

Re-imagining health and well-being in low resource African settings using an augmented AI system and a 3D digital twin

Deshendran Moodley^{1,2,*,†}, Christopher Seebregts^{1,2,3,†}

¹University of Cape Town, 18 University Avenue, Rondebosch, Cape Town, South Africa

²Centre for Artificial Intelligence Research, Cape Town, South Africa

³Jembi Health Systems, Tokai, Cape Town, South Africa

Abstract

This paper discusses and explores the potential and relevance of recent developments in artificial intelligence (AI) and digital twins for health and well-being in low-resource African countries. We use the case of public health emergency response to disease outbreaks and epidemic control. There is potential to take advantage of the increasing availability of data and digitization to develop advanced AI methods for analysis and prediction. Using an AI systems perspective, we review emerging trends in AI systems and digital twins and propose an initial augmented AI system architecture to illustrate how an AI system can work with a 3D digital twin to address public health goals. We highlight scientific knowledge discovery, continual learning, pragmatic interoperability, and interactive explanation and decision-making as essential research challenges for AI systems and digital twins.

Keywords

Augmented AI, AI in Health, Digital Twin, PHEOC, Pragmatic Interoperability

1. Introduction

Public health aims to prevent disease, promote health, and prolong life among the population. In the Alma-Ata declaration [1], the World Health Organization (WHO) reaffirmed that attaining the highest possible level of health is the most significant worldwide social goal. Public health programs strive to optimise the conditions in which people can achieve health and well-being.

Low-resource countries bear much of the world's disease burden and outbreaks [2] [3]. They also lack essential infrastructure and need to maximize the utility of the available resources. Despite these challenges, African and low-resource countries have made admirable progress in digitising health systems, resulting in a significant number of surveillance and routine health information systems. This opens up novel opportunities to leverage advanced computing techniques such as AI to advance health and well-being in low-resource African countries

Submitted to Workshop on AI for Digital Twins and Cyber-physical applications at IJCAI 2023, August 19–21, 2023, Macau, S.A.R

*Corresponding author.


✉ deshen@cs.uct.ac.za (D. Moodley); chris@jembi.org (C. Seebregts)

🌐 <https://www.cs.uct.ac.za/~deshen> (D. Moodley); <https://www.jembi.org> (C. Seebregts)

🆔 0000-0002-4340-9178 (D. Moodley); 0000-0002-8212-9363 (C. Seebregts)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

[4][5].

Modern approaches to health systems and public health in low-resource settings increasingly focus on preventative in addition to curative health approaches. This leads to a more holistic perspective on health and well-being. However, preventative health requires a better understanding of the factors that cause and influence disease onset and the positive and negative determinants of disease and adverse health conditions. Understanding and promoting healthy behaviours and lifestyles is a primary goal of digital health systems in African countries.

1.1. One Health and epidemic surveillance

The WHO recommends a One Health approach [6] to design and implement health programmes, policies, legislation and research. The One Health approach proposes an integrated, unifying approach to balance and optimize the health of people, animals and the environment. It is essential to prevent, predict, detect, and respond to global health threats and emergencies such as the recent COVID-19 pandemic [7]. The One Health approach mobilizes multiple sectors, disciplines and communities at varying levels of society to work together to achieve better public health outcomes. This way, new and better ideas are developed that address root causes more holistically and create long-term, sustainable solutions to predict adverse health conditions before they occur. Supporting the One Health approach presents new challenges to digital health systems [8]. Public health decision-makers integrate several different data sources from diverse digital and paper-based systems to assist with complex decision-making around outbreak management and epidemic control. These include routine health data systems, surveillance systems, household and bio-behavioural surveys, and research data. If it is paper-based, the data must first be digitised, then transformed and compiled into dashboards to optimize public health situation detection, response and prediction. A One Health digital platform must deal with new levels of interoperability between different jurisdictions and disciplines to connect vast amounts of heterogeneous data and systems. The platform must support real-time situation analysis, predictive modelling, and proactive decision-making to reach its full potential.

1.2. Health monitoring and digital data management

Surveillance of public health threats is one of the main functions of public health, e.g. in a public health emergency operations centre (PHEOC) [9] [10]. Typical activities include monitoring public health threats in a country, e.g. surveillance of public health threats followed by situation analysis and decision-making [11]. Similar challenges are encountered in other vertical disease areas, such as the attainment of HIV Epidemic Control [12] [13]. HIV Epidemic Control Rooms have been developed to monitor progress towards achieving goals and performance targets. Traditional information and data sources are changing due to digitization and shifting emphasis to using information from routine health information systems to derive, monitor and manage epidemiological targets [14].

Data management processes supporting monitoring, such as semantic annotation and data integration, are established in practice. In the case of situation analysis, well-established surveillance methods support detection. Support for prediction from digital health and AI is emerging with examples such as the digital predictive tool developed for forecasting the

occurrence of diseases based on historical weather and health data in Uganda [4].

1.3. Relevance and potential of AI and digital twins

Powerful data-driven AI models can be used for dynamic data fusion, situation analysis, predictive modelling, and scientific knowledge discovery. These are all valuable functions in public health and can lead to a new understanding of complex dynamic processes. An AI system working in parallel with a digital twin can augment and amplify interactive decision-making and scientific knowledge discovery in public health decision-making.

In this position paper, we explore the relevance and potential of artificial intelligence and digital twins as two key emerging technologies that push the boundaries of the One Health vision in low-resource public health systems. Our perspective is informed by over a decade of research and real-world implementation of advanced digital health technologies in multiple African countries [15]. Based on our experience, we propose an abstract architecture for an augmented AI system incorporating a Digital Twin, highlighting the potential for AI and digital twins to realise the One Health vision. We focus on the potential of AI systems and digital twins to support scientific knowledge discovery, interactive decision making and pragmatic interoperability.

The paper is structured as follows. In section 2, we describe digital twins in health care and then review emerging trends and critical challenges in AI systems in section 3. Section 4 provides a preliminary proposal for an abstract architecture for an augmented AI system incorporating a digital twin. In section 5, we offer a brief discussion and conclusions.

2. Digital Twins in Health Care

A digital twin is a virtual model of a physical entity, with dynamic, bi-directional links between the physical entity and its corresponding twin in the digital domain [16]. While they have been traditionally explored in the manufacturing sector, digital twins are increasingly being explored in medicine and health care [17, 18, 19]. Applied to medicine and public health, digital twin technology has been proposed to drive a much-needed radical transformation of traditional electronic health/medical records (focusing on individuals) and their aggregates (covering populations) to make them ready for a new era of precision (and accuracy) medicine and public health [16].

2.1. Cognitive digital twins

A digital twin augmented with AI capabilities can be called a Cognitive Digital Twin (CDT) [20, 21]. Semantic Web technologies, such as ontologies and knowledge graphs, can be incorporated within DTs to support reasoning and deliberation [21]. Abburu et al. describe a broader vision and architecture for a CDT in the process and manufacturing industry [20]. Their CDT architecture incorporates cognitive features that will enable sensing complex and unpredicted behaviour and reason about dynamic strategies for process optimization, leading to a system that continuously evolves its own digital structure and behaviour.

The CDT evolves into a self-learning and proactive system that will optimize its own cognitive capabilities over time based on the data it will collect and the experience it will gain. It will find new answers to emerging questions by combining expert knowledge with the power of Digital Twins. A CDT will thus achieve synergy between the DT and the expert and problem-solving expertise. Unlike the process industry, which can be viewed as a closed and bounded system, in public health, a plethora of digital twins may be developed by different organisations and in other domains with potentially different perspectives. A more recent proposal by Ricci et al. [22, 23] explores and proposes a distributed and open architecture and platform for a Web of Digital Twins. They describe an open ecosystem of multiple digital twins possibly belonging to different domains and organisations. The open-distributed system perspective aligns well with the One Health perspective for low-resource African countries. Population, clinical, vital statistics, e.g. births and deaths, and geospatial and environmental data may be stored in different systems across different government departments, possibly with other systems at the facility, city, district, state and country levels.

2.2. Digital twins for healthcare and well-being

Several recent studies have explored digital twins for healthcare [17, 24, 25]. The most common approach is to use DTs for precision medicine [17, 18] where healthcare practice shifts from being primarily reactive to a more preventative approach. Many studies highlighted the higher complexity of designing digital twins for health care compared to digital twins in manufacturing and industry.

Ahmadi-Assalemi et al [17] highlighted different types of support that DTs must provide at varying levels of the healthcare system, i.e. individual level, healthcare professional level, and health system level. At the individual level, a patient's complex needs may influence real-time behaviours, feelings, adherence to the set targets and the utilization of healthcare services. They propose a more ambitious definition of health as the state of complete well-being, including the physical, mental, and social aspects in addition to the biomedical one. This proactive patient care aims to preempt the disease through preventative medicine and early detection, which could change the societal culture by empowering individuals to prevent their own disease. They note that many factors that affect a person's well-being and condition, including environmental, demographic, socioeconomic or biological, in a constantly changing landscape, are not detected during routine health screening.

Kamel-Boulos et al suggested a role for Digital twins in precision public health and disease outbreaks potentially integrated as part of a health city [16]. The suggested systems could include advanced systems for case management and case finding as well as determining risk levels in particular geographic areas [16].

Hassani et al [24] highlight the complexity of healthcare and propose different digital twins for the various life stages of a person. They argue that digital twins can be used to combat healthcare inequality, improve operational efficiencies of healthcare facilities and accelerate advances in healthcare research. They also highlight the need for further research on the interactions between digital twins and AI. More in-depth reviews of digital twins in healthcare can be found in [26, 27, 28].

3. AI systems

In this section, we provide an overview of emerging trends and future directions for AI systems in the context of the One Health vision. Next-generation AI systems must support scientific knowledge discovery, complex social and behavioural modelling, decision making and simulation.

3.1. Augmented AI systems

Our perspective of AI aligns closely with the notion of augmented AI [29, 30], which more recently is being referred to as Hybrid Intelligence [31, 32]. In augmented AI, the AI system amplifies human cognition rather than replacing it. As such the human user, in this case the public health practitioner, works interactively and cooperatively with the system.

An augmented AI system can be characterised as an adaptive and cognitive system. The adaptive characteristic presents broader challenges beyond merely dealing with changes emanating from dynamic environments. To better describe this interaction, we use the word agent as a personification of the system from the user's perspective. From the human user perspective, the agent forms the communication interface between the system and the user. For human-agent interaction to be effective, the user must be able to specify their knowledge about the environment and decision-making context to the agent.

The agent's knowledge base and reasoning processes constitute a shared model reflecting both the system's and the user's knowledge, reasoning and decision-making processes. This shared model serves as a basis for communication and interaction between the user and the system. The user may initially specify goals and objectives which are incomplete and vague. These will also change and evolve naturally over time as the agent adapts to the user and the user to the agent. Even though the agent may incorporate different inference algorithms, the agent must be able to follow reasoning patterns compatible with and understandable by the human user within their application context. In this way, the rationale and analysis of any decision recommendation can effectively be communicated and explained to the user. The human user can ask the question, "Why would I do this?" or "How did you arrive at this conclusion?". As such, an agent must convince the user that the action it recommends is the best one to take after analysing all the information at hand.

3.2. Scientific knowledge discovery

The One Health approach proposes an integrated, unifying approach to balance and optimize the health of people, animals and the environment. It will leverage scientific understanding and theories in the earth, biological and behavioural sciences. Physical, biological and social processes are constantly evolving and depend on dynamic micro-scale and meso-scale environments [33]. Processes may vary significantly at different locations and typically change over time. Subsequently, theoretical positions differ significantly and may even be contradictory. An integral part of data analysis and scientific enquiry is the application of different theories to analyse incoming observations. Well-founded and established theories can be incorporated into modelling, simulation and decision-support tools. In contrast, new theories can be constantly

adapted and evaluated against incoming observations to test their validity. Moreover, powerful Artificial Intelligence (AI) data mining and pattern analysis techniques can be continuously applied to data over time to capture new patterns, formulate new theories and refine existing theories. Ultimately the scientific community should publish dynamic and formal models of their theories and validation data online for immediate dissemination and reproducibility. This has the potential to accelerate and reduce the cost and effort for scientific discovery and lower the barrier for decision-makers to access the latest data analysis and simulation models to improve their planning and inform policy.

AI-driven knowledge discovery systems will have the capacity to pursue scientific research, collect measurements, find regularities, form hypotheses, and gather additional data to test them [34]. Like humans these continual learning systems [35] will learn incrementally and cumulatively. Key functions will include taxonomy formation, descriptive law induction and explanatory model construction [34]. These systems would engage in all of these scientific activities, each of which can involve detecting and responding to anomalous observations. While these systems may act autonomously [34], a more compelling vision is for them to work in tandem with human researchers as a fully functional member of a research team. There are three fundamental functional areas where AI can contribute to new scientific understanding [36]. First, AI can act as an instrument revealing properties of a physical system that are difficult or even impossible to probe. Humans then use these insights to formulate new scientific understanding. Second, AI can act as a source of inspiration for new concepts and ideas that are subsequently understood and generalized by human scientists. Third, AI acts as an agent to create new understanding. AI reaches new scientific insight and can transfer it to human researchers.

3.3. Knowledge based approaches

Generating hypotheses and maintaining a consistent body of knowledge in science is a formidable task due to the vast number of hypotheses generated and maintained, and the complexity, non-monotonicity, uncertainty and unreliability of knowledge and data published [37]. Logic-based concept ontologies emanating from the Semantic Web community have been widely used to represent, and reason about concepts and relations in a domain [38]. Ontology languages like OWL are typically based on monotonic logic such as Description Logics and provide limited support to deal with uncertainty and non-monotonicity. While they can effectively capture both syntax and semantics to enable communication, they are limited in terms of capturing and representing pragmatics, i.e., the context in which the knowledge is used [39, 40]. Ontologies can enable semantic interoperability, but not pragmatic interoperability between agents [41]. Casanovas [42] argue that social context is crucial and that meaning is only fully realised in actual situations assuming an interactional and dynamic notion of context. A further limitation is the lack of explicit support for representing decision-making processes.

3.3.1. Bayesian Networks and frameworks

Bayesian Networks (BNs) are highly effective causal models that aim to capture human reasoning and decision-making in uncertain situations [43]. They support three different types

of explanations [44]. For the explanation of the evidence, an abduction inference process is typically used to obtain the most probable explanation (MPE) explanation of the evidence. The purpose of this kind of explanation is to offer a diagnosis for a set of observed anomalies. Explanation of the model consists of displaying the information contained in the knowledge base and allows experts to easily navigate the knowledge. Explanation of reasoning provides justification for inferences made by the system. It allows users to understand and check the correctness of the reasoning process and the inferences made by the system. Compared to a rule-based expert system with IF-THEN rules, BNs are more compatible with present thinking about explanation [45]. However, it is neither clear how explanations in BNs can capture pragmatics, nor how to operationalise explanatory virtues in the context of BNs [45].

In philosophy, the conventional rational choice model of decision-making (based upon expected utility maximization), is widely discussed and used [46]. Bayesian Decision Networks (BDN), which are based on expected utility maximisation, are widely used for representing decision-making processes. Fareh [47] discusses how Bayesian Networks (BN) can be combined with ontologies to represent incomplete knowledge and uncertainty to predict liver cancer. BDNs also incorporate uncertainty and can recommend decisions where there is incomplete knowledge. However, rational choice can be viewed as a prescriptive model, a way of specifying how individuals ought to behave, rather than how they actually behave. Human behaviour can also be impacted by emotion, mood, personality, needs and subjective well-being. Abaalkhail [48] reviews current ontologies for affective states (emotion and mood). These ontologies can be used to build decision-support systems that consider emotion and mood. In a recent review paper, Barthès [49] explored decision-making mechanisms for ethical decision-making in cognitive agents. He highlights uncertainty and incomplete knowledge as key issues that must be taken into account. He provides a preliminary computational framework that takes into account different factors that influence human decision-making, including emotion and mood. This can be used as a foundation to design and analyse empathetic agents that are better able to support human decision-making.

3.3.2. Explanation in AI systems

Knowledge discovery, cognition, communication, decision making and explanation are integrally linked. In their recent papers, Chari et al [50, 51] provide an excellent analysis of explanation in knowledge-based AI systems. They identify nine types of explanation [50] in AI systems and highlight the need for AI systems to present personalized, trustworthy, and context-aware explanations to users. Our view is that a 3D virtual twin can provide a visual explanation in both space and time and is grounded in a pseudo-reality. We consider this as a new type of explanation. This is linked to the notion of pragmatic and social world communication. A 3D virtual digital twin can provide a rich context for communicating predictions, theories and beliefs posed by the AI system. It provides the outputs and deliberations of the AI system to be rendered in a pseudo-reality for further analysis.

3.4. Continual learning

Continual learning, also known as incremental learning or lifelong learning, is an increasingly relevant area of study that asks how artificial systems might learn sequentially, as biological systems do, from a continuous stream of data [35, 52, 53]. Unlike, traditional Deep Neural Network (DNN) tasks that rely on fixed datasets and stationary environments, the problem of continual learning is defined by a sequential training protocol and by the features expected from the solution. In contrast to the common machine learning setting of a static dataset or environment, the continual learning setting explicitly focuses on non-stationary or changing environments. Continual learning systems must deal with temporal generalisation, change detection and mechanisms for a model update, specifying and aligning expected features and loss functions, and appropriate validation and data partitioning methods. While some of these issues have been explored, e.g. prequential or walk-forward validation in time series forecasting [54], and more recently temporal generality in language models [55], this is still an emerging area with many open challenges.

Spatial Temporal Graph Neural Networks (ST-GNN) are a new wave of advanced DNN techniques, which have emerged recently to model and predict flow in complex systems [56]. Typical characteristics are high frequency and noisy observations from multiple sensors with complex and often latent spatial and temporal dependencies. While the canonical application is for traffic flow in a city, these techniques have wider applications for modelling dynamic systems in general, for example, weather modelling [57]. Weather prediction has a higher complexity than the traffic problem. Unlike traffic flow prediction which models a single variable, i.e. traffic speed at different points in a city, weather prediction involves many different weather variables at different temporal and spatial scales. A key feature of ST-GNNs is that they dynamically learn complex spatial-temporal dependencies inherent in the data and capture this in an adjacency matrix. These dependencies can be used as the foundation for constructing causal theories in the domain. We believe that ST-GNNs are at the frontier of learning-based approaches for knowledge discovery, predictive modelling and explanation in dynamic environments.

3.5. Architectures for scientific knowledge discovery systems

AI systems for scientific knowledge discovery are introduced in section 3.2. In this section we explore some of the architectures for such systems. Systems that can contribute to new scientific understanding are certainly on the frontier of AI [34, 36, 37]. A compelling vision for the field is to design an “AI scientist” that works in tandem with human researchers as a fully functioning member of a research team. This overlaps with the notion of hybrid intelligence and augmented intelligence systems where the machine works in collaboration with and cooperatively with the human user [58, 31]. Hybrid intelligence systems for knowledge discovery have been explored recently by Gil in her position paper on “Thoughtful Artificial Intelligence Systems” [59] where she proposes seven principles for such systems and sets an agenda for research in this area.

3.5.1. Ontology driven systems

One of the first examples of knowledge discovery systems is the Robot Scientist [60]. The Robot Scientist is a semantic architecture that incorporates the LABORS ontology. More recent

efforts are DISK [61] and HELO [62]. In the HELO study [62] the representation of uncertainty is identified as a key limitation of previous ontology-driven knowledge discovery systems. HELO explicitly targets this limitation and Bayesian probability is incorporated into the HELO ontology to represent the current belief or uncertainty of different research statements or hypotheses. The uncertainty of research statements is updated as new evidence arises, forming the basis for decision-making during the research process. This aligns with our previous work where we found that a knowledge engineering approach that uses ontologies alone has substantial limitations in real-world applications. Based on our experience in four studies in South Africa across diverse domains and user communities, i.e. earth observation [33], health [63], biodiversity [64], and finance [65], we found that an approach that combines both ontologies and Bayesian decision networks to be highly effective to deal with these challenges. While ontologies are ideal for representing and structuring domain knowledge, Bayesian decision networks (BDN) provide explicit support for reasoning with uncertainty, and for capturing and reasoning about decision-making processes.

3.5.2. Emerging architectures based on Philosophy of Science

Understanding scientific method is essential for designing knowledge discovery systems. Various theories of scientific method have been proposed in Philosophy of Science, for example, the inductive theory and the hypothetico-deductive theory of methods. The scientific Abductive Theory of Method (ATOM) [66, 67] is a recent theory that is more encompassing and detailed and is compatible with current thinking around knowledge discovery in the AI community [34, 36, 37] (described above). ATOM consists of two overarching processes, namely i) phenomenon detection where novel or anomalous patterns are detected in data, and ii) theory construction where plausible theories are generated and evaluated to explain detected phenomena. In recent work [68], we explored a generic architecture for a knowledge discovery agent based on ATOM. There are many advantages for grounding the architecture in ATOM. It will more closely align with the process that researchers actually use making it more human-centred. It provides an abstract architecture that incorporates both data-driven and knowledge-based AI approaches described above. It also provides an entry point to the large body of recent work emanating from the Philosophy of Science community on theory construction. Haig suggests the use of analogical reasoning for theory construction and for assessing explanations [66]. Confirmation based on analogical inference using a Bayesian framework is currently being explored in the Philosophy of Science community [69, 70]. Analogical inference is related to the notion of transfer learning in the machine learning community. This is particularly important in low-resource settings where large clinical data sets may not be available. Machine learning models may be training initially on large data sets from diverse populations in the developed world and then fine tuned on smaller data sets in low resource African countries. An initial exploration of transfer learning to detect heart disease in diverse populations from ECG signals can be found in [71].

4. An architecture for an AI-driven digital twin system

We propose an initial abstract architecture for an augmented AI system in figure 1. The architecture draws from and extends the ideas presented in Wanyana and Moodley's KDE Agent architecture [68], Russell and Norvig's learning agent [72] and Ricci et al's WoDT [22]. The AI system augments the 3D digital twin and uses it as a key platform for its simulations, for knowledge discovery and ultimately to augment the human user, in this case, the public health decision maker.

4.1. System overview

The system will continuously interrogate vast quantities of heterogeneous observational data, automatically generate and maintain internal models to explain evolving phenomena, evaluate these models through simulation in the DT and in cooperation with the user, discover a new understanding of evolving phenomena and provide possible courses of actions to achieve the goals set by the human user.

The system maintains two views of the world which connect it to the user and the physical world. The Cognitive view maintains a model of the user, i.e. the beliefs, goals and decision paths of the public health decision-maker. The Physical view encompasses the 3D digital world, which is used to model a pseudo-reality of the physical world, including properties of physical entities, and interactions and behaviours of individuals and populations.

Adaptation and cognition are two essential functional areas of the system. The adaptive characteristic is concerned with rapid learning and model updates when the world changes, communicating new information about the environment to the user and adjusting to changes in the user's decision-making goals and preferences as the user learns and adjusts to the system. The cognitive characteristic involves the use of representing prior knowledge, world dynamics and appropriate reasoning mechanisms for sense-making and belief update and revision that is aligned to human reasoning. The cognitive aspect facilitates the overall interactions with the user and provides explicit support for providing explanations and allows the user to interrogate courses of action recommended by the system.

The abstract architecture consists of three layers, i.e. the Monitoring Layer, the Situation Analysis Layer and the Augmentation Layer

4.2. Monitoring Layer

At the Monitoring Layer, ontologies will be used for semantic data annotation and data fusion. Data will include earth observation data, clinical data, population data, behavioural data and clinical data. Existing semantic architectures from our previous work in earth observation [33] and biodiversity [64] will be leveraged for semantic mediation and dynamic data fusion.

The Sensor Web Agent Platform (SWAP) architecture [33] proposes a three-layered open distributed system architecture and a conceptual knowledge representation and reasoning framework that integrates ontologies and Bayesian networks for developing distributed Sensor Web applications. The SWAP semantic architecture provides modular ontologies for reasoning about space, time, theme (the entity and properties being observed) and uncertainty and supports




Figure 1: An architecture for an AI-driven 3D digital twin

semantic annotation and dynamic data fusion from heterogeneous sensor data. Coetzer et al [64] explored a semantic architecture for a knowledge-based system that uses expert knowledge to generate ecological interaction networks from distributed and heterogeneous natural history occurrence data. The system builds on the SWAP architecture and introduces a biodiversity mapping ontology that supports semantic annotation and fusion of heterogeneous data from three biodiversity databases hosted by different museums across South Africa.

The 3D digital twin provides a pseudo-reality of the physical world. Observational data stored in databases in existing systems together with real-time sensor observations will be semantically annotated, fused and rendered in the 3D digital twin. The fused observations are contextualised within the 3D world and must be harmonised and logically consistent with other observations, entities and processes that are unfolding in the 3D world. The 3D world thus provides context for and guides data fusion and is integral to achieving pragmatic interoperability.

4.3. Situation Analysis

This layer consists of two components. Situation detection can use ontologies, semantic rules, fuzzy rules and Bayesian networks for detecting current situations. Semantic technologies are increasingly being used for situation analysis in personal health monitoring systems that incorporate wearable sensors [73]. Bayesian networks can be used for situation analysis and

root cause analysis to support clinical diagnosis [74]. They can be supplemented by machine learning, e.g. ST-GNNs for flow prediction for early detection of adverse events.

Ontologies and Bayesian networks can be combined for behavioural modelling and risk analysis, based on an approach we used in our previous work on analysing the risk of non-adherence behaviour for medical treatment for tuberculosis patients [63].

4.4. Augmentation Layer

The Augmentation Layer has three functional components. Similarly to the CDT, the system augments the DT and in turn augments the user. This tridirectional augmentation is a crucial aspect of the architecture. Each environment evolves separately, but feedback loops between the user, the DT and the AI system allow for synchronisation and alignment.

The system maintains the current beliefs in a knowledge base. It uses continual learning, e.g. STGNNs, for discovering novel spatial-temporal patterns or new phenomena. This triggers the theory construction module, which attempts to find explanations for the phenomena. If there is no suitable explanation, it uses 3D simulation in the DT to better understand the nature and possible root causes of the novel phenomena. These possible explanations form new unconfirmed theories that must be further explored. The latest theories are exposed to the user in the DT for further analysis and input. The user can prioritise theories or propose alternate theories not discovered by the system. The critic module, similar to the critic module in Russell & Norvigs Learning Agent [72], provides internal evaluation and analysis of new theories. It prioritises theories aligned with the user's goals and uses the DT to evaluate theories by integrating knowledge from different contexts and perspectives.

The key notion is explicitly modelling the human decision maker's mental models, i.e. the Cognitive view. We see the digital twin as capturing physical notions, but the specification of beliefs, goals, decisions, and knowledge are abstract notions which we separate from the physical representation. The physical representation (Physical view) not only models properties of physical entities but also interactions and behaviours of individual people and populations.

The architecture for an AI-driven digital twin system shown in Figure 1 aligns with the public health decision-making process in a PHEOC, described in section 1.2. An AI-driven digital twin can represent and support different activities and, in addition, provide enhanced support for predictive analytics, knowledge discovery and interactive decision-making.

5. Discussion and conclusion

This paper describes an ambitious and bold vision to harness the potential of artificial intelligence and digital twins to re-imagine public health in African countries. A future digital platform for One Health will have to overcome both social and technical challenges. An open and distributed One Health platform must mediate between many heterogeneous systems which manage population, clinical, vital statistics, geospatial and environmental data across different government departments at the facility, city, district, state and country levels.

The application of Digital Twins to public health shows much promise in LMICs. An opportunity exists to use a 3D digital twin to mediate and contextualise different data sources, models and systems from a public health perspective. This could include dynamic fusion and

harmonisation of multiple information sources and models in an open distributed environment with different contributors and users. Equipped with AI, the system could also provide predictive modelling and risk profiling to mitigate disease outbreaks and support early warning systems for adverse health conditions. The digital twin provides a powerful visual analysis and simulation tool for investigating and explaining the root causes of events of interest in public health. However, it will be important to understand the ethical implications of these approaches.

To truly unleash the power of these emerging technologies would require a continental-scale effort and buy-in across multiple African countries, ministries of health, international donors, universities and non-governmental organisations. Establishing standards and platforms and an ecosystem to support this vision will indeed be challenging, especially in low-resource environments. The structure of the ecosystem must balance long-term sustainability, provide a clear vision, and must be resilient to changing technologies and organisational changes in the health system. Based on our previous experience in the public health sector our view is that a middle-out strategy [75] should be considered for structuring such an ecosystem. In such a middle-out approach, a central authority, such as the national government, manages and provides leadership in setting standards, the platform architecture and the overall direction while still allowing autonomy for participating organisations to implement and manage their own systems.

References

- [1] WHO, Declaration of Alma-Ata, Technical Report, 1978. URL: <https://www.who.int/publications/i/item/WHO-EURO-1978-3938-43697-61471>.
- [2] WHO, World Health Statistics 2021 (2021). URL: <https://apps.who.int/iris/bitstream/handle/10665/342703/9789240027053-eng.pdf>.
- [3] M. M. Coates, M. Ezzati, G. R. Aguilar, G. F. Kwan, D. Vigo, A. O. Mocumbi, A. E. Becker, J. Makani, A. A. Hyder, Y. Jain, D. Cristina Stefan, N. Gupta, A. Marx, G. Bukhman, Burden of disease among the world's poorest billion people: An expert-informed secondary analysis of Global Burden of Disease estimates, *PLoS ONE* 16 (2021). doi:10.1371/journal.pone.0253073.
- [4] B. Natifu, Climate Change and Health in Sub-Saharan Africa: The Case of Uganda, Technical Report, 2020. URL: https://www.cif.org/sites/cif_enc/files/knowledge-documents/final_chasa_report_19may2020.pdf.
- [5] T. Ciecierski-Holmes, R. Singh, M. Axt, S. Brenner, S. Barteit, Artificial intelligence for strengthening healthcare systems in low- and middle-income countries: a systematic scoping review, 2022. URL: <https://www.nature.com/articles/s41746-022-00700-y>. doi:10.1038/s41746-022-00700-y.
- [6] The Lancet, One Health: a call for ecological equity, 2023. doi:10.1016/S0140-6736(23)00090-9.
- [7] K. E. Worsley-Tonks, J. B. Bender, S. L. Deem, A. W. Ferguson, E. M. Fèvre, D. J. Martins, D. M. Muloi, S. Murray, M. Mutinda, D. Ogada, G. P. Omondi, S. Prasad, H. Wild, D. M. Zimmerman, J. M. Hassell, Strengthening global health security by improving

- disease surveillance in remote rural areas of low-income and middle-income countries, 2022. URL: [https://www.thelancet.com/journals/langlo/article/PIIS2214-109X\(22\)00031-6/fulltext](https://www.thelancet.com/journals/langlo/article/PIIS2214-109X(22)00031-6/fulltext). doi:10.1016/S2214-109X(22)00031-6.
- [8] Working together for the health of humans, animals, plants and the environment. One Health Joint Plan of Action (2022–2026), *Bulletin de l’OIE* 2022 (2022) 18–19. doi:10.20506/bull.2022.2.3324.
- [9] World Health Organisation, Framework for a Public Health Emergency Operations Centre, World Health Organisation, 2015. URL: <https://www.who.int/publications/i/item/framework-for-a-public-health-emergency-operations-centre>.
- [10] WHO, Handbook for Public Health Emergency Operations Center Operations and Management., 2021. URL: <https://africacdc.org/download/handbook-for-public-health-emergency-operations-center-operations-and-management/>.
- [11] Ethiopia Public Health Institute (EPHI), National Public Health Emergency Operation Center Handbook, Technical Report, 2022. URL: https://ephi.gov.et/wp-content/uploads/2022/07/EPHI_cPHEM_EWISMD_PHEOC_Handbook_V1.pdf.
- [12] UNAIDS, Fast-Track Targets, Technical Report, 2020. URL: https://www.unaids.org/sites/default/files/media_asset/201506_JC2743_Understanding_FastTrack_en.pdf.
- [13] PEPFAR, Reimagining PEPFAR’s Strategic Direction, Technical Report, 2022. URL: https://www.state.gov/wp-content/uploads/2022/09/PEPFAR-Strategic-Direction_FINAL.pdf.
- [14] B. Rice, A. Boulle, S. Baral, M. Egger, P. Mee, E. Fearon, G. Reniers, J. Todd, S. Schwarcz, S. Weir, G. Rutherford, J. Hargreaves, Strengthening routine data systems to track the HIV epidemic and guide the response in Sub-Saharan Africa, *JMIR Public Health and Surveillance* 4 (2018). doi:10.2196/publichealth.9344.
- [15] D. Moodley, A. Pillay, J. Seebregts, Establishing a health informatics research laboratory in South Africa, in: *Digital Re-imagination Colloquium 2018: Preparing South Africa for a Digital Future through e-Skills*, NEMISA, 2018, pp. 16–24. URL: <https://www.cair.org.za/sites/default/files/2020-02/HeAL-NEMISA-2018.pdf>.
- [16] M. N. Kamel Boulos, P. Zhang, Digital twins: From personalised medicine to precision public health, 2021. doi:10.3390/jpm11080745.
- [17] G. Ahmadi-Assalemi, H. Al-Khateeb, C. Maple, G. Epiphaniou, Z. A. Alhaboby, S. Alkabi, D. Alhaboby, Digital twins for precision healthcare, in: *Advanced Sciences and Technologies for Security Applications*, 2020. doi:10.1007/978-3-030-35746-7_{_}8.
- [18] G. Coorey, G. A. Figtree, D. F. Fletcher, J. Redfern, The health digital twin: advancing precision cardiovascular medicine, 2021. doi:10.1038/s41569-021-00630-4.
- [19] A. Fuller, Z. Fan, C. Day, C. Barlow, Digital Twin: Enabling Technologies, Challenges and Open Research, *IEEE Access* 8 (2020) 108952–108971. doi:10.1109/ACCESS.2020.2998358.
- [20] S. Abburu, A. J. Berre, M. Jacoby, D. Roman, L. Stojanovic, N. Stojanovic, COGNITWIN - Hybrid and Cognitive Digital Twins for the Process Industry, in: *Proceedings - 2020 IEEE International Conference on Engineering, Technology and Innovation, ICE/ITMC 2020*, 2020. doi:10.1109/ICE/ITMC49519.2020.9198403.
- [21] X. Zheng, J. Lu, D. Kiritsis, The emergence of cognitive digital twin: vision, challenges and opportunities, *International Journal of Production Research* 60 (2022) 7610–7632. doi:10.1080/00207543.2021.2014591.

- [22] A. Ricci, A. Croatti, S. Mariani, S. Montagna, M. Picone, Web of Digital Twins, *ACM Transactions on Internet Technology* 22 (2022) 1–30. doi:10.1145/3507909.
- [23] A. Ricci, A. Croatti, S. Montagna, Pervasive and Connected Digital Twins - A Vision for Digital Health, *IEEE Internet Computing* 26 (2022) 26–32. doi:10.1109/MIC.2021.3052039.
- [24] H. Hassani, X. Huang, S. MacFeely, Impactful Digital Twin in the Healthcare Revolution, *Big Data and Cognitive Computing* 6 (2022). doi:10.3390/bdcc6030083.
- [25] D. Chen, N. A. AlNajem, M. Shorfuzzaman, Digital twins to fight against COVID-19 pandemic, *Internet of Things and Cyber-Physical Systems* 2 (2022) 70–81. doi:10.1016/j.iotcps.2022.05.003.
- [26] N. P. Rocha, A. Dias, G. Santinha, M. Rodrigues, A. Queirós, C. Rodrigues, Smart Cities and Public Health: A Systematic Review, in: *Procedia Computer Science*, volume 164, Elsevier B.V., 2019, pp. 516–523. doi:10.1016/j.procs.2019.12.214.
- [27] A. Khan, M. Milne-Ives, E. Meinert, G. E. Iyawa, R. B. Jones, A. N. Josephraj, A Scoping Review of Digital Twins in the Context of the Covid-19 Pandemic, *Biomedical Engineering and Computational Biology* 13 (2022). doi:10.1177/11795972221102115.
- [28] S. Elkefi, O. Asan, Digital Twins for Managing Health Care Systems: Rapid Literature Review, 2022. doi:10.2196/37641.
- [29] N. n. Zheng, Z. y. Liu, P. j. Ren, Y. q. Ma, S. t. Chen, S. y. Yu, J. r. Xue, B. d. Chen, F. y. Wang, Hybrid-augmented intelligence: collaboration and cognition, 2017. doi:10.1631/FITEE.1700053.
- [30] K. L. A. Yau, H. J. Lee, Y. W. Chong, M. H. Ling, A. R. Syed, C. Wu, H. G. Goh, Augmented Intelligence: Surveys of Literature and Expert Opinion to Understand Relations between Human Intelligence and Artificial Intelligence, *IEEE Access* 9 (2021) 136744–136761. doi:10.1109/ACCESS.2021.3115494.
- [31] Z. Akata, D. Balliet, M. De Rijke, F. Dignum, V. Dignum, G. Eiben, A. Fokkens, D. Grossi, K. Hindriks, H. Hoos, H. Hung, C. Jonker, C. Monz, M. Neerincx, F. Oliehoek, H. Prakken, S. Schlobach, L. Van Der Gaag, F. Van Harmelen, H. Van Hoof, B. Van Riemsdijk, A. Van Wynsberghe, R. Verbrugge, B. Verheij, P. Vossen, M. Welling, A Research Agenda for Hybrid Intelligence: Augmenting Human Intellect with Collaborative, Adaptive, Responsible, and Explainable Artificial Intelligence, *Computer* 53 (2020) 18–28. doi:10.1109/MC.2020.2996587.
- [32] W. Zhang, H. Ning, L. Liu, Q. Jin, V. Piuri, Guest Editorial: Special Issue on Hybrid Human-Artificial Intelligence for Social Computing, 2021. doi:10.1109/TCSS.2021.3049702.
- [33] D. Moodley, I. Simonis, J. Tapamo, Architecture for managing knowledge and system dynamism in the worldwide sensor web, *International Journal on Semantic Web and Information Systems* 8 (2012). doi:10.4018/jswis.2012010104.
- [34] P. Langley, Agents of exploration and discovery (2021). doi:10.1609/aaai.12021.
- [35] R. Hadsell, D. Rao, A. A. Rusu, R. Pascanu, Embracing Change: Continual Learning in Deep Neural Networks, 2020. doi:10.1016/j.tics.2020.09.004.
- [36] M. Krenn, R. Pollice, S. Y. Guo, M. Aldeghi, A. Cervera-Lierta, P. Friederich, G. dos Passos Gomes, F. Häse, A. Jinich, A. K. Nigam, Z. Yao, A. Aspuru-Guzik, On scientific understanding with artificial intelligence, *Nature Reviews Physics* 4 (2022) 761–769. doi:10.1038/s42254-022-00518-3.

- [37] H. Kitano, Nobel Turing Challenge: creating the engine for scientific discovery, 2021. doi:10.1038/s41540-021-00189-3.
- [38] P. Hitzler, Semantic Web: A Review Of The Field (2020). URL: <https://doi.org/10.1145/nnnnnnnn>. doi:10.1145/nnnnnnnn.
- [39] M. Schoop, A. De Moor, Dietz Jan L G, pragmatic-web-2006-manifesto, Communications of the ACM 49 (2006).
- [40] M. Singh, The Pragmatic Web, IEEE Internet Computing 6 (2002) 4–5. URL: <http://computer.org/internet/>.
- [41] F. W. Neiva, J. M. N. David, R. Braga, F. Campos, Towards pragmatic interoperability to support collaboration: A systematic review and mapping of the literature, Information and Software Technology 72 (2016) 137–150. doi:10.1016/j.infsof.2015.12.013.
- [42] P. Casanovas, V. Rodríguez-Doncel, J. González-Conejero, The role of pragmatics in the web of data, in: Pragmatics and Law, Springer, 2017, pp. 293–330. doi:10.1007/978-3-319-44601-1{_}12.
- [43] K. B. Korb, A. E. Nicholson, Bayesian artificial intelligence, CRC press, 2010.
- [44] C. Lacave, F. J. Díez, A review of explanation methods for Bayesian networks, 2002. doi:10.1017/S026988890200019X.
- [45] M. Tesic, U. Hahn, Explanation in AI systems, in: Human-Like Machine Intelligence, Oxford University Press, 2021, pp. 114–136. doi:10.1093/oso/9780198862536.003.0006.
- [46] Briggs R A, Normative theories of rational choice: Expected utility. Available online: , 2014. URL: <https://plato.stanford.edu/entries/rationality-normative-utility/>.
- [47] M. Fareh, Modeling incomplete knowledge of semantic web using Bayesian networks, Applied Artificial Intelligence 33 (2019) 1022–1034. doi:10.1080/08839514.2019.1661578.
- [48] R. Abaalkhail, B. Guthier, R. Alharthi, A. E. Saddik, Survey on Ontologies for Affective States and Their Influences, Technical Report, ????. URL: <http://tomdrummond.com/leading-and-caring->.
- [49] J. P. A. Barthes, Cognitive Agents and Ethical Behavior in Collaborative Teams, in: Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics, volume 2020-October, Institute of Electrical and Electronics Engineers Inc., 2020, pp. 3776–3781. doi:10.1109/SMC42975.2020.9282936.
- [50] S. Chari, D. M. Gruen, O. Seneviratne, D. L. McGuinness, Directions for Explainable Knowledge-Enabled Systems (2020). URL: <http://arxiv.org/abs/2003.07523>.
- [51] S. Chari, D. M. Gruen, O. Seneviratne, D. L. McGuinness, Foundations of Explainable Knowledge-Enabled Systems (2020). URL: <http://arxiv.org/abs/2003.07520>.
- [52] D. L. Silver, Q. Yang, L. Li, Lifelong Machine Learning Systems: Beyond Learning Algorithms, Technical Report, 2013. URL: www.aaai.org.
- [53] G. M. van de Ven, T. Tuytelaars, A. S. Tolias, Three types of incremental learning, Nature Machine Intelligence 4 (2022) 1185–1197. doi:10.1038/s42256-022-00568-3.
- [54] V. Cerqueira, L. Torgo, I. Mozetič, Evaluating time series forecasting models: an empirical study on performance estimation methods, Machine Learning 109 (2020) 1997–2028. doi:10.1007/s10994-020-05910-7.
- [55] A. Lazaridou, A. Kuncoro, E. Gribovskaya, D. Agrawal, A. Liška, T. Terzi, M. Gimenez, C. de Masson, T. Kocisky, S. Ruder, D. Yogatama, K. Cao, S. Young, P. Blunsom, Mind the Gap: Assessing Temporal Generalization in Neural Language Models, in: Advances in Neural

- Information Processing Systems, volume 34, Advances in Neural Information Processing Systems, 34, pp.29348-29363., 2021, pp. 29348–29363. URL: <http://data.statmt.org/>.
- [56] K. H. N. Bui, J. Cho, H. Yi, Spatial-temporal graph neural network for traffic forecasting: An overview and open research issues, *Applied Intelligence* 52 (2022) 2763–2774. doi:10.1007/s10489-021-02587-w.
- [57] M. Davidson, D. Moodley, ST-GNNs for Weather Prediction in South Africa, in: *Communications in Computer and Information Science*, volume 1734 CCIS, Springer Science and Business Media Deutschland GmbH, 2022, pp. 93–107. doi:10.1007/978-3-031-22321-1{_}7.
- [58] D. Dellermann, P. Ebel, M. Söllner, J. M. Leimeister, Hybrid Intelligence, *Business and Information Systems Engineering* 61 (2019) 637–643. doi:10.1007/s12599-019-00595-2.
- [59] Y. Gil, Thoughtful Artificial Intelligence: Forging A New Partnership for Data Science and Scientific Discovery, Technical Report, ????
- [60] R. D. King, J. Rowland, W. Aubrey, M. Liakata, M. Markham, L. N. Soldatova, K. E. Whelan, A. Clare, M. Young, A. Sparkes, others, The robot scientist Adam, *Computer* 42 (2009) 46–54.
- [61] Y. Gil, D. Garijo, V. Ratnakar, R. Mayani, R. Adusumilli, H. Boyce, P. Mallick, Automated hypothesis testing with large scientific data repositories, in: *Proceedings of the Fourth Annual Conference on Advances in Cognitive Systems (ACS)*, volume 2, 2016, p. 4.
- [62] L. N. Soldatova, A. Rzhetsky, K. De Grave, R. D. King, Representation of probabilistic scientific knowledge, *Journal of Biomedical Semantics* 4 (2013). doi:10.1186/2041-1480-4-S1-S7.
- [63] O. Ogundele, D. Moodley, A. Pillay, C. Seebregts, An ontology for factors affecting tuberculosis treatment adherence behavior in sub-Saharan Africa, *Patient Preference and Adherence* 10 (2016). doi:10.2147/PPA.S96241.
- [64] W. Coetzer, D. Moodley, A. Gerber, A knowledge-based system for generating interaction networks from ecological data, *Data & Knowledge Engineering* 112 (2017) 55–78. URL: <https://www.sciencedirect.com/science/article/pii/S0169023X17300459>. doi:10.1016/j.datak.2017.09.005.
- [65] R. Drake, D. Moodley, INVEST: Ontology driven Bayesian networks for investment decision making on the JSE, in: *Proceedings of the Second Southern African Conference for Artificial Intelligence Research*, 6-10 Dec. 2021, South Africa, 2022, pp. 252–273.
- [66] B. D. Haig, An abductive theory of scientific method, *Psychological Methods* 10 (2005) 371–388. doi:10.1037/1082-989X.10.4.371.
- [67] B. D. Haig, An abductive theory of scientific method, in: *Method matters in psychology*, Springer, 2018, pp. 35–64.
- [68] T. Wanyana, D. Moodley, An agent architecture for knowledge discovery and evolution, in: *German Conference on Artificial Intelligence (Künstliche Intelligenz)*, 2021, pp. 241–256. doi:10.1007/978-3-030-87626-5{_}18.
- [69] R. Dardashti, S. Hartmann, K. Thébault, E. Winsberg, Hawking radiation and analogue experiments: A Bayesian analysis, *Studies in History and Philosophy of Science Part B - Studies in History and Philosophy of Modern Physics* 67 (2019) 1–11. doi:10.1016/j.shpsb.2019.04.004.
- [70] C. J. Feldbacher-Escamilla, A. Gebharter, Confirmation Based on Analogical Inference:

Bayes Meets Jeffrey, *Canadian Journal of Philosophy* 50 (2020) 174–194. doi:10.1017/can.2019.18.

- [71] S. Aarons, D. Moodley, M. Nzomo, A Generalizable Hybrid Deep Learning Algorithm for the Detection of Atrial Fibrillation from Diverse Electrocardiogram Data, Technical Report, University of Cape Town, Cape Town, 2022. URL: https://projects.cs.uct.ac.za/honsproj/cgi-bin/view/2022/aarons_fisher_rosenthal.zip/deliverables/finalpapershai.pdf.
- [72] S. Russell, P. Norvig, *Artificial intelligence: A modern approach*, Prentice-Hall, Englewood Cliffs, NJ, 2009.
- [73] M. Nzomo, D. Moodley, *Semantic Technologies in Sensor-Based Personal Health Monitoring Systems: A Systematic Mapping Study (2023)*. URL: <http://arxiv.org/abs/2306.04335>.
- [74] T. Wanyana, M. Nzomo, C. S. Price, D. Moodley, Combining Machine Learning and Bayesian Networks for ECG Interpretation and Explanation, in: *International Conference on Information and Communication Technologies for Ageing Well and e-Health, ICT4AWE - Proceedings*, Science and Technology Publications, Lda, 2022, pp. 81–92. doi:10.5220/0011046100003188.
- [75] T. Mudaly, D. Moodley, A. Pillay, C. Seebregts, Architectural frameworks for developing national health information systems in low and middle income countries, in: *Proceedings of the 1st International Conference on Enterprise Systems, ES 2013*, 2013. doi:10.1109/ES.2013.6690083.