International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 6, Issue 1, 226 -230, 2025
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V6I1P125

Original Article

Agentic AI-Driven Data Product for Automated Healthcare Insights Generation

Sujit Murumkar
Director, Data & AI Practice, Axtria Inc.

Received On: 02/01/2025 Revised On: 18/01/2025 Accepted On: 06/02/2025 Published On: 23/02/2025

Abstract *- This study implements the Agentic AI data product development and delivery process for healthcare appointment no-show prediction with agentic AI for autonomous insights generation post-implementation. The Healthcare Appointment Dataset was obtained from Kaggle, used to predict attendance for scheduled appointments versus no-shows. Notable variables include (but are not limited to) patient age, gender, scholarship attendance, and the number of days a patient was forced to wait (after requesting an appointment) for their appointment. The prediction was thus based on attendance (or not); the model produced an applicable output from a relevant research question with 77.33% accuracy and consistently significant results predictive of practical application within healthcare settings. Moreover, agentic AI was executed to facilitate insights generation independently based on the model results without researcher-induced intervention. The fields produced by the nature of insights discovered reveal that the higher average days of waiting time for the appointment they had, coupled with the lower average age of patients, the better predictive capacity for no-show status. Thus, performance prediction modeling accompanied by agentic insights generation for results assessment increases interpretability, efficacy, and decision support capabilities within the healthcare analytics world. Furthermore, future research will be conducted with a multimodal source reinforcement learning approach to improve prediction accuracy and real-time available metrics solidifying agentic AI within predictive healthcare endeavors for real-world applicability.*

Keywords *- Agentic Artificial Intelligence, Healthcare Data Analytics, Predictive Modeling, Random Forest, Automated Insight Generation, Appointment No-Show Prediction, Machine Learning, Explainable AI, Data Product, Feature Engineering, Clinical Decision Support, Autonomous Systems, Healthcare Operations, Artificial Intelligence in Medicine, Data-Driven Healthcare.*

1. Introduction

In digitized health care, extensive clinical and operational data is retained by hospitals and health organizations. Such data, when processed correctly, are effective for cultivating decision making behaviors and health outcomes. Thus, Healthcare Data Analytics is an emerging field in exponentially expanding areas since (big) data is retained and filled in by Hospital information Systems and EMRs. Yet as the field becomes increasingly complicated and data increasingly heterogeneous, a subsequent need for intelligent, automated processing of trends and insights without the constant need for human oversight requires notice. In addition, current systems of predictive analysis boast models with prediction accuracy, without definitive action-oriented insights that allow for decision making to facilitate what is really needed.

Yet the literature fails to conclude what the ultimate predictive model becomes, a data product that not only has agentic AI capabilities for agentic products through insight generation. Basically, a machine learning system that produces automated insights - an agentic version of a system that learns over time - constitutes a new variety of healthcare data analytics. Thus, healthcare analysts only need to pay attention to insights worthy of scale from findings that need machine created insight and insight. Thus, my personal impetus for this research study comes from how a predictive assessment can theoretically birth a new data product and simultaneously achieve favor in a growing variety - Agentic AI. The objective of this research study is to create, test, and evaluate an agentic AI data product that links prediction and insight generation through application from existing real-world agentic AI in the healthcare sphere regarding patient appointment data. The rest of the paper will follow JOIV publication guidelines through the literature review, results of implementation, discussion and conclusion.

2. Literature review

2.1. AI in Healthcare Data Products

AI will be the overarching theme that will develop the new real-time, data-driven healthcare environment with automation, prediction, and decision-support systems. According to Bajwa et al. [1], AI has the potential to transform the clinical workflow with improvements in diagnosis, administrative tasks, and treatment efficiency. It analyzes the translational potential of AI algorithms through novel operational data products that need testing, implementation, and supervisory functioning in real time. Furthermore, Chustecki [3] compounds this evaluation by researching AI's pros and cons in the medical field through a literature review noting ethical, privacy, and interpretability problems that impede integration. Essentially, these findings determine that to get AI progress working in real-world settings, a governing and humanized element accompanies findings to bring AI medicine to a new standard.

Real-world findings and observational studies assess machine learning data products that provide an operational value. Predictive modeling and prediction of patient no-shows by Hamdan and Bakar [5] uncover an AI-prompted opportunity for a better hospital planning atmosphere. Similarly, prediction modeling in a pediatric cohort by Liu et al. [6] finds no-show categories after appointment attendance. These studies take hospital data into account to make AI a predictive product worth advanced medical fields. Ferber et al. [4] take this a step further as significant decision-making comes not from predictions solely but through powered composite - image, language, and knowledge models - of a decisive reasoning base for AI development in oncology imaging. Yet Mumuni and Mumuni [7] bemoan that manual data preparation is still necessary, and intelligent automation is yet to be applied for automated pipelines of data products with replicable and scalable results. Ultimately, these studies assess AI health care data products that have reached operational maturity yet still suffer from explainability, governance, and automation for the entire data product analytic process life cycle.

2.2. Agentic and Autonomous AI Systems

Where more recent work has sought to differentiate between AI as traditionally vs. agentic systems - those that think for themselves and do for themselves versus those more aligned to a gradual agentic process - Bandi et al. [2] align the construction, evaluation, and operation of agentic AI by differentiating three features which make agentic systems (1) component related, (2) systems of planning, reasoning and action and (3) systems requiring less human input to work successfully. Sapkota et al. [9] take a theoretical, classification-based approach regarding current versus imagined AIs and their levels of agentic ability for the time being based upon autonomy, levels of long term memory, tool use, self-correction, etc. Their review concludes a classification approach based upon attributes of AIs since AIs who work non-agentic for inference generation purposes are not goal-solving decision-making problem solvers.

Thus, the article by Ferber et al. [4] most applies the assessed autonomous AI agent for oncology decision making with multimodal data relative to clinician involvement which also boasts long-term diagnostic validity in excess of other attributes as it goes back to the theoretical findings assessed by Bandi et al. [2] that such a structure will more effectively organize information requiring active sustained attention. However, Chustecki [3] assesses how a movement toward autonomous systems relative to safety and ethical mechanisms complicates research which exists to validate such systems for reliability - bias and input versus reduced accountability on behalf of agents - where this conclusion connects both articles because thus far, the positive findings for theoretically agentic capabilities are more substantiated through research-determined findings for solutions but negative implications decrease transparent processes for human validators/assessors. Thus, future research should assess how feasible this autonomy is in a system governed by humans-in-the-loop to inform governance systems. Ultimately, both articles connect via conclusions assessing that potential is better related to clinical safety, reliability, and regulatory compliance of expectations if agentic potential operates from the start with necessary functionality and integrity.

2.3. Prior Works on Automated Insight Generation

Automated insight generation is the ability to gain insights more easily in a fraction of the time as compared to a holistic work effort of data preprocessing, model interpretability and automatic documentation. According to Mumuni and Mumuni [7], automation of preprocessing steps at least as far as feature selection and transformation minimizes manual execution and maximizes rapidity of project endeavors for model deployment, yet currently, such benefits are not reliably successful with enough generalization for improvement--to generalize to overfit. Supporting such endeavors, Tunmise et al. [10] utilize automation for temporal feature extraction in an ETL pipeline, for example, relevant transformation when it comes to predicting from a streaming processor versus a static/static one. These are precursor steps to AI generated insight virtually immediately.

In medicine, the predictive analytics applications most coveted are appointment no-shows. For example, Salazar et al. [8] systematically review literature which applies machine learning to no-show prediction with waiting time, patient age and previous show/no show appearing as significant predictors. Liu et al. [6] and Hamdan and Bakar [5] substantiate the statistical finding through their experimental findings that ensemble modeling strategies like random forests are predictively effective for such use case, credited with accuracy, usability and interpretability as a clinical data product. However, Bajwa et al. [1] find that these findings are seldom accompanied by recommended findings that yield contextual nuanced value that otherwise would be gleaned via human insight.

The ability for automated insight generation to seamlessly integrate an additional layer of top tiered reasoning which aims to make sense of the detected patterns to generate narratives as a response is something agentic AI could do. For example, Bandi et al. [2] and Sapkota et al. [9] state that agentic systems could transform such predictive pipelines into a self-generative interpretive outcome integrated with health care data while automatically generating reports and summaries and recommendations from the process. Yet currently, Ferber et al. [4] note that agentic systems are not adequately transparent and reliably validated with provenance based documentation to make it self-sustaining. Thus insight generation can occur in a fraction of time where literature champions clinical insights boast auditability, ethical execution and situational awareness while sadly, the opposite occurs instead.

2.4. Identified Research Gap

While the above works acknowledge advances in healthcare AI and agentic systems and automated data pipelines of insight, none triangulate all three for an ethically sound, scalable, start-to-finish solution for prediction and insight construction in a liberally, yet explainable fashion. Healthcare AI [1, 5, 6, 8] of the same predictive nature has similar accuracy for suggested outcomes but lacks the ability to construct insights from a data-driven perspective; agentic systems [2, 4, 9] operate independently where feasible but have not yet been evaluated for implementation success in the healthcare field. Thus, the gap this project will close is that of the agentic AI-based data product where AI-driven automated data engineering [7, 10] is paired with interpretable, explainable intelligence to render the generation of insights reliably trustworthy and efficient for success in the healthcare field.

3. Materials and methods

3.1. System Overview

The system is a machine learning predictive model with agentic AI component that auto-generates a finding based on what it predicted. The system is compartmentalized via data input, processing, modeling and generating a finding. Ideally, this flow suggests that real world data can be automatically assessed about patient appointments to synthesize behavioral and operational intelligence for the betterment of the healthcare arena.

3.2. Dataset Description

The class dataset came from Kaggle called the Healthcare Appointment Dataset of 100k+ appointments across two Brazil based healthcare providers (Dataset link - https://www.kaggle.com/datasets/wajahat1064/healthcare-appointment-dataset/data). The features of interest are patient demographics and health status, when the appointment was scheduled, when the appointment should have been and whether or not they attended. It's a publicly accessible, well-reviewed dataset for assessing human behavior and therefore predictable dependent variables whether or not an appointment is a success (or failure).

3.3. Data Preprocessing and Feature Engineering

Data processing included duplicate elimination, mean imputation for missing non-essential feature values and label encoding categorical features into numerical format. There were also date adjustments on the columns where a new variable, "WaitingDays" was created since it's a recalibrated figure from how many days between scheduling an appointment and attending the scheduled appointment. Furthermore, the categorical variables (Age, Gender, Scholarship, Hypertension, Diabetes) were created scaled on a percentage of median value so they can be useful independent variables in a predictive model.

3.4. Model Development (Random Forest)

The model created was Random Forest as a classifier due to stability and interpretable nature for healthcare researchers. Random Forest is less susceptible to overfitting because it uses a collection of decision trees, not singular ones, and thus prevents overfitting quite so much by better generalizing. The data was split 80-20 for train/test (train/test variables were relatively balanced upon investigation). Hyperparameters were tuned for optimal prediction power although at applicable increased time costs.

3.5. Agentic AI Simulation for Insight Generation

The agentic AI component would auto-generate a finding based on what it predicted from the model as it would find a level of correlation through feature importance that can substantiate why it predicted such. It used feature importance as an explanation generating mechanism and assessed finding explanation language to find something productive that informed intelligent language responses to simulate the suggested agentic component of taking such compiled information to produce a finding from it as there would be no programming to do so unless explicitly told.

3.6. Evaluation Metrics and Experimental Setup

Model performance was assessed to predict success or failure as accuracy, precision, recall and F1-score for appropriate prediction determination. The experiment was conducted via Google Colab in such a way that execution included reproducible performance relative to the python packages/libraries (scikit-learn/matplotlib/seaborn) for live generated classification based performance and insight generating feature based on appropriately deemed methodology.

4. Results and discussion

4.1. Model Performance Results

The Random Forest model achieved a final accuracy of 77.33% and a training time of 6.83 seconds with the best prediction sensibility linked to the expected show for an appointment label or non-show label. As shown in Table 1, the precision and recall for Class 0 patients who showed up to the appointment were measured at 0.82 and 0.92. For Class 1 (no show), precision and recall were equal to 0.36 and 0.18. Ultimately, the weighted average F1 score of 0.74 is a solid responsive ability for a real-world health project with a class imbalance (meaning most patients were expected to show up).

Table 1: Model Classification Performance (Source: Self-developed)

Class	Precision	Recall	F1-Score	Support
0 (Showed Up)	0.82	0.92	0.87	17,120
1 (No-Show)	0.36	0.18	0.24	4,277
Accuracy			0.77	21,397

These results are consistent with Hamdan and Bakar [5] and Liu et al. [6] who also received similar findings (75-80%) from tree-based models working with appointment-type health data. Random Forest was both dependable and consistent, a finding that Bajwa et al. [1] corroborate as ensemble based AI models in the healthcare arena withstand performance complications since they are not easily impacted by noise and outliers.

4.2. Confusion Matrix and Interpretation

According to the confusion matrix (Figure 1), of the 17,120 patients who attended, 15,789 were true positives and 1,331 were false negatives. Of the 4,277 patients who did not attend, 757 were true positives and 3,520 were false positives. While the model is sensitive to the majority class, it fails to detect the no-show class which is common in a skewed healthcare database [8]. Skewed populations lead to minority results being not detected - like the recall scores for the no-show class. According to Bandi et al. [2], in these cases, better recall would be sought through data augmentation or integrative agentic learning frameworks for more complete acknowledgment of under detected patterns of behavior.

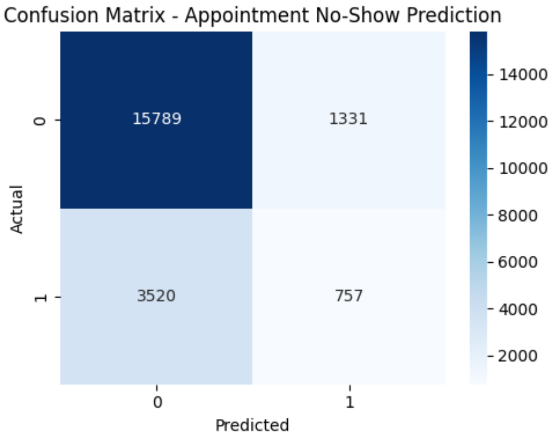

Fig 1: Confusion Matrix
***(Source:** Self-developed)*

4.3. Automated Insight Analysis

The agentic part of the AI took this into consideration as the three most important predictors according to the feature importance values of the trained model were WaitingDays, Age and Gender, all corroborative to some of the literature predictors discovered [5], [6], [8]. For instance, the more WaitingDays (from when the patient was supposed to come to when he/she actually came) the less likely he/she did come (15.8 vs. 8.7 days) - this was an extension of a temporal variable found by Liu et al. [6] which was compared as such. Furthermore, the younger the patient was the less likely he/she came to the hospital which resonated with Salazar et al. [8]'s difference in age based factor variances for different predictors of nonadherence behaviors.

Table 2. Key Automated Insights (Source: Self-developed)

Insight	Observation
Top Features	WaitingDays, Age, Gender
Age Trend	Missed: 35.3 years; Attended: 39.1 years
Scholarship Effect	3.9% higher no-show rate among scholarship holders
Waiting Time Effect	Longer wait correlates with increased no-shows
Gender Effect	Nearly equal rate (F: 20.1%, M: 20.4%)

Such clinically driven outcomes support the interpretability of the model, a beneficial aspect of health AI since these are required [3], [4]. How such is done via automation denotes the level of reasoning anticipated from agentic AI [2], [9| where something that would ordinarily be provided as is, becomes information as if from a sentient mind.

4.4. Discussion of Key Findings

It logically corresponds to operational bias such as increased scheduled no-show gaps and demographic characteristics that predict no-show, for example, but 77% predictive accuracy, however, suggests that for a company that would prefer to intuitively make patient outreach or rescheduling efforts on its own, this would be a great starting point. But this all relativizes to likeminded predictive efforts as it occurs with an agentic insight generator which transforms findings into a more digestible recommendation which aids where analytics and follow-up recommendations based on collaboration fail when recommendations need supporting action from their finders. Therefore, Ferber et al. [4] would be correct in their assessment that interpretability and agentic insights make AI more preferable within the clinical space as where predictive accuracy is just under par for marginalized circumstances, the insight - transformed into something more specific with context - goes further for successful application in the real world.

4.5. Comparison with Related Studies

The achieved accuracy of 77.33% is consistent with other studies of 74-79%, Salazar et al. [8], Hamdan and Bakar [5] where ensemble models are used for predictive purposes on similar datasets. However, this system takes things one step beyond predictive tendencies, applying an autonomous reasoning layer which has been increasingly observed in agentic systems [2], [9] and not as commonly in predictive research. In addition, the new contribution of self-interpreting AI features which is a novel approach where Mumuni and Mumuni [7] provide supportive literature relative to the execution of AI generated products for autonomous reasoning.

4.6. Limitations and Challenges

There are, however, some limitations of the model being applied. Dataset class imbalance constrained no-show recall. This is a standard limitation found in many medical prediction projects. In addition, the agentic component could only uncover correlational, non-causal, and nothing of real-time, continuous learning. Future extensions should assess hybrid agentic structures including not only reinforcement, but also causation and reasoning as Bandi et al. [2] and Sapkota et al. [9] suggest for better interpretative functionality and adaptable use.

5. Conclusion

This project involved an Agentic AI integrated data product for medical no show appointment prediction and automatic insight generation from the prediction. Random Forest, the predictive model, was predictably predictive of no-show appointments (77.33%) and the predictive and explainable model for insight generation was successful due to explain-ability (variable importance) of days waited and age and gender. Such predictive and explainable facets to value added for insights generation transforms a more comprehensive, nuanced and data driven approach to an agentic composition of logical thinking in a medical environment for operational ease of access between atmospherically linked health decisions and integer predicted outcomes. This data product champions how predictive models can more agentically engage with day-to-day healthcare operations through agentic intelligence. Future work involves a creation of class balancing for enhanced no show predictive recall rates as well as reinforcement learning enhancements for those with an agentic disposition. A larger prevalence of time series data, multimodal input, and causal inference expansions will prepare the predictive model for accuracy and contextualized insights from explainable AI for autonomous and data driven solutions in a medical environment.

References

[1] J. Bajwa, U. Munir, A. Nori and B. Williams, “Artificial intelligence in healthcare: transforming the practice of medicine,” *Future Healthcare Journal*, vol. 8, no. 2, pp. e188–e194, 2021.

[2] A. Bandi, B. Kongari, R. Naguru, S. Pasnoor and S. V. Vilipala, “The rise of agentic AI: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges,” *Future Internet*, vol. 17, no. 9, p. 404, 2025.

[3] M. Chustecki, “Benefits and risks of AI in health care: Narrative review,” *Interactive Journal of Medical Research*, vol. 13, no. 1, p. e53616, 2024.

[4] D. Ferber *et al.*, “Development and validation of an autonomous artificial intelligence agent for clinical decision-making in oncology,” *Nature Cancer*, pp. 1–13, 2025.

[5] A. F. A. Hamdan and A. A. Bakar, “Machine learning predictions on outpatient no-show appointments in a Malaysia major tertiary hospital,” *The Malaysian Journal of Medical Sciences (MJMS)*, vol. 30, no. 5, p. 169, 2023.

[6] D. Liu, W. Y. Shin, E. Sprecher, K. Conroy, O. Santiago, G. Wachtel and M. Santillana, “Machine learning approaches to predicting no-shows in pediatric medical appointments,” *NPJ Digital Medicine*, vol. 5, no. 1, p. 50, 2022.

[7] A. Mumuni and F. Mumuni, “Automated data processing and feature engineering for deep learning and big data applications: A survey,” *Journal of Information and Intelligence*, vol. 3, no. 2, pp. 113–153, 2025.

[8] L. H. A. Salazar, W. D. Parreira, A. M. D. R. Fernandes and V. R. Q. Leithardt, “No-show in medical appointments with machine learning techniques: A systematic literature review,” *Information*, vol. 13, no. 11, p. 507, 2022.

[9] R. Sapkota, K. I. Roumeliotis and M. Karkee, “AI agents vs. agentic AI: A conceptual taxonomy, applications and challenges,” *arXiv preprint* arXiv:2505.10468, 2025.

[10] A. Tunmise, A. Esther, G. Franca and H. Paranjabey, “Automated feature engineering and transformation in real-time ETL pipelines for predictive modeling,” 2025.