Attention Based Forecasting Coupled with Reinforcement Learning Policies for Adaptive, Sustainable Groundwater Management

Ouiam Kmich[1]*, Hicham Moutachaouik[1], Mustapha Hain[1], Badr Hirchoua[1]

[1]AIMCE Laboratory, ENSAM, University of Hassan II, Casablanca, Morocco
**Corresponding author: ouiamkmich86@gmail.com*

ABSTRACT

Groundwater systems worldwide are under increasing pressure from climate change, population growth, and rising water demand, creating a critical need for management approaches that can simultaneously ensure economically viable extraction and long-term aquifer sustainability. This study proposes a novel decision-support framework that couples a Transformer-based time-series forecasting model with reinforcement learning (RL) agents to generate adaptive extraction policies. A comprehensive dataset of 2,967,615 daily groundwater-level observations collected from 692 monitoring stations across California between 1970 and 2025 was analyzed to characterize regional temporal trends and to select a high-quality case study. The modeling and policy-learning experiments were conducted on station 21N02E26E005M, which provides a continuous 54-year record comprising 19,895 daily measurements, serving as a representative benchmark for evaluating the proposed methodology. The custom Transformer architecture was trained on this station's data to predict future water-surface elevations (WSE) from temporal features, lag variables, and engineered seasonal indicators. The model achieved exceptional predictive performance on a held-out test set, with a coefficient of determination of $R^2 = 0.9913$, a root-mean-square error (RMSE) of 0.2472 ft, and a mean absolute error (MAE) of 0.0963 ft. To translate these forecasts into actionable management decisions, a Gymnasium environment was constructed to simulate aquifer dynamics, incorporate safety bounds (50–200 ft), and apply a multi-objective reward function balancing economic extraction, environmental sustainability, and trend stability. Two state-of-the-art RL algorithms, Proximal Policy Optimization (PPO) and Deep Q-Network (DQN), were trained for 50,000 timesteps and evaluated over multiple 365-day episodes. Both agents learned policies that maintained groundwater levels within the prescribed safe range for 100% of the simulation period, thereby achieving perfect sustainability scores. The DQN agent, however, extracted substantially more water (approximately 16,222 units) while maintaining the same sustainability outcome, resulting in a higher average total reward compared with the more conservative PPO agent, which extracted only approximately 179 units. These findings demonstrate that integrating attention-based predictive models with RL-driven policy optimization can produce robust, adaptive groundwater governance strategies that reconcile economic and environmental objectives. The framework's capacity to forecast long-range temporal dependencies and continuously adjust extraction actions in response to evolving hydrogeological conditions offers a scalable pathway for real-time, climate-aware water resource management. Future work will extend this approach to incorporate spatial heterogeneity, more detailed aquifer physics, and stakeholder preferences, further enhancing its applicability to diverse hydrosocial contexts.

Keywords: Groundwater Management, Transformer Networks, Reinforcement Learning, Sustainable Water Governance, Predictive Modeling.