White Paper

The Health Digital Twin

A new paradigm for decision making
in healthcare

siemens-healthineers.com

SIEMENS
Healthineers

Challenges and trends driving the need for a Digital Twin in healthcare

Healthcare is undergoing significant change. Many new technologies are emerging that promise to help healthcare organizations meet growing demand and efficiently operate to deliver better patient care. These technologies promise to address several clinical needs and challenges, such as improving outcomes, lowering healthcare costs, addressing the lack of healthcare professionals and driving standardization of practice while offering personalization of care.

In parallel, the digitalization of healthcare is proceeding at an exponential pace. In 2020, it is expected that medical data will double every 73 days[1], and an estimated 80% of this data will be unstructured by 2025[2]. The volume, variety, velocity and veracity associated with this data explosion are creating several challenges. While medical professionals have access to more data than ever before, the ability to extract insight and value from this data is becoming increasingly difficult. The volume and complexity are simply too much for them to absorb, analyze, and use in ways that are meaningful for them and their patients. In addition, the growth and pace of medical research, clinical trials, and treatment options are staggering. It is estimated that it would take a doctor 150 hours each week to read every piece of content published in their field of interest[3].

The increasing availability of digital medical data is driving the need and opportunity for solutions to process and interpret this data. Advances in computational power, coupled with advances in technology areas such as artificial intelligence, computational modeling and simulation, are poised to help with complex decision making across all healthcare stakeholder groups. One exciting new technology trend that could revolutionize healthcare is the "digital twin". This technology is well established in manufacturing and industry, where it is used to model and simulate real-world assets like a machine or even entire factory before they are built[4]. These digital twins are used to predict failure, plan maintenance strategies, or even guide operations in a changing environment. The digital twin concept is now beginning to appear in healthcare where it could become a key component of the new wave of personalized and precision medicine. These digital twins could be used to allow patients to manage their health proactively, enable physicians to optimize therapies and predict outcomes, and much more.

For example, what if we could create a digital representation of a person's heart, that functions and beats like the real organ? What if the virtual heart could react to physiological stresses, diseases, medications and interventions as if they were real? This digital model could be used to predict disease, test therapies, plan interventions and monitor the cardiovascular health of the patient. This type of technology could bring significant value not only to the patient, provider and other stakeholders across the continuum of care, but also change the practice of care itself, potentially leading to better outcomes and lower costs.

Contents

What is a Health Digital Twin?

A digital twin is a digital replica of a physical asset[5]. In healthcare, as applied to a patient, a digital twin (hereafter referred to as a "health digital twin", HDT) can be defined as one or more computational models that provide a dynamic digital representation of a real-world biological target ("artifact") or some aspect of a person's physical state[6]. While a comprehensive digital model of an entire person is not yet a reality, a variety of HDT models have already been developed and tested today. This includes simulation of specific body parts or organs (such as the heart or the lung), models of disease states (such as lung cancer) or simplified "avatar" representations of a person's physical body. HDT models are typically constructed to provide decision support around defined clinical questions. HDT models can be developed for specific clinical use cases or applications or can be designed to be grouped together as a comprehensive library of patient-specific decision support tools. The models can be used to simulate a variety of scenarios *in silico*, virtually test therapies or devices, and predict trends and outcomes. A key attribute of HDT models is that they can be dynamically controlled by a user and can simulate scenarios to explore "what if" questions. The models can also be dynamically updated and can evolve as new clinical knowledge and new data become available. HDT models can be used at all clinical stages, from disease prevention to diagnosis, therapy, monitoring and home care. Because HDT models can be individualized with the patient's own data, they can offer personalized decision support and facilitate precision medicine approaches that are best suited to that patient.

Development of Health Digital Twin models

When developing HDT models for use in a clinical environment, a clear understanding of the intended use case and supporting workflow is required. This includes when, where, how and by whom the models will be used and the IT infrastructure that is required for deployment. The models would be designed to minimize workflow disruption, but their use may require some process re-engineering to ensure the proper fit and acceptance by end users. Additionally, the nature of the available IT infrastructure, including methods of connectivity, data accessibility, quality and structure and data privacy and security are important and sometimes limiting considerations.

Some of the key aspects related to the development of HDT model are depicted in Figure 1. Creation of HDT models often involves a reverse engineering of the human body or components of it. This requires scientific understanding, including knowledge of the structural and/or functional characteristics of the target artifact to be modeled. This understanding is used to design and implement a biophysical, biochemical or other logical representation of the target (depending on the type of model that is desired). This understanding is often limited, since physiological artifacts can be complex, highly non-linear, and contain many feedback loops and interactions. Furthermore, they can contain a variety of spatial and temporal scales, ranging from the molecular level to the level of the whole body, and from nanoseconds to years. Interactions between the target artifact and its surrounding environment must also be considered. This includes such things as where the artifact exists *in vivo*, any stresses that act on the artifact and wear or degradation of the artifact materials (e.g. due to age or disease). Ultimately, the complexity and functionality of an HDT model depends on its intended use. That said, the level of understanding that exists will in some cases affect the operation of the model, which can limit potential applications.

Development of Health Digital Twin models
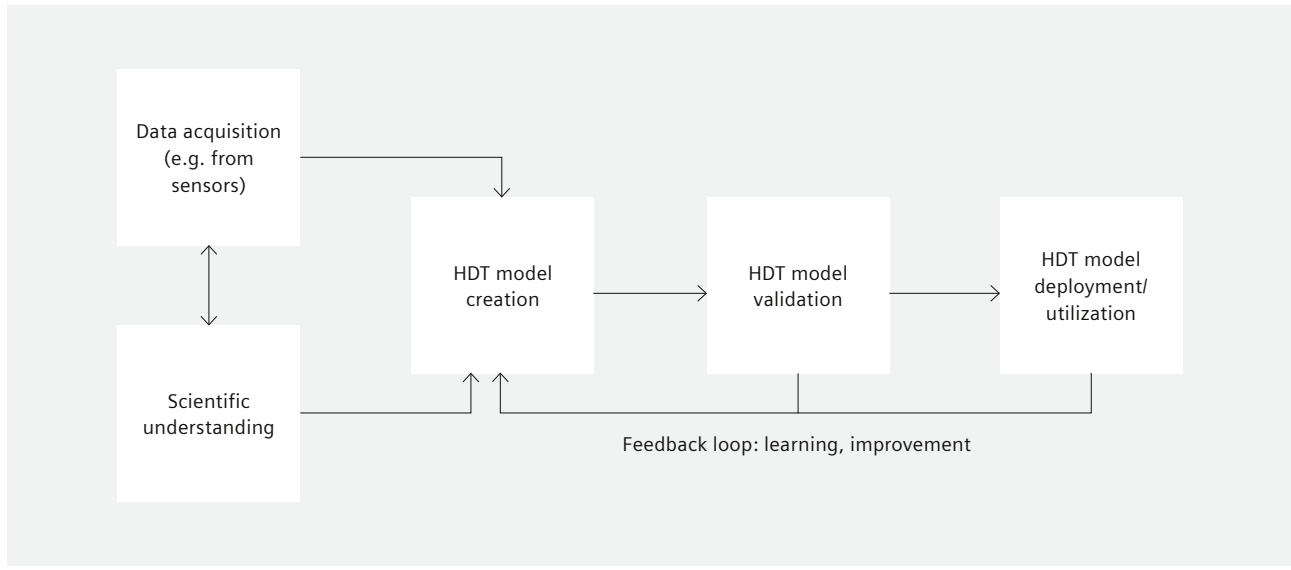
Figure 1:
Development of HDT models.

Obtaining suitable data is also a key requirement when building an HDT model. Data can come from a variety of sources such as sensor-based systems, including imaging scanners, lab tests, wearable or continuous monitoring devices, or from observations or interpretations. Data could be derived from empirical or experimental sources, such as from statistical analysis of historical or dynamic data that is gathered from the patient mathematical models, or clinical and scientific observations. Data could be related to a specific event (like tests that are performed before, during and after a clinical procedure) or collected continuously over the lifetime of a patient.

As data can often come in unstructured form, it must first be converted into a structured (machine interpretable) form before it can be processed and used in HDT models. Computational models built from the data can vary in scale and complexity from predictive algorithms that utilize a limited number of data points to full organ models with multiple layers of simulation (e.g. the VirtualHeart – see Figure 2.) The frequency at which the HDT model is updated is also dependent on the use case and on the availability and speed of sensor measurements and other necessary data. Model utilization can create a feedback loop that can dynamically update and improve the model over time.

Development of Health Digital Twin models
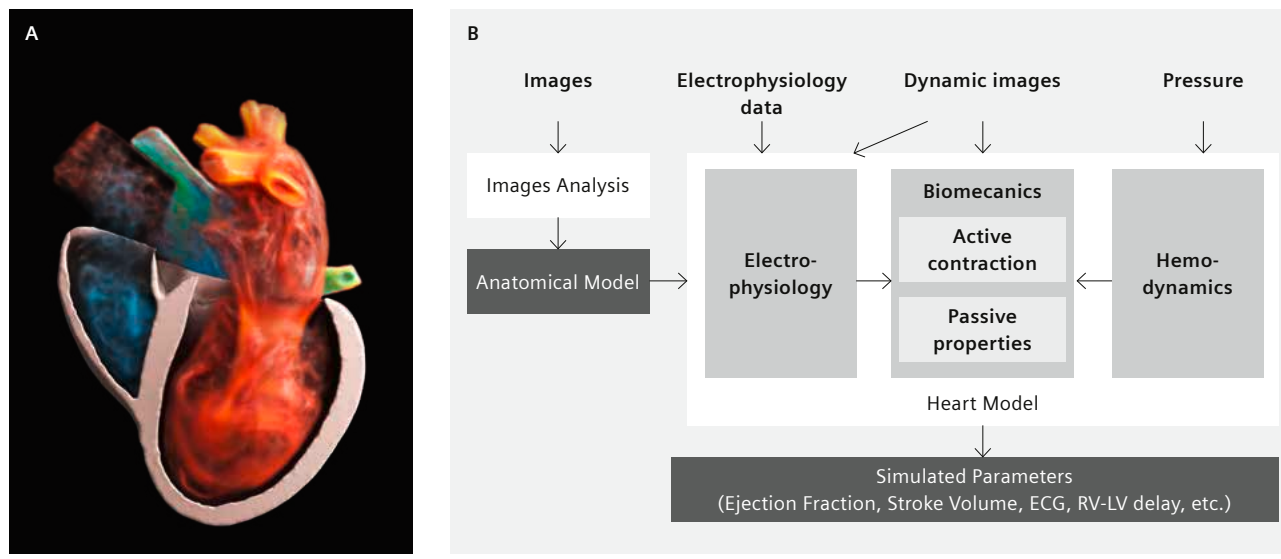
Figure 2: Siemens Healthineers' VirtualHeart.
The VirtualHeart is a prototype of a HDT model developed by Siemens Healthineers[7,8]. It is a multiscale, physiological model of a patient's heart. It has similar dimensions, electrical signal activation, muscle contraction, ejection fraction and pressure dynamics as the real heart. The model is built using both mechanistic and statistical approaches, with data collected from standard of care devices like echocardiography, electrocardiogram, cuff pressure. Since the model can be manipulated by a clinician, it can be used to virtually test assumptions, like a therapy for instance, and identify what treatment strategy is best for the patient. A. Graphical representation of a cross-sectional view of a VirtualHeart model. B. Elements comprising a typical VirtualHeart model. Upper boxes: input data; white boxes: processing units; dark grey boxes: output data. Arrows denote data flow.

Another important consideration is that an HDT model will also need to be individualized so that it can be used to provide an accurate simulation of function for a specific patient. This requires the parameters of the model be adjusted such that it replicates how the target artifact works within a specific patient. This includes potential defects, effects of previous events and physical properties. This individualization task may need to be performed continuously, to keep the digital twin up to date.

Once a HDT model has been created, it needs to be validated. This is especially true if the model is intended for use in a clinical environment. Often, data sets used to build prototype HDT models are limited to a small number of patients from a single institution, or other small cohort. Models must be validated with additional pools or retrospective or prospective data (ideally with data from a different source than used to construct the models) to ensure that they work as anticipated.

Validation is also an essential step to obtaining regulatory approval which is required for use in clinical practice. Key regulatory questions are related to how the HDT models will be used and the associated risk to the patient, and confirmation of safety and efficacy of the models. Regulators around the world have begun to outline guidelines for AI-based medical products, which include HDT models[9,10].

Rapid advancements in the field of artificial intelligence (AI) are also playing an important role in the development of HDT models. Fueled by an exponential increase of digital health data and computational power, AI approaches can be used to learn from available data, helping to identify and interpret patterns in ways that surpass the capabilities of the human brain. Because of this, AI could foster new ways of integrating available data and physiological models, help to uncover new mechanisms of biological function, add new levels of decision support and automate creation of more comprehensive HDT models.

Deployment of Health Digital Twin models
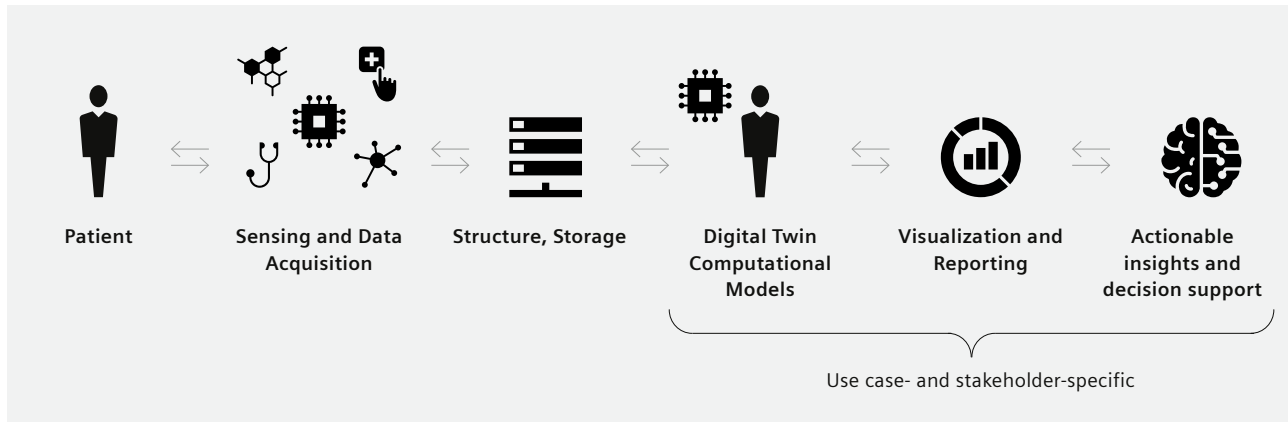
Figure 3:
Main components required for deploying HDT models.

Deployment of HDT models relies on a digital infrastructure that facilitates the creation, storage, transmission, analysis, visualization and reporting of the models and the data used to create them (see Figure 3). This includes the use of sensors or other means of capturing data related to the patient (e.g. from observations made by clinicians or patients themselves). Connectivity to devices or other storage repositories (such as electronic medical records, PACS, RIS, LIS, etc.), where healthcare related data required by the models resides, is also required. Access to this data relies upon standards and interoperability, which are still evolving. Furthermore, the infrastructure should also provide sufficient levels of security and privacy, since healthcare data can include sensitive information with restricted access (e.g. in order to comply with regulations such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States, and General Data Protection Regulation (GDPR) in the European Union).

Based on the clinical use case HDT models can be deployed in several ways. For example, the models could be embedded into a device (such as an imaging machine) to enhance the device's functionality. The models could also be provided as cloud-based services or deployed on a local network. Repositories of models could be constructed for a variety of applications at different stages of a clinical workflow. The use of the models could also extend outside of the clinical environment, for example for use in a patient's home or by other stakeholders in the healthcare continuum (e.g. payers, pharmaceutical, medical device companies, research labs).

The growing demand for clinical data has raised several critical questions about ownership and management of this data, which could also affect the deployment of HDT models. This includes the legal frameworks and regulations that address privacy and confidentiality of patient data, and the associated security that is required. The question of data ownership is similarly complex. Current laws can vary greatly by locality and are seen by many experts as being inadequate[11]. Some argue that rather than ownership, it is a matter of control, privacy and consent as well as value and intended use[12]. There is an increasing demand for more open and decentralized approaches with secured platforms and data access points which are managed in partner networks. Furthermore, new models for health data management are expected in the upcoming years, where health data ownership and access may be regulated by patients themselves.

The evolution of the HDT and how the technology is used will be driven by customer needs and supported by technological advancements. In the near term, the development of HDT models will target select, high-value use cases and will enable an additional level of 'intelligence' in existing healthcare products and services. In the future, the HDT concept could expand to include a modular, potentially interconnected, system of models with increased functionality that enables many more use cases and applications. This will fuel the emergence of novel products and services. Looking further out, the HDT may be represented by even more holistic models enabled by modular architecture and high interconnectivity fueled by large scale digitalization and availability of health data. This will facilitate expansion of the HDT out of the clinical arena into new markets and customers.

Value of the Digital Twin

	Healthy living	Prevention	Diagnosis	Treatment	Home care
Patient	Wellness and lifestyle coaching	Risk assessment, prognosis	Understanding diagnoses, related resources	Explanation of treatment options, risks/benefits, etc.	Monitoring recovery, ensuring compliance
Provider	Recommendations for wellness; Support GP in consulting patients.	Risk-stratification, early warning, prognosis	Early diagnosis, more detailed diagnosis	Therapy selection, planning, guidance, monitoring	Telemedicine, alerts, monitoring recovery
Payer	Promotion of healthy lifestyle; Population health management	Suggestions based on level of risk; more effective screening programs	Determine eligibility for re-imbursement	Comparative effectiveness of eligible treatments; performance metrics	Monitoring effectiveness of treatments, outcome reporting, value-based outcomes
Pharma/Med Device	Targeted marketing of nutraceuticals, supplements; Continuous monitoring devices	Preventative medicines or devices based on risk.	Companion diagnostics	Therapy response prediction, targeted therapy selection; Enabler of robotic treatments; *in silicio* clinical trials	Clinical trial monitoring, adverse side effect detection

Stakeholder Perspectives

Figure 4:
Examples of HDT value to stakeholders across the healthcare continuum.

HDT models aim to provide actionable insights and facilitate decision support across the healthcare continuum. HDT technology will enable mass personalization of disease prevention, diagnosis, treatment and monitoring, through the delivery of high-value, precision care.

As depicted in Figure 4, HDT models can play a role at all stages of the clinical value chain, as well as be used externally to the clinical environment (e.g. to promote healthy living, disease prevention and home care).

For patients, HDT models could provide a better understanding of their health or disease, enabling them to make better health-related decisions. This could help to facilitate more informed discussions between patients and their doctors, simplifying the communication and interpretation of diagnostic or therapeutic results and options that are often complex. For example, for a patient with coronary artery disease, a HDT model of their heart would enable them to review and understand with their clinician where the blockages are, what therapies are available, and the risks and outcomes that could be associated with each approach. HDT models driven by data from wearable or implantable sensors, or other continuous monitoring devices, could also serve to coach patients on lifestyle improvement. This could include providing personalized diet and exercise regimens for healthy individuals and also for rehabilitation. In short, HDT technology would empower patients to be more involved and take greater control of their health or manage disease.

Value of the Digital Twin

For clinicians, HDT models could help to assimilate all of the clinical data available to them and identify patterns or other associations which could provide them with a deeper understanding of a patient's health. This could enable them to diagnose disease earlier or offer preventative guidance to patients, understand disease prognosis and offer appropriate treatment and monitoring options. Clinicians could also use HDT models to stratify patients based on risk, plan a therapeutic intervention or simulate the effect of medications. For example, for a patient with liver cancer, an HDT model of their liver that contains the location of a tumor could be used by a surgeon to plan a tumor ablation surgery. This could provide the clinician with an understanding of critical parameters, such as where the ablation probe should be placed and how much energy should be applied, in order to optimize the destruction of the tumor while retaining as much surrounding healthy tissue as possible. The HDT models could provide insights and decision support to supplement their clinical experience and allow them to better anticipate the results of their actions.

For a hospital, the incorporation of HDT models into clinical and operational workflows could have many potential impacts. Use of the models by clinicians and other staff across the enterprise could streamline and optimize processes, enabling faster and more standardized decision making from complex data and interpretation of clinical guidelines.

At the same time, the models could also provide more personalized and precision diagnostics, therapeutics and patient monitoring since the models would be tailored to the individual patient. HDT technology could also help to alleviate potential bottlenecks in patient handling or shortages in staffing. This could include assisting with patient triage (e.g. admissions in an emergency room) or automated monitoring and alerting for patients recovering in an intensive care unit. Overall, the use of HDT technology could optimize outcomes and lower costs at an institutional level by enabling doctors to make better, more informed decisions, both at the patient level as well as at the population level.

Use of HDT technology could also benefit other stakeholders, such as payers and pharmaceutical and medical device companies. By helping to better stratify patients, HDT models would ensure that patients are receiving care that is safer and more effective. For payers, this would help to lower costs through the reduction or elimination of unnecessary procedures as well as lower the risk of adverse results. For pharma and medical device companies, this could help to better select and recruit patients for clinical trials, enabling the development of more targeted and precise therapies. These impacts would feed back to patients and providers, improving the delivery of care and management of health and disease.

A new paradigm for decision making in healthcare

The ability to simulate human physiology *in silico* using HDT models represents a new paradigm for utilizing available digitized health data from a patient. These models could be used to derive personalized, actionable insights, which could in turn lead to optimized clinical processes and improved outcomes. For companies like Siemens Healthineers who operate in the healthcare space, this technology could be infused into future products, serving to enhance their functionality. It could also lead to the creation of new services and related business models, that will enable mass personalization of healthcare delivery.

In the nearer term, HDT models could provide additional levels of functionality to products such as imaging scanners and lab test equipment. Utilizing HDT models infused with artificial intelligence, these devices would require less human intervention, leading to increased automation and optimization of processes. These models could also enable testing of different scenarios as well as predicting results outcomes before clinical decisions are made, improving treatment efficacy and patient safety. HDT models could also leverage available integrated diagnostics data gathered from multiple modalities both *in vitro* and *in vivo*, to build more comprehensive decision support solutions.

Moving towards the future, advances in scientific understanding of human physiology, combined with advances in key technology areas such as computing, AI, sensing and digital infrastructures will lead to the growth and evolution of HDT models. The models will become more comprehensive and complex and will facilitate a wider range and depth of decision support. This will be fueled by an ever-increasing and widespread availability of digital health data, which will enable the validation and refinement of the HDT models.

While comprehensive HDT models that provide a detailed representation of the entire human body may still be several years away, the underlying concept is still poised to make a significant impact in the nearer term. It is predicted that as more HDT models are generated and validated, challenges to the widespread implementation of the technology will decrease and adoption will increase. While this journey may take significant development efforts, it will be disruptive to the healthcare system at large, fundamentally changing the way healthcare decisions are made.

References

[1]   Marr, B. The 9 Biggest Technology Trends that will Transform Medicine and Healthcare in 2020. Forbes, November 1, 2019.

[2]   Reinsel, D., Gantz, J., Rydning, J. Data Age 2025: Don't Focus on Big Data; Focus on the Data That's Big. IDC, April 2017.

[3]   Ahmed, N. M., Toor, A. S., O'Neil, K. et al. Cognitive Computing and the Future of Health Care. IEEE Pulse, May/June 2017.

[4]   Werner Kritzinger, W., Karner, M., Traa, G. et al. Digital Twin in manufacturing: A categorical literature review and classification. IFAC-PapersOnLine, 51(11): 1016-1022, 2018.

[5]   El Saddik, A. Digital Twins: The Convergence of Multimedia Technologies. IEEE MultiMedia, 25(2): 87-92, April-June, 2018.

[6]   Bruynseels, K., Santoni de Sio, F., van den Hoven, J. Digital Twins in Health Care: Ethical Implications of an Emerging Engineering Paradigm. Fronteirs in Genetics, 9 (31): 1-11, 2018.

[7]   Kayvanpour, E., Mansi, T., Sedaghat-Hamedani, F., et al. Towards Personalized Cardiology: Multi-Scale Modeling of the Failing Heart. PLoS ONE, 10(7): e0134869, 2015.

[8]   Mansi, T., Tiziano, P, Comaniciu, D. Artificial Intelligence for Computational Modeling of the Heart. Academic Press, November 25, 2019.

[9]   Merritt, R. Healthcare Regulators Open the Tap for AI. Nvidia Blog, December 2, 2019

[10]   Murgia, M. and Shrikanth, S. How governments are beginning to regulate AI. Financial Times, May 29, 2019.

[11]   Andanda, P. Towards a Paradigm Shift in Governing Data Access and Related Intellectual Property Rights in Big Data and Health-Related Research. IIC, 50: 1052–1081, 2019.

[12]   Lovett, L. Patient data: Who owns it? Who can improve it? What's it worth? Mobihealthnews, February 18, 2019.

On account of certain regional limitations of sales rights and service availability, we cannot guarantee that all products included in this brochure are available through the Siemens sales organization worldwide.

Availability and packaging may vary by country and is subject to change without prior notice. Some/All of the features and products described herein may not be available in the United States.

The information in this document contains general technical descriptions of specifications and options as well as standard and optional features which do not always have to be present in individual cases.

Siemens reserves the right to modify the design, packaging, specifications, and options described herein without prior notice. Please contact your local Siemens sales representative for the most current information.

Note: Any technical data contained in this document may vary within defined tolerances. Original images always lose a certain amount of detail when reproduced.