IJEAST

INTERNATIONAL JOURNAL
OF ENGINEERING APPLIED SCIENCE
AND TECHNOLOGY

VOLUME : 11 ISSUE : 01 Print / Issue Publication Date: 02-Jun-2026

ISSN : 2455-2143

DOI : 10.33564/IJEAST.2026.v11i01.020

International Journal of Engineering Applied Sciences and Technology, 2026
Vol. 11, Issue 01, ISSN No. 2455-2143, Pages 140-146
Published Online May 2026 in IJEAST (http://www.ijeast.com)

HEALTHCARE INSIGHTS FROM HOSPITAL PATIENT RECORDS USING TABLEAU

P. Abhinav Rao, Pargi Rohan Singh, K. Rahul
Student, Department of Artificial Intelligence and Data Science,
Methodist College of Engineering and Technology,
Hyderabad, India

Sana Mateen
Assistant Professor, Department of Computer Science and Engineering,
Methodist College of Engineering and Technology,
Hyderabad, India

Abstract: **This dataset brings together hospital management data from multiple branches patient demographics, doctor specializations, appointment schedules, treatment types, billing amounts, payment methods, and insurance details all in one place rather than spread across separate systems. The focus is on what actually shapes hospital performance day to day: no-show rates, treatment costs, billing recovery, and payment patterns. Small things, but they add up fast. Hospitals that track these closely tend to catch problems early, a billing gap here, a scheduling pattern there, before they turn into something harder to fix. The dataset is designed for administrators, managers, and clinical staff who want to dig into the why behind the numbers, not just read a summary someone else put together.**

Keywords- **Healthcare, Hospital revenue, Administrator, Insurance, Appointments schedules.**

I. INTRODUCTION

Healthcare is one of the highest producers of data but not all of it is useful when making operational and clinical decisions. In this project, we take a look at a Hospital Management dataset and break it down into five tables relating to Patients, Doctors, Appointments, Treatments and Billing. Each one sheds light on what happens in a hospital on a day-to-day basis. The tables which we collected cover almost all aspects of a hospital activity. They include: patient data, patient insurance data, doctor data, type, branch and availability of doctor, doctor appointment information, status of doctor appointment type of treatment (MRI scan, biopsy, chemotherapy, and so forth), billing information and details of any dues on the part of the patient.

The project seeks to apply data visualization in Python and Tableau to extract business perception from the data. Specifically, we aim to analyze the trends in appointment no-shows and cancellations, identify treatment categories with greater revenues, analyze doctor workloads across different branches of the hospital, and identify incomplete payments, which affect revenues. By transforming raw data into interactive dashboards and charts, this project seeks to close the divide between raw hospital data and informed decision-making, eventually helping in better resource planning, patient care, and financial management.

II. LITERATURE SURVEY

Data visualization is a essential part in running hospitals. The data hospitals generate daily, from patient records to billing to scheduling, is too dense to read in raw form. Tableau addresses that by converting it into charts, dashboards, and reports that staff can actually use. Its adoption in healthcare analytics is driven by three practical factors: speed, ease of use, and reliability. A billing manager can pull up payment trends without writing a query. A department head can track bed occupancy across shifts without calling IT. That kind of access, without the technical overhead, is what makes it useful in a clinical environment where time is short and decisions carry real consequences.[1]

Chang analyzed hospital patient records using Tableau, by focusing on the collectives like gender and clinical metrics like length of stay. The study found that clinical staff could explore data patterns themselves, no programming knowledge needed. Tableau's drag-and-drop interface is simple enough that data-driven decisions don't have to sit with IT or analysts alone [2].

Dewar, Head of Information at St. George's Healthcare NHS Trust in London, said Tableau gave the hospital a way to work with operational data that had previously sat in separate systems with no easy way to combine it. The Arrivals dashboard tracked patient arrivals by date and specialty. Decision-makers could read demand patterns in about 30 minutes. Before Tableau, that same task couldn't be done within a single working day [3].

Massachusetts General Hospital reduced hospital-acquired infections by 85% after using Tableau. The Cleveland Clinic applied it differently, identifying patients most likely to end up in the emergency room unnecessarily, then contacting them before that happened. In both cases, the analysis pulled from patient records, appointment histories, and billing data to produce results that were clinical and financial at the same time [4].

Ala reviewed appointment scheduling problems across healthcare systems and found that a well-designed system gets more out of expensive staff and facilities while cutting patient wait times. The harder part is what any scheduling system has to account for: emergency department flow, outpatient capacity, and how resources get allocated across a hospital. Those aren't abstract problems. They show up directly in hospital management datasets covering appointments, doctors, patients, billing, and treatment records [5].

Niu looked at 1,144 studies on appointment scheduling optimization, all published between 2016 and 2023. And if you just look at the numbers, the growth is hard to ignore. Back in 2016,only few articles were published on this topic. But by 2022, that number had shot up to 224.. Scheduling isn't glamorous, but it's a real problem that affects how hospitals, clinics, and businesses actually function day to day.[6]

Kumar analyzed the doctor ratings based on patient waiting time using a Random Forest Regressor as part of a broader predictive framework for outpatient scheduling. The finding was practical: better scheduling gets more out of hospital time and space, which affects both patient satisfaction and revenue. Those are measurable outcomes, and billing and appointment datasets in Tableau give you a direct way to track them [7].

Tableau isn't just for crunching clinical numbers in healthcare. Hospitals actually use it to pull together appointment scheduling, billing, and financial records all in one place — so administrators aren't jumping between five different systems just to get a clear picture. It connects the patient care side with the financial side, which is something a lot of hospitals have struggled with for years.[8]

A study shows that private medical practice and a university hospital found that 22.8% of appointments were booked online after the scheduling system was introduced. No-show rates dropped, resource use improved, and the structured appointment data made it possible to compare booking patterns directly between the two settings. That kind of comparison is difficult without clean, consistently structured data [9].

Billing and claims data tell payers what care patterns actually look like and give doctors something concrete to base decisions on. Payment monitoring on the provider side catches duplicate charges for hospital services before they go through. That whole process depends on structured billing datasets that record treatment costs and payment status for each visit [10].

III. MATERIALS & METHODS

Dataset: Hospital Management System

Link:https://www.kaggle.com/datasets/kanakbaghel/hospital-management-dataset

The Hospital Management dataset on Kaggle has real hospital operational data across five CSV files: patient demographics, appointments, doctor records, billing, and treatment history. It's the kind of dataset that's actually useful for Tableau work because the files connect to each other. You can link a patient's demographics to their treatment, compare billing amounts by doctor, or track which departments had appointment spikes in Q3.

Most people use it to answer operational questions. Which doctors are overloaded? Where is billing leaking? Are certain treatment types correlated with longer stays or higher costs? The dataset doesn't answer those on its own, but it gives you enough to build dashboards that do. The five-file structure also makes it a decent test for anyone learning relational data modeling in a BI tool.

This dataset provides comprehensive information about the operational activities of a hospital, organized across five interrelated CSV files. The primary goal is to analyze appointment patterns, billing performance, doctor efficiency, and patient demographics to derive operational insights and support data-driven decision-making in hospital management.

Table-1: Hospital Management Dataset Structure

File Name	**Key Fields**	**Role**
patients.csv	Patient_ID, Name, Age, Gender, DOB, Contact	Patient dimension
appointments.csv	Appointment_ID, Patient_ID, Doctor_ID, Date, Status	Central fact table (hub)
billing.csv	Billing_ID, Appointment_ID, Amount, Payment_Status	Billing dimension
doctors.csv	Doctor_ID, Name, Specialization, Department	Doctor dimension
treatments.csv	Treatment_ID, Appointment_ID, Treatment_Type, Description	Treatment dimension

Target Variable: The target variable is CNT (Appt_ID) , a count of confirmed appointments per doctor, time period, or patient. It's the core activity metric: how many appointments happened, and where.

On its own, the number is limited. Paired with billing amounts and treatment types, it tells you more. a doctor seeing 80 patients and billing $200 each is running a completely different practice than one who sees 20 patients but charges $1,500 a pop. Same hospital, totally different reality. That kind of gap doesn't just show up on a spreadsheet and disappear. It ripples into how many staff you need, how

rooms get scheduled throughout the day, and where the bulk of the hospital's money is actually coming from.
Appointment counts over time also surface problems before they get bad. If confirmed appointments in one department drop 15% over six weeks, something changed, whether that's a doctor leaving, a scheduling issue, or a referral problem. Hospital administrators rarely catch that kind of shift without a dashboard watching the numbers.

Stakeholders: Everyone's looking at the same data, but nobody's asking the same question. Administrators want to know if the workload actually makes sense across the hospital. When one department has been slammed for months and another is practically sitting idle, that shows up in appointment numbers long before anyone officially raises a flag. Billing teams are hunting for where money is getting stuck. When unpaid balances keep piling up around the same treatment type or the same patient group, that's not bad luck — something in that process is broken. Could be a wrong billing code, a messy insurance workflow, or nobody following up properly. Either way, it's not random. Department heads use appointment volume to fight for their teams. Walking into a budget meeting and saying "we're overwhelmed" is easy for people to brush off. Walking in with six months of data showing one doctor absorbed 30% more appointments than everyone else on the team — that's a different conversation entirely. No-show rates carry their own story. If patients from a certain age group or neighborhood keep missing appointments, slapping a generic reminder on it won't do much. You have to actually know who's not showing up before you can even start figuring out why. That one matters more than people give it credit for. Some treatments bring in steady volume at margins you can count on. Others eat up resources and barely cover their own costs. That's the kind of thing you really want to know before you decide to grow a department.

Software: Tableau
The software used for this project is Tableau Desktop, a powerful data visualization tool widely used in the field of Business Intelligence. Tableau enables users to transform raw data into interactive and understandable visual formats such as bar charts, line graphs, heatmaps, and dashboards. Its drag-and-drop interface is highly intuitive, making it accessible even for those with no programming experience. The software supports connections to various data sources and offers tools for data cleaning, transformation, and visualization.

IV. DATA VISUALIZATION

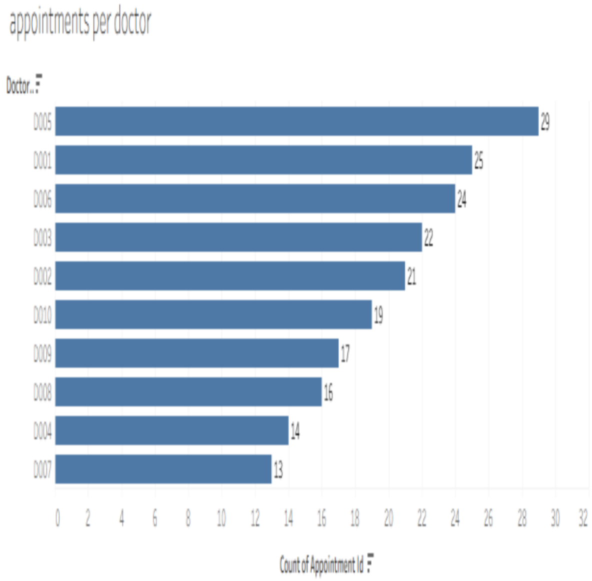

Figure-1: Appointments per doctor

This horizontal bar chart shows the D005 is carrying 29 appointments. D007 has 13. That's not a small gap — that's more than double the workload for one doctor versus another on the same team. D001 and D006 are also up there, both north of 24. So it's not like D005 is some outlier anomaly. There's a cluster of doctors doing the heavy lifting while a few others have a noticeably lighter plate. Honestly, that kind of imbalance doesn't usually happen by design. It creeps in — patients requesting specific doctors, scheduling gaps that never got fixed, someone just being easier to book. Whatever the cause, it's probably worth a conversation before the busier doctors burn out or patient wait times start.

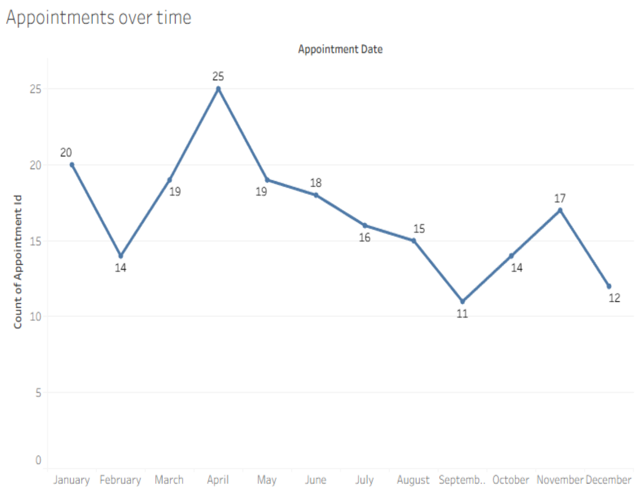

Figure 2: Appointments over time

International Journal of Engineering Applied Sciences and Technology, 2026
Vol. 11, Issue 01, ISSN No. 2455-2143, Pages 140-146
Published Online May 2026 in IJEAST (http://www.ijeast.com)

This line chart represents the monthly trend in the total number of the patient appointments throughout the year. April had 25 appointments. Nothing else came close. January was fine — 20, a decent start. Then February dropped to 14, which is odd. One month in and the numbers are already sliding. The post-April stretch is where it gets bleak. May through October just sort of drifts. In September the trend hits to 11 and stays there long enough to feel like the new normal. When November comes it bounces to 17, which looks promising until December lands at 12 and you realize it wasn't really a recovery at all. What stands out to me is how long the slow stretch runs. It's not just a summer dip — it goes May through October with nothing above moderate. If staffing or outreach is built around the April peak, the back half of the year probably feels pretty thin in comparison.

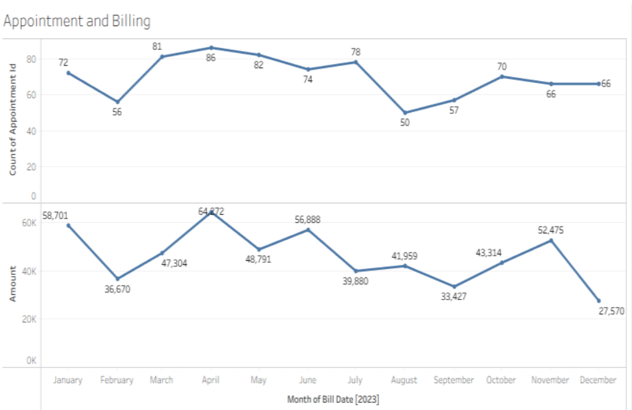

Figure-3: Appointments and billing

This dual axis graph compares the two key operational metric first one is the monthly count of the appointments and second is the total billing amount. April and March were the busiest months — 86 and 81 appointments respectively. August was the slowest at 50. That part's fairly predictable. The billing line is where it gets more interesting. April topped out at $64,272, which tracks. But December had 66 appointments — a pretty normal number — and still managed to pull in just $27,570, the lowest revenue of the year. That's not a volume problem. Something about what patients came in for that month was just cheaper, whether that's seasonal procedure types, more discounts, or a concentration of follow-ups rather than new treatments. September has a similar issue. 57 appointments, $33,427 in billing. That's a worse revenue-per-appointment ratio than most other months, and it's worth understanding why before assuming the slow months are the only problem. The takeaway isn't just "more appointments equals more money" though mostly it does. It's that two months with reasonable patient traffic still underperformed on revenue, and that gap probably deserves more attention than the slow months do.

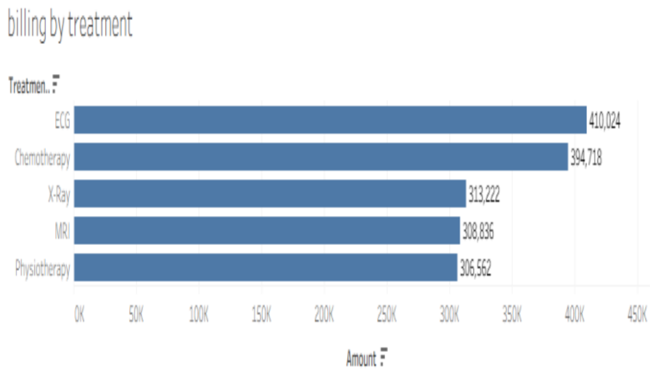

Figure-4: Billing by treatment

This bar chart shows the total billing amount generated by each treatment type offered in the hospital. ECG is the top earner at $410,024. Chemotherapy isn't far behind at $394,718. After that, X-Ray, MRI, and Physiotherapy are all clustered between $306K and $314K — which is a weirdly tight band for three treatments that have almost nothing in common procedurally. The gap between the top two and the bottom three is the real story here. ECG and Chemotherapy are pulling in roughly $100K more than anything else. That's not a small margin. Whether that's because they're priced higher, used more often, or both — it's worth knowing, because if either of those departments hits a slow stretch, the revenue impact won't be subtle.

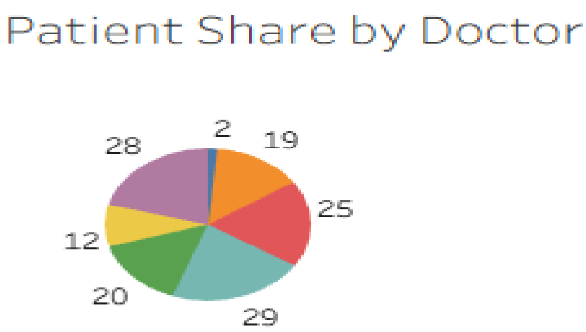

Figure-5: Patients Share by Doctor

Most doctors are somewhere in the middle — 20s, teens, a few at 12. Then there's whoever has 2 appointments. That's not a rounding error, that's basically inactive. The busiest doctor is carrying 29. The quietest has 2. On the same team, presumably under the same scheduling system. That gap is hard to explain away as just patient preference or specialty differences. A pie chart for this kind of data usually looks fairly balanced — lots of roughly equal slices. This one doesn't. Some slices are genuinely thin, and that's worth asking about before assuming the distribution is intentional.

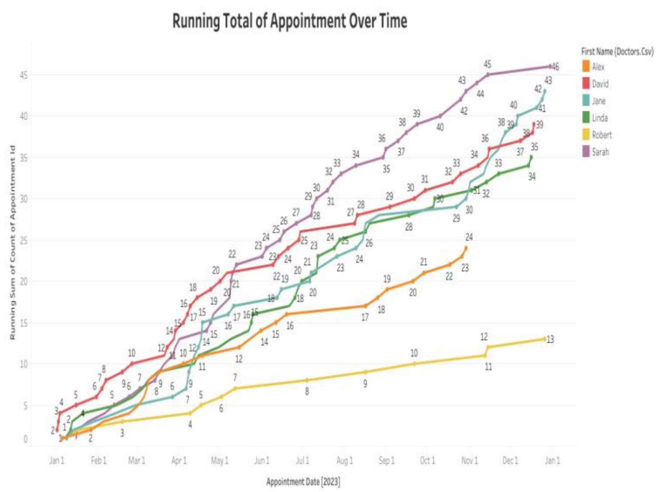

Figure-6: Running Total of Appointment Over Time

This graph shows Sarah handled the most appointments across 2023. Robert had the least. Everyone else fell somewhere in between, and all of them grew steadily through the year — no major gaps or flat stretches, just consistent patient flow. What the cumulative view adds is the shape of the gap. It doesn't start small and widen dramatically. The spread between Sarah and Robert was there early and mostly stayed that way. Through this it says that the imbalance isn't something that it built up over time — it was probably tells that how appointments were being assigned from the start.

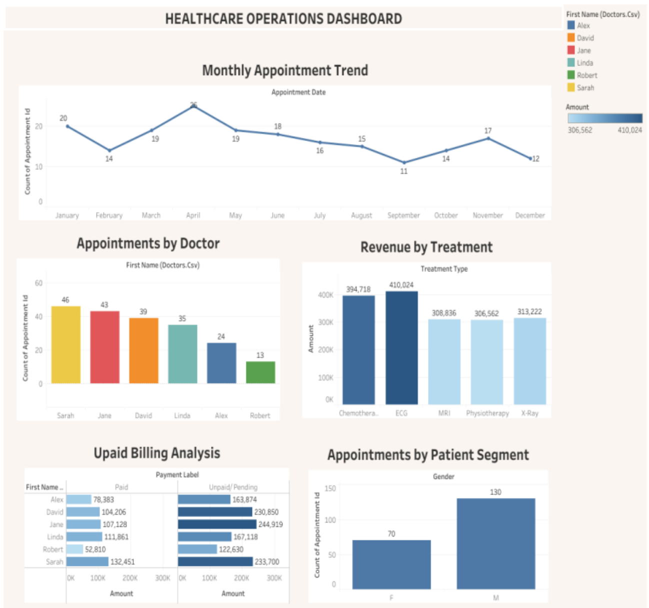

Figure-7: Health Operations Dashboard

The "Healthcare Operations Dashboard" provides a comprehensive overview of hospital operations as the April was the peak appointments were up, revenue was up, things were moving. By September, both had drifted to their lowest points. That seasonal rhythm shows up clearly when you look at the full year. Sarah carried the heaviest patient load. Robert had the lightest. The gap between them wasn't small, and it wasn't explained by specialty alone — which makes the scheduling logic worth a second look. ECG and Chemotherapy led on revenue, not by a little. MRI, Physiotherapy, and X-Ray were all clustered closer together, steady but unremarkable. The top two treatments are doing a lot of the financial heavy lifting. The unpaid billing is probably the part that deserves the most immediate attention. It's spread across multiple doctors, which means it's not one person's problem it's a process problem. That kind of thing doesn't fix itself. Male patients made up the larger share of appointments. No context is given for why, which makes it hard to know whether that's expected for this hospital's catchment area or something worth digging into. Taken together, the dashboard doesn't reveal anything shocking. What it does is put the revenue gaps, workload gaps, and billing gaps side by side and that proximity makes them harder to look past than they would be in five separate reports.

V. RESULT AND DISCUSSION

From the Hospital Management dataset it shows clear patterns across appointments, billing, doctor performance, and patient distribution. Monthly appointment volumes are uneven. April had the highest count at 25; September fell to just 11. That's a consistent first-half surge that the second half doesn't match. Seasonal illness, elective treatment timing, and school or work schedules all probably contribute. The practical problem is that fixed annual staffing plans can't account for this. Scheduling needs to track actual monthly demand, not assumptions about it. The 2023 dual axis chart tells a similar story. April hit 86 appointments and $64,272 in billing. December managed just $27,570 even with 66 appointments on the books. September billed $33,427 across 57 visits.

Volume doesn't equal revenue. Some months clearly run heavy on low-cost treatments, discounted services, or bills that never got collected. Before reading a busy month as a good month, it's worth checking what those appointments actually brought in, and whether billing recovery is leaving money on the table.

The revenue of ECG and Chemotherapy is $410,024 and $394,718 respectively. Physiotherapy and MRI came in lower but contributed consistently. The bottom three treatments are close enough in revenue that demand across diagnostic and rehab services looks stable rather than skewed toward any one procedure. ECG's gap over everything else is significant. Equipment downtime or staffing gaps there would hurt billing faster than anywhere else in the mix.

D005 handled 29 appointments. D007 handled 13. That gap isn't small. Doctors at the top of that range carry real burnout risk. Doctors at the bottom may simply be getting fewer referrals because of how appointments get assigned, not because patients don't need them. Either way, the distribution isn't random and it's not self-correcting. Fixing it means looking at how appointments get routed in the first place.

One doctor handled 29 patients. Another handled 2. The pie chart makes that gap hard to ignore. An overburdened doctor is a care quality risk. An underused one is a revenue problem. Neither fixes itself. The booking system needs to route based on actual capacity and specialty fit, not just whoever's next in the queue.

Spreadsheet rows don't show you that D005 is carrying twice the load of D007, or that April billings outpace December by more than 2x despite similar appointment counts. The charts do. Putting patient flow, billing, and staff workload on one dashboard means the connections between them are actually visible. That's harder to achieve when the same data lives across separate reports.

VI. CONCLUSION

Five tables related to hospital management dataset were used in tableau. The analysis ran across six modules and covered patient flow, billing, staff workload, and treatment revenue. April had the most appointments. September had the fewest. Billing mostly followed that pattern, but December and September brought in less than their appointment counts would suggest. That probably comes down to billing inefficiencies or more low-cost treatments in those months. Doctor workload is skewed. Some doctors are carrying significantly more than others, and that's a problem for both care quality and staff retention. ECG and Chemotherapy generated the most revenue by a visible margin. The other treatments were consistent but well behind those two.

Building the data model around appointments as the central fact table worked well. Linking patients, doctors, billing, and treatments to it kept the relationships clean without the data getting messy or duplicated which sounds basic, but it's the kind of thing that quietly breaks dashboards when it's not done right. The calculated fields helped too. Grouping billing into categories, flagging appointment outcomes, tagging treatments by cost range, these aren't flashy additions, but they're what made the analysis actually usable rather than just technically correct. Real-time data would change how the dashboard gets used day-to-day less retrospective review, more live operational awareness. Predicting no-shows before they happen is a genuinely hard problem and probably the highest-value thing to tackle, since missed appointments hit both revenue and patient care at the same time. Anomaly detection on billing is worth exploring too, though that one needs clean, consistent data to work well — which is its own challenge in most hospital environments.

VII. REFERENCES

[1]. Sira Consulting Inc. (2025). How Tableau is Transforming Patient Data into Actionable Insights in Healthcare, Sira Consulting. Available at: Sira Consulting Inc.

[2]. Chang H. (2017). Interactive Visualization of Healthcare Data Using Tableau, Healthcare Informatics Research, Vol. 23, No. 4, (pp. 349–354). doi:10.4258/hir.2017.23.4.349

[3]. Landi A. (2018). How Software Maker Tableau Tames Data for Doctors, Fortune Magazine. Available at: Fortune Article

[4]. Ala M., Alam U.F., Sasangohar S. et al. (2022). Appointment Scheduling Problem in Complexity Systems of the Healthcare Services: A Comprehensive Review, Journal of Healthcare Engineering, Article ID 5819813. doi:10.1155/2022/5819813

[5]. Niu T., Lei B., Guo L., Fang S., Li Q., Gao B., Yang L., and Gao K. (2023). A Review of Optimization Studies for System Appointment Scheduling, Axioms, Vol. 13, No. 1, p.16. Available at: MDPI Axioms Article

[6]. Kumar S. (2020). Optimizing Appointment Scheduling for Outpatients and Income Analysis for Hospitals Using Big Data Predictive Analytics, Journal of Ambient Intelligence and Humanized Computing, Springer. doi:10.1007/s12652-020-02118-4

[7]. Schmi A. (2025). Efficient Patient Care in the Digital Age: Impact of Online Appointment Scheduling in a Medical Practice and a University Hospital on the No-Show Rate, Frontiers in Digital Health. doi:10.3389/fdgth.2025.1567397

[8]. Data Science Dojo (2025). Data Exploration: Healthcare Data Visualization with Tableau, Data Science Dojo Blog. Available at: Data Science Dojo Blog

[9]. Fang Z. (2023). Research on Equipment Management System of Smart Hospital Based on Data Visualization, in Advanced Hybrid Information Processing (ADHIP 2022), Lecture Notes of the Institute for Computer Sciences, Vol. 468, Springer, Cham. doi:10.1007/978-3-031-28787-9_29

[10]. Zigler J. (2024). Visualization Analysis in Healthcare: A Review of Techniques and Applications, The Visual Computer, Springer Nature. doi:10.1007/s00371-024-03586-x

[11]. Rabiei B., and Almasi S. (2022). Requirements and Challenges of Hospital Dashboards: A Systematic Literature Review, BMC Medical Informatics and Decision Making, Vol. 22, No. 1, p.287. doi:10.1186/s12911-022-02032-x

[12]. Helminski D. (2024). Development, Implementation, and Evaluation Methods for Dashboards in Health Care: Scoping Review, JMIR Medical Informatics, Vol. 12, p.e59828. doi:10.2196/59828

[13]. Suarez A. (2022). Machine Learning Approaches to Predicting No-Shows in Pediatric Medical Ap-

pointments, npj Digital Medicine. doi:10.1038/s41746-022-00594-w

[14]. Moreira R. (2023). A Robust Deep Learning Technique for No-Show Prediction in Hospital Appointments, in Proceedings of the 9th International Conference on Advanced Intelligent Systems and Informatics (AISI 2023), Lecture Notes on Data Engineering and Communications Technologies, Vol. 184, Springer, Cham. doi:10.1007/978-3-031-43247-7_1

[15]. Bachtiar N.A. (2023). Machine Learning Predictions on Outpatient No-Show Appointments in a Malaysia Major Tertiary Hospital, BMC Health Services Research / PMC. Available at: PMC Article on No-Show Prediction

[16]. Office of the National Coordinator for Health Information Technology (ONC) (2025). Hospital Trends in the Use, Evaluation, and Governance of Predictive AI, 2023–2024, U.S. Department of Health and Human Services. Available at: ONC Predictive AI Report

[17]. Grossi C. (2022). How Can Big Data Analytics Be Used for Healthcare Organization Management? Literary Framework and Future Research from a Systematic Review, BMC Health Services Research / PMC. Available at: PMC Big Data Analytics Article

[18]. Ramachandran A. (2024). The Evolution of Automated Medical Billing with Artificial Intelligence: A Review with a Global and Saudi Perspective, PMC. Available at: PMC Medical Billing AI Article

IJEAST

INTERNATIONAL JOURNAL OF ENGINEERING APPLIED SCIENCE AND TECHNOLOGY

ABOUT IJEAST

International Journal of Engineering Applied Science and Technology (IJEAST) is a peer-reviewed, open access journal that publishes high-quality research papers in the field of Engineering, Applied Science and Technology.

IJEAST aims to provide a platform for researchers, academicians, and professionals to share their innovative ideas, research findings, and practical experiences with the global scientific community.

FOCUS AREAS

- Engineering
- Applied Science
- Technology
- Innovation & Development
- Interdisciplinary Studies

PEER REVIEWED
All submissions are rigorously peer reviewed to ensure quality.

OPEN ACCESS
Free and unrestricted access to research for all.

GLOBAL REACH
Connecting researchers and professionals worldwide.

TIMELY PUBLICATION
We ensure a swift and efficient publication process.

For more information, visit our website
www.ijeast.com

INTERNATIONAL JOURNAL
OF ENGINEERING APPLIED SCIENCE AND TECHNOLOGY

editor@ijeast.com
www.ijeast.com
India

2455-2143