IN FOCUS

The Virtual Child

Richard J. Gilbertson[1,2], Sam Behjati[3,4,5], Anna-Lisa Böttcher[6,7], Marianne E. Bronner[8], Matthew Burridge[9], Henrick Clausing[10], Harry Clifford[11], Tracey Danaher[12], Laura K. Donovan[13], Jarno Drost[14,15], Alexander M.M. Eggermont[14], Chris Emerson[11], Mona G. Flores[16], Petra Hamerlik[17], Nada Jabado[18], Andrew Jones[9], Henrick Kaessmann[19], Claudia L. Kleinman[20,21], Marcel Kool[6,7,14], Lena M. Kutscher[6,22], Gavin Lindberg[23], Emily Linnane[1,2], John C. Marioni[1,3,24], John M. Maris[25,26], Michelle Monje[27], Alexandra Macaskill[28], Steven Niederer[29], Paul A. Northcott[30], Elizabeth Peeters[31], Willemijn Plieger-van Solkema[32], Liane Preußner[33], Anne C. Rios[14], Karsten Rippe[34], Peter Sandford[9], Nikolaos G. Sgourakis[26,35], Adam Shlien[36], Pete Smith[37], Karin Straathof[38,39], Patrick J. Sullivan[40], Mario L. Suvà[41,42], Michael D. Taylor[43], Emma Thompson[44], Roser Vento-Tormo[3], Brandon J. Wainwright[45], Robert J. Wechsler-Reya[46], Frank Westermann[6,47], Shannon Winslade[44], Bissan Al-Lazikani[48], and Stefan M. Pfister[6,7]

Summary: We are building the world's first Virtual Child–a computer model of normal and cancerous human development at the level of each individual cell. The Virtual Child will "develop cancer" that we will subject to unlimited virtual clinical trials that pinpoint, predict, and prioritize potential new treatments, bringing forward the day when no child dies of cancer, giving each one the opportunity to lead a full and healthy life.

WHAT IS THE CHALLENGE?

Cancer remains the leading cause of death-by-disease among children in the developed world. Around 50,000 children are diagnosed with cancer each year in the United States and Europe, and over 7,700 die of the disease (1). This mortality varies with cancer type. Ninety percent of children with acute lymphoblastic leukemia–a once uniformly fatal disease–can now be cured. In stark contrast, solid tumors, particularly those arising in the nervous system, have resisted

[1]CRUK Cambridge Institute, University of Cambridge, Cambridge, United Kingdom. [2]Department of Oncology, University of Cambridge, Cambridge, United Kingdom. [3]Wellcome Sanger Institute, Hinxton, United Kingdom. [4]Cambridge University Hospitals NHS Foundation Trust, Cambridge, United Kingdom. [5]Department of Paediatrics, University of Cambridge, Cambridge, United Kingdom. [6]Hopp Children's Cancer Center (KiTZ), Heidelberg, Germany. [7]Division of Pediatric Neurooncology, German Cancer Research Center (DKFZ) and German Cancer Consortium (DKTK), Heidelberg, Germany. [8]Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, California. [9]Amazon Web Services, London, United Kingdom. [10]SELECT SouthWest, Germany. [11]NVIDIA, Cambridge, United Kingdom. [12]Children's Cancer Foundation, Australia. [13]University College London Great Ormond Street Institute of Child Health, United Kingdom. [14]Princess Máxima Center for Pediatric Oncology, Utrecht, the Netherlands. [15]Oncode Institute, Utrecht, the Netherlands. [16]NVIDIA, Santa Clara, California. [17]University of Manchester, Manchester, United Kingdom. [18]Department of Paediatrics, Research Institute of the McGill University Health Centre, Montreal, Quebec, Canada. [19]Center for Molecular Biology of Heidelberg University (ZMBH), Heidelberg, Germany. [20]Lady Davis Research Institute, Jewish General Hospital, Montreal, Quebec, Canada. [21]Department of Human Genetics, McGill University, Montreal, Quebec, Canada. [22]Developmental Origins of Pediatric Cancer Junior Research Group, German Cancer Research Center (DKFZ), Heidelberg, Germany. [23]The EVAN Foundation, Maryland. [24]European Molecular Biology Laboratory, European Bioinformatics Institute (EMBL-EBI), Wellcome Genome Campus, Cambridge, United Kingdom. [25]Division of Oncology and Center for Childhood Cancer Research, Children's Hospital of Philadelphia, Philadelphia, Pennsylvania. [26]Perelman School of Medicine at the University of Pennsylvania, Philadelphia, Pennsylvania. [27]Department of Neurology and Neurological Sciences, Stanford University, Stanford, California. [28]AstraZeneca, Oncology R&D, Cambridge, United Kingdom. [29]Turing Research and Innovation Cluster in Digital Twins (TRIC: DT), The Alan Turing Institute, London, United Kingdom. [30]Department of Developmental Neurobiology, St Jude Children's Research Hospital, Memphis, Tennessee. [31]b.r.a.i.n.child, SickKids Hospital, Toronto, Ontario, Canada. [32]Dutch Childhood Cancer Organization (VKKN), the Netherlands. [33]BioNTech SE, Mainz, Germany. [34]German Cancer Research Center (DKFZ) Heidelberg, Division of Chromatin Networks, Heidelberg, Germany. [35]Department of Pathology and Lab Medicine, Children's Hospital of Philadelphia, Philadelphia, Pennsylvania. [36]Genetics and Genomics Program, The Hospital for Sick Children, Toronto, Canada. [37]Hula Therapeutics, Philadelphia, Pennsylvania. [38]University College London Cancer Institute, London, United Kingdom. [39]Great Ormond Street Hospital for Children, London, United Kingdom. [40]Team Finn Foundation, Vancouver, British Columbia, Canada. [41]Department of Pathology and Center for Cancer Research, Massachusetts General Hospital and Harvard Medical School, Boston, Massachusetts. [42]Broad Institute of MIT and Harvard, Boston, Massachusetts. [43]Texas Children's Cancer Center, Hematology-Oncology Section and Department of Pediatrics - Hematology/Oncology and Neurosurgery, Baylor College of Medicine, Houston, Texas. [44]Brain Tumour Charity, Fleet, United Kingdom. [45]The University of Queensland Frazer Institute, Translational Research Institute, Woolloongabba, Queensland, Australia. [46]Herbert Irving Comprehensive Cancer Center, Columbia University Medical Center, New York, New York. [47]Division of Neuroblastoma Genomics, German Cancer Research Center (DKFZ), Heidelberg, Germany. [48]Department of Genomic Medicine, The University of Texas MD Anderson Cancer Center, Houston, Texas.

Corresponding Authors: Richard Gilbertson, Cancer Research UK Cambridge Centre, University of Cambridge, CRUK Cambridge Institute, Li Ka Shing Centre, Robinson Way, Cambridge CB2 0RE, United Kingdom. E-mail: Richard.Gilbertson@cruk.cam.ac.uk; Bissan Al-Lazikani, Department of Genomic Medicine, Division of Cancer Medicine, The University of Texas MD Anderson Cancer Center, Houston, TX 77030. E-mail: BOverington@mdanderson.org; and Stefan Pfister, Hopp Children's Cancer Center Heidelberg (KiTZ), Division Head Pediatric Neurooncology, German Cancer Research Center (DKFZ), German Cancer Consortium (DKTK), Deputy Section Head KiTZ Clinical Trial Unit, Heidelberg University Hospital, Im Neuenheimer Feld 280, Heidelberg 69120, Germany. E-mail: s.pfister@kitz-heidelberg.de

Cancer Discov 2024;14:663–8

doi: 10.1158/2159-8290.CD-23-1500

attempts at cure and are the leading cause of childhood cancer–related death (2). In low- and middle-income countries, only 20% of children with cancer survive: a stark inequality that demands serious attention.

The rare and heterogenous nature of childhood tumors has attracted limited intellectual and fiscal investment, restricting the development of childhood cancer–specific therapies. Therefore, childhood cancers have been treated through the extrapolation of adult cancer therapies. But childhood tumors arise in embryonic tissues as early as the first trimester of pregnancy (3) and are therefore intrinsically different from cancers in adults. They are typically not caused by carcinogen exposure, are rarely driven by the mutant kinase targets that have dominated adult cancer drug development, and are likely to possess different therapeutic vulnerabilities.

In contrast to the many hundreds of agents approved for treating adult malignancies, the FDA has approved only 52 drugs for pediatric use, most of which are cytotoxic chemotherapies. Of mechanism-of-action–based therapies, only five have been approved for use in pediatric brain tumors. Furthermore, conventional surgery, radiation, and cytotoxic chemotherapy are used to the limits of tissue tolerance in children. Because these modalities are broadly toxic to developing organ systems–especially the nervous system–it is not surprising that more than two thirds of survivors suffer lifelong, life-changing side effects (4).

The only reasonable conclusion from these experiences is that our approach to developing treatments for children with cancer is deeply flawed. Unless we radically change the way we study and treat these diseases, they will remain a leading cause of childhood morbidity and death for decades to come.

WHAT IS THE BIG IDEA?

We have assembled an international, multidisciplinary team of patient advocates, industry partners (including Amazon Web Services, NVIDIA, BioNTech, and AstraZeneca), drug discoverers, and fundamental and clinician scientists from three continents to transform the way the world understands and treats childhood cancer. This Virtual Child Team is deploying its unique combination of skills in biology, multi-omics, advanced computing, and artificial intelligence/machine learning to create the Virtual Child: the world's first spatially and temporally aware *in silico* model of normal and malignant development (Fig. 1). The model will be used to test the overarching hypothesis that the interconnected, communicating components of the normal and malignant developing nervous system can be modeled computationally and interrogated by deep learning, unmasking targets of highly effective, relatively nontoxic childhood cancer therapies. It will form the first testable multiscale component of the ultimate ambition of a pediatric digital twin. This computer model of normal and cancerous human development at the level of each individual cell will draw on billions of datapoints generated from millions of human and animal cells. We are beginning with the nervous system because neural tumors kill more children than any other cancer. The Virtual Child will provide the world with a shared data and modeling environment that accelerates biological discovery, therapeutic invention, and clinical trial research, leading to the development of kinder curative therapies for children with cancer.

By developing "digital cancers," the Virtual Child can be subject to unlimited numbers of "virtual, *in silico* clinical trials" to pinpoint, predict, and prioritize potential new targets and treatments before they are ever developed preclinically and ultimately given to children. This will include the invention of completely new immunotherapeutic and small-molecule treatments for children with neural tumors. All this information (successes and failures) will be used to continually educate the Virtual Child, further increasing its predictive power. We are inviting the participation of other organ system teams from around the world to add to the Virtual Child, including the sites of other common childhood cancers, e.g., bone and muscle, ultimately transforming the way the world understands and treats all childhood cancer.

MASSES OF UNTAPPED DATA: THE BUILDING BLOCKS OF THE VIRTUAL CHILD

Over the last two decades, Virtual Child Team members have led collaborative networks across >40 countries that have defined the genetic, epigenetic, and transcriptomic landscapes of multiple childhood neural tumors (5–10). These studies have generated the great majority of multi-omic data (e.g., bulk and single-cell RNA sequencing, proteomic, chromatin marks) available for normal and malignant pediatric neural tissues; identified disordered development as the root cause of childhood brain tumors; guided the production of accurate genetically engineered mouse models; underpinned the World Health Organization classification of brain tumors; produced the free online childhood brain tumor methylome diagnostic platform used >135,000 times by health care professionals across the world (www.molecularneuropathology.org); and led to new therapeutic approaches. In addition, this effort is seeking to redress the white racial bias in cancer genome databases through the Molecular Neuropathology Outreach Consortium recruiting tumor material and patient information from Asia (India, Indonesia, Thailand, Pakistan, Jordan, Qatar), Africa (Egypt, South Africa), and South America (Argentina, Brazil, Chile).

However, these massive multi-omic datasets and associated mouse models remain a largely untapped resource. They are distributed across an array of public and private archives in different formats, limiting access by the global research community for analysis. Therefore, as a first step to generating the Virtual Child, we are working closely with both NVIDIA and Amazon Web Services to establish a Virtual Child data lake within a tiered, trusted research environment. This cloud-based resource will be geographically distributed across the different regulatory locations of the Virtual Child Team to minimize data movement across regulatory boundaries and facilitate federated learning.

IN SILICO FRAMEWORK FOR MODELING: MULTISCALE KNOWLEDGE GRAPHS

On the basis of extensive experience, we are adopting graph-based methods to represent the dynamic data and their connections into knowledge graphs (KG; refs. 11–14). Each KG

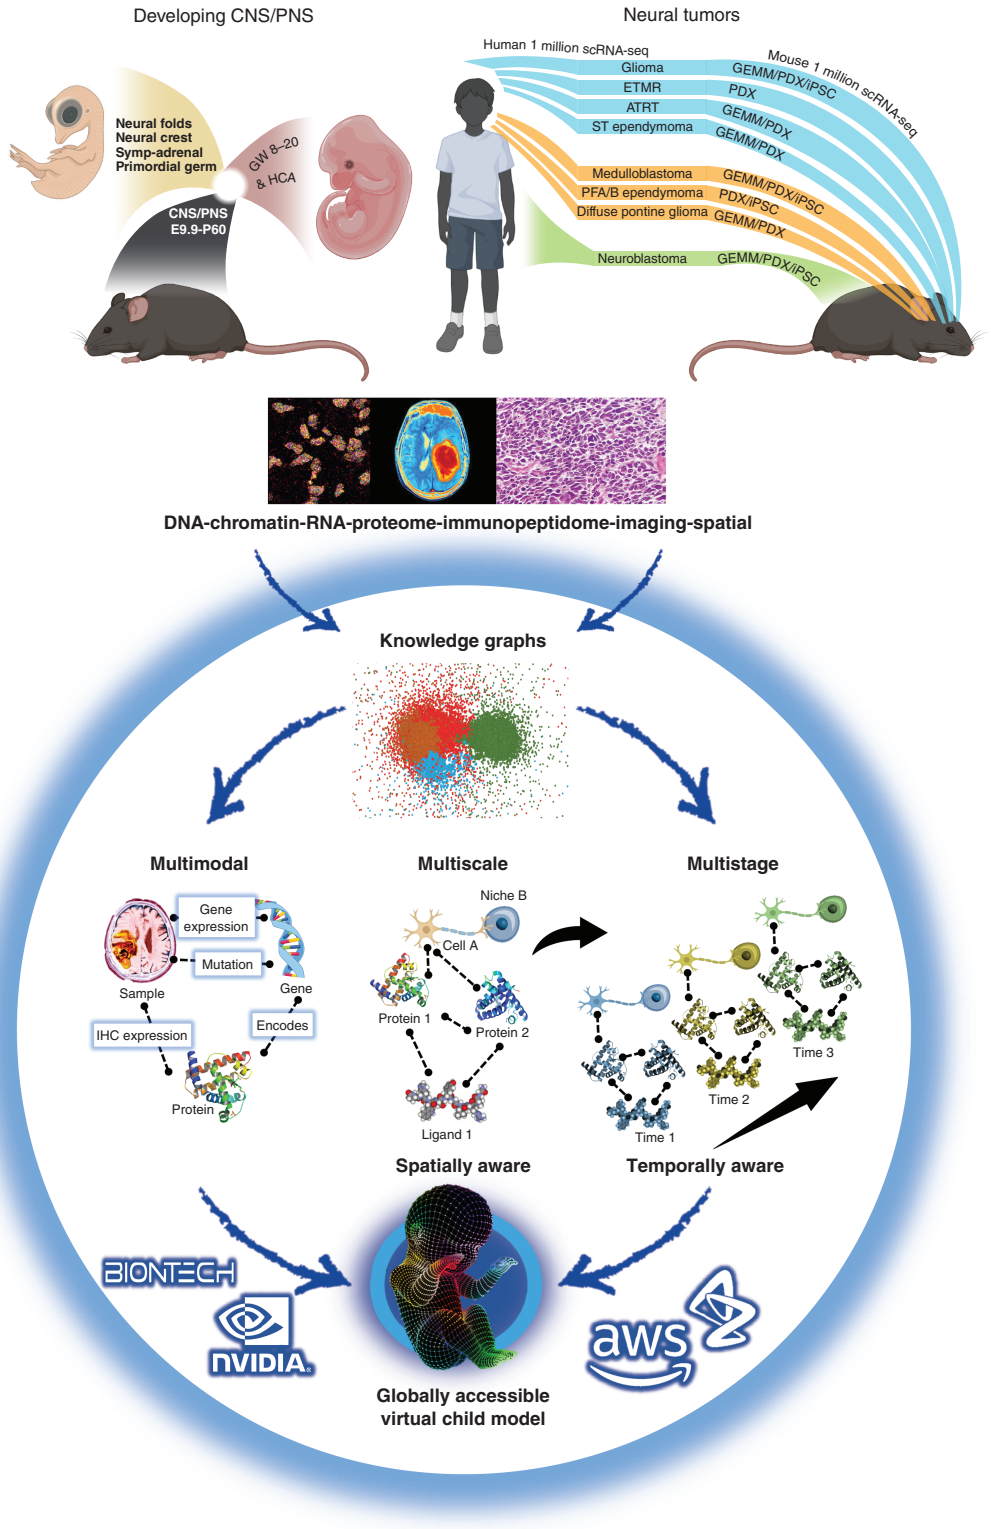

Figure 1. Top, multi-omic, multispecies, multimolecular, and multidimensional data to feed the Virtual Child will be collected from both normal (left) and malignant (right) tissues. Much of these data are available already as detailed in the text. Bottom, these data will be built successively into single mode and then multimode knowledge graphs within the context of a trusted research environment in the cloud, supported by industrial partners Amazon Web Services and NVIDIA. Virtual clinical trials conducted using the model will identify and prioritize new therapeutic targets to be developed into small-molecule and immune therapies with industry partners BioNtech and AstraZeneca. The entire model and targets in development will be made available to the global research community. Figure was created in part using BioRender. GEMM, genetically engineered mouse models; PDX, patient-derived xenografts; scRNA-seq, single-cell RNA sequencing; iPSC, induced pleuripotent stem cell; ETMR, embryonal tumor with multilayered rosettes; ATRT, atypical teratoid rhabdoid tumour; ST, supratentorial.

will represent the components of the interconnecting nervous system (e.g., genes, proteins, cells, cellular niches) and the communications between them. KGs are the most versatile and informative frameworks for this task because they can represent complex interactions between key components and are a useful input into deep learning methods such as neural networks. KGs are also "forgiving" of incomplete data and allow estimation of confidence based on existing data. They also naturally lend themselves to the application of complex deep learning models such as graph neural networks. Finally, KGs can be used to represent temporal changes to a system to enable the modeling of normal development and compare that to aberrant development leading to malignancy.

For example, a KG of single-cell RNA-sequencing data of a specific time point in normal development organizes genes in two-dimensional space according to the expression pattern of each gene relative to all other genes in the transcriptome, across all cells in the tissue. Genes with tightly correlated expression, and therefore potentially related function, cluster closely together within subnetworks, unmasking key biological processes. Additional single mode KGs can be built from proteomic, methylomic, assay for transposase-accessible chromatin with sequencing (ATAC-seq), and even spatial and imaging data. Using published neural networking and novel algorithmic approaches we can take random walks across each KG to understand the context of the constituent components. Using Natural Language Processing (NLP) and other algorithmic approaches, such as Word2Vec–a neural network that learns the context of words (genes) from a large corpus of text (KGs)–we can then learn the context of each gene within different KGs. This provides a rich semantic view of each tissue's transcriptome, pinpointing biological processes that are retained, or lost, following the transition of tissues during normal and malignant development.

MULTIMODEL KNOWLEDGE GRAPHS: KNITTING TOGETHER THE VIRTUAL CHILD

To create a comprehensive multimodal model of the developing child, we are mapping extensive multimodal preclinical and clinical data onto multimode KGs. Together, these represent each stage of the developing nervous system, both normal and tumor prone, as well as each neural tumor type. We have used these approaches successfully to show that multimode KG structures and communication patterns hold intrinsic information that can be analyzed using biology-agnostic graphical mathematics, demonstrating that machine learning of graph communication patterns can predict therapeutic targets of cancer and other human diseases (11); map key biological processes in solid tumors (12); and predict single agent drug response for individual patients (14).

In addition to multi-omic data, we are integrating spatially aware morphologic and molecular data (derived from histologic and imaging analysis) into the Virtual Child. The integration of these multifaceted datasets will be crucial to gain a full understanding of how specific cell lineages in developing neural tissues are perturbed during tumorigenesis. We have already provided proof-of-principle of this approach by integrating immunohistochemical and RNAscope histopathology data together with whole-exome sequencing, total and single-cell RNA-sequencing data, and chromatin immunoprecipitation sequencing into a unified KG of uveal melanoma. The analysis identified patient subtypes and predicted therapeutic vulnerabilities, including those of metastatic uveal melanoma that have been experimentally validated.

Multimodal KGs of each development and disease/normal state across species will be used to evaluate and optimize the most informative mathematical constructs that "pass messages" between nodes within KGs. "Message passing" across the KG will be evaluated to provide clear metrics that capture the magnitude of interaction between connected nodes, the consequence of this impact (e.g., activating, inhibiting, synergistic), and useful inputs for further deep learning methods. The overarching aim will be to identify molecular vulnerabilities and to prioritize these according to relative importance in the model as well as whether they are actionable (targets of known drugs) or tractable (amenable to drug discovery/immune-oncology approaches).

THE CONCEPT OF VIRTUAL CLINICAL TRIALS

In silico perturbation and simulation of multimode models within the Virtual Child will be performed to predict (i) the divergence of cancer from normal development, pinpointing the timing, cellular, and molecular origins of each tumor type; (ii) target molecules with the greatest impact on this and other state transitions, for example, primary to metastatic state; (iii) the tumor subtype and "location/identity" of the target and impact (e.g., medulloblastoma subtype, cancer cell, immune cell, endothelium); (iv) measurable impacts (phenotypes and biomarkers) of target perturbation; (v) human tumor subtypes and model systems that recapitulate target biology and impact with the greatest fidelity; and (vi) information added in multimode relative to the single-mode models. These outputs will be fed forward for experimental testing. Data generated will also be used in iterative learning to refine the Virtual Child model and prioritize/deprioritize targets. We envisage a major output of this work to be the *in silico* reconstruction of the developing nervous system and associated tumors, represented as a "tapestry" of multimodal and multiscale models. This will provide a comprehensive view of the biology of tumor cells and their microenvironment. We envisage that these later insights will be of particular value in predicting therapeutic efficacy, for example, interaction between immune therapies and the microenvironment and the status of the blood–tumor barrier for small-molecule distribution.

Virtual clinical trials will include the capacity to generate virtual patients for each "trial arm." Test-virtual patients will comprise entire multimode Virtual Child models that contain tumors with the specific vulnerability to be targeted; control-virtual patients will comprise entire multimode Virtual Child models that contain tumors that lack the vulnerability to be targeted. We will run perturbation and deep learning experiments on these multimode models to predict the relative value of the virtual treatment in each virtual patient. Attention will be paid not only to the impact on the virtual tumor but also on the entire multimode nervous system, providing understanding of potential "on-target" toxicities in other neural tissues. These virtual trials will initially involve an adversarial framework where the perturbation generates a

set of new KGs that compete to mimic a cancer-kill or normal development KG versus disease KG. Scoring functions will reward perturbations that generate informative KGs.

In addition to training algorithms to identify biological patterns and key druggable molecular vulnerabilities, the Virtual Child will model uncertainty. For each prediction, a level of confidence will also be delivered to inform decisions and also guide experimental requirements to fill the data gaps required to increase confidence.

HOMING IN ON BETTER TREATMENTS OF CHILDHOOD CANCER

Cancer treatments that show promise in preclinical studies often fail in the clinic. The cause of this poor predictive power is multifactorial but results in large part from inadequate fidelity of preclinical models; limited evidence of efficacy and pharmacokinetics (e.g., penetration of the blood–brain/tumor barrier); poor understanding of target inhibition in specific disease contexts; and failure to predict and mitigate treatment resistance. The data emerging from the Virtual Child are designed specifically to overcome these challenges. We will subject these data, including extensive virtual trials, to a target product profile process to ensure that candidates promoted to preclinical study reach minimum standards in predicted impact, tumor specificity, targetability (small molecule, biologic and/or immunotherapy), and availability of models. For example, promoted targets of immunotherapies will be expressed homogeneously at medium to high density in tumors but not normal tissues; present within the surfaceome if targeted by conventional chimeric antigen receptor T cells; associated with model systems that replicate most closely protein expression (IHC/multiplexed flow cytometry of surface markers/mass spectrometry and tumor-subtype distribution); and complemented with biomarker negative controls.

Children with neural tumors are treated with complex combinations of surgery, radiation, and chemotherapy, whereas preclinical studies typically test novel therapies as single agents. These differences impede efforts to implement new treatments into standard of care. Furthermore, as conventional treatments are rarely tested in preclinical models, then predicting the relative therapeutic value of new treatments is challenging. In addition, most drugs are designed to stay outside the blood–brain barrier to minimize potential neurotoxicity, reducing their value for treating brain tumors. The Virtual Child Team will subject potential new targets and associated immuno- and/or small-molecule therapies prioritized by the Virtual Child *in silico* models to extensive *in vitro* and preclinical validation studies. Candidate therapies will be tested both as monotherapies and in combination with conventional surgery, radiation, and/or chemotherapy. These preclinical trials will explore the dosing and scheduling of test therapies including the additional challenge of the blood–brain barrier. They will reveal how and when to deliver these alongside standard of care, potential therapeutic value, and toxicity and guide the further development of biomarkers of treatment response initially identified by the Virtual Child. With this approach we will select and design clinical candidates and optimized trial designs for translation to patients.

AN EVER-LEARNING AND EVERY-GROWING RESOURCE

A key component of the Virtual Child is its ever-learning nature. All multi-omic and phenotypic data generated as part of our interrogation of immune-oncology, epigenetic, and other small-molecule vulnerabilities will be compared with the predicted lineage architecture and/or communication consequences of "targeting" these vulnerabilities in virtual clinical trials. These iterative analyses will be used to refine targeting of vulnerabilities and will be fed back into the KGs and Virtual Child learning algorithms. Thus, over time, the Virtual Child will grow in capability and sophistication and form a key component of a future pediatric digital twin.

The Virtual Child and all associated data will be made available to the worldwide research community as an unrestricted vehicle for research and education. The Virtual Child and its associated advocate–academic–industry community will thereby transform the treatment of childhood cancer, bringing forward the day when no child dies of cancer, giving each one the opportunity to lead a full and healthy life.

Authors' Disclosures

R.J. Gilbertson reports nonfinancial support from AstraZeneca during the conduct of the study; personal fees from AstraZeneca outside the submitted work. J.C. Marioni is an employee of Genentech since September 2022. M. Monje reports holds equity in MapLight Therapeutics and CARGO Therapeutics. L. Preußner reports other support from BioNTech SE outside the submitted work. M.L. Suva reports other support from Immunitas Therapeutics outside the submitted work. F. Westermann reports grants from German Research Foundation (DFG) during the conduct of the study. B. Al-Lazikani reports grants from Cancer Prevention and Research Institute of Texas, Lyda Hill Foundation, Cancer Research UK, Wellcome, and Commonwealth Foundation for Cancer Research during the conduct of the study; other support from Exscientia Plc and Novo Nordisk; personal fees from GSK; and personal fees from Astex pharmaceuticals outside the submitted work. No disclosures were reported by the other authors.

Acknowledgments

R.J. Gilbertson holds a Quest for Cures grant from the Brain Tumour Charity; CRUK Children's Brain Tumour Centre of Excellence and Core funding; and NCI P01CA096832 and U54 CA2431525. N. Jabado holds a Canada Research Chair (CRC) Tier 1 in Paediatric Oncology from the Canadian Institutes of Health Research. R.J. Wechsler-Reya holds grants R35 NS122339 and U01 CA253547.

Published first April 4, 2024.

REFERENCES

1. Kyu HH, Stein CE, Boschi Pinto C, Rakovac I, Weber MW, Dannemann Purnat T, et al. Causes of death among children aged 5–14 years in the WHO European Region: a systematic analysis for the Global Burden of Disease Study 2016. Lancet Child Adolesc Health 2018;2:321–37.
2. Siegel RL, Miller KD, Fuchs HE, Jemal A. Cancer statistics, 2022. CA Cancer J Clin 2022;72:7–33.
3. Behjati S, Gilbertson RJ, Pfister SM. Maturation block in childhood cancer. Cancer Discov 2021;11:542–4.
4. Bhakta N, Liu Q, Ness KK, Baassiri M, Eissa H, Yeo F, et al. The cumulative burden of surviving childhood cancer: an initial report from the St Jude Lifetime Cohort Study (SJLIFE). Lancet 2017;390:2569–82.

5. Johnson RA, Wright KD, Poppleton H, Mohankumar KM, Finkelstein D, Pounds SB, et al. Cross-species genomics matches driver mutations and cell compartments to model ependymoma. Nature 2010;466:632–6.
6. Northcott PA, Jones DT, Kool M, Robinson GW, Gilbertson RJ, Cho YJ, et al. Medulloblastomics: the end of the beginning. Nat Rev Cancer 2012;12:818–34.
7. Hovestadt V, Jones DT, Picelli S, Wang W, Kool M, Northcott PA, et al. Decoding the regulatory landscape of medulloblastoma using DNA methylation sequencing. Nature 2014;510:537–41.
8. Schwartzentruber J, Korshunov A, Liu XY, Jones DT, Pfaff E, Jacob K, et al. Driver mutations in histone H3.3 and chromatin remodelling genes in paediatric glioblastoma. Nature 2012;482:226–31.
9. Capper D, Jones DTW, Sill M, Hovestadt V, Schrimpf D, Sturm D, et al. DNA methylation-based classification of central nervous system tumours. Nature 2018;555:469–74.
10. Parker M, Mohankumar KM, Punchihewa C, Weinlich R, Dalton JD, Li Y, et al. C11orf95–RELA fusions drive oncogenic NF-κB signalling in ependymoma. Nature 2014;506:451–5.
11. Mitsopoulos C, Schierz AC, Workman P, Al-Lazikani B. Distinctive behaviors of druggable proteins in cellular networks. PLoS Comput Biol 2015;11:e1004597.
12. Wedge DC, Gundem G, Mitchell T, Woodcock DJ, Martincorena I, Ghori M, et al. Sequencing of prostate cancers identifies new cancer genes, routes of progression and drug targets. Nat Genet 2018;50:682–92.
13. Mitsopoulos C, Di Micco P, Fernandez EV, Dolciami D, Holt E, Mica IL, et al. canSAR: update to the cancer translational research and drug discovery knowledgebase. Nucleic Acids Res 2021;49:D1074–82.
14. Coker EA, Stewart A, Ozer B, Minchom A, Pickard L, Ruddle R, et al. Individualized prediction of drug response and rational combination therapy in NSCLC Using artificial intelligence–enabled studies of acute phosphoproteomic changes. Mol Cancer Ther 2022;21:1020–9.