

PROJECT TYPE

AI platform

TECHNOLOGIES

Python (Django), FastAPI, React, TypeScript, PostgreSQL, Redis, Scikit-learn, TensorFlow, HL7v2, FHIR APIs, AWS (EC2, RDS, S3, Lambda), Docker, Kubernetes, Terraform, GitLab CI/CD, Twilio 10DLC SMS/voice, SendGrid email, ELK Stack, Prometheus, Grafana, HIPAA-compliant ePHI encryption (AES-256), role-based access control (RBAC), SSO

DURATION

16 weeks

METHODOLOGY

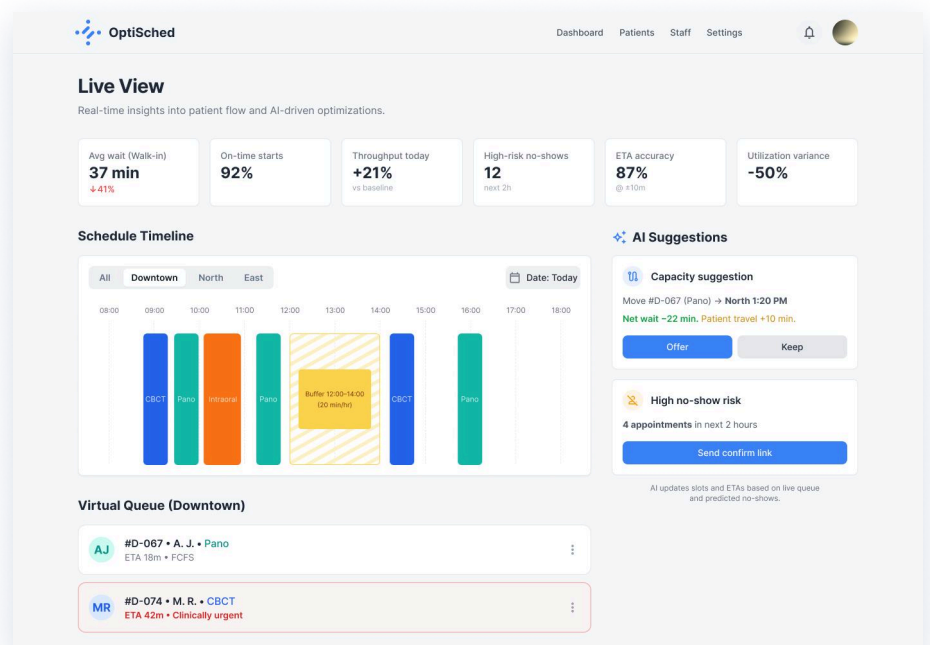
Scrum

TEAM

- 1 Project Manager
- 1 Business Analyst
- 1 Solution Architect
- 1 Frontend Developer
- 1 Backend Developer
- 1 ML Engineer
- 1 QA Engineer
- 1 UX/UI Designer
- 1 DevOps

HIPAA-compliant AI-powered patient management platform for a dental imaging provider

An AI platform to reduce patients' wait times and increase daily throughput of the Client's branches. The platform leverages predictive scheduling for managing first-come-first-served queues, inter-branch balancing, and smart reminders to reduce no-shows.



About the Client

The Client is an independent dental imaging provider with three branches across the city (downtown, north, and east). It offers cone beam computed tomography and panoramic intraoral X-ray by referrals from general dentists, orthodontists, periodontists, endodontists, and oral and maxillofacial surgery practices. The Client's issues were long waits and uneven utilization across branches, driven by unpredictable walk-ins and no-shows.

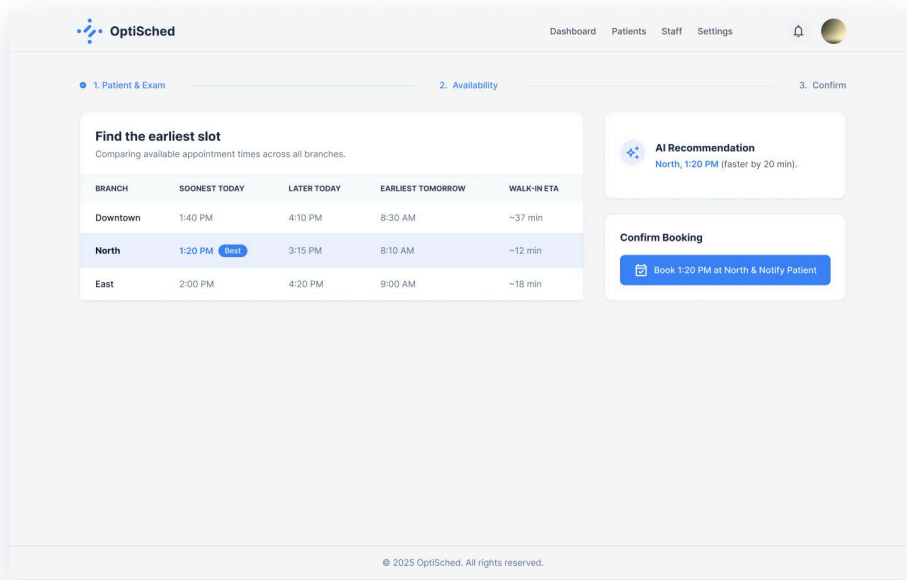
Business challenge

Deliver an AI-enabled and HIPAA-aligned platform to:

- ✓ Optimize appointments and capacity of the branches,
- ✓ Coordinate walk-in patients,
- ✓ Send omnichannel reminders to patients and provide analytics to branch administrators.

Our solution

We built an AI-powered platform that reduced patient wait times due to precise planning and managed flow. Predictive AI models identify no-shows and walk-in peaks, dynamically redistribute slots, and balance the load between three branches. Patients receive the closest available appointment slots, and administrators get a real-time status of queues. According to the HIPAA requirements, all components operate with ePHI encryption.



Optimization of appointments and capacity management

The basis is a system for predicting no-shows and smart slot planning. The AI model calculates the probability of a no-show for each appointment, taking into account the patient's behavioral history, such as confirmations or transfers, the type of examination, time of day, day of the week, and weather factors.

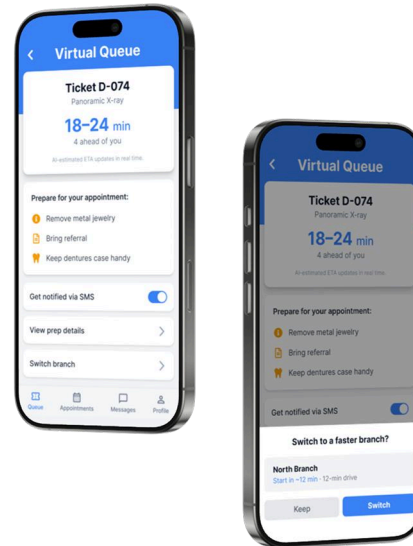
On this basis, the schedule is rebuilt: slots freed up due to high-risk appointments are offered in advance to the waiting list and pre-confirmed "standby" patients. Patients can self-register with offers of alternative slots and branches with confirmation and rescheduling in just one click.

Depending on the risk profile and the proximity of the visit, the tone and frequency of notifications (via SMS, e-mail, or a voice call bot) change: from a soft reminder 72 hours before active contact and offers of jump-in slots available within the next two to twenty-four hours.

Orchestration of walk-in flow

To fairly integrate walk-in patients, we developed an AI forecasting of their influx by branches and hours, according to seasonality, days of the week, weather, and local events. Based on the forecast, the system reserves buffers in the grid for walk-ins and dynamically expands and shortens them during the day, focusing on the actual queue length, staff availability, and expected no-shows according to the schedule.

Walk-ins are processed through a virtual queue via a kiosk at the reception: the patient receives a number, estimated time, and preparation instructions (removal of jewelry, etc.). In case of overload, the system offers to move to another branch, taking into account travel time and current waiting time, maintaining the first-come-first-served principle within each queue. Clinical exceptions (trauma, acute pain) are raised in priority according to transparent rules.



Load balancing between branches

A single AI optimizer takes into account both the appointment schedule and walk-in queues. It offers walk-in patients the closest options in neighboring branches, maintaining the SLA for scheduled visits and not violating the first-come-first-served rules. To smooth out peaks, scenarios are provided: automatically opening “micro-slots” later in the day, calling a replacement technician, and redistributing short appointments to underloaded points in the network.

Communications and reducing no-shows

Communications are built around risk scoring: high-risk appointments receive earlier and more frequent alerts. Notifications are 10DLC-compatible, contain smart links for confirmation and rescheduling, and preparation instructions.

Integrations and security

The platform integrates with the scheduler and radiological information system according to HL7v2. All components are designed in accordance with HIPAA: ePHI encryption, role-based access control, single sign-on, and activity auditing. Historical data for model training is de-identified, with model drift monitoring and retraining on a scheduled basis.

Operational transparency is provided by ETA accuracy, occupancy by branches, average waiting time, no-show and late-cancel share, on-time start execution, volume of redirections between branches, and share of patients who left without waiting.

All critical actions remain under the control of employees: proposals and rearrangements are confirmed by administrators, with the patient always finalizing the choice.

Orchestration of walk-in flow

To fairly integrate walk-in patients, we developed an AI forecasting of their influx by branches and hours, according to seasonality, days of the week, weather, and local events. Based on the forecast, the system reserves buffers in the grid for walk-ins and dynamically expands and shortens them during the day, focusing on the actual queue length, staff availability, and expected no-shows according to the schedule.

Walk-ins are processed through a virtual queue via a kiosk at the reception: the patient receives a number, estimated time, and preparation instructions (removal of jewelry, etc.). In case of overload, the system offers to move to another branch, taking into account travel time and current waiting time, maintaining the first-come-first-served principle within each queue. Clinical exceptions (trauma, acute pain) are raised in priority according to transparent rules.

Customer's benefits

- ✓ Average waiting time before examination was reduced by 32-38%.
- ✓ Daily throughput increased by 18-24%.
- ✓ No-shows and late cancellations reduced by 22-30%
- ✓ ETA accuracy for walk-in $\geq 85\%$ within ± 10 minutes.
- ✓ Variability of branch occupancy reduced by 50%; the rate of patients who left without waiting reduced by 40-55%.