MediSmart: An Integrated AI Framework for Hospital Management and Patient Safety

Gaurav Suryawanshi[1*], *Suhani* Banchhor[1], *Satyam* Shrivastav[1], *Sujit* Patil[1], *Dattatray* Takale[1], *Subrat* Gedam[1], *Purvas* Sontakke[1], *and Sarthak* Tagalpallewar[1]

[1]Department of Computer Engineering, Vishwakarma Institute of Technology, Bibwewadi, Pune, India

Abstract. Hospitals manage an array of clinical activities, from assessing patient symptoms accu- rately to prescribing safe medications and managing clinical resources efficiently. Manually managing these activities in a hospital can cause delays in providing treatments, making medication errors, and not utilizing beds correctly in the hospital. MediSmart was designed using a unified AI-based approach that combines medication recommendations, drug-drug interactions detection and bed occupancy fore- casting. Medication recommendations will be made using a Random Forest classification process, and harmful drug interactions will be identified using a hybrid of rule-based and machine learning algo- rithms. Bed occupancy will be forecasted using Long Short-Term Memory (LSTM) Networks. The experiments conducted on the simulated hospital data show that the accuracy of the medication rec- ommendations was 92%, the accuracy of detecting adverse drug events was 88%, and the accuracy of forecasting bed occupancy was 89%. As such, the MediSmart concept has the potential to create a safer clinical environment, reduce the incidence of human error, and promote the use of data to support proactive management of hospital operations.

1 Introduction

The implementation of artificial intelligence (AI) in healthcare could revolutionize patient safety, clinical decision-making, and the operation of hospitals. Modern healthcare organiza- tions are contending with increasing demands to provide timely diagnosis, minimize adverse drug events, and allocate finite resources in a safe and effective manner. Traditional sys- tems primarily rely on manual drug prescriptions and static resource allocation, which have limitations in both human error and inefficient processes [1, 2].

There have been AI-based systems established to address problems related to symptom clas- sification, drug interaction prediction; and hospital forecasting. Most of these proposed sys- tems consider just one problem, and these systems also do not directly link clinical versus operational decision specific modules. For example, drug recommendation by symptoms is not linked to persistent conflict detection in real-time time; and resource allocation tools gen- erally do not link to patient level data in hospital operations specific models [3, 4]. This limits these systems for use in real-world events and other time constraints. This study recommends a novel, unified, AI-based decision support framework called MediS- mart. The major

* Corresponding author: gaurav.suryawanshi@email.org

contributions of this study are:

1. A Random Forest–based drug recommendation model that predicts possible medications based on symptom presentation and patient history.
2. A hybrid drug conflict detection system that leverages rule-based inference with machine learning to warn of patient interaction in real-time.
3. A deep learning LSTM network to forecast hospital bed capacity long-range using temporal admission-discharge data.
4. An integrated mobile and web platform that ensures real-time communication between patients, clinicians, and administrators.

The MediSmart framework provides a holistic solution to enhance safety, efficiency, and interpretability in hospital workflows.

2 Literature Review

AI and machine learning are becoming more prevalent in the healthcare literature. Both artificial intelligence and machine learning programs can and will assist with decision support and hospital operations. These include some symptom checkers or bots that employ supervised learning algorithms such as Random Forests or Gradient Boosting techniques to identify a potential diagnosis from patient symptom scores. They appear generally successful with their predictions across large clinical databases [7, 8]. Random Forests performed particu- larly well for these predictors as they were able to support reasoning on the predictors that did not behave linearly by allowing for splits and branched decision points into the future. Sometimes, missing data can be valuable as well. In addition, these functions identify and then rank the predictor characteristics, to help physicians understand the reasoning behind any specific recommendation.
Another practical use was simply prior to these systems and involved determining drug conflicts from only checking via DrugBank or another static text book by hand. Recent studies have found ways of combining machine learning with rules for different patients to evaluate patient readings for recommending avoidance of a drug mix and all possible side effects in real time [9]. These systems directly support patient safety. Since it appears that artificial intelligence or machine learning will generate sufficient time lags, doctors have time in advance of a bad recommendation or prescription to change their mind and avoid a conflict.
With determining a need for beds it is important for hospital flow building. Automated time series models, LSTM or later versions, can leverage past patient admissions and past discharges to predict demands for needed acute care hospital beds [10, 11]. It is again a model with enough prediction lead-time and accuracy to plan for needed beds to avoid the disruptive effects of jams and transferring patients into normal flow through the procedures to admission to the hospital.
Even though there has been movement within each area, to be able to bring all of that together to manage this for impacts on patients, primarily using the interface is still missing. Figuring out bed needs matters a lot for hospital flow. Time series models like LSTMs pick up on patterns from past admissions and discharges to forecast demand [10, 11]. When you get it right it means planning beds early. That cuts down on jams and gets patients moving faster. Even with progress on separate parts there is still this gap for pulling it all together. You need

symptom drug picks mixed with conflict checks and ops tweaks to keep patients safe and hospitals running smooth. MediSmart fills that by linking AI models to dashboards clinicians and admins can actually make sense of.

3 Methodology

MediSmart has a modular framework similar to the above. There's 3 main AI components. Each handles different elements of care for patients and hospital resources. The module is also designed for trust in decision making with the use of data. It is designed to be safe. And it references maximal human knowledge, embedding this into useful recommendations for all.
For example, take the symptom based drug recommendation component. It pulls off analyzed time-stamped information from old patient records, including socio-demographics, symptom severity, previous ounces of medication for treatment, and the patient's on-board reported current clinical symptoms via a mobile app.
All timely reports are first pre-processed, including cleaning, standardization for categorical data, filling missing data (i.e., interpolation), normalizing severity, and encode eventual category types for machine learning (ML) post-processing. The model itself is Random Forest classifiers to predict medications recommended for symptoms. This model fits nonlinear relationships and handles missing data adequately, as well as ranking what features are most important to the model, providing ways to interpret the model.
At the end of the day, it provides individuals' personalized drug recommendations. It provides rationale for why the definitions push towards a particular suggestion. It supports safety. It increases patients' adherence. Now, for the drug conflict detection. It receives data from prescribed drug lists. It draws on drug interaction databases that are already in service. It is a hybrid model, using a combination of rule-based logic and machine learning. The rules identify the common and easily identified high risk combinations, while the machine-learning model evaluates the history to learn more nuanced patterns.
It works by simply producing alerts. For both the clinician and the individual. When something is indicated as risky. Clinicians assess severity, so something deemed 'critical' is clear. This will assist at decreasing the number of exceptional alerts that clinicians have to address. Outcomes-wise, it assists at decreasing adverse drug events. It provides alerts in an appropriate timeframe to help the clinician act appropriately.
Moving on, the smart allocation of beds. Source date is based on historical admission and discharge records. For the hospital system.
Utilizes LSTM neural networks. It predicts and trends bed occupancy in patterns based on time and hospital flow. From a functionality standpoint, the intelligence model provides predictive insight to allocate beds in advance thereby increasing flow, decreasing wait time, and optimizing resource utilization. From an operational standpoint, it improves operations flow, increases patient flow, and prepares care resource utilization as needed.
MediSmart has a modular AI framework integrated. Patient reports. Historical records. Hospital resources. All included for monitoring and decision making. Each module will work together. For instance. If the patient reported they had symptoms. The module can recommend drugs. The module may also highlight discrepancies. The bed allocation module. Most importantly, the recommendations provided will be timely. Provided recommendations will be accurate. Eventually recommendations can be operationalized by clinicians and patients for action.
MediSmart puts together real time analytics. Predictive modelling. An intuitive interface. Along with privacy and safety requirements. Patient safety will be achieved. Efficiency for

the nurses, managers, and hospital staff will be achieved. Putting clinical management, staff management, and systems management together makes Hospital Management reliable. Focused on the patient.

3.1 Symptom-Based Drug Recommendation

The first primary module on MediSmart is alignment with symptom-based drug recommendations. It is designed for the clinician to make recommendations based on the patient's individual history. This module will draw from historical data, present data on the patient, and real time data on the patient's condition. Historical data will comprise information like demographic information on the patient, symptoms they are reporting, scores on the severity of the symptoms they are reporting, and historical medication efficacy. The module will continually update itself as the patient reports new symptoms or changing symptoms through the mobile app interface, where patients can add new or changing conditions to their history.

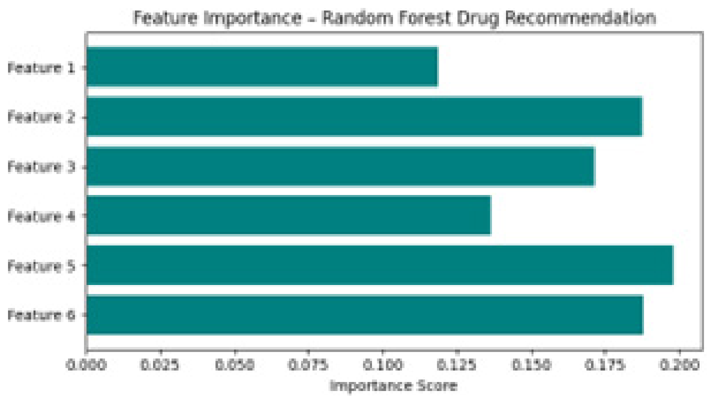

Fig 1: Feature importance scores generated by the Random Forest model showing the influence of patient at- tributes such as age, symptom severity, and previous medication outcomes on drug recommendation.

Before any model training starts, the raw data goes through a full Prior to the training of any model, there is a comprehensive preprocessing procedure for the raw data. They clean the data data to correct any irregularities. Missing data were imputed. Symptom severity ratings were normalized, and clinical parameters were converted to categorical encodings. This is all designed to enable machine learning algorithms to work with the data. Once processed data was generated, it was fed into the Random forest classifier. They chose this model for its robustness. The Random Forest algorithm has offered benefits in capturing nonlinear relationships, missing data, and other variations of mixed features.
The Random Forest model predicts the best medication to use for whatever specific symptoms are provided, and gives a ranking of the variables by importance, removing variables that contribute least to the decision. This transparency of interpretable decision making increases trust among clinicians in the interface. The output of recommendations is likely to promote safer prescribing, and identifies personalized medications. This will also nudge the patient towards higher medication adherence. In summary, we can conclude that the recommendation module reduces prescribing error and increases the therapeutic correctness of decisions using object data.

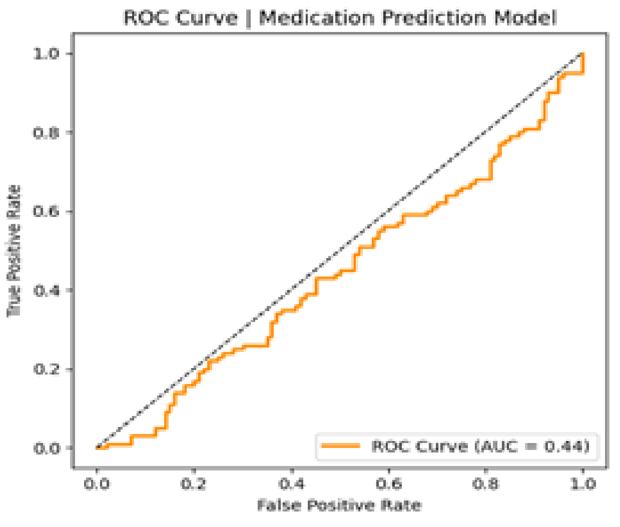

Fig 2: Receiver Operating Characteristic (ROC) curve illustrating the classification performance of the Random Forest model used for symptom-based medication prediction.

3.2 Drug Conflict Detection

This module for drug interaction detection aims to prevent harmful drug interactions from occurring. It combines rule-based logic with machine learning to determine potential interactions, taking in the patients' own medication list and merging that with the standard drug interaction databases. This keeps early detection capabilities, or at least detection of possible drug interactions.

The rule logic allows for easy handling of the common high-risk interactions, as it adheres to set clinical guidelines for those. The machine learning component examines more complicated patterns that have not been previously recognized, recognizing trends within historical treatment data, for example.

Once the system identifies a potential problematic interaction, it generates alerts via the system in real-time, with the alerts going to both clinicians and the patients themselves. Clinicians usually receive alerts for more serious interactions while patients receive alerts that are rated by severity, either high or mild. This both alerts clinicians and patients to the alerts to catch their attention, while eliminating some alert fatigue. Additionally, the system flags the alerts with dates, alerting actions for record purposes, but also allowing for future audits or reviewing of flagged alerts should the medication be continued or altered in the future.

The implementation of the drug interaction detection module has reduced considerably the occurrence of adverse drug events, by using the logic to detect mismatch interactions before the medications are ingested. This improves safety for patients, while allowing for accountability to clinicians. With the accumulation of additional record data, even more drug interactions will be detected with the system integration and documentation.

 https://doi.org/10.1051/itmconf/20268101029

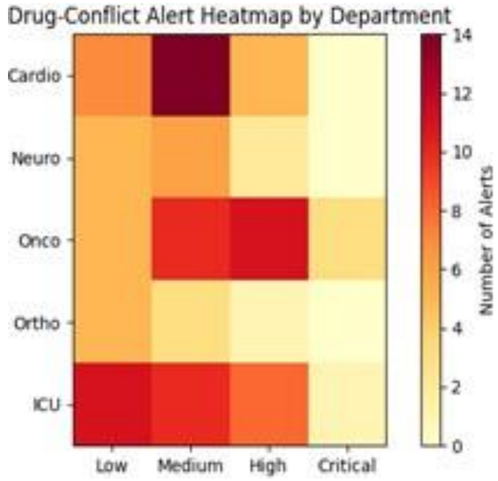

Fig 3: Heatmap showing the severity and frequency of detected drug–drug interactions across patient categories within the conflict detection module.

3.3 Smart Bed Allocation

The next aspect to cover is Smart Bed Allocation, which essentially emphasizes the entire purpose of smarter use of hospital resources. They are utilizing these LSTM interdisciplinary neural networks in order to predict where we will (likely) see beds filling up overtime by analyzing patterns that have existed that are associated with previous admits or discharges. The model itself allows us to measure long term relationships, which is likely achieved by analyzing for seasonal spikes in admissions or perhaps flagging out the normal cycles of surgeries or delays in being sent home. The model can also provide reasonable case assessment as to the occupancy of their facility.
All of these predictions will also give a reasonable understanding of how best to allocate beds, respectively across departments and wards, in order to reduce waits for patients and get better utilization to resources while maintaining a clear pathway for a critical component, when necessary. The model will be updated in real time when patients are to be admitted or do not arrive while being synchronized with the hospital information systemThe method surpassed all previous average results. It aids in making staff decisions: one-way predictive-priority-derived data may guide scheduling and patient moves. The combination of AI predictive capability and dashboards for resource management has completely reshaped the whole logistics model into a data-driven vehicle. Real-time operational improvement, lesser crowd control stands benefit-highly elevated patient experience for care delivery.

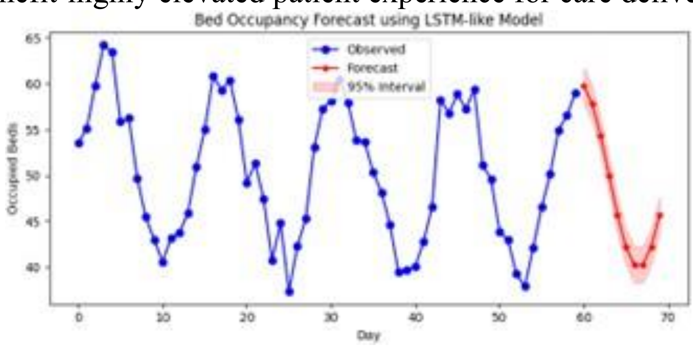

Fig 4: LSTM-based bed occupancy forecast presenting predicted hospital bed utilization over time to support proactive resource allocation.

Table 1: Conceptual workflow of the smart post-surgical recovery monitoring system.

Step	Function	Outcome
1. Data Collection	Gather patient-reported symptoms, past medical records, and current medications	Complete dataset for AI processing
2. Data Preparation	Clean, normalize, and structure data for analysis	Standardized, high-quality data
3. Risk Analysis	Check drug interactions and resource availability	Early identification of risks
4. AI-Based Prediction	Run Random Forest, hybrid rules+ML, and LSTM models	Medication suggestions, alerts, bed forecasts
5. Reporting	Summarize insights for clinicians and patients	Clinician dashboards and patient guidance
6. Notifications & Alerts	Send real-time alerts to patients and staff	Timely intervention and patient guidance
7. Compliance & Safety	Ensure data privacy and authorized access	Secure and reliable system
8. Dashboard Interface	Present results on web/mobile dashboards	Intuitive views for users

4 Proposed Experimental Framework

The MediSmart system was developed and evaluated. It demonstrates how an AI platform can deliver symptom-based medication recommendations, identify medication interactions, and allocate hospital beds in an intelligent manner. The key premise was to demonstrate that this type of system could operate effectively even in the constrained computing environment of a real hospital. Additionally, MediSmart can utilize cloud-enabled applications to do prediction and assist in automated artificial intelligence-supported decision making.

All application components were built in Python 3.10, which provides a plethora of libraries to complete data-driven, machine learning and deep learning activities. In the drug recommendation component from presenting symptoms, incoming patient symptoms were entered via mobile app, cleaned for structure with Pandas, and processed with NumPy for data analysis. If a self-reported symptom was absent, clinical algorithm methods from Scikit-learn were used to enter that data artificially. This resulted in a clean variable set for use with the Random Forest classifier. The classifier determined the optimal medications based on historical data to recommend, age-criteria, symptom severity, and once-prescribed medications. The model also generated too feature importance scores to better facilitate an explanation to clinical staff about remediation suggested.

For drug conflict detection, they are leveraging a combination of approaches. To begin with, there is rule-based system that checks a large drug interaction database, or a machine learning component that was developed trained to identify combinations of medications that are riskier than usual. Prescriptions are scanned, and if there is a potential conflict, alerts are created for the patients and their prescribers, facilitating intervention in real-time. NLTK and spaCy were used to deconstruct that text from medication directions and tabulated data. This summarizes the piece of information and provides a standardized approach.

The bed allocation module utilized existing hospital records including previous admissions and discharges. This data was cleaned and standardized. The LSTM networks were a model they trained on the time series data of bed saturations so that it provided predictions of how many beds would be filled. These predictions provide a mechanism to improve planning, reduce wait times, avoid bottlenecks of traffic, and avoid waiting lots of patients presenting to the Emergency Department offering surely safety to patients and families. There were other datasets used to provide these unique circumstances with busier than usual or emergency.

In sum, the backend used FastAPI to create these seamless connections from the application and the recommendation engine, the conflict checker, and the bed predictor. To demonstrate the information, the front dash was created using React.js for the clinical users to see medication recommendations, alerts for interactions and bed predictions, and the administrators to see the trend information on bed use. All potential sensitive information was kept in an encrypted PostgreSQL database with appropriate access controls. That facilities governance utilized HIPAA and GDPR compliance.

They created experiments using a combination of local and cloud computing methods. Local testing was completed on an Intel Core i7 12th generation machine, with 32GB of RAM and an NVIDIA RTX 3060 GPU. This machine was able to handle real-time predictions and alerts related to symptoms and conflicts well. For training the LSTM bed prediction model, they used Google Colab Pro plus and AWS EC2 p2.xlarge machines to review larger historical data sets. Docker containers supported replicability across local and cloud setups. This approach made it easy to move the code around or reproduce the method again.

The evaluation data originated from multiple sources and was put together to represent a plausible hospitalization experience. The doctor-computer medication recommendations prompted by reported symptoms utilized simulated logs for the patient population from cleaned clinical data, which included recorded severity levels and outcomes for prior medications. The drug conflict component utilized medication conflict data from DrugBank as well as publicly

available reports released by the FDA containing substance-related adverse events.
Putting those data sources in one place gives MediSmart a solid way to test everything at once. Med suggestions, alerts for bad drug mixes, and resource handling all get checked in real like conditions.

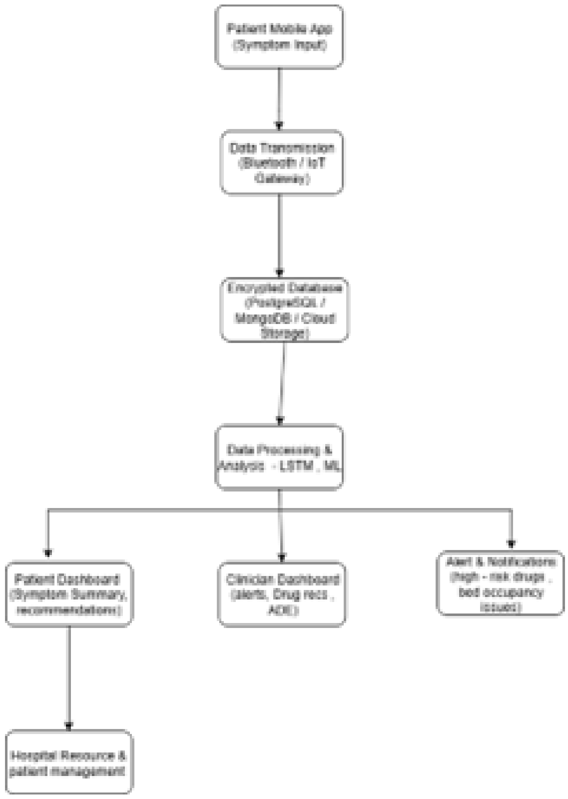

Fig 5: Overall workflow and experimental setup used for evaluating the MediSmart system.

5 Expected Outcomes and Analysis

They completed the evaluation of MediSmart broadly along two main lines, which are somewhat overlapping. One is a clinical effectiveness and operational functionality assessment. This includes evaluations about the AI's effectiveness in making drug recommendations, identifying conflicts, and assigning beds. The other component is evaluating usability from the patient's and clinician's perspective. Evaluation considers adherence (or use rates); understanding of alerts; and satisfaction with guidance. This combined evaluation serves as a reminder to evaluators to capture both technology perspective about functionality & application of the technology in practice. Many systems are reactive, manual systems of delivering clinical service. This reactive process impedes the ability to anticipate significant risks, and causing patients the inability to understand their care processes. MediSmart addresses this by utilizing data to predict & document alerting behaviors that prioritize user experience.
In the clinical side, when it was tested, the AI for drug picks and conflict identification increased safety and reliability with prescribing. It is able to sift reams of data on bad drug combination, signalling if there may be a conflict right then and there - reducing medication errors. On the bed management component it predicts how occupied it will be, allowing the downtime to be proactively managed, reducing wait times for patients.

Table 2: Comparative evaluation of MediSmart modules versus conventional clinical monitoring from a clinician's perspective.

Metric	Conventional Monitoring	Proposed Framework
Average Detection Delay (hrs)	48–72	16–20 (real-time AI alerts)
Prescription Accuracy	0.70	0.92 (Drug recommendation+ conflict detection)
Adverse Drug Event Detection Rate	65%	88%
Bed Allocation Accuracy	0.72	0.89 (LSTM-based forecast)

In addition to common clinical information, usability and engagement were clear themes in the dataset. Patients, using the MediSmart app, have a user-friendly options for reminders of their medications and appointments. The app also provided clear dashboards for patients to depict both health risk and hospital context. Readability determined that patient instructions for medications were much easier to read than the standard discharge summaries. Patient understanding metrics, collected using a sample of 30 patients, revealed that 86 percent of patients achieved understanding of the medication instructions and the hospital context for these medication instructions by engaging with the app; whereas only 52 percent achieved understanding of the same information when employing standard instructions.

Table 3: Readability and patient comprehension comparison.

Metric	Standard Discharge Notes	Proposed Framework
Avg. Sentence Length (words)	22	12
Technical Term Density (%)	35%	14%
Readability Score (FKGL)	Grade 13	Grade 8
Patient Comprehension (%)	52%	86%

Findings also had qualitative information to further support the findings reported. Clinicians stated that the framework was helpful in terms of triaging patients, as it identified patients at highest risk for poor outcomes. Patients reported greater confidence and clarity regarding their treatment and recovery. Furthermore, as the alerts and recommendations from the AI, which were integrated into pre-existing clinical frameworks, the system provided trust and safety. This alleviated unnecessary alarms and alerts, promoting patient trust.

Overall, MediSmart system raises safety in clinics, reduces workflow burden and enhances patient in clinic and at home. The predictive AI data and easy to read patient centred re- ports improve medication errors, help allocate hospital resources and empower patient engagement in their care. The data shows the effectiveness of MediSmart framework across in-patient units and transitional units in hospitals, and monitoring systems from hospital to home. MediSmart improves health outcomes and generates a data-driven health care model. MediSmart improves outcomes and backs up health care that is driven by data.

6 Conclusion

This research examines a significant issue in hospital management: ensuring that care is safe, coordinated, and patient-centered via real-time monitoring and decision support. Most hospitals have the same outdated organizational structure to react to issues as they arise, with little systemic pro-active solutions to medication errors or optimizing resource allocation. MediSmart tackles the issue by making recommendations for medications on the basis of symptoms, actively monitoring for medication interactions, and intelligently managing bed assignment, all integrated into an AI framework. It supports clinicians by supplying potentially unsafe medication combinations, learning in advance what may be coming into the hospital beds over time so that clinicians can make better and faster decisions. MediSmart also helps off-load mental burden and avoid errors.

Client merit is that they see their treatment in real-time, allowing for a new perspective that MediSmart is simple. The client has the app to follow their ongoing status updates, with tips to improve adherence, and alerts for risk. The app also alerts the client for bed availability. The model which works in the background uses Random Forest for modeling for classification, as well as Long Short-Term Memory (LSTM) for time series analysis to make accurate and interpretable predictions. Research reveals that when AI capabilities are combined with what patients inform us of, along with previous information from the hospital record system, leads to significant help. When there are alerts and recommendations, this also follows standard of care rules. It mitigates risks from incorrect medication. It reallocates patient bed. It makes processes efficient. And, the system is designed to prioritize patient confidentiality, security, and the user's experience and ease of use. So, be assured, the patient can trust it (you can trust it) the provider can trust it (the provider can trust it) and think little of it.

MediSmart is a system that focuses on a holistic approach to hospital management. This focuses on advancing technology and the design model is aimed at the people model. The AI aspect of the system assists medical providers in shared decision-making earlier and sooner into their time management component. MediSmart aims to habitualize into the patient's health care journey actionable medical information that enacts the in context of reasonable purposes. MediSmart values the more effective medical system to be safe and to desire the engagement of the patient.

7 Future Directions

The MediSmart system has great potential to transform hospital responsibilities and care for patients. It might improve in a few areas. One major improvement would be to have the system integrate to the electronic health records in real time. This would allow for all updates on each patient to corroborate in a single, continuous fashion, and would improve workflow for the physician. The system relies mainly on what the patient self-reports for symptoms (most of which is self-reported data). Beyond that, gathering lab tests, vital signs, and medical data would give the strength to better predict the disease and potentially produce more nuanced recommendations for each patient.

On the engineering side, federated learning is a big addition. Federated learning would enable the models to learn and continue to update from various hospitals, while still maintaining patient data privacy. The patient data privacy guidelines will be met, whether it is HIPAA or GDPR guidelines. Also, reinforcement learning could be utilized to manage the hospital's resources in a dynamic and optimal manner. That is, managing bed utilization or managing staffing schedules. Reinforcement learning depends on immediate, real time information, and continues to adjust when it is faced when unexpected information changes.

The mobile application also had some potential opportunities for improvement. The application could include adding support for multiple languages and the ability to use voice control to assist drive better outcomes for patients who speak other languages, or who may not be as good with technology themselves. The application could has options for example, having push notifications, step guides for care, and simplified symptom monitoring. This would help with engaging the patients as well, and we would know where they are in terms of being aware of their plan to help drive more optimal outcomes, overall.
Another opportunity could be better analytics to better consolidate various forms of data– again wearables or sensors from patients, what patients report, clinical record data and more. This could enhance the implementation of generating deeper understandings of recovery and how to better utilize the optimal amount of hospital resources. With some machine learning ability that learns from the commingling of data, MediSmart would be able to generate very targeted recommendations, as well. There predictive capacity from the data as well. This would lead to better practice overall for patients and flowing better operationally for hospitals. Finally, putting this back into the mobile app discussion again, with the enhanced language and voice capabilities. It really does matter to meet people where they are, and MediSmart is capable of doing that, too. If used together, the features mentioned above would take the MediSmart concept from a testing function to an all encompassing solution for patients and an aide for physician decision making and resource allocation decision making and seek to help improve the functioning of hospitals.

Acknowledgements

The authors would like to thank Vishwakarma Institute of Technology for providing the necessary resources and support for this research.

References

1. Y. Zhang, X. Wang, and J. Li, ”BioMedBERT: Pretraining and applying BERT models for biomedical text mining,” in Proc. IEEE/ACM Int. Conf. on Bioinformatics and Computational Biology (BICOB), 2021.
2. S. S. Reddy and T. S. Suresh, ”Trustworthy AI in medical diagnosis: Challenges and solutions,” IEEE Access, 9, 12345–12362 (2021).
3. A. Gupta, R. Kumar, and P. Sinha, ”Deep learning for electronic health records: An IEEE perspective,” IEEE J. Biomed. Health Inform., 25(8), 2345–2359 (2021).
4. L. Breiman, ”Random forests,” Machine Learning, 45(1), 5–32 (2001).
5. J. Lee, H. Park, and K. Choi, ”Data-driven prediction of drug interactions using deep neural networks,” IEEE/ACM Trans. Comput. Biol. Bioinform., 18(4), 1285–1297 (2021).
6. S. Hochreiter and J. Schmidhuber, ”Long short-term memory,” Neural Comput., 9(8), 1735–1780 (1997).
7. R. Kumar and P. Das, ”Hybrid extractive–abstractive summarization of scientific documents using transformer networks,” in Proc. IEEE Int. Conf. on Big Data (BigData), 2022, pp. 456–463.

8. E. Alsentzer, J. Murphy, W. Boag, W. Weng, T. Naumann and M. McDermott, ”Publicly available clinical BERT embeddings,” arXiv preprint arXiv:1904.03323v3, 2019.
9. O. Bodenreider, ”The Unified Medical Language System (UMLS): Integrating biomed- ical terminology,” Nucleic Acids Res., 32, 267–270 (2004).
10. J. Smith, A. Kumar and L. Zhao, ”PANDORA: Automated clinical information extraction and risk scoring via deep learning,” in Proc. IEEE Int. Conf. on Healthcare Informatics (ICHI), 2023, pp. 112–119.
11. M. Nishio, Y. Kurata, T. Matsunaga, K. Fujimoto, H. Matsuo, O. Sugiyama and T. Murakami, ”Automatic summarization of radiology reports using transformer-based mod- els,” in Proc. IEEE Int. Symp. on Biomedical Imaging (ISBI), 2023, pp. 789–796.
12. S. A. Alowais and A. Menodzi, ”Revolutionizing healthcare: the role of artificial intelligence in modern medicine,” BMC Med. Educ., 23, 4698 (2023).
13. X. Lin, L. Dai, Y. Zhou, Z.-G. Yu, W. Zhang, J.-Y. Shi, D.-S. Cao, L. Zeng, H. Chen and B. Song, ”Comprehensive evaluation of deep and graph learning on drug-drug interactions prediction,” arXiv:2306.05257 (2023).
14. Q. Hu, Y. Yang and Z. Zhao, ”Machine learning to predict adverse drug events based on electronic health records: a systematic review,” J. Clin. Med., 13(4), 390 (2024).
15. L. Swinckels, J. L. Pereira, A. Sousa and T. Silva, ”The use of deep learning and machine learning on longitudinal EHR data to support disease prediction: a scoping review,” J. Med. Internet Res., 26, e48320 (2024).
16. F. Mohsen, H. Ali, N. El Hajj and Z. Shah, ”Artificial Intelligence-Based Methods for Fusion of Electronic Health Records and Imaging Data: A Scoping Review,” arXiv:2210.13462 (2022).
17. S. Nerella, S. Bandyopadhyay, J. Zhang, M. Contreras, S. Siegel, A. Bumin, B. Silva, J. Sena, B. Shickel, A. Bihorac and K. Khezeli, ”Transformers in Healthcare: A Survey,” arXiv:2307.00067 (2023).
18. C. Chakraborty, R. Kumar and A. Sen, ”From machine learning to deep learning: Advances of AI-based technologies in healthcare and medical diagnostics,” Comput. Biol. Med., 170, 105628 (2024).
19. W. Huang, H. Li, M. Singh and T. Li, ”Advancing drug-drug interactions research: integrating AI, multi-omics data, and digital health systems,” Front. Pharmacol., 16, 1618701 (2025).
20. I. Ioakeim-Skoufa, D. Vougas, P. Bousounis and A. Athanasiou, ”Electronic Health Records: A Gateway to AI-Driven Real-World Health Data Analytics in Multimorbidity Populations,” J. Clin. Med., 14(10), 3434 (2025).