predicting future fluids displacements and, finally, for defining a prescriptive policy for optimal reservoir management, both in case of oil production as well as in case of CO_2 injection and storage.

The structure of this paper is related to the workflow shown in the conceptual scheme of Figure 1. First, I recall the basics of Electric Resistivity Tomography (ERT). This is aimed at defining the fluid model and its continuous update (in almost real time). Then, I discuss how to apply effective techniques based on Artificial Neural Networks and Recurrent Networks in particular, for probabilistic prediction of fluids movements. Finally, I explain how Reinforcement Learning algorithms can support the crucial phase of actions prescription, using the entire set of information about fluid distribution.

After the methodological section, I discuss several synthetic 2D and 3D tests in order to clarify the details of each step, and to highlight benefits and limitations of my approach.

2. Methodology

In this section, I discuss the three main methodological aspects of the workflow. These concern, respectively, the techniques to map the distribution of fluids over time, the methods to predict their displacement in future, and the approaches to optimize the management policy of the reservoir.
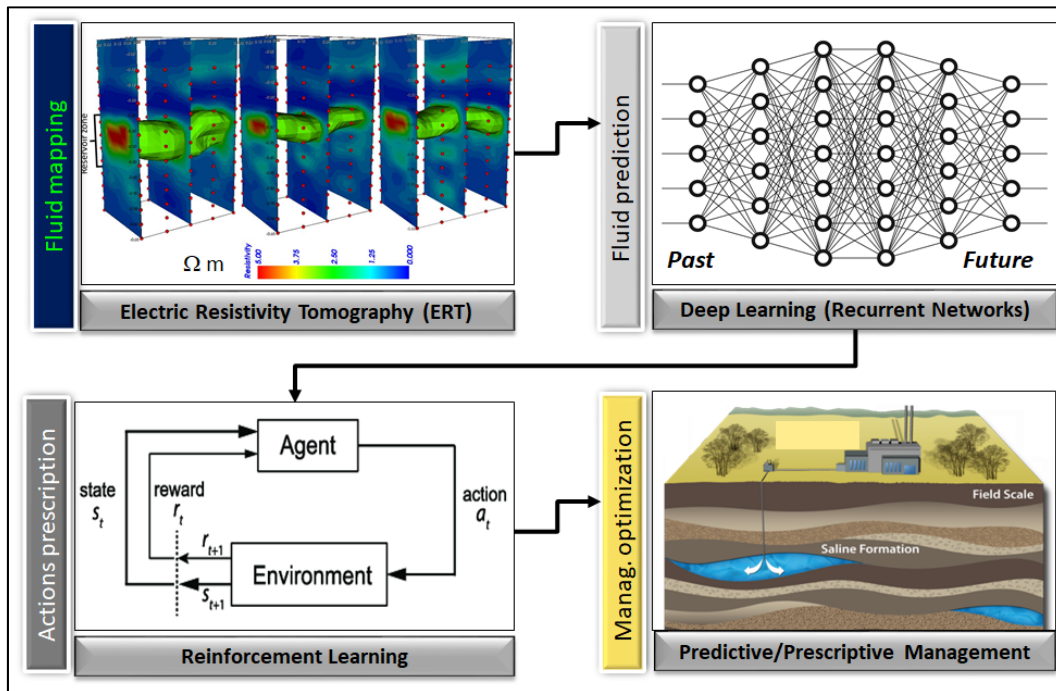
Figure 1. The main blocks of the workflow for predictive and prescriptive management, in case of oil production as well as in case of CO_2 injection and storage.

2.1. Fluid mapping using Electric Resistivity Tomography

Electric Resistivity Tomography (ERT) is a consolidated approach to retrieve, by appropriate inversion algorithms, a resistivity model from surface and/or borehole measurements of electric potentials taken along multichannel arrays of electrodes. Different rock formations and materials generally show different resistivity values. Therefore, the retrieved ERT models are commonly used in geosciences as well as for supporting engineering studies. Often, resistivity tomography is successfully applied in hydrogeological studies for monitoring fluids movements over time [14–16]. In a time-lapse ERT survey (sometimes indicated also as "DC time-lapse survey", where DC means "Direct Current"), electrodes are installed at fixed locations during the period of monitoring. First, a "base resistivity data set" is collected. The inversion of this initial data set produces a "base resistivity model" to be used as a "reference model". Then, one or more "monitor surveys" are repeated during the period of monitoring. The purpose of time-lapse ERT is to detect possible variations of resistivity distribution caused by fluid displacements from one survey to another. These variations can be retrieved using various inversion approaches. The simplest one is by subtracting the inverted monitor model from the inverted base model. A more sophisticated inversion strategy is known as "Difference inversion" where we invert the difference between the monitor and the base data sets. This approach allows minimizing the coherent inversion artifacts, and the output is a "difference image"; this represents the spatial distribution of resistivity differences, generally expressed in percent, between base and monitor models.

In ERT method, it is easy to acquire multiple data sets over time (monitor surveys), even on daily basis. Inverting the differences between consecutive data sets, allows obtaining a dynamic model of variations of electric resistivity caused by changes in fluid saturation and hydraulic permeability. These can be seasonal changes of the water table, or salty water invasion in coastal areas.

The same approach can be applied also to monitor fluids distribution in hydrocarbon reservoirs [17,18] for monitoring the process of CO_2 storage and for other purposes [19–22].

Furthermore, the ERT approach can be used also in conjunction with other methods based on electric and electromagnetic measurements. For instance, Dell'Aversana et al. [23,24] and Bottazzi et al. [25] proposed to combine electric dipoles and magnetic coils permanently installed on the well completion for detecting reservoir fluids movements in real time during oil production and water injection.

2.2. Predictive methods

In this second methodological sub-section, I summarize the methods that I use for predicting the ERT response from the data acquired through a sequence of time-lapse surveys. For that purpose, I apply several techniques of time series forecasting. A time series is a sequence of data (such as experimental measurements) over time. An important objective in time series analysis is to forecast its future values, at a given time t. Under certain assumptions, we can hope to make probabilistic predictions using the information given by the past observations.

Time-lapse ERT method can produce large volumes of data in the form of time series. In the part of this paper dedicated to synthetic examples, I will show how and in what measure electric resistivity changes when fluids move in the subsurface over time. These changes can show different trends depending on spatial distribution of fluids, rock types, permeability and many other factors. If we take multiple measurements of electrical potentials through many monitor surveys (eventually combined with other types of measurements, like seismic data, pressure, and temperature), we can create a "historical data set". This can be used for predicting the evolution of fluids in the subsurface by applying forecasting algorithms.

There are many approaches to the problem of time series forecasting, with variable level of complexity and effectiveness. For example, a simple class of methods includes Exponential Smoothing that is based on a weighted sum of past observations; the model uses an exponentially decreasing weight for past observations [26]. Single Exponential Smoothing, SES for short, is effective for univariate data without a trend or seasonality. Double Exponential Smoothing is an extension to Exponential Smoothing that explicitly adds support for trends in the univariate time series. Triple Exponential Smoothing is a further extension of the previous approaches that explicitly includes the possibility to take into account for seasonality in the univariate time series.

Approaches that are more complex are based on Multivariate Analysis (MVA). This is useful in problems where multiple measurements are made on each experimental unit and the relations among these measurements and their structures play an important role. In my workflow, I apply both univariate-based methods (including Exponential Smoothing) as well as MVA based approaches. In fact, ERT data sets are internally correlated, for reasons that I am going to explain.

Among the various predictive approaches, I use Neural Networks [27,28] that allow predicting the trend of time series by learning from the past historical trends measured in the data. In that frame, a specific effective approach for predicting the time evolution of time series is through Recurrent Neural Networks (RNNs) [29]. The basic idea in RNNs is to add a time delay unit and a feedback connection in order to use information from a previous state for inferring a model in the subsequent state. We can imagine the simplest RNN as consisting of just three layers: one input layer with d input neurons, one hidden layer with c hidden neurons and one output layer with q output neurons. At each

time point t, the network contains three weight matrices, W_{ih}, W_{hh} and W_{ho} that connect, respectively: input to hidden layer; recurrent hidden layer to hidden layer; hidden layer to output. These weight matrices represent the parameters of the RNNs. The hidden layer captures information about what happened in the previous time step. It can be considered as the "memory" of the network. In such a way, RNNs make use of sequential information. In traditional Neural Networks, all inputs are assumed independent from the outputs; instead, in RNNs we assume that there is some significant correlation between input and output. For instance, let us suppose that we desire to predict the next word in a sentence. For that goal, we cannot ignore the words that came before it. RNNs have been developed for purposes like this (and in fact, they are widely applied in computational linguistics).

We can think about RNNs as they had a sort of memory state that is used for predicting a model of the future states. For that reason, we can expect that RNNs are useful when we want to retrieve a dynamic model based on time-lapse geophysical measurements. This is the case in which we desire to predict the fluids movements within a reservoir (hydrocarbons, water or CO_2) using the data acquired in the past through multiple monitor ERT surveys. Indeed, my method applies RNNs to time-lapse ERT data for predicting how resistivity distribution changes over time because of fluids movements. In the section dedicated to synthetic tests, I will discuss illustrative examples of this predictive approach.

2.3. Prescriptive methods

2.3.1. Overview

In the previous sub-sections, I have discussed how we can take electric measurements for mapping fluids displacement in the reservoir. Furthermore, we can apply various techniques of time-series analysis for predicting the evolution of the fluids distribution over time. The next step consists in the prescriptive part of the workflow. My goal is to use all the real and predicted data for optimizing the decision flow and the future actions for various purposes. As discussed in the introduction, monitoring and predicting water movements in the reservoir represents an important aspect in production management optimization. Furthermore, in the CCS frame, monitoring and predicting the movements of the CO_2 plume during and after injection, is crucial for maximizing the sweep efficiency and storage capacity of CO_2 in geologic formations. Both types of prescriptive aspects (addressed to oil production and CO_2 storage optimization) can be supported by a Reinforcement Learning approach.

2.3.2. Reinforcement learning

The basic idea of Reinforcement Learning is to emulate the natural learning behavior in machines and automata, generally called "agents", operating in their "environment" [30–32]. In simple words, in every particular state of its environment, the agent performs some action and receives some type of positive or negative feedback about the appropriateness of its behavior. Let us suppose that the agent and the environment are in a certain state s_t. The agent selects an action a_t to take at time t in its environment. That action will affect some aspects of the environment where the agent itself is operating. The effects of the agent's action will be returned by the modified environment in the form of a reward (or a punishment) and a new state s_{t+1}. The final objective of the

agent is to maximize the total reward accumulated during its lifetime (cumulative reward). In other words, in the Reinforcement Learning framework, the agent has to learn how to optimize its actions for a long-term horizon. The cumulative reward that is achievable from a given state *s*, is called the Value function, *V(s)*. Commonly, a "discount factor", *γ*, is applied to the various rewards in order to assign higher weight to rewards received near than rewards received further in the future. Based on these rewards and/or punishments received from its environment, the agent tries to learn progressively a "policy" for maximizing the cumulative long-term reward. The policy, often denoted by the symbol π, is expressed by a function that takes the current environment state *s* belonging to the set *S* of all possible states to return an action *a* belonging to the set *A* of all possible actions.

$$\pi(s): S \rightarrow A \tag{1}$$

Finding the best policy can be a very difficult problem to solve when the agent acts in complex environments with a wide range of possible discrete or continuous states and actions. Moreover, complexity increases further if the rewards is sparse and delayed. There are many different approaches falling in the class of Reinforcement Learning methods. Among the various algorithms, I selected the Q-Learning method, for the possibility to combine its simplicity with the effectiveness of deep neural networks (Deep Q-Learning). This method takes its name from the so-called *Q-function*. This represents the "quality" of a certain action given a state. It is a function of a state-action pair and returns a real value.

$$Q(s, a) = S \times A \rightarrow R \tag{2}$$

The optimal *Q-function*, *Q*(s, a)*, represents the expected total reward received by an agent starting in the state *s*, picks the action *a*, and then will behave optimally afterwards. The importance of finding the optimal *Q-function* is that we can use it to extract the optimal policy *π** by choosing the action *a* that gives maximum reward *Q*(s, a)* for state *s*. We have:

$$\pi^*(s) = \arg\max_a Q^*(s, a) \qquad \forall s \in S. \tag{3}$$

The Bellman equation tells us that the maximum future reward is the reward *r* the agent received for entering the current state *s* plus the maximum future reward for the next state *s'*.

$$Q(s, a) = r + \gamma \max_{a'} Q(s', a'), \tag{4}$$

where *γ* is the discount factor that reduces the contribution of future rewards with respect the immediate reward. We can use a recursive formula for finding *Q(s,a)*:

$$Q_{t+1}(s_t, a_t) = Q_t(s_t, a_t) + \alpha \left[r_{t+1} + \gamma \max_a Q_t(S_{t+1}, a) - Q_t(s_t, a_t) \right] \tag{5}$$

The factor *α*, called "*learning rate*", gives us the possibility to control how much we desire to take into account the difference between the previous *Q* value and the new one. With this approach, we start by using random values (or any guess value) for the *Q-function*. Then, when the agent proceeds exploring its environment, the initial *Q* values progressively converge to the optimal ones, based on the positive and/or negative feedbacks the agent receives after its actions.

2.3.3. Q-Learning and management optimization: basic principles

The Q-Learning approach can fit well our purpose of finding an efficient action policy for maximizing oil production, as well as for optimizing CO_2 storage. In order to apply this method, in both cases I need to clarify what I mean by an *agent*, what is the *environment* where it acts, what is a *state* and what is an *action*, and how we estimate the *reward*. Finally, I need to clarify quantitatively what I intend for *best policy*. It means that I need to formalize some type of objective function that should converge towards an optimal policy addressed to oil reservoir management and/or CO_2 storage.

Let us start from the case of oil production. As I said, an important aspect of oil recovery maximization is to mitigate the risk of undesired phenomena of water invasion into the production wells, such as coning and cresting. Our goal is to regulate properly the production and/or injection parameters (pressure, production rate, number and location of injection wells, and so forth) in order to have high production of oil and, at the same time, low risk of water breakthrough. This goal can be achieved if we can continuously monitor the distance of the waterfront from the production well(s). In fact, it is of primary importance to have the possibility to control the complex dynamics of fluids movements (using the ERT method, in my approach) in order to prevent the invasion of water into the well(s). Consequently, we can learn how to regulate our production parameters in order to reach an optimal balance between production rates, water injection rates and water breakthrough risk. This is a typical optimization problem, and it can be addressed with the Q-Learning method. In fact, in this context, every production or injection well (or every production compartment of any individual well) can represent an agent. The state of each agent is defined by its production parameters, together with the correspondent water breakthrough risk. The action of each agent consists of regulating the production/injection variables. A positive reward consists of high oil production and, at the same time, low risk of water breakthrough; this risk is inversely proportional to the distance of the waterfront from the production well (estimable through ERT).

Having these basic concepts in mind, in order to apply the Q-Learning method, we start by constructing an action-state *M x M* matrix (Table 1). This matrix allows representing the transition rules and the correspondent rewards in our environment. It consists of *M* possible states (rows) that can change into another state (action or transition). Hence, in each state, the agent has *M* potential actions (columns) which it can choose. Depending on the particular state in which the agent is and the subsequent action which the agent then takes, there is a unique distribution for transition to a next state. This means that the transition probabilities to any next state depend on (only) the previous state as well as the action then taken (i.e. we are assuming the conditions of a *Markovian Decision Process*, briefly MDP).

After each of the agent's *M* actions, the agent will get a reward R_{ij}. (i.e. the reward obtained when moving to the sate *j* from the state *i*). As I said earlier, we need to define a *Reward Function* that represents, for each state in our process, the way to estimate the expected value of the reward for an action taken in that state. I set this function to be proportional to the oil production rate and to its increment when moving from a state to another. In our case, the state transition consists of regulating the production parameters (for instance, increasing oil production rate in a certain well, or in a certain sector of a well). Furthermore, the rewards must be decreased by introducing a penalization term in case the oil rate increment causes a water displacement, increasing the risk of water breakthrough in the production well. Such a risk is related to the waterfront distance from the well:

high distance (based on ERT models) corresponds to low water breakthrough risk, and vice versa. I set the Reward function using the following simple formula:

$$R(s, s', a) = P_{cur}(s) + \Delta P(s, s', a) - \Delta r(s, s', a) \tag{6}$$

where $P_{cur}(s)$ is the oil production in the current state s, $\Delta P(s, s', a)$ is the variation of oil production when moving to the next state s' (through the action a), and $\Delta r(s, s', a)$ is the correspondent variation of water breakthrough risk. That risk term is inversely proportional to the average distance variation, Δd_{water}, (observed or predicted) of the water table from the production well (or, better, from the sector of the well equipped with the electrodes). That distance variation is determined through the ERT survey performed at the state s, together with the predictive method discussed above and applied to the state s'. We have:

$$\Delta r(s, s', a) \sim - 1/\Delta d_{water}(s, s', a) \tag{7}$$

For example, if the agent w_i (a certain production well) is in a state s_i (i.e. it is characterized by certain production parameters and a certain water breakthrough risk), and it takes an action a_k (for instance it increases production rate without affecting significantly the waterfront distance), then we assign a positive reward to that action-state transition, without any risk penalization term. The reason is that, after moving into the new state, our well has improved oil production without increasing significantly the risk of water breakthrough. Instead, if the agent w_i is in a state s_i and it takes another action a_l that increases production rate, but with the undesired effect of increasing significantly the risk of water breakthrough, then we add the penalization term (7) to the total reward. After assigning a total reward to each action-state transition, we obtain the *Reward Matrix* or, briefly, the *R-Matrix*. Table 1 shows an example of this matrix obtained from a synthetic test (discussed in the following) where I simulated some schematic transitions from states characterized by different production and risk values. I simulated a scenario where the production increment causes the displacement of water towards the production well. Consequently, the reward of each transition is penalized every time the waterfront advances towards the well itself.

For instance, the transition from the "Low Production—Low Risk" state (LP-LR in the matrix of Table 1) to the "High Production—Medium Risk" state (HP-MR) has a Reward of 0.6 (highlighted in red). The normalized production and risk values used for the rewards computation are shown above and on the left of the table, for each row and each column.

Table 1. Example of R-Matrix (or R-Table) with assigned rewards for each state transition. Each reward value R_{ij} is assigned as explained in the text.

PROD.	RISK			PROD.	0,1	0,1	0,1	0,5	0,5	0,5	1	1	1			
				RISK	0,1	0,5	1	0,1	0,5	1	0,1	0,5	1			
PROD.	RISK				LP-LR	LP-MR	LP-HR	MP-LR	MP-MR	MP-HR	HP-LR	HP-MR	HP-HR		LEGEND	
0,1	0,1		LP-LR		0,1	-0,3	-0,8	0,5	0,1	-0,4	1	0,6	0,1		LP:	LOW PRODUCTION
0,1	0,5		LP-MR		0,5	0,1	-0,4	0,9	0,5	0	1,4	1	0,5		MP:	MEDIUM PRODUCTION
0,1	1		LP-HR		1	0,6	0,1	1,4	1	0,5	1,9	1,5	1		HP:	HIGH PRODUCITON
0,5	0,1		MP-LR		0,1	-0,3	-0,8	0,5	0,1	-0,4	1	0,6	0,1			
0,5	0,5		MP-MR		0,5	0,1	-0,4	0,9	0,5	0	1,4	1	0,5		LR:	LOW RISK
0,5	1		MP-HR		1	0,6	0,1	1,4	1	0,5	1,9	1,5	1		MR:	MEDIUM RISK
1	0,1		HP-LR		1	0,6	0,1	1,4	1	0,5	1,9	1,5	1		HR:	HIGH RISK
1	0,5		HP-MR		1	0,6	0,1	1,4	1	0,5	1,9	1,5	1			
1	1		HP-HR		1	0,6	0,1	1,4	1	0,5	1,9	1,5	1			

Creating an R-Matrix is just the starting point of the Q-Learning optimization process. In fact, the R-Matrix is a schematic representation of the rewards that we receive from the environment where our agent acts. I remind that that our objective is to optimize the entire process of action-state transitions in the agent's environment; in our case, the final goal is to reach, as soon as possible, a stable condition of maximum oil production with minimum (or, at least, with controlled) risk of water breakthrough. Such an optimum condition is obtained through the iterative process of updating the Q values through the Bellman equation (4) and the associated recursive formula (5). As said earlier, initially we set random (or null) values for the Q-function. Then, the agent proceeds in its exploration of the environment and acts in it, getting positive and/or negative feedbacks from it. In this process, the agent always tries to optimize its rewards by taking the best possible action that it learns. In such a way, the initial Q values are progressively updated through the recursive Bellman formula in order to converge to the optimal ones. The result will be an updated Q-Table (or Q-Matrix). Such matrix allows defining the optimal policy in our environment, in terms of i,j steps with the maximum long-term reward. For illustrative purposes, I prepared a tutorial Python code showing in detail all the Q-Learning steps mentioned above [33].

A crucial question is how to implement these theoretical concepts into a pragmatic approach. In order to allow practical applications, the optimisation Q-Learning process must be translated into an "automatic and iterative tuning mechanism" of the production/injection variables. Such a tuning procedure is possible if it can take profit from the continuous information available from our agents (the production/injection wells) about their dynamic environment. This type of information is the one that we can get in terms of oil production rate and water table distance using ERT data. In fact, assuming to have installed in each well a continuous and permanent monitoring/mapping ERT system, we can obtain a real-time update about the movements of fluids in a wide range of distances around and between the wells. Furthermore, we have seen in section 2.2 (*Predictive methods*) that we can also forecast the fluid movements. It means that we can associate a value of actual and future water-breakthrough risk to each state/action during the process of oil production. In other words, we can construct our R-Matrix and update it when necessary during the production history of our wells. Finally, the agents (our wells) learn through the Q-Learning algorithm how to set an optimal policy for maximizing the production and minimizing the water breakthrough risk at each moment of the production history. Consequently, all the production/injection parameters are regulated based on such an optimal policy. The scheme of Figure 2 summarizes the entire approach described above.
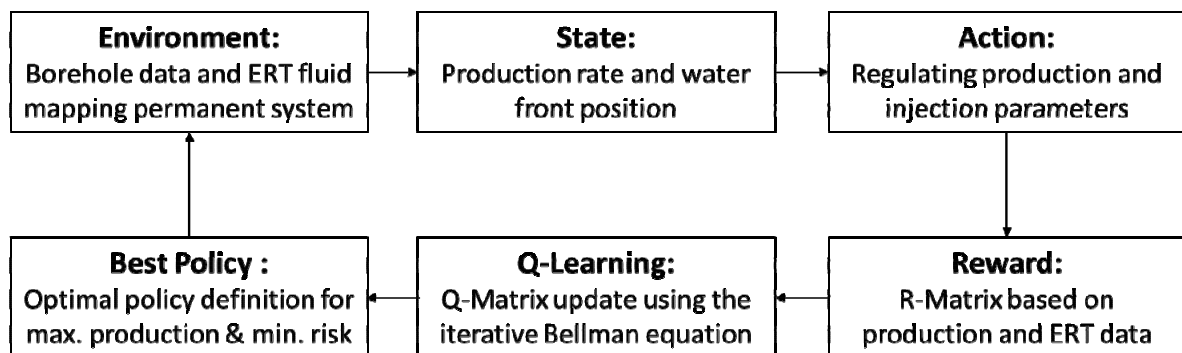

Figure 2. Schematic block diagram for an optimal policy definition supported by ERT permanent fluid mapping system and Q-Learning iterative method.

2.3.4.   Q-Learning for CO_2 optimal storage

An analogous workflow, with some necessary modifications, can be applied for defining a best policy addressed to optimal management of CO_2 geological storage. Following Shamshiri and Jafarpour [13], in this case the goal is to maximize the sweep efficiency and storage capacity of CO_2 in geologic formations by a controlled injection. That goal can be reached by adjusting injection rate allocations by accounting for aquifer heterogeneity. In this case, "the objective is to maximize the sweep efficiency by mobilizing the injected CO_2 uniformly in all directions and exposing a larger volume of fresh brine to the CO_2 plume, leading to an enhanced solubility and residual trapping" .

In order to formalize this idea, the authors cited above introduce the concept of *pseudo production wells*. These are *virtual producers* with negligible total fluid production rate. They are introduced in the simulation just with the purpose to indicate the arrival of CO_2 (and its fractional production rate). That information (the arrival time of CO_2) is used to equalize production rates at all pseudo-producers and to delay CO_2 breakthrough. Finally, Shamshiri and Jafarpour formalize their optimization problem by minimizing a term of an objective function that penalizes the difference between CO_2 production rates at different pseudo production wells. Furthermore, a second term is introduced in the objective function to delay the CO_2 breakthrough at the same pseudo-production wells.

I use the same approach in my problem formalization. However, differently from Shamshiri and Jafarpour, I do not need to introduce the concept of pseudo-production wells. In fact, as explained earlier, we can monitor the movements of the injected CO_2 through real measurements. Indeed, we can use permanent ERT layout installed in the injection wells (as well as in additional monitoring wells) in order to control (and to predict) the CO_2 movements in all directions.

An additional fundamental difference with respect to the approach of Shamshiri and Jafarpour is in my technique for optimizing the injection control parameters. Differently from Shamshiri and Jafarpour, my idea is to adjust these parameters through the Q-Learning approach. In other words, I search for a best policy through the iterative application of the Bellman equation, as explained in the previous sub-section. This time, the goal is that CO_2 movement in the reservoir is as uniform as possible, despite the existing heterogeneities and preferential flow paths. Figure 3 shows the block diagram addressed to optimize CO_2 geological storage. It is conceptually similar to the workflow addressed to oil production optimization (Figure 2), but with the necessary adaptations to the problem of CO_2 storage (highlighted in red).
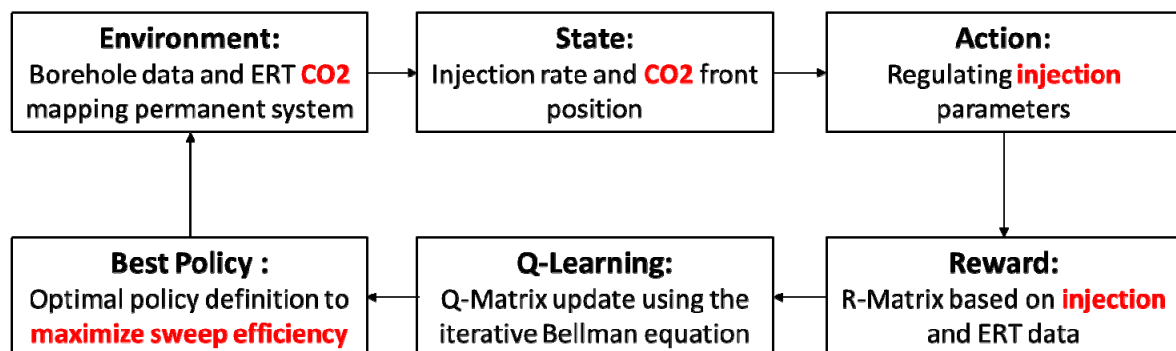

Figure 3. Schematic block diagram for an optimal policy definition for CO_2 geological storage, supported by ERT permanent CO_2 mapping system and Q-Learning iterative method.

3. ERT synthetic tests

In this section, I discuss 2D and 3D ERT time-lapse synthetic tests. The discussion is split into two sub-sections. First, I show how we can use ERT method for monitoring the CO_2 plume during injection. Next, I show how we can map oil displacement by water during production. Finally, I show synthetic tests concerning the prediction of future fluid displacements, using the predictive methods introduced in section 2.2.

3.1. 2D borehole ERT for CO_2 monitoring

The Carbon Capture and Storage (CCS) workflow includes CO_2 capture from large stationary sources, transport to selected locations and, finally, injection and storage into porous geological formations, with the final goal to isolate CO_2 from the atmosphere [34]. Unfortunately, the possibility of leakage from the storage reservoirs cannot be excluded, and this can cause serious consequences. In fact, leaked fluids can migrate along faults or other discontinuities, accumulating in formations overlying the storage formations. Consequently, they can interfere with another subsurface activity (such as natural gas storage, deep-well injection of wastes and so forth), or can cause groundwater contamination. Finally, CO_2 can reach the surface and escape into the atmosphere. In order to mitigate the leakage risk, it is crucial to have one or more effective methods for monitoring the diffusion of the CO_2 plume in the subsoil over time. As I have said earlier, ERT tomography offers significant chances to map the movements of CO_2 in the subsurface, if we assume that it significantly affects the bulk resistivity after injection in the storage formation. That assumption is reasonable and supported by modelling as well as proven by real experiments of CO_2 storage and monitoring. An effective way to apply ERT for this purpose is to use borehole arrays of densely spaced electrodes in multiple injection/monitoring wells. With this approach, we can expect to perform permanent geo-electric monitoring of the CO_2 movements with relatively high spatial resolution in the rock volume between the wells.

For illustrative purposes, I start with a simple 2D time-lapse inversion test of synthetic cross-well ERT data. My goal is to show benefits and limitations (in terms of spatial accuracy and model uncertainties) that we can expect when we acquire and invert ERT data using borehole layouts [35]. I simulate a simplified scenario of CO_2 injection, where there is just one injection well and one monitoring well at a distance of 250 m. I assume that the wells cross an aquifer in a geological formation. In this test, the inversion algorithm inverts the electric potentials and tries to retrieve directly the resistivity distribution. Fluid saturation can be estimated through a separate step, using for instance the second Archie's law. The uncertainty propagation from electric potential data to resistivity model space is estimated by assigning, in the input data file, an error-vector describing the standard deviation of reciprocal measurements for each electric quadrupole [36]. The uncertainties propagation from the resistivity model to the saturation model depends on the uncertainties on porosity and on the parameters of the petro-physical relationship(s) used for the resistivity-saturation conversion, as discussed in Dell'Aversana et al. [37]. In these synthetic tests, I assume that porosity and the parameters of the Archie's law are perfectly known.

For each inversion, I set the average background resistivity at 1 Ωm. The depth scale in the reference figures is a relative depth, ranging from 0 m to 400 m. I assume that the CO_2 injection in

this aquifer formation creates a significant resistivity contrast (two orders of magnitude) between the background geology and the rock formation interested by the CO_2 plume.

For this specific test, I have used the open source "R2" software package based on Python libraries [38–41]. Furthermore, I have used the "pyres" Python wrapper package for R2. This creates a powerful and flexible programmatic interface for modeling and inversion of ERT datasets using the NumPy, SciPy, and Matplotlib Python packages. The R2/pyres software allows forward/inverse solution for 3D or 2D current flow in a quadrilateral or triangular mesh. The mesh is made up of a set of elements. Electrodes are specified at node points corresponding with the corners of the elements. The boundary conditions along all four boundaries of the mesh are Neumann conditions (zero flux). R2 software uses a mesh of finite elements to compute the model, in both forward and inverse mode. The mesh is based on a quadrilateral (structured or unstructured) or triangular (unstructured) mesh of finite elements. Node points on its perimeter define each element, and the voltage is computed at each node point given a dipole current source applied at a pair of nodes. As we will see in the following test results (shown in Figure 4), the elements of the mesh gradually increase in size laterally and vertically outside the region of investigation. The inverse solution is based on a regularized objective function combined with weighted least squares (an Occam's type solution) as defined in Binley and Kemna [39].

In the test here described, I put electrodes in both wells at a constant distance of 5 m, for a thickness of 120 m, assuming that this is the total thickness of the aquifer. Furthermore, I added some electrodes below the aquifer zone (from 120 m to 250 m, with a larger spacing than in the electrodes above). I simulated the acquisition of a base survey at time T0 (first row in the figure) and of two monitor surveys at times T1 and T2 (rows 2 and 3). I inverted the synthetic response for each simulated scenario, using a uniform starting model of 1 Ωm. The first column of panels shows the "true models" at each time. We can see a high resistivity plume growing over time caused by the expanding CO_2 plume and progressively moving from the injection well towards the monitoring well. The white dotted line has the scope of better showing the boundaries of the plume, although in the reality we can expect a smoothed transition zone in terms of resistivity values. The second column of panels show the inverted model of the synthetic data associated to each individual survey from T0 to T2. Finally, the third column of panels show the same inverted models together with the inversion mesh. A zoomed image is shown in Figure 5.
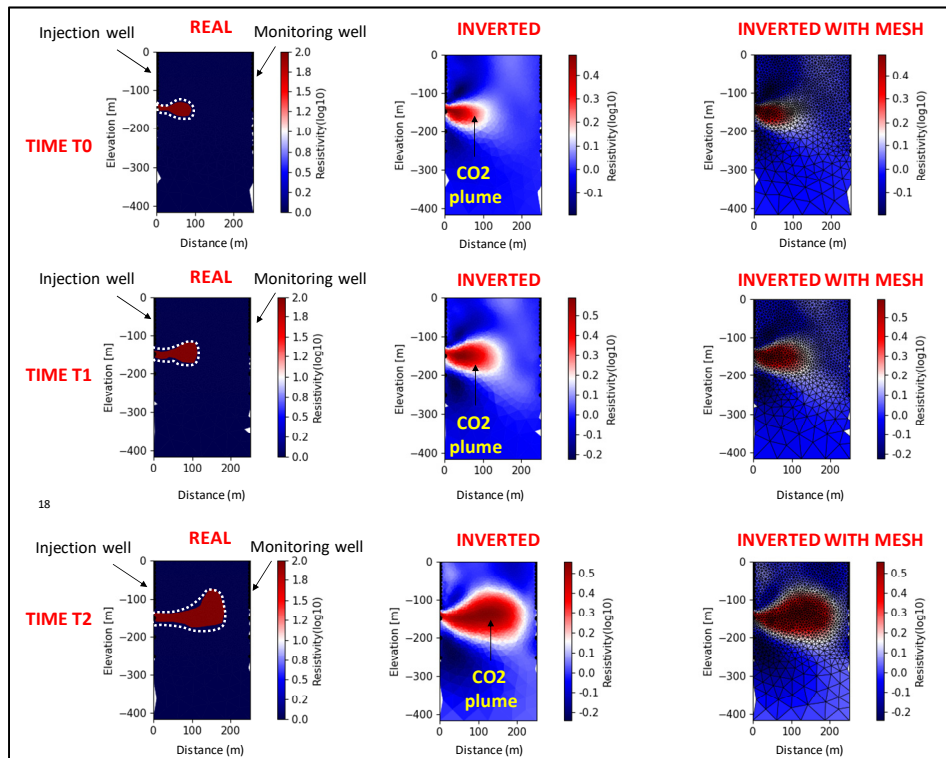
Figure 4. Example of 2D sequential inversions of synthetic cross-well ERT data.
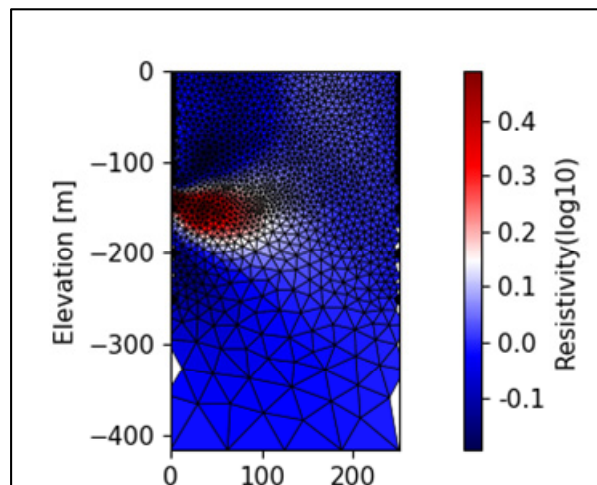

Figure 5. Zoomed image of the inverted model at time step T0 showing also the unstructured finite-element mesh.

The first general comment that we can do is that the "real models" are reproduced quite well through inversion, despite I added 5% of noise in the input data. The evolution of the CO_2 plume (the resistivity anomaly) over time appears clearly in terms of growing size and direction of migration. Of course, there are unavoidable approximations and uncertainties on the inverted models. The "true shape" is not reproduced with accuracy everywhere and it appears generally overestimated in size by the effects of the smoothing operator included in the inversion algorithm. Consequently, the resistivity of the CO_2 plume is underestimated. However, these limitations can be mitigated by a better use of the mesh/inversion parameters that can be fixed optimally in the input files of the R2 software. For instance, we can improve the mesh using different parameters for the element number

and size. In my test, I tried to find an optimal compromise between accuracy and computation times. The right panels of Figure 4 clearly show that I have used a detailed mesh in the area where electrodes are densely spaced, whereas the mesh elements become larger in the area below. Of course, the mesh and the smoothing parameters can be better tuned based on the user's experience. Finally, the inversion result can be improved depending on the available computation resources and on the electrodes layout.

The key message of this test is that the ERT method can provide, in principle, an effective approach for mapping the space-time evolution of the CO_2 plume during and after injection. The problem of resistivity underestimation can be mitigated by the fact that, in time-lapse applications, we are interested mainly to the relative variations of resistivity rather than its absolute values. This point will be better illustrated in the next 3D simulation test. The main limitation of this approach is that the wells where we put the electrodes cannot be too far. Based on modelling and real experiments, we can say that the spacing between wells should not be larger than 300–400 m, also depending on the length of the borehole layout. In fact, longer layouts allow longer investigation distances. However, the investigated volume can be expanded if the borehole ERT layout is complemented by a 3D surface layout.

3.2. 3D borehole ERT for oil reservoir monitoring

The following 3D test simulates a scenario of a significant resistivity change in an oil reservoir and below it, caused by the water table approaching four horizontal production wells. This test is described with additional details in a previous paper. Here, I recall just the key aspects of this simulation.

In the production scenario of this test, I simulated the response of a DC cross-hole acquisition survey, where a multi-channel electrode system is deployed in four parallel horizontal wells. These are located at a mutual constant distance of 250 m. The wells are at a constant depth of 2340 m below the surface, in correspondence of the top of the reservoir. In each well, I simulated an acquisition layout consisting of 15 electrodes, 25 m spaced.

Figure 6 shows the grid used for the modelling and the four horizontal wells where the electrodes (black dots) are installed for acquisition. The grid consists of irregular rectangular cells. Their size changes in function of the electrodes spacing. The maximum expected spatial resolution of the inverted model parameter (resistivity, in this case) corresponds to the minimum half-spacing between the electrodes.
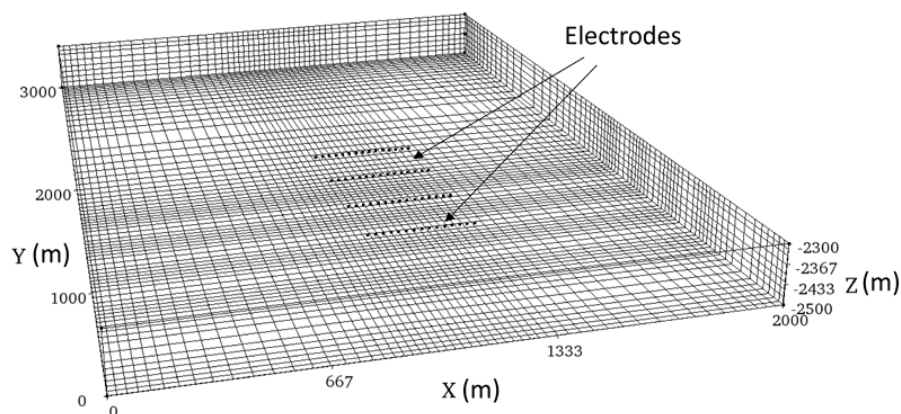

Figure 6. Modelling grid and electrodes' layout in the four horizontal wells.

I used the "PUNQ-S3 reservoir model" developed by the Imperial College of London [42]. It represents a small-size industrial reservoir scenario consisting of 19 × 28 × 5 grid blocks. A South and East fault system bounds the modelled hydrocarbon field. Moreover, there is an aquifer bounding the reservoir to the North and West. In order to generate the porosity/permeability fields, a geostatistical model based on Gaussian Random Fields has been used. A synthetic history has been simulated using the reservoir simulator ECLIPSE. I selected a few significant scenarios of variations in terms of fluids saturation, in order to simulate a base-model and a few monitor-models. Then, I applied the Archie's second law [43] for transforming the porosity and saturation distributions into the corresponding resistivity model. Finally, I simulated the DC acquisition using the derived resistivity model. This consists of five levels (having a thickness of 10 m each) with variable resistivity. The left panel of Figure 7 shows an example view of the base resistivity model of the reservoir; the orange/red colors indicate resistivity of the order of 100 Ωm, while the blue color indicates about 3–4 Ωm. The resistivity of the background (transparent in the figure) is set to 1 Ωm.
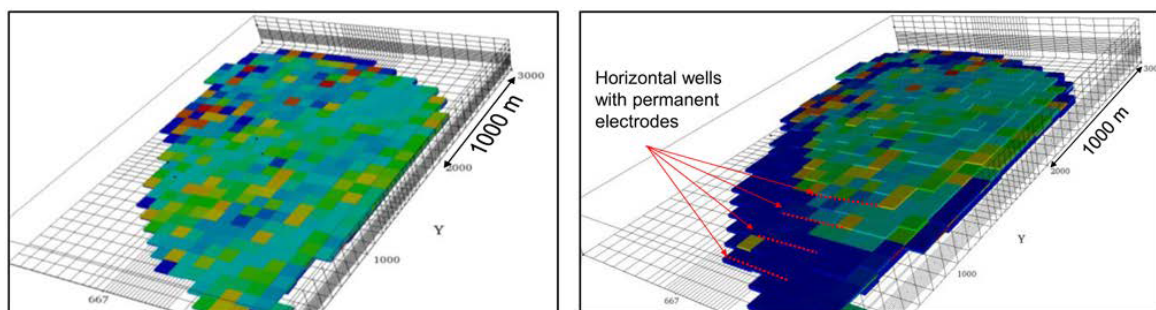

Figure 7. Maps of the base reservoir model (left) and monitor reservoir model (right), when the waterfront is approaching the wells.

The reservoir is discretized by square cells of 100 m × 100 m and with a thickness of 10 m. The right panel of the same figure shows the monitor-reservoir model. The acquisition layout with electrodes installed in the four wells is highlighted in red. The monitor model shows a significant variation in conductivity (or, equivalently, in resistivity). In this simulation, I assumed that the movement of the waterfront approaches the wells from the South-West side (bottom left, in Figure 7) and from below. This waterfront is responsible of the decrease of resistivity observed below the wells. I simulated the acquisition in both base and monitor models, before and after the movement of water, using a "Mixed-dipole gradient" array, recording 2145 electric potentials in total. In order to simulate noisy data, I contaminated the synthetic response with 5% of Gaussian noise. Finally, I inverted the data using a Robust inversion algorithm that minimizes the absolute difference between the measured and calculated apparent resistivity [44].

I subtracted the simulated response in the two scenarios (base and monitor models). The result represents the Difference data vector, i.e. the input for the time-lapse Difference inversion. The inverted model represents the Difference conductivity model. This is the 3D spatial distribution of the changes of conductivity from the base to the monitor model. As an example of the inversion results, Figure 8 shows three displays of the anomaly retrieved through Difference inversion at three different depths. It is expressed in terms of percentage variation of electric conductivity, where the green color is the reference 0% of variation.
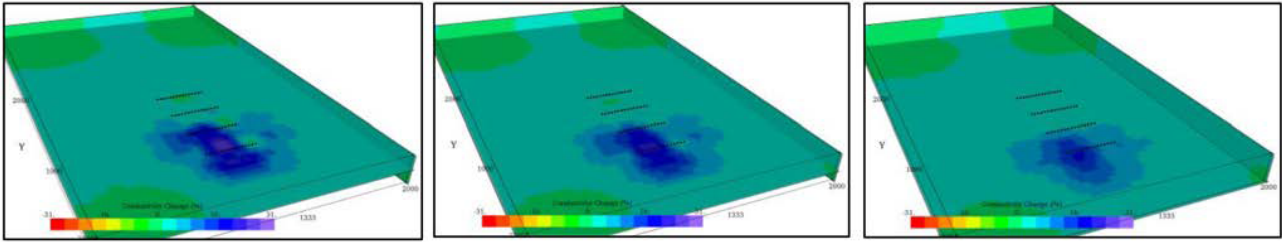
Figure 8. Variation of the conductivity distribution retrieved by Difference inversion at 60, 90 and 160 m below the wells (from left to right, respectively). Color scale: Green = 0%; Blue = 15–20%; Violet = 30%. Yellow, Orange and Red indicate negative variations in the color scale.

This test confirms that we can *detect, in* principle, significant variations of electric conductivity associated to waterfront movements, using a cross-well electrodes layout permanently installed in production wells and performing time-lapse DC tomography. As already remarked earlier, one of the main benefits of DC tomography is the simplicity and the rapidity by which we can acquire repeated surveys and invert time-lapse data sets in relatively short time (for instance, on daily basis). This allows performing a quasi-real time reservoir monitoring. This is an effective approach for continuous control of how the waterfront responds to oil production. Figure 8 shows just an example of the geo-electrical response that we can expect from time-lapse DC tomography in case of significant movements of the waterfront below the production wells. Finally, the time-space correlation of water movements in response to oil production is a fundamental information for optimizing decisions and actions during production.

Of course, before moving from modelling results to real applications we need to solve practical, operational and engineering issues. In particular, it is crucial to define appropriate constructive solutions for the well completion. We discussed those important issues and the possible solutions in dedicated papers, including significant results of laboratory experiments.

3.3. A synthetic test about fluid prediction

Figure 9 shows a schematic 2D scenario where the water table (resistivity = 1 Ωm) is moving upward (dashed lines), following an irregular dynamic trend. We can imagine that this is a typical situation where oil over-production or a waterflooding technique is causing an undesired "cresting" phenomenon. In this simulation, I deployed an ERT permanent layout (electrode spacing: 10 m) into a horizontal well (relative elevation = 0 m). My goal is to monitor the water movements in order to avoid the possible invasion in the producing well. Doing many monitor surveys like this, I created a dynamic database of electric potential measurements over time. Finally, I used these measurements for training my RNN predictive algorithm.

Figure 10 shows the apparent resistivity pseudo section obtained through the reference ERT survey at time T0, when the water table is marked by the black dashed curve in the previous figure.

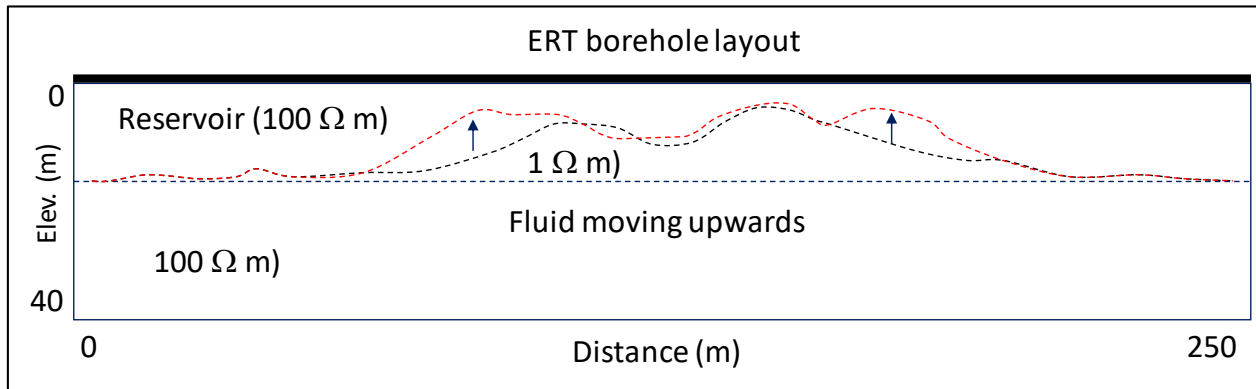**Figure 9.** Simulation scenario. Black dashed curve: waterfront at time step T0. Red dashed line: waterfront at time step T1. Vertical axis represents a relative depth with respect to the horizontal well (thick black line).
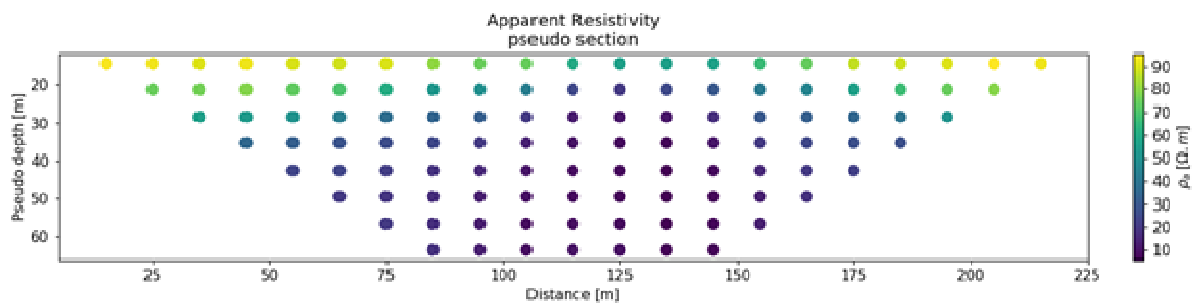

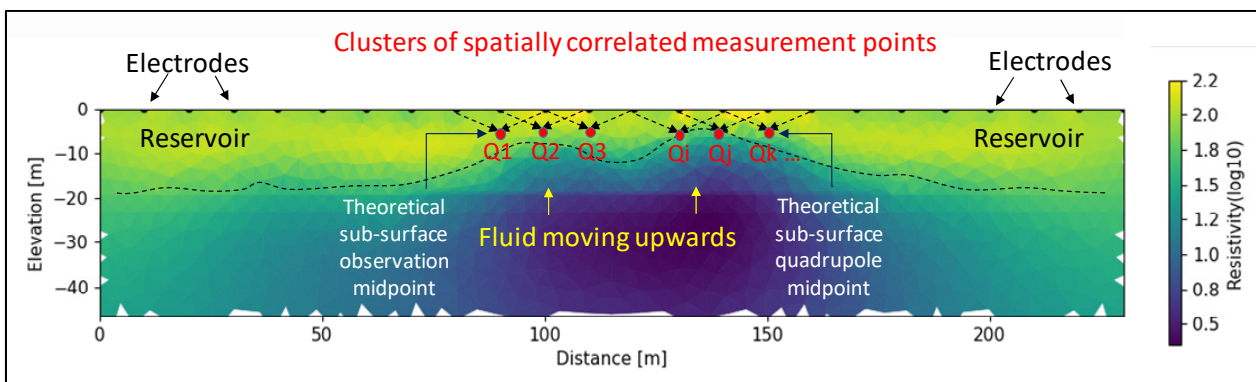**Figure10.** Apparent resistivity pseudo-section at reference time step T0.

Figure 11. Example of ERT model (obtained through inversion of the pseudo-section of Figure 10), in a scenario of detection of fluids below the measurement points.

Figure 11 shows the ERT section obtained by inversion of the apparent resistivity pseudo-section of Figure 10. Looking at this resistivity model, it is clear that the fluid movements in a certain area will affect the response of a cluster of observation points. Consequently, each prediction should be based on the history of a cluster of correlated measurements collected over time in the same "zone". For example, looking at Figure 11, we expect that the measurements at Q1 and Q3 be correlated with measurements at Q2 (where Q1, Q2 and Q3 indicate three examples of theoretical image points illuminated by the ERT response). Such a spatial correlation between clusters of measurements in the data set justifies my multivariate approach.

Using a large database of apparent resistivity pseudo-sections (like that of Figure 10) obtained through a sequence of ERT monitor surveys, we can try to predict the water table position at the next time step. For that purpose, I applied a well-known predictive approach: the Multivariate LSTM Forecast Model [45]. LSTM means "Long-Short Term Memory". It is a type of Recurrent Neural Network allowing the network to retain long-term dependencies at a given time from many time steps before. My objective is to fit the LSTM predictive algorithm on the ERT time-lapse data.

Let us explain the key steps of the workflow. First, I split the historical dataset of multiple monitor surveys (apparent resistivity or, equivalently, electric potentials pseudo-section) into train and test sets. Next, I run the training phase on these past data. It is possible to monitor the training results by plotting the loss function vs. the number of epochs (Figure 12).
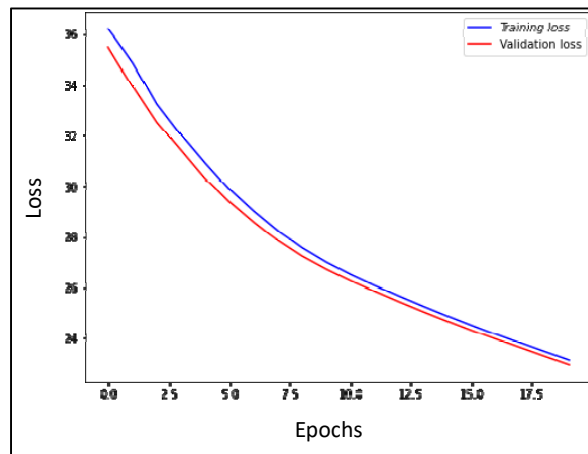
Figure 12. Trend of training and validation loss functions vs. number of iterations.
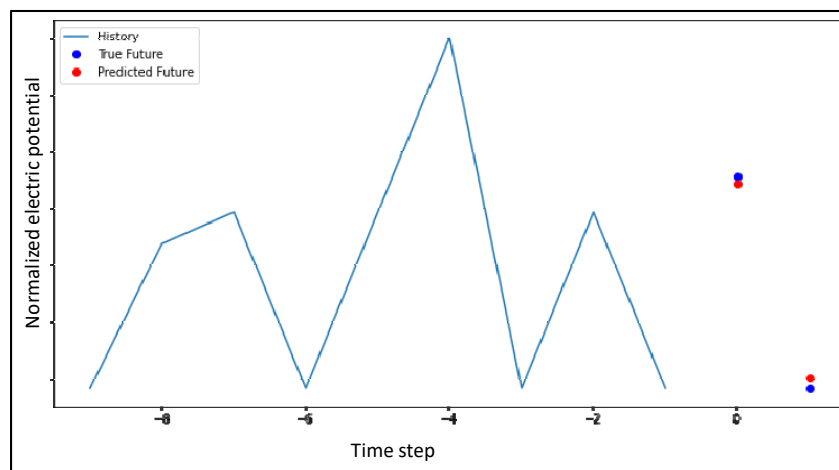
Figure 13. Example of multiple step predictions. Time step 0 corresponds to the first future monitor survey, negative time steps correspond to the past surveys.

After the LSTM model is fit, I use it to forecast the future data. This means predicting the future electric potentials or, equivalently, the apparent resistivity pseudo-section, in a changed scenario of the water table (such as the one represented by the red dashed line in Figure 9).

Figure 13 shows the comparison between "true" future and predicted future data at two time steps in the future at just one single measurement point. This is, for instance, one measurement point of a pseudo-section like the one of Figure 10. In this example, I know in advance the "true" future,

because this is a synthetic test. This allows me checking the reliability of my LSTM Forecast Model. We see that the forecasts are very accurate. Such accuracy is possible because the prediction is based on the past measurements using a multivariate approach. In other words, the LSTM Forecast Model, if trained properly on past correlated data, can provide robust predictions.

Combining the predictions of electric potentials at all measurements points, I create a table of electric potentials (Table 2). It includes the data for all the measurements points along the entire acquisition layout observed at the past time steps (i.e. past monitor surveys) and predicted for future time steps (i.e., future monitor surveys).

Finally, the predicted potentials are used for creating a predicted apparent resistivity pseudo-section (Figure 14). Inverting it, I obtain a predicted resistivity model (Figure 15).

To mention an additional illustrative example, I applied the same predictive approach to the 2D data set discussed in the section 3.1 (2D borehole ERT for CO_2 monitoring). Figure 16 shows the inverted model from the predicted data just one time-step ahead. The inverted resistivity model represents the predicted CO_2 plume at the "future" time step T3, where the synthetic ERT data at time T3 have been predicted through the LSTM Forecast Model.

Table 2. Electric potentials at the various measurements points (here indicated as "Quadrupole" mid points), acquired through past monitor surveys (data columns relative to negative time steps) and used for predictions (last column).

| Acquis. Point.	El. Pot. (mv)	El. Pot. (mv)	El. Pot. (mv)	El. Pot. (mv)	El. Pot. (mv)	El. Pot. (mv)
QUADRUPOLE	STEP −9	STEP −8	STEP −7	STEP −6	...	PREDICTIONS
Q1	1,4900	37,5606	43,9962	1,4698	...	1,5203
Q2	1,4799	36,6813	44,2806	1,4889	...	1,5117
Q3	1,5104	36,4429	44,0380	1,4663	...	1,4322
Q4	1,4643	36,3216	44,2720	1,4774	...	1,3501
Q5	1,4845	35,7880	44,1606	1,4583	...	1,0806
Q6	1,4675	35,4775	44,1187	1,4397	...	0,3414
Q7	1,4682	34,9042	44,0458	1,3957	...	-0,3313
Q8	1,4780	34,1099	44,0022	1,2836	...	-0,8040
Q9	1,4649	33,1987	44,1670	1,1961	...	-0,3267
Q10	1,4848	33,2913	44,1957	1,1803	...	-0,0450
...	...	...	...	...	...	...
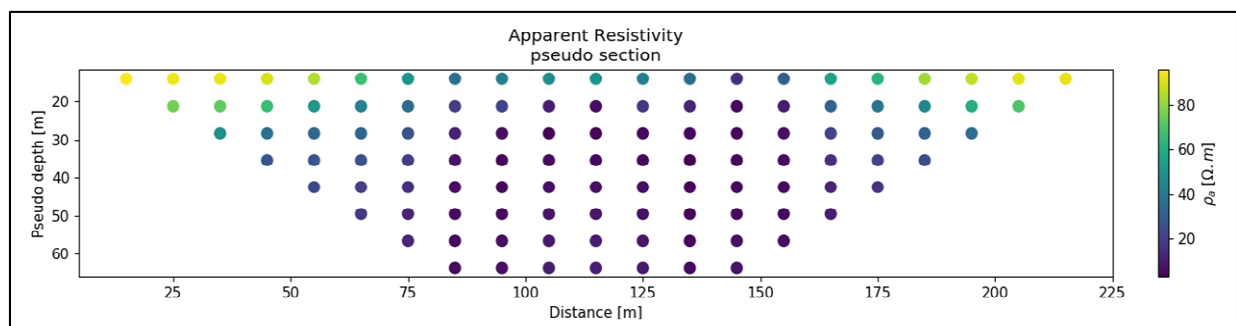
Figure 14. Predicted apparent resistivity section.
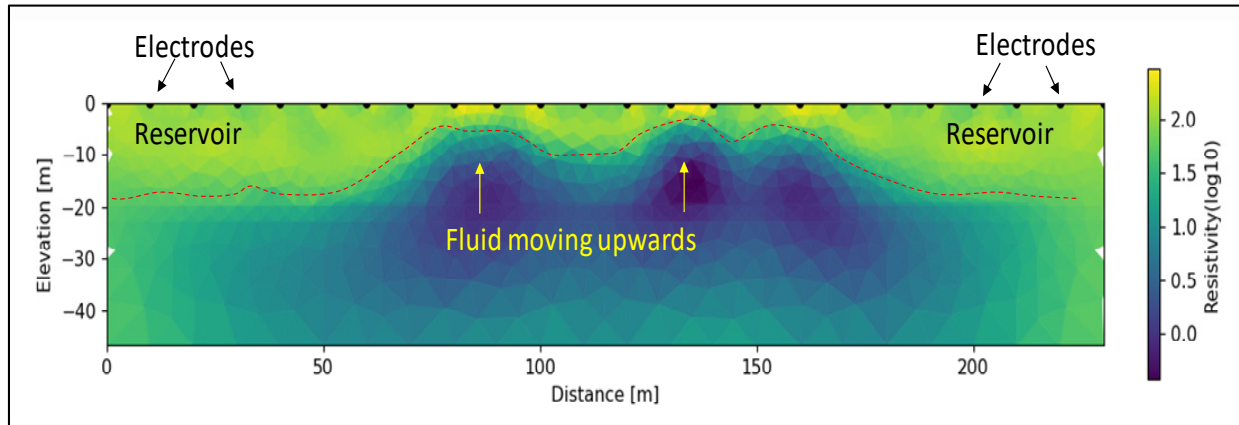
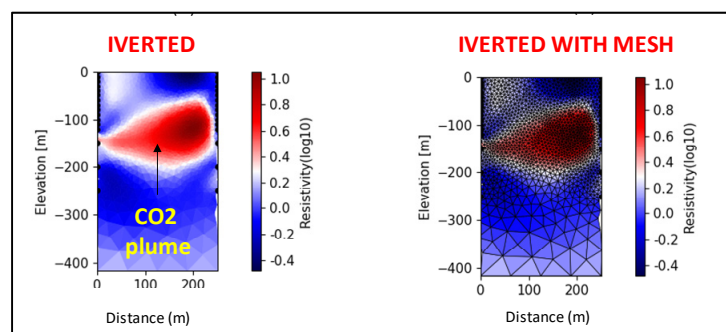**Figure 15.** Predicted ERT model at time step T1.

Figure16. CO_2 plume prediction at time T3 (compare with Figure 4).

4. Conclusions

Despite the obvious differences in the case of optimizing oil production or CO_2 storage, there are common issues and possible similar solutions. Many crucial problems are associated to the fact that fluids generally show high mobility in the subsoil. Hence, possible solutions depend on our ability to map in real time, and possibly to predict, the movements of these fluids in the reservoir. In this paper, I have introduced a complete workflow to optimize the key steps in both processes of oil production and safe CO_2 storage and monitoring. I have discussed the benefits (and the limitations too) offered by electric (and electromagnetic) methods in the crucial phase of mapping fluids over time, where fluids can be brine, hydrocarbons and CO_2, depending on the context.

The main conclusions can be summarized as following: 1) Electric Resistivity Tomography (ERT) is a valid tool for mapping fluids' displacement over time, during oil production as well as CO_2 injection. Synthetic tests show that this method is effective for mapping reservoir fluids. Resolution increases if densely spaced electrodes are deployed in the production/injection wells. These wells must be spaced at distances never larger than few hundred meters. 2) Future displacement of fluids can be predicted using Recurrent Neural Networks applied to historical data recorded by Electric Resistivity Tomography layouts through multiple monitor surveys. 3) Q-Learning represents a valid approach for setting an optimal prescriptive policy for managing oil production as well as CO_2 injection. Such an optimal policy is based on the time-lapse Resistivity Electrical historical data combined with the Recurrent Neural Network predictions.

Conflict of interest

The author declares no conflicts of interest in this paper.

References

1. Chaperon I (1986) Theoretical Study of Coning Toward Horizontal and Vertical Wells in Anisotropic Formations: Subcritical and Critical Rates. *SPE Annu Tech Conf Exhib*, 5–8.

2. Chierici GL, Ciucci GM, Pizzi G (1964) A Systematic Study of Gas and Water Coning By Potentiometric Models. *J Pet Technol* 16: 923–929.

3. Wheatley MJ (1985) An Approximate Theory of Oil/Water Coning. SPE Annual Technical Conference and Exhibition, Las Vegas, Nevada, USA, 22–26.

4. Al-Sikaiti SH, Regtien J (2008) Challenging Conventional Wisdom, Waterflooding Experience on Heavy Oil Fields in Southern Oman. *World Heavy Oil Congr*: 10–12.

5. 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.

6. Jiang X (2011) A review of physical modelling and numerical simulation of long-term geological storage of CO_2. *Appl Energy* 88: 3557–3566.

7. Jung JY, Huh C, Kang SG, et al. (2013) CO_2 transport strategy and its cost estimation for the offshore CCS in Korea. *Appl Energy* 111: 1054–1060.

8. 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.

9. González-Nicolás A, Cihan A, Petrusak R, et al. (2019) Pressure management via brine extraction in geological CO_2 storage: Adaptive optimization strategies under poorly characterized reservoir conditions. *Int J Greenhouse Gas Control* 83: 176–185.

10. Pongtepupathum W, Williams J, Krevor S, et al. (2017) Optimising Brine Production for Pressure Management During CO_2 sequestration in the Bunter Sandstone of the UK Southern North Sea. *Soc Pet Eng*.

11. 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.

12. Liao C, Liao X, Mu L,et al. (2017) Improving water-alternating-CO_2 flooding of heterogeneous, low permeability oil reservoirs using ensemble optimisation algorithm. *Int J Global Warming* 12: 242–260.

13. Shamshiri H, Jafarpour B (2010) Optimization of Geologic CO_2 Storage in Heterogeneous Aquifers Through Improved Sweep Efficiency. SPE International Conference on CO_2 Capture, Storage, and Utilization held in New Orleans, Louisiana, 10–12.

14. 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.

15. Goldman M, Kafri U (2006) Hydrogeophysical applications in coastal aquifers. In: Vereecken H, Author, *Applied Hydrogeophysics*, Eds., Springer: Dordrecht, The Netherlands, 233–254.

16. Kuras O, Pritchard J, Meldrum P, et al. (2009) Monitoring hydraulic processes with automated time-lapse electrical resistivity tomography (ALERT): *Compt Rendus Geosci* 341: 868–885.

17. Dell'Aversana P, Rizzo E, Servodio R (2017) 4D borehole electric tomography for hydrocarbon reservoir monitoring. EAGE Conference and Exhibition, 2017: 1–5.

18. McNeice GW, Colombo D (2018) 3D inversion of surface to borehole CSEM for waterflood monitoring. *SEG Int Expo Ann Meet*, 878–880.

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.

20. Bergmann P, Schmidt-Hattenberger C, Labitzke1 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.

21. 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.

22. 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.

23. Dell'Aversana P, Servodio R, Bottazzi F, et al. (2019_a) Asset Value Maximization through a Novel Well Completion System for 3d Time Lapse Electromagnetic Tomography Supported by Machine Learning. *Abu Dhabi Int Pet Exhib Conf.*

24. Dell'Aversana P, Servodio R, Bottazzi F, et al, (2019_b). Asset Value Maximization through a Novel Well Completion System for 3d Time Lapse Electromagnetic Tomography Supported by Machine Learning. *Soc Pet Eng J.*

25. 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. *Int Pet Technol Conf.*

26. Brown RG (1956) Exponential Smoothing for Predicting Demand. Cambridge, Massachusetts: Arthur D. Little Inc 15.

27. Saad EW, Prokhorov DV, Wunsch DC (1998) Comparative study of stock trend prediction using time delay, recurrent and probabilistic neural networks. *IEEE T Neural Network* 9: 1456–1470.

28. Tealab A (2018) Time series forecasting using artificial neural networks methodologies: A systematic review. *Future Comput Inform J* 3: 334–340.

29. Sherstinsky A (2020) Fundamentals of Recurrent Neural Network (RNN) and Long Short-Term Memory (LSTM) Network. *Phys D* 404: 132306.

30. Kaelbling LP, Littman ML, Moore AW (1996) Reinforcement Learning: A Survey. *J Artif Intell Res* 4: 237–285.

31. Raschka S, Mirjalili V (2017) *Python Machine Learning: Machine Learning and Deep Learning with Python, scikit-learn, and TensorFlow*, 2nd Ed., PACKT Books.

32. Russell S, Norvig P (2016) *Artificial Intelligence: A Modern approach, Global Edition,* Pearson Education, Inc., publishing as Prentice Hall.

33. Machine learning workflows. Multidisciplinary applications using Python, 2021. Available from: https://www.researchgate.net/publication/348741974_Q_Learning_generic.

34. Benson SM, Surles T (2006) Carbon dioxide capture and storage: an overview with emphasis on capture and storage in deep geological formations. *Proc IEEE* 94: 1795–1805.

35. Christensen NB, Sherlock D, Dodds K (2006) Monitoring CO_2 injection with cross-hole electrical resistivity tomography. *Explor Geophys* 37: 44–49.

36. LaBrecque DJ, Miletto M, Daily W, et al. (1996) The effects of noise on Occam's inversion of resistivity tomography data. *Geophysics* 61: 538–548.

37. Dell'Aversana P, Carbonara S, Vitale S, et al (2011) Quantitative estimation of oil saturation from marine CSEM data: A case history. *First Break* 29.

38. Befus KM (2017) Pyres: A Python Wrapper for Electrical Resistivity Modeling with R2. *J Geophys Eng* 15.

39. Binley A, A Kemna (2005) *Electrical Methods*, In: Hydrogeophysics by Rubin and Hubbard Eds., Springer, 129–156.

40. Binley A (2015) *Tools and Techniques: DC Electrical Methods*, In: Treatise on Geophysics, 2nd Ed., Schubert: Elsevier, 233–259.

41. Binley A (2016) R2 version 3.1 Manual. Lancaster, UK. Available from: http://es.lancs.ac.uk/people/amb/Freeware/freeware.htm.

42. PUNQ-S3 reservoir model, Imperial College of London. Available from: https://www.imperial.ac.uk/earth-science/research/research-groups/perm/standard-models/.

43. Archie GE (1950) Introduction to petrophysics of reservoir rocks. *AAPG Bulletin* 34: 943–961.

44. Claerbout JF, Muir F (1973) Robust modeling with erratic data. *Geophysics* 18: 826–844.

45. Sagheer A, Kotb M (2019) Time series forecasting of petroleum production using deep LSTM recurrent networks. *Neurocomputing* 323: 203–213.