Optimizing NMR Spectroscopy Pulse Sequencing with Reinforcement Learning for Soil Atomic Abundance

Rohan Shenoy
University of California, Berkeley
rohan_sh@mit.edu

Hans Gaensbauer
Massachusetts Institute of Technology
hgaens@mit.edu

Elsa A. Olivetti
Massachusetts Institute of Technology
elsao@mit.edu

Evan A. Coleman
MIT Climate Project
ecol@mit.edu

Abstract

Determining the amount of sequestered carbon in soils and monitoring soil health in farmlands is an important climate change problem. Motivated by the lack of scalable and inexpensive techniques for in-situ soil health monitoring, we focus on low-voltage nuclear magnetic resonance (NMR) spectroscopy as a promising new approach and develop a reinforcement learning technique to modulate NMR pulses for rapid atomic abundance assessment of soils. Our preliminary results derived using Monte Carlo sampling and parallelized OpenAI Gym training show the promise of our RL-based approach.

1 Introduction

Maintaining and monitoring soil health is a crucial challenge in the battle against climate change, particularly in the agricultural sector [1]. Soil degradation diminishes soil fertility and can lead to desertification, compelling farmers to resort to methods such as deforestation to meet growing food demand [2]. Furthermore, soil degradation and erosion lead to the release of sequestrated soil carbon stocks, which make up 2-3 times the carbon content of the atmospheric carbon pool [3]. Hence, even small changes in soil carbon stocks can have disastrous effects on the atmospheric CO_2 concentration. Important challenges in the monitoring of soil health include determining soil carbon concentration, assessing soil fertility, and identifying heavy metal contaminants. These challenges are often considered separately, but they share a common goal: accurately determining the atomic abundance (i.e., atomic concentration) of key soil elements.

Current soil testing techniques are able to identify these concentrations through combusting or reacting soil extracts and observing the reactions. However, these techniques have fundamental limitations that hinder their effectiveness for broader climate objectives. Firstly, due to the nature of these tests, they must occur in controlled laboratory settings, leading to costs (in USD) of \$10-\$50/sample for basic fertility and contaminant testing and up to \$3000/sample for more comprehensive testing [4]. It has been estimated that to develop a reliable soil profile, soil testing should occur every few meters [5], meaning that this type of soil testing becomes prohibitively expensive even on the scale of a single farm. These scalability concerns hinder large-scale data collection efforts, which are important to identify and detect changes in atomic abundance. Being able to accurately detect these changes enables us to quantify the effects of different sustainable farming practices [6, 7], as well as identify opportunities for their implementation. Furthermore, soil carbon quantification methods incur measurement errors when processing samples, [8], which undermine the reliability of soil carbon crediting programs, an integral part of the larger carbon market model for decarbonization.

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2022.

These limitations of current common soil testing techniques motivate the design of a reliable, scalable method for rapid atomic abundance measurement in soils. In this work, we use low-voltage NMR spectroscopy as a promising new approach for in-situ soil monitoring and present a novel reinforcement learning framework for modulating low-voltage NMR to perform rapid in-situ atomic abundance assessment. First, we present a fast, robust simulator for generating large quantities of NMR spectroscopy data able to simulate the spin dynamics of different soil samples in parallel. We then demonstrate how to utilize this simulator to train a reinforcement learning agent capable of modulating an NMR pulse sequence for the purpose of determining the carbon concentrations of various coffee samples, which we use as a soil simulant.

2 Background

We start with some necessary background for NMR spectroscopy. Nuclear magnetic resonance is a physical phenomenon wherein the nuclear spins of atomic nuclei contained in a magnetic field are pulsed with electromagnetic radiation, causing the spins to precess before returning back in line with the magnetic field. This phenomenon is analogous to a spinning top being knocked over, precessing before eventually returning to rotation around the vertical axis. Nuclear magnetic resonance spectroscopy is a measurement technique based on this phenomenon, in which a substance is placed in a strong magnetic field and is exposed to a fixed sequence of radio-frequency pulses causing the atomic nuclei to emit energy at different frequencies according to their nuclear spin characteristics. The electromagnetic emission is measured at an axis perpendicular to the direction of the magnetic field, giving a signal known as the free induction decay (FID) [Figure 2a]. The Fourier transform of this signal is computed to obtain the NMR spectra [Figure 2b], and regression-based techniques are used to compare this spectra to previously sampled spectra for identification and analysis. NMR spectroscopy is most commonly used in medicine for a procedure known as Magnetic Resonance Imaging (MRI) to produce high quality images of human anatomy, but also has many prevalent applications in food science, environmental monitoring, and various engineering fields [9, 10, 11, 12]. A current area of emerging NMR research is centered around developing and applying low-voltage NMR systems. These systems are of interest for soil monitoring because of their relative low-cost, as well as their ability to determine atomic abundance *in situ*. This is a crucial extension to traditional measurement techniques for atomic abundance, since it addresses many of the scalability concerns mentioned above.

3 Methods

In this work, we aim to perform atomic abundance assessment using Nuclear Magnetic Resonance spectroscopy as a candidate procedure. As noted earlier, this assessment is useful for measuring soil carbon, fertility assessment, and heavy metal contaminant identification, which are all crucial aspects of maintaining soil health. Traditional NMR spectroscopy relies on applying a fixed pattern of radio-frequency pulses (which can be seen in Figure 1c) to the sample of interest and then performing regression-based techniques between the measured spectra and laboratory measurements. Regression-based techniques have been shown to fail when applied to samples outside of the training distribution [13], which often occurs in practice when considering the vast array of chemical compositions of different soil samples [14]. Owing to a lack of data, we propose to use reinforcement learning to learn the best policy for application of radio-frequency pulses in determining atomic abundance. Our reinforcement learning framework has 3 components: a Monte Carlo sampling procedure to represent a soil sample as a set of atomic spins, a parallelizable training environment where the agent can modulate the radio-frequency pulses, and a candidate reward model that will enable us to determine atomic abundance. We detail each component below.

3.1 Monte Carlo sampling

To train the agent in simulations, we must first create a representation for a soil sample that the agent can interact with. We represent a soil sample as a collection of atomic spins, where each spin is characterized by it's gyro-magnetic ratio γ, T_1 relaxation time, and T_2 relaxation time (see Appendix A). To calculate a representative set of spins for a given soil sample, a Monte Carlo sampling procedure is used. We begin with the NMR spectra of the soil sample, and propose 3 distributions from which the gyro-magnetic ratio, T_1 relaxation time, and T_2 relaxation time can be

sampled to determine a random spin. As the initial distributions are up to our determination, the agent can learn the optimal policy for any specified spin distribution. After a spin is sampled, the NMR spectra is computed for the current set of spins, and the difference between the target spectra and the sampled spectra is computed using a weighted average of the absolute mean squared error and phase mean squared error. If this difference is below a set threshold, the spin is accepted into the set, and otherwise it is rejected and removed from the set. This threshold is set depending on the physical parameters of the experiment and decreases as more spins are accepted, which makes sense in practice because as spins are accepted into the sample, our target spectra should continue to approach the desired spectra. This procedure is repeated until a desired number of spins are accepted, or a sampling limit is reached.

3.2 Parallel training environment

Given a set of spins, we built an NMR simulator (see Appendix A) that is a modified OpenAI Gymnasium environment, capable of simulating spin dynamics given the pulse sequence that is applied. The state space of our simulator is the measured magnetization in the X and Y directions, M_x, M_y respectively, as well as the current maximum transverse magnetization $M_{\max_t}$, which is calculated by $\max_t \sqrt{M_{x_t}^2 + M_{y_t}^2}$. It is important to note that the state does not include any information about the underlying spins, as would be the case in a real NMR deployment. The action space is a continuous variable with relatively high magnitude with respect to the larger magnetic field, as would be in a low-voltage NMR set up. To train an agent capable of handling the wide distribution of soils, we need the agent to interact with many different spinsets and simulator characteristics such as the temperature T_K, magnetic field strength B_0, and measurement noise. This motivated an important step in the training pipeline of our agent, parallelizing the NMR simulator so that the agent learns from interactions with a variety of spinsets and simulator configurations simultaneously. We utilized the `SubprovVecEnv` class built into OpenAI Gymnasium to vectorize our Markov Decision Process (MDP) – the vectorized MDP formalized as a tuple $(S^n, A^n, P_{sa}, \gamma, R)$. Our state and action spaces have been transformed to be n independent states and actions, however, our probability distribution matrix and reward function remain the same over all environments and are thus unchanged in the vectorized MDP. This step is crucial to making this approach feasible, as it allows us to parallelize training (i.e., process many soil samples simultaneously). This parallelism drastically improves the speed at which the reinforcement learning model converges to the optimal policy, as well as the stability of the training due to the averaging of noise across different environments where the same sample was processed.

3.3 Reward model

An important step in the training pipeline is designing the reward function in a manner such that the pulse sequence learned by the agent gives information that is valuable for an end-user attempting to calculate atomic abundance. In this section, we present the design of the reward model used in the training of our agent, and delve into how this model helps us achieve our final goal. The reward model for each episode is of the form $R = \sum_{t=1}^{T} \gamma^t \times M_{\max_t}$, where T is the total time in the episode, γ is a discount factor, and $M_{\max_t}$ is the maximum observed transverse magnetization at time t. The goal of this reward model is to have the learned pulse sequence obtain the maximum possible transverse magnetization by knocking as many spins into the the transverse direction as possible. Since our radio-frequency pulse is tuned to only interact with the atomic nuclei of the element of interest, the maximum magnetization we can achieve is a function of the abundance. Thus, if the maximum transverse magnetization achieved over the episode is monotonic with respect to the atomic concentration, then atomic abundance can be learned directly from $M_{\max_T}$.

4 Results and discussion

We trained a standard PPO agent for 5,000,000 timesteps across a parallelized environment containing 50 spinsets derived from a fixed distribution (see Appendix B). An example pulse sequence that is learned by the agent can be seen in Figure 1c. The agent tends to utilize the extremes of the action space to obtain the maximum transverse magnetization. We then utilize our Monte Carlo sampling procedure on real NMR spectra obtained from a serial dilution test of caffeine (see Appendix B), which we use as a soil simulant due to the controllable nature of the underlying atomic concentrations.

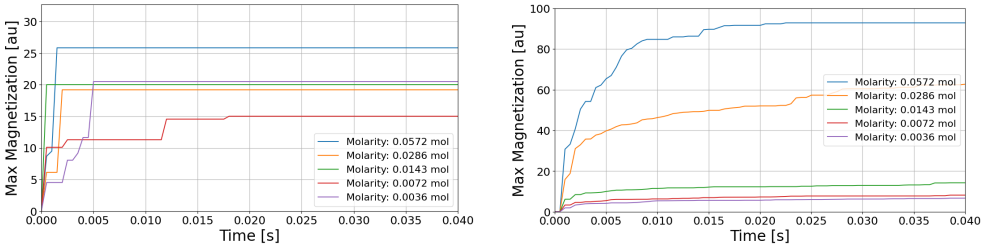

(a) Standard Pulse Max Magnetization

(b) Model Max Magnetization

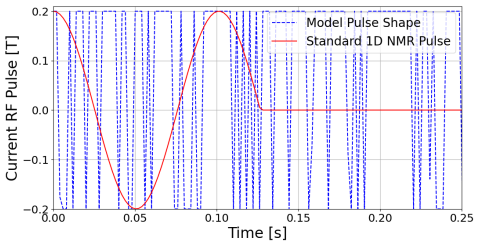

(c) NMR Pulse Sequences

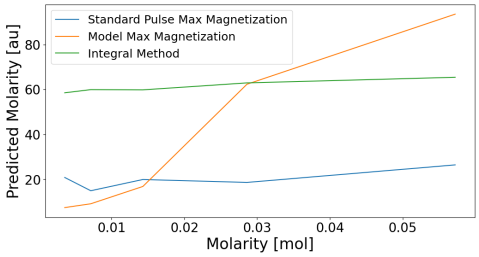

(d) Predictions of Atomic Abundance

Figure 1: Results from RL Validation Experiments

We then compared the maximum observed transverse magnetization achieved by the reinforcement learning agent to that achieved by the standard 1D NMR pulse sequence across the spinsets generated by the Monte Carlo sampling procedure. We can see in Figure 1a that the maximum observed magnetization achieved by the reinforcement learning agent is monotonically increasing with respect to the caffeine concentration, however, the same cannot be said for that achieved by the standard pulse sequence (1b). Another common method for atomic abundance calculation in standard 1D NMR is calculating the integral of the observed transverse magnetization, however, in Figure 1c it can be seen that this method also fails to achieve monotonicity in the same set up that our agent achieves monotonicity in.

The success of our agent in achieving monotonicity in cases where standard atomic abundance techniques fail motivates further exploration into machine-learning for dynamic pulse sequencing. Furthermore, the relatively low magnitude of the external magnetic field compared to the applied pulsing field highlights the feasibility of implementing this approach in a low-voltage NMR setup. While we validate our approach on a simple data set, where coffee is used as a soil simulant, further work can be done to extend this approach to a larger distribution of real soils. Furthermore, for the model itself, questions remain about the optimal choices of training hyper-parameters, alternative reinforcement learning models, and monitoring the spoiling of the underlying spins so that statistically independent samples can be collected in practice.

Acknowledgments and Disclosure of Funding

This work was completed with the support of the MIT Climate & Sustainability Consortium.

References

[1] V. Girija Veni, Ch. Srinivasarao, K. Sammi Reddy, K.L. Sharma, and Ashish Rai. Chapter 26 - soil health and climate change. In *Climate Change and Soil Interactions*, pages 751–767. 2020. ISBN 978-0-12-818032-7. doi: https://doi.org/10.1016/B978-0-12-818032-7.00026-6. URL `https://www.sciencedirect.com/science/article/pii/B9780128180327000266`.

[2] Afroz Alam. Soil degradation: A challenge to sustainable agriculture. *International Journal of Scientific Research in Agricultural Sciences*, 1:50–55, 08 2014. doi: 10.12983/ijsras-2014-p0050-0055.

[3] Mengmeng Ai, Yingying Sun, Bo Yan, and Yang Wei. A summary of the impact of land degradation on soil carbon sequestration. In *IOP Conference Series: Materials Science and Engineering*, volume 394, page 052028, 2018. doi: 10.1088/1757-899X/394/5/052028. URL https://doi.org/10.1088/1757-899X/394/5/052028.

[4] Katy Willis. How much does it cost to test soil and what does it include?, November 2022. URL https://www.angi.com/articles/how-much-does-testing-soil-cost.htm. Accessed: 2024-08-20.

[5] Somya Sharma, Swati Sharma, Rafael Padilha, Emre Kiciman, and Ranveer Chandra. Domain adaptation for sustainable soil management using causal and contrastive constraint minimization. *CoRR*, abs/2401.07175, 2024. URL https://arxiv.org/abs/2401.07175.

[6] Rattan Lal. Soils and sustainable agriculture. a review. *http://dx.doi.org/10.1051/agro:2007025*, 28, 03 2008. doi: 10.1051/agro:2007025.

[7] Zerina Kapetanovic, Ranveer Chandra, Tusher Chakraborty, and Andrew Nelson. Farmbeats: Improving farm productivity using data-driven agriculture. *SIAM News*, July 2019. URL https://www.microsoft.com/en-us/research/publication/farmbeats-improving-farm-productivity-using-data-driven-agriculture/. https://sinews.siam.org/Details-Page/farmbeats-improving-farm-productivity-using-data-driven-agriculture.

[8] R. J. Even, M. B. Machmuller, J. M. Lavallee, T. J. Zelikova, and M. F. Cotrufo. Large errors in common soil carbon measurements due to sample processing. *EGUsphere*, 2024:1–28, 2024. doi: 10.5194/egusphere-2024-1470. URL https://egusphere.copernicus.org/preprints/2024/egusphere-2024-1470/.

[9] Massimo F. Marcone, Shuhua Wang, Wafaa Albabish, Dhara Somnarain, Anna Hill, and Shuming Nie. Diverse food-based applications of nuclear magnetic resonance (nmr) technology. *Food Research International*, 51(2):729–747, 2013. doi: 10.1016/j.foodres.2013.01.051.

[10] Panagiotis Dais and Emmanuel Hatzakis. Quality assessment and authentication of virgin olive oil by nmr spectroscopy: A critical review. *Analytica Chimica Acta*, 765:1–27, 2013. doi: 10.1016/j.aca.2012.12.003.

[11] Ronald Soong, Myrna J. Simpson, and André J. Simpson. In vivo nmr spectroscopy: toward real time monitoring of environmental stress. *Magnetic Resonance in Chemistry*, 53(9):700–708, 2015. doi: 10.1002/mrc.4237.

[12] J.J. van der Klink and H.B. Brom. Nmr in metals, metal particles and metal cluster compounds. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 36(2):89–201, 2000. ISSN 0079-6565. doi: https://doi.org/10.1016/S0079-6565(99)00020-5. URL https://www.sciencedirect.com/science/article/pii/S0079656599000205.

[13] Evan Coleman, Sujay Nair, Xinyi Zeng, and Elsa Olivetti. Structured spectral reconstruction for scalable soil organic carbon inference. *Proceedings of the Climate Change AI Workshop at ICLR 2024*, 2024. URL https://www.climatechange.ai/papers/iclr2024/2.

[14] Lehmann, Kleber Johannes, and Markus. The contentious nature of soil organic matter. *Nature*, 528(7580):60–68, 2015. doi: 10.1038/nature16069.

A Bloch Equations

The time evolution of the magnetization components $M_x(t)$, $M_y(t)$, and $M_z(t)$ in the presence of a magnetic field $\mathbf{B}(t)$ can be described by the following differential equations:

$$\frac{dM_x(t)}{dt} = \gamma \left(\mathbf{M}(t) \times \mathbf{B}(t)\right)_x - \frac{M_x(t)}{T_2} \tag{1}$$

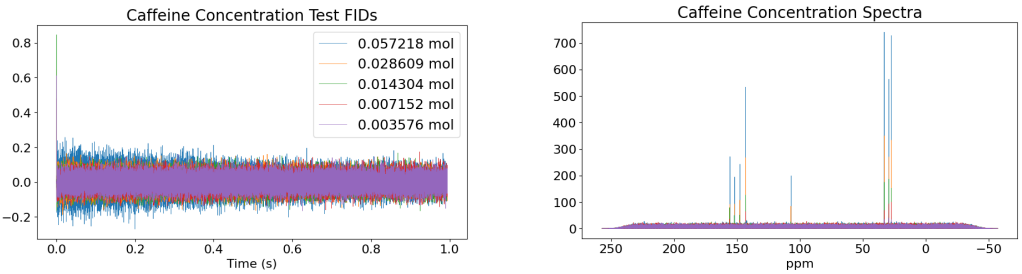

(a) Free Induction Decay Signal of Caffeine Samples (b) NMR Spectra of Caffeine Samples

Figure 2: Real NMR Data of the Serial Dilution Caffeine Dataset

$$\frac{dM_y(t)}{dt} = \gamma \left(\mathbf{M}(t) \times \mathbf{B}(t)\right)_y - \frac{M_y(t)}{T_2} \tag{2}$$

$$\frac{dM_z(t)}{dt} = \gamma \left(\mathbf{M}(t) \times \mathbf{B}(t)\right)_z - \frac{M_z(t) - M_0}{T_1} \tag{3}$$

These macroscopic equations, known as the Bloch equations, detail the underlying calculations that are done to calculate how the spins evolve with respect to time. The gyromagnetic ratio γ determines how fast the spin revolves around the external magnetic field. $B(t)$, the overall magnetic field, is the sum of the vectors detailing the external magnetic field, B_0, and B_x, the magnetic field applied by the model in the x direction. In our experiment, the ratio of B_0 to the amplitude of the pulsed B_x is 1 to 0.2, which is small in terms of typical NMR spectroscopes, but reasonable for a low-voltage NMR spectroscope employing a permanent magnet. T_1 and T_2 are the relaxation variables, which detail how fast the spin returns back to being in line with the external magnetic field after the pulse sequence is ended. Our NMR simulator uses these equations to determine the state of our environment, however, the state does not contain information about the spins themselves, but rather the magnetization that is being observed, as would be the case in practice.

B Validation Data Collection

The dataset used for the validation portion of our experiment consists of NMR scans of a series of caffeine samples which were repeatedly diluted from one original sample. The raw signals of these scans can be seen in Figure 2a, which was taken over 8 hours and then averaged to obtain a one second interval. This was done in a JEOL NMR and demonstrates the characteristic signal-to-noise ratio for a highly controlled NMR setup. The Fourier transform of these signals is computed to obtain the NMR spectra of these samples, which can be seen in figure 2b. Our initial distribution for the Monte Carlo sampling procedure is obtained from the 99th percentile of points in the sum of the NMR spectra. This is reasonable in practice as one can gradually collect small soil samples as this practice is used *in situ*, and occasionally send a composite sample back to a lab to update the sampling distribution.