

Perspective

Digital twins and the ethics of health decision-making concerning children

Matthias Braun^{1,*} and Jenny Krutzinna²¹Friedrich-Alexander-Universität Erlangen-Nürnberg, Research Group Ethics and Governance of Emerging Technologies, Kochstraße 6, 91054 Erlangen, Germany²Centre for Research on Discretion and Paternalism, Department of Administration and Organization Theory, University of Bergen, Bergen, Norway*Correspondence: matthias.braun@fau.de<https://doi.org/10.1016/j.patter.2022.100469>

THE BIGGER PICTURE There is a substantial need in child protection to design the decision-making process in a way that is in the best interests of the child. The solution to this problem will not lie in new technology alone but also in new techniques and technologies that are urgently needed to make children and their interests more visible and to integrate them in decision-making processes. In the health context, this concerns particularly better knowledge of the health status of those children who are especially dependent on the vicarious decisions of others. In doing so, however, we are confronted with an ethical dilemma: on the one hand, children are a particularly vulnerable group, dependent on empowerment and opportunities for genuine participation. In this regard, digital twins (DTs) may provide a substantive opportunity to empower children by providing better and more precise information on their behalf. On the other hand, DT is a technology with great potential to add new forms of vulnerability through its constant, real-time, and *ad personam* predictions. Consequently, we argue that DTs hold significant potential for a positive contribution to these processes provided that critical concerns regarding vulnerability, recognition, and participation are adequately addressed.



Concept: Basic principles of a new data science output observed and reported

SUMMARY

In this review, we explore from an ethical perspective the opportunities and challenges for decision-making concerning children if digital twins (DTs) were to be used to provide better information about their health status as a basis for proxy decision-making. We note a sense of urgency due to the speed of progress and implementation of this advancing technology and argue that bringing a solid conceptual basis into the development process is of utmost importance for the effective protection of children's rights and interests.

INTRODUCTION

There is currently an emerging technological development, the so-called "digital twin" technology. While still in its nascency, it is already perceived as one of the top ten future technology trends.¹ Broadly, digital twins (DTs) are high-precision simulations that map and model events or objects in real time.² In the health context, a DT essentially comprises an artificial intelligence (AI)-driven, real-time simulation of a person's health status created with the goal of predicting future developments.³ One use of DTs deserves special attention because of its promise to contribute new solutions to a public health problem that is currently facing major challenges: children's health monitoring and the resulting decisions that are made on their behalf.

In this review, we explore from an ethical perspective the opportunities and challenges for decision-making concerning

children if DTs were to be used to provide better information about their health status as a basis for proxy decision-making. We note a sense of urgency due to the speed of progress and implementation of this advancing technology and argue that bringing a solid conceptual basis into the development process is of the utmost importance for the effective protection of children's rights and interests.

To further understand what benefits and drawbacks are associated with the development and use of DTs in the context of child welfare, we will explore our hypothesis in four steps. First, we briefly analyze the fundamental challenges of proxy decisions that must be made because of particular extrapolations of vulnerability. A key challenge with regard to proxy decisions in child welfare is them being made on the basis of often incomplete or even faulty information about the child's health status. In a second step, we therefore take a closer look at current



Table 1. Comparison of different decision-making approaches

	Real-time	Use of big data	Ongoing data input	Algorithmic data processing	Physical embodiment
Conventional	Y	N	Y	N	N
Prediction-based	N	Y	N	Y	N
Digital twins	Y	Y	Y	Y	Y

Illustrated is the conventional approach to decision-making compared with two alternative approaches (prediction-based and DTs) to the use of health data.

technological approaches to incorporate health status information into the decision-making process and briefly analyze their advantages and disadvantages. Against the background of these two steps, we then look at a third step as a new approach, the so-called DTs. Finally, we ask what ethical conditions can be defined for future development and use in child welfare in light of the analyses.

MAKING DECISIONS ON BEHALF

There is a substantial need in child protection to design the decision-making process in a way that is in the best interests of the child. (We are using the terms “child protection” and “child welfare” interchangeably and in a broad sense, referring to states’ general mandate to safeguard children. As such, country- and/or system-specific policies have no bearing on our analysis, which provides an ethical assessment of DTs in decision-making concerning children to include children’s health status.). The solution to this problem will not lie in new technology alone but also in new techniques and technologies that are urgently needed to make children and their interests more visible and to integrate them in decision-making processes. In the health context, this particularly concerns better knowledge of the health status of those children who are especially dependent on the vicarious decisions of others. In child protection, this is even more important because parenting insufficiencies often mean that the welfare of the children is not safeguarded unless an agent of the state (such as a child welfare worker) intervenes. In doing so, however, we are confronted with an ethical dilemma: on the one hand, children are a particularly vulnerable group, dependent on empowerment and opportunities for genuine participation. In this regard, DTs may provide a substantive opportunity to empower children by providing better and more precise information on their behalf. On the other hand, DT is a technology with great potential to add new forms of vulnerability through its constant, real-time, and *ad personam* predictions. Consequently, we argue that DTs hold significant potential for a positive contribution to these processes provided that critical concerns regarding vulnerability, recognition, and participation are adequately addressed.

A hypothetical case exemplifies the types of decisions typically arising in a child-protection context: Maria works as hotline screener at the frontlines of child welfare/child protection. Her workday consists of taking calls from people who, for various reasons, worry about the well-being of a particular child or children.⁴ Maria and the callers may not have much in common, but during this moment of decision-making, they are all clearly acting on behalf of—or even in the name of—the child they are concerned about. There is, however, another side to the same

coin: the child themselves is not directly present in this multi-layered decision-making process.^{5,6} They do not have their own say and cannot be heard with their own voice. While this form of absence is structurally inherent in decision-making systems, there is another form of absence that is by no means necessarily present: many decision-making situations lack precise knowledge of the child’s current state of health. Consequently, when Maria has to decide whether she should start an intervention to secure and foster the well being of the respective child, she acts as if she has knowledge of the child’s concrete health conditions as well as their living conditions.

This describes a typical situation for decision-makers in the child welfare system. We refer to to these sorts of decision-making as conventional (see Table 1). It illustrates the importance of precision and robustness of the underlying decision-making process since stakes are high in that any decision made will have serious consequences for the child concerned. The risk inherent in the decision-making process is that in the absence of reliable facts on all relevant aspects, the decision-maker (here, Maria) will rely predominantly on her experience and instinct when making a decision. While this is undoubtedly an important perspective, it cannot be the sole source for a decision, especially one made on behalf of another, who in this situation is a child in a vulnerable situation. Solid decision-making processes must have an objective basis, including facts and evidence, so that their interpretation and conclusion can be justified. In practice, however, the decision-making situation described above is characterized by a lack of comprehensive information about the child concerned, meaning that knowledge about the child’s (health) condition is rudimentary at this moment in time, combined with a clear sense of urgency. Decision-makers like Maria are under enormous pressure to make the right decision but all too often lack adequate resources, a fact reflected in personnel shortages and very high levels of staff fluctuation in the child welfare system.^{7–10}

MAKING GOOD DECISIONS CONCERNING CHILDREN

Any good decision concerning children will first and foremost respect children and their rights. Children are not merely small human adults¹¹ and, thus, have been afforded special legal rights, most prominently through the Convention on the Rights of the Child (CRC).¹² The CRC, ratified in all countries except the USA, seeks to ensure, *inter alia*, that decisions concerning children focus primarily on the children themselves through the principle of the child’s best interests and the child’s right to participation. States are to make the best interests of the child a primary consideration “in all actions concerning children” (Article 3 CRC), a fundamental provision relevant to the

interpretation of all other CRC rights. Children are also granted participatory rights: to express their views freely, to have their views given appropriate weight in all matters affecting them, and to be heard in proceedings affecting them (Article 12 CRC).

Despite these strong safeguards, children's rights are not always respected. The best-interests principle has been criticized as vague (e.g., Kelly¹³), and interpretations of children's age- and maturity-related capacities have affected their inclusion and participation in child welfare proceedings.¹⁴ Consequently, younger children are often excluded from direct participation based on an interpretation of their capacities and maturity, even when age thresholds suggest they should be included. These deficits remain despite the development of conceptual models to aid implementation in practice^{15,16} and ongoing efforts to raise awareness for the need to improve child centrism.¹⁷

Thus, respecting the rights of children, addressing their needs, and ensuring their health and well being remains a particular challenge.⁴ Oftentimes, children are unable to stand up for their needs themselves or point out grievances regarding their health and well being, whether due to young age or impairment or because of existing power imbalances or their "oppression" as children (see Barth and Olsen¹⁸). Instead, children remain dependent on adult decision-making about and for them: parents, doctors, judges, or social workers act on behalf of the child they are concerned with. While certain circumstances in a person's life clearly justify proxies, acting on behalf of someone is not legitimate *per se*. Legitimacy hinges on two conditions: first, an actual limitation of the person's (here, child's) self-determination regarding decision-making, and second, due consideration of the represented person's needs and interests as seen from that person's perspective. Consequently, surrogate decision-making should be limited to instances where it is strictly necessary. Where it is unavoidable, participatory approaches are called for that include the children themselves in the decision-making process. Currently, the child themselves is absent from the multi-layered decision-making processes,^{5,6} without having a say or being heard directly. An example is child-protection interventions determining if a care order should be made to protect a child's well-being. In such cases, the decision-maker acts as if they have concrete health and social information, but this is exclusively provided by third parties from their adult perspective and not the child themselves.¹⁹ While DTs will not solve this problem alone, they may provide a useful additional step toward its resolution, as we explain below.

ATTEMPTS TO GAIN MORE HEALTH INFORMATION

Child-protection professionals are required to base their decision on the child's individual circumstances and with the child's best interests being a primary consideration.²⁰ This requires knowledge about typical health trajectories for children born with withdrawal symptoms following maternal antenatal substance misuse and comprehensive knowledge about the specific circumstances the child was born into, as well as information about the child themselves. For a thorough risk assessment based on the child's best interests, contextual information regarding the child's social and family background is as relevant as the child's developmental condition,²¹ which, besides the withdrawal symptoms, may be affected by premature birth,

low birth weight, or congenital conditions such as heart defects. Solid decision-making processes must undoubtedly be built on an objective basis, including empirical evidence, so that their interpretation and conclusion can be justified. In practice, however, in the initial decision-making situation described above, only partial knowledge about the child is available, for instance, concerning the child's (health) condition. Additionally, the respective decision-making processes often take place under great time pressure. Thus, the conventional approach of social workers using their previous experience and training and combining it with knowledge about the specific case to make a decision has obvious shortcomings in that the information available is likely to be patchy and time pressure to be high (see Table 1).

Improving both decision-making quality and efficiency is a priority, and a need for change in the current system has been recognized.²² Recently, aided by the wider digitization movement, decision-making using predictive algorithms to assist decision-making processes has emerged.²³ This, what we call prediction-based decision-making (see Table 1), raises multiple technical, legal, and ethical concerns.²⁴ Indeed, predictive analytics may be used to optimize proxy decision-making processes and to develop more standardized protocols and routines, but beside the familiar issues of algorithmic bias, black boxes, and concerns relating to justice and fairness,²⁵ these technologies provide predictions based on correlations in population-based samples. This approach is not only error prone and socially problematic, as demonstrated by the Allegheny County Office of Children, Youth, and Families case where a lack of data on actual maltreatment led to the risk model predicting which families get reported by the community rather than which children were likely to be maltreated and essentially equated "parenting while poor" as "poor parenting";²⁶⁻²⁸ there is also a general concern that the mass of data collected may simply "begin to speak for children."²⁹ For instance, empirical research analyzing written decisions in child-protection cases reveals that information concerning the children themselves is often weak and that the emphasis commonly is on parental (mal)functioning,¹⁷ which means that predictive technology based on these data would inherit this "oblivion" to the child and would be based entirely on historical samples. Thus, current technologies still struggle with bringing real-time information about the respective child into the proxy decision-making process as well as allowing bidirectional ways of interaction between the technical system and the agents who have to make decisions.

TRANSFORMING DECISION-MAKING: DTs AS A SOLUTION?

While the ambivalences of acting on behalf of others remain a challenge, DTs seem to offer a new (more or less visionary) option to address the challenge of missing health information in decision-making concerning children.

Specifically, a DT is a mirror of a physical process, articulated alongside that process and—in most cases—exactly matching the operation of that physical process in real time.^{30,31}

DT has its roots in the simulation of simple (physical) objects. In recent years, the level of complexity has massively increased and now includes buildings, factories, and entire production

processes with the aim of predicting future developments, optimizing processes better and faster, and preventing possible errors before they occur in the physical world. DT systems are equipped with various forms of AI and machine learning and are highly automated.

Different types of data can be used to describe a physical object. Real-time data, collected by wearables, for example, should ensure that the respective health simulation correctly depicts the current state. Also required are data on previous illnesses or possible chronic diseases (historical data) as well as classic medical data on physical parameters such as height, weight, etc. Depending on the type of use, the integration of other data such as genetic data would also be possible.

In particular, the supposed predictive power of such systems makes them very interesting for the simulation of living objects or systems. DTs are thus also becoming increasingly interesting in the context of health. In the last few years, massive developments have taken place that enable different areas of application for DTs in the context of health. For instance, a DT simulation may detect a heart condition in a child via prediction before it even arises. This prediction is made based on the digital artifact's simulated trajectory; the simulation may then go on to suggest a prognosis and recommend an optimal course of treatment.³² The DT may also directly interact and communicate with the physical person whom the twin simulates to warn her or to suggest relevant lifestyle changes.³³ Furthermore, it would be conceivable to use DTs to test and validate specific drugs or treatments³⁴ and their repercussions for physical and/or mental health.³⁵ The use of DTs with the aim of public-health monitoring has recently been discussed by Laubenbacher et al.³⁶ The basic idea is to use DTs to better predict viral infections and pandemics and to develop appropriate measures.

Compared with conventional approaches (see Table 1), DTs promise two key features: first, providing health information and predictions in real time, and second, enabling bidirectional interactions between DTs, decision-makers, and, depending on the design of the models, the child themselves. Compared with predictions-based models (see Table 1), DTs may ensure two further assets: firstly, predictions are not made solely on the basis of correlation-based models, but the actual state of health of the respective child can be depicted. Secondly, such a simulation could take place in real time by collecting data via sensors and wearables, which in many cases are already worn and used by the children. But how exactly would this work? To better understand how DTs can be a possible solution to improve proxy decision-making processes and to better explore the ethical challenges involved, we need to briefly define what exactly are the conditions and requirements for using DTs to simulate the health conditions of a respective person.

A first requirement is that there is a corresponding DT for a real, physical person, which is provided with an unmistakable identification. This identification could also be used to log in to one's virtual representation and to make settings as to which health parameters should be recorded and who should be able to view them.

Secondly, it must be ensured that the DT is created correctly and that the data are verified. Likewise, there needs to be an institutional anchor for the creation of a DT, which is also authorized to create such a DT. Similar to decisions on newborn

screening, one could think of different models of how the coordination between clinicians, parents, and others can take place. To maintain privacy, passwords, fingerprints, iris recognition, authorized login, blockchain encryption, authentication, identification, and other technology can be used as encryption to prevent tampering. For children, the account and password can be controlled by their parents or guardians but also by themselves. When necessary, doctors or professionals can be authorized to inquire about or modify them. What we describe here is not a complete list and needs further elaboration. Central to our investigation here are two things: first, that the focus on securing privacy and preventing misuse of data, and especially the created applications, is already and should be a central focus in the further development of such health simulations. Second, however, it is equally relevant that the privacy of children in vicarious decisions is at risk either way. Having more information about their health status as a basis for making proxy decisions may be an intrusion upon children's privacy. However, not having it available and making vicarious decisions based on a poorer level of health information also represents a restriction of privacy.

Thirdly, a constant exchange of data between the sensors and wearables on the one hand and the DT on the other is required. This would ensure that changes and dynamics in the health status of the respective person are also reflected in the DT. For this purpose, techniques such as blockchain technologies would have to be used in the future, for example, in order to maintain the accuracy and precision of the simulation, even if the data flow is interrupted or irregular.

Fourthly, there is a need to define precisely in which area DTs should be used and what they should define. Depending on this are the type of data, for example, text (diagnoses records), numbers (weight, blood pressure, heart rates), or images (electrocardiogram), that are to be transferred and how they are to be processed. It will be crucial to define the interfaces at which the different data can be fed in and transferred.

A fifth requirement is to determine to what point which form of feedback must be given by the DT. It is conceivable that feedback is given to the simulated person themselves, to the responsible doctor, or, as in Maria's case, to the decision-maker in charge. The results of the health assessment will be fed back to the real person and their doctors in a hospital as predictions, suggestions, guidance, alarms, or treatment schedules. This information allows the real person to make certain improvements, such as increasing exercise, improving nutrition, seeking further treatment, etc.

Many of these technical requirements have not yet been met. In particular, the calculation and modeling of such large amounts of data still pose significant problems. At the same time, further developments in the field of quantum computing and AI indicate that this problem could be solved sooner than it currently seems.³⁷

One could take this as an argument to stop thinking about the ethical requirements for the possible use of DTs in the field of decision support systems concerning children. However, two reasons immediately speak against this: in the context of AI, but also in the context of other technologies such as CRISPR, we have to learn how challenging, if not impossible, it is to think about the normative conditions and consequences only when a technology is already in use.³⁸ A good example of this is the current

debates about facial recognition and its use in different areas of society.³⁹ Of course, one can try to recapture this through appropriate governance approaches, as is currently being done in several debates. But in doing so, one misses the fact that social and societal transformations associated with the introduction of any technology should not only be taken into account from the beginning but also that different possibilities for shaping the further development of the technology have to be developed. Especially when the fundamental vulnerability of human life as a whole and of concrete injuries in the respective proxy decision-making process is taken seriously is it important to ask which goals should be pursued in the development of a technology. Thinking about how DTs can be developed from an ethical perspective in such a way that they eliminate existing grievances in proxy decision-making, and at the same time make the respective representatives more visible, is central to responsibly shaping the future.

Thus, we need to ask under which conditions and with what aims DTs could be developed and used to make children more visible in decision-making processes concerning them to ensure the safe implementation of this technology in child welfare. Given that DTs are currently being developed, one does not need be a prophet to anticipate their future implementation in medicine, which makes the consideration of both positive and negative consequences of such technology a pressing issue. To uncritically hail this technology as the solution to the problems within proxy decision-making processes in child welfare would mean succumbing to technological solutionism,⁴⁰ violating the rights of those who should be protected. Nevertheless, merely maintaining the status quo of current actions and processes will not lead to better solutions for those children who are fundamentally dependent on somebody acting on their behalf in certain circumstances and who has the best possible health information to do so.

POTENTIAL RISKS AND BENEFITS OF DTs IN CHILD PROTECTION

As analyzed above, a central key challenge in child welfare decision-making stems from the complex circumstances surrounding such decisions, which carry high stakes and serious consequences for the child and their family, making appropriate child representation and participation crucial. Involving the child can, however, be a delicate issue due to the urgency and limitations of current modes of participation, which is where DTs provide a unique opportunity to improve decision-making.

First, when child-protection agents receive a notification of concern for a child, they have to act immediately and carefully evaluate potential next steps. Here, they will typically have limited information available, requiring them to first engage in cumbersome information gathering from health professionals (and others) to understand the child's situation. A DT representation of the child would allow the simulation to function as a surrogate for the simulated child with regard to their health by providing a consolidated and up-to-date record of all information on the particular child collected by multiple professionals/services. This means that for doctors, judges, and social workers, the DT could, for example, provide a provisional health status of the physically absent child, thus putting decisions on whether

to intervene on a perhaps more solid ground and helping with daily decisions concerning several children.

Second, current modes of participation stipulate talking and listening to the child directly. While this can clearly be seen as the “gold standard,” it may sometimes be necessary to carefully manage how and by whom the child is addressed directly. DTs would enable this direct evidence to be integrated and made available to all professionals or decision-makers with a legitimate interest. This would increase the utility of the child's testimony by increasing its reach but would also serve to reduce the number of challenging conversations the child has to have with adults while still revealing the child's needs and preferences. Such conversations with adults may be intimidating, the physical space may be child-unfriendly, or past experiences and trauma may render the “hearing” potentially harmful to the child. In cases where direct testimony is impossible (e.g., due to young age) or undesirable (e.g., due to its harmful effects), DTs could still be used to integrate health-status information into the decision-making process in a constant, timely, and comprehensive way. This will not solve the dilemma that children too rarely have a direct say in decision-making processes, but if DTs are used in such a way that they provide reliable information about the respective health status when needed, this would result in fewer cases where very far-reaching decisions for children's futures are made based on very vague evidence. Returning to the case of Maria and her decisions on behalf of the respective child, whose name is Lucy: a possibility would be to use DTs to create a simulation of Lucy's recent health status and use it to get more and better information about, for example, Lucy's health or to measure the stress level Lucy is exposed to in her current family constellation. These data could, for instance, be used to simulate predictions about different health trajectories depending on conceivable family constellations. Any input from Lucy herself would also be incorporated.

Thirdly, a central ethical dilemma is associated with the use of DTs in decision-making concerning children. The very nature of this decision-making enhancement means that risks relating to representation may also be enhanced.⁴¹ We must thus be aware of the fundamental vulnerability of human beings arising from the dependence on others, an experience common to all humans, not only children, and which may thus be considered part of the human condition.^{42–44} In the child-welfare context, the risk of inadvertently violating children's rights by failing to recognize this vulnerability in how children are represented and involved in decision-making processes needs to be tackled head on.⁴⁵

This tension can be seen very concretely in questions of privacy and surveillance. To be clear, the use of DTs would inevitably lead to a form of surveillance that curtails the rights to privacy. There is no getting away from that at this point. However, the crux of the matter is that the existing procedures also lead to massive encroachments on privacy and the right to bodily integrity.⁴⁶ The question that arises is whether there is an acceptable level of surveillance that, when applied and used correctly, leads to greater respect for the right to bodily integrity. If we look at this question more closely, we notice that there are many examples where we decide this with good reasons. In addition to traffic navigation, for example, the current coronavirus disease 2019 (COVID-19) public-health monitoring is particularly worthy of mention in the context of health. We (necessarily) use information

about individual health conditions in order to increase the health protection of both a larger number of people and the respective individuals themselves. Surveillance and privacy are not simply mutually exclusive from this perspective, but vulnerability dilemmas can arise in which a higher level of surveillance leads to a higher level of respect for the right to bodily integrity.

That this is the case, however, is not easily guaranteed. Instead, three dimensions have to be addressed to involve children more directly in matters concerning their health and well-being: first, there must be a mandatory focus on the child as a person with their own rights and fundamental desires (person-centric). Second, children must be encouraged and enabled to reflect on the procedures and structures that affect them and to implement these reflections on their experience in the design and development of child welfare and health systems (enabling reflection). Third, children and their perspectives must be introduced into the modes of control and evaluation (allowing control and evaluation).

CONCLUSIONS

A first condition is to design and develop the DTs to be dynamically adjustable to changing preferences and changing health parameters. This means that the DT must be able to map health conditions in real time. This is an important precondition to ensure the proxy decision also reflects actual and possibly changing preferences and health conditions of the respective child. To make such a dynamic possible, a transparent standard for the inclusion and exclusion criteria of the recommended decisions is needed. Social workers but also the judges and, not least, the doctors consulted must be able to understand and evaluate the respective basis for the representative simulation and whether it is appropriate and compliant with applicable laws in the respective jurisdiction.

Second, dynamic development of these systems also implies that children must be able to withdraw information from the decision-making system or change their preferences. In this respect, the condition to develop the twins dynamically is more-over linked to the fact that these systems allow real participation. Genuine participation also includes an easily accessible way to give the child control over the use of simulated representation. The degree of control could be linked to certain developmental steps of the child. At this point, however, it becomes clear that the DTs cannot replace the responsibility of human actors in the decision-making process. On the contrary, the great opportunity of using such digital technologies in child welfare decision-making is precisely to make visible hidden—but for decision-making purposes very important—medical information or individual preferences of the respective child.

Thirdly, one way of dealing with this challenge could be to integrate certain institutional safeguards. For example, as in the area of vaccination, a commission could be institutionally anchored that constantly deals with the security of the data created and that ensures that quick action can be taken in the case of possible violations of the children's privacy or possible structural errors in the procedures. One possibility to strongly regulate access could be to encrypt the data and simulations in such a way that their results are only used in certain suspicious cases or for pending decisions.

ACKNOWLEDGMENTS

M.B. is grateful for funding from the German Research Foundation (EmpkinS, Research Collaboration Center 1483, ID 442419336). The funder has no role in design and results of this research. J.K. has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement no. 724460). Publications from the projects reflect only the authors' views, and the funding agency is not responsible for any use that may be made of the information contained therein.

AUTHOR CONTRIBUTIONS

M.B. and J.K. have equally contributed to all parts of the article. M.B. and J.K. equally share first authorship.

DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCES

- Rathore, M.M., Shah, S.A., Shukla, D., Bentafat, E., and Bakiras, S. (2021). The role of AI, machine learning, and big data in digital twinning: a systematic literature review, challenges, and opportunities. *IEEE Access* 9, 32030–32052.
- Saddik, A.E. (2018). Digital twins: the convergence of multimedia technologies. *IEEE MultiMed.* 25, 87–92.
- Jones, D., Snider, C., Nassehi, A., Yon, J., and Hicks, B. (2020). Characterising the Digital Twin: a systematic literature review. *CIRP J. Manuf. Sci. Technol.* 29, 36–52.
- Berrick, J.D. (2018). *The Impossible Imperative* (Oxford University Press).
- Rafeedie, J., Hudson, S.M., Deavenport-Saman, A., Rao, S., Rogers, K., and Roberts, S. (2019). Decision-making in foster care: a child-centered approach to reducing toxic stress in foster children. *Child. Youth Serv. Rev.* 96, 10–16.
- Coyne, I., Hallström, I., and Söderbäck, M. (2016). Reframing the focus from a family-centred to a child-centred care approach for children's healthcare. *J. Child Health Care* 20, 494–502.
- Châteauneuf, D., Turcotte, D., and Drapeau, S. (2018b). The relationship between foster care families and birth families in a child welfare context: the determining factors. *Child Fam. Soc. Work* 23, 71–79.
- Châteauneuf, D., Pagé, G., and Decaluwé, B. (2018a). Issues surrounding post-adoption contacts in foster adoption: the perspective of foster-to-adopt families and child welfare workers. *J. Publ. Child Welfare* 12, 436–460.
- Berrick, J.D., Dickens, J., Pösö, T., and Skivenes, M. (2016). Time, institutional support, and quality of decision making in child protection: a cross-country analysis. *Hum. Serv. Organ. Manag. Leadersh. Govern.* 40, 451–468.
- Juhász, I., and Skivenes, M. (2018). Child welfare workers' experiences of obstacles in care order preparation: a cross-country comparison. *Eur. J. Soc. Work* 21, 100–113.
- Archard, D.W. (2018). Children's rights. *Stanf. Encycl. Philos.*
- UN General Assembly (1989). *Convention on the rights of the child. Treaty Ser.* 1577.
- Kelly, J.B. (2005). The best interests of the child: a concept in search of meaning. *Fam. Court Rev.* 35, 377–387.
- McEwan-Strand, A., and Skivenes, M. (2020). Children's capacities and role in matters of great significance for them. *Int. J. Child. Right* 28, 632–665.
- Krutzinna, J. (2022). Who Is 'the Child'? Best Interests and Individuality of Children in Discretionary Decision-Making. *The International Journal of Children's Rights* 30, 120–145.

16. Lundy, L. (2007). 'Voice' is not enough: conceptualising article 12 of the United Nations convention on the rights of the child. *Br. Educ. Res. J.* 33, 927–942.
17. Kríž, K., Krutzinna, J., Pösö, T., Skivenes, M. The Invisible Child: A Comparative Study of Newborn Removal Judgments from a Child Equality Perspective (CEP). *The International Journal of Children's Rights*.
18. Barth, R.P., and Olsen, A.N. (2020). Are children oppressed? The timely importance of answering this question. *Child. Youth Serv. Rev.* 110, 104780.
19. Krutzinna, J., and Skivenes, M. (2020). Judging parental competence: a cross-country analysis of judicial decision makers' written assessment of mothers' parenting capacities in newborn removal cases. *Child Fam. Soc. Work.* <https://doi.org/10.1111/cfs.12788>.
20. UN Committee on the Rights of the Child (2013). General comment No. 14 (2013) on the right of the child to have his or her best interests taken as a primary consideration (art. 3, para. 1). <https://www.refworld.org/docid/51a84b5e4.html>.
21. Ward, H., Brown, R., and Westlake, D. (2012). *Safeguarding Babies and Very Young Children from Abuse and Neglect* (Jessica Kingsley Publishers).
22. Cuccaro-Alamin, S., Foust, R., Vaithianathan, R., and Putnam-Hornstein, E. (2017). Risk assessment and decision making in child protective services: predictive risk modeling in context. *Child. Youth Serv. Rev.* 79, 291–298.
23. Gillingham, P. (2019). Can predictive algorithms assist decision-making in social work with children and families? *Child Abuse Rev.* 28, 114–126.
24. Munro, E. (2019). Predictive analytics in child protection. *CHES Work. Pap.*
25. Keddell, E. (2019). Algorithmic justice in child protection: statistical fairness, social justice and the implications for practice. *Soc. Sci.* 8, 281.
26. Eubanks, V. (2018a). A child abuse prediction model fails poor families. *Wired* <https://www.wired.com/story/excerpt-from-automating-inequality/>.
27. Eubanks, V. (2018b). A response to Allegheny county DHS. Va. Eubanks <https://virginia-eubanks.com/2018/02/16/a-response-to-allegheny-county-dhs/>.
28. Hurley, D. (2018). Can an algorithm tell when kids are in danger? *The New York times magazine*. <https://www.nytimes.com/2018/01/02/magazine/can-an-algorithm-tell-when-kids-are-in-danger.html>.
29. Swist, T., and Collin, P. (2017). Platforms, data and children's rights: introducing a 'networked capability approach'. *New Media Soc.* 19, 671–685.
30. Batty, M. (2018). Digital twins. *Environ. Plann. B: Urban Anal. City Sci.* 45, 817–820.
31. Voigt, I., Inojosa, H., Dillenseger, A., Haase, R., Akgün, K., and Ziemssen, T. (2021). Digital twins for multiple sclerosis. *Front. Immunol.* 12, 1556.
32. Ahmed, H., and Devoto, L. (2020). The potential of a digital twin in surgery. *Surg. Innovat.* 28, 509–510.
33. Chinesta, F., Cueto, E., Abisset-Chavanne, E., Duval, J.L., and Khaldi, F.E. (2018). Virtual, digital and hybrid twins: a new paradigm in data-based engineering and engineered data. *Arch. Comput. Methods Eng.* 27, 105–134.
34. Björnsson, B., Borrebaeck, C., Elander, N., Gasslander, T., Gawel, D., Gustafsson, M., Jörnsten, R., Lee, E.J., Li, X., Lilja, S., et al. (2019). Digital twins to personalize medicine. *Genome Med.* 12, 4.
35. Bagaria, N., Laamarti, F., Badawi, H., Albraikan, A., Martinez, R., and El Saddik, A. (2020). Connected Health in Smart Cities.
36. Laubenbacher, R., Sluka, J.P., and Glazier, J.A. (2021). Using digital twins in viral infection. *Science* 371, 1105.
37. Marx, V. (2021). Biology begins to tangle with quantum computing. *Nat. Methods* 18, 715–719.
38. Braun, M., and Dabrock, P. (2018). Mind the gaps! towards an ethical framework for genome editing. *EMBO Rep.* 19, 197–200.
39. Smith, M., and Miller, S. (2021). The ethical application of biometric facial recognition technology. *AI Soc.* 1–19. <https://doi.org/10.1007/s00146-021-01199-9>.
40. Morozov, E. (2014). *To Save Everything, Click Here: Technology, Solutionism, and the Urge to Fix Problems that Don't Exist* (Public Affairs).
41. Braun, M. (2021b). Represent me – please! towards an ethics of simulated twins in medicine. *J. Med. Ethics* 47, 394–400.
42. Braun, M. (2017). *Zwang und Anerkennung*.
43. ten Have, H. (2016). *Vulnerability: Challenging Bioethics* (Routledge).
44. Braun, M. (2020). Vulnerable life: reflections on the relationship between theological and philosophical ethics. *Am. J. Bioeth.* 20, 21–23.
45. Krutzinna, J. (2021). Simulating (some) individuals in a connected world. *J. Med. Ethics* 47, 403–404.
46. Braun, M. (2021a). Ethics of digital twins: four challenges. *J. Med. Ethics.* <https://doi.org/10.1136/medethics-2021-10767>.

About the authors

Matthias Braun is an Assistant Professor and leader of the young researcher's group "Ethics and Governance of Emerging Technologies" at Friedrich-Alexander University of Erlangen-Nuremberg, Germany. He is also a visiting scholar at Maastricht University, the Netherlands. His research addresses ethical and governance challenges of emerging technologies as well as questions of political ethics as, for example, research on vulnerability, solidarity, and justice.

Jenny Krutzinna holds a PhD in Bioethics and Biomedical Jurisprudence. She is Senior Researcher at the Centre for Research on Discretion and Paternalism, University of Bergen. Previously, she was postdoctoral researcher in Digital Ethics at the Oxford Internet Institute, University of Oxford. Her research focuses on how injustices caused by systemic inequalities and power imbalances affect people's well being. She seeks to enhance the participation of marginalized groups by giving them visibility and a voice and demonstrating how their rights are affected by public and private institutional actors' policies and decisions.