Towards Responsible Digital Twins

Zoran Milosevic [1[0000-0002-1364-7423]] and Pieter van Schalkwyk [2[0009-0009-3982-9346]]

[1] Deontik, Australia,
[2]XMPRO, USA/Australia
zoran@deontik.com,pvs@xmpro.com

Abstract. A digital twin is a virtual representation of real-world entities and processes, synchronized at a specified frequency and fidelity. The capability of digital twins is continually evolving from simple decision support, via decision augmentation for end users to autonomous decision automation. This evolution is enabled by increasingly sophisticated technologies used by digital twins, e.g. advanced analytics, IoT and AI. In many applications, multiple digital twins can be used to address different system functionality, and composed as required, leading to potentially quite complex technical systems. Digital twins further increasingly require explicit consideration of *socio-economic* factors, to ensure building *responsible* digital twin solutions, minimizing potential harm for the users. This paper discusses how such socio-economic factors, particularly the enterprise, legal and ethics policies and various value constraints, can be gradually transformed into a set of governance rules for building, operating and evolving responsible digital twin solutions and ecosystems. These policies include voluntary type of rules, e.g. digital ethics norms, as well as regulatory policies, which impose formal legal obligations, e.g. legislative and regulatory mechanisms. We use two application domains at two ends of the complexity spectrum, namely personalized health care and renewable energy, to illustrate our approach.

Keywords: Digital Twins, Machine Learning, Artificial Intelligence, complex system, computable policy, obligations, prohibitions, accountability, ethics.

1 Background

A digital twin (DT) is a virtual representation of real-world entities and processes, synchronized at a specified frequency and fidelity [1]. DTs can be used to study, monitor, and optimize the composition and functions of their physical counterpart. This emerging field has witnessed a meteoric rise, with an impressive growth rate of 71% between 2020 and 2022 [2]. This trend is projected to continue upward, with a forecasted leap from USD 10.1 billion in 2023 to USD 110.1 billion by 2028, representing a compound annual growth rate (CAGR) of 61.3% [3].

The concept of digital twins, initially introduced within the manufacturing sector by Grieves and Vickers [4], has since evolved and permeated a broad spectrum of industries. This includes aerospace and defense, agriculture, food and beverage, architecture and construction, financial services, healthcare and life sciences, mobility and transportation, natural resources, and telecommunications, as reflected in The Digital Twin Consortium's diverse working groups [1].

The rapid growth of DTs highlights their importance in the digital transformation of various industries. This growth prompts a discussion on the responsibilities tied to their development and use. As DTs evolve from simple to complex systems, the digital transformation process becomes more intricate.

A discrete DT represents a single entity, like a robot arm in a factory. When multiple discrete DTs are combined, they form a composite DT, representing a larger system comprising various components. For instance, a production cell's DT is a composite of the DTs of the devices within the cell. This process of combining discrete DTs into composite ones illustrates how DTs can represent increasingly complex systems.

DTs improve decision-making through real-time data and context-specific information. They provide users with detailed data visualizations, aiding in informed operational decisions. With the integration of AI and advanced analytics, they can extract hidden insights from large datasets, a task challenging for manual processing. This leads to decision augmentation, providing users with prescriptive recommendations.

In the future, DTs will move beyond decision augmentation to decision automation (Fig. 1). They will make strategic decisions based on AI, analytics, and business rules, enabling 'lights-out' operations and driving an algorithmic business model [6]. As such, their responsible use will become increasingly important, with the need to identify and integrate governance policies, as responsible features of digital twins. These policies can be voluntary, e.g. digital ethics norms as well as regulatory and legislative policies, with formal legal obligations, as also suggested in [16].

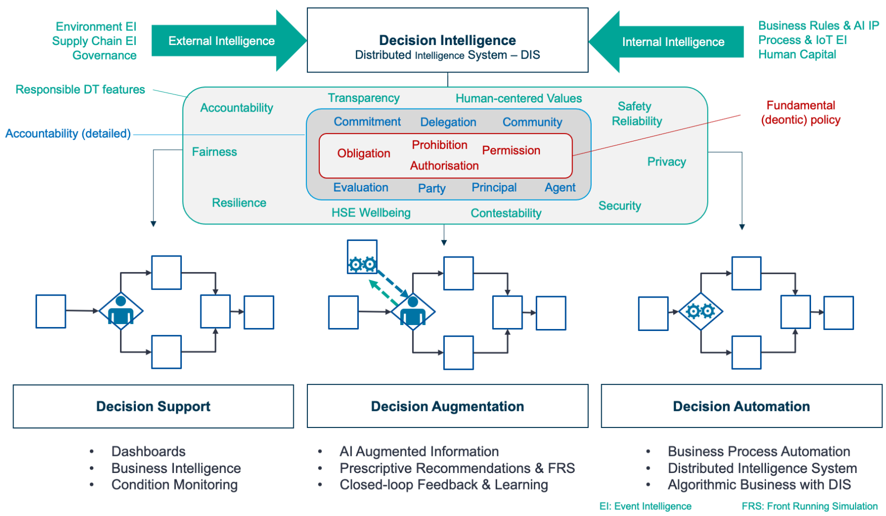

Figure 1: From decision support to decision automation

The Digital Twin Consortium's (DTC) Digital Twin Capabilities Periodic Table [7] and Reference Architecture [8] exemplify efforts to provide technical and architectural guidance for the development and implementation of DTs. Additionally, the Industrial Digital Twin Association (IDTA), a German-led initiative under the Industrie 4.0 umbrella [9], offers specific guidance on the technological implementation of digital twins within the manufacturing sector. These initiatives collectively contribute to the evolving body of knowledge and best practices in the field of DT technology.

While instrumental in advancing DT implementation, the technology-centric approach presents a challenge because it does not adequately address the socio-economic impact of utilizing DT technology. Furthermore, it does not sufficiently consider the implications for a DT's responsible and ethical use in facilitating effective digital transformation, highlighting the need for a more comprehensive approach that integrates technological advancements with socio-economic impact and ethical considerations.

The next section delves into the challenges and problems associated with the responsible use of DTs and the potential repercussions of failing to address this issue. Section 3 introduces a Responsible Digital Twins (RDT) framework. This framework aims to address ethical and socio-economic considerations in a standards-based, machine-readable format of policy expressions, providing a comprehensive approach to DT technology's responsible and ethical use. Section 4 discusses two application domains that represent different ends of the DT complexity spectrum. These serve to illustrate the diverse applications, potential impacts of DT technology in real-world scenarios and our RDT proposal. Section 5 provides concussions and future work directions.

2 Problem

While DTs present substantial opportunities to influence the digital transformation of organizations profoundly, it is important to acknowledge that, akin to other technologies such as AI and IoT, they can be utilized for both beneficial and detrimental purposes. This dual potential extends to impacts on humans, the environment, and institutional sectors such as healthcare and finance.

Enterprise or socio-economic rules, considered as constraints on behaviour of various actors in the DT ecosystems, are crucial in developing and deploying DTs across the spectrum from simple discrete through to complex composite DT systems. These rules can be described in terms of the primitive policy concepts, i.e. obligations, permissions, prohibitions, and authorizations, also known as deontic concepts [12], which in turn, can be combined to express more complex, accountability concepts modelling enterprise, legal or legislative policies. These rules, help ensure that DTs are aligned with enterprise policy, ethical norms and legislative/regulative policies. A comprehensive approach to integrating such rules in developing and deploying DTs is important to address interoperability challenges and conflicts related to the precedence or prioritization of the business and societal impacts of certain DT use cases [10].

This paper seeks to address a specific set of challenges for expressing enterprise or deontic rules in a machine-readable format. The primary objective is to develop a methodology that is both standards-based and scalable, capable of accommodating a wide range of applications, from simple, discrete systems to large-scale, ultra-complex system-of-systems configurations. It involves navigating the complexities of translating enterprise rules into a format that can be readily interpreted by DTs, while adhering to established standards. The ultimate goal is to create a practical framework that can effectively address the increasing complexity of DT applications, accommodating socioeconomic factors, thereby enhancing their utility and impact in various sectors.

3 Solution Approach

3.1 Motivations

Digital Twins are made possible by the combination of technology such as event-based processing and analytics, modelling and simulation, machine learning, and AI, as introduced above. DTs are essentially *technical systems* but involve close interactions and synchronization with *human actors*, in many respects like the SCADA systems, used in support of industrial systems, such as power, irrigation and water systems.

Enterprise and social policies – governance of the synchronization points
The close interactions between technical systems and actions of humans and the effect of decisions made by automated systems on humans, require careful analysis of *synchronization points* between physical and digital systems – with the aim of identifying the *enterprise or social policies* that need to be respected at these points and beyond. This would need to apply to the design, implementation, testing, operation and updates stages of DT components life-cycle, while in compliance with the organizational, regulative, legislative and policies reflecting safety and ethical standards and norms.

We propose the term 'responsible digital twins' to signify the explicit integration of these policies at DT's life cycle. Our approach is influenced by the increasing recognition of a need for supporting 'responsible AI' technologies, while adding the specifics arising from the broader set of DT technical and engineering characteristics.

We propose the following characteristics of DTs (Fig. 2), as their 'responsible' properties, where the first group below is influenced by the AI ethics principles [13][16]:

- Human, societal and environmental wellbeing - capturing the fact that DT systems should benefit individuals, society and the environment.
- Human-centred values - emphasizing the fact that DT systems should respect human rights, diversity, and the autonomy of individuals.
- Accountability – referring to the actions of parties involved in developing and deploying DT systems, including their responsibility for any harm that is caused by their systems.

- Transparency – referring to the ability of users to understand the operation of a complex system, such as a black box AI system, how their data are used and how decisions are made; referred to as explainability in the AI context.
- Contestability - enables consumers to challenge the output of the AI algorithm when it impacts them.
- Fairness - DT systems should be fair in their treatment of all users, regardless of their race, gender, religion, or other personal characteristics.
- Privacy - DT systems should respect the privacy of users and should not collect or use personal data without their consent.

We also believe that the 'responsible' properties cover engineering principles of:

- Security - DT systems should enable precise access control over specific data, resources, and actions in your deployment, also supporting privacy property.
- Reliability and Safety – DT systems should operate in accordance with their purpose and should not pose a risk to users' physical or psychological well-being; they should function well for people across different use conditions and contexts, including ones it was not originally intended for.
- Resilience - DT should absorb disturbance and reorganize while undergoing change thus retaining the same function, structure, identity, and feedbacks; this includes adaptation to system changes as technology and society evolve.

These characteristics (in outer layer of policies in Fig 1), are guiding principles, similar to ethics [13][16] or the Gemini principles [18]. Note that 'trustworthiness' is sometimes used to refer to security, privacy, safety, reliability, and resilience [9].

We also note that current consumer laws define safety and quality requirements for goods or services to minimise harm for consumers, but the specifics of consumer-facing uses of AI such as generative AI have not yet been considered by a court [16].

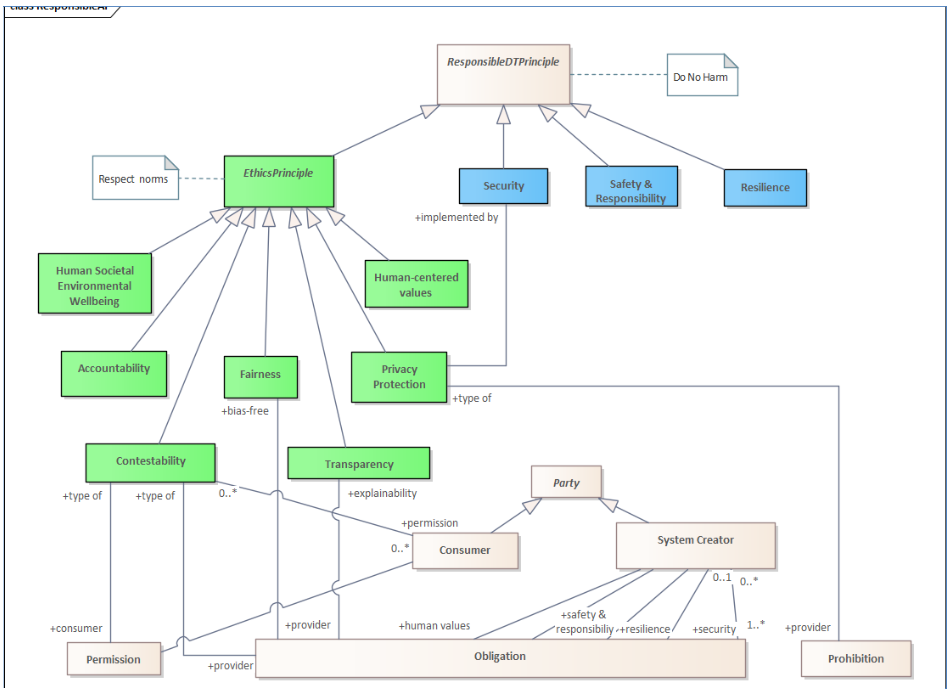

Figure 2: Classification of DT responsible principles and mapping onto deontic policies

Operationalization of principles

We need an approach for operationalizing high-level principles mentioned above. One option, similar to the AI responsible patterns [15], would be to develop a catalogue of reusable responsible DT patterns. They would provide reusable solutions to common problems occurring in system developments. Such approach was also used in the early interoperability framework in Australia [19] for identifying interoperability patterns.

Another approach is to express the enterprise and social policies associated with each of the responsible DT characteristics, using computable (machine readable) expression of foundational, i.e. deontic, policies of authorizations, obligations, prohibitions, permissions and their violations. This approach is introduced in [13] which describes a detailed mapping of digital ethics principles into deontic policy concepts. This was further elaborated in modelling the consent as an authorization policy [20]. This is the approach that we adopt for the operationalization of responsible DT principles, where the mapping of each of the principles can be refined into the fundamental deontic concepts (Fig 2). The approach is presented next.

3.2 Standard-based specification framework

Our solution for a computable expression of policies is based on the ISO/IEC 15414 Enterprise Language (EL) standard, from the family of Reference Model for Open Distributed Processing (RM-ODP) standards [12], augmented with related work from ontology research, in particular the conceptual modelling of legal relations [23] . The Enterprise Language standard provides precise expression and guidelines in the expression of *foundational* (deontic) policies of obligation, prohibition, permission, and authorisation, and the expression of *accountability* policies of the parties involved in the system, whether using, controlling or interacting with it. The accountability concepts are an element of responsible characteristics, and derived from deontic concepts, shown as middle layer of policies in Fig. 1. Both the foundational and accountability concepts can be used to express constraints on the actions of parties filling various roles in a system, be they humans or automated systems.

This standard thus provides foundations for computable expression of enterprise specifications for a system, which in our case is a digital twin ecosystem. This specification would typically involve defining:

- the purpose of a digital twin system in terms of behaviour of the system
  - individual components, their interactions, compositions etc
- policies that capture further restriction on the behaviour
  - between the system and its environment, or
  - within the system itself, related to the business decisions by the system owners
- explicit description of ecosystems that can span multiple policy domains (e.g. federation) and are not owned by a single party.

This specification style places greater emphasis on the expression of correct or normal behaviour and on the chain of responsibility involved in achieving it [12]. This in turn supports the expression of business rules and behaviour that clearly describe obligations, permissions, authorisation and prohibitions (the so-called deontic concepts), as well as the accounatbility of each of the objects involved in the specification, as explained next (Note: an object can represent an IT system or a natural person).

Deontic concepts

The EL standard includes the concepts of obligations, prohibitions and permissions, stating the constraints for actions that are obliged, prohibited or permitted. In addition, the standard provides concepts for modelling the dynamics of deontic constraints i.e. when they become applicable to the actions of parties and how they are passed among parties. These are needed for the governance, compliance and management of interactions between autonomous decision-making components and humans in a system. This is achieved by introducing a special type of enterprise object, called *deontic token*, which captures deontic assertions. The deontic tokens are held by the parties involved and holding one controls their behaviour [12]. Deontic tokens can be manipulated as objects while deontic constraints (e.g. obligation) cannot. There are three types of deontic tokens: *burden*, representing an obligation, *permit* representing permission and *embargo*, representing prohibition. In the case of a burden, an active enterprise object holding the burden must attempt to discharge it either directly by performing the specified behaviour, or indirectly by engaging some other object to take possession of the burden and performing the specified behaviour. In the case of permit, an object holding the permit is able to perform some specified piece of behaviour. In the case of embargo, the object holding the embargo is inhibited from performing the behaviour.

Another concept introduced to support modelling the dynamics of deontic constraints is *speech act*, Fig. 4. This is a special kind of action used to modify the set of tokens held by an active enterprise object. The name was chosen by analogy to the linguistic concept of speech act, which refers to something expressed by an individual that not only presents information but performs an action. Thus, a speech act intrinsically changes the state of the world in terms of the association of deontic tokens with active enterprise objects. This concept fits well with the nature of AI enabled applications, as it allows the speech act to be performed by people and AI systems, yet distinguish them when needed to establish links with ethics, legal and social norms.

Accountability concepts

The deontic modelling framework is further extended to support traceability of obligations of parties, according to their broader responsibilities derived from ethical, social or legal norms, referred to as a set of accountability concepts [12]:

Principal is a party that has delegated something (e.g. authorisation or provision of service) to another. *Agent* is an active enterprise object that has been delegated something (e.g. authorisation, responsibility of provision of service) by, and acts for, a party.

Delegation is an action that assigns something (e.g. authorisation, responsibility of provision of service) to another object, e.g. agent.

Additional action types, capture important business events in any organisational system, and model how responsibilities evolve.

Commitment, is an action resulting in an obligation by one or more participants in the act to comply with a rule or perform a contract. This effectively means that they will be assigned a burden. Examples are commitments by clinicians to deliver safe, reliable and effective healthcare to patients.

Declaration, is as an action by which an object makes facts known in its environment and establishes a new state of affairs in it. For example, an AI system (or a party managing it) may inform the interested parties about change of some legal rule.

Evaluation, is an action that assesses the value of something. Value can be considered in terms of various variables e.g. importance, preference and usefulness, such as performance parameters to express administrative performance, or accuracy or reliability measures associated with research findings or to assess the fairness of training data.

Prescription, is an action that establishes a rule. Prescriptions provide a mechanism for changing the system's business rules at runtime, enabling its dynamic adaptation to business changes, such as creation of new policies reflecting new legislations for AI.

Authorisation, is an action indicating some empowerment, through which an enterprise object issues a required and will itself undertake a burden to facilitate the behaviour. For example, the contestability is an authorisation for the consumer to challenge AI decisions, through a permit by the AI system which has the burden to enable it.

Deontic and accountability concepts are constraints over *actions* of the parties or systems (Fig.4), making it possible to define computable constraints over the actions, thus supporting real-time monitoring and downstream discretionary or non-discretionary enforcements. There are several policy languages for expressing such constraints, which is beyond the limits of this paper, and are for example discussed in [20].

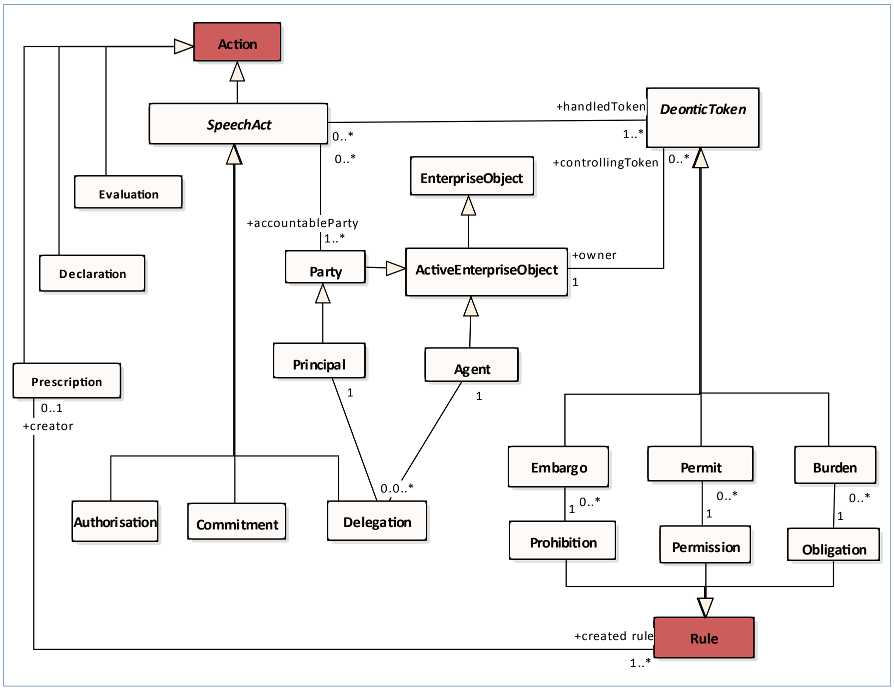

Figure 3: ODP Enterprise Language: deontic and accountability concepts

3.3 Practitioners-friendly

The computable policy framework provides a way of translating high-level characteristics of responsible digital twins (RDT) into a set of computable policies that can be integrated in the design, implementation and operation of DT systems.

In order to make this rather technical framework practitioner friendly, we are proposing a mapping framework as shown in Fig. 5. Physical entities are virtually represented by DTs that synchronise twinning information at a certain frequency and fidelity. This requires technical integration that is governed by technical architecture and integration capabilities for each digital twin use case. The collection or combination of all the digital twin use cases interoperates to provide operational, tactical, and strategic decision support. The digital twin of a complex system is not a single twin, but the harmonious operation of the collection of digital twin use cases. The framework introduces a Business Process Logic layer that enables stakeholders to manage how these use cases interoperate by providing operational rules, optimisation and AI models, and enterprise or socio-economic rules. This framework allows stakeholders to influence business outcomes by adjusting and prioritising the different rules and models.

This solution is result of our experience in architecting and building many industrial, financial or health systems, reflecting the needs to support customer specific or government specific requirements.

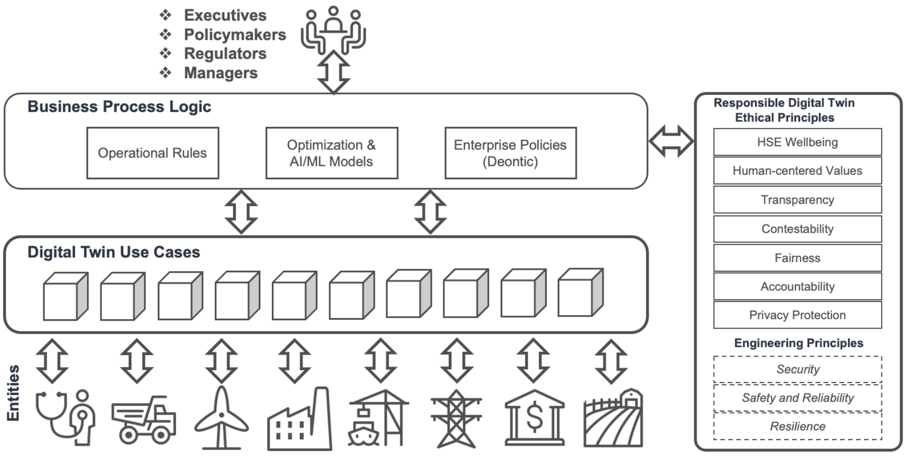

Figure 4 Adding Responsible DT layer to a DT system

4 Two application domain considerations

4.1 Digital health – personalized medicine

DTs are used to address several healthcare challenges, including surgery, pharmacy, cardiology and operating theatres [21]. They can be also used in support of personalized medicine where the bi-directional data flow between a patient and its virtual replica, (Fig. 6) allows for real-time continuous updating of the virtual model and, conversely, targeted interventions on the patient based on predictive simulations performed on DTs [22].

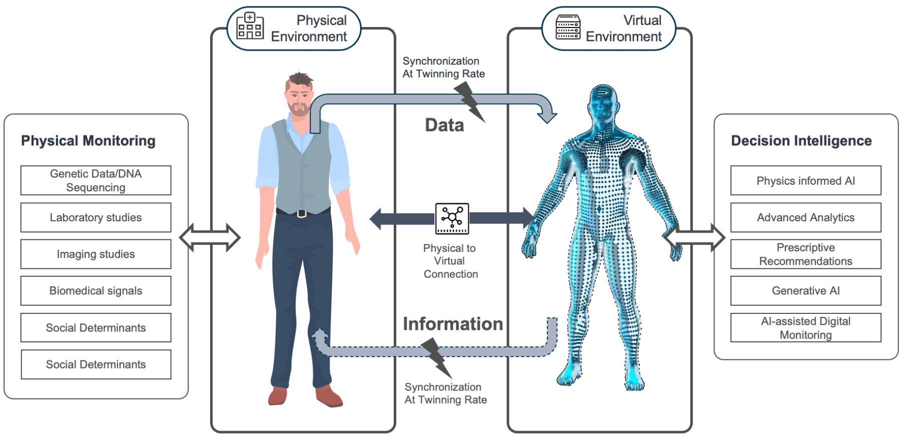

Figure 5: Single DT for a person

There are also DT applications across more complex healthcare systems which would involve multiple DTs, focusing on different aspects of healthcare. These would be for example DTs that are used as part of large hospital information systems, aimed at improving organizations' workflow and resource management. Such DTs can reference virtual copies of individual patients which can then be integrated into higher-level representation of clinical workflows in the hospital, effectively expressing interoperability between these twins, from the application perspectives. In the case of responsible DTs, these workflows need to explicitly embed rules that govern responsible DTs, reflecting rules associated with clinical care, but also rules emanating from organizational policies or external regulatory policies as mentioned in section 3. In all these settings, DT applications could improve the study and monitoring of highly complex systems characterized by many interacting components and thousands of variables that can be difficult to characterize with traditional approaches.

One interesting solution for single digital twin is their use in enabling better precision and personalized of dementia care [14]. In this case clinician first enters patient data to a mobile Decision Support System (DSS), which is linked to a server running the Machine Learning or Deep Learning Algorithm. The algorithm connects to the database containing data about past dementia cases to find one or more past cases that best match the data of the present patient. Algorithm then constructs the appropriate DTs through union of best matching cases and the DT and all related information are shown to the end-user via the mobile DSS. This helps clinician performs a more informed and precise diagnosis and treatment planning decision, after which the details and outcomes about new patient get recorded as new data for future reference.

4.2 Renewable energy

Renewable or distributed energy resources are increasingly employing digital twins to address many use cases across these complex assets' value chains and life cycle. A prime example of this is the energy grid challenge associated with introducing renewable or distributed energy resources into power systems networks.

The electricity grid represents an ultra-large scale complex cyber-physical system that merges engineering automation and control technology with emerging digitalization capabilities, such as digital twins. These digital twins provide decision intelligence, as shown in Fig 7 to facilitate decision support, augmentation, and automation. This integration of digital twins into complex systems underscores the transformative potential of digitalization in addressing contemporary challenges in the energy sector.

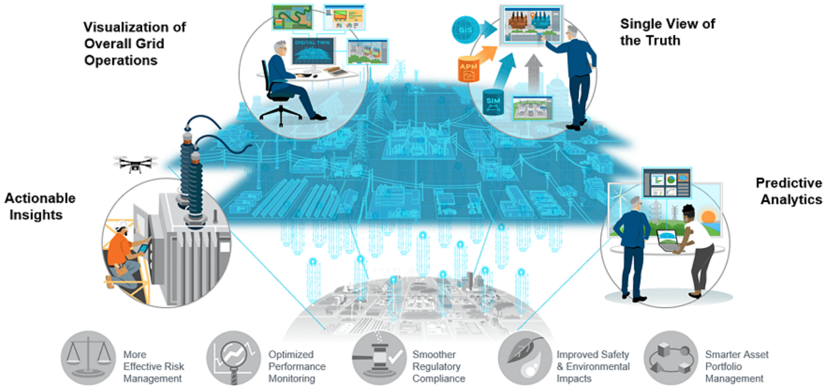

Figure 6 Digital Twins in a Sustainable Energy Grid [24]

DTs, with their extensive applications across the value chain, establish a complex system within the renewable energy sector, including applications such as wind farms, hydroelectric power, biomass, and green hydrogen production. They facilitate process and network optimization, thereby augmenting efficiency in wind turbine operations or hydroelectric power generation. They employ predictive analytics to forecast asset performance and maintenance requirements, reducing downtime in biomass processing facilities or hydrogen production plants. They contribute to quality management and monitor Environmental, Social, and Governance (ESG) factors, endorsing sustainable operations across all renewable energy applications. They identify potential operational risks and enhance safety through scenario simulations, improving the safety protocols in wind farms or hydroelectric plants. Moreover, they optimize supply chain operations and aid in workforce management by providing real-time visibility and predictive insights, streamlining operations across the renewable energy sector.

The combination of different use cases in each of the application areas described across the overall value stream of generation, transmission, and distribution for power systems results in a complex digital twin system of systems.

The development and deployment of DTs in the renewable energy sector present an even more complex task, given the extensive continuum of decision support required across the value chain through the combination of traditional and emerging technology. The temporal scope of use cases can vary significantly, from long-term energy market optimization to near real-time operational responses. This variability necessitates consistent and coherent policy frameworks across diverse use case categories and time considerations.

Operational rules across these use cases are often encapsulated as first-order logic rules, readily expressed in a machine-readable format for digital twins. These rules facilitate decision support, decision augmentation, and in some instances, business process automation based on the model as shown in Fig 5 that depicts adding a responsible DT layer to a DT system.

Statistical and mathematical models for simulation, optimization, and prediction, typically code-based, are also machine-readable and can be embedded within digital twin use cases. However, deontic and accountability rules, are often overlooked during the requirements-gathering phase. This oversight may lead to interoperability challenges and potential conflicts regarding the precedence or prioritization of the business and societal impacts of digital twin use cases.

Contrary to operational rules and AI/ML models, deontic and accountability rules are infrequently presented in an explicit machine-readable format. This highlights a potential area for improvement in the responsible use of DTs, as proposed by the authors, underscoring the need for a more comprehensive approach to incorporating such rules in DT decision-making processes for renewable energy use cases.

The following scenarios illustrate conflicting operational and enterprise rules:

- An operational rule may stipulate that additional battery storage is required to stabilize the national grid. However, this could lead to increased mining activities for the raw materials necessary for battery production, subsequently resulting in a rise in reportable (scope 1, 2, and 3) emissions. This outcome could conflict with a public commitment to reduce the reportable emissions on a year-by-year basis towards a 2030 goal. Such a policy statement might be communicated through press releases and investor briefings, but it is not typically presented in a machine-readable format that could inform the operational digital twin of the potential conflict, which our proposal can address.
- An AI model optimizing a renewables-based energy grid may necessitate consumer data regarding behaviour and energy utilization patterns, potentially indicating anomalies for certain users. This use case may interoperate with a revenue optimization DT use case, making anomalous event data accessible to service agents. These agents could then utilize this information to target specific individuals or organizations. However, without explicit machine-readable privacy deontic policies, such applications of DT could occur irresponsibly, without the consent of both organizational and user stakeholders. This underscores the critical need for comprehensive privacy policies in the deployment and operation of DTs, to ensure their responsible and ethical use.
- The World Economic Forum [27] underscores how recycling to conserve metals utilized in wind turbines represents one pathway toward a circular economy. The decommissioning and disposal of such assets constitute significant asset management use cases for DTs within the overall value chain, particularly towards the end of the physical asset's life cycle. Operational rules may dictate that the most cost-effective approach would be to abandon or dispose of these assets in landfills rather than invest in an active recycling policy. However, implementing such a policy could override operational decisions made through these DTs, highlighting the potential for DTs to contribute to sustainable practices in the renewable energy sector.

The digital health and renewable energy scenarios underscore the imperative for a responsible DT (RDT) framework, accommodating RDT principles, and including machine readable and standards-based policy expressions, both of which provide a comprehensive guide and support for the ethical development, deployment, and operation of DTs, as shown with the examples in Table 1 next. An RDT framework will ensure that DTs contribute positively to individuals, society, and the environment.

Table 1: Responsible DT (RDT) principles – Healthcare and Renewable examples

RDT principle	**Health**	**Renewable**
Human-centered	Dementia DT should respect human rights, diversity, and the autonomy of individuals, e.g. obtaining and recording patient consent for using their data as part of the Dementia DT system, taking into account their specific demographic data as well as their autonomy in making decisions - or relying on a delegated person who make decision on behalf of patient	The model should not be used to target specific individuals or organizations in a discriminatory way. For example, the model should not be used to target low-income households or minority communities for higher energy prices.
Accountability	Developers are obliged to develop and deploy the DT system, which support patient preferences, and their changes over time; they are obliged to minimise any harm that is caused by their systems; clinicians are obliged to check the recommendation by the virtual DT and make final professional decision before they are permitted to prescribe a medication	It is important to have clear processes in place to hold those responsible for developing and deploying the AI model accountable for their actions. This could include having a board of directors or an ethics committee that oversees the development and deployment of the model.
Transparency	Dementia DT should provide mechanisms to both patients and clinicians showing how ML/AI/DSS components arrive at a decision treatment, and also state clinical risks and benefits for the person in question, taking into account their medical history, demographics, and other parameters	Model workings should be transparent to individuals and organizations. This could include providing information about the data that was used to train the model, the algorithms used to make decisions, and the potential biases that could be present in the model.
Fairness	Dementia DT systems should be fair in their dementia care support to all users, regardless of their race, gender, religion, or other personal characteristics.	Model should not discriminate against some groups. For example, it should not target low-income households or minority communities for higher energy prices.
Privacy	Dementia DT should respect the privacy of users, and should not collect or use personal data without their consent; consent should be regarded as a combination of permis-	The data about consumer behavior and energy utilization patterns could be sensitive and used to identify individuals or organizations. Clear privacy policies

	sion (for patients) and obligation (for clinicians' respecting patient preferences) of accessing patient health records, i.e fine-grained data access support	should explain how data will be collected, used, and shared. There should be mechanisms for individuals to control their privacy settings and opt out of data collection.
Safety and reliability	Dementia DT should operate in accordance with their purpose, to support dementia patients and should not pose a risk to users' physical or psychological well-being	The model should not make decisions that could endanger people or property. The model should also be resistant to hacking and manipulation.

5 Conclusions and Future Work

This paper presents our proposal for starting new inquiry into responsible DTs, by explicitly positioning of enterprise and social policies in the context of DT technologies. The aim is to help practitioners with designing, building, operating and evolving responsible DTs which embed computational expression of such policies, while balancing value proposition and risks. We find that there are still some legal and ethics ambiguities about the chain of responsibilities involving humans and automated decision makers, and it is the level of risks that can determine best governance mechanisms for responsible DTs, as also discussed in [16]. We argue that the concept of responsible DTs may have been overlooked so far due to the focus on technical issues such as interoperability and composability. We believe that our proposal for computational expression of policies, implementing socioeconomic constraints for responsible DT principles, can create interest by developers involved in building tools for DT solutions and provide valuable guidelines to practitioners helping end-users on this specific digital transformation journey. We are also hoping that the deontic-based formalism used in our approach, and based on the ODP Enterprise Language, can provide a new perspective on the formalization of architecture principles in general [28], and RDT principles in particular.

In future, we are planning to develop detailed proof of concept prototypes involving end users in renewable energy, digital health, but also manufacturing, finance, supply chain. These would need to include expressive and machine-readably policy language to operationalize the RDT principles, which we presented elsewhere [13][20], but the elaboration of which was beyond the space limitations of this paper. Such a policy language could be implemented by a separate DT instance, further allowing simulation of complex policy interactions, to detect policy conflicts. This in turn could support real-time monitoring and enforcement of actions of parties according to their legal, ethical or professional policies.

Another area of investigation would be how to model various value functions to capture business and social objectives, and use them to resolve conflicts of policies when composing DTs across complex systems. This would help in addressing ethics dilemmas and provide further support to humans when dealing with conflicting policies, including how to implement complexities associated with monitoring of obligations and prohibitions in case of trade-offs between the compliant behaviour and cost of violations, as discussed in [26].

Another topic is to bring in elements of legal concepts and their relationships, captured through respective ontologies as discussed in [23]. For example, the concept of rights, signifies permission of some actors, but puts an obligation on others, on opposing side of the relationship, and these 'correlatives' may provide more tighter accountability expressions over our accountability modeling concepts. Further, the concept of liability, such as one discussed in [27], can be related to an obligation of a party who is expected to perform some action, which it fails to perform (i.e. violation), but may also have links to the concept of power.

References

1. Digital Twin Consortium. (2023). Working Groups. Retrieved from https://www.digitaltwinconsortium.org/workinggroups.htm
2. IoT Analytics (2023). Decoding Digital Twins: Exploring the 6 main applications and their benefits. Retrieved from https://iot-analytics.com/6-main-digital-twin-applications-and-their-benefits/
3. Markets and Markets (2023). Digital Twin Market by Application, Industry, Enterprise, and Geography – Global Forecast to 2028. Retrieved from https://www.marketsandmarkets.com/Market-Reports/digital-twin-market-225269522.html
4. 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.
5. Packt. (2023). Building Industrial Digital Twins. Retrieved from https://www.packtpub.com/product/building-industrial-digital-twins/9781839219078
6. XMPro. (2023). Why Decision Intelligence with Digital Twins is Kinda Like DCS for Automation and Control. Retrieved from https://xmpro.com/why-decision-intelligence-with-digital-twins-is-kinda-like-dcs-for-automation-and-control/
7. Digital Twin Consortium (2022). Capabilities Periodic Table. Retrieved from https://www.digitaltwinconsortium.org/initiatives/capabilities-periodic-table/
8. Digital Twin Consortium (2023). Platform Stack Architectural Framework: An Introductory Guide Form. Retrieved from https://www.digitaltwinconsortium.org/platform-stack-architectural-fram-formework-an-introductory-guide-form/
9. IDTA (2023). Submodels. Retrieved from https://industrialdigitaltwin.org/en/content-hub/submodels
10. McKinsey (2022). Digital twins and the enterprise metaverse. Retrieved from https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/digital-twins-the-foundation-of-the-enterprise-metaverse.
11. W. Chen, Z. Milosevic, F. A. Rabhi and A. Berry, "Real-Time Analytics: Concepts, Architectures, and ML/AI Considerations," in IEEE Access, vol. 11, pp. 71634-71657, 2023, doi: 10.1109/ACCESS.2023.3295694.
12. ISO/IEC 15414, *Information technology: Open distributed processing, Reference model – Enterprise Language*, 3rd ed, 2015

13. Z. Milosevic, "Ethics in Digital Health: A Deontic Accountability Framework," *2019 IEEE 23rd International Enterprise Distributed Object Computing Conference (EDOC)*, Paris, France, 2019, pp. 105-111, doi: 10.1109/EDOC.2019.00022.
14. Nilmini Wickramasinghe and others, Digital twins to enable better precision and personalized dementia care, *JAMIA Open*, Volume 5, Issue 3, October 2022, ooac072, https://doi.org/10.1093/jamiaopen/ooac072
15. Responsible AI Pattern Catalogue, CSIRO/Data61, https://research.csiro.au/ss/science/projects/responsible-ai-pattern-catalogue/
16. Safe and responsible AI in Australia, Discussion paper, June 2023, Australian Government, Department of Industry, Science and Resources
17. S. V. Nath, P. van Schalkwyk and D. Isaacs, Building Industrial Digital Twins: Design Develop and Deploy Digital Twin Solutions for Real-World Industries Using Azure Digital Twins, Birmingham, U.K.:Packt, 2021.
18. Gemini Principles, University of Cambridge, 2022, https://www.cdbb.cam.ac.uk/DFTG/GeminiPrinciples
19. Interoperability Framework 2.0, https://developer.digitalhealth.gov.au/resources/interoperability-framework-v2-0
20. Milosevic, Z., Pyefinch, F. (2022). Computable Consent – From Regulatory, Legislative, and Organizational Policies to Security Policies. EDOC 2022. Lecture Notes in Computer Science, vol 13585. Springer, Cham. https://doi.org/10.1007/978-3-031-17604-3_1
21. Hassanzadeh, Hamed; Boyle, Justin; Khanna, Sankalp. Digital twins in healthcare: An operating theatre flow modelling case study. In: Digital Health Institute Summit; Brisbane. AIDH; 2022. -. csiro:EP2022-2887. http://hdl.handle.net/102.100.100/444330?index=1
22. Digital Twins: The New Frontier for Personalized Medicine?, Appl. Sci. 2023, 13(13), 7940, https://www.mdpi.com/2076-3417/13/13/7940
23. C Griffo, JPA Almeida, G Guizzardi, Conceptual modeling of legal relations, Proc. 37th International Conference, ER 2018,
24. Bentley (2022). Digital Twins for a Sustainable Grid. Retrieved from https://blog.bentley.com/digital-twins-for-a-sustainable-grid/
25. WEF (2022). *3 ways the circular economy is vital for the energy transition*. Retrieved from https://www.weforum.org/agenda/2022/02/3-ways-circular-economy-renewables-energy-transition/
26. P.F. Linington, Z. Milosevic, A. Tanaka and A. Vallecillo, *Building Enterprise Systems with ODP, An Introduction to Open Distributed Processing*, Chapman & Hall/CRC Press, 2011.
27. *Report on the safety and liability implications of Artificial Intelligence, the Internet of Things and robotics*, EUR-Lex, 19/02/2020 https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020DC0064
28. D. Greefhorst and H. A. Proper. Architecture principles - the cornerstones of enterprise architecture. The Enterprise Engineering Series. Springer, Heidelberg, Germany, 2011. ISBN: 978-3-642-20278-0