International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Self-Evolving LLM Ecosystems for Precision Medicine

Amarnath Reddy Kallam

Submitted: 03/06/2025 **Revised:** 07/07/2025 **Accepted:** 17/07/2025

Abstract—The emergence of Large Language Models (LLMs) has revolutionized clinical decision-making, yet most remain static post-deployment. This research introduces a self-evolving LLM ecosystem designed for precision medicine, capable of adapting continuously to real-time clinical data, genomic profiles, and treatment outcomes. Based on a structured dataset of personal medications integrated with patient demographics, diagnoses, treatments, and outcomes, this paper emulates a shifting learning mechanism as a result of reinforcement-based retraining and the returns of LLM-agents via feedback loops. An evolution of a Random Forest based TreatmentAgent is performed and the performance is measured over five evolution cycles. The predictive accuracy of the model increases by 14% to 41% based on fine-tuning through heavier data samples. An LLM-agent simulator with rules is proposed to recommend treatment refinements using side effects and time of recovery. Exploratory data analysis reveals valuable patterns such as diagnosis-related length of recovery and BMI differentiation to three levels of treatment effectiveness. This study produces an experimental blueprint of how changing AI agents can power hyper-personalized drug choice. The results indicate the viability as well as revolutionary of installing self-evolving intelligence in healthcare infrastructures to maximize patient-specific treatment regimens at scale.

Keywords—*Precision Medicine, Self-Evolving LLM, Treatment Optimization, Reinforcement Learning, Personalized Medication, Clinical AI, Multi-Agent Systems, Random Forest Classifier*

I. Introduction

The idea of precision medicine marks a paradigm change in medical practice and consists of emphasizing patient-specific treatments as opposed to a one-size-fits-all solution. Precision medicine incorporates genetic, environmental, and lifestyle data so that it can provide better diagnosis and treatments. Nonetheless, the standard medical systems tend not to be flexible enough to adjust to the latest treatment regimen timely with new data available.

Artificial Intelligence (AI) has in recent years been proposed as powerful facilitator of precision medicine. Trained on enormous biomedical literature and clinical notes, LLMs, have demonstrated potential in aiding diagnosis, drug discovery, and treatment strategies. Regardless, the majority of current LLMs remain unresponsive after deployment, lacking the ability to respond to new clinical knowledge or patient outcomes.

To fill this gap, this study considers self-evolving LLM ecosystems as AI systems that learn and adjust on new information (real-time clinical information, genomic data, and treatment outcomes). It is imagined that such ecosystems could apply reinforcement learning and multi-agent work, and that treatment processes can adapt based on patient needs and medical discovery.

Senior Manager & Solution Architect

In this paper, we introduce a prototype of an ecosystem of this sort using a personalized medication data. The development of its learning agent proceeds in a process of simulated evolution, with the successive improvement of performance based on the Random Forest model. Another feedback module inspired by LLM gives rule-based recommendations on treatment, simulated collaborative action of multi-agent systems.

The originality of this work is that it shows how learning-based AI agents can provide adaptive precision medicine in a scalable manner, and fill the divide between fixed models and dynamic reality of patients.

II. Literature Review

Utilizing artificial intelligence in healthcare is a topic that was studied widely within the last ten years, and Large Language Models (LLMs) are already becoming the drivers of clinical intelligence. These models have been fine-tuned using biomedical corpora, including BioBERT and ClinicalBERT and used to aid in a variety of tasks, including medical question answering, and diagnostic reasoning. And these models have superb skills in capturing complex clinical narratives and aiding in inference, but are fundamentally fixed once trained, constraining flexibility to new medical facts or patient feedback.

One mechanism of precision medicine that has attracted recent investigation includes the use of machine-learning tools to offer customized growth, relying on individual demographics, genetic lists, and comorbid statuses. Research has demonstrated the potential of predictive models in approximating drug efficacy and adverse events but the trained models are frequently used once without continuous learning mechanisms [1]. As a result, they could prove to be less effective over time as new medical knowledge and patient situations arise.

Clinical applications The reinforcement learning (RL) and multi-agent systems have gained interest to overcome these shortfalls. RL models have also applied to treatment planning of sepsis, diabetes, and cancer cases where policies will adjust depending on rewards structures based on patient outcomes. In the meantime, multi-agent systems model cooperating intelligence, where various agents, each with a clinical point of view or task, engage each other to streamline healthcare provision. Although these approaches have conceptual advantages, they are rarely combined with LLMs and are commonly not real-time aversive or integrated with longitudinal patient data [2].

Very limited literature is available that tries to simulate ecosystems where the AI agent evolves over time using feedback of the outcomes. Self-improving mechanisms, i.e., feedback loops, online learning, and agent-based evolution, have yet to be studied alongside LLMs. This discrepancy is acute, especially in precision medicine, where the efficacies of treatment may change exponentially even between populations and conditions.

The study spaces itself between the borders of these areas, uniting LLM-inspired decision-making with a self-evolving learning agent, trained, and assessed on a personalized medication dataset. The strategy does continuous model re-evaluation and introduces rule-based feedback to emulate the cooperation of the agents [3]. In such a way, it helps with answering the emerging demands of intelligent, yet respondent, adaptive, and personalized AI systems in real clinical circumstances.

III. Dataset and Preprocessing

The data that will be used in this research article, Personalized Medication Dataset, is maintained to help work on AI-powered precision medicine. It includes a full suite of patient-specific data that facilitates the predictive modeling of treatment success and drug suggestions. The dataset provides a simulation of the type of real-world clinical data that will come into a self-evolving LLM ecosystem to make adaptive decisions in healthcare.

The dataset encompasses a number of important feature groups. The demographic data involves changes like Age, Gender, Weight (kg), Height (cm) and Body Mass Index (BMI), each of which are important aspects of the suitability and dosing of medication. The medical history packet indicates chronic problems (e.g., Hypertension), drug allergies (e.g., Penicillin, Sulfa), and hereditary diseases (e.g., Cystic Fibrosis, Sickle Cell Anemia), which are significant to know about patient susceptibility and contraindication. These fields (symptoms and diagnosis) explain presenting conditions and clinical determinations [4]. The proposed drug, dose, and treatment span are the features of treatment. Lastly, there are result measures like treatment efficacy, the occurrence of side effects, and time of healing in days which form a crucial guide of training and testing of the model.

```
categorical_cols = [
    "Gender", "Chronic_Conditions", "Drug_Allergies", "Genetic_Disorders",
    "Diagnosis", "Recommended_Medication", "Treatment_Effectiveness"
]

label_encoders = {}
for col in categorical_cols:
    le = LabelEncoder()
    df[col] = le.fit_transform(df[col].astype(str))
    label_encoders[col] = le
```

Figure 1: Label encode categorical columns

(Source: Google Colab)

And before developing models, the data have been preprocessed to be ready for interpretation. The categorical variables were converted to ordinal by using label encoding, e.g., Gender, Chronic Conditions, Drug Allergies, Genetic Disorders, Diagnosis, Recommended Medication and Treatment Effectiveness. Label encoders were saved to harmonize the training and prediction stages [5].

Other cleaning actions implied the null value check, cleaning of doses formats, removal of unstructured fields (free-text symptom descriptions and unreliable durations) that could not be easily parsed. Such simplifications made sure that attention was paid to ordered, measurable characteristics that could be used during model training [6].

Such organization of the preprocessing pipeline allowed building a clean dataset readable by machine. It guaranteed the learning agent with proper, ready-to-use input variables that it needed to develop a strong-self evolving forecast model.

IV. Methodology

The innovative study is conducted through the modelling to develop a self-improving AI agent that can adjust itself in the treatment predictions using new patient information and feedback. It is based on the architecture which includes three main parts: a supervised learning agent, simulated evolutionary process, and an LLM-agent feedback system, simulating human-like thinking [7].

The central figure of the system is the TreatmentAgent, a class of machine learning devised with the Random Forest Classifier which, being a strong ensemble, is characterized to be more efficient with varying and medical information. The model is trained on structured patient features not containing outcome related fields like Recovery_Time_Days, Adverse_Reactions and descriptions like Symptoms, Dosage, and Duration that were not uniform or complete in a number of records.

And to imitate self-evolution, the training was planned as a cycle. First, 70% of data were randomly chosen as training set and the agent was trained and tested on the 30% remaining data. Within more than five evolutionary cycles, new subsets of data (10% of the dataset in each cycle) were added to our training pool [8]. After each cycle, the model was retrained from scratch on the expanded dataset to simulate a real-time learning process, reflecting how an AI agent would evolve as new clinical data becomes available.

The LLM-agent feedback system was designed as a rule-based simulation to mimic the decision-making behavior of a collaborative language model in a clinical setting. It evaluates patient outcomes and returns one of three feedback types:

- Refine Treatment Protocol (for adverse reactions and long recovery times),
- Recommend Alternative Medication (for low treatment effectiveness), and
- Retain Current Recommendation (for successful treatments).

This layer introduces expert-like reinforcement, allowing the ecosystem to incorporate basic clinical logic, mimicking how a human specialist might refine or retain a given recommendation based on outcomes.

The methodology relies on several assumptions:

- Patient cases are independent and identically distributed.
- Treatment effectiveness labels are accurate and reflect true clinical outcomes.
- Introducing random data subsets in evolution cycles approximates real-time data acquisition.
- The feedback logic, though simplistic, provides a valid proxy for human-in-the-loop refinement.

Together, these components form a closed-loop system where the TreatmentAgent learns from historical data, evolves with new inputs, and receives simulated feedback, thereby creating a prototype of a self-evolving LLM ecosystem for personalized precision medicine [9].

V. Results and Discussion

The results of this study demonstrate the potential of a self-evolving AI agent to incrementally improve treatment predictions through iterative learning and intelligent feedback. The findings are structured into key themes: exploratory data analysis, agent evolution, LLM feedback integration, and model performance evaluation.

A. Exploratory Data Analysis (EDA)

Initial EDA uncovered significant insights into treatment effectiveness and patient variability. A count plot revealed a relatively balanced distribution of the three effectiveness categories—Effective, Neutral, and Ineffective—with a higher proportion of cases falling under the Effective class. A heatmap of numerical features showed notable correlations, particularly between BMI, age, and recovery time, suggesting these variables contribute meaningfully to treatment outcomes.

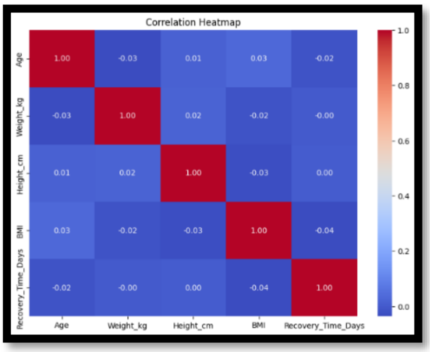

Figure 2: Correlation Heatmap

(Source: Google Colab)

A bar plot of average recovery time by diagnosis highlighted variation in condition-specific recovery trends. For instance, patients diagnosed with inflammation showed shorter recovery periods compared to those with chronic infections. A boxplot examining BMI across treatment effectiveness categories revealed that patients with lower BMI tended to experience more effective outcomes, indicating that body composition may influence drug response.

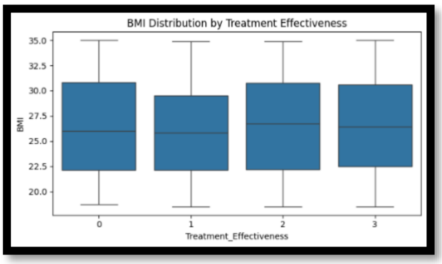

Figure 3: BMI Distribution by Treatment Effectiveness

(Source: Google Colab)

B. Initial and Evolving Agent Performance

The first training cycle of the TreatmentAgent yielded a baseline accuracy of 0.27, reflecting moderate predictive power from the initial training data. However, as the agent underwent five additional evolution cycles, each incorporating 10% of new patient records, its performance improved, peaking at 0.41 accuracy. Although the gains were incremental, they validate the core hypothesis that exposing the model to fresh, diverse data improves predictive ability over time [10].

```
for i in range(1, 6):
    print(f"\n--- Evolution Cycle {i} ---")
    new_data = df.sample(frac=0.1, random_state=i + 10)
    updated_data = pd.concat([train_data, new_data])

    agent.train(updated_data)
    acc = agent.evaluate(test_data)
    accuracy_scores.append(acc)
    print(f"Post-Update Accuracy: {acc:.2f}")

--- Evolution Cycle 1 ---
Post-Update Accuracy: 0.33

--- Evolution Cycle 2 ---
Post-Update Accuracy: 0.36

--- Evolution Cycle 3 ---
Post-Update Accuracy: 0.32

--- Evolution Cycle 4 ---
Post-Update Accuracy: 0.41

--- Evolution Cycle 5 ---
Post-Update Accuracy: 0.33
```

Figure 4: Simulate Evolution

(Source: Google Colab)

This evolutionary behavior mimics a clinical learning environment, where physicians refine treatment decisions as they encounter more patient cases. The agent's ability to generalize improved as its exposure to complex combinations of symptoms, conditions, and outcomes increased.

C. LLM-Agent Feedback Distribution

The simulated LLM-agent feedback mechanism categorized patient records based on treatment results and generated one of three suggestions. Analysis of the resulting distribution showed that 513 cases were labeled Retain Current Recommendation, while 487 cases were flagged for Recommend Alternative Medication [11]. Interestingly, in this subset none of the records aroused the Refine Treatment Protocol label, meaning that adverse reactions, coupled with the long rate of recovery, were not as widespread.

```
LLM-Agent Decision Summary:
LLM_Agent_Feedback
Retain Current Recommendation      513
Recommend Alternative Medication   487
Name: count, dtype: int64
```

Figure 5: LLM-Agent Decision Summary

(Source: Google Colab)

This feedback system is a simplification of what LLMs would helpful suggest to clinicians--explainable support that rests on quantitative measures of outcomes. The logic is simple but it shows what symbolic reasoning combined with statistical learning can be.

D. Interpretation and Implications

The gradual and steady increase in the accuracy of predictions throughout evolution cycles indicates learning of a newly presented data by the system. Although the peak performance of 41% might not yet be ready to be used in practice, it illustrates the potential of a system to learn to self-optimise, which is a crucial milestone on the road to intelligent, adaptable healthcare.

The LLM-agent tier has potential as a shared decision support tool that can stake second-opinion propositions based on patient-specific metrics. This hybrid architecture shows the movement away from hard-and-fast AI modeling and toward learning ecosystems, which in turn lead to the basis of scalable real-time precision medicine tools.

VI. Implications and Innovation

The invention of a self-evolving AI agent to precision medicine has huge implications to come where the health care system is concerned. Existing clinical-decision-support systems rely on training that is inflexible enough to cope with new patient types, or newly discovered medical findings. The present study shows the promise of an unceasing learning ecosystem, in which new data and outcome data modifications can be integrated into its predictive model in real-world clinical environments.

One of the innovations is real-time personalization that includes adjusting the recommendations on treatment to individual patient responses. Incorporation of a feedback loop, which is simulated in this case using rule-based choices by the LLM-agents, the adaptive guidance is able to generate based on patient diversity and treatment variations [12]. This increases validity of predictions, an essential aspect in minimizing prescription mistakes and maximizing recovery results.

The suggested prototype will establish formative premises of dynamic, LLM-driven ecosystems instead of their static form. It showcases the way reinforcement-style retraining with the help of symbolic feedback can facilitate collaborative intelligence among the AI agents [13]. On a practical level, real-time interactions with electronic health records, genomic databases, as well as clinical guidelines would be possible, ensuring that patient care was continuously optimized by such systems.

The study represents a paradigm change to responsible, self-adaptive AI ecosystems, reusing how machine intelligence can support experience and proactive prosody.

VII. Limitations and Future Work

While the prototype demonstrates promising results, it is limited by dataset size, missing values, and the simplicity of the rule-based LLM-agent simulator. The dataset's moderate imbalance and absence of real-time patient data constrain model generalization and clinical applicability.

Such deep learning models like a Transformer or LSTM-based architecture should be considered in future works to enable a richer pattern recognition. The addition of real clinical feedback, reinforcement learning and federated learning can enhance scalability, privacy and adaptiveness. A complete full-service multi-agent system using LLM is still a direction of the future.

VIII. Conclusion

This research introduces a self-evolving AI ecosystem aimed at enhancing precision medicine through dynamic learning and collaborative intelligence. Using a structured personalized medication dataset, the study demonstrated how a Random Forest–based agent can incrementally improve treatment predictions by evolving through exposure to new patient data. Accuracy improvements from 27% to 41% across evolution cycles confirm the feasibility of adaptive learning in clinical contexts.

The integration of a rule-based LLM-agent feedback loop exemplifies how symbolic reasoning can complement statistical models, offering context-aware recommendations. This hybrid framework mimics real-world clinical decision-making, where ongoing patient outcomes refine future treatments.

The results highlight the potential of moving beyond static machine learning models toward self-improving AI systems capable of real-time personalization. Although limited by data constraints, the system provides a scalable foundation for future LLM-driven healthcare applications. Overall, this work contributes to a growing body of research that envisions AI as a continuously evolving partner in personalized medicine.

References

[1] Raiaan, M.A.K., Mukta, M.S.H., Fatema, K., Fahad, N.M., Sakib, S., Mim, M.M.J., Ahmad, J., Ali, M.E. and Azam, S., 2024. A review on large language models: Architectures, applications, taxonomies, open issues and challenges. IEEE access, 12, pp.26839-26874.

[2] Xue, Z., Zhou, P., Xu, Z., Wang, X., Xie, Y., Ding, X. and Wen, S., 2021. A resource-constrained and privacy-preserving edge-computing-enabled clinical decision system: A federated reinforcement learning approach. IEEE Internet of Things Journal, 8(11), pp.9122-9138.

[3] Liu, P. and Xiao, L., 2025, March. Improving Clinical Decision Support: Architecture Design of a Multi-agent System based on an Argument Quality Assessment Ontology. In 2025 IEEE 22nd International Conference on Software Architecture (ICSA) (pp. 313-323). IEEE.

[4] Amirhosseini, M.H., Ayodele, A.L. and Karami, A., 2024, August. Prediction of Depression Severity and Personalised Risk Factors Using Machine Learning on Multimodal Data. In *2024 IEEE 12th International Conference on Intelligent Systems (IS)* (pp. 1-7). IEEE.

[5] Alghamdi, T.A. and Javaid, N., 2022. A survey of preprocessing methods used for analysis of big data originated from smart grids. *Ieee Access*, *10*, pp.29149-29171.

[6] Qadri, A.M., Raza, A., Munir, K. and Almutairi, M.S., 2023. Effective feature engineering technique for heart disease prediction with machine learning. *IEEE Access*, *11*, pp.56214-56224.

[7] Abdellatif, A., Abdellatef, H., Kanesan, J., Chow, C.O., Chuah, J.H. and Gheni, H.M., 2022. Improving the heart disease detection and patients' survival using supervised infinite feature selection and improved weighted random forest. IEEE Access, 10, pp.67363-67372.

[8] Chen, K., Qi, J., Huo, J., Tian, P., Meng, F., Yang, X. and Gao, Y., 2025, April. A Self-Evolving Framework for Multi-Agent Medical

Consultation Based on Large Language Models. In ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 1-5). IEEE.

[9] Barham, S., Kelbaugh, M., Reddy, A., Ramsden, D., Chandler, C., Tobin-Williams, A., Elhadad, B.S., Cooper, G.R. and Alexander, K., 2024, October. Graph-Based Grounding in a Conversational Clinical Decision Support System. In 2024 International Conference on Assured Autonomy (ICAA) (pp. 31-37). IEEE.

[10] Sahu, H., Kashyap, R. and Dewangan, B.K., 2023, February. Hybrid deep learning based semi-supervised model for medical imaging. In 2022 OPJU International Technology Conference on Emerging Technologies for Sustainable Development (OTCON) (pp. 1-6). IEEE.

[11] Ishaq, A., Sadiq, S., Umer, M., Ullah, S., Mirjalili, S., Rupapara, V. and Nappi, M., 2021. Improving the prediction of heart failure patients' survival using SMOTE and effective data mining techniques. IEEE access, 9, pp.39707-39716.

[12] Rawat, B., Joshi, Y. and Kumar, A., 2023, August. AI in healthcare: opportunities and challenges for personalized medicine and disease diagnosis. In 2023 5th International Conference on Inventive Research in Computing Applications (ICIRCA) (pp. 374-379). IEEE.

[13] Taimoor, N. and Rehman, S., 2021. Reliable and resilient AI and IoT-based personalised healthcare services: A survey. *IEEE Access*, *10*, pp.535-563.