

# Engineering Longevity: A Hybrid AI-driven Simulation Framework for Optimizing National Health and Economic Prosperity

Daniel Ioan CHIS<sup>1</sup>, Ioan DUMITRACHE<sup>2</sup>

<sup>1</sup> National University of Science and Technology Politehnica Bucharest, 313 Splaiul Independentei, 060042, Bucharest, Romania  
chisdanielioan@gmail.com (\*Corresponding author)

<sup>2</sup> Romanian Academy, 125 Calea Victoriei, 010071, Bucharest, Romania  
ioan.dumitrache@acad.ro

**Abstract:** The non-linear relationship between population health and economic prosperity represents a new frontier in modelling and optimization. This paper presents a novel hybrid computational framework which was designed for crafting optimal health policies through empirically validated simulations, integrating a top-down macroeconomic framework with a bottom-up agent-based simulation (ABM) for a dynamic, heterogeneous population. The proposed framework employs a robust, four-step pipeline. First, an agent-based model generates a large dataset, which will train a high-speed surrogate model. Further on, an AI-based optimization engine leverages this surrogate model in order to efficiently discover the optimal policy strategies, which are subsequently validated with regard to their statistical robustness. This framework was validated against a post-menopausal osteoporosis case study, identifying a policy that significantly outperforms the benchmark strategy across a ten-run validation process. To sum up, this paper provides a new, empirically grounded blueprint for a “Digital Twin” for public health, combining complex systems modelling with AI-driven optimization to help engineer longevity at a national level.

**Keywords:** Agent-based simulation, AI-driven, Optimization.

## 1. Introduction

While life sciences have dramatically extended the human lifespan, a new challenge has emerged: increasing the population’s healthspan, the period of life spent in good health, in an economically optimized manner. In many developed societies, life expectancy now surpasses 80 years, but this longevity is accompanied by a rising tide of non-communicable diseases. This trend, compounded by declining birth rates, is creating a looming demographic and economic crisis where a shrinking active population must support a growing elderly population that requires a majority of the healthcare resources.

At the individual level, the objective is not merely to live longer but to improve one’s quality of life by reducing the time spent with disabilities or chronic conditions.

Engineering longer, healthier, and more productive lives stands as one of this century’s greatest challenges.

The established Human Capital Theory (Schultz, 1961) affirms the close link between a nation’s health and its economic prosperity. A healthier population is more productive, innovative, and resilient, which in turn drives long-term economic growth. This growth stems from individuals working for more years with a greater productivity,

while their associated healthcare costs are both reduced and delayed.

Despite this understanding, modeling and optimizing this complex relationship to guide strategic investment remains a difficult problem. The central challenge lies in its complexity: how a society with finite resources can best invest in health to reduce the burden of chronic diseases (measured in DALYs), increase the healthy life expectancy (HALE), and, in doing so, create a virtuous cycle of sustainable economic growth?

Traditional linear models are insufficient for this purpose because they cannot capture the feedback loops, non-linearities, and emergent behaviors that characterize the interplay between a nation’s economy and the health of its population (Diez Roux, 2011).

To address this gap, this paper proposes and validates a novel framework designed for modelling this complex adaptive system.

The remainder of this paper is organized as follows. Section 2 provides a review of the current state of the art, while Section 3 showcases the proposed theoretical models for the national health-economy landscape as a single, integrated entity. Section 4 details the novel methodology,

including the AI-driven simulation pipeline and a case study on post-menopausal osteoporosis. Subsequently, Section 5 analyzes the broader implications of this work, proposing a “Digital Twin for Public Health”. Finally, Section 6 provides the concluding remarks.

## 2. State of the Art

Despite rapid advancements, the contemporary health system solutions remain largely fragmented, reflecting a reductionist approach that is ill-equipped for handling systemic complexity (Ahn et al., 2006).

This work breaks down traditional silos by integrating principles from complex systems science, computational modeling, and AI-driven policy optimization (Chiş et al., 2025).

AI in clinical practice: Artificial intelligence, particularly deep learning, has proven to be highly effective in clinical settings, assisting with diagnosis and individual risk prediction (Russell & Norvig, 2021). However, these tools are typically designed for narrow clinical tasks, such as analyzing medical images, and are not integrated into strategic, population-level resource allocation (van der Schaar et al., 2021). Furthermore, the “black-box” nature of many models creates barriers to clinical trust (Ghassemi et al., 2020). This approach feeds high-quality, structured data from the emerging ‘interpretable-by-design’ frameworks directly into the proposed macro-level models.

Computational modeling in epidemiology and health economics: Fields like computational epidemiology and health economics have embraced dynamic modeling. Agent-Based Models (ABMs) effectively simulate disease transmission and behavioral interventions, while Markov models are the standard for cost-effectiveness analyses (Epstein, 2009; Ganda et al., 2013). However, these methods are often limited to assessing specific interventions and are not designed to explore a continuous landscape of policy strategies or model the dynamic feedback between health and economic outcomes (Rutter et al., 2017). The validation of such complex models against real-world data also remains a critical challenge (Wilensky & Rand, 2015).

Digital twins for the health policy: The concept of “Digital Twin” - a virtual replica of a physical system - is an emerging paradigm in healthcare, with applications ranging from modeling human organs to optimizing hospital workflows (Sturmberg & Martin, 2013; Caramihai et al., 2022). While the vision of a “Digital Twin for Public Health” is gaining traction for urban planning and pandemic preparedness (Pammi et al., 2025; Kou et al., 2021), these applications rarely scale to the national level so as to co-optimize health and economic outcomes. Moreover, population-level digital twins raise significant ethical challenges regarding data privacy, consent, and governance that must be proactively addressed (Kamel Boulos & Zhang, 2021).

AI, ethics, and health economics: The use of AI for resource allocation necessitates a robust ethical framework. A primary concern is the risk of algorithmic bias entrenching or amplifying the existing health disparities (Obermeyer et al., 2019; Leslie, 2019). This has fuelled critical research into fairness, accountability, and transparency in digital health (Morley et al., 2020). This work aligns with their research by proposing an “Ethical Sandbox” to proactively stress-test the AI-generated policies for fairness, moving beyond a reactive “human-in-the-loop” model (Chiş & Dumitrache, 2025). This validation is essential for building public trust and ensuring that AI policy tools serve the entire population justly (Jobin et al., 2019).

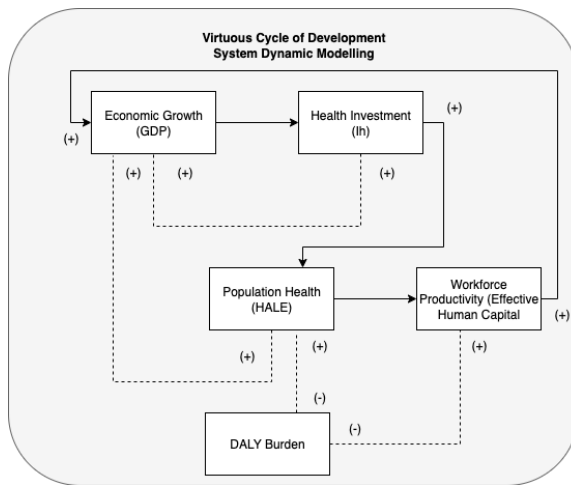
Therefore, this work proposes a unified framework for bridging these fields by integrating a bottom-up, epidemiologically-grounded ABM with a top-down macroeconomic structure, all orchestrated by an AI optimization engine. This framework creates a population-level Digital Twin designed not merely for analysis purposes, but for a proactive and ethically-informed policy discovery.

## 3. A Complex Systems Framework for the Society

The proposed approach is founded on the principles of complex adaptive systems (CAS), considering the national health-economy landscape as a single, integrated entity (Diez Roux, 2011; Galea et al., 2010) (Figure 1).

This framework is defined by the following core characteristics:

- **Emergence:** Macro-level phenomena, such as GDP growth or shifts in life expectancy, aren't the result of a central plan but are the emergent outcomes of innumerable micro-level interactions among individuals and organizations (Holland, 1998);
- **Agent heterogeneity:** The system is composed of diverse, autonomous agents (e.g. citizens, healthcare providers, policymakers), whose varied attributes and behaviors drive the system's overall dynamics;
- **Non-linearity:** The relationship between actions and outcomes is often disproportionate. A small, strategic policy change can lead to substantial, system-wide transformations over time (Plsek & Greenhalgh, 2001);
- **Feedback loops:** The system's behavior is regulated by feedback mechanisms. The focus is on modeling the "Virtuous Cycle of Development," a positive feedback loop where health investments improve the economic productivity, which then enables further investment in health (Burge et al., 2007).



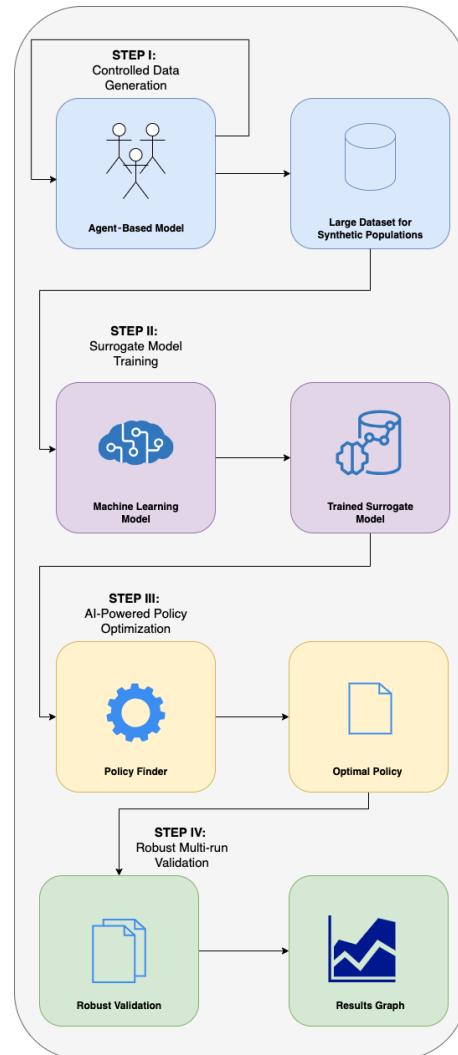
**Figure 1.** The "Virtuous Cycle of Development" Framework (Chiş et al., 2025)

This diagram illustrates the core theoretical model, showing how Economic Growth provides resources for Health Investment, which improves Population Health. This, in turn, boosts the Workforce Productivity, fueling further Economic Growth and creating a positive feedback loop.

#### 4. Methodology: A Hybrid AI-driven Simulation Pipeline

To empirically validate the proposed theoretical framework, a multi-stage computational pipeline

was engineered. This section provides a detailed account of its technical implementation and statistical validation. The pipeline, depicted in Figure 2, serves to transform policymaking from a static analysis into a dynamic discovery process.



**Figure 2.** The AI-Driven Four-Step Simulation and Optimization Pipeline; this diagram illustrates a revised, empirically-focused methodology (Chiş et al. 2025)

#### 4.1 The hybrid AI-driven Simulation Pipeline

The proposed pipeline includes four distinct steps:

1. **Controlled Data Generation:** To ensure that the obtained results are not skewed by random variations, first a single, consistent agent population is generated. A large number of ABM simulations ( $N = 2,000$  for this study) are then run, testing a different policy strategy in each run against an identical copy of this initial population. This process creates a clean, high-quality dataset where the differences between outcomes can be directly attributed to the policy being tested;

2. **Surrogate Model Training:** Because full ABM simulations are computationally slow, the dataset from Step I is used for training a high-speed surrogate model. This machine learning model (specifically, a XGBoost regressor (Chen & Guestrin, 2016)) learns to approximate the full simulation results. It can predict the final Wellbeing Score for a given policy in milliseconds, allowing for a rapid exploration;
3. **AI-Powered Policy Optimization:** With this fast surrogate model, the vast landscape of possible policies can be efficiently explored. An AI optimization engine (using Bayesian optimization) rapidly queries the surrogate model thousands of times to identify the policy predicted to yield the highest National Wellbeing Score;
4. **Robust Multi-run Validation:** The policy identified by AI is a promising candidate, but it must be validated. This candidate policy is run through the full, slow ABM simulation 10 times, each time with a new, randomly generated initial population. The final performance is then reported as the mean and standard deviation across these runs, providing a statistically stable measure of the strategy's expected real-world performance.

## 4.2 Macro-Level Simulation Engine

The system is defined by equations that link economic policies to health outcome metrics. The standard Cobb-Douglas Production Function is adapted by introducing an “Effective Human Capital” term, which incorporates public health metrics directly into the economic model. This is inspired by the Grossman model, which treats health as a capital asset (Grossman, 1972).

- **Population Income Identity & Investment:** This equation showcases the fundamental economic trade-off. The National Output ( $Y_t$ ) can be expressed as the sum between the Consumption ( $C_t$ ) component, the investment in traditional Physical Capital ( $I_k$ ), and the strategic Investment in Health ( $I_h$ ).

$$Y_t = C_t + I_k + I_h \quad (1)$$

- **Future Economic Output:** This equation models the core hypothesis of the proposed framework. The future economic output  $Y_{t+1}$  is a function of „Effective Human Capital“. The output is directly improved by the return on health investment, denoted as  $g(I_h)$ , which represents the total reduction in DALYs achieved by the health policy.  $T$  represents the technological factor, while  $K$  and  $I_k$

represent the initial capital and the investment in traditional Physical Capital, respectively.

These factors are assumed not to be directly influenced by the health policy.

$$Y_{t+1} = f(K_t + I_k, T, (P_{total} \cdot (\frac{HALE_t + g(I_h)}{LE}))) \quad (2)$$

- **Return on Health Investment ( $g(I_h)$ ):** The ABM component computes the true return on an investment by simulating the cascading benefits. This “multiplier effect of investing in health” represents the total, system-wide reduction in *DALY*s across the entire disease network ( $k$ ) resulting from a targeted investment in a single area ( $j$ ):

$$g(I_{h,j}) = \sum_k \Delta DALY_k \quad (3)$$

### Population Wellbeing Score (The Goal of AI):

This equation is the primary optimization metric used for the carried out simulation. It represents a utility function that balances economic prosperity with public health, rewarding the strategies that increase the number of healthy life years but also economic efficiency, while reducing the burden of chronic diseases:

$$WellbeingScore = AVGHealthyYears - DALYperCapita \times w_1 + ScaledGDPperCapita \times w_2 \quad (4)$$

where  $w_1$  and  $w_2$  are weighting factors.

- **Health Outcome Metrics (DALY & HALE):** These are the standard public health measures calculated by the model:

$$DALY = \sum YLL + \sum YLD \quad (5)$$

$$HALE = Age_{start} + \frac{\sum Healthy_{years}}{N_{population} \cdot T_{simulation}} \quad (6)$$

## 4.3 Micro-Level Simulation Engine

At the core of the proposed simulation pipeline there is an Agent-Based Model that simulates a diverse population over time, based on system dynamics principles (Epstein & Axtell, 1996; Wilensky & Rand, 2015).

- **Population Dynamics:** The simulation starts with a realistic population. Each year, the model simulates births and deaths to reflect the natural population changes;
- **Detailed Agent Logic:** Agents go through different health states (e.g. healthy, at-risk, diseased, deceased) based on probabilities derived from real-world epidemiological



data. An agent's economic productivity is directly linked to their current health;

- **Policy Implementation:** The health investment (Ih) directly influences the model by changing these health state probabilities. For instance, funding for lifestyle campaigns reduces the chance that healthy agents will transition to an at-risk state.

#### 4.4 Case study: Optimizing Post-menopausal Osteoporosis Management

In order to validate the proposed framework, it was applied to post-menopausal osteoporosis, a chronic disease affecting around 20% of the women over 50 in the EU (Kanis et al., 2021). As a “silent disease” that is often asymptomatic until a fracture occurs, it is an ideal test case for proactive, preventive strategies (Compston, 1995).

##### 4.4.1 Model Configuration

- **Population:** 1,000 female agents, aged 51-80 at the start;
- **Time Horizon:** 35 years;
- **Policy Vector:** The AI optimizer allocated the health budget across three pillars: Lifestyle campaigns (Primary Prevention), Screening programs (Secondary Prevention), and Treatment (Tertiary Prevention).

##### 4.4.2 Key Findings and Statistical Validation

After running the full pipeline, the AI optimizer discovered a non-intuitive optimal strategy: dedicating 85% of the investment to lifestyle campaigns, 5% to screening, and 10% to treatment.


To validate this strategy, a 10-run validation was performed (Step IV of the simulation). Table 1 compares the averaged results obtained for the four employed metrics below with regard to the benchmark scenario (no new investment) and the AI-Identified Scenario. The results are rendered as a mean  $\pm$  standard deviation.

The statistical results confirm that the AI-identified strategy is both superior to the benchmark scenario and stable.


It produced a 5.78% increase in the Wellbeing Score and an 18.09% reduction in the total burden of disease (DALYs).

This health improvement powers the “Virtuous Cycle,” leading to a statistically significant 7.45% increase in GDP over the 35-year simulation period.

Figures 3, 4 and 5 show the averaged results from the 10-run validation, visually confirming the proposed strategy's consistent ability to reduce mortality and boost the long-term economic output.



**Figure 3.** The mortality rate for the agent population (Avg. over 10 runs); this graph shows that the AI-identified strategy (the green line) leads to a consistently lower mortality over the long term in comparison with the baseline model




**Figure 4.** The population size over time (Avg. over 10 runs); as a direct consequence of lower mortality, the total population size becomes larger under the optimized health policy

**Table 1.** Key Metrics Comparison in the Benchmark Scenario and the AI-Identified Scenario

Metric	Benchmark Scenario	AI Identified Scenario	Improvement (%)
Wellbeing Score	42.35 $\pm$ 0.45	44.80 $\pm$ 0.38	+5.78%
Total DALYs (cumulative)	1636.96 $\pm$ 55.2	1340.84 $\pm$ 41.9	-18.09%
HALE (years)	43.03 $\pm$ 0.21	44.83 $\pm$ 0.19	+4.18%
Final GDP (€)	(2.28 $\pm$ 0.05)M	(2.45 $\pm$ 0.04) M	+7.45%

The 10-year lag before the GDP curves diverge provides a critical insight, illustrating that preventive health investments require patience before yielding economic returns.



**Figure 5.** GDP Trajectory Comparison (Avg. over 10 runs); this graph clearly illustrates the economic benefits of the health policy

The improved population health as shown in Figures 3 and 4 translates into a more productive workforce, leading to a superior long-term economic growth and demonstrating the “Virtuous Cycle” in action.

The AI-identified strategy, with its heavy emphasis on upstream lifestyle interventions, proved to be the most effective. This demonstrates that preventive investments can generate the highest long-term systemic return.

## 5. A “Digital Twin” for Public Health

In order to turn the analysed framework into a practical tool, this paper proposes a cloud-based platform that acts as a “Digital Twin for Public Health,” which would serve as an adaptable “sandbox” for policymakers.

The proposed platform builds on the existing e-Health concepts (Caramihai et al., 2022), but it scales them to the national population level for supporting long-term strategic planning.

The proposed multi-layer architecture, illustrated in Figure 6, is designed for scalability and modularity and it includes four layers:

- **User Interface (UI) Layer:** This is the primary access point for policymakers. The UI includes a Configuration Module for defining the simulation parameters (e.g. disease, population demographics, policy goals etc.), a Simulation Panel for launching and monitoring experiments and a Results

Dashboard for visualizing and comparing the outcomes of different policy strategies in an intuitive manner;

- **Cloud Backend (Application Logic Layer):** This layer acts as the system’s brain, orchestrating all operations. The AI Policy Optimizer manages the entire optimization workflow. The Simulation Controller interfaces with the simulation engine for executing the experiments, while the History and Simulations Controller logs all simulation runs in a database for auditing, reproducibility, and comparative analysis purposes;
- **Core Engines (AI Models & the Simulation Engine):** This layer contains the core computational components. The Simulation Engine runs the ABM, including the Population Dynamics Engine, to generate the raw simulation data. The AI Models sublayer includes the Health Policy Optimizer, which leverages various ML Algorithms (like the XGBoost surrogate model) for analyzing the data and discovering the optimal strategies;
- **Data Layer:** The foundation of the entire system, this layer aggregates and manages all the necessary data, including the baseline Model Data (epidemiological parameters), the dynamic Agent and Population Data, and specific Disease Data (e.g. transition probabilities, costs etc.).

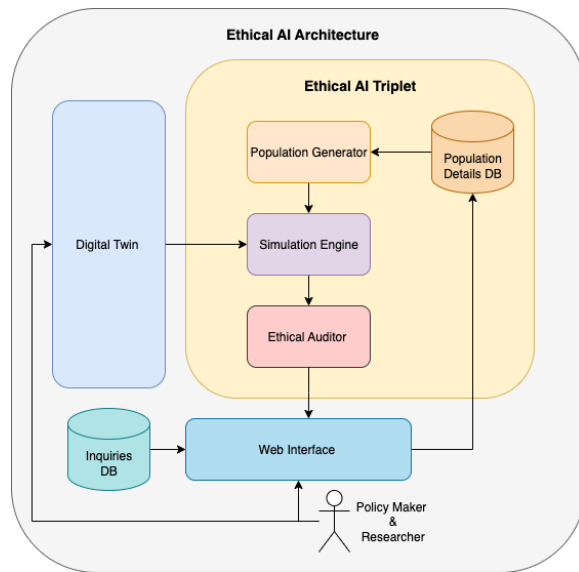
This layered architecture ensures that the system is both powerful and accessible, allowing policymakers to leverage complex simulations without needing a deep technical expertise.

### 5.1 Conceptual and Model Limitations

First, any simulation is necessarily a simplification of reality. The proposed ABM, although detailed, cannot fully model the richness of human behavior — such as cultural beliefs or social cohesion — which significantly influences the health outcome metrics. Second, the model’s parameters were derived from the existing literature, and their applicability may vary across different populations. Finally, the “National Wellbeing Score” is itself a value-laden construct; the choice of weights is a normative decision that requires public deliberation, a process that the proposed framework is designed to inform, not automate.

## 5.2 Technical and Data Limitations

The validity of this framework is critically dependent on high-quality, granular data for calibration - a “garbage in, garbage out” constraint. In many real-world settings, such data is fragmented or unavailable (Tracy et al., 2018). Furthermore, the computational cost of generating the initial training data, while a one-time effort for a given problem, remains significant and requires high-performance computing.



**Figure 6.** Digital twin architecture related to population health and economic prosperity

## 5.3 Implementation and Ethical Considerations

A primary ethical concern is the algorithmic bias. If the calibration data reflects the existing health disparities, the AI could amplify those inequities (Obermeyer et al., 2019; Leslie,

2019). This risk underscores the need for a robust ethical governance layer and internal auditing mechanisms (Raji et al., 2020), such as the “Ethical AI Triplet” framework proposed in the related work of Chiş & Dumitrache (2025), to stress-test policies for fairness. Ultimately, deploying such a Digital Twin is not just a technical challenge but also a socio-political one, requiring a sustained political will, inter-agency collaboration, and public trust.

Despite the aforementioned limitations, the proposed framework represents a significant step forward. By making the trade-offs between health, economics, and equity explicit and quantifiable, it can transform policymaking into a more transparent, evidence-based, and deliberative process.

## 6. Conclusion

This paper presented a novel, consolidated approach to public health policy design, namely an empirically validated methodology which combined a complex systems perspective with an AI-powered agent-based modeling technique. The post-menopausal osteoporosis case study demonstrated the proposed framework’s capability to uncover non-intuitive, high-leverage policies that can improve population health while driving economic prosperity. While significant challenges remain, this new hybrid, multi-scale approach provides a blueprint for a new generation of data-driven, evidence-based tools. In this sense, it shows how experts could engineer a healthier and more prosperous future by integrating complex health and economic data in a robust simulation environment.

## REFERENCES

- Ahn, A.C., Tewari, M., Poon, C.-S. et al. (2006) The Limits of Reductionism in Medicine: Could Systems Biology Offer an Alternative? *PLoS Medicine*. 3(6), Art. ID. e208. <https://doi.org/10.1371/journal.pmed.0030208>.
- 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. <https://doi.org/10.1359/jbmr.061113>.
- Caramihai, S. I., Dumitrache, I., Moisescu, M. A. et al. (2022) Decision Support Collaborative Platform for e-Health Integration in Smart Communities Context. *Procedia Computer Science*. 214, 1152–1159. <https://doi.org/10.1016/j.procs.2022.11.290>.
- Pammi, M., Shah, P. S., Yang, L. K. et al. (2025) Digital twins, synthetic patient data, and in-silico trials: can they empower paediatric clinical trials?, *The Lancet Digital Health*. 7(5), Art. ID 100851. <https://doi.org/10.1016/j.landig.2025.01.007>.
- Chen, T. & Guestrin, C. (2016) XGBoost: A Scalable Tree Boosting System. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*, 13–17 August 2016, San Francisco, California. New York, Association for Computing Machinery. pp. 785–794.

- Chiş, D. I. & Dumitrache, I. (2025) Ethical AI Triplet: A Framework for Stress-Testing Fairness in Digital Twins in Healthcare. *Control Engineering and Applied Informatics (CEAI)*. 27(3). <https://doi.org/10.61416/ceai.v27i3.9745>.
- Chiş, D. I., Dumitrache, I., Chiş, T. et al. (2025) Engineering Longevity: A Multi-Scale Systems Framework Using Digital Twins and AI to Achieve Sustainable Public Health. *Baltica*. 38(10), 2-25. <https://doi.org/10.61586/77dMe>.
- Compston, J. (1995) Assessment of fracture risk and its application to screening for postmenopausal osteoporosis (WHO Technical Report Series No 843). *Annals of the Rheumatic Diseases*. 54(7), 548. <https://doi.org/10.1136/ard.54.7.548>.
- Diez Roux, A. V. (2011) Complex Systems Thinking and Current Impasses in Health Disparities Research. *American Journal of Public Health*. 101(9), 1627-1634. <https://doi.org/10.2105/AJPH.2011.300149>.
- Epstein, J. M. & Axtell, R. L. (1996) *Growing Artificial Societies: Social Science from the Bottom Up*. Washington, Brookings Institution Press.
- Epstein, J. M. (2009) Modelling to Contain Pandemics. *Nature*. 460, Art. ID. 687. <https://doi.org/10.1038/460687a>.
- Galea, S., Riddle, M. & Kaplan, G. A. (2010) Causal thinking and complex system approaches in epidemiology. *International Journal of Epidemiology*. 39(1), 97-106. <https://doi.org/10.1093/ije/dyp296>.
- Ganda, K., Puech, M., Chen, J. S. 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. <https://doi.org/10.1007/s00198-012-2090-y>.
- Ghassemi, M., Naumann, T., Schulam, P. et al. (2020) A Review of Challenges and Opportunities in Machine Learning for Health. In: *AMIA Joint Summits on Translational Science Proceedings*. 2020, 191.
- Grossman, M. (1972) On the Concept of Health Capital and the Demand for Health. *Journal of Political Economy*. 80(2), 223-255. <https://doi.org/10.1086/259880>.
- Holland, J. H. (1998) *Emergence: From Chaos to Order*. Oxford, Oxford University Press.
- Jobin, A., Ienca, M. & Vayena, E. (2019) The global landscape of AI ethics guidelines. *Nature Machine Intelligence*. 1(2), 389-399. <https://doi.org/10.1038/s42256-019-0088-2>.
- Kanis, J.A., Norton N, Harvey N. C. et al. (2021) SCOPE 2021: a new scorecard for osteoporosis in Europe. *Archives of Osteoporosis*. 16(1), 82. <https://doi.org/10.1007/s11657-020-00871-9>.
- Kou, L., Wang, X., Li, Y. et al. (2021) A multi-scale agent-based model of infectious disease transmission to assess the impact of vaccination and non-pharmaceutical interventions: The COVID-19 case. *Journal of Safety Science and Resilience*. 2(4), 199-207. <https://doi.org/10.1016/j.jnlssr.2021.08.005>.
- Leslie, D. (2019) *Understanding artificial intelligence ethics and safety*. The Alan Turing Institute.
- Morley, J., Machado, C. C. V., Burr, C. et al. (2020) The ethics of AI in health care: A mapping review. *Social Science & Medicine*. 260, Art. ID 113172. <https://doi.org/10.1016/j.socscimed.2020.113172>.
- Obermeyer, Z., Powers, B., Vogeli, C. et al. (2019) Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 366(6464), 447-453. <https://doi.org/10.1126/science.aax2342>.
- Plsek, P. E. & Greenhalgh, T. (2001) Complexity science: The challenge of complexity in health care. *BMJ*. 323(7313), 625-628. <https://doi.org/10.1136/bmj.323.7313.625>.
- Raji, I. D., Smart, A., White, R. N. et al. (2020) Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT\* '20)*, 27 – 30 January 2020, Barcelona, Spain. New York, Association for Computing Machinery pp. 33-44.
- Rutter, H., Savona, N., Glonti, K., et al. (2017) The need for a complex systems model of evidence for public health. *The Lancet*. 390(10112), 2602-2604. [https://doi.org/10.1016/S0140-6736\(17\)31267-9](https://doi.org/10.1016/S0140-6736(17)31267-9).
- Russell, S. J. & Norvig, P. (2021) *Artificial Intelligence: A Modern Approach*. 4th ed. New Jersey, Pearson Education, Inc..
- Schultz, T. W. (1961). Investment in Human Capital. *American Economic Review*. 51(1), 1-17. <https://www.jstor.org/stable/1818907>.
- Sturmberg, J. P. & Martin, C. M. (2013) *Handbook of Systems and Complexity in Health*. New York, Springer Science & Business Media.
- Tracy, M., Cerdá, M. & Keyes, K. M. (2018) Agent-Based Modeling in Public Health: Current Applications and Future Directions. *Annual Review of Public Health*. 39, 77-94. <https://doi.org/10.1146/annurev-publhealth-040617-014317>.
- van der Schaar, M., Alaa, A.M., Floto, A. et al. (2021) How artificial intelligence and machine learning can help healthcare systems respond to COVID-19. *Machine Learning*. 110(1), 1-14. <https://doi.org/10.1007/s10994-020-05928-x>.
- Wilensky, U. & Rand, W. (2015) *An Introduction to Agent-Based Modeling: Modeling Natural, Social, and Engineered Complex Systems with NetLogo*. Cambridge, USA, MIT Press.
- Kamel Boulos, M.N. & Zhang P. (2021) Digital Twins: From Personalised Medicine to Precision Public Health. *Journal of Personalized Medicine*. 11(8), Art. ID 745. <https://doi.org/10.3390/jpm11080745>.





This is an open access article distributed under the terms and conditions of the Creative Commons Attribution-NonCommercial 4.0 International License.