HEALTH CARE INDUSTRY

How AI Is Helping to Heal Patients—and Hospitals

By María López, Nacho Hafner, Enzo Lago, Luis Peral, German Adrian Garcia, Alfonso Abella, Alejandro Fernandez Piqué, and Alejandro Fernandez Alvarez

ARTICLE JANUARY 05, 2026 8 MIN READ

Expectations for AI run high in health care, especially for providers. As with many other sectors, however, results have lagged, and most health systems have little to show for their sometimes substantial investments, other than scattered impact from often siloed pilot projects.

Until now.

One major European health system is producing a winning trifecta of results from AI—improved care, expanded capacity, and reduced costs—and is on its way to achieving all three at scale. In the process, it is demonstrating how AI can reshape the health care landscape, reimagining the core of key processes end to end, blurring the lines between provider, innovator, and tech player, and opening the door for providers to redefine competitive advantage.

Longstanding Challenges…

The health system in question operates more than 200 hospitals and care centers and is a leader in clinical quality, patient experience, and operational efficiency. Nevertheless, like many other providers, it faced operational challenges around clinical care and the staff and patient experience. Also like many health systems, this provider had pursued AI pilots but with limited scale and success. (This is not unusual: recent BCG research into AI value and maturity found that health care providers rank below the cross-industry average in creating measurable value from AI.)

Meanwhile, the health system was continuing to operate in an environment where demand outpaces capacity, quality of care varies (sometimes unacceptably), and long wait times and high operating costs undermine the patient experience and provider margins. Vast amounts of data (from electronic health records, medical imaging, and contact center interactions, for example) remained siloed, limiting both clinical decision support and operational orchestration.

…Meet Today’s End-to-End Solutions

The health system’s management set out to embed AI (predictive, generative, and agentic) into its clinical and patient support operations, partnering with BCG and BCGX (our tech build and design unit) along the way. The strategy was based on accessing unstructured data at scale and translating it into standardized, clinically validated, and AI-powered workflows. The goal was to drive efficiency but, even more important, to support enhanced and more personalized care. The provider’s hope was that by automating routine tasks and guiding clinical decisions, AI would reduce the tension between scaling up care delivery and tailoring care to individual needs.

The new approach was to start with high-volume, high-friction workflows and redesign them through operating model changes codeveloped with physicians and functional experts. An initial analysis pinpointed key issues to address, all typical sources of frustration for providers:

- Manual and time-consuming processes for care pathway recommendations
- Variability in care delivery stemming from inconsistent pathway adoption by clinicians
- A stressful patient experience caused by delays in diagnosis and treatment
- Reduced diagnostic accuracy and efficiency because of resource constraints
- Inefficient and inconsistent telephone call handling that was unable to cope with rising demand
- Lack of a complete view of the patient across systems, limiting service quality

The analysis suggested three high-value arenas for AI-driven transformation. One was a core clinical domain: evidence-based care pathways to reduce unwarranted variation. The second was a vital diagnostic service: medical imaging. Third was an often overlooked function central to the patient experience: the ability of the call center to handle patient needs.

Collectively, the three programs would have a broad impact and address or fix issues affecting different stakeholders (patients, clinicians, radiologists, contact center teams), data types (electronic health records, notes, labs, images, knowledge bases), and decision cycles. In addition, since all three require shared data and AI foundations as well as an adoption-first operating model, they would help set the stage for a broader AI-driven transformation. In each case, the health system sought to build and train AI models with integration and scalability in mind, measuring impact end to end across quality, the patient experience, speed, and efficiency.

Evidence-Based Care at Scale

Matching a patient’s clinical needs with the best pathway for care is no small feat. Never mind doing so thousands or even millions of times a year for the same or similar diseases or conditions. A top priority for any health care provider is delivering a care pathway that integrates a patient’s characteristics, such as diagnostics, age, and medical history, with the best evidence-based clinical practice.

Many providers make big investments in evidence-based care capabilities, but their adoption remains limited. Manual monitoring of patient and physician data from multiple sources is one big bottleneck. Rigid care pathway “triggers” that are unable to adequately differentiate among relevant variations in seemingly similar individuals and their circumstances are another roadblock to accurate pathway assignment at scale. This lack of flexibility can lead to mistakes that reduce clinician confidence in the reliability of the pathways.

Enter GenAI, specifically, a combination of workflows and tools that ingest large amounts of patient data from multiple sources (while ensuring anonymization) and use a logic engine to match the data against validated, evidence-based clinical pathways. The logic engine proposes the optimal pathway, entry point, and next best step for each individual patient. Take the example of two patients with osteoarthritis of the knee. AI analyzes the relevant data for both and proposes either surgery or a nonsurgical pathway, with the recommended option depending on when each patient was diagnosed or some other clinically relevant factor (such as past interventions, test results, lifestyle, and age). Once the treatment pathway is validated, the workflows execute the plan automatically through the required scheduling actions and follow-ups.

The diagnosis requires validation by an expert clinician, of course, who always remains in control. But the health system's results have regularly achieved a score of 90% or higher on the F1 machine learning evaluation metric.[1]

AI-based diagnosis has so far had a significant impact for patients and the health system, which has seen an increase in clinical capacity of 10% to 20%, an improved patient and provider experience, better patient adherence to treatment plans, and more consistent, more efficient, higher-quality care. Thanks to the increased capacity, operating margins have also improved. In addition, by focusing on outcomes and reducing unwarranted clinical activity, AI has paved the way for a more constructive relationship with payers.

Smart Triage and Faster Diagnosis

The health system faced the quandary of rising demand for imaging services and limited radiologist capacity. A lack of automation across the prescription–scheduling–study–results process led to bottlenecks and long waits for diagnosis and treatment. Because it can speed up workflows, increase productivity, and boost clinical outcomes, AI helped the system break the imaging logjam. The technology helps clinicians by:

- Interpreting medical images and highlighting potential pathologies (such as a meniscus tear or bone edema) with region-of-interest overlays
- Triaging for urgency
- Expediting the diagnostics process through smart assignment of medical-imaging reviews
- Generating structured preliminary reports

In practice, the AI-augmented imaging ingests thousands of images and historical radiologist reports from electronic health records and picture-archiving and communications systems, which are labeled and used to train models for specific pathologies. Each model is powered by an algorithm that can interpret the relevant medical images. The model highlights the potential pathologies found, states its confidence in the diagnosis, and provides any other relevant findings. It also triages the case based on its urgency and generates a preliminary report for the radiologist to review and approve.

AI shows potential to significantly improve the health system's imaging function. It has already delivered faster turnaround times, better report consistency, and material cost and productivity gains. Based on an initial scope of use involving knee injuries, AI showed a potential reduction in reporting time of 20% to 40% when fully deployed. Remarkably, it has also allowed for earlier identification and specialist review of suspected serious malignancies, which could greatly improve clinical outcomes for patients.

Reimagined Patient Engagement

Like many providers, our system's contact centers faced rising inbound call volume with little automation and representatives able to achieve only a fragmented view of patients' needs and histories. These issues were increasing average call-handling time and undermining service levels.

The solution was found in a suite of voice and text GenAI agents that can handle call management end to end, including patient authentication, scheduling, appointment cancellations and changes, frequently asked questions, and an efficient handover to human agents when necessary. These virtual agents are equipped with advanced language and intent detection models, and they are integrated with patient master data and appointment APIs. A control panel and several dashboards enable human real-time monitoring of relevant KPIs, such as share of automated calls handled, average call-handling time, adherence to service-level agreement criteria, and call abandonment.

The AI capability has achieved a 20% reduction in average call-handling time, improved service levels, reduced operating costs, and greater consistency in the patient experience. The increases in efficiency and effectiveness have translated directly into operating margin improvements while freeing human agents to intervene when human knowledge or touch is required.

Lessons Learned

Our health system's experience highlights four principles that can help providers achieve the desired impact:

- **Prioritize adoption and strategic alignment over algorithms.** Engage clinicians early on, run weekly reviews of impact achieved, and maintain clear human override authority and triggers. Using review flags, audit trails, and escalation guardrails, this approach accelerates adoption while ensuring alignment on the clinical implications and risk tradeoffs.
- **Design to scale and industrialize, rather than run single pilots.** Embed AI capabilities in existing processes (such as access to electronic health records, picture archiving and communication systems, and contact centers), while leveraging shared data pipelines. Look for opportunities to reuse models, safety guardrails, and other components across use cases.
- **Institutionalize AI through a build-operate-transfer approach.** Build with cross-functional pods (teams of clinicians and experts in operations, data and machine learning, and engineering). Operate in real settings with telemetry-driven iteration and reliability strengthening. Transfer product ownership and capabilities to everyday staff via upskilling and playbooks to sustain autonomous scale.
- **Measure end-to-end value.** To understand and publicize the value achieved, use shared and standard KPIs across quality, experience, and cost (such as adherence, time to treatment, average call-handling time, and freed-up capacity).

AI is disrupting the traditional health care value chain, enabling providers to transcend their historical roles and tap new strategic opportunities. Once primarily consumers of medtech solutions, hospitals and health systems can now become technology creators, developing proprietary algorithms from clinical data that rival offerings from technology vendors. By leveraging predictive analytics and patient data, providers can design structured, personalized clinical pathways that go beyond treatment plans to include tailored prevention strategies. AI also enables delivery of end-to-end patient engagement experiences without relying on third-party call centers.

The next era of care will be led by health care institutions that own their AI agenda, not by those that wait for the ecosystem to shape it for them.

Authors

María López

Managing Director & Partner
Madrid

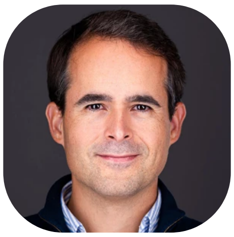

Nacho Hafner

Managing Director & Partner
Madrid

Enzo Lago

Managing Director & Partner
Madrid

Luis Peral

Project Leader
Madrid

German Adrian Garcia

Lead Data Scientist
Barcelona

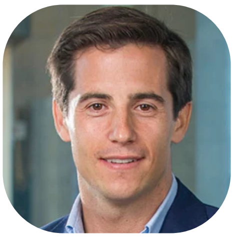

Alfonso Abella

Managing Director & Senior Partner, BCG X; Global Lead, Deep Customer Engagement AI
Madrid

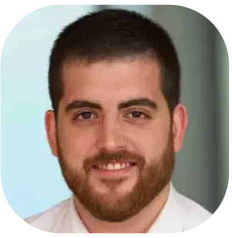

Alejandro Fernandez Piqué

Lead Data Scientist
Barcelona

Alejandro Fernandez Alvarez

Lead Data Scientist
Barcelona

ABOUT BOSTON CONSULTING GROUP

1 F1 is used to measure performance when a model needs to both identify positive cases and avoid misclassifying negative cases.