Secure Medical Digital Twins: A Use-Case Driven Approach

Salmah Ahmad[1,3][0000-0002-3084-535X], Bianca Bartelt[1,3][0009-0009-5531-3737], Matthias Enzmann[1,2,4][0000-0002-8871-8767], Jörn Kohlhammer[1,2,3][0000-0003-1706-8979], Stefan Wesarg[1,2,3][0000-0002-2073-6010], and Ruben Wolf[1,4][0009-0001-0493-2927]

[1] ATHENE — National Research Center for Applied Cybersecurity, Darmstadt, Germany

[2] Fraunhofer Cluster of Excellence Immune-Mediated Diseases CIMD, Frankfurt am Main, Germany

[3] Fraunhofer Institute for Computer Graphics Research IGD, Fraunhoferstr. 5, 64283 Darmstadt, Germany

{salmah.ahmad,bianca.bartelt, joern.kohlhammer,stefan.wesarg}@igd.fraunhofer.de

[4] Fraunhofer Institute for Secure Information Technology SIT, Rheinstr. 75, 64295 Darmstadt, Germany

{matthias.enzmann,ruben.wolf}@sit.fraunhofer.de

Abstract. Medical digital twins (MDTs) can support health professionals in many ways, such as in achieving a more precise diagnosis or in developing better treatment plans. For that, MDTs will need to process patients' data from multiple sources and will have to govern access to that data by health professionals from multiple institutions, and at the same time support collaboration and data sharing between them. This raises questions regarding trust and the security of processing. We illustrate two medical use cases where MDTs can have a significant impact and discuss the security scope of our MDTs.

Keywords: Digital patient twins · Data protection · Access control · Cyber security.

1 Introduction

Medical digital twins (MDTs) denote a comprehensive model of a patient's history and health status. Ideally, it is based on anatomical, physiological, diagnosis- and treatment-related information from multiple sources. While often incomplete, it may also provide derived information about the patient based on known relations between the medical data. A medical digital twin also comprises information that is among the most critical and sensitive data according to the EU GDPR (General Data Protection Regulation), motivating a dedicated approach to secure such MDTs, which is further described in this paper. This approach goes beyond the current state of technology, in which mostly the transmission of medical data and the perimeter around medical data repositories are secured, while the medical data itself remains readable and exploitable.

It is essential to be able to transparently share the data contained in MDTs in a fine-grained way to give medical experts in different institutions access to certain parts of the data. The use of MDTs can be distinguished into four different groups. First, patients should be able to grant access to all data related to a specific disease for which they are treated. It should, second, also be possible to grant this access related to a treatment process across several treatment facilities. In research use cases, it is also necessary to receive access to data to build patient cohorts based on, third, certain patient features or, fourth, on organ-specific features. All use cases should be supported by easy and usable ways to grant access to specific data attributes and to use the MDTs on the medical side.

A digital twin that processes medical data of a patient requires extra care in its operation due to the sensitive nature of the data being processed. On a technical and organizational level, "extra care" translates to the implementation of measures to ensure, inter alia, authorized access, data integrity, authenticity, and confidentiality. Those are classic IT security goals which apply to the operation of a medical digital twin, too.

For every security goal, many technical solutions exist to protect the data, including access control frameworks, schemes for message authentication, digital signatures, and encryption. The challenge is to make them work together and enforce the security goals within an organization, and even more challenging beyond it, e. g., when patient data is to be shared back and forth between a physician's office and a hospital. MDTs are expected to flourish the most when medical data from different sources, like different institutions, are available to enhance the analysis of a patient's condition, to achieve a more precise diagnosis, and to allow health professionals to develop better treatment plans. We share that vision and therefore consider the flow of medical data as an important building block for an MDT. However, uncontrolled data flows may erode trust in that concept. Therefore, any solution should involve patients so that they have a say in where their medical data goes and technically enforce those decisions, e. g., by encrypting data for specific recipients. However, there can be a fine line between giving patients control over their data and overburdening them with micro-management of who gets to process which of their data. In addition, a solution should also take into account that patients may want to delegate that control to a trusted entity, e. g., a family doctor, a relative, caregiver, or a medical center. Finally, the sharing of MDTs for medical research should not be overlooked in the implementation of access control mechanisms and the access control process should support both denying and granting access in a usable way.

The contributions of our paper are the following: We provide an overview of the current research landscape on MDTs and usable interfaces for patients and medical experts. We contribute a new definition for secure MDTs that goes beyond the current perimeter- and transfer-oriented security concepts. We show new possibilities for the secure use of MDTs that are not easily feasible with traditional approaches and provide two example use cases in the areas of Parkinson's disease and inflammatory bowel disease.

2 Related Work and Background

The essential concept of a Digital Twin (DT) was first described by Grieves for manufacturing organizations [12]. In his concept, he described a physical object (PO), like a manufacturing machine, and a corresponding instance of a digital model/process (i. e., the DT) that mirrors the state of the PO using a bidirectional connection between them. The connection between the two entities should allow data to flow from the PO to the DT and vice versa. Data from the PO updates the DT's state model of the PO while data from the DT can adjust the PO's settings. However, the DT can be more than a remote control as it may process data from additional sources, e. g., data from environmental sensors, to find and propagate settings that may further improve the PO's performance. In [26] and [19], an overview of recent background and details on DTs in general can be found. Of course, DTs are not limited to manufacturing but have also been discussed for other fields like agricultural production [27], medicine [3, 8, 25], or smart cities [4].

A medical digital twin (MDT) can be described by observations of and measurements taken from a patient [9]. By processing a patient's data, an analytical model can be built that provides a holistic view of the patient [13] which may support health professionals in achieving timely diagnoses, possibly simulate different diagnostic scenarios, and to develop personalized treatment plans [11].

If we view an MDT as a holistic medical data model of a patient, then only a part of it will typically be required by a health professional. That part, say medical data on the patient's heart condition, can be regarded as state data which can be processed on its own, e. g., by a heart MDT. An MDT "factored out" like that can be seen as a separate DT instance (DTI) that may have connections to other DTIs, e. g., a lung MDT. Following that, a DT aggregate – comprised of DTIs – comes to mind that "allow[s] for a larger and more complete dataset regarding the operation of a type of physical object" [16]. Similarly, Haße et al. speak of a sub model of a DT when it comes to sharing parts of a DT [14]. Shared digital twins are of particular interest in the medical field because they can facilitate more effective collaboration among health professionals [25].

Information security is paramount for MDTs as a compromised MDT may have direct or indirect consequences for patients – and that includes more than their health. In its foresight report for 2030 [10], ENISA expects the "exploitation of e-health (and genetic) data" to be among the top threats by 2030 and says that such sensitive data "may be exploited or used by criminals to target individuals or by governments to control populations, e. g., using diseases [. . .] as a reason for discriminating against individuals". In addition, a compromised MDT may draw wrong conclusions about a patient's current state of health. In [18], general attack principles (e. g., eavesdropping) and defence mechanisms (e. g., encryption) for industrial cyber-physical systems are discussed that also apply to MDTs. Many defence mechanisms work well within an organization, like role-based access control (RBAC) [21]. However, when data leaves the control sphere of the organization, like with a shared MDT, the same or similar level of control over the data is no longer guaranteed. Data encryption is always an option for protecting shared data, but it is not a panacea, as implementing cryptography-based access control across organizations can be complex [29].

3 Medical Digital Twins

In the following, we use the term medical digital twin (MDT) in two ways. Firstly, we refer to an MDT as a collection of the medical data of some patient and secondly, to relate to the system that holds, processes, and provides access to that data (see Fig. 1).

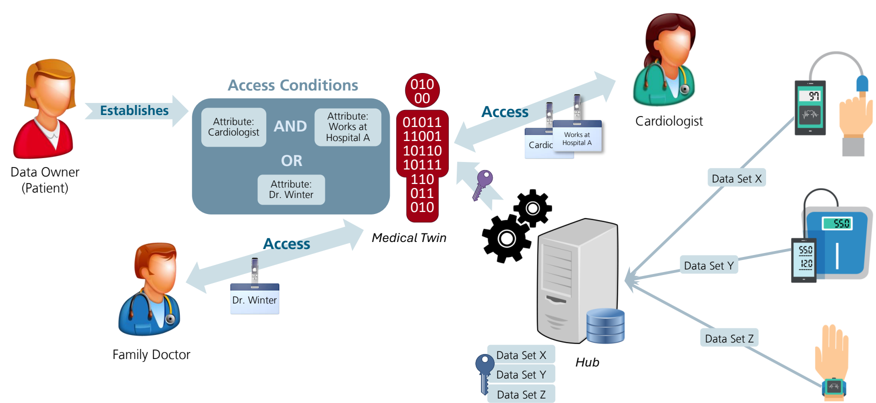

Fig. 1. Usage of a secure MDT: Data owner sets access conditions. Medical practitioners who fulfil these conditions can access the encrypted data. [17]

3.1 Shadow of a Medical Digital Twin

A digital twin can be considered as a numerical model of a physical process or an object. In case of digital patient twins, it quickly becomes obvious that a comprehensive modelling of all physiological processes going on in such an organism seems infeasible. Therefore, a medical digital twin in our concept always focuses on a subset of relevant aspects. This can be organ-driven (e. g., a digital twin of the heart), focused on a specific disease and the related therapy process, or more generally on a patient journey covering patients' pathways through treatment and the hospital.

An MDT – no matter what it specifically represents – acts as a data repository and therefore requires well-defined data structures. This allows for an aggregation of medical and health-related data. An MDT can contain information like attributes from an electronic health record, additional attributes that are relevant for medical research, sensor parameters as well as patient diary entries. For storing and accessing those data entries, data needs to be managed accordingly.

When considering specific use-cases of an MDT, a shadow [2] of the comprehensive MDT needs to be created. That means a specific subset of data structures is used for representing the case-related data, like relevant lab parameters, image-derived biomarkers, medication information, and many more. In addition, that data needs to be made accessible to the involved physicians, patients, medical nurses as well as other relevant stakeholders. Another important aspect is the integration of analysis capabilities that work on the data contained in the medical patient twin. As an example, a shadow focusing on a cardiology-related use case might integrate algorithms for detecting abnormalities in ECG data.

3.2 Collaboration, Data Sharing, and Security

A crucial aspect to be considered is that an MDT contains personal data and thus, in Europe, falls under the General Data Protection Regulation (GDPR) which governs the prerequisites, safeguards, and rules for any processing of such data. Of course, additional regulations and national laws regulating the health sector may exist but, in this paper, we focus on the security aspects of data protection. To be practically useful, an MDT needs to consider both the medical application and information security. For that, secure methods for populating the medical patient twin as well as accessing the stored data need to be provided. Our approach aims for a solution where even if an unauthorized person gains access to the MDT infrastructure the data stored therein cannot be read.

From a pure network-security point of view, MDTs may just look like another distributed system for which plenty technical and organizational measures come to mind that may be employed to achieve overall security, see for instance [18]. While an MDT is an application that heavily relies on networks, and thus surely benefits from network defence mechanisms, it is also about sharing and processing medical data. This requires more than encrypted data transmissions, where data is typically encrypted during transport only. Once the data is received, it is unclear how the data will be protected afterwards, e. g., how confidentiality and access control are implemented at the receiver's end, if at all. In our view, an MDT is about collaboration and hence, an MDT may be accessed by different user groups, e. g., the patient, health professionals from different departments and different organizations, or medical researchers. Naturally, each group will literally and technically have its own view on the MDT. In the following, we will refer to a selected portion of the detailed medical data contained in an MDT as a shadow (of the MDT).

A shadow may be plain data, data derived from an MDT, or both. From an operational and from a security perspective, a particular shadow only needs to be accessible to a certain group. For instance, if a patient is referred to a specialist, the specialist will only need access to a specific shadow of the MDT, i. e., to a limited data set. Likewise, if an MDT provides a research shadow, e. g., selected pseudonymized or anonymised patient data, medical researchers from different organizations may be given access to that data. Still, access to an MDT must only be given to qualified people, i. e., individuals that received an authorization for accessing the MDT. For instance, a specialist with access to some patient's MDT may want to discuss the condition of the patient with a colleague and for that may want to share a shadow of the MDT with their colleague, who may or may not be working at the same practice. Likewise, a patient may see a new physician and wants to share their medical history with them. For all that, the specialist/patient needs a way to authorize access to a shadow of the MDT for another health professional. In addition, if a shadow needs to be processed by another organization, it must be securely transferred to and securely stored by that organization. Employing encrypted transmissions and encrypted storage, like full disk encryption (FDE), are a bare minimum to achieve that.

FDE does not protect individual shadows but all of them at once. At first sight, this seems to solve the problem of secure storage, but it does not. The problem is that "protecting all at once" in that case also implies "compromising all at once", which may happen when an online attacker gains access to the system that in turn has access to the encrypted storage. To that system, and hence to the attacker, the encrypted storage is "logically unencrypted", since the system automatically decrypts the stored data before it passes the data to the medical application that wants to process it. To mitigate such attacks, we aim for individually encrypted shadows of an MDT such that each (transferred) shadow is protected by its own encryption keys. Corresponding decryption keys are attached to the MDT and will only be made available after a user has demonstrated their authorization and will only be retained as long as the authorized user works with the shadow. That approach is related to the Zero Trust paradigm [5] where trust needs to be earned on every access to a resource, rather than being presumed based on weak assumptions, like "if data is accessed from within our organization's network then access is allowed". In our model, trust is earned per data set, i. e., the user must satisfy all access conditions associated with the requested data on *every* access to get the decryption key of the data.

3.3 Approach

Our approach makes use of attribute-based encryption (ABE) [23]. In ABE, more specifically in ciphertext-policy ABE, the encryptor defines an access policy in terms of attributes that a prospective decryptor must have in order to satisfy the policy, which in turn allows him to decrypt a ciphertext, i. e., the encrypted data. Each attribute is associated with a public encryption key and a secret decryption key per user. Both are issued by a so-called attribute authority that also validates if someone requesting a certain attribute is allowed to hold it. In our use cases, we expect different stakeholders to share data, and thus user groups will not necessarily be from the same organization, e. g., Hospitals A and B. For more flexibility, we use multi-authority ABE [6, 20] where several attribute authorities can be established that only share an initial set of global (security) parameters. Every attribute authority may define and issue its own attributes, and the global parameters make sure that those attributes are compatible, such that they can be combined in an access policy. Larger organizations, like hospitals, may establish their own attribute authority while smaller ones, like a general practice, may employ established trust service providers.

The ABE approach is related to conventional role-based access control (RBAC) where a system-wide infrastructure controls and enforces access permissions to resources. In RBAC, resources can have security labels that determine which roles may access a certain resource (e. g., a health record), and every user in the system is assigned to roles (e. g., cardiology, administration), and roles are assigned access permissions for resources. In that sense, roles in RBAC can be compared to attributes in ABE. The major difference between RBAC and ABE is that ABE is independent of an enforcement infrastructure that controls access to data. This is advantageous if data is expected to be shared between infrastructures of different organizations. With RBAC, those infrastructures will likely have different controls, e. g., access rules and security labels. Hence, the original access conditions (roles and permissions) from the source system can no longer be relied on because they would have to be translated to the target system's, which may or may not be comparable. In contrast, access in ABE is exclusively controlled by the access policy which is an integral part of the ciphertext. And only that access policy determines the decryption key of the ciphertext and remains exactly the same, irrespective of the target system.

3.4 Usage

In Fig. 1, we have three attributes "Cardiologist", "Works at Hospital A", and "Dr. Winter" that are combined in an access policy. Although each attribute is personalized to its holder, the latter attribute is a kind of "anomaly" in ABE because it may only be held by a single person, which is why we call it a singleton. The other attributes represent groups and hence may be held by a number of persons. Personalisation of attributes is an important aspect in ABE because it prevents collusion. That is, say, two persons – one with attribute A and another with attribute B – cannot combine their attributes in order to satisfy an access policy requiring "A and B", which none of them could have satisfied alone.

We expect every encryptor to only use the most recent attributes for access policies, i. e., encryptors follow the authorities' public key updates, which are accessible through a public directory. If a user is no longer eligible for a certain attribute, he is not permitted to renew his secret attribute key with the issuing authority. Renewal of a secret key becomes necessary whenever an attribute's corresponding public key is updated. Without renewal, a user will not be able to decrypt ciphertexts that were encrypted with recently updated public attribute keys. The time frame for public-key updates is up to the authority/organization and may depend on the organization's risk tolerance. For instance, a risk-averse organization may update an attribute key whenever the group of users is reduced, e. g., when someone leaves the organization. Others with more risk appetite (or possibly less sensitive data) may wait until a certain threshold of changes is reached, or some fixed time frame has passed, before they update an attribute key. Note that ciphertexts encrypted under some access policy will not change if the public key of an attribute from the policy is updated, i. e., someone with "old" keys can still decrypt them. For that, our current approach is to have lazy updates of key encryption keys (KEKs = ABE keys). That is, we use hybrid encryption, where data encryption keys (DEKs) are used to encrypt the "actual data" and KEKs to encrypt the DEKs. Now, "lazy" means that ciphertexts of DEKs are updated to the current ABE keys with the next write operation, i. e., DEKs are re-encrypted using current KEKs. Another approach would be to update a DEK ciphertext whenever an attribute of its access policy changes. So again, this is a matter of risk tolerance, but also of work load because existing DEK ciphertexts would have to be constantly re-encrypted in the background.

Still, if an attacker manages to break into the system of an authorized user that is currently working with a shadow, the attacker may learn what is contained in that shadow. However, the attacker will not gain access to all other shadows because their decryption keys are unavailable without separate authorizations. The challenge of the cryptographic approach is that keys on several access levels need to be generated and efficiently managed to provide correct handling for the different stakeholders.

A final requirement is the necessity to integrate population norms into the MDT. This allows to compare individual patient parameters to those that can be expected in a healthy person or to those specific for the staging of a disease. The other direction is the derivation of cohorts when aggregating the information stored in several shadows. This requires a partial access to patient-individual parameters that needs to be granted to researchers for verifying hypotheses or preparing patient studies.

4 Use Cases

Medical digital twins (MDTs) have vast potential in the field of healthcare, offering innovative solutions for the treatment of various medical conditions. In the following sections, we explore two use cases where MDTs can have a significant impact: neurodegenerative diseases and inflammatory bowel disease.

4.1 Neurodegenerative Diseases

Increasing life expectancy is a key driver of the growing incidence of neurodegenerative diseases (NDDs) such as Alzheimer's, Parkinson's and ALS [24]. These diseases lead to a progressive loss of cognitive and motor functions, which severely affects quality of life. As NDDs cannot yet be cured, physicians focus on relieving the symptoms, which is usually very complex, time-consuming and cost-intensive [28]. The integration of secure MDTs, as defined above, offers a promising approach to address these challenges.

In the context of NDDs, our MDT serves as both a structured, disease-specific repository of a patient's health data and the system that manages, processes, and provides secure access to this data. The MDT captures and organizes heterogeneous, multimodal health data most relevant to NDDs, such as longitudinal cognitive assessments, motor function tests, medication records, patient diary entries and neuroimaging data (e. g., MRI, PET, SPECT scans). For example, the MDT can be continuously updated with imaging results that reveal disease-specific changes, such as dopaminergic neuron loss in Parkinson's or amyloid plaques in Alzheimer's. This enables physicians to monitor disease progression, detect early biomarkers and personalize treatment strategies.

Access to the MDT is strictly controlled, as outlined in our approach. Only authorized stakeholders, such as neurologists, caregivers or family members, can access specific shadows of the MDT. For example, a neurologist may access longitudinal MRI scans and cognitive test results to monitor disease progression, while a caregiver may only view medication schedules and daily activity logs. As NDDs progress and the patient may no longer be able to make health decisions for themselves, the MDT allows for secure delegation of access to trusted family members or legal guardians. This ensures continuity of care while maintaining strict data protection. Building on these capabilities, our proposed MDT securely brings together key NDD data, enabling personalized care, clear updates for patients and caregivers and ultimately contributes to an overall improvement in the treatment of NDDs.

4.2 Inflammatory Bowel Disease

Inflammatory bowel disease (IBD) is a chronic condition that is costly to both patients and the healthcare system. It is characterized by relapsing-remitting symptoms, which cause inflammation of the gastro-intestinal tract [1, 22]. Well-known varieties of IBD are Crohn's disease or ulcerative colitis. Symptoms include diarrhea, abdominal pain, and, in the case of ulcerative colitis, rectal bleeding. The progression of the disease frequently extends over several decades, and its prevalence is increasing at a global scale [7].

Managing this condition is often challenging, primarily due to limited access to comprehensive patient records. This restricts informed decision-making and can result in a significant amount of manual data processing if the complete medical history of a patient is to be considered. By helping physicians analyze the patient's medical history, symptoms, and response to previous treatments, MDTs can support insights that enable personalized care. Treatment planning is also informed by medical imaging results. Endoscopies are a diagnostic tool that can provide more qualitative information about the location and extent of intestinal inflammation. Echographies, on the other hand, are used to assess disease severity, such as in the *Limberg score*. The availability of these results within the MDT facilitates the identification of correlations with other components of the data set. This comprehensive monitoring can assist in detecting early indications of disease flare-ups or complications, enabling timely intervention and averting severe consequences.

Moreover, patients have the capacity to record symptoms themselves within the MDT, thereby becoming active participants in their own care by monitoring their health parameters. As IBD is a lifelong disease, being able to easily and quickly record daily symptoms is very helpful, as it provides real-time data on various parameters such as inflammation levels, bowel movements and medication adherence. Within our secure MDT approach, the management of patient data pertaining to IBD (including symptoms, laboratory results and imaging biomarkers) is accomplished through individually encrypting each data entry. This ensures that only authorised users are able to decrypt the precise information they require.

Overall, our MDT approach transforms IBD management by combining continuous patient-reported data, endoscopic and imaging results, and laboratory biomarkers in a secure, collaborative framework – supporting personalized treatment insights, reducing manual record reconciliation, and safeguarding patient trust through robust cryptographic controls.

5 Future Work

Further research and development are necessary to implement our proposed medical digital twin (MDT) in personalized healthcare. Our goal is to establish a solid foundation for the creation and operation of MDTs. This involves collecting medical data and diagnosis- and treatment-related information from trusted sources to create comprehensive models of a patient's history and health status. Throughout the development process, we maintain close contact with potential users from both the medical community and patients to ensure a user-centered approach. We are currently conducting a survey aimed at practicing physicians to better understand the current landscape of medical data exchange. In close collaboration with medical experts, we will further develop practical use cases. This will include the development of user-friendly interfaces that allow medical experts to request access to and utilize specific data within the MDT. Medical personnel will be able to populate the MDT with new examination results. On the other side, patients will be able to intuitively grant, revoke or change access rights to specific data while also being able to integrate personal data into the MDT. For that, we will implement encryption-based access concepts to ensure the security of processing. We are working on interoperability with the FHIR standard and HL7[15]. The newest version of FHIR (v6.0.0), which will probably be released in the coming months, will support an attribute-level rights management, which may allow us to integrate our encryption concept with the HL7/FHIR approach.

6 Conclusion

This paper introduces the concept of secure medical digital twins as a new means to support important medical use cases that require the sharing of medical data. By securing the data and the access to the data in a decentralized way, we foster the secure sharing of medical data which can otherwise be cumbersome with the more traditional perimeter- and transfer-oriented security concepts. This has benefits for patients, who have full transparency of the grants of access to their medical data, and for medical experts and institutions, who can rely on these patient approvals during the treatment process and for research purposes. We showed in two example use cases how these benefits manifest themselves during treatment planning for Parkinson's disease and for inflammatory bowel disease. We are working on several prototypical implementations that will provide the basis for future medical information systems as well as tele-medical applications.

Acknowledgments. This research work was supported by the National Research Center for Applied Cybersecurity ATHENE. ATHENE is funded jointly by the German Federal Ministry of Research, Technology and Space and the Hessian Ministry of Science and Research, Arts and Culture. The work was partly funded by the Fraunhofer Cluster of Excellence Immune-mediated Diseases CIMD.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Baumgart, D.C., Carding, S.R.: Inflammatory bowel disease: cause and immunobiology. Elsevier pp. 1627–1640 (2007)
2. Becker, F., Bibow, P., Dalibor, M., Gannouni, A., Hahn, V., Hopmann, C., et al.: A Conceptual Model for Digital Shadows in Industry and its Application. In: Ghose, A., Horkoff, J., Silva Souza, V.E., Parsons, J., Evermann, J. (eds.) Conceptual Modeling, ER 2021, pp. 271–281. Springer (2021)
3. Björnsson, B., Borrebaeck, C., Elander, N., Gasslander, T., Gawel, D.R., et al.: Digital twins to personalize medicine. Genome Medicine **12**(4) (2020)
4. Bujari, A., Calvio, A., Foschini, L., Sabbioni, A., Corradi, A.: IPPODAMO: a Digital Twin Support for Smart Cities Facility Management. In: GoodIT '21. pp. 49–54. ACM (September 2021)
5. Bundesamt für Sicherheit in der Informationstechnik (BSI): Positionspapier Zero Trust 2023. Tech. rep., Bundesamt für Sicherheit in der Informationstechnik (BSI), Bonn (2023)
6. Chase, M.: Multi-authority attribute based encryption. In: Vadhan, S.P. (ed.) Theory of Cryptography. pp. 515–534. Springer Berlin Heidelberg (2007)
7. Collaborators, G.B.D., et al.: Global, regional, and national incidence, prevalence, and years lived with disability for 354 diseases and injuries for 195 countries and territories, 1990-2017: a systematic analysis for the Global Burden of Disease Study 2017. University of Leicester (2018)
8. Coorey, G., Figtree, G.A., Fletcher, D.F., Snelson, V.J., Vernon, S.T., Winlaw, D., et al.: The health digital twin to tackle cardiovascular disease - a review of an emerging interdisciplinary field. npj Digital Medicine **5**(126) (October 2022)
9. Erol, T., Mendi, A.F., Doğan, D.: The digital twin revolution in healthcare. In: Proc. of 4th International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT). pp. 1–7. IEEE (October 2020)
10. European Union Agency for Cybersecurity (ENISA): Identifying Emerging Cyber Security Threats and Challenges for 2030 (2023)
11. Farsi, M., Daneshkhah, A., Hosseinian-Far, A., Jahankhani, H. Jahankhani, H. (eds.): Digital Twin Technologies and Smart Cities. Springer, Cham (2020)
12. Grieves, M.W.: Product lifecycle management: the new paradigm for enterprises. Int. J. Product Development **2**(1), 2 (2005)

13. Haleem, A., Javaid, M., Singh, R.P., Suman, R.: Exploring the revolution in healthcare systems through the applications of digital twin technology. Biomedical Technology **4**, 28–38 (December 2023)
14. Haße, H., van der Valk, H., Möller, F., Otto, B.: Design Principles for Shared Digital Twins in Distributed Systems. Business & Information Systems Engineering **64**(6), 751–772 (April 2022)
15. Health Level Seven International, Inc.: HL7 FHIR (2025), https://hl7.org/fhir/ [Accessed: 2025-07-24]
16. Human, C., Basson, A.H., Kruger, K.: A Design Framework for a System of Digital Twins and Services. Computers in Industry **144** (2023)
17. Icons: DB server by OpenClipArt (CC0 1.0); Free Large Boss Icon Set by www.aha-soft.com; Health device icons designed by macrovector_official / Freepik
18. Jiang, Y., Wu, S.M.R., Liu, M., Luo, H., Kaynak, O.: Monitoring and Defense of Industrial Cyber-Physical Systems Under Typical Attacks: From a Systems and Control Perspective. IEEE Transactions On Industrial Cyber-Physical Systems. **1** (2023)
19. Jones, D., Snider, C., Nassehi, A., Yon, J., Hicks, B.: Characterising the Digital Twin: A systematic literature review. CIRP Journal of Manufacturing Science and Technology **29**, 36–52 (May 2020)
20. Lewko, A., Waters, B.: Decentralizing attribute-based encryption. In: Paterson, K.G. (ed.) Advances in Cryptology – EUROCRYPT 2011. pp. 568–588. Springer Berlin Heidelberg (2011)
21. National Institute of Standards and Technology (NIST): American National Standard for Information Technology - Role Based Access Control, INCITS 359-2012 (R2022) (May 2012)
22. Piovani, D., Danese, S., Peyrin-Biroulet, L., Bonovas, S.: Inflammatory bowel disease: estimates from the global burden of disease 2017 study. Wiley Online Library pp. 261–270 (2020)
23. Sahai, A., Waters, B.: Fuzzy identity-based encryption. In: Cramer, R. (ed.) Advances in Cryptology – EUROCRYPT 2005. pp. 457–473. Springer Berlin Heidelberg (2005)
24. Su, D., Cui, Y., He, C., Yin, P., Bai, R., Zhu, J., Lam, J.S.T., Zhang, J., Yan, R., Zheng, X., Wu, J., Zhao, D., Wang, A., Zhou, M., Feng, T.: Projections for prevalence of parkinson's disease and its driving factors in 195 countries and territories to 2050: modelling study of global burden of disease study 2021. BMJ **388** (2025)
25. Vallée, A.: Digital twin for healthcare systems. Frontiers in Digital Health **5** (September 2023)
26. VanDerHorn, E., Mahadevan, S.: Digital Twin: Generalization, characterization and implementation. Decision Support Systems **145** (June 2021)
27. Verdouw, C., Tekinerdogan, B., Beulens, A., Wolfert, S.: Digital twins in smart farming. Agricultural Systems **189** (April 2021)
28. Zahra, W., Rai, S., Birla, H., Singh, S., Dilnashin, H., Rathore, A., et al.: The Global Economic Impact of Neurodegenerative Diseases: Opportunities and Challenges. In: Bioeconomy for Sustainable Development. Springer, Singapore (2020)
29. Zhang, L., Kan, H., Huang, H.: Patient-Centered Cross-Enterprise Document Sharing and Dynamic Consent Framework using Consortium Blockchain and Ciphertext-Policy Attribute-Based Encryption. In: CF '22, Proceedings of the 19th ACM International Conference on Computing Frontiers.; 2022. pp. 58–66. ACM (May 2022)