Contextual intelligence: Integrating multimodal data for complex, multimorbid patients



ELSEVIER

Advancing human progress together

Executive summary

Overview

The Elsevier European Physician Advisory Board examined how AI can integrate multimodal and longitudinal data to deliver real-time, context-aware recommendations for patients with multimorbidity. Drawing on a structured evidence review and the combined clinical experience of board members across European health systems, the group identified five areas where AI's potential and its barriers are most consequential.

The problem

Fragmentation is the root failure of current multimorbidity care. Patients are managed across disconnected specialties, each working from partial data, without visibility of longitudinal disease trajectories or cross-condition interactions. Multimorbidity introduces non-linear complexity — a decision in one domain can destabilize another — and context-aware care remains the exception rather than the norm.

What AI can offer

AI's primary value is contextual integration: unifying multimodal inputs — clinical notes, laboratory results, imaging, medications, patient-generated data — into a continuously updated patient narrative. The greatest near-term impact lies at high-risk, time-sensitive decision points: deterioration alerts, polypharmacy flagging, and discharge risk stratification. Longer term, the board points toward agentic, trajectory-aware systems that monitor evolving clinical states in real time, moving beyond episodic decision support toward continuously adaptive intelligence.

Key barriers

The board's overarching finding is that barriers are organizational, not technical. Most health systems lack the infrastructure for multimodal integration, and clinical data in practice is routinely incomplete, inconsistently coded, and fragmented. Interpretability remains a critical unresolved challenge: without explainable outputs, AI will not gain clinical traction — particularly in multimorbidity, where treatment trade-offs must be transparent to both clinicians and patients. No consensus yet exists on minimum explainability thresholds.





Recommendations

1. Data infrastructure — establish unified longitudinal records

Prioritize unified longitudinal records aggregating multimodal data across care settings for high-risk cohorts, within interoperable data lake infrastructure. Without this foundation, AI will remain limited to fragmented insights.

2. Deployment focus prioritize 2–3 high-risk decision points

Avoid broad, unfocused strategies. Deploy targeted pilots at high-impact clinical moments — deterioration alerts, polypharmacy checks, discharge stratification — with clear, pre-defined success metrics.

3. Interpretability — define minimum explainability standards

Establish minimum interpretability requirements for AI in clinical use, including visibility of key decision drivers and confidence levels. Co-design explainability features with clinicians rather than deriving them from technical defaults.

4. Data quality — invest in standardization for priority pathways

Invest in data standardization for priority pathways while adopting models designed to tolerate real-world imperfection. Begin with a single bounded domain — such as cardiometabolic care — before scaling.

5. Longitudinal capability — build trajectory-based intelligence incrementally

Incorporate longitudinal data streams starting with high-frequency data such as labs and vital signs. Develop simple trend-alert models before progressing to complex predictive systems.

6. Governance — move from exploration to governed execution

Establish governance frameworks covering liability, equitable model performance, and post-deployment monitoring for drift. Launch controlled pilots with embedded feedback loops to ensure learning is captured and scaled systematically.



Background and rationale

Elsevier is a global information and analytics company specializing in scientific, technical, and medical knowledge. Founded in 1880 and headquartered in Amsterdam, it provides journals, books, research databases, and digital solutions used by researchers, clinicians, educators, and policymakers in more than 170 countries. As part of RELX, Elsevier combines peer-reviewed scientific content with advanced data analytics and digital technologies to support evidence-based research and clinical decision-making. Within healthcare, Elsevier has expanded its focus on the integration of artificial intelligence with trusted medical knowledge to support clinical practice, research, and education. Its health division publishes hundreds of journals and thousands of books each year and delivers widely used digital tools for clinical decision support and medical learning.

In recent years, Elsevier has increasingly explored how AI can enhance the delivery of reliable clinical information at the point of care. By combining curated medical content with emerging AI capabilities, Elsevier aims to support clinicians in navigating growing volumes of medical evidence while maintaining high standards of clinical reliability and transparency.

Examining AI in healthcare is particularly important now as health systems face increasing pressures from workforce shortages, rising patient demand, and rapidly expanding medical knowledge. AI technologies are developing quickly and are beginning to influence clinical workflows, diagnostics, and decision support, making timely evaluation essential to ensure safe, effective, and responsible adoption.

The Elsevier European Physician Advisory Board (EPAB) brings together physicians from across Europe who provide clinical insight into emerging technologies and healthcare innovation. Drawing on their diverse experience across European healthcare systems, the group selected the topic of AI for clinical decision support in multimorbid acute and complex care to reflect their shared, pan-European perspective on how AI is shaping clinical practice and decision-making for clinicians today.

Process

Topics and questions for this white paper were selected through a structured, multi-step process designed to ensure transparency, evidence-informed discussion, and clinical relevance. First, the EPAB aligned on the overall objectives of the white paper, focusing on thought leadership and practical insights for clinicians. Members were then presented with six potential thematic areas related to artificial intelligence and healthcare delivery.

In the first round of voting, each member selected two priority topics, with aggregated results shared with the group; this process identified AI for decision support in acute and complex care (multimorbidity) as the focus. Subsequently, 8–10 potential discussion questions were developed using literature review and expert input.

A second round of voting selected the final key questions guiding the white paper. Following a feasibility assessment, a single, multi-part question was identified as the focus for this inaugural white paper.

Following the above process, the question: How can AI integrate multimodal and longitudinal clinical data to deliver real-time, context-aware, and interpretable recommendations for multimorbid patients?

A literature search was subsequently carried out and key papers reviewed by the EPAB for relevance and compared with the combined direct clinical exposure to the topic identified. A non-exhaustive checklist of ideas to consider for narrative synthesis and screening of evidence was used to assist the evidence review (Appendix 1).



Translating knowledge into clinical practice

After review, our Advisory Board have drawn the following five conclusions for the next phase of implementing AI in multimorbid care, along with a specific call to action for each section:

Building context-aware care from fragmentation

"Multimorbidity is breaking current care models that were never designed to manage interacting diseases over time."

There is a clear and consistent recognition that **fragmentation is a key root failure point in modern healthcare systems in effectively managing multimorbidity.** Patients with multiple chronic conditions are routinely treated across siloed specialties, each operating with partial datasets and narrow clinical objectives. This exists both as a structural limitation within organisational designs, as well as in an informational capacity. Clinicians are forced to synthesise disparate data sources under significant time pressure, often without visibility of longitudinal disease trajectories or cross-condition interactions.

Our board emphasises that multimorbidity introduces **non-linear complexity:** treatment decisions in one condition can destabilise another, and risk is distributed rather than isolated. AI's value proposition, therefore, is not incremental efficiency, but **contextual integration** — the ability to unify multimodal data (clinical notes, labs, imaging, medications, patient-generated data) into a coherent, continuously updated patient narrative, giving context where none presently exists.

The Advisory Board believes that the challenge of embedding AI (in contextual integration) is not conceptual but operational. While multimodal integration is technically feasible, most health systems lack the infrastructure to support real-time synthesis across care settings. The Advisory Board notes that data fragmentation persists even within single institutions, let alone across networks. As a result, clinicians continue to rely on incomplete snapshots rather than longitudinal insight.

The practical implication is that **context-aware care is currently the exception, not the norm.** AI has the potential to change this by enabling dynamic prioritization — surfacing what matters most for a patient at a given moment, based on evolving clinical context rather than static guidelines. This points toward a broader evolution: from AI as a static decision support tool — consulted episodically and applied to discrete problems — toward continuously adaptive clinical intelligence frameworks that operate under dynamic uncertainty, updating their outputs as patient states and contextual factors change in real time.

> Health systems must move beyond digitization toward **true data integration.** The next pragmatic step is to **establish unified, longitudinal patient records that aggregate multimodal data across settings,** even if initially limited to high-risk cohorts. These records should be securely housed and systematically curated within interoperable multimodal data lake infrastructure that enable scalable access and analysis. Without this foundational architecture, AI will remain constrained to fragmented insights rather than delivering meaningful clinical value.

Real-time decision support in high-risk scenarios: Making moments matter

"The real opportunity is not more data — it is better decisions at the moments that matter most."

A dominant theme across the literature is that **AI's greatest clinical impact lies in high-risk, time-sensitive decision points.** Task-specific AI applications — rather than broad, generalized systems — currently offer the most immediate path to adoption and value.[6] This view is echoed by our experience. To a clinician, focused, actionable recommendations offer more direct benefit presently than complex, multi-layered outputs.

Multimorbidity amplifies risk because clinical decisions are rarely isolated. For example, a medication change may influence renal function, cardiovascular stability, and metabolic control simultaneously. Multimodal AI



systems can detect patterns across these domains, enabling earlier identification of deterioration or adverse interactions.[2] Such systems outperform single-data-stream models in predictive accuracy, particularly in complex patient populations.[4]

However, there is strong consensus across the Advisory Board that **accuracy alone is insufficient.** Members emphasized that decision support must be **timely, relevant, and operationally usable.** Alerts that arrive too late, or recommendations that lack clarity, risk being ignored. Clinicians prioritise tools that directly influence outcomes — such as preventing deterioration or avoiding hospital readmission.

In practice, this should translate into AI systems that operate at critical junctures:

- Identifying early signs of clinical deterioration
- Flagging unsafe medication combinations in real time
- Supporting discharge decisions with risk stratification
- Prioritizing patients for intervention based on dynamic risk

The emphasis is clear: **AI must intervene where clinical uncertainty and risk are highest,** not where data is easiest to analyze.

> Organizations should prioritize **2–3 high-risk clinical decision points** where AI can demonstrably improve outcomes. The next step is to **deploy targeted pilots in these areas (e.g., deterioration alerts, polypharmacy checks) with clear success metrics,** rather than pursuing broad, unfocused AI strategies.

Interpretability as a clinical requirement: Building trust in recommendations

"If clinicians cannot understand the recommendation, they will not trust it — regardless of how accurate it is."

Interpretability emerges as one of the most critical — and unresolved — issues across the evidence base. Whilst multimodal modelling has made significant advances in recent times, increasing complexity often reduces transparency.[1,3] This creates a fundamental tension: **the most powerful models are often the least explainable.**

The Advisory Board is unequivocal that without interpretability, AI will struggle to gain traction in clinical environments. Trust is not built on performance metrics alone — it is built on understanding. This reflects a broader distinction that the Advisory Board considers increasingly important: the gap between technical model performance — accuracy, sensitivity, specificity measured in controlled conditions — and real-world cognitive integration, meaning whether a tool is actually usable within clinical workflows, under time pressure, and alongside the judgements clinicians are already making. As AI deployment scales, this gap may prove one of the most consequential differentiators between tools that transform care and those that are quietly abandoned. Clinicians must be able to interrogate recommendations, assess their validity, and explain them to patients and colleagues.

This is especially true of shared decision-making. In multimorbidity, where treatment decisions often involve trade-offs (e.g., symptom relief versus side effects, short-term versus long-term outcomes), **transparency becomes essential for aligning clinical and patient priorities.**[3] AI systems must therefore provide not only recommendations, but also the reasoning behind them — highlighting key drivers, uncertainties, and alternative options.



However, the Advisory Board recognizes that **full transparency may not always be achievable** in highly complex models. This creates an unresolved question: what level of explainability is sufficient for safe and effective use? The Advisory Board agrees there is a lacuna in current evaluation frameworks, with limited consensus on standards or thresholds.[4]

> Health systems should define **minimum interpretability requirements for AI tools in clinical use,** including visibility of key drivers and confidence levels. The immediate step is to **co-design explainability features with clinicians,** ensuring outputs align with real-world decision-making needs rather than technical abstractions.

Handling imperfection: The "messy data" reality

"AI does not fail in theory — it fails in the reality of incomplete, inconsistent, and fragmented data."

The literature is clear: a significant gap exists between controlled research environments and real-world clinical data.[5] Across institutions, data is frequently:

- Missing or incomplete
- Inconsistently coded
- Distributed across multiple systems
- Temporally fragmented

This reality is reinforced by our Board's personal experience. Models trained on curated datasets often perform poorly when exposed to real clinical workflows. The lack of interoperability between electronic medical records is a critical barrier to scale.

The implication is that **data quality is not a technical detail — it is a strategic constraint.** Even the most advanced AI models cannot compensate for fundamentally flawed input data. Many high-performing models are validated on retrospective datasets that do not reflect operational complexity.[4]

However, there is also a shift in thinking. Rather than attempting to "clean" data to perfection, there is growing recognition that AI systems must be **robust to imperfection.** This includes:

- Handling missing data without failure
- Adapting to variability in coding and documentation
- Learning from heterogeneous data sources
- Operating across institutional boundaries

At the same time, this does not eliminate the importance of data quality at the point of entry. Where feasible, quality checks at source remain the gold standard for quality assurance. AI should not replace good data practices but complement them — handling imperfection while helping to improve data quality over time.

This represents a move from idealized models to **resilient, real-world systems.**

> The priority is not perfect data, but **usable data at scale.** Health systems should invest in **data standardization and interoperability for priority pathways,** while simultaneously adopting AI models designed to tolerate real-world data limitations. A pragmatic first step is to **focus on one clinical domain (e.g., cardiometabolic care) and improve data consistency within that scope.**

Multimodal and longitudinal intelligence: Beyond snapshots

"The real value of AI lies not in single data points, but in understanding how conditions evolve and interact over time."

Multimorbidity is inherently dynamic. Patients experience fluctuating symptoms, evolving disease states, and changing responses to treatment. Yet much of current clinical decision-making remains **snapshot-based,** relying on isolated encounters rather than continuous trajectories.

Multimodal AI has the potential to move beyond this limitation. By integrating longitudinal data — such as trends in laboratory values, imaging changes, medication history, and patient-reported outcomes — AI can provide a **temporal understanding of disease progression.**[1,4] This is particularly important in multimorbid patients, where interactions between conditions may only become apparent over time.

The Advisory Board emphasizes the importance of continuous monitoring and adaptive systems. **Agentic AI,** where systems do not simply analyze static data, but actively track and respond to evolving clinical states should be applied, alongside multimodal integration to improve diagnostic accuracy and treatment planning.

However, significant challenges remain. Longitudinal data is often incomplete, inconsistently recorded, and difficult to align across modalities. Temporal fragmentation is a major barrier, with data captured at irregular intervals and across disconnected systems.[5]

Despite these challenges, the direction of travel is clear: **future AI systems must be trajectory-aware, not snapshot-driven.**

> Health systems should begin to **incorporate longitudinal data streams into AI use cases,** starting with high-frequency data such as labs and vital signs. The next step is to **develop simple trajectory-based models (e.g., trend alerts) before progressing to more complex predictive systems,** building capability incrementally.



Conclusion

Consensus, uncertainty, and strategic direction: From promise to practice

"The challenge is no longer whether AI can support multimorbidity care — it is how to implement it safely, effectively, and at scale."

Across the literature and the Advisory Board's personal experience, there is a striking level of consensus on the direction of travel. AI is widely recognized as a necessary enabler of improved multimorbidity care, particularly through multimodal integration and longitudinal analysis. There is also agreement that **the primary barriers are organizational, rather than technical.**

Key areas of consensus include:

- The importance of multimodal data integration
- The value of longitudinal, trajectory-based insights
- The need for targeted, high-impact use cases
- The centrality of workflow integration and clinician trust
- The critical role of governance and interpretability

However, significant uncertainties remain. These include:

- How to scale AI across heterogeneous health systems
- How to define and measure explainability in complex models
- How to manage liability in AI-assisted decisions
- How to ensure equitable performance across populations
- How to monitor and maintain model performance over time

Additionally, there are concerns around the lack of attention to model drift and post-deployment monitoring, as well as a disconnection at times between theoretical models and practical implementation.

For healthcare leaders, this creates a dual imperative: **act now on areas of consensus, while actively managing areas of uncertainty.**

> The strategic priority is to **move from exploration to execution.** Health systems should establish **clear governance frameworks, prioritize high-value use cases, and commit to iterative deployment with continuous evaluation.** The next pragmatic step is to **launch controlled, real-world pilots with embedded monitoring and feedback loops,** to ensure that learning is captured and scaled systematically.
>
> AI can transform multimorbidity care — but only if it is **grounded in clinical reality, embedded in workflows, and built on trustworthy data foundations.**
>
> The opportunity is immediate, but so is the responsibility to implement carefully.

References

1. Ardic, N., & Dinc, R. (2025). Emerging trends in multi-modal artificial intelligence for clinical decision support: A narrative review. *Health Informatics Journal, 31(3),* 1–16.
2. Deng, W., Zhang, L. Y., Yue, J. R., & Huang, X. L. (2026). The future of multimorbidity management in the older adults: Transforming AI-enabled precision medicine. *Frontiers in Public Health, 13,* Article 1691682.
3. Kulkarni, R., et al. (2025). Rethinking complex care using participatory medical cognition and user-driven learning amidst multimorbidity: Participatory action research study. *Journal of Participatory Medicine, 17,* Article e81950.
4. Schouten, D., et al. (2025). Navigating the landscape of multimodal AI in medicine: A scoping review on technical challenges and clinical applications. *Medical Image Analysis, 105,* Article 103621.
5. Wasi, A. T., & Ridoy, S. Z. (2025). Position: Real-world clinical AI requires multimodal, longitudinal, and privacy-preserving corpora. *Advances in Neural Information Processing Systems (NeurIPS 2025).*
6. Zhang, R., Chen, Y., Yue, W., Zhang, Y., Li, X., Feng, S., Yuan, F., & Luo, M. (2026). Multimodal artificial intelligence in medicine: A task-oriented framework for clinical translation. *Frontiers in Medicine, 12,* Article 1736272.

Glossary

Agentic AI	AI systems that do not passively wait to be queried, but actively and autonomously monitor data, track evolving states, and initiate actions or alerts based on what they detect — without requiring a human prompt at each step. In clinical contexts, this means systems that continuously observes a patient's condition and respond to changes in real time.
AI	Artificial Intelligence. The use of computer systems to perform tasks that would typically require human intelligence, such as recognizing patterns, making predictions, or generating recommendations. In healthcare, AI refers to algorithms and models that analyze clinical data to support diagnosis, risk stratification, and decision-making.
AI use case	A specific, defined application of AI technology to a concrete problem or task. In healthcare, a use case might be "deterioration alert" or "polypharmacy check" — a bounded scenario where AI is deployed to produce a particular output that supports a clinical decision.
Cognitive integration	The degree to which an AI tool's outputs can be understood, trusted, and acted upon by a clinician within the flow of their normal work. Distinct from technical accuracy — a tool may perform well in benchmarks but fail at cognitive integration if its recommendations are poorly timed, hard to interpret, or incompatible with how clinicians actually make decisions.
Contextual integration	The capacity of a system to unify data from disparate sources into a coherent, continuously updated whole — providing meaningful context rather than isolated data points. In healthcare AI, this refers to assembling a complete patient narrative from fragmented inputs across care settings.
Data lake	A large-scale digital storage system designed to hold raw data from multiple sources in its native format until it is needed for analysis. Unlike traditional databases, data lakes can accommodate structured data (e.g., lab values) and unstructured data (e.g., clinical notes or images) side by side, making them well-suited to multimodal healthcare AI.

Decision support (clinical)	Tools or systems that provide clinicians with relevant information, recommendations, or alerts at the point of care to assist in making decisions. Traditional decision support is typically rule-based and static; AI-powered decision support can be adaptive and context-sensitive.
Dynamic prioritization	The real-time identification and surfacing of the most clinically urgent or relevant information for a specific patient at a specific moment. Unlike static protocols, dynamic prioritization adapts as a patient's condition evolves, ensuring clinician attention is directed where it is most needed.
Electronic medical record (EMR) / Electronic health record (EHR)	Digital systems used by healthcare institutions to store patient health information, including clinical notes, diagnoses, medications, test results, and correspondence. Interoperability between EMR/EHR systems is a major barrier to multimodal AI integration across care settings.
Explainability / explainability threshold	Explainability refers to the capacity of an AI system to provide a human-understandable account of how it reached a particular output. An explainability threshold is the minimum level of explanation considered sufficient for safe clinical use — a standard that, as this document notes, does not yet have consensus definition.
Interoperability	The ability of different digital systems, software platforms, or institutions to exchange and use data with one another seamlessly. In healthcare AI, interoperability is a prerequisite for multimodal and cross-setting integration — without it, data remains locked in disconnected silos.
Interpretability	Interpretability in AI refers to the degree to which a human can understand the cause of a decision or prediction made by a machine learning model. It is the ability to map the internal workings of a model — its architecture, parameters, and feature combinations — into a form that is transparent and comprehensible to humans.

Longitudinal insight	An understanding derived from a complete record of data over time, as opposed to a snapshot from a single point. In a technical context, this requires systems capable of storing, linking, and analyzing data across extended timeframes and multiple sources.
Model drift	The gradual decline in an AI model's performance over time as the real-world data it encounters diverges from the data it was trained on. In healthcare, drift can occur as patient populations, clinical practices, or coding standards change. Post-deployment monitoring is required to detect and correct for drift.
Multimodal data	Health information drawn from multiple distinct source types — such as clinical notes, laboratory results, medical imaging, medication records, and patient-generated data. Integrating these sources gives a more complete picture of a patient's health than any single data type alone.
Patient-generated data	Health-relevant information recorded or reported by the patient themselves, outside of formal clinical encounters. This may include data from wearable devices, home monitoring equipment, symptom tracking apps, or patient-reported outcome measures.
Polypharmacy	The concurrent use of multiple medications by a single patient. In multimorbid patients, polypharmacy is common and increases the risk of adverse drug interactions, which AI systems can help identify and flag in real time.
Post-deployment monitoring	Ongoing evaluation of an AI system's performance after it has been introduced into clinical use. This includes tracking for model drift, unintended consequences, equity of performance across patient groups, and alignment with real-world outcomes — as distinct from the controlled conditions under which a model was originally validated.
Real-time synthesis	The continuous, automated integration of data from multiple sources as it is generated, producing an up-to-date and unified view of a patient's clinical status across settings and time points.

Risk stratification	The process of categorizing patients according to their likelihood of experiencing a particular adverse outcome — such as hospital readmission or deterioration. AI systems can perform dynamic risk stratification by continuously updating a patient's risk score as new data becomes available.
Trajectory-aware (systems)	AI systems designed to track how a patient’s condition changes over time, rather than assessing their state at a single point. Trajectory-aware models identify trends and patterns across data collected over weeks, months, or years, enabling earlier detection of deterioration or disease progression.

Appendix 1: Literature search strategy

- Define multimorbidity/complex patients
- Outline data types to integrate
    - Comment on QA of these
    - Clinical notes
    - Lab results
    - Medication
    - maging
    - Genomics
    - Miscellaneous Investigations (e.g., ECG)
    - Wearables
    - Psychosocial
- Integration methods (technical)
    - Interoperability with EHR systems
    - Human in loop
    - Alerts
        - Interruptive versus non-interruptive
        - Fatigue
    - Workflow
- Real-time clinical decision support versus longitudinal modeling

Search strategy

“multimorbid”

“multiple chronic conditions”

“complex patient”

“polypharmacy”

“clinical decision conflict”

“multimodal learning”

“dynamic risk prediction”

“real-time clinical decision support”

“context-aware AI”

“workflow-integrated AI”

“EHR-embedded AI”

“adaptive CDS”

Appendix 2: Biographies

Authors

Dr. Richard Daniels

Dr. Richard Daniels is a practising Paediatrician at St Mary's Hospital, part of Imperial College Healthcare NHS Trust in London, bringing frontline clinical insight to his work across the digital health landscape.

With a parallel career spanning agile start-ups to large-scale national public sector infrastructure programmes, Richard has developed a deep commitment to the safe and effective implementation of technology-augmented care — from community and primary care settings through to acute and emergency medicine. He is a recognised voice on the responsible deployment of AI in clinical practice, with a focus on how emerging technologies can strengthen healthcare systems without compromising patient safety.

Richard sits on the Elsevier European Physician Advisory Board and holds advisory roles across the industry and within his own Trust.

Dr. Mehmet Tarik Baran

Dr. Mehmet Tarık Baran is an Assistant Professor and Head of the Department of General Surgery at Mardin Artuklu University, with a background in surgical sciences and a strong focus on integrating artificial intelligence (AI) into healthcare.

His academic journey began with a medical degree from Istanbul University and specialization in General Surgery from Recep Tayyip Erdoğan University, where his thesis explored advanced imaging modalities in breast cancer surgery. He has published extensively on AI applications in medicine, including machine learning, deep learning, and generative AI, and holds numerous international certificates in these fields.

Currently, he actively researches and applies AI tools — such as computer vision, predictive modeling, and large language models — to enhance diagnostic accuracy, surgical outcomes, and personalized treatment strategies.

Dr. Baran is also involved in interdisciplinary administration and training programs to advance AI-driven innovations in healthcare, aiming to transform surgical practice and patient care through cutting-edge technological solutions.

Scientific Editor

Dr. Ximena Alvira

Dr. Ximena Alvira is a physician-scientist and Clinical & Research Manager at Elsevier Health. With more than 25 years of experience in clinical medicine, neuroscience research, medical writing, and healthcare innovation, she advises healthcare organizations worldwide on the evaluation and implementation of evidence-based clinical technologies, including artificial intelligence solutions.

Dr. Alvira holds a medical degree from Universidad El Bosque (Colombia) and a PhD in Neuroscience from Universidad Autónoma de Madrid (Spain), and completed clinical training at leading U.S. institutions, including Massachusetts General Hospital and Beth Israel Deaconess Medical Center. She collaborates with healthcare and academic institutions across Europe, the Middle East, Asia-Pacific, and Latin America on research strategy, clinical AI evaluation, and scientific publishing. Over her career, she has delivered more than 350 research and publishing workshops, reaching over 36,000 healthcare professionals worldwide.

Contributors

Dr. Edouard Battegay
Head, International Center for Multimorbidity & Complexity in Medicine (ICMC)
UZH Healthy Longevity Center, University Zurich

Dr. Sheeba Bhatia
General Adult Psychiatry Registrar, NHS
Greater Manchester West Mental Health NHS Foundation Trust

Dr. Wouter De Corte
Visiting Professor in AI for Healthcare, KU Leuven Belgium
Medical Lead, AZ Groeninge University Innovation Hub

Dr. Wassilissa Golowatjuk
MBA in Digital Business and AI
FAU Erlangen

Dr. Stefan Jovinge
President of Research, Education, Development, and Innovation
Scania University Hospitals

Dr. Claudio Ronco
Professor of Medicine
Director, International Renal Research Institute of Vicenza
Department of Nephrology, Dialysis, and Transplantation, San Bortolo Hospital, Vicenza



Copyright © 2026 Elsevier B.V.