RESEARCH-ARTICLE

Human Digital Twin for Healthcare Applications: a White Label Digital Twin Implementation

LEONARDO MICELLI, University of Urbino Carlo Bo, Urbino, PU, Italy

SARA MONTAGNA, University of Urbino Carlo Bo, Urbino, PU, Italy

Published: 03 September 2025

Citation in BibTeX format

GoodIT '25: International Conference on Information Technology for Social Good
September 3 - 5, 2025
Antwerp, Belgium

Conference Sponsors:
SIGCAS

GoodIT '25: Proceedings of the 2025 International Conference on Information Technology for Social Good (September 2025)
https://doi.org/10.1145/3748699.3749778
ISBN: 9798400720895

Human Digital Twin for Healthcare Applications: a White Label Digital Twin Implementation

Leonardo Micelli
Department of Pure and Applied Sciences
University of Urbino Carlo Bo
Urbino, Italy
leonardo.micelli@uniurb.it

Sara Montagna
Department of Pure and Applied Sciences
University of Urbino Carlo Bo
Urbino, Italy
sara.montagna@uniurb.it

Abstract

The concept of digital twin has become a pillar of healthcare research over the last decade. By mirroring real systems through a range of functionalities–from tracking the state of physical assets to providing cognitive services–it enables real-time analyses, thereby supporting informed decision-making.

In this paper we are specifically interested in the human digital twin as one of the main component into the personalised healthcare framework. Even though literature is full of examples discussing this issue, a unified reference model and supporting technology are still lacking. As a result, a substantial gap remains between theoretical developments and practical implementation, with real-world applications still limited.

Accordingly, this paper presents a human digital twin model, along with a practical implementation on the White Label Digital Twin platform, designed to be compliant to the reference standards and general enough to be specified and instantiated for different healthcare domains. As a proof of concept, we present an application in the context of hypertensive patient care where data are acquired from two different data sources and integrated into a unique standardised model.

CCS Concepts

• **Applied computing** → **Health informatics**; *Health care information systems*; • **Human-centered computing** → *Ubiquitous computing*.

Keywords

Human Digital Twin, Personalised Health, White Label Digital Twin

ACM Reference Format:
Leonardo Micelli and Sara Montagna. 2025. Human Digital Twin for Healthcare Applications: a White Label Digital Twin Implementation. In *International Conference on Information Technology for Social Good (GoodIT '25), September 03–05, 2025, Antwerp, Belgium.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3748699.3749778

1 Introduction

The concept of Digital Twin (DT) [4] – namely virtual representation of a physical system continuously updated with real-time data – has been defined and applied across various domains, including manufacturing, aerospace and, more recently, healthcare [18, 19].

In particular, DT that mirrors human being, *i.e.* Human Digital Twins (HDTs), have been referenced in several contexts. For instance, in Industry 5.0 they are introduced as enablers of human-centric, intelligent systems [24].

In healthcare, the emergence of HDT reflects the growing interest in exploiting DTs to model and monitor the human body for personalised medicine and improved clinical decision-making [1]. In this work, we focus specifically on the use of HDTs within this domain, with a particular emphasis on the acquisition and integration of multimodal data that span multiple layers associated with human physiology, pathology and psychology. These include genomic data, various types of medical imaging (e.g., X-ray, MRI), laboratory results, clinical data from electronic medical records, patient histories, physical examination findings, behavioural and cognitive tests.

Despite the increasing number of publications on HDTs in healthcare, most existing contributions remain generalist in nature. The current literature often introduces the concept without providing a precise definition, a robust reference model, or a standardised framework to represent and manage the complexity of the human organism in a digital environment. To address this gap, this article aims to propose a reference model for the integration of multimodal healthcare data, emphasising the importance of preprocessing and standardisation. We then introduce a design framework for HDTs with a focus on healthcare and well-being applications. We finally show an implementation on the White Label Digital Twin (WLDT) platform. Our goal is to go beyond conceptual discussions and move towards actionable frameworks that support the development of interoperable, scalable, and clinically useful HDTs.

As a proof of concept, we present a case study focused on hypertensive patient management, demonstrating the feasibility and utility of the proposed model through the design in a real-world clinical scenario. Through this use case, we show how the HDTs model we devised can be flexibly specified according to the available data which, in this specific scenario, are acquired from two different sources. In particular, we implemented a full-stack HDTs that acquires vital sign measurements from wearable technologies and user conversations from a chatbot. Conversations are appropriately preprocessed to extract mood information, which is then integrated with the other data into our unified framework based on the reference FHIR standard.

GoodIT '25, Antwerp, Belgium

ACM ISBN 979-8-4007-2089-5/25/09
https://doi.org/10.1145/3748699.3749778

2 Background & Motivation

The first seed for DT applications was planted in the 1970s in the NASA Apollo program, where a first digital replica of an actual system was created for remote monitoring and simulation purposes [12]. One of the first areas where the DT concept was adopted was Industry 4.0, which has been the main driver for the growth in research and implementation of DT-based solutions [14, 22]. Current research has shown that DT is an enabling technology for automation and intelligence [8], and it is increasingly being adopted in other domains such as transportation [6].

The concept of DT has been introduced also in healthcare [3, 5, 9], where it holds the premise to revolutionise the way healthcare services are managed and delivered —e.g. by modelling healthcare facilities assets and processes, to support prediction, optimisation and failure analyses.

Moving from assets to humans, they have been proposed as a technique to support the prevention and treatment of diseases. Much of the existing research explores the use of DTs in personalised medicine [1], where patient-specific models are developed to integrate real-time data, simulate disease trajectories, and define individuals' unique therapeutic strategies [20]. In this context, DTs can be valuable by allowing to model human systems, enabling remote monitoring, simulation and AI model training. For instance, [21] adopted a DT-based system to manage hypertension among people with Type 2 diabetes (T2D). In this study, 319 participants with T2D were divided into two groups: one with DT-assisted treatment and the other with standard care (SC). For the group that received DT assisted treatment, a specific set of parameters was collected, including age, sex, weight, systolic and diastolic blood pressure (monitored through IoT data). The study demonstrated that the DT-supported group achieved significant reduction to systolic and diastolic blood pressure compared to the SC group, showing great promise for DT-assisted treatment. Moreover, the literature reports several examples of DTs for different individual organs or specific physiological systems, which are highly detailed yet very use-case specific. For example, some research efforts have successfully created heart models derived from echocardiography scans, which simulate the structure and function of the heart [15]. These heart models are used to predict outcomes such as risk of heart failure, enabling personalised treatment plans and advancing the potential of precision medicine.

However, such models still represent isolated elements of the human body and do not contribute to creating a holistic perspective. The complex nature of human beings is a concrete challenge in modelling and engineering DTs that include all the necessary information to be useful in medical contexts. Moreover, most of the work is focused on the ancitipatory performances of DTs. These studies primarily aim to enhance the cognitive component of DT, evaluating its accuracy to predict future specific health conditions or anticipate the effects of treatments. As such, research tends to focus on the decision-making capabilities or the algorithmic side of DTs, where mostly machine learning (ML) algorithms and simulation are used. However, for a comprehensive and effective decision-making, the challenge lies not only in predictive power but also in how multimodal data are preprocessed and effectively integrated and represented in the DT model itself.

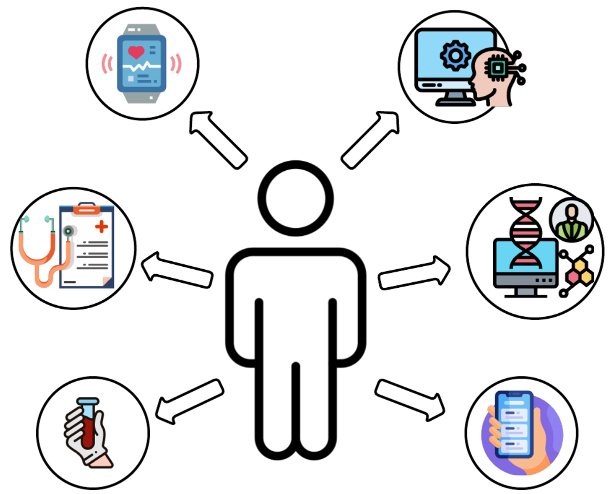

Figure 1: An effective HDT framework that handles multimodal data. Icons from flaticon.com

Accordingly, while previous work contribute significantly to the decision-making process in healthcare, they miss the discussion on the actual modelling and integration aspects that are essential to fully envision the HDT. As discussed in [23], this highlights the need for an interoperable model for HDT that is generic enough to facilitate adoption and collaboration in a wide range of domain, and also extensible enough to allow for highly use-case-specific customisation. Furthermore, the implementation of DTs in clinical practice has yet to move beyond theoretical frameworks. As a result, although DTs show promise for healthcare applications, particularly in personalised medicine, there remains a significant gap in terms of fully integrated, real-world implementation that support the holistic view of the HDT.

3 Human Digital Twin

The contribution of this work is to propose a unifying and generic framework to define and build DT-based ecosystems in the context of human health, allowing multimodal data collection, harmonisation and standardisation from different data sources – as shown in Figure 1 – and embedding them into an extensible model for digital human representation that builds on a set of predefined main components, The idea is that the same model and architecture can be adopted for different healthcare applications by appropriately extending from our proposed model. Moreover, the framework should be flexible enough to include a set of configurable cognitive tools that can be adopted to build health-related services, such as remote monitoring and decision making, in different healthcare domain.

Accordingly, in this section, we propose a general framework to model and build HDTs. First, we are presenting the framework architecture and then the conceptual model to represent a human being in the context of DTs.

In the literature, there are several reference models and metamodels. In this paper, we adopted the Web of Digital Twins metamodel as a reference [16], as it allows us to capture the level of flexibility

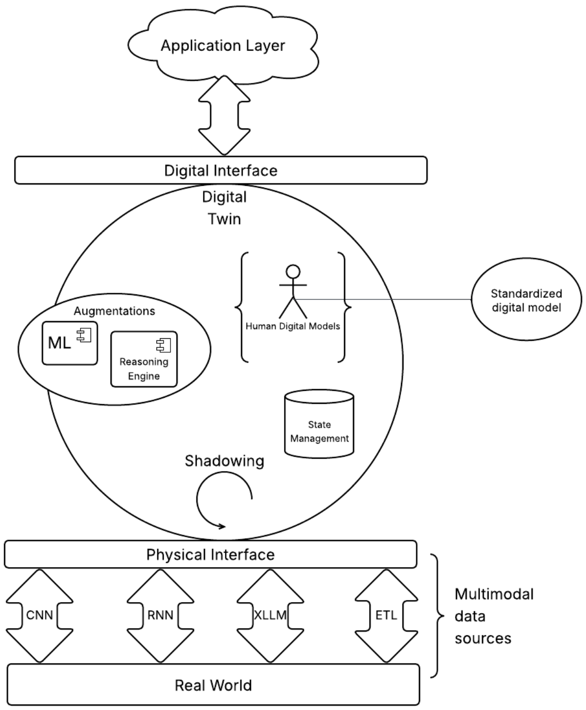

Figure 2: High level architecture.

and heterogeneity required for a HDT. Accordingly, the proposed framework adopts the abstractions of the WLDT model [13], which is grounded in the aforementioned metamodel. WLDT is a JVM-based open source general purpose library that allows developers to create DTs in terms of modular, adaptable, and interoperable software agents [1]. It was built in response to the lack of standards or common agreements for DT design and development that has led to the proliferation of several platform-specific solutions. Instead, WLDT aims to offer a flexible and extensible framework where a DT is an independent software agent implementing all the features and functionalities of its physical counterpart and it can be deployed and executed in the cloud or at the level of edge computers.

Accordingly, the proposed HDT framework is designed to be as generalisable as possible, allowing its adaptation and specification for any healthcare setting.

3.1 HDT Framework Architecture

Following the result of the survey reported in [8], the general structure of the proposed HDT framework architecture includes the main components of a DT, as for the representation available in Figure 2. The architecture is composed of three main elements:

3.1.1 ***Physical Interface***. This layer is responsible for communicating with the real world: here real-time, raw data are collected and pre-processed from different types of sources into a standardised format. In particular, since the goal is to collect data from a wide range of sources —including medical devices, electronic health records, clinical notes, and wearable sensors, the architecture is designed to include a dedicated layer for integrating diverse pre-processing algorithms. This modular layer allows for the incorporation of specialised tools according to the nature of each data source. For example, large language models (LLMs) can be adopted to convert unstructured text from clinical notes or conversations, e.g. Q-A chatbots that allow patient-physician interactions, into structured clinical data. Similarly, Convolutional Neural Networks (CNNs) can be employed to extract clinically relevant features from diagnostic imaging, and Recurrent Neural Networks (RNNs) from biomedical signals. All contribute to potentially capturing critical information necessary for constructing a comprehensive HDT model. The freshly processed data are then propagated to the upper layer, the Digital Twin layer, in the form of physical events e_{PT} that are processed to update the status of the DT and reflect the collected information.

3.1.2 ***Digital Twin Layer***. This is the core layer of the architecture, and is composed by three main elements:

The model set M. It is responsible to define a digital representation of the physical human. Since this is a set, the proposed framework allows to have multiple valid digital representations of the human. This allows to represent and store information with different abstraction levels and standard formats, adapting the framework to many different use cases.

The storage set S. It is responsible to manage the storage of the HDT state. Again, the framework allows for multiple storage managers, such as $\forall m \in M$, $\exists S_m \subseteq S$. This means that the framework allows the developers to have multiple storages for each digital model defined, or not at all.

Augmentations. This set encompasses any augmented capability the DT can have. This includes predictive models, reinforcement learning models that are fed with the physical human actual data, reasoning engines that utilises a semantic digital model that is kept up to date by the shadowing function. These augmentations are essentially special functions that are triggered by the shadowing process, and the trigger can be either upon receiving an event e_{DT} from the digital interface, a timed event firing periodically, or a reactive trigger that fires whenever a particular state of a model is updated.

The Shadowing Function. It is responsible for listening to the Physical Interface for any physical event e_{PT} and updating the state of the DT accordingly, via one or more $s \in S$, having also the responsibility to keep the information stored in different models to be consistent. The shadowing function also listens for any digital event e_{DT} coming from the Digital Interface. These events can trigger different actions from the DT: for example exposing the history of the status of the twin via API or invoke an augmentation.

3.1.3 ***Digital Interface***. This layer is responsible for exposing the Digital Twin API to external services. It allows applications to: actively request information about the DT status in a rest api-like fashion, receive information asynchronously from the DT, and invoke actions on it, if available. Upon request, the digital interface

[1] https://wldt.github.io

Parameter Type	Derived Parameter	Description
Behavioural Parameter	Social Interaction	Describes human interaction in the physical world
	Activity	Describes human physical activity (movement, facial expressions, running, etc.)
	Lifestyle	Describes behaviours that take place throughout the day
	Cognitive performance	Describes human competencies in terms of, for instance, knowledge, skills, abilities and aptitudes
Physical Parameter	External Parameter	Observable characteristics such as age, weight, etc.
	Physiological Parameter	Parameters measured using devices (e.g., heart rate, temperature)
Environmental Parameter	–	Describes the context in which the human operates and lives

Table 1: Parameter types, their derived parameters, and corresponding descriptions.

generates an event e_{DT} that the shadowing function will process to validate the request and generate a response.

3.2 HDT Domain Model

In the healthcare context, a human being can be modelled using a set of measurable parameters. According to [8], these parameters can be part of three different categories:

(1) *Physical Parameters*: describes quantifiable and measurable aspects of the human body;
(2) *Behavioural Parameters*: describes the human's habits, behaviour, mood and social interactions;
(3) *Environmental Parameters*: describe the context the human operates and lives in;

These categories can be further broken down into more detailed subcategories: for example, physical parameters can be broken down into *External Parameters* such as gender, age, and height and *Physiological Parameters* such as heart rate, systolic and diastolic blood pressure, etc. These categories are sufficiently general to be useful in different healthcare contexts, from telemedicine to operational clinics, and can be extended further to suit specific use cases, as reported in Table 1.

Since this domain model focuses on capturing the multifaceted nature of the real world, and it is designed to handle heterogeneous knowledge and information, potentially related to different categories, we need sufficiently expressive forms of knowledge representation, grounded in reference ontologies. Moreover, in order for the HDT internal model to be fully interoperable with other applications, there is the need to have a representation of this model that is expressed through an internationally recognised standard.

As such, knowledge graphs (KGs) can be adopted to represent semantic relations among data by structuring complex, interconnected data, thus improving integration, reasoning, and scalability. For example, a KG can integrate data from electronic health records, genetic studies, diagnostic imaging and wearable devices, enabling more accurate and personalised diagnoses. Additionally, KGs help identify relationships between diseases, treatments, and medications, supporting clinical decisions based on a comprehensive view of the patient [2]. While existing knowledge structures and ontologies provide a foundation, our approach remains open to the flexible integration into the KG of emerging aspects as new knowledge becomes available. Furthermore, the model can be customised and extended by incorporating domain-specific reference ontologies. On the other side, the FHIR standard[2] can be exploited to formalise the model and interoperate with other healthcare services. Moreover, the LOINC standard[3] can be adopted for identifying health measurements, observations, and documents.

However, one key element of the proposed framework is that we can store different digital models, all representing the same physical human. This mechanism allows us to represent the same human following any healthcare standard, maximising adaptability and interoperability. In particular, we envision at least two models: The first model focuses on the acquisition and organisation of raw data, such as vital signs, activity levels, or other domain-specific measurements. These data points are collected in real time and serve as the foundational layer, representing the direct interface with the physical world. This model, for instance, is the most convenient to operate with machine learning algorithms. The second model builds on this raw data by modelling into a KG the relationships and interdependencies among the variables, using formal semantics or ontologies to support interoperability and enabling reasoning over the domain.

3.2.1 HDT. In WLDT, the DT's model allows capturing and representing the Physical Asset (PA) at an appropriate level of abstraction, i.e., avoiding irrelevant aspects for its purpose and modelling only domain-level information rather than technological ones. This is achieved through the following elements:

(1) *adapters:* In WLDT, adapters can either be *Physical*, encapsulating an interaction logic between the DT and the PA, or *Digital*, encapsulating the interaction between the DT and the digital world. The set of all the physical adapters a DT has forms the Physical Interface of the DT, whereas the set of digital adapters forms the Digital Interface of the DT. In our HDT framework different physical adapters may be used to model data acquisition from different sources (i.e. IoT data, image processing, clinical notes analysis), whereas different digital adapters can be used to model the services that the HDT exposes to the external world (i.e. REST APIs, asynchronous data streaming etc.);
(2) *properties:* represent the observable attributes of the corresponding PA as labelled data whose values can dynamically change over time, in accordance with the evolution of the PA's state. In the HDT context, properties are used to reify one or more models of the physical human, enabling the modelling of medical concepts such as blood pressure, age, weight and heart rate as represented in Table 2. Accordingly, in WLDT the concept of *property* reifies for all the parameter categories listed in Table 1;

[2]https://www.hl7.org/fhir/overview.html

[3]https://loinc.org

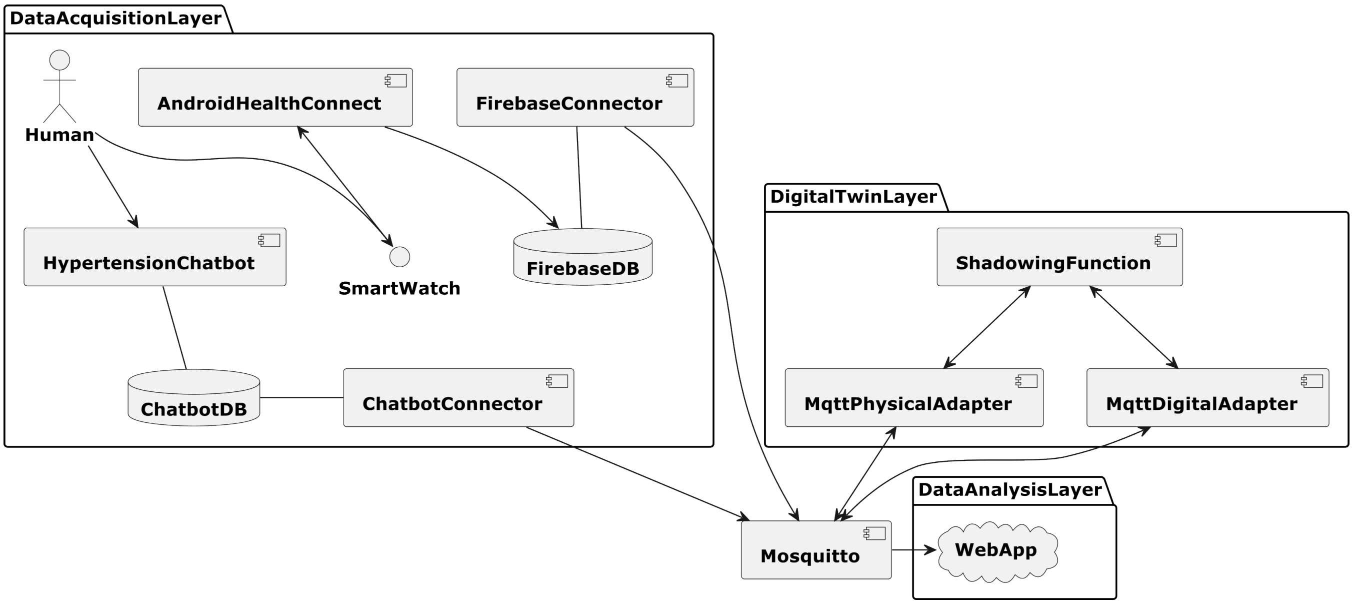

Figure 3: Use Case system architecture.

(3) *events:* in WLDT, events are used as the base form of communication between each component of the framework, mainly between the *Shadowing Function*, the adapters and the augmentations (which will be discussed later). Events can also be used to represent the domain-level physical events e_{PT} that can be observed in the PA: in the HDT context, this can mean modelling concepts like "seizures" or "being at hypertension risk". These events can be automatically fired by the HDT to asynchronously notify, via its digital interface, the correct caretakers.

(4) *relationships:* represent the links that exist between the modelled PA and other PAs of the organisations through links to their corresponding Digital Twins. Like properties, relationships can be observed, dynamically created, and change over time, but unlike properties, they are not properly part of the PA's state: instead, they are part of its operational context. In the HDT framework this mechanism allows to model the relationship between the physical human and its medical devices, or used to create more high-level digital twin ecosystems similarly to what is envisioned in [17];

(5) *actions:* represent the actions that can be invoked on the PA through interaction with the DT or directly on the DT if they are not directly available on the PA (the DT is augmenting the physical capabilities). *Digital Actions*, i.e. actions triggered from the digital interface of the DT, can model the request of AI-based prediction of the next state of the HDT coming from an HTTP request;

(6) *augmentation functions:* represent special actions that augment the DT beyond it's classical capabilities of data collection and state mirroring. In the HDT context this can include AI models that can perform state prediction based on the collected history of data and past states of the HDT, reasoning engines that leverage a semantic model of the HDT to perform inference and semantic reasoning.

4 Use Case

Chronic disease management faces different challenges, including the integration of large and multimodal health data (such as medical records, genomics, and real-time vital signs), low patient adherence to treatment and lifestyle indications, and limited access to specialised care—especially in remote or low-resource settings. These issues often lead to fragmented care, delayed diagnoses and, especially, poor long-term outcomes [7].

In this section, we present a full-stack system based on the presented framework that enables remote monitoring, data collection and analysis of different information related to the management of hypertensive people. In particular, the first prototype [4] we are presenting in this paper is focused on the hypertension disease and enables the automatic acquisition, tracking and analysis of various health-related measurements from two different sources: wearable devices, including blood pressure, and chatbot conversations.

4.1 Architecture and Implementation

The architecture of the system implemented for the use case is represented by Figure 3, following the three-layered structure of the main framework.

4.1.1 Data Acquisition Layer. All the components responsible for gathering raw data from the human are implemented in this layer. In particular, in this use case, data are primarily derived from two main sources: sensor readings collected by a wearable device, and

[4]https://github.com/lm98/hdt-monitor

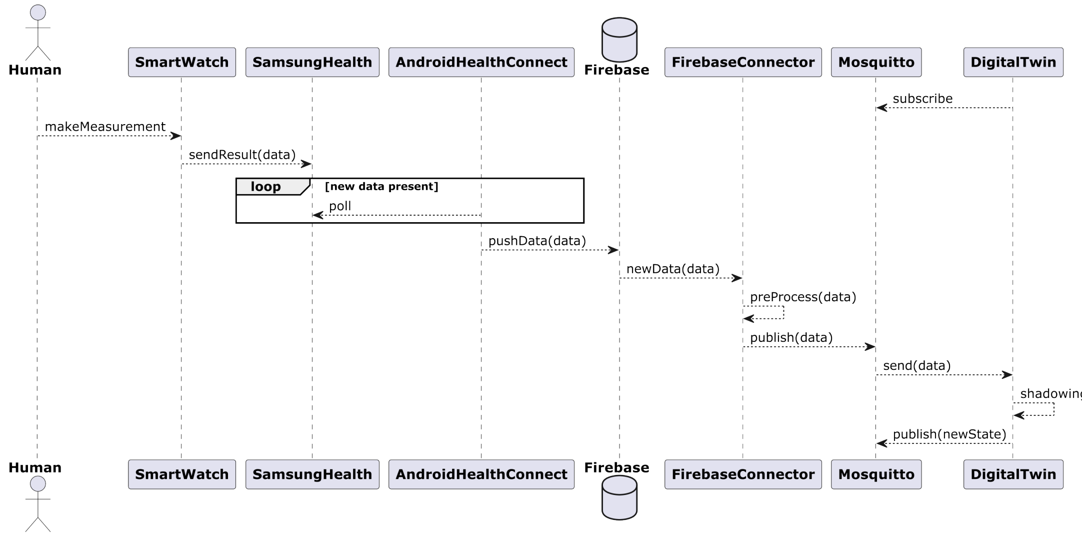

Figure 4: Sequence diagram showing the flow of calls and events originating from the wearable device and transmitted to the DT for model updates.

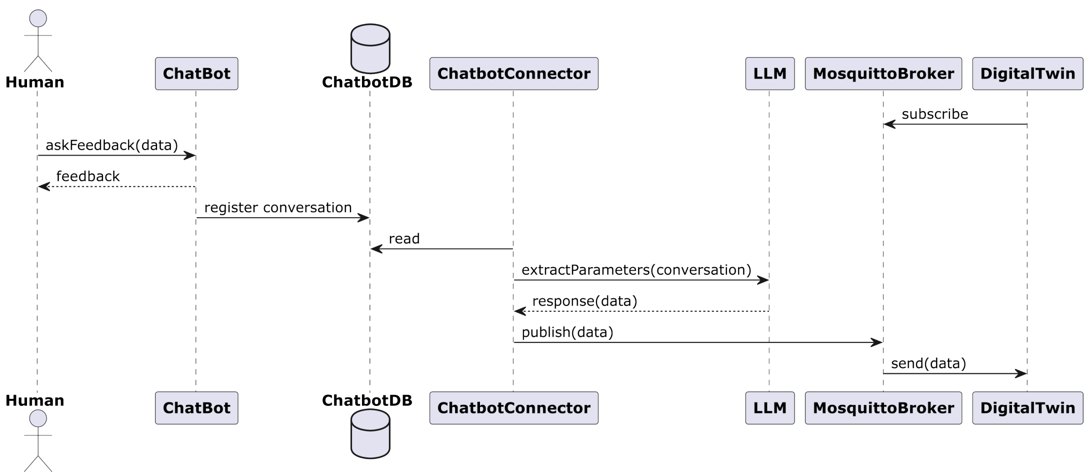

Figure 5: Sequence diagram showing the flow of calls and events originating from the chatbot and transmitted to the LLM for preprocessing and then to the DT for model updates.

unstructured textual data obtained from conversations conducted via a chatbot.

Wearable Data. A *Samsung Galaxy Watch 5* was used to gather blood pressure data. The measurements are then saved inside a Samsung smartphone via Samsung Health Connect [5]. In the smartphone is also installed an Android app *Android Health Connect*, which periodically exports measurements from Samsung Health and pushes them to a Firebase real-time database. The *FirebaseConnector* component is a kotlin process that:

(1) listens for updates on the Firebase real-time db;
(2) pre-processes the data to adhere to the DT's internal model;
(3) publishes the data over MQTT via Mosquitto broker;

[5]https://developer.samsung.com/health/blog/en/accessing-samsung-health-data-through-health-connect

The behaviour of this component can be represented by the sequence diagram in Figure 4.

Chatbot Data. In [10, 11], we presented *AI 4 HyperTension*, a Telegram chatbot supporting hypertensive patients in managing their condition. The chatbot has been designed to *(i)* support the acquisition of patient data, while ensuring privacy and security, and *(ii)* provide trustworthy answers to in-topic queries. *ChatbotConnector* collects conversations from chatbot and then uses them to perform a sentiment analysis to extract relevant information on patient mood, as shown in Figure 5.
The ChatbotConnector component is composed of several entities:

(1) The chatbot *AI 4 HyperTension* [10, 11] interacts with the end user.
(2) A Mongo DB is used to store conversations with users.
(3) A Python script that leverages Google Gemini [6] API to perform the sentiment analysis on the chat extract.

The Python script reads the extract of the chat with the chatbot and prompts the LLM via Gemini API to extract a structured evaluation of the user's general feeling and then publishes it via MQTT to the central broker, where the DT's PhysicalInterface listens for new data. The following is an example of the prompt used:

Data Extraction Prompt
Analyse Leo's energy levels (from 0 to 100) and general mood based on the following chat.
Respond only by providing a JSON object with three keys: `'energyLevel'` (Double value 0.0–100.0), `'mood'` (`'Calm'`, `'Nervous'`, `'Stressed'`, `'Anxious'`), `'timestamp'` (the date of the last message in the chat, in numeric format).

The desired format is:

```
{
    mood: value
    energyLevel: value
    timestamp: value
}
```

–Chat begins here–

Given the following conversation, for instance:

```
"Blood Pressure 120 77"
"Perfect, all clear! I will enter your data right away. Remember to send me a new measurement tomorrow. The average on blood pressure is 119/76 which is within the optimal health parameters. Very good!"
"I feel a little tired today"
"I'm sorry to hear you feel this way. Insufficient sleep, lack of exercise, and a poor diet can contribute to fatigue. If this condition persists or worsens, I recommend you see your doctor for further investigation"
"I slept 8 hours today but it is quite hot outside"
"Drinking plenty of water can help combat heat-related fatigue. Also be sure to rest in a cool, quiet place to help cool down your body."
```

the mood extracted is:


```
{
    'mood': 'Calm',
    'energyLevel': 30.0,
    'timestamp': 1749130500000
}
```


4.1.2 Digital Twin Layer. This layer contains an instance of *WLDT Digital Twin*. Here are three main components:

(1) MqttPhysicalAdapter: listens for sensor data over MQTT and generates physical events for the ShadowingFunction to update the DT's state;
(2) ShadowingFunctions: listens for physical events, updates the DT's state accordingly and generates an *UpdateStatusEvent*;
(3) MqttDigitalAdapter: listens for UpdateStatusEvents and propagates them to subscribers via MQTT.

Figure 6 shows the model of the acquired data. On the left, it presents a representation using the FHIR standard, while on the right, it displays the KG, illustrating how the various concepts and measures are interconnected. As shown in Table 2, each acquired parameter is modelled as a specific FHIR resource and has a corresponding LOINC code, as well as a corresponding property within the WLDT implementation.

4.1.3 Application Layer. The services that interact with the DT layer are included in the Application Layer. In particular, this use case presents a web application that subscribes to the MqttDigitalAdapter as shown in Figure 3 to asynchronously update a real-time chart for the patient's measures. A screenshot of the user interface is shown in Figure 7 illustrating on the left the values of a set of blood pressure measurements collected over a 24-hour period using the *Samsung Galaxy Watch 5*. These data are displayed within the web application. On the right, Figure 7 shows the energy level extracted from a conversation with the *AI 4 HyperTension* chatbot, conducted within the same time window.

5 Conclusion & Future Work

In this paper, we presented a three-layers framework for defining and implementing HDTs. The first layer is responsible for enabling the interaction with the real world. The second layer is the DT itself where a general model for each HDT has been introduced: it includes all the necessary abstractions to define a digital representation of a human. These abstractions, when appropriately instantiated, can be applied to any specific case study. The third layer provides services to inspect the HDT state and perform any kind of analysis on data acquired. A skeleton implementation on the WLDT platform has also been provided and is available to anyone who wishes to implement an HDT in their own case study.

A reference example is then provided to demonstrate the feasibility of the proposed framework for hypertension management. In this case study, the internal state of each HDT is autonomously updated in real time to mirror the human based on a set of vital signs, acquiring data from wearable sensors. Moreover, to demonstrate the ability of the proposed approach to integrate, in a unique reference framework and model, different types of data from different sources, we integrated mood data extracted from a chatbot conversation by querying an LLM. Acquired information are formalised with the FHIR standard and then visualised on the top-layer web application.

Although crucial in the future, in this paper we are not focusing on data analysis and are not discussing how to validate extracted data, which will be the subject of future work. In future work, we expect to extend the work in several directions: *(i)* additional vital signs and activity measures will be acquired from wearable devices and validated against measurements performed with certified medical devices; *(ii)* LLM sentiment analysis task will be subject of an ad-hoc evaluation; *(iii)* data will be acquired on a large population sample to enable different analyses related to the dynamic behaviour of blood pressure, its correlation with other parameters

[6] https://gemini.google.com/app?hl=it

```
{
  "resourceType": "Patient",
  "id": "leonardo-micelli",
  "name": [{
    "use": "official",
    "family": "Micelli",
    "given": ["Leonardo"]
  }]
}

{
  "resourceType": "Observation",
  "id": "blood-pressure",
  "status": "preliminary",
  "category": [{ "coding": [{ "system":
"http://terminology.hl7.org/CodeSystem/observation-category",
"code": "vital-signs" }] }],
  "code": { "coding": [{ "system": "http://loinc.org",
"code": "85354-9", "display": "Blood pressure panel with all
children optional" }] },
  "subject": { "reference": "Patient/leonardo-micelli" },
  "component": [
    {
      "code": { "coding": [{ "system": "http://loinc.org",
"code": "8480-6", "display": "Systolic blood pressure" }] },
      "valueQuantity": { "value": 120, "unit": "mmHg" }
    },
    {
      "code": { "coding": [{ "system": "http://loinc.org",
"code": "8462-4", "display": "Diastolic blood pressure" }] },
      "valueQuantity": { "value": 80, "unit": "mmHg" }
    }
  ]
}
```


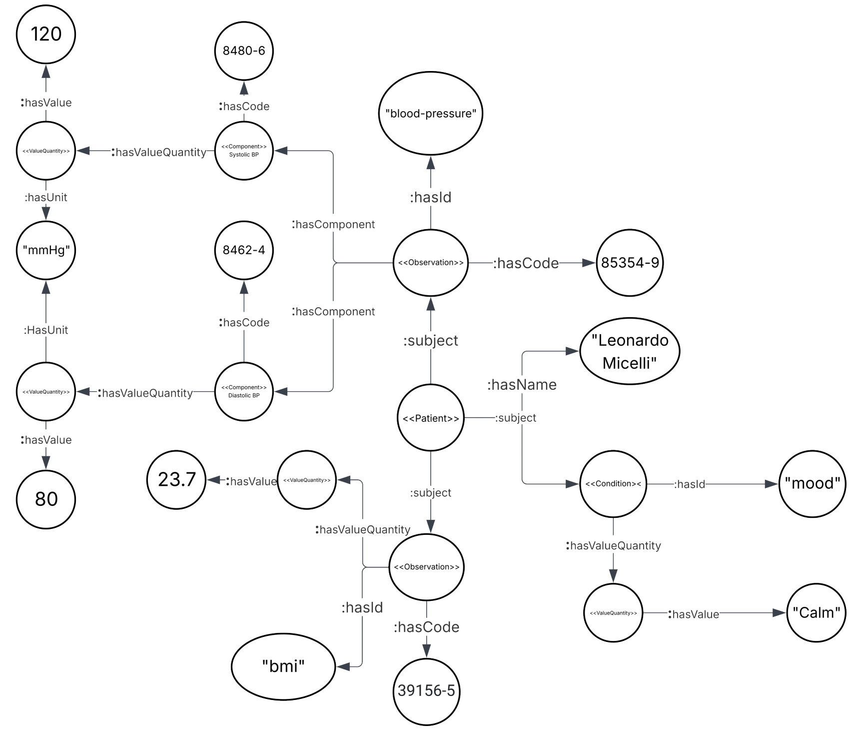

Figure 6: An extract of the Knowlege Graph.

Parameter	FHIR Resource	LOINC Code	Description	WLDT Property
Age	Derived	-	Calculated from birthDate	Age
BMI	Observation	39156-5	Calculated from weigth and height	Weight and Height Properties
Blood Pressure Panel	Observation	85354-9	Blood pressure panel (includes components)	BloodPressure
Systolic BP	Component	8480-6	Systolic blood pressure	–
Diastolic BP	Component	8462-4	Diastolic blood pressure	–
Patient Mood	Condition	80296-7	Captures a patient's mood	Mood

Table 2: FHIR concepts with corresponding resources and LOINC codes plus mapping from conceptual model attributes to WLDT properties.

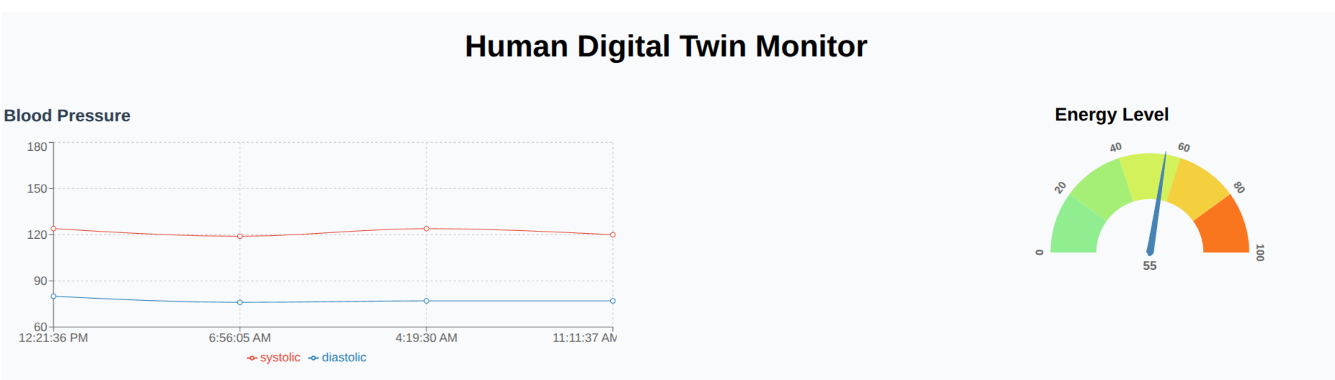

Figure 7: Blood Pressure Measures acquired from the wearable device *Samsung Galaxy Watch 5* and the energy level extracted from LLM-based sentiment analysis performed on chatbot conversation in the same time window.

for the identification of risk factors; *(iv)* additional assets will be included in the ecosystem of DTs to evaluate relations among HDTs and with other PAs; *(v)* privacy issues will be addressed to ensure compliance with the reference regulations.

Acknowledgments

This work has been funded by the Italian Ministry of University and Research (MUR) programme "PRIN 2022 Scorrimento" - grant number 20222X4M4T - I-DASY - CUP: J53C24002660006

References

[1] Bergthor Björnsson, Carl Borrebaeck, Nils Elander, Thomas Gasslander, Danuta R. Gawel, Mika Gustafsson, Rebecka Jörnsten, Eun Jung Lee, Xinxiu Li, Sandra Lilja, David Martínez-Enguita, Andreas Matussek, Per Sandström, Samuel Schäfer, Margaretha Stenmarker, X. F. Sun, Oleg Sysoev, Huan Zhang, Mikael Benson, and on behalf of the Swedish Digital Twin Consortium. 2019. Digital twins to personalize medicine. *Genome Medicine* 12, 1 (2019), 4.

[2] Antonella Carbonaro, Alberto Marfoglia, Filippo Nardini, and Sabato Mellone. 2023. CONNECTED: leveraging digital twins and personal knowledge graphs in healthcare digitalization. *Frontiers in Digital Health* Volume 5 - 2023 (2023). doi:10.3389/fdgth.2023.1322428

[3] Angelo Croatti, Matteo Gabellini, Sara Montagna, and Alessandro Ricci. 2020. On the integration of agents and digital twins in healthcare. *Journal of Medical Systems* 44, 9 (2020), 161.

[4] Michael Grieves and John Vickers. 2017. *Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems*. Springer International Publishing, Cham, 85–113.

[5] Evangelia Katsoulakis, Qi Wang, Huanmei Wu, Leili Shahriyari, Richard Fletcher, Jinwei Liu, Luke Achenie, Hongfang Liu, Pamela Jackson, Ying Xiao, Tanveer Syeda-Mahmood, Richard Tuli, and Jun Deng. 2024. Digital twins for health: a scoping review. *npj Digital Medicine* 7, 1 (2024), 77. doi:10.1038/s41746-024-01073-0

[6] Krešimir Kušić, René Schumann, and Edouard Ivanjko. 2023. A digital twin in transportation: Real-time synergy of traffic data streams and simulation for virtualizing motorway dynamics. *Advanced Engineering Informatics* 55 (2023), 101858. doi:10.1016/j.aei.2022.101858

[7] Caixia Li, Yina Zhao, Yang Bai, Baoquan Zhao, Yetunde Oluwafunmilayo Tola, Carmen WH Chan, Meifen Zhang, and Xia Fu. 2025. Unveiling the Potential of Large Language Models in Transforming Chronic Disease Management: Mixed Methods Systematic Review. *Journal of Medical Internet Research* 27 (16 Apr 2025), e70535. doi:10.2196/70535

[8] Yujia Lin, Liming Chen, Aftab Ali, Christopher Nugent, Ian Cleland, Rongyang Li, Jianguo Ding, and Huansheng Ning. 2024. Human digital twin: A survey. *Journal of Cloud Computing* 13, 1 (2024), 131.

[9] Ying Liu, Lin Zhang, Yuan Yang, Longfei Zhou, Lei Ren, Fei Wang, Rong Liu, Zhibo Pang, and M Jamal Deen. 2019. A novel cloud-based framework for the elderly healthcare services using digital twin. *IEEE access* 7 (2019), 49088–49101.

[10] Sara Montagna, Stefano Ferretti, Lorenz Cuno Klopfenstein, Antonio Florio, and Martino Francesco Pengo. 2023. Data Decentralisation of LLM-Based Chatbot Systems in Chronic Disease Self-Management. In *Proceedings of the 2023 ACM Conference on Information Technology for Social Good* (Lisbon, Portugal) *(GoodIT '23)*. Association for Computing Machinery, New York, NY, USA, 205–212. doi:10.1145/3582515.3609536

[11] Sara Montagna, Stefano Ferretti, Lorenz Cuno Klopfenstein, Michelangelo Ungolo, Martino Francesco Pengo, Gianluca Aguzzi, and Matteo Magnini. 2025. Privacy-preserving LLM-based chatbots for hypertensive patient self-management. *Smart Health* 36 (2025), 100552. doi:10.1016/j.smhl.2025.100552

[12] R Piascik, John Vickers, Dave Lowry, Steve Scotti, Jeff Stewart, and Anthony Calomino. 2010. Technology area 12: Materials, structures, mechanical systems, and manufacturing road map. *NASA Office of Chief Technologist* (2010), 15–88.

[13] Marco Picone, Marco Mamei, and Franco Zambonelli. 2021. WLDT: A general purpose library to build IoT digital twins. *SoftwareX* 13 (2021), 100661. doi:10.1016/j.softx.2021.100661

[14] Flávia Pires, Ana Cachada, José Barbosa, António Paulo Moreira, and Paulo Leitão. 2019. Digital twin in industry 4.0: Technologies, applications and challenges. In *2019 IEEE 17th international conference on industrial informatics (INDIN)*, Vol. 1. IEEE, 721–726.

[15] Shuang Qian, Devran Ugurlu, Elliot Fairweather, Laura Dal Toso, Yu Deng, Marina Strocchi, Ludovica Cicci, Richard E. Jones, Hassan Zaidi, Sanjay Prasad, Brian P. Halliday, Daniel Hammersley, Xingchi Liu, Gernot Plank, Edward Vigmond, Reza Razavi, Alistair Young, Pablo Lamata, Martin Bishop, and Steven Niederer. 2025. Developing cardiac digital twin populations powered by machine learning provides electrophysiological insights in conduction and repolarization. *Nature Cardiovascular Research* 4, 5 (2025), 624–636. doi:10.1038/s44161-025-00650-0

[16] Alessandro Ricci, Angelo Croatti, Stefano Mariani, Sara Montagna, and Marco Picone. 2022. Web of Digital Twins. *ACM Transactions on Internet Technology* 22, 4, Article 101 (November 2022), 30 pages. doi:10.1145/3507909

[17] Alessandro Ricci, Angelo Croatti, and Sara Montagna. 2021. Pervasive and connected digital twins—a vision for digital health. *IEEE Internet Computing* 26, 5 (2021), 26–32.

[18] Alessandro Ricci, Angelo Croatti, and Sara Montagna. 2022. Pervasive and Connected Digital Twins—A Vision for Digital Health. *IEEE Internet Computing* 26, 5 (2022), 26–32. doi:10.1109/MIC.2021.3052039

[19] Roberto Saracco. 2019. Digital Twins: Bridging Physical Space and Cyberspace. *Computer* 52, 12 (2019), 58–64. doi:10.1109/MC.2019.2942803

[20] Kaan Sel, Andrea Hawkins-Daarud, Anirban Chaudhuri, Deen Osman, Ahmad Bahai, David Paydarfar, Karen Willcox, Caroline Chung, and Roozbeh Jafari. 2025. Survey and perspective on verification, validation, and uncertainty quantification of digital twins for precision medicine. *npj Digital Medicine* 8, 1 (2025), 40. doi:10.1038/s41746-025-01447-y

[21] Paramesh Shamanna, Shashank Joshi, Mala Dharmalingam, Arun Vadavi, Ashok Keshavamurthy, Lisa Shah, Shambo Samrat Samajdar, and Jeffrey I Mechanick. 2024. Digital twin in managing hypertension among people with type 2 diabetes: 1-year randomized controlled trial. *JACC: Advances* 3, 9_Part_2 (2024), 101172.

[22] Abhimanyu Sharotry, Jesus A Jimenez, David Wierschem, Francis A Mendez Mediavilla, Rachel M Koldenhoven, Damian Valles, George Koutitas, and Semih Aslan. 2020. A digital twin framework for real-time analysis and feedback of repetitive work in the manual material handling industry. In *2020 Winter Simulation Conference (WSC)*. IEEE, 2637–2648.

[23] Venkatesh Upadrista, Sajid Nazir, and Huaglory Tianfield. 2025. Blockchain-enabled digital twin system for brain stroke prediction. *Brain Informatics* 12, 1 (2025), 1. doi:10.1186/s40708-024-00247-6

[24] Baicun Wang, Huiying Zhou, Xingyu Li, Geng Yang, Pai Zheng, Ci Song, Yixiu Yuan, Thorsten Wuest, Huayong Yang, and Lihui Wang. 2024. Human Digital Twin in the context of Industry 5.0. *Robotics and Computer-Integrated Manufacturing* 85 (2024), 102626. doi:10.1016/j.rcim.2023.102626