Quantum Walk for Moisture Dispersion in Soil: Towards Open Systems and a Real-World Hamiltonian

Martin Castellanos-Cubides*
Applied Physics, Columbia University
(Dated: September 12, 2025)

In our previous work we simulated moisture dispersion in soil using a quantum walk framework, modeling soil patches as nodes on a qubit lattice with dynamically updated coupling constants based on local gradients. The previous paper on the subject served as a proof of concept, establishing the quantum algorithm and examining the limitations of classical simulations in higher dimensions. However, it assumed a closed system in which total moisture was conserved. Accounting for environmental interactions such as loss through evaporation or gains from watering systems, and the development of a Hamiltonian to emulate real-world behavior was left for future work. This paper addresses both of these challenges.

I. INTRODUCTION

Controlled-environment agriculture, such as greenhouses and indoor farms, offers a promising path toward sustainable and localized food production. These systems allow for precise regulation of light, air temperature, and humidity. However, soil moisture remains one of the most difficult variables to control. Unlike air or water vapor, which follow relatively simple thermodynamic laws, soil is a complex porous medium where other factors such as capillary action, gravity, evaporation, and plant root uptake interact in nonlinear and spatially varying ways. These effects are coupled and often depend on local gradients, making accurate modeling of soil moisture a computationally difficult task. Simulations must capture both global dynamics and local interactions, a requirement that scales poorly in classical computation. Accurately predicting how moisture disperses through soil is essential for optimizing irrigation, conserving water, and ensuring crop health.

A. Beyond Closed-System Models

In our previous work, we introduced a quantum walk framework to simulate moisture dispersion, representing soil as a lattice of qubits with dynamically updated coupling constants based on local gradients. That model treated the soil as a closed system, where total moisture was conserved. While this inherent conservation of the quantum walk was physically reasonable for an idealized isolated system, and even provided initial motivation, it fails to capture essential real-world processes such as moisture loss through evaporation or gains from irrigation. These interactions are inherently open-system effects, requiring the exchange of water between the soil and its surroundings.

* mc5534@columbia.edu

B. The Need for a Realistic Hamiltonian

The lack of real-world data limited our ability to develop and validate a Hamiltonian that accurately reproduced observed moisture behavior. Without this, the simulation could not be directly connected to known physical processes or reliably predict how the system would evolve under agricultural conditions. Establishing such a Hamiltonian is essential for systematically modeling how moisture flows, dissipates, and interacts with external reservoirs.

C. Scope of This Work

In this work, we extend our previous closed-system quantum walk model by incorporating environmental interactions that break moisture conservation. We introduce a modification to the algorithm that allows for moisture to be gained or lost dynamically, representing real-world processes such as evaporation and irrigation. To support and validate this model, we designed and built a controlled experimental system capable of monitoring moisture levels across discrete soil patches over time. By initializing known moisture distributions and tracking their evolution for a week, we gathered time-resolved data describing how moisture naturally spreads and dissipates. This data was then used to fit a first-order Hamiltonian, enabling a direct comparison between theoretical predictions and real-world behavior.

II. METHODOLOGY

A. Extending the Quantum Walk to an Open System

Our original algorithm modeled soil moisture as a closed system, where total moisture was conserved. In this framework, each soil patch is represented by a qubit, and moisture spreads through unitary diffusion between neighboring qubits. While this formulation is mathematically consistent and ideal for purely theoretical simulations, it fails to capture real-world processes, where moisture is continuously lost and gained by the system.

To address this limitation, we introduce an ancilla-based open-system extension to the original algorithm that allows each node to dynamically exchange moisture with an external reservoir. This modification breaks conservation by enabling moisture to flow into or out of the system during each timestep.

1. *Ancilla Registry Structure*

In this extended model, each qubit representing a soil patch is paired with a dedicated ancilla qubit. The primary qubit encodes the current moisture level of that soil patch, while the ancilla partner qubit acts as a local environmental reservoir, initialized either in the $|0\rangle$ or $|1\rangle$ state.

- $|0\rangle \rightarrow$ External source ready to add moisture (irrigation, rain)
- $|1\rangle \rightarrow$ External sink ready to remove moisture (evaporation, root uptake)

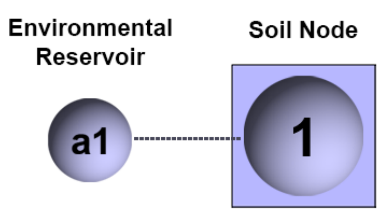

FIG. 1. Each soil patch is represented by a primary qubit paired with an ancilla qubit acting as a local environmental reservoir. Moisture exchange occurs between the two at the start of each timestep.

2. *Algorithm Step Breakdown*

Each timestep of the open-system simulation consists of three key stages:

Stage 1: Ancilla Connection via Hamiltonian Interaction Term

At the start of each timestep, the primary qubit and its ancilla are connected through a Hamiltonian interaction term, which governs the flow of moisture between the soil node and its external reservoir. The strength of this coupling can be defined through an arbitrary equation, making it highly flexible. For instance, the coupling term can depend on the current state of the soil qubit or on external variables such as time, environmental conditions, or targeted irrigation patterns. By adjusting this term, different environmental scenarios can be simulated within the same framework.

$$H_{int} = g\left(t, \psi_i, ...\right)\left(\sigma_i^+ \sigma_{i_a}^- + \sigma_{i_a}^+ \sigma_i^-\right) \tag{1}$$

If we wish to maintain a closed system for a given timestep, where no moisture is added or removed, the coupling strength is simply defined to be $g\left(t, \psi_i, ...\right) = 0$. In this case, the interaction term has no effect, and the ancilla qubit can effectively be ignored, reducing the system to the original closed quantum walk formulation.

Stage 2: Addition or Removal of Moisture During a Timestep

Once the coupling strength $g\left(t, \psi_i, ...\right)$ between the soil qubit and its ancilla has been defined, the system evolves using the standard timestep operator. The direction of moisture exchange depends on the initial state of the ancilla qubit, as previously defined in *Ancilla Registry Structure* ($|0\rangle \rightarrow$ drain, $|1\rangle \rightarrow$ add).

Stage 3: Measurement and Reset of Ancilla Registry

After allowing moisture exchange during a single timestep, we measure and reset the ancilla registry to its initial state, either $|0\rangle$ or $|1\rangle$. The measurement and realignment allows for a net change in moisture, as it escapes/enters through the ancilla qubits.

To provide intuition, the ancilla qubits can be thought of as buckets that either receive water from the soil or pour water into it:

- **Removal case ($|0\rangle$):** Imagine the ancilla as an empty bucket placed beneath its partner soil node. At each timestep some water falls into the empty bucket from the soil node. Measurement and realignment would then be dumping out the water to return to a completely empty bucket for the next timestep. The water that was in the bucket has now exited the system through the realignment, resulting in a net loss of moisture.
- **Addition case ($|1\rangle$):** Here the ancilla starts as a full bucket. During each timestep some of that bucket's water trickles into soil node. Measurement and realignment is then refilling the bucket to return to a completely full bucket for the next timestep, resulting in a net moisture gain.

B. Experimental Setup and Data Collection

To validate the open-system quantum walk framework and provide real-world data for modeling, we designed and executed a controlled soil moisture experiment. This experiment produced a continuous, time-resolved dataset that was later used to fit the Hamiltonian model and evaluate how well it captured real-world moisture dispersion dynamics.

1. System Overview

The physical experiment was carried out using a custom-built soil monitoring box divided into four equally sized cubes arranged in a 2x2x1 grid. Each patch was treated as a discrete spatial region that directly corresponded to one node in the computational model described in Section A.

Two of the patches were initially watered, while the other two were left dry, creating a clear initial gradient for moisture to diffuse across. This layout was chosen to provide a simple, interpretable scenario that still exhibited meaningful dynamic behavior. The exact initial moisture levels and final observed distributions are presented in Results.

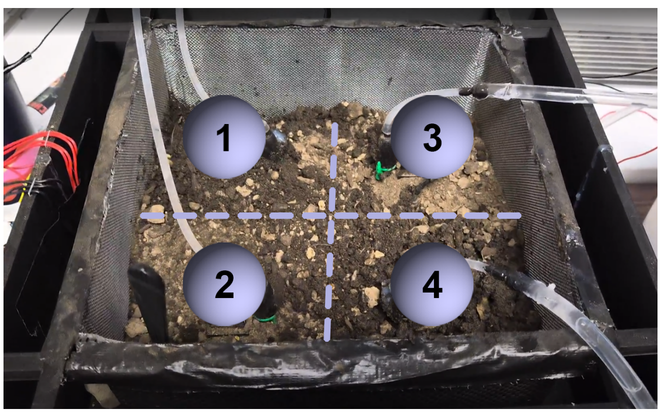

FIG. 2. Experimental box layout. Each of the four soil patches corresponds to a single node in the simulation.

2. Sensors and Data Logging

We used capacitive soil moisture sensors (AITRIP ASIN B094J8XD83) for reliable, corrosion-resistant moisture measurement across the soil patches. These sensors estimate volumetric moisture content based on the dielectric properties of the soil, offering improved durability compared to resistive probes.

Each sensor was embedded near the center of a soil patch and connected to a micro controller for automated data logging. The system recorded both analog voltage outputs and timestamps, allowing for high-resolution time series of moisture content over a 7-day experimental period.

3. Environmental Considerations

The experiment was conducted indoors to minimize uncontrolled environmental variables such as direct sunlight, wind, or rain. While small variations in room temperature and humidity naturally occurred, these effects were assumed to be uniform across the system and were therefore not explicitly modeled.

Because the primary goal was to study moisture diffusion and loss, no active irrigation or additional water input was provided after the initial setup. This allowed the natural interplay between soil patches and evaporation to dominate the system dynamics, providing a clean test case for fitting the Hamiltonian model.

C. Model Fitting and Hamiltonian Extraction

The second component of the methodology involved developing and fitting a Hamiltonian model to describe the moisture dynamics observed in the physical experiment.

1. Hamiltonian Formulation

To represent the system, we defined a Hamiltonian $H(t)$ with two primary contributions:

1. **Loss Term (Zeroth Order):** Represents evaporation or drainage of moisture from each soil patch to the environment. This term models how moisture naturally decreases over time, even without interactions between patches.

$$H_{i_{\mathrm{loss}}} = -\lambda \left(\sigma_i^+ \sigma_{i_a}^- + \sigma_{i_a}^+ \sigma_i^-\right)$$

   Every patch shares a common weight λ for environmental loss with their ancilla partner, limited to be independent of the gradient.

2. **Diffusion Term (First Order Gradient):** Represents moisture transfer between neighboring patches. This term captures the tendency of moisture to flow from wetter patches to drier ones, simulating the natural diffusion process in soil.

$$H_{ij_{\mathrm{diffusion}}} = -\Gamma \nabla_{ij}(t) \left(\sigma_i^+ \sigma_j^- + \sigma_j^+ \sigma_i^-\right)$$

   Every patch shares a common weight Γ for transitions between neighbors, limited to a first order dependence on the gradient.

2. Fitting Procedure Formulation

The values of λ and Γ were determined by fitting the model to the experimental dataset collected over six days. And evaluating success by minimizing the sum of the difference between the experimental and simulation end points.

3. Scope of the Model

This Hamiltonian was intentionally limited to a first-order diffusion term and a uniform zeroth-order loss term. While higher-order terms and spatially varying evaporation could be included, these were left for future work to maintain tractability and focus on validating the basic framework.

By keeping the model simple, we established a clear baseline for comparing theoretical predictions to real-world measurements, while still capturing the dominant physical behaviors of moisture dispersion and environmental loss.

III. RESULTS

A. Initial Moisture Distribution

At the start of the experiment, two patches were watered while the other two were left dry, creating a clear gradient for moisture diffusion. This initial condition allowed for both spatial spread of moisture between patches and overall loss due to evaporation to be observed over the course of the six-day run. The initial values observed for moisture were $\psi_i = [0.413, 0.175, 0.165, 0.476]$, normalized with the sensor's measurement of free space.

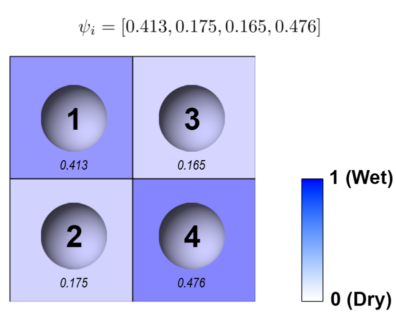

FIG. 3. 2x2 diagram of the initial moisture distribution, where darker shading represents higher moisture content.

B. Experimental Data Over Time

The final values observed at the end of the six days were were $\psi_f = [0.319, 0.141, 0.116, 0.401]$.

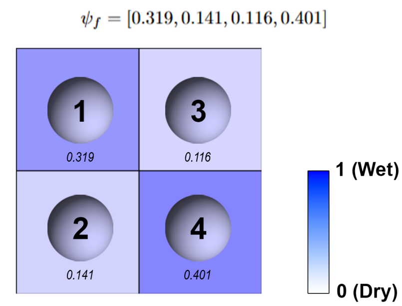

FIG. 4. 2x2 diagram of the final moisture distribution, where darker shading represents higher moisture content.

Moisture levels were recorded once per hour over six consecutive days for a total of 138 entries.

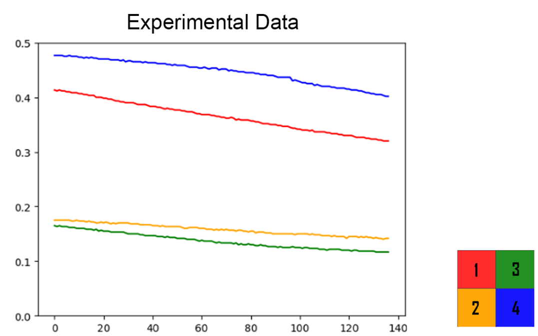

FIG. 5. Line graph showing moisture levels for each patch over six days.

C. Fitted Hamiltonian Solution

Without applying a formal optimization algorithm because of simulation costs, an approximate solution was found manually with time step $dt = 0.1$, as well as loss and dispersion factors $\lambda, \Gamma = [1, 0.1]$.

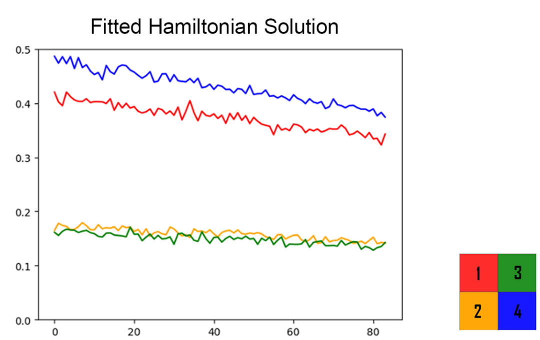

FIG. 6. Line graph showing the simulated moisture levels for each patch over 80 entries, where each simulated entry is scaled by a factor 40/69 to an experimental one hour entry.

IV. DISCUSSION

A. Interpretation of Results

The simplified Hamiltonian model was able to replicate the dominant moisture dynamics observed in the experiment, demonstrating that even a minimal formulation can capture essential features of soil-water behavior. In particular, nodes 1, 2, and 3 displayed approximately linear decreases in moisture throughout the seven-day period. This pattern was well matched by the combination of a uniform zeroth-order loss term and a first-order diffusion term, suggesting that these two mechanisms are sufficient to explain drying behavior in regions where evaporation and direct lateral moisture transfer dominate. The agreement in these three patches indicates that the core assumptions of the model are valid for much of the system under the controlled conditions of this experiment.

However, node 4 deviated significantly from this linear trend. While the other nodes declined at nearly constant rates, node 4's moisture decay was distinctly nonlinear, showing that the simplified Hamiltonian does not fully capture all processes at play. This discrepancy suggests that additional mechanisms, such as spatial heterogeneity or nonlinear environmental interactions, are influencing the dynamics of that region. The inability to reproduce node 4's behavior highlights the need for extending the model beyond its current formulation, as higher-order terms or more complex structures will be necessary to account for the richer physical behavior observed in the data.

B. Limitations of the Approach

There are several limitations that constrain both the accuracy and scope of this study. First, the Hamiltonian itself was deliberately restricted to only two terms: a uniform zeroth-order loss representing evaporation and a first-order gradient term representing diffusion between neighboring patches. While this provided clarity and tractability, it also meant that more complex spatial effects, heterogeneous loss rates, and time-dependent environmental factors were excluded from consideration. These omissions inevitably reduce the model's ability to replicate certain behaviors, such as the nonlinearity observed in node 4.

Another key limitation lies in how the fit between the model and experimental data was evaluated. In this work, the parameters λ and Γ were adjusted informally, with the model essentially tuned by eye to match the initial and final moisture levels. Although this approach was sufficient for demonstrating proof of concept, it lacks the rigor needed for precise quantitative validation. A systematic fitting procedure, such as minimizing the mean squared error across all nodes and timesteps, would provide a more objective measure of model performance and allow for direct comparison between different Hamiltonian formulations.

Finally, the experiment itself was highly constrained. The study was limited to a single $2{\times}2$ grid of soil patches under uniform, controlled conditions. While this made it possible to directly map the physical system to the quantum walk framework, it also represents a significant simplification relative to real-world agricultural environments. In practice, soil moisture behavior is influenced by a wide range of variables, including soil composition, plant root structure, uneven irrigation, and environmental variability. These factors were not represented in this initial experiment, limiting the generalizability of the findings.

C. Toward More Realistic Models

Despite these limitations, this work provides a clear path forward for building more realistic and comprehensive models. The strong agreement between the simplified Hamiltonian and three of the four nodes demonstrates that the core framework is sound, but the deviation observed in node 4 highlights where future development must focus. Introducing higher-order spatial terms into the Hamiltonian will be a critical next step for accurately capturing nonlinear dynamics and patch-specific behaviors. These terms would allow the model to represent more complex moisture flows that go beyond simple uniform diffusion.

Equally important will be the adoption of a more rigorous parameter estimation process. Rather than adjusting parameters by eye, future work should implement formal optimization techniques to determine the most accurate values of λ , Γ, and any additional coefficients introduced by higher-order terms. Such methods would allow for statistically meaningful comparisons between alternative models and provide confidence in the predictive power of the final formulation.

Finally, expanding the experimental system will be essential for testing scalability and realism. Moving beyond a $2{\times}2$ grid to larger arrays of patches, introducing nonuniform soil properties, and allowing for time-varying environmental inputs such as periodic irrigation or rainfall will all bring the experiment closer to real agricultural conditions. These improvements will create richer datasets and challenge the quantum walk framework to model increasingly complex and realistic scenarios.

V. CONCLUSION

This work extends our previous closed-system quantum walk framework to model soil moisture dispersion under realistic conditions where moisture is continuously gained and lost through environmental interactions. By introducing an ancilla-based modification to the original algorithm, we enabled dynamic moisture exchange between soil patches and external reservoirs, providing a mechanism to represent processes such as evaporation and irrigation within the quantum walk paradigm.

To support and validate this approach, we built a controlled experimental setup consisting of a four-patch soil box monitored continuously over six days. This experiment produced a detailed time series of moisture levels, capturing both the diffusion of water between patches and its gradual depletion to the environment. Using this dataset, we formulated a simplified Hamiltonian containing only two terms: a uniform zeroth-order loss term and a first-order diffusion term. These terms were adjusted to match the observed data, providing a direct connection between the simulation and measured soil behavior.

The model reproduced key behaviors observed in the experiment, including the approximately linear decline in moisture seen in three of the four patches and the gradual equalization of moisture across the grid. However, one patch exhibited nonlinear behavior that the current Hamiltonian could not capture, highlighting the need for higher-order terms and more complex formulations. This finding demonstrates both the strength and the limitations of the simplified model: it is sufficient to describe the dominant processes in many cases, but incomplete when faced with spatial variability or nonlinear dynamics.

While the fitting process here was informal, relying on visual matching rather than rigorous optimization, the results clearly demonstrate proof of concept. Future work will focus on introducing higher-order spatial terms, implementing systematic parameter estimation techniques, and scaling up the experimental system to larger and more realistic soil networks. These advances will allow the framework to be tested under increasingly complex conditions and evaluated with greater statistical rigor.

VI. REFERENCES

This is a continuation of previous work Ref. [1]

[1] M. Castellanos-Cubides, Quantum walk for moisture dispersion in soil: in one and two dimensions, MartinQQ , 5 (2025).