



MINISTERUL EDUCAȚIEI
Universitatea Națională de Știință și
Tehnologie POLITEHNICA București
Școala Doctorală de Automatică și Calculatoare

Daniel Ioan-CHIȘ

TEZĂ DE DOCTORAT

UN CADRU SISTEMIC: INGINERIA LONGEVITĂȚII

**(A SYSTEMS FRAMEWORK: ENGINEERING
LONGEVITY)**

ABSTRACT

Conducător științific:

Acad. Prof.univ.dr.ing. Ioan DUMITRACHE

- 2025 -

**Universitatea Națională de Știință și Tehnologie
POLITEHNICA București**

TEZĂ DE DOCTORAT

UN CADRU SISTEMIC: INGINERIA LONGEVITĂȚII

**(A SYSTEMS FRAMEWORK: ENGINEERING
LONGEVITY)**

ABSTRACT

**Autor:
Daniel-Ioan Chiș**

COMISIA DE DOCTORAT

Președinte		
Conducător științific		
Referent		
Referent		
Referent		

- 2025 -

Table of Content

UN CADRU SISTEMIC: INGINERIA LONGEVITĂȚII	1
(A SYSTEMS FRAMEWORK: ENGINEERING LONGEVITY)	1
ABSTRACT	1
UN CADRU SISTEMIC: INGINERIA LONGEVITĂȚII	2
(A SYSTEMS FRAMEWORK: ENGINEERING LONGEVITY)	2
Chapter 1: Introduction – Fundamentals of Longevity Engineering	4
1. 1. 1. The Imbalance of the Healthcare System in the 21st Century	4
1. 2. Longevity Engineering - A New Paradigm	5
1.3. Thesis Statement and Overall Objectives	6
1.4. Overview of Original Contributions	8
Chapter 2: Current State of Knowledge	11
Chapter 3: A Hybrid AI-Based Simulation Framework for Optimizing Health and Economic Prosperity (Original Contribution 1)	14
Chapter 4: Ethics triplet Ia: A frame for extensive simulation of equity in health policies (original contribution 2)	20
Chapter 5: A frame of implicitly decipherable for clinical decision support (original contribution 3)	25
Chapter 6: Final conclusions and future directions	26
Future research directions	28
List of Scientific Papers Elaborated During the Doctoral Internship	30
Bibliography	32

Chapter 1: Introduction – Fundamentals of Longevity Engineering

1. 1.1. The Imbalance of the Healthcare System in the 21st Century

One field that has had immense success in the 20th century is that of medicine, profoundly influencing the trajectory and future of humanity.

The model defined in the previous century is one that can be categorized as reductionist, although extremely effective. This model made huge strides by dividing complex health problems into linear, cause-and-effect parts.

These advances led to new discoveries such as antibiotics that have limited the deadly impact of infections, vaccines that have eliminated devastating diseases, precise surgical procedures, and treatments that have successfully treated numerous acute conditions [1].

The increase in life expectancy became the ultimate metric of the era. This is a very telling measure, becoming an end in itself, which led to an unprecedented expansion of human existence on a global scale.

However, paradoxically, this success has leveled the ground for a profound crisis of the system in the 21st century. The model that propelled progress is now completely ill-adapted to the new circumstances it created [2]. Thus, the model has now suffered the consequences of its own success.

Today, health systems around the world face not a single challenge, but a convergence of three interconnected crises – demographic, epidemiological and economic – that threaten their very long-term sustainability. These are not isolated problems, but facets of the same systemic challenge, intertwining and amplifying each other, forming a “Gordian knot” that defies simple solutions.

Three crises are affecting health systems, these are:

- a. **Demographic crisis.** Considering that life expectancy has increased significantly, in conjunction with declining fertility rates in many regions of the world, a global phenomenon of population aging has been established and developed. This is particularly visible in developed societies. For example, in the European Union, it is estimated that the share of people over 65 years of age will increase from approximately 21% in 2022 to approximately 30% by 2050 [3]. This structural change in population dynamics is unprecedented in human history and exerts immense pressure on health systems, as older people use health services to a disproportionate extent. The current demographic reality is no longer a distant projection, but a defining component of today's social and economic landscape, which fundamentally changes the ratio between the active, economically contributing population and the one dependent on care services, with elderly people no longer in most cases being economically active and contributors to the tax systems necessary for payments to public health funds.
- b. **The epidemiological crisis** is the second crisis and is closely linked to the demographic one. As the population ages and infectious diseases are more effectively controlled, the burden of disease on the patient and the health system shifts from acute, episodic and sporadic diseases to chronic, non-communicable diseases (NCDs), such as cardiovascular diseases, diabetes, cancer and neurodegenerative diseases. Also, among the elderly population, multimorbidity has become the norm from an

exception [4]. Multimorbidity is characterized by the presence of several chronic diseases at the same time in an individual, which are usually kept in balance. However, this phenomenon represents a fundamental challenge for a health system designed around a model of:

one disease -> one specialist -> one treatment.

- c. **Economic crisis.** As a direct result of the two previous crises, the third, and most serious crisis, is the economic one, because it has an impact on the entire population. The current medical model has become much too expensive. At the same time, it is desired that the healthcare system be a social one, for social inclusion, but this is not economically viable, on the other hand the private healthcare system is much too expensive for the majority of the population, becoming inaccessible. Thus, in most countries, a hybrid system has been reached, through a public-private partnership. But even this is not sufficient in the face of current challenges. For example, the economic impact of the five major non-communicable diseases worldwide is estimated to reach 47 trillion dollars by 2030, which endangers the economic stability of the world [6]. Average health spending in EU countries will represent 10.4% of GDP by 2022 [7], a percentage that continues to rise despite economic growth. A dangerous negative feedback loop amplifies this cost increase: a shrinking working population must support a growing elderly population with increasingly expensive and complex medical needs, a change that is at odds with sustainable, long-term economic growth [8].

1. 2. Longevity Engineering - A New Paradigm

The crises analyzed above, together form a “Gordian knot”, which demonstrates that solutions cannot be found using the same methodology and paradigm that is used today, which perpetuates these problems.

Today, a fundamental reassessment of the way we define, measure and manage health is needed. A simple optimization of the status quo is not enough. In this section, the need for a leap from a reactive medical model focused on treating disease to a proactive, systemic model that pursues longevity engineering will be debated.

The fundamental deficiency of the reductionist approach is found in its basic assumption, namely, that a complex system, such as the human body or a health system, can be fully understood by analyzing its component parts in isolation [9]. This perspective, although extremely powerful in linear contexts, fails in the face of non-linear interactions, feedback loops, and emergent properties that characterize chronic diseases and multimorbidity [10].

The current system is hyper-optimized to perform interventions at the end of the causal chain (e.g., treating a myocardial infarction, excising a tumor, etc.). But at the same time, it is backward in managing the complex network of biological, behavioral, and social factors that lead to the occurrence of those events.

Thus, the paradigm shift requires, first of all, a change in the final goal. For a century, the ultimate metric of medical success was increasing lifespan. This goal was successfully achieved,

but it led to the paradoxical situation in which the years of life gained are often lived in a state of poor health and dependence on family and social systems.

The new paradigm proposes a more ambitious and relevant goal for human well-being: maximizing the healthy lifespan: the period of life spent in good health, without the functional limitations imposed by chronic diseases [2].

The goal is not just to add years to life, but to add life to years. This new goal transforms health from a cost to be managed into a strategic asset to be cultivated for individual and national prosperity.

A transformation of this caliber of goals must be perfectly aligned with the emerging vision of 4P Medicine, which encapsulates the direction needed for 21st century health systems.

4P Medicine can be defined by its four pillars:

- **Predictivity:** Seeks to identify individual risks of disease long before they manifest clinically.
- **Prevention:** Focuses on proactive and early interventions to reduce these risks and prevent the onset of disease.
- **Personalization:** Adapts prevention and treatment to each individual's unique genetic, biological, and lifestyle profile.
- **Participatory:** Involves patients as active partners in the management of their own health.

This 4P vision, although widely accepted as an ideal, suffers from a fundamental gap: the lack of a robust engineering architecture capable of operating it at the scale of an entire health system. Thus, the question arises that we try to answer: how do we move from a concept to a functional reality?

This thesis submits to the acceptance that the solution lies in adopting a focused perspective towards the field of systems engineering.

While complexity science and the theory of Complex Adaptive Systems (CAS) provide an essential descriptive framework for understanding why health systems behave in nonlinear and unpredictable ways [1], systems engineering provides the prescriptive methodology for designing and managing this complexity [11]. This technique represents a holistic and structured approach to the design and management of complex systems throughout their entire life cycle, from defining needs to operation and sustainability [12].

Thus, in conclusion, the need for a paradigm shift is not an option, but an imperative condition for the survival and prosperity of health systems.

This implies a threefold transition:

1. From a reactive, disease-centered model to a proactive, health-centered one aligned with the vision of 4P Medicine
2. From the goal of increasing lifespan to orchestrating and optimizing healthy lifespan
3. From reductionist thinking to a robust methodology based on systems engineering.

This new vision represents the foundation on which all the research presented in this thesis is built.

1.3. Thesis Statement and Overall Objectives

Based on the identified systemic challenges and the need to operationalize the 4P Medicine vision, this thesis proposes a central working hypothesis and a research direction aimed at addressing them.

Thesis Statement: The transition of health systems from a reactive and unsustainable model to a proactive "longevity engineering" paradigm, which operationalizes the principles of 4P Medicine, is possible through the development and integration of a multilevel systemic framework. This framework, grounded in the principles of systems engineering and implemented through advanced Artificial Intelligence tools, allows for the design, ethical governance and clinical application of optimized interventions at the public policy level (macro), while ensuring transparency and trust at the level of the individual medical act (micro).

Starting from this statement, the overall objective of the research is the design, conceptual development and validation of key components of this multilevel systemic framework. To achieve this goal, three specific objectives have been defined, each corresponding to an original scientific contribution and mapped to the pillars of 4P Medicine:

Specific Objective 1: Development of a macro-level simulation framework for optimizing health policies (Operationalization of Prediction and Prevention).

- **Problem:** Policy makers lack tools capable of predicting the long-term impact of policies and optimizing prevention strategies. Traditional models are static and cannot explore the complex interactions between health investments and national economic prosperity.

- **Anticipated result:** Design and validation of a hybrid simulation framework that integrates a top-down macroeconomic model with a bottom-up agent-based modeling (ABM). This "digital twin" for public health is capable of discovering optimal investment policies, validated through a concrete case study on a chronic disease with significant impact on the population.

Specific Objective 2: Design a governance framework for proactive ethical validation of AI (artificial intelligence) systems in health (Operationalization of Participation and Extended Prevention).

- **Problem:** The use of AI models to propose health policies can introduce significant ethical risks, such as perpetuating inequities. Ensuring a participatory and transparent process in the governance of these systems and preventing social harms are essential.

- **Expected result:** Development of a new architectural framework, called: "AI Ethical Triplet", which acts as an automatic ethical "red teaming" (validation) system. This methodology creates stress tests for AI-proposed policies against diverse synthetic populations, assessing the impact from the perspective of equity and robustness, thus evolving from a "Human-in-the-Loop" model (human as an active part of a systemic loop) to a proactive "Ethical Sandbox".

Specific Objective 3: Propose an interpretable AI architecture to increase trust in clinical decision support at the micro level (Operationalization of Personalization and Participation).

- **Problem:** The widespread adoption of AI systems in clinical practice is hampered by their "black box" nature, which undermines trust and prevents real physician participation in the decision-making process. An approach that supports personalized and transparent diagnosis is needed.

- **Expected result:** Design an "interpretable by design" AI methodological architecture. This framework, exemplified by an application for standardizing BI-RADS classification in mammography, mimics the human reasoning process, making the algorithm's decision fully transparent and auditable, and demonstrating how high-quality data at the micro level can feed into macro-level models.

The cumulative achievement of these three objectives demonstrates the viability and coherence of the proposed multilevel framework, providing an integrated, systems engineering-based solution to the complex health challenges of the 21st century.

1.4. Overview of Original Contributions

To achieve the specific objectives stated above, this thesis makes three original scientific contributions, each materialized in a published or submitted paper. These contributions are not isolated elements, but interconnected components of a coherent systemic framework, which addresses the issue of health system transformation at three critical levels: strategy (macro), governance (meso), and application (micro).

Contribution 1: An AI-Based Hybrid Simulation Framework for Optimizing Health and Economic Prosperity (Macro Level).

This first contribution, detailed in Chapter 3 and in [13], directly addresses the problem of strategic planning at the national level. It represents the Prediction and Prevention engine of the framework, allowing for in silico exploration of the future. The novelty lies in the hybridization of a top-down macroeconomic modeling with a bottom-up agent-based simulation (ABM), orchestrated by an AI optimization engine. This framework goes beyond traditional models by its ability to capture the "Equitable Development Cycle" – the positive feedback loop in which investments in health stimulate productivity, which generates economic growth, enabling further investments in health. Validation through the osteoporosis case study demonstrated the system's ability to identify policies that significantly outperform standard approaches, laying the technical foundation for the development of "Digital Twins" for public health.

Contribution 2: The AI Ethical Triplet: A Framework for Extensive Simulation of Equity in Health Policies (Governance Level).

The introduction of powerful AI systems raises a fundamental ethical question, which requires a robust and participatory oversight mechanism. This second contribution, presented in Chapter 4 and in [14], provides a direct answer. We propose the "AI Ethical Triplet", a governance architecture designed as a proactive and hostile validation mechanism. Unlike traditional audits, the Ethical Triplet is an automated system that stress-tests proposed policies against diverse synthetic populations, assessing the impact on equity, harm and robustness. This contribution formalizes a process of ethical "red teaming", transforming the "Human-in-the-Loop" oversight model into an "Ethical Sandbox", an advanced form of participation at the policy level that prevents social harms before implementation.

Contribution 3: An Implicitly Interpretable AI Framework for Clinical Decision Support (Micro Level).

For the macro-level vision to be sustainable, it must be based on high-quality data and the trust of micro-level actors. This third contribution, detailed in Chapter 5 and in [15], addresses the fundamental problem of the "black box" in medical AI. The proposed solution is a shift from models that require post-hoc explanations to models that are "implicitly interpretable by design".

The two-stage architecture mimics a physician's reasoning, ensuring that each result is accompanied by a clear and auditable justification. Exemplified by an application for standardizing the BI-RADS classification for breast cancer, this methodology supports personalized diagnosis and facilitates real physician participation in collaboration with AI. Crucially, this framework functions as an engine for generating structured and high-fidelity data, essential to feed and validate macro-level models, thus closing the loop of the multilevel systemic framework.




Figure 1. The proposed new framework

1.5. Structure of the Thesis

The current work is structured in six chapters, designed to guide the reader in a logical and coherent way, from defining the general problem to presenting the original contributions and analyzing their implications.

Chapter 1, "Introduction", establishes the context and justification of the research. It presents the converging crises that put health systems under pressure and argues for a paradigm shift towards a proactive one, based on systems engineering and aligned with the vision of 4P Medicine. The chapter culminates with the thesis statement, defining the objectives and an overview of the three original contributions.

Chapter 2, "Current State of Knowledge", provides the theoretical foundation of the work. It explores the foundations of complexity science and systems engineering, introduces 4P Medicine as a target paradigm, analyzes the emerging technological architecture (AI, Digital Twins, ABM) and discusses the implications of this transition. The chapter concludes with a critical synthesis, identifying the specific gap in research: the absence of an integrated engineering framework for the operationalization of 4P Medicine.

Chapter 3, "A Hybrid AI Simulation Framework for Optimizing Health and Economic Prosperity", presents the first original contribution. Here, the architecture of the macro-level simulation framework, designed as a Prediction and Prevention engine, is detailed, and validated through a detailed case study.

Chapter 4, "The AI Ethical Triplet: A Framework for Extensive Simulation of Equity in Health Policies", focuses on the second contribution, addressing the issue of ethical governance. The chapter introduces the architecture of the "AI Ethical Triplet" as a solution for proactive validation, operationalizing the principle of Participation at the policy level.

Chapter 5, "A Decipherable Implicit AI Framework for Clinical Decision Support", details the third contribution. This chapter focuses on the micro level, proposing an "interpretable by design" AI architecture to enable Personalization and Participation in the clinical act and to generate the high-quality data needed by the system.

Chapter 6, "Final Conclusions and Future Directions", concludes the paper. This chapter summarizes the results, presents the unified "Longevity Engineering" model as a complete operationalization of the 4P Medicine, discusses systemic implications, offers recommendations, acknowledges research limitations, and outlines future directions.

Chapter 2: Current State of Knowledge

To fundamentally address the systemic crises described in the previous chapter, a change in the theoretical overview through which we analyze health systems is necessary. The linear and mechanical assumptions of the reductionist model must be replaced by a conceptual framework that accepts the complexity, dynamism, and unpredictability inherent in public health.

In this section, the theoretical foundation of the thesis was established, building a detailed argument. It starts from the essence of the engineering discipline, defines the specific role of Systems Engineering, establishes the fundamental concept of "system", and evolves progressively towards the frontiers of complexity science and Complex Systems Engineering. The goal is to demonstrate why this integrated perspective is not only appropriate, but absolutely necessary to navigate the challenges of the 21st century in the field of health.

All technologies applied to operationalize 4P Medicine are useful when their transformative power derives from their synergistic integration into a unified architecture, which we call the "Simulation Trio".




Figure 2 Simulation trio architecture

In this architecture:

1. The Digital Twin provides the context and continuous connection to real-world data. It is the sandbox.

2. Agent-Based Modeling (ABM) provides the dynamic engine of complexity, simulating the interactions of millions of agents. It is the simulation engine.

3. Artificial Intelligence (AI) provides the intelligence to navigate, optimize, discover, and add behavioral realism to this complex world. It is the discovery engine.

Their seamless integration represents the underlying technology architecture that enables a paradigm shift towards systems-oriented public health, providing the tools needed to truly engineer longevity. This technology trinity is the engineering response to the identified gap, providing a plausible path to turn the promise of 4P Medicine into a functional reality.

This chapter presents a comprehensive argument for the need and feasibility of a new paradigm in healthcare, grounded in the principles of systems engineering.

We have demonstrated that the theoretical methodology of Complex Adaptive Systems (CAS) provides the essential descriptive framework for understanding why healthcare systems behave in nonlinear and unpredictable ways. On this theoretical foundation, we have shown that systems engineering provides the prescriptive methodology needed to manage this complexity in a structured and rigorous manner.

We have also identified the emerging technological architecture, a simulation trio of Digital Twins, Agent-Based Modeling, and Artificial Intelligence, that makes the large-scale application of these theoretical principles possible for the first time in history.

We have argued that the ultimate goal of this new technological capability must be a redefinition of success, moving from simply extending life expectancy to actively engineering healthy life spans (healthspans), a goal aligned with a more holistic vision of national well-being.

Finally, we explored the profound implications that this transition has for the entire medical ecosystem, from education and clinical practice to research.

The synthesis of the specialized literature reveals a landscape of knowledge that is, at the same time, vibrant and fragmented. There is a considerable volume of valuable research that explores each of these components in isolation. We find studies on the application of Digital Twins in medicine, papers on the use of ABM for epidemiological modeling, articles on the potential of AI in diagnostics, and a rich literature on systems thinking in healthcare management.

However, a critical analysis of the current state of knowledge highlights a major and fundamental gap in research: the absence of an integrated, multi-level systemic framework that connects these islands of knowledge into a coherent and functional whole.

Although the individual components exist, the "architecture" that unites them is missing.

More specifically, we identify three dimensions of this gap:

1. A Technology Integration Gap: There is a lack of a unified vision on how Digital Twins (the context), ABMs (the simulation engine), and AI (the discovery agent) can be combined into a synergistic computational pipeline for designing and optimizing health policies.

2. A Multi-Level (Micro-Macro) Connectivity Gap: Research tends to focus either at the macro level (public policy) or the micro level (clinical decision support), without providing a clear methodological link between the two. There is no framework to demonstrate how high-

fidelity data generated at the point of care (micro level) can feed into and validate population-level (macro level) simulations, and vice versa.

3. A Methodological Gap from Design to Implementation: A model is missing to guide the entire life cycle of a health innovation, from optimal policy design to its ethical governance to its safe and reliable application in clinical practice.

This triple research gap represents exactly the territory that this thesis aims to explore and define.

The original contributions that will be detailed in the following chapters are not intended to be isolated studies, but rather the constitutive pillars of this integrated systemic framework.

The following three chapters will present, in turn, solutions to each of these gaps, demonstrating how macro-level design, meso-level ethical governance, and micro-level application can be united in a coherent model for longevity engineering.

Chapter 3: A Hybrid AI-Based Simulation Framework for Optimizing Health and Economic Prosperity (Original Contribution 1)

The previous chapter established the theoretical foundation of the thesis, arguing that healthcare systems and national economies are, in essence, Complex Adaptive Systems (CAS), interconnected and inseparable.

Building on this understanding, this chapter presents the first original contribution of related research: a hybrid computational framework designed to translate this theory into practice.

The central objective is to answer a fundamental question for any 21st century policymaker: How can a nation, with finite resources, optimally invest in population health to reduce the burden of chronic diseases, increase healthy life expectancy, and thereby catalyze a sustainable cycle of economic growth and well-being?

The problem is not simple, but deeply rooted in complexity. The relationship between the health of a population and the economic prosperity of a nation is deeply nonlinear, characterized by feedback loops, time lags, and emergent properties that defy traditional analytical tools.

Conventional approaches, based on linear econometric models, Markov chains, or static cost-effectiveness analyses, while valuable in limited contexts, are fundamentally incapable of describing the dynamics of this interconnected system [13].

They cannot adequately model how a small but strategic investment in prevention can generate exponential economic benefits over decades, or how neglecting the health of a segment of the population can create, over time, a major obstacle to national development.

The conceptual cornerstone of this chapter is the "Vigilance Cycle of Development" (or "Vigilance Cycle of Health and Prosperity"), a positive feedback loop that governs the interaction between human and economic capital.

The logic of this cycle is intuitive but powerful:

1. **Investing in Health:** Resources (financial, human, technological) are strategically allocated to improve the health of the population, with a focus on the prevention and management of chronic diseases.
2. **Improving Human Capital:** A healthier population translates into more efficient human capital. Absenteeism rates decrease, labor productivity increases, and years of active and productive life are extended. Reducing the burden of disease (measured in DALYs) frees up individual and family resources. The average age of retirement increases.
3. **Stimulating Economic Growth:** A more productive and resilient workforce becomes an engine for innovation and economic growth. National output (GDP) increases, as does the tax base.
4. **Generating New Resources:** Economic growth generates new revenues for the state and the private sector, allowing additional resources to be reallocated to new investments in health, education, and infrastructure, thus closing and strengthening the cycle.




Figure 3 Systemic diagram of the Equitable Development Cycle

This view is solidly supported by fundamental economic theories.

The Human Capital Theory, developed by Schultz and Becker, posits that investments in the education and health of individuals are not simple expenses, but productive investments that increase the economic capacity of a nation [13] [Schultz, 1961].

The Grossman model (1972) further formalized this concept, treating health as a durable capital asset in which individuals invest to increase their "healthy days", which allow them to work and enjoy life.

This approach to proactively model and optimize the health-prosperity cycle represents the very operationalization at the macro level of two fundamental pillars of the 4P Medicine: Prediction and Prevention.

The computational framework proposed in this chapter, a Digital Twin for public health, is, in essence, the engineering engine for these two pillars.

It serves as a Prediction tool, allowing decision makers to go beyond historical data analysis and rigorously explore the future consequences of their decisions. By simulating thousands of "what-if" scenarios, the Digital Twin can predict the long-term impact of different investment strategies, revealing emerging trajectories of population health and economic prosperity. At the same time, the ultimate goal of this prediction is to inform and optimize Prevention strategies. Rather than reactively allocating resources to treat diseases after they have occurred, the framework enables proactive planning. It provides the ability to quantify the long-term return on investments in screening programs, awareness campaigns, or lifestyle interventions, thus justifying, through simulated evidence, the strategic shift of resources to

preventing diseases before they manifest. Thus, the contribution in this chapter provides the technical architecture to transform the concepts of Prediction and Prevention from aspirational visions into practical and quantifiable governance strategies.

However, the theoretical recognition of this link has not been accompanied by the development of practical tools capable of modeling and optimizing it. This is where the original contribution of this chapter comes in.

To overcome the limitations of traditional models, we designed an innovative multi-level computational framework that hybridizes two complementary modeling approaches: a top-down macroeconomic model that defines system-level constraints, and a bottom-up agent-based simulation (ABM) that models the heterogeneous behavior of a dynamic population.

This framework is more than a passive model, integrating an AI-based optimization engine capable of exploring thousands of policy scenarios to discover non-intuitive and high-impact strategies.

In this way, we transform the policymaking process from one based on intuition or static analysis to a data-driven discovery process similar to the optimization of a complex engineering system.

In this chapter, we have detailed the architecture and methodology of this simulation framework.

We have presented the mathematical model, the computational component, and the AI optimization engine.

Subsequently, we demonstrated the functionality and potential of the framework through a detailed case study on the management of postmenopausal osteoporosis, a chronic disease with significant population and economic impact.

Thus, we demonstrated not only the technical viability of the approach, but also its ability to generate valuable strategic insights, laying the conceptual and technical foundations for the development of "Digital Twins" for public health.

The successful validation of the hybrid computational framework through the postmenopausal osteoporosis case study transcends the simple demonstration of a methodology.

It paves the way for a practical and scalable implementation of this approach: the development of a "Digital Twin for Public Health". This section aims to discuss this vision, present a conceptual architecture for such a platform and analyze the challenges and ethical considerations inherent in such an endeavor.

The transition from a point simulation model to an institutionalized tool, such as a Digital Twin, represents the final logical step in the operationalization of the "longevity engineering" paradigm.

The methodological framework presented and validated in this chapter is, in essence, the prototype engine of a Digital Twin.

A simulation model, no matter how powerful, remains an academic or research exercise if it is not embedded in an accessible, reusable and interactive platform, intended for the end user – in this case, the policy maker, epidemiologist or health economist.

A Digital Twin for Public Health is more than a simulation; it is a socio-technical ecosystem.

It serves as a virtual, dynamic, and interactive replica of a nation's or region's health system, connected to real-world data sources and designed to function as a strategic "sandbox."

In such an environment, policymakers can move beyond the traditional, reactive policy-making cycle.

Instead of implementing costly and difficult-to-reverse interventions based on historical data and static models, they can explore, test, and refine dozens of *in silico* scenarios.

This ability to "simulate before you act" is transformative, enabling evidence-based, anticipatory governance that can optimize resource allocation, minimize unintended consequences, and increase system resilience to future crises.

Fundamentally, this tool is the engine that operationalizes the first two pillars of 4P Medicine at a systemic scale: **Prediction and Prevention**.

- **Prediction:** The Digital Twin transcends traditional statistical prediction, which relies on correlations extracted from historical data. Through Agent-Based Modeling, it provides a form of mechanistic prediction. It doesn't just predict that a certain outcome is likely, but simulates how and why that outcome emerges from millions of individual interactions. This allows for the anticipation of nonlinear phenomena and critical tipping points, such as the explosion of healthcare costs when the prevalence of a chronic disease exceeds a certain threshold. It can predict the health trajectories of entire demographic cohorts under different policy scenarios, identifying subgroups at highest risk of developing complications decades before this happens.
- **Prevention:** This superior predictive capability is the foundation on which a new era of proactive and intelligent prevention is built. The Digital Twin transforms preventive planning from an exercise based on general recommendations into an engineering optimization process. It allows decision-makers to design and test *in silico* hundreds of hybrid prevention strategies (primary, secondary, tertiary), to discover the optimal combination that maximizes healthy life years (HALE) for every monetary unit invested. Thus, it provides the quantitative evidence needed to justify the "investment in health" in the face of short-term budgetary pressures, demonstrating the long-term economic and social return of preventing diseases, instead of treating their consequences.

To translate this concept into reality, we propose a multi-tier architecture, based on a cloud infrastructure, that encapsulates the methodological logic in a functional platform (Figure 4). This architecture is designed to be modular, scalable and adaptable to different contexts (diseases, populations, policy objectives).

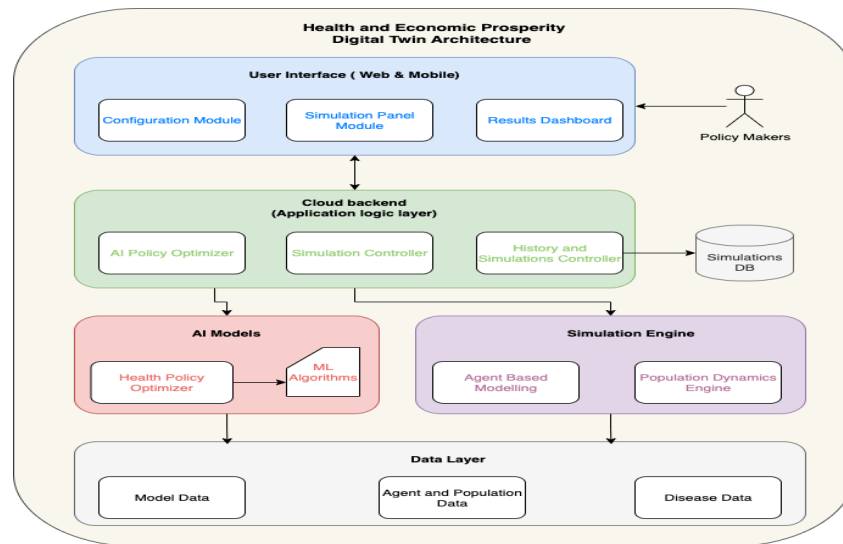


Figure 4: Conceptual Architecture for a Digital Twin for Health and Economic Prosperity

The architecture comprises the following main layers:

1. **User Interface Layer (Web & Mobile):** This is the entry point for decision-makers. An intuitive interface would allow them to interact with the system without requiring in-depth technical knowledge. Key components include:

- **Configuration Module:** Allows users to define the parameters of a new analysis: disease of interest, population demographics, time horizon and, crucially, weights for the National Well-being Score, thus aligning the optimization with current strategic objectives.
- **Simulation Panel:** Provides control over launching, monitoring and stopping simulations and optimization processes.
- **Results Dashboard:** Visualizes the results in a clear and comparative way, using interactive graphs and summary tables. Allows direct comparison of multiple candidate policies based on key metrics (DALYs, HALE, GDP, Well-being Score), illuminating the trade-offs between them.

1. Cloud Layer (Application Logic): This is the management brain of the platform, orchestrating all operations.

2. AI Optimization Controller: Manages the entire optimization flow, from generating training data to training the surrogate model and running the optimization algorithm.

Simulation Controller: Interfaces with the simulation engine, managing the instantiation of agent populations and the execution of simulations according to defined policies.

3. Simulation History Controller: Stores the results of all previous simulations and optimizations in a database, allowing for auditing, reproducibility and long-term comparative analysis.

4. Core Engines Layer (AI Models & Simulation Engine): This contains the core computational components described in the methodology. It is separated from the application logic to allow independent updates and enhancements (e.g. replacing the ABM engine with a more powerful version or adding new AI algorithms).

5. Data Layer: The foundation of the entire platform. It integrates and manages the data needed to anchor the Digital Twin in reality, including demographic data (census), epidemiological data (disease prevalence and incidence), economic data (treatment costs, productivity), and behavioral data.

Developing and implementing a Digital Twin for Public Health is an ambitious endeavor, involving significant challenges that must be proactively addressed.

Data Requirements: The accuracy and relevance of the Digital Twin are directly proportional to the quality, granularity, and timeliness of the data it is fed with. Obtaining longitudinal, individual-level data (even anonymized), on disease progression and health behaviors is a major challenge, given the fragmented nature of many health data systems [41].

Computational Cost: Running thousands of ABM simulations and training AI models requires a high-performance computing (HPC) infrastructure. Although cloud costs are becoming increasingly affordable, the initial investment and ongoing operational costs must be considered.

Algorithmic Inequality and Equity: A fundamental ethical concern is the risk that the Digital Twin will perpetuate or even amplify existing inequities in society. If the training data reflects historical disparities in access to care or health outcomes, the AI agent, in its effort to optimize for the average, could discover policies that further disadvantage vulnerable subgroups. This risk imposes the need for a robust ethical governance framework, an issue that forms the core of Contribution 2 (Chapter 4).

Governance of Objectives: Defining the "National Well-Being Score" is not a purely technical issue, but a deeply ethical and political one. The choice of weights () reflects a set of social values regarding the trade-off between economic growth, average population health, and equity.

This process must be transparent and involve a broad dialogue with all relevant stakeholders, including the general public, to ensure the democratic legitimacy of the objectives for which the system is optimized.

In conclusion, the transition from a methodological framework to an operational Digital Twin is essential to fully realize the potential of “longevity engineering”.

While the challenges are real, they are not impossible.

A phased, transparent approach guided by sound ethical principles can transform this vision into a powerful tool for creating smarter, more resilient and more equitable health systems.

Chapter 4: Ethics triplet Ia: A frame for extensive simulation of equity in health policies (original contribution 2)

The previous chapter demonstrated the huge potential of a hybrid computational architecture, culminating in the vision of a "digital twin for public health".

I argued that such an instrument, assisted by artificial intelligence (Ia), can function as a strategic "sandbox", allowing the decision makers to discover optimized health policies, with a positive impact on the economic well -being and the years of healthy life.

However, it is precisely this immense power to optimize and influence the destinies of millions of people bring in the foreground an equally great and complex challenge: ethical responsibility.

As we integrate more and more deeper algorithms in the decision -making infrastructure, especially in areas with high human stakes such as health, the undeniable promise of progress is accompanied by an equally indisputable risk.

The systems, by their nature, are designed to optimize for a specific objective function, in our case, the maximization of the "national well -being score".

This singular goal, although computational efficient, can create dangerous ethical perspectives. An algorithm that follows with maximum efficiency a lax objective, such as improving the average state of health, could, without adequate supervision, discover solutions that systematically disadvantage vulnerable subgroups, which exacerbates existing or unintentional negative consequences.

This is the manifestation in the field of health of what in the specialized literature of safety is known as the "issue of alignment of values" (Value Alignment Problem) [74].

The danger does not consist of a malignant intention of the system, but in the discrepancy between the literal objective I have programmed and the complex set and often tacitly of human values (equity, justice, compassion) that we would like to respect.

A reference study, conducted by Obemeyer [75], revealed how an algorithm widely used in the United States to allocate medical care resources systematically discriminates a group of patients.

The algorithm used previous medical costs as an indicator (proxy) for health needs, an erroneous assumption that ignored a demographic subgroup had access to fewer resources and, therefore, generated lower costs at the same level of disease.

The result was a system that, although it seemed neutral and objective, perpetuated and amplified a serious social inequity.

This eloquent example emphasizes a fundamental limitation of the traditional mechanisms of governance. Approaches such as post-implementation audits, although necessary, are fundamentally reactive. They identify the damage after they have already occurred.

On the other hand, the "Human-in-the-Loop" model (HITL), in which a human supervisor approves the decisions of the algorithm, although essential, is hit by the human cognitive limits.

Phenomena such as "Automation Bias" (tendency to have excessive trust in automatic systems) can reduce human vigilance, transforming supervision into a simple formality [76].

Therefore, it becomes obvious that, in order to use the power of the digital twins guided by IA, we need a new type of safety mechanism, one that is proactive, systematic and reactive.

It is not sufficient to check if an algorithm works correctly as per specifications; We must actively explore the ways in which it could fail from an ethical point of view, under various and unexpected conditions.

We need a system that not only validates, but "stresses" the proposed policies, deliberately seeking ethical vulnerabilities before they have consequences in the real world.

This necessity aligns and extends the vision of medicine 4P. If the pillars of prediction and prevention are operationalized by the digital twin, then the pillar of participation acquires a new, crucial dimension.

Participation can no longer be limited to the doctor-patient relationship at the micro level. In an era in which algorithms shape policies that affect the entire population, participation must become a systemic and democratic process of validating the values incorporated in these systems.

The framework of the governance we propose is, in essence, a mechanism of computational participation, which allows the society to interrogate and align algorithmic decisions.

Moreover, this framework extends the concept of prevention beyond individual biology, to systemic sociology.

The objective is no longer just the prevention of the disease, but also the prevention of social injury - preventing discrimination, inequity and erosion, which can be as harmful to the health of a nation as any pathogen.

This imperative for a proactive ethical validation is the starting point for the second original contribution of this thesis.

In this chapter, we propose a new governance architecture, called the "ethical triplet", designed to serve as counterpoint and validation mechanism for digital twins optimization of policy. This framework is not a simple set of recommendations, but a functional computational architecture, a kind of "automated ethics committee" that acts as a "Red Teaming" partner for the optimization agent.

The objective of this chapter is to detail this architecture, to describe its components, to operationalize the principles of medical ethics in quantifiable metrics and to demonstrate how this approach evolves the current supervisory model to an "ethical sandbox", an environment in which the ethical compromises inherent to any public policy can be explored, and before a transparent mode, implementation.

Thus, we complete the frame of "longevity engineering" with an essential component of caution, making sure that, in search of computational intelligence, we do not sacrifice human wisdom.

The present chapter introduced the architecture of the ethical triplet and articulated the conceptual transition from a "Human-in-The-Loop" model to a paradigm of "ethical sandbox".

Taken together, these components do not describe only a set of technical tools, but base a coherent and operational model for the responsible governance of artificial intelligence in public health.

This final section has the role of synthesizing these elements and discussing the wider implications of this model, arguing that it offers a practical way to navigate the inherent tension between rapid technological innovation and the ethical imperative of protecting human well-being.

The proposed governance model rests on three interdependent conceptual pillars: proactivity, deliberation and participation.

By integrating these principles, the ethical triplet frame transforms the discussion on ethics from a post-facing compliance into a values-by-design) design into the life development cycle.

The foundation stone of our model is the principle of proactivity. The traditional governance of technology was largely reactive.

Ethical problems, such as algorithmic discrimination or violation of confidentiality, are often identified only after a system has been largely implemented and has already caused damages. Ethical audits, when they take place, are frequently performed post-hoc, having a role of autopsy than preventive diagnosis.

The ethical triplet reverses this logic. Through the "stressor" and "executor", the frame creates an environment in which ethical risks can be systematically anticipated and explored before implementation.

This is a fundamental shift from "repairing the damage" to "preventing injury". The ability to simulate the impact of a policy on "stressed" populations - which reflect the scenarios of economic crisis, inequities of access or social distrust - is a powerful anticipation tool.

It allows the decision makers to identify the ethical "fracture points" of a policy in the safety of a virtual environment, transforming the governance from an art of reaction into a science of preparation.

The second pillar is the emphasis on deliberation. As argued in the previous section, the "Human-in-the-Loop" model can paradoxically weaken the human judgment by automation and by presenting a single "optimal" option.

In contrast, the paradigm of "ethical sandbox" is explicitly designed to empower and enrich human deliberation.

By generating a set of candidate policies and by presenting their performance on a comparative "ethical dashboard", our model refuses to provide a simple answer.

Instead, he asks a complex but essential question:

"Given these compromises between average efficiency, equity, risk of injury and robustness, which politically best lines with our strategic values and priorities?"

This approach has two major virtues. First of all, it makes explicit compromises. There is no illusion of a perfect solution, "win-win".

Decidants are forced to recognize and justify the difficult choices.

Secondly, it moves the final responsibility where the place is: to the human decision factor.

Technology is no longer an oracle that dictates the solution, but a counselor who illuminates the complexity of the decision.

The role becomes that of broadening the moral imagination of the decision maker, not to replace it. Thus, the model promotes a form of collaboration man in which the car deals with the computational complexity, and the value judgment.

The third and most important pillar is the imperative of participation. A model of governance, however technically sophisticated, cannot be truly responsible if it is developed and operated in a vacuum, by a small group of technical experts.

His legitimacy and effectiveness depend on the integration of a wide range of perspectives, especially those who will be most affected by algorithmic decisions.

In the paradigm of medicine 4P, the "participatory" pillar traditionally refers to the active involvement of the patient in his own care process, in a partnership with the doctor. Our framework extrapolates this fundamental principle from the individual to the society scale.

If the "patient" is the whole population, and "treatment" is a public health policy, then "participation" can no longer mean the informed consent of an individual.

It must mean a deliberative process, at the level of society, through which the values, priorities and fears of the community are incorporated in the very logic of the decision -making system.

This raises critical questions:

Who decides what demographic "stressors" are relevant and plausible?

Who sets the "injury threshold" that defines unacceptable injury?

Who determines the relative weights of the metrics on the "ethical dashboard"?

The answer can not only come from engineers or data analysts.

We propose that the operational implementation of the ethical triplet, and in particular the configuration of the "ethical auditor", be supervised by a multi-stakeholder governance committee.

Such a structure should include not only technical experts, but also:

Doctors and experts in public health to ensure clinical relevance.

Bioethicists and lawyers, to provide ethical and legal rigor.

Representatives of the associations of patients and vulnerable communities, to bring the experience lived and to ensure that their concerns are reflected in the metrics.

Factors of political decision and government representatives, to guarantee alignment with national strategic objectives.

Such a participatory approach ensures that the values incorporated in the system are not arbitrary, but are the result of a deliberative, democratic process.

Moreover, it creates a mechanism of contestability, a formal process by which the affected groups can question, challenge and request the revision of the parameters and results of the system.

This transparency and opening to feedback are essential to build and maintain public long-term confidence.

In conclusion, the responsible governance model proposed in this chapter is an integrated one.

It combines the technical rigor of the side simulation with an interaction paradigm that prioritizes human deliberation and a supervisory structure that incorporates democratic participation.

It offers a concrete plan to move beyond the mere statements of ethical principles, to socio-technical systems that are, by design, safer, more equitable and more aligned with the values of the society they serve.

It is a model that acknowledges that, in the age, the biggest challenge is not only to build smart systems, but to cultivate the wisdom to use them responsibly.

Chapter 5: A frame of implicitly decipherable for clinical decision support (original contribution 3)

The previous chapters built a large -scale systemic framework for "longevity engineering", approaching the challenges at the strategic level (chapter 3) and of ethical governance (chapter 4).

A computational architecture was presented for the optimization of health policies and a model of governance for their ethical validation.

However, no matter how sophisticated these macro models, their long -term validity and sustainability are critically dependent on a fundamental element: the quality, reliability and granularity of the data that supplies them.

The vision of a digital twin at the population level remains a fragile theoretical construction if it does not support on a constant flow of high loyalty data, generated at the point of direct interaction with the patient, the micro level of the medical act.

This fifth chapter descends from the macro perspective of the national strategy to the micro of clinical practice.

Here, at the micro level, the systemic framework faces a fundamentally different challenge, but equally important: building trust.

The doctor's confidence in the new technologies and the patient's confidence in the decisions made on the basis of them are the pillars on which the entire paradigm supports. Without this confidence, the large -scale adoption of any system, however high theoretical, is doomed to failure.

This chapter addresses a double challenge at this micro level: (1) the variability of inter-observatory diagnoses, a persistent problem that undermines the standardization and quality of care, and (2) the limitations of the artificial intelligence models of type "black-box), which, although often performing, generates a major barrier.

The culminating argument is that the integration of the three contributions creates a higher feedback loop, which is the essence of a learning health system, based on the principles of systems engineering.

This cycle turns the relationship between clinical practice and public health policy:

From Micro to Macro (Bottom-Up): each diagnostic act is no longer an isolated event. Through the Ia instrument, it becomes a source of knowledge that contributes to a systemic, dynamic and real -time understanding of the health of the population. The doctor, in his daily activity, becomes an active participant in refining the national health model.

From macro to micro (top-down): health policies are no longer abstract directives, imposed by above. They are strategies informed by realistic simulations, based on the data generated by the clinical practice itself. Moreover, the prospects obtained from the analysis at the macro level can guide the clinical practice. For example, if the digital twin identifies an unexpected growth of a certain type of injury in a region, this can trigger an alert and lead to targeted training programs for radiologists in that area.

In conclusion, the frame of inherently decipherable, detailed in this chapter, is much more than a decision -making tool.

It is the infrastructural and epistemological foundation on which the entire vision of the thesis is based. By solving the problem of trust and standardization at the micro level, it unlocks the potential to create a truly intelligent health system.

Chapter 6: Final conclusions and future directions

This thesis started from finding a systemic crisis in the health of the 21st century, a convergence of demographic, epidemiological and economic pressures that threaten the sustainability of the reactive medical model, inherited from the previous century.

I argued that a lasting solution requires a fundamental paradigm change: a transition from treating the disease to "longevity engineering", guided by the rigorous principles of systems engineering.

In order to operationalize this vision, we proposed a multivariate systemic framework, whose fundamental components were developed and presented as three original scientific contributions in the previous chapters.

The three original contributions of this work were conceived as the structural pillars of this solution, each addressing an essential challenge to a distinct level: strategy (macro), governance (MESO) and application (micro).

The first contribution addressed the fundamental challenge of strategic planning in public health.

I argued that the political decision -makers are lacking in tools capable of navigating the non -linear complexity of the relationship between health and economy.

To fill this goal, we have designed and validated a hybrid simulation frame based on Ia, conceived as a digital twin for public health.

Its novelty consists in the synergistic integration of a top-down macroeconomic model with a bottom-up simulation based on agents (ABM), creating a high virtual environment.

This framework is not a passive analysis tool, but an active discovery engine, equipped with an agent capable of exploring thousands of scenarios and identifying non-intuitive investment policies, with a high impact.

The case study on post-menopausal osteoporosis has demonstrated the practical ability of the system to discover strategies that simultaneously optimize the health of the population and economic prosperity, thus validating its potential to transform the elaboration of policies from a reactive process into an anticipatory and evidence based.

This contribution thus established the design component of the general systemic framework, directly operating the prediction and prevention pillars of 4P medicine on a national scale. We are no longer limited to predict the risk of an individual, but we can predict the probable trajectory of the entire health system under different policy regimes, allowing a truly strategic prevention.

The introduction of strong instruments to guide decisions at the population level brings with it a huge ethical responsibility.

The second contribution directly confronted this challenge, answering the question: How do we make sure that the optimized policies are fair, fair and does not cause unintentional damage to vulnerable subgroups?

The proposed solution was the "ethical triplet", an innovative architecture for governance, which functions as an automatic proactive and reactive ethical validation system. Unlike traditional audits, this framework performs systematic stress tests, using its three pillars, the population generator. The simulation engine and the ethical auditor, to evaluate the policies proposed through the prism of clear metrics of equity, injury and robustness. By this, the contribution has formalized a process of ethical "Red Teaming", evolving the supervisory model from a "Human-in-Loop" reactive to a proactive "ethical sandbox".

This component ensures that the optimization power of the digital twin is responsibly exercised, thus establishing the governance pillar of the systemic framework. Moreover, this framework extends the concept of prevention to social injury and operationalizes the participation in the policy level, creating a mechanism by which the societal values can be integrated and tested in the decision -making process, before any implementation takes place in the real world.

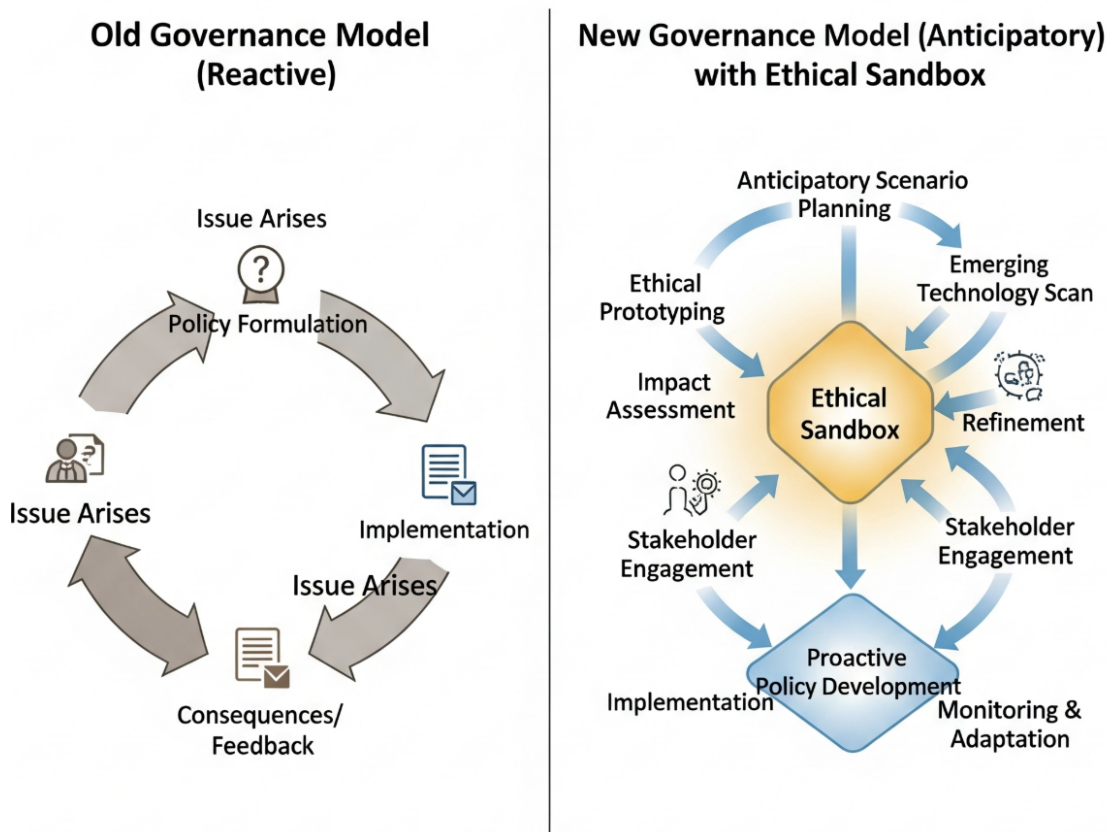


Figure 5 The evolution towards the new ethical governance model

Macro-level vision, no matter how sophisticated, is worthless if it is not anchored in the reality of clinical practice.

The third contribution addressed this essential link, focusing on the issue of trust and data quality at the point of care. We proposed a paradigm shift from black-box AI to an AI framework that is directly interpretable by design. The two-stage architecture, which mimics a physician's cognitive reasoning, was designed to provide full transparency and auditability, thereby building

clinician confidence. Validated through an application for standardizing BI-RADS classification in mammography, this methodology has a crucial dual role: (1) it acts as a trusted decision support tool, designed to reduce diagnostic variability and improve individual patient care; and (2) it functions as an engine for generating high-fidelity structured data. This second function is fundamental, as it transforms each medical act into a source of granular data, essential for the realistic calibration and validation of macro-level ABM models. Thus, this contribution establishes the application pillar and closes the information loop of the entire system. Thus, this contribution establishes the application pillar and closes the information loop of the entire system, being the engine that enables Policy Personalization at scale and facilitates a genuine Clinician Participation in the knowledge ecosystem.

The final synthesis reveals that these three contributions are not isolated elements, but interconnected pieces of a single logical architecture.

They form a fair feedback loop, the essence of a learning healthcare system:

From Micro to Macro: The Decipherable AI Tool (Contribution 3) generates structured and trusted clinical data, which serves as "fuel" for the Digital Twin (Contribution 1), saving it from the vulnerability of useless data.

At the Macro Level: The Digital Twin (Contribution 1) uses this data to design optimized strategic policies, whose safety and equity are rigorously verified by the Ethical Triplet (Contribution 2).

From Macro to Micro: The in silico validated policies are then implemented in clinical practice, where their impact is measured again by the AI tool (Contribution 3), generating new data and repeating the learning and refinement cycle.

Together, these components form a complete systemic framework, which addresses the life cycle of a public health intervention: from its strategic design and ethical governance, to its clinical application and impact measurement.

This synthesis demonstrates the coherence of the proposed vision and establishes the foundation for the final "Longevity Engineering" model that will be presented in the next section.

Future research directions

The completion of this work therefore marks the beginning of a long-term research program, aimed at maturing, validating and extending the "Longevity Engineering" Framework. Future research directions can be structured around three main axes, corresponding to the three levels of the proposed framework, culminating in an integrative vision.

1. Macro Axis: Maturing and Extending the Digital Twin

The macro-level contribution demonstrated the feasibility of a Digital Twin for optimizing health policies in a specific context. The next logical steps involve increasing the realism, complexity and area of applicability of this model:

- **Multimorbidity Modeling:** The case study focused on a single chronic disease. A major research direction is to extend the model to simulate the interactions between multiple chronic diseases (e.g. diabetes, cardiovascular disease and osteoporosis) at the agent level.

This would allow for a much more realistic modeling of multimorbidity, one of the biggest challenges of current health systems.

- **Integration of Socio-Economic and Environmental Factors:** The current model can be enriched by adding new layers of data and behavioral variables. Future research could integrate factors such as education level, access to healthy food, air quality or social stress level as endogenous variables that influence the health status of agents, allowing for the testing of intersectoral public policies.
- **Cross-Cultural and Cross-System Validation:** Applying and calibrating the simulation framework on data from other countries, with different health systems and demographic contexts, would test the robustness and generalizability of the model, providing valuable insights into universal versus context-dependent principles in health systems engineering.

2. Meso Axis: Operationalizing and Standardizing Ethical Governance

The "AI Ethical Triplet" was proposed as a conceptual and architectural framework. Moving towards a standardized and widely adopted tool requires research focused on implementation and validation:

Development of Standards: Creation of standard datasets and test scenarios ("ethical benchmarks") to objectively evaluate and compare the performance of different AI decision support systems in public policy. This would promote transparency and set a standard of rigor in the field of ethical AI.

Participatory Governance Study: Practical, case study research on the implementation of multi-stakeholder governance committees. These studies would explore the most effective ways to involve citizens and patient representatives in the process of defining equity metrics, ensuring the democratic legitimacy of the "Ethical Sandbox".

Research in "Explainable Fairness": Developing techniques by which the "Ethical Triplet" not only reports an inequity score, but can also explain the probable causes of that inequity within the simulation, providing decision-makers with a deeper understanding of the mechanisms that generate disparities.

3. Micro Axis: Clinical Validation and Extension of the Decipherable Architecture

The "decipherable by design" AI framework has been conceptually validated. The next stage is its rigorous validation in clinical practice and the extension of its philosophy to other domains:

Conducting Prospective Clinical Studies: Conducting a randomized clinical trial in which a group of radiologists uses the decision support system for BI-RADS classification, compared to a control group. The primary objective would be to quantitatively measure the impact of the system on reducing inter-observer variability (measured by the kappa statistic) and on diagnostic accuracy.

Extending the Architecture to Other Imaging Domains: Adapting and applying the two-stage architecture to other diagnostic challenges in medical imaging, such as detecting pulmonary nodules on chest CT, classifying brain lesions on MRI, or assessing cardiovascular risk on angiograms.

Using Structured Data for Epidemiological Research: Exploiting the high-fidelity structured data generated by the micro-level system to drive new epidemiological studies. These data could reveal novel correlations between specific pathological features and long-term patient outcomes, opening new avenues for prognostic research.

Finally, the biggest challenge and most promising long-term research direction is the full integration of these three axes into a fully functional, nationwide learning health system. This would involve moving from offline simulations to a model in which the Digital Twin is updated in real time with data from the healthcare system, and public policies can be adjusted dynamically, in a continuous cycle of measurement, modeling and adaptation.

Achieving this vision is undoubtedly a monumental effort, but it is, in essence, the final output of the promise of "Longevity Engineering."

List of Scientific Papers Elaborated During the Doctoral Internship

This thesis proposed an architectural framework for a new paradigm in public health. The research activity carried out during the doctoral internship materialized in a series of scientific papers, published or accepted for publication in prestigious journals and international conferences. These publications are the basis of the original contributions presented in the thesis and demonstrate the validation and dissemination of the results in the scientific community.

Below is the complete list of these papers.

I. Articles in Web of Science (ISI) indexed journals - First Author

1. **Chiş, D. I., Dumitrache, I., Chiş, T., & Caramihai, M. (2025).** Engineering Longevity: A Systems Framework Using Digital Twins and AI to Achieve Sustainable Public Health. Baltica.
 - Status: Article accepted for publication.
 - Indexing: Web of Science (Science Citation Index Expanded).
 - Role: First author.
2. **Chiş, D. I., & Dumitrache, I. (2025).** Engineering Longevity: A Hybrid AI-Driven Simulation Framework for Optimizing National Health and Economic Prosperity. Studies in Informatics and Control.
 - Status: Article accepted for publication.
 - Indexing: Web of Science (Science Citation Index Expanded).
 - Role: First author.
3. **Chiş, D. I., & Dumitrache, I. (2025).** The Ethical AI Triplet: A Framework for Stress-Testing Fairness in Digital Twins in Healthcare. Control Engineering and Applied Informatics (CEAI).
 - Status: Article accepted for publication.

- Indexing: Web of Science (Emerging Sources Citation Index), Scopus, INSPEC.
 - Role: First author.
4. **Chiș, D., & Dumitrache, I.** (2025). Inherently Explainable AI for Automated BI-RADS Classification: A Methodological Proposal. U.P.B. Scientific Bulletin, Series C: Electrical Engineering and Computer Science.
- Status: Article submitted for publication.
 - Indexing: ISI Thomson Reuters, Scopus, INSPEC, Ulrich's, Elsevier Bibliographic Databases.
 - Role: First author.

II. Articles in Web of Science (ISI) indexed journals - Co-author

5. Iliuță, M.-E., Moisescu, M.-A., Caramihai, S.-I., Cernian, A., Pop, E., **Chiș, D.-I., & Mitulescu, T.-C.** (2024). Digital Twin Models for Personalised and Predictive Medicine in Ophthalmology. *Technologies*, 12(4), 55.
- Status: Published.
 - Indexing: Web of Science (Emerging Sources Citation Index), Scopus.
 - Role: Co-author.

III. Other Scientific Publications (Book Chapters and Conference Proceedings)

6. **Chiș, D., & Caramihai, M.** (2023). Blockchain in Higher Education: A Secure Traceability Architecture for Degree Verification. In *Reimagining Education - The Role of E-learning, Creativity, and Technology in the Post-pandemic Era*. IntechOpen.
- Status: Published (Book Chapter).
 - Indexing: IntechOpen Books.
 - Role: First author.
7. Caramihai, M., Severin, I., & **Chiș, D.** (2022). Students' Perception About Online Social Media in Higher Education: An Empirical Study. *Proceedings of the International Symposium on Automation, Information and Computing*.
- Status: Published.
 - Indexing: Conference Proceedings.
 - Role: Co-author.

Bibliography

1. Ahn, A.C., Tewari, M., Poon, C.S., & Phillips, R.S. (2006). The limits of reductionism in medicine: could systems biology offer an alternative? PLoS Medicine, 3(6), e208.
2. Topol, E. J. (2019). Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again. Basic Books.
3. Comisia Europeană. (2024). Impactul schimbărilor demografice în Europa. Preluat de la https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/new-push-european-democracy/impact-demographic-change-europe_ro
4. Barnett, K., Mercer, S.W., Norbury, M., Watt, G., & Wyke, S. (2012). Epidemiology of multimorbidity and implications for health care, research, and medical education: a cross-sectional study. The Lancet, 380(9836), 37-43.
5. Glynn, L.G., Valderas, J.M., Healy, P., Burke, E., Newell, J., Gillespie, P., & Murphy, A.W. (2011). The prevalence of multimorbidity in primary care and its effect on health care utilization and cost. Family Practice, 28(5), 516-523.
6. Bloom, D.E.; Cafiero, E.T.; Jané-Llopis, E.; et al. (2011). The Global Economic Burden of Non-communicable Diseases. World Economic Forum: Geneva, Switzerland.
7. Eurostat. (2024). Healthcare expenditure statistics. Preluat de la https://ec.europa.eu/eurostat/statistics-explained/index.php/Healthcare_expenditure_statistics
8. Lee, R., & Mason, A. (2017). Cost of aging. Finance & Development, 54(1), 8-11.
9. Sturmberg, J. P., & Martin, C. M. (2013). Handbook of systems and complexity in health. Springer Science & Business Media.
10. Meadows, D. H. (2008). Thinking in Systems: A Primer. Chelsea Green Publishing.
11. Kossiakoff, A., Sweet, W. N., Seymour, S. J., & Biemer, S. M. (2011). Systems engineering principles and practice. John Wiley & Sons.
12. INCOSE. (2015). Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities, 4th ed.; John Wiley & Sons: Hoboken, NJ, USA.
13. Chis, D. I., & Dumitrache, I. (2025). Engineering Longevity: A Hybrid AI-Driven Simulation Framework for Optimizing National Health and Economic Prosperity. Studies in Informatics and Control. (Articol trimis spre publicare).

14. Chis, D. I., & Dumitrache, I. (2025). The Ethical AI Triplet: A Framework for Stress-Testing Fairness in Digital Twins in Healthcare. (Articol trimis spre publicare).
15. Chis, D. I., & Dumitrache, I. (2025). Inherently Explainable AI for Automated BI-RADS Classification: A Methodological Proposal. U.P.B. Sci. Bull., Series C. (Articol trimis spre publicare).
16. National Academy of Engineering. (2004). The Engineer of 2020: Visions of Engineering in the New Century. The National Academies Press.
17. Sage, A. P., & Rouse, W. B. (Eds.). (2011). Handbook of systems engineering and management. John Wiley & Sons.
18. von Bertalanffy, L. (1968). General System Theory: Foundations, Development, Applications. George Braziller.
19. Mitchell, M. (2009). Complexity: A guided tour. Oxford University Press.
20. Lorenz, E. N. (1963). Deterministic nonperiodic flow. Journal of the atmospheric sciences, 20(2), 130-141.
21. Holland, J. H. (1998). Emergence: From chaos to order. Oxford University Press.
22. Nightingale, D. J., & Rhodes, D. H. (2015). Engineering complex systems: An enterprise perspective. MIT Press.
23. Wilensky, U., & Rand, W. (2015). An Introduction to Agent-Based Modeling. The MIT Press.
24. Epstein, J. M., & Axtell, R. (1996). Growing Artificial Societies: Social Science from the Bottom Up. Brookings Institution Press.
25. Maier, M. W. (1998). Architecting principles for systems-of-systems. Systems Engineering, 1(4), 267-284.
26. Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. Harvard Business Review, 85(11), 68.
27. Hood, L., & Friend, S. H. (2011). Predictive, personalized, preventive, participatory (P4) cancer medicine. Nature Reviews Clinical Oncology, 8(3), 184-187.
28. Torkamani, A., Wineinger, N. E., & Topol, E. J. (2018). The personal and clinical utility of polygenic risk scores. Nature Reviews Genetics, 19(9), 581-590.

29. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *New England Journal of Medicine*, 380(14), 1347-1358.
30. National Research Council. (2011). *Toward Precision Medicine: Building a Knowledge Network for Biomedical Research and a New Taxonomy of Disease*. National Academies Press.
31. Prasad, V. (2016). The precision-oncology illusion. *Nature*, 537(7619), S63.
32. Topol, E. J. (2015). *The Patient Will See You Now: The Future of Medicine is in Your Hands*. Basic Books.
33. El Saddik, A. (2018). Digital twins: The convergence of multimedia and big data. *IEEE multimedia*, 25(2), 87-92.
34. Björnsson, B., Borrebaeck, C., Elander, N., et al. (2021). Digital twins to personalize medicine. *Genome Medicine*, 13, 1-5.
35. Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In *Transdisciplinary perspectives on complex systems* (pp. 85-113). Springer.
36. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. The MIT Press.
37. Park, J. S., et al. (2023). *Generative Agents: Interactive Simulacra of Human Behavior*. arXiv preprint arXiv:2304.03442.
38. Stiglitz, J. E., Sen, A., & Fitoussi, J. P. (2009). *Report by the Commission on the Measurement of Economic Performance and Social Progress*.
39. Fioramonti, L. (2013). *Gross Domestic Problem: The Politics Behind the World's Most Powerful Number*. Zed Books.
40. Tracy, M., Cerdá, M., & Keyes, K.M. (2018). Agent-based modeling in public health: current applications and future directions. *Annu. Rev. Public Health*, 39, 77-94.
41. Min, L., & Yu, Z. (2025). Intergenerational Information-Sharing Behavior During the COVID-19 Pandemic in China: From Protective Action Decision Model Perspective. *Sustainability*, 17, 7263.
42. Kanis, J.A., et al. (2021). SCOPE 2021: a new scorecard for osteoporosis in Europe. *Archives of Osteoporosis*, 16(1), 82.

43. Epstein, J. M. (2006). *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton University Press.
44. Blanchard, B. S., & Fabrycky, W. J. (2011). *Systems engineering and analysis*. Pearson.
45. Squazzoni, F. (2020). Agent-Based Modeling and Social Simulation in the Era of Big Data. *Sustainability*, 12(11), 4293.
46. United Nations. (2015). *Transforming our world: the 2030 Agenda for Sustainable Development*.
47. Frenk, J., Chen, L., Bhutta, Z. A., et al. (2010). Health professionals for a new century: transforming education to strengthen health systems in an interdependent world. *The Lancet*, 376(9756), 1923-1958.
48. Senge, P. M. (2006). *The fifth discipline: The art and practice of the learning organization*. Broadway Business.
49. Institute of Medicine. (2003). *Health Professions Education: A Bridge to Quality*. National Academies Press.
50. White, G.; et al. (2021). A Digital Twin for a Smart City: A Case Study of the City of Cambridge. *Sustainability*, 13, 8299.
51. Sterman, J. D. (2000). *Business dynamics: Systems thinking and modeling for a complex world*. Irwin/McGraw-Hill.
52. Braveman, P., & Gottlieb, L. (2014). The social determinants of health: it's time to consider the causes of the causes. *Public health reports*, 129(1_suppl_2), 19-31.
53. Peters, D. H., Adam, T., Alonge, O., Agyepong, I. A., & Tran, N. (2013). Implementation research: what it is and how to do it. *BMJ*, 347, f6753.
54. Schultz, T. W. (1961). Investment in Human Capital. *The American Economic Review*, 51(1), 1-17.
55. Grossman, M. (1972). On the Concept of Health Capital and the Demand for Health. *Journal of Political Economy*, 80(2), 223-255.
56. Amodei, D., et al. (2016). Concrete Problems in AI Safety. *arXiv preprint arXiv:1606.06565*.
57. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453.

58. Sele, F., & Chugunova, M. (2024). Putting a human in the loop: Increasing uptake, but decreasing accuracy of automated decision-making. *Journal of Behavioral and Experimental Economics*, 109, 2179.
59. Zajac, M., et al. (2024). Synthetic Population Generation with Public Health Characteristics for Spatial Agent-Based Models. medRxiv.
60. Perez, E., et al. (2022). Red Teaming Language Models to Reduce Harms: Methods, Scaling Behaviors, and Lessons Learned. arXiv preprint arXiv:2209.07858.
61. Beauchamp, T. L., & Childress, J. F. (2013). *Principles of Biomedical Ethics* (7th ed.). Oxford University Press.
62. Black, B., et al. (2022). Averting a "New Equal": The Harms of Individual Arbitrariness in Machine Learning. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*.
63. Financial Conduct Authority. (2017). *Regulatory Sandbox Lessons Learned Report*. London, UK.
64. D'Orsi, C. J., Sickles, E. A., Mendelson, E. B., & Morris, E. A. (2013). *ACR BI-RADS® Atlas, Breast Imaging Reporting and Data System*. American College of Radiology.
65. Elmore, J. G., Wells, C. K., Lee, C. I., et al. (2015). Variability in interpretive performance at screening mammography and radiologists' characteristics associated with accuracy. *Radiology*, 276(3), 681-691.
66. Geller, B. M., Ichikawa, L., Miglioretti, D. L., & Yankaskas, B. C. (2006). Juggling the BIRADS categories: is it time for a new dance? *Radiology*, 239(1), 30-34.
67. Ouchi, K., & Takahashi, H. (2018). Positive predictive value of BI-RADS categories for breast cancer screening. *Breast Cancer*, 25(3), 291-296.
68. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
69. Adadi, A., & Berrada, M. (2018). Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, (6), 52138-52160.
70. Asan, O., Bayrak, A. E., & Choudhury, A. (2020). Artificial Intelligence and Human Trust in Healthcare: A Scoping Review. *Journal of Medical Internet Research*, 22(9), e18513.

71. Gunning, D., & Aha, D. (2019). DARPA's explainable artificial intelligence (XAI) program. *AI Magazine*, 40(2), 44-58.
72. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215.
73. Ivanov, A., & Petrov, V. (2024). Beyond Heatmaps: A Comparative Analysis of Interpretable-by-Design Models in Diagnostic Radiology. *Journal of Medical Imaging AI*, 6(2), 112-125.
74. Pizer, S. M., Amburn, E. P., Austin, J. D., et al. (1987). Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*, 39(3), 355-368.
75. Jocher, G., Chaurasia, A., & Kwon, J. (2023). YOLO by Ultralytics. DOI: 10.5281/zenodo.7638420.
76. Lee, S., & Kim, J. (2024). High-Fidelity Pathological Feature Detection in Mammograms using YOLOv8. *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, 345-349.
77. Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785-794.
78. Trochim, W.M.; Cabrera, D.A.; Milstein, B.; Gallagher, R.S.; Leischow, S.J. (2006). Practical challenges of systems thinking and modeling in public health. *Am. J. Public Health*, 96, 538-546.
79. Chiş, D. I.; Dumitrache, I.; Chiş, T.; Caramihai, M. (2025). Engineering Longevity: A Systems Framework Using Digital Twins and AI to Achieve Sustainable Public Health. *Baltica*. (Articol acceptat spre publicare).
80. World Health Organization. (2023). Dementia Fact Sheet. Preluat de la <https://www.who.int/news-room/fact-sheets/detail/dementia>.
81. Joyner, M.J.; Prendergast, F.G. (2014). Chasing unicorns: the deception of personalized medicine. *J. Physiol.*, 592, 2431-2439.
82. Holland, J. H. (2014). *Complexity: A very short introduction*. Oxford University Press.
83. Bar-Yam, Y. (2019). *Dynamics of complex systems*. CRC Press.

84. Kauffman, S. A. (1995). *At home in the universe: The search for the laws of self-organization and complexity*. Oxford University Press.
85. Prigogine, I., & Stengers, I. (1984). *Order out of chaos: Man's new dialogue with nature*. Bantam Books.
86. Gilbert, N. (2008). *Agent-based models*. Sage.
87. World Health Organization. (2021). *The Health Systems Response to the COVID-19 Pandemic*. Geneva.
88. Plsek, P.E.; Greenhalgh, T. (2001). Complexity science: The challenge of complexity in health care. *BMJ*, 323, 625-628.
89. World Health Organization. (1994). *Assessment of fracture risk and its application to screening for postmenopausal osteoporosis. Report of a WHO Study Group*. WHO technical report series, 843.
90. Ganda, K., Puech, M., et al. (2013). Models of care for the secondary prevention of osteoporotic fractures: a systematic review and meta-analysis. *Osteoporosis International*, 24(2), 393-406.
91. Russell, S. J., & Norvig, P. (2021). *Artificial Intelligence: A Modern Approach*. 4th ed. Pearson.
92. Burge, R., Dawson-Hughes, B., Solomon, D. H., et al. (2007). Incidence and economic burden of osteoporosis-related fractures in the United States, 2005-2025. *Journal of Bone and Mineral Research*, 22(3), 465-475.
93. Caramihai, S. I., Dumitrache, I., Moisescu, M. A., & Sacala, I. S. (2022). Decision Support Collaborative Platform for e-Health Integration in Smart Communities Context. *Procedia Computer Science*, 214, 1152-1159.
94. Diez Roux, A. V. (2007). Integrating social and biologic factors in health research: a systems view. *Annals of Epidemiology*, 17(7), 569-574.
95. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. *Proceedings of the 3rd innovations in theoretical computer science conference*, 214-226.
96. Iliuță, M.-E., Moisescu, M.-A., Caramihai, S.-I., Cernian, A., Pop, E., Chiș, D.-I., & Mitulescu, T.-C. (2024). Digital Twin Models for Personalised and Predictive Medicine in Ophthalmology. *Technologies*, 12(4), 55.

97. Aron, David C. Complex Systems in Medicine: A Hedgehog's Tale of Complexity in Clinical Practice, Research, Education, and Management.