

# AI becoming “normal”: Federico Cabitza’s vision for the future of digital health

February 4, 2026

**As he begins his term as head of the Center for Digital Health and Wellbeing (DHWB), Federico Cabitza outlines a clear mission: to take artificial intelligence out of the lab and turn it into an established, everyday practice.**

**Professor Cabitza, your appointment as Director of the Center for Digital Health and Wellbeing begins on February 1. What are the key issues and priorities you intend to address in your first year?**

A digital health research center today faces increasingly demanding challenges. We live in an interconnected world—across peoples and cultures, but also between humans, animals, and the environment—with rising health expectations and longer lifespans. Yet the gap between life expectancy and years lived in good health remains wide. This issue, in my view, is still too often overlooked and should instead be central: extending healthy lifespan and preventing the final “thousand weeks” of life from becoming a long period of increasing dependence and care needs, which ultimately puts pressure on the sustainability of our universal, equitable health system. A research center like ours has all the ingredients needed to operate in this complex context: talent, expertise, experience, and a mission to produce research that can be shared with the international scientific community and transferred to the communities it serves. But this does not happen automatically. I strongly believe in teamwork and collective intelligence—a topic I have also researched in the past. My initial focus will therefore be on creating the conditions that allow the Center’s collective intelligence to fully express itself. The first step is getting to know each team member, sharing my vision, and understanding how each person sees their role and contribution. My vision for AI in healthcare can be summed up in one word: “normalization”. From the very first year, I want to move beyond innovation as something exceptional—demos, special projects, isolated pilots—and bring AI into everyday healthcare work and professional practice. My goal is to define a path toward *standardization*, so that AI becomes a “normal” technology: part of established routines and an environment that enables good work and remains truly human-centered. For this to happen, three elements are essential. First, we must cultivate a *culture of evaluation*. It is not enough to know whether a model works in theory or in the lab; we need to understand whether it is reliable, robust, and useful in real contexts, and to clearly identify its

strengths, limits, costs, and external effects—such as added complexity, changes in responsibility, or technological dependence. Second, we need strong and accountable *data governance*. In healthcare, AI succeeds or fails—often silently—based on the quality, meaning, and traceability of data and on how that data is produced and reused. The third concerns attention to the dimension of *adoption*: a phase that European regulation labels *deployment*, as if it were an autonomous, almost mechanical process. I prefer, instead, to understand it as a voluntary, conscious, and informed act of acceptance and development—much like welcoming an idea or a new member into a community—directed toward something that requires solicitude, attention, and care in order to grow: the proper use of digital tools. This requires training, usage protocols, transparent *auditing* criteria, and clear allocation of responsibilities. If these three aspects do not become a priority in our project-planning work, the risk is to produce “apparent innovation”: something that shines and gains the ephemeral attention of the news, but that does not stand the test of time and does not truly become *routine*, the normality that bears fruit over time.

**You succeed Stefano Forti. How do you intend to balance the continuity of projects already launched by the Foundation with the innovative drive that characterizes your vision?**

I am fortunate: Stefano Forti has done an outstanding job, and I inherit the results of years of careful research and development rooted in the needs of citizens and patients in the Autonomous Province of Trento. His work on digital technologies for prevention, health, and well-being is both a foundation and a strong incentive to continue building solutions that make AI’s potential available to as many people as possible. On the other hand, I believe that continuity and innovation are not mutually exclusive, if you are clear and honest about what you want to preserve. What I want to preserve is human and organizational capital: relationships with the local communities, trusted projects, internal expertise, functioning infrastructures, and hard-earned experience. Innovation, to me, means changing the unit of project design—not adding “one more algorithm” or “one more app,” but building stronger socio-technical systems that genuinely place people and their work at the center. It means designing solutions to fit into workflows over time, measuring their impact on outcomes and processes, managing exceptions and failures, and maintaining skills and knowledge. In other words, the innovation we can and want to deal with does not only concern technology but—precisely because we are creators of ideas and technologies in the broadest sense of the term—the design process and the ability to make these solutions manageable and sustainable over time. And for me, this is also a respectful way of valuing what has been built before.

**The DHWB Center is one of the three pillars of TrentinoSalute4.0. How do you intend to develop this synergy with PAT and ASUIT to ensure that digital healthcare generates tangible value for the region?**

Real value emerges when innovation follows a short, repeatable path: starting from a concrete clinical need; precisely defining use cases; co-designing with professionals; testing rigorously; measuring impacts beyond just numbers; and only then scaling and transferring solutions. With the Autonomous Province of Trento, I see a key enabling role: ensuring shared digital health infrastructure—standards, interoperability, access rules, data security, and quality—as a public asset rather than fragmented, proprietary solutions. With the University of Trento Healthcare System I see an operational role: integrating solutions into real care processes, adapting them to workflows, training users, monitoring adoption, and documenting outcomes for the scientific

community. Crucially, this chain must extend beyond hospitals to include primary care, community services, nurses, and professionals responsible for continuity of care—often working with fewer resources and less visibility. If AI works there, it becomes a structural capability of the region, not just a collection of projects.

**In a [recent article published in Agenda Digitale](#), you state that AI is not just software to be installed, but a new element in the workplace that directly impacts workers. Your analysis leads us to believe that, when it comes to human-AI interaction, you focus not only on “machines,” but also on the real consequences they have on people. Together with Luciano Floridi, you investigated the relationship with the ‘new machines’. What role do you see for generative AI in the healthcare of the future? Will it be decision support or will it radically change the paradigm of care?**

I just want to go back to what I just said. Neither generative nor predictive AI will radically change the paradigm, but they will be agents of change in the “connective tissue” of work: the way documentation is produced and consulted; the methods of coordination and communication between professionals—allowing for more efficient handoffs and promoting the sharing of information (even from unstructured texts or informal conversations); exchanges between professionals and patients, making them more direct and clear, a real resource for care. The goal is to make AI a normal, almost invisible technology —integrated into tools and processes rather than disruptive.

Let me give you some examples. Imagine systems that listen to and transcribe clinical conversations (with all the necessary safeguards) and extract notes from them to increase patient adherence to treatment; or systems that interview patients both before and after visits to produce summaries for doctors or assess patients’ level of understanding and stress, respectively; systems that compile medical records or reports to be validated from these conversations and interviews, or that double-check consistency and completeness and verify: “Is anything missing?”, “Is there a discrepancy with the patient’s history?”, “Is this therapy compatible with that data?”; and then, of course, they are available to any operator who is entitled to answer questions about the patient, to quickly retrieve guidelines, relevant scientific papers, or to suggest (only suggest) courses of action and alternatives, explaining the pros and cons of the main options. This is very close to the ideal of *disappearing computation*: *computation* that does not impose itself, but supports the work without invading or distorting it. And this leads to a natural goal: AI should help healthcare professionals work better *together, collaborate*. Not just “empowering individual doctors” (augmenting them, as they say in jargon), but facilitating collaboration and multidisciplinary: bringing out different points of view, integrating skills, coordinating interventions, and above all, getting specialties that currently interact too little or too late to talk to each other. If technology makes that conversation more informed, more timely, and more traceable, then it’s already tangibly improving care. In the care of complex patients with multimorbidity—whose numbers continue to grow with an aging population—the quality of care increasingly depends on the quality of teamwork, as several studies have shown. And this applies not only in highly specialized settings, but also where the burden is daily and resources are more strained, with less visibility and recognition: family doctors, nurses, pharmacists, continuity-of-care professionals—anyone involved in responding to a citizen’s care and assistance needs. When this happens, AI becomes ‘standard’.

**Given your experience at Galeazzi and San Raffaele—large hospitals in Milan—how do you intend to bridge the gap between basic FBK research and clinical application *at the patient’s bedside?*”**

I would broaden my perspective: AI is not only about the critical phase of illness—the “patient’s bedside”—but also about prevention, continuity of care, and risk reduction. Bridging the gap between research and practice therefore means, first of all, designing with the clinic and not for the clinic; with doctors and not for doctors; and with patients—not treating them as passive users, but as co-creators of meaning and producers of value (for example, through their data, preferences, and lived experience). Basic research is indispensable, but translation often fails in the last mile: integration, usability, maintenance, accountability, training, and exception handling. I would like to systematize a pathway—an *FBK method for digital health*: selecting use cases in which the benefit is measurable and the need clearly defined; requiring external, multi-context validation as a prerequisite; and adopting clinically meaningful metrics. In medicine, errors do not all weigh equally; statistical distributions change; rare cases count and affect people in the flesh. If we continue to talk only about average accuracy, we risk speaking a language that does not serve clinicians and that, in the long run, erodes confidence.

And then there is an organizational aspect: research must also produce “adoption artifacts”, not just computational models. Usage protocols, operational guidelines, monitoring tools, governance models. This is what transforms a technology into a capability. On this front, I see it as strategic for multiple FBK centers to speak with one another and coordinate—to share lessons learned and *best practices* from the field, to circulate and combine complementary and convergent skills, and to respond to the needs of the complex and heterogeneous landscape that defines contemporary health care. Ideally, this would happen through orchestrated and synergistic interventions at the design level, but also through the simple desire to do good research together.

**Is there a question you haven’t been asked yet, but that you would very much like to answer in order to better convey who Federico Cabitza is?**

The best way to get to know me is to talk or work with me — even better: do research with me, because that’s where I feel most comfortable. If I had to choose one question, it would be: “What makes you satisfied in your work?”

I come from a family of doctors—only children, for four generations. My father and grandfather were highly respected, nationally known physicians; my paternal great-grandfather, whom I never knew, was a leading doctor, beloved throughout the country where he practiced. They were all people who cared for other people, in the fullest sense of the term. I chose engineering, but my aspiration has never been to “make technology.” It has been to contribute—through the tools I develop and the evidence I help generate—to making people feel better, and to making care, as the capacity of a community and a system, more effective, more equitable, and more sustainable.

For this reason, one of the basic principles of the medical profession applies to me: *primum non nocere*. Over twenty years of work at the intersection of IT and healthcare, I have come to understand that one of the most fragile—and most precious—things is the relationship between those who provide care and those who receive it, together with the quality of the clinical reasoning that sustains it. Technology should enhance communication, access to the best available knowledge, and decision-making; but in doing so, it must not jeopardize either the care relationship

or the quality of clinical reasoning. So, if artificial intelligence reduces the routine and documentation burden and gives time back to the caring relationship and empathic communication, I see this as a concrete good—and as my professional goal. But it must not do so by creating dependency, diminishing vigilance, or shifting responsibility in opaque ways. This idea led me to see AI not as a machine that “gives answers,” but as a device that *opens up* reasoning—one that lays alternatives on the table, argues for and against them, makes uncertainties and *trade-offs* explicit, and facilitates informed, competent discussion. Generative AI is also powerful because it is persuasive; predictive AI often is because it carries a reputation for objectivity and accuracy—not always deserved. Learning to rely on these tools appropriately—to achieve and sustain what the literature calls *appropriate reliance*, that is, calibrated and justified trust—is one of the most important skills an AI user can develop, and a key objective for those who design these systems, so they generate value without causing harm.

I'll close with one final thought: even those who design these systems must recognize their limits. Engineers are not omnipotent. And no research center, however ambitious, can “solve” all the problems I mentioned above. What I can promise—what I would like to leave—is a direction and a method: an *FBK method* for digital health grounded in evidence of effectiveness, an *evidence-based design* that combines vision with concreteness, imagination with rigor, international openness with close ties to the local communities. If one day I can say, sincerely, that I have contributed to this—and that I have passed it on to those who work with us and to those who will come after us—then I can say, with serenity, “I am satisfied.” That’s my idea of satisfaction.

#### PERMALINK

<https://magazine.fbkeu/en/news/ai-becoming-normal-federico-cabitzas-vision-for-the-future-of-digital-health/>

#### TAGS

- #artificialintelligence
- #digitalhealthwellbeing
- #generative AI
- #human-centered
- #innovation
- #sanità digitale
- #trentinosalute4.0

#### AUTHORS

- Marzia Lucianer