

This PDF is available at <http://nap.edu/25620>

SHARE



Frontiers of Engineering: Reports on Leading-Edge Engineering from the 2019 Symposium (2020)

DETAILS

140 pages | 6 x 9 | PAPERBACK

ISBN 978-0-309-49981-1 | DOI 10.17226/25620

CONTRIBUTORS

National Academy of Engineering

GET THIS BOOK

FIND RELATED TITLES

SUGGESTED CITATION

National Academy of Engineering 2020. *Frontiers of Engineering: Reports on Leading-Edge Engineering from the 2019 Symposium*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/25620>.

Visit the National Academies Press at NAP.edu and login or register to get:

- Access to free PDF downloads of thousands of scientific reports
- 10% off the price of print titles
- Email or social media notifications of new titles related to your interests
- Special offers and discounts



Distribution, posting, or copying of this PDF is strictly prohibited without written permission of the National Academies Press. ([Request Permission](#)) Unless otherwise indicated, all materials in this PDF are copyrighted by the National Academy of Sciences.

Copyright © National Academy of Sciences. All rights reserved.

FRONTIERS OF **ENGINEERING**

Reports on Leading-Edge Engineering from the 2019 Symposium

NATIONAL ACADEMY OF ENGINEERING

THE NATIONAL ACADEMIES PRESS

Washington, DC

www.nap.edu

THE NATIONAL ACADEMIES PRESS • 500 Fifth Street, NW • Washington, DC 20001

NOTICE: Publication of signed work signifies that it is judged a competent and useful contribution worthy of public consideration, but it does not imply endorsement of conclusions or recommendations by the National Academy of Engineering (NAE). The interpretations and conclusions in such publications are those of the authors and do not purport to represent the views of the council, officers, or staff of the NAE.

Funding for the activity that led to this publication was provided by The Grainger Foundation, The Boeing Company, National Science Foundation, Defense Advanced Research Projects Agency, Department of Defense ASD(R&E) Research Directorate—Laboratories Office, Air Force Office of Scientific Research, Microsoft Research, Amazon, Cummins Inc., and individual donors. This material is based on work supported by the Air Force Office of Scientific Research under award number FA9550-19-1-0333. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force. This material is based on work supported by the National Science Foundation under Grant No. EFMA-903556. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

International Standard Book Number-13: 978-0-309-49981-1

International Standard Book Number-10: 0-309-49981-X

Digital Object Identifier: <https://doi.org/10.17226/25620>

Additional copies of this report are available from The National Academies Press, 500 Fifth Street, NW, Keck 360, Washington, DC 20001; (800) 624-6242 or (202) 334-3313; <http://www.nap.edu>.

Copyright 2020 by the National Academy of Sciences. All rights reserved.

Printed in the United States of America

Suggestion citation: National Academy of Engineering. 2020. *Frontiers of Engineering: Reports on Leading-Edge Engineering from the 2019 Symposium*. Washington, DC: The National Academies Press. doi: <https://doi.org/10.17226/25620>.

The National Academies of
SCIENCES • ENGINEERING • MEDICINE

The **National Academy of Sciences** was established in 1863 by an Act of Congress, signed by President Lincoln, as a private, nongovernmental institution to advise the nation on issues related to science and technology. Members are elected by their peers for outstanding contributions to research. Dr. Marcia McNutt is president.

The **National Academy of Engineering** was established in 1964 under the charter of the National Academy of Sciences to bring the practices of engineering to advising the nation. Members are elected by their peers for extraordinary contributions to engineering. Dr. John L. Anderson is president.

The **National Academy of Medicine** (formerly the Institute of Medicine) was established in 1970 under the charter of the National Academy of Sciences to advise the nation on medical and health issues. Members are elected by their peers for distinguished contributions to medicine and health. Dr. Victor J. Dzau is president.

The three Academies work together as the **National Academies of Sciences, Engineering, and Medicine** to provide independent, objective analysis and advice to the nation and conduct other activities to solve complex problems and inform public policy decisions. The National Academies also encourage education and research, recognize outstanding contributions to knowledge, and increase public understanding in matters of science, engineering, and medicine.

Learn more about the National Academies of Sciences, Engineering, and Medicine at www.nationalacademies.org.

SYMPOSIUM ORGANIZING COMMITTEE

JENNIFER WEST (*Chair*), Fitzpatrick Family University Professor, Department of Biomedical Engineering, Duke University

LI CHANG, Technical Fellow, Boeing Commercial Airplanes, The Boeing Company

TARIK DICKENS, Assistant Professor, Department of Industrial and Manufacturing Engineering, Florida A&M University and Florida State University

CHARLES GERSBACH, Associate Professor, Department of Biomedical Engineering, Duke University

CHRISTOFFER HECKMAN, Assistant Professor, Department of Computer Science, University of Colorado Boulder

PETR NOVOTNY, Research Staff Member, Enterprise Solutions, Thomas J. Watson Research Center, IBM

HAE-JONG SEO, Autonomous Driving R&D, NVIDIA

ELAINE SHI, Associate Professor, Department of Computer Science, Cornell University

RENEE WEGRZYN, Program Manager, Biological Technologies Office, DARPA

Staff

JANET HUNZIKER, Director, Frontiers of Engineering Program

SHERRI HUNTER, Meetings Coordinator

Preface

This volume presents papers on the topics covered at the National Academy of Engineering's 2019 US Frontiers of Engineering Symposium. Every year the symposium brings together 100 highly accomplished early-career engineering leaders to share their cutting-edge research and innovations in selected areas. The 2019 symposium was hosted by Boeing in North Charleston, South Carolina, September 25–27. The intent of this book is to convey the excitement of this unique meeting and to highlight innovative developments in engineering research and technical work.

GOALS OF THE FRONTIERS OF ENGINEERING PROGRAM

The practice of engineering is continually changing. Engineers must be able not only to thrive in an environment of rapid technological change and globalization but also to work on interdisciplinary teams. Today's research is being done at the intersections of engineering disciplines, and successful researchers and practitioners must be aware of developments and challenges in areas that may not be familiar to them.

At the annual 2½-day US Frontiers of Engineering Symposium, 100 of this country's best and brightest early-career engineers—from academia, industry, and government and a variety of engineering disciplines—learn from their peers about pioneering work in different areas of engineering. The number of participants is limited to 100 to maximize opportunities for interactions and exchanges among the attendees, who are chosen through a competitive nomination and selection process. The symposium is designed to foster contacts and learning among promising individuals who would not meet in the usual round of professional meetings.

This networking may lead to collaborative work, facilitate the transfer of new techniques and approaches, and produce insights and applications that bolster US innovative capacity.

The four topics and the speakers for each year's meeting are selected by an organizing committee of engineers in the same early-career cohort as the participants. Speakers describe the challenges they face and communicate the excitement of their work to a technically sophisticated but nonspecialist audience. They provide a brief overview of their field of inquiry; define the frontiers of that field; describe experiments, prototypes, and design studies (completed or in progress) as well as new tools and methods, limitations and controversies; and assess the long-term significance of their work.

THE 2019 SYMPOSIUM

The topics covered at the 2019 symposium were (i) advanced manufacturing, (ii) engineering the genome, (iii) self-driving cars: technology and ethics, and (iv) blockchain technology.

Industry 4.0 refers to the fourth industrial revolution in which the industrial and information revolutions have merged, creating systems that are smart, agile, resilient, and customizable. The first speaker in the session, *Advanced Manufacturing in the Age of Digital Transformation*, focused on applications of data analytics, autonomy, model-based engineering, and machine learning to manufacturing. This was followed by a presentation on computational modeling used in the sequencing of digital manufacturing in the domain of additive manufacturing. The next speaker discussed new directions for legged robots and their future applications in the manufacturing sector. The session concluded with a talk on the impact of the digital twin in boosting efficiency, slashing costs, and revealing problems before production.

The next session, *Engineering the Genome*, described how the growth of genome engineering tools have the potential to alter any DNA or RNA sequence, leading to an almost limitless range of applications in treating human genetic diseases, developing industrial biotech products, improving crop and livestock productivity, and addressing conservation and invasive species challenges. Talks detailed the application areas available through CRISPR-Cas9, the impact of genome engineering on ecosystems such as mosquito transmission of diseases, industrial scale-up for manufacturing molecules for various applications, and the need for standards and data sharing in this rapidly evolving field.

Although self-driving cars are now on the roads, the ramifications of this trend are complex due to the potential effects on infrastructure, the economy, and society and the ways transportation factors into daily life. The first presentation in the session titled *Self-Driving Cars: Technology and Ethics* provided an overview of the challenges and opportunities provided by self-driving vehicles. The next speaker described how self-driving cars are being developed at scale.

This was followed by a discussion on the philosophy and ethics surrounding the aggressive development of self-driving cars and the technologies that support them. The session closed with a talk about humans' interactions with autonomous and intelligent systems, including research on design and learning algorithms that influence humans' actions for better safety and coordination.

Blockchain—the underlying technology on top of which Bitcoin and other applications are implemented—was the topic of the final session. The presentations introduced the history and key concepts of blockchain and provided an overview of the major platforms and applications including Bitcoin, Ethereum, and Hyperledger; discussed the domain of private and permissioned blockchain platforms that are building blocks of networks among groups of enterprises; and delved into the economic and social research that leverages blockchain technologies.

In addition to the plenary sessions, the attendees had many opportunities for informal interaction. On the first afternoon of the meeting, break-out sessions provided an opportunity for attendees to share their research and technical work so that they could get to know more about each other relatively early in the program. On the second afternoon, Boeing arranged a tour of the plant where Boeing's 787 Dreamliners are manufactured.

Every year a distinguished engineer addresses the participants at dinner on the first evening of the symposium. Ms. Joan Robinson-Berry, vice president and chief engineer for Boeing Global Services, spoke about the breadth of research and engineering—"from freezer to flight"—at Boeing South Carolina. She also issued a call to action for increased diversity in the engineering workforce and encouraged the attendees to be drivers of change by bringing people from outside their communities to the table.

The NAE is deeply grateful to the following for their support of the 2019 US Frontiers of Engineering Symposium:

- Boeing
- The Grainger Foundation
- Defense Advanced Research Projects Agency
- Air Force Office of Scientific Research (This material is based on work supported by the Air Force Office of Scientific Research under award number FA9550-19-1-0333. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.)
- Department of Defense ASD(R&E) Research Directorate–Laboratories Office
- National Science Foundation (This material is based on work supported by the National Science Foundation [NSF] under grant EFMA-1903556. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.)

- Microsoft Research
- Amazon
- Cummins Inc.
- Individual contributors

We also thank the members of the Symposium Organizing Committee (p. iv), chaired by Dr. Jennifer West, for planning and organizing the event.

Contents

ADVANCED MANUFACTURING IN THE AGE OF DIGITAL TRANSFORMATION

Introduction	3
<i>Li Chang and Tarik Dickens</i>	
Computational Materials for the Design and Qualification of Additively Manufactured Components	5
<i>Christopher G. Lang</i>	
Robots That Walk: What the Challenge of Locomotion Says About Next-Generation Manufacturing	11
<i>Christian Hubicki</i>	
The Digital Twin Concept	17
<i>Pamela A. Kobryn</i>	

ENGINEERING THE GENOME

Introduction	25
<i>Charles Gersbach and Renee Wegrzyn</i>	
Genome Editing with Precision and Accuracy	27
<i>Krishanu Saha</i>	

Using CRISPR to Combat Human Disease Vectors <i>Omar S. Akbari</i>	35
Microbes and Manufacturing: Moore's Law Meets Biology <i>Patrick Boyle</i>	45
Empowering Genome Editing Through Standards <i>Samantha Maragh</i>	55

SELF-DRIVING CARS: TECHNOLOGY AND ETHICS

Introduction <i>Christoffer Heckman and Hae-Jong Seo</i>	65
Perceptions of Low-Cost Autonomous Driving <i>Tae Eun Choe, Xiaoshu Liu, Guang Chen, Weide Zhang, Yuliang Guo, and Ka Wai Tsoi</i>	67
Why Everyone Has It Wrong About the Ethics of Autonomous Vehicles <i>John Basl and Jeff Behrends</i>	75
Influencing Interactions Between Human Drivers and Autonomous Vehicles <i>Dorsa Sadigh</i>	83

BLOCKCHAIN TECHNOLOGY

Introduction <i>Petr Novotny and Elaine Shi</i>	93
Blockchain Beyond Cryptocurrency: An Overview <i>Hong Wan</i>	95
Cryptocurrencies as Marketplaces <i>Jacob Leshno</i>	103

<i>CONTENTS</i>	<i>xi</i>
-----------------	-----------

APPENDIXES

Contributors	111
Participants	115
Program	125

ADVANCED MANUFACTURING IN THE AGE OF DIGITAL TRANSFORMATION

Advanced Manufacturing in the Age of Digital Transformation

LI CHANG
Boeing

TARIK DICKENS
Florida A&M University-Florida State University College of Engineering

The synergy of past manufacturing-enabling revolutions has made pathways for highly efficient systems of the future. The industrial and information revolutions have merged into a fourth industrial revolution known as the digital age, in which information is ubiquitous. The culmination of these two streams is called “Industry 4.0.” The expanse of this fourth industrial revolution is signified in the National Academy of Engineering’s Grand Challenges, through which wide applicability and engineering solutions can transform the societal landscape. This is a new frontier to imagine the possibilities of 22nd century manufacturing, with smart, more agile, resilient, and customizable systems to meet the needs of the world’s growing populations.

The first speaker, Gabriel Burnett (Boeing), introduced the future of Boeing’s production system with a focus on applications of data analytics, autonomy, model-based engineering, and machine learning.¹ Boeing is moving away from a document-centric system to become a model-based enterprise. Hardware advances are enabling data collection, and computing advances allow the data to be leveraged across the value stream to enhance business value. The next speaker, Christopher Lang (National Aeronautics and Space Administration Langley Research Center), discussed the use of computational modeling in the sequencing of digital manufacturing, focusing on metal additive manufacturing. Christian Hubicki (Florida State University and Florida A&M University) described novel directions of bioinspired robots for the future of digital manufacturing. Finally, Pamela Kobryn (Air Force Research Laboratory) presented the digital twin

¹ Paper not included in this volume.

concept as applied to aircraft to boost efficiency, reduce costs, increase agility, and better manage maintenance schedules according to both past and projected operational needs.

Computational Materials for the Design and Qualification of Additively Manufactured Components

CHRISTAPHER G. LANG
NASA Langley Research Center

NASA is developing next-generation computational materials capabilities to support the qualification of additively manufactured metallic structural components for aerospace applications. The quality of these parts directly depends on a wide range of process parameters, including build conditions and feedstock properties. Computational materials research aims to develop a fundamental understanding of the dependence of the part properties and performance on the process parameters and to apply that understanding to efficient qualification practices.

Integrated multiscale modeling methods allow prediction of the process–structure–property relationships, including the effect of defects. This paper primarily focuses on the powder bed fusion process and its application to aerospace flight systems, with discussion of in situ monitoring, process-to-microstructure linkages including residual stress, and microstructure-to-performance linkages. Computational materials research for additive manufacturing (AM) processes will enable efficient and accurate design, manufacture, and certification of future aerospace flight systems.

INTRODUCTION

Although AM technology has recently experienced considerable growth and publicity for its potential to significantly transform the manufacturing industry, its promise is limited in application because of a lack of confidence in part quality. Improvements in material properties, consistency, and process control are necessary for AM to realize the advertised potential of enhanced performance, reduced cost, and increased manufacturing speed; for example, the application of AM to fracture-critical flight components requires extensive qualification efforts.

AM encompasses a variety of materials (e.g., metals, polymers, and ceramics) and processes (e.g., powder bed, blown powder, wire fed, laser, and electron beam). Part quality and consistency depend on numerous process-specific parameters that are selected or adjusted for each component.

LASER POWDER BED FUSION

I focus on the laser powder bed fusion (LPBF) process for metallic AM, although many of the approaches are applicable to a wide range of materials and manufacturing processes. The LPBF parameter space consists of laser power, scan speed, laser spot size, scanning strategy, feedstock, part geometry, and machine conditions. The selection of process parameters determines the resulting microstructure and component properties.

Various libraries of process parameters for a given machine and material have been determined through physical testing by AM suppliers or individual laboratories, with additional testing required for each new part geometry or powder supply. An integrated computational materials engineering (ICME) approach reduces the amount of physical testing and informs design engineers about detrimental performance expected for specific process parameters (Turner et al. 2015).

NASA is developing AM rocket engine components for human spaceflight. To address the immediate need for a consistent framework specific to the production and evaluation of LPBF processes, standards have been released by the Marshall Space Flight Center (MSFC 2017a,b) for materials, process control, personnel training, inspection, and acceptance requirements. Concurrently, an ICME approach to the design and qualification of aerospace AM materials and their components is being developed at NASA and provides a path toward rapid manufacturing and qualification.

Improved control and understanding of the AM process offer improved consistency and more complex design such as multiple alloys and functionally graded material components. When combined with in situ process monitoring, computational modeling enables the development and integration of manufacturing process capabilities and constraints as well as qualification considerations such as inspection requirements in the component design.

COMPUTATIONAL MODELING OF THE AM PROCESS

Process modeling is used to develop an understanding of the relationship between the process parameters, feedstock, microstructural and porosity evolutions, and resulting mechanical properties by solving the governing equations for the physics of the process. Determination of the temperature history, deformations due to residual stress, microstructure evolution, and porosity are among the goals of current process simulation efforts.

Physics

Modeling of the AM process requires a multiscale approach to accurately account for the physics at multiple length scales from microstructure to component. An accurate temperature history and melt pool geometry are necessary to understand the microstructure, defect formation, and residual stress formation. The temperature history is predicted by numerical models at different levels of fidelity. Various physics—melting, evaporation, fluid flow, recoil pressure, powder packing density, and surface tension—are incorporated to improve the model accuracy. To accommodate accuracy and computational resource requirements, thermal models are generally restricted to a low number of scan tracks and powder layers.

Simulation of residual stress formation requires a scale-up to efficiently account for the numerous layers in an AM build. A promising approach for predicting residual stress is the modified inherent strain method, which computes the strain at the scan track scale and imposes the strains in a layer-by-layer fashion to a part scale mechanical analysis (Liang et al. 2018). Phase-field and kinetic Monte Carlo models are used to simulate grain structures dependent on feedstock and temperature history.

Porosity

Two sources of porosity during the LPBF process are lack of fusion and keyholing. The melt pool transitions from conduction mode to keyhole mode for increased laser power and reduced scan speed. Keyhole mode occurs when a vapor cavity forms with a high aspect ratio of depth to width as compared to conduction mode (Trapp et al. 2017). In contrast, lack of fusion porosity occurs when insufficient power and overlap of successive melt pools are applied to fully melt the powder. A balance for avoiding lack of fusion and keyhole porosity is determined by the selected process parameters (Tang et al. 2017).

Porosity cannot be completely avoided, and its impact on part performance is application dependent. Micromechanical simulations quantitatively characterize the influence of porosity and other heterogeneities in the microstructure on the mechanical behavior of parts produced by LPBF. Porosity is embedded in process-specific microstructure models, and the heterogeneous strain localization in the vicinity of the porosity is solved as a function of the pore shape, size, density, and proximity to the free surface.

IN SITU PROCESS DATA

For the design and qualification of AM components, experimental data are required to capture critical events and behavior during the manufacturing process. To that end,

- Powder bed systems are being equipped with sensors and measuring devices to record data during the manufacturing process.
- System monitoring provides critical data necessary for understanding process events, performing feedback control, diagnosing machine operation, and validating computational models.
- Key process measurements include temperature history, melt pool dimensions, and defect formation.

Collection of in situ data provides a component build history that can be used to identify critical events during the process that may affect part quality.

Dynamic X-ray radiography (DXR) at the Argonne National Laboratory Advanced Photon Source provides high-speed cross-section videos of the LPBF process (Zhao et al. 2017). The real-time imaging yields data relative to the laser position, including melt pool dimensions, keyhole behavior, solidification rate, and porosity formation. DXR data help characterize the melt pool and solidification behavior for various feedstock compositions and baseplate material as well as varying laser parameters.

SUMMARY

Computational modeling supports the qualification efforts necessary to realize the full potential of AM for designing and manufacturing aerospace components. A large design space exists for AM, and an ICME approach to process and component design will support qualification efforts through improved process understanding and control for application- and material-specific needs. Simulation tools that assist in choosing parameters for process control and designing AM-specific components will lead to microstructures that help attain and even exceed design specifications.

Micromechanical simulations characterize part performance for process-specific microstructures including the effect of defects. Integrated computational modeling and in situ process monitoring efforts provide a path toward accelerated design and qualification of aerospace components.

REFERENCES

- Liang X, Cheng L, Chen Q, Yang Q, To AC. 2018. A modified method for estimating inherent strains from detailed process simulation for fast residual distortion prediction of single-walled structures fabricated by directed energy deposition. *Additive Manufacturing* 23:471–86.
- MSFC [Marshall Space Flight Center]. 2017a. Standard for additively manufactured spaceflight hardware by laser powder bed fusion in metals. Technical Standard MSFC-STD-3716. Huntsville AL.
- MSFC. 2017b. Specification for control and qualification of laser powder bed fusion metallurgical processes. Technical Standard MSFC-SPEC-3717. Huntsville AL.
- Tang M, Pistorius PC, Beuth JL. 2017. Prediction of lack-of-fusion porosity for powder bed fusion. *Additive Manufacturing* 14:39–48.

- Trapp J, Rubenchik AM, Guss G, Matthews MJ. 2017. In situ absorptivity measurements of metallic powders during laser powder-bed fusion additive manufacturing. *Applied Materials Today* 9:341–49.
- Turner JA, Babu SS, Blue C. 2015. Advanced simulation for additive manufacturing: Meeting challenges through collaboration. ORNL/TM-2015/324. Oak Ridge TN: Oak Ridge National Laboratory.
- Zhao C, Fezzaa K, Cunningham RW, Wen H, De Carlo F, Chen L, Rollett AD, Sun T. 2017. Real-time monitoring of laser powder bed fusion process using high-speed X-ray imaging and diffraction. *Scientific Reports* 7(1):3602.

Robots That Walk: What the Challenge of Locomotion Says About Next-Generation Manufacturing

CHRISTIAN HUBICKI

*Florida A&M University–Florida State University (FAMU-FSU)
College of Engineering*

Why study walking and running robots, and bipedal robots in particular? Because robots can go where people can't or shouldn't go because it isn't safe, such as malfunctioning nuclear power plants, or the staircases and corridors of burning buildings—places designed to be navigated by bipedal humans.

In addition to disaster response and exploration applications, the study of bipedal robotics informs the design and control of assistive devices. Robotic prosthetic legs and exoskeletons can enable walking and running for patients who have lost the ability to do so, or perhaps even superhuman performance (e.g., lifting or conveying immensely heavy loads).

More broadly, the challenge of engineering bipedal robots yields general lessons across autonomy-related fields including manufacturing. While manufacturing has excelled in extreme repeatability of its processes in controlled environments (famously measuring quality in “sigmas”), legged robotics grapple with uncontrolled environments. If manufacturing is to push its capabilities or venture into remote and uncontrolled locations, it may be beneficial to learn from the field of legged robotics.

This paper summarizes four lessons that were essential to recent advances in walking and running robots and illustrates how they can be applied to manufacturing.

- Bipedal roboticists have designed robots with underactuation and compliance for improved agility and efficiency, and automated manufacturing facilities can similarly lower capital and energy costs.
- Walking and running controllers are being designed for robustness to unknown terrain, and manufacturing can use similar robust control to operate on parts of unknown shape and softness.

- Locomoting robots that rely on self-stable and emergent behaviors provide an example of decreased reliance on heavy computation and ways to find surprising effective behaviors not programmed by their control engineers.
- And finally, the need to change locomotion behaviors on the fly—for example, to change speeds or tasks—has driven the development of task-flexible control algorithms that can make robots more versatile (e.g., for tasks like carrying packages).

UNDERACTUATION AND COMPLIANCE

Legged robots have made enormous strides in terms of agility and stability—Boston Dynamics’ Atlas (figure 1) can do backflips—but energy economy is a longstanding challenge in legged robotics. Humanoid robots typically require an order of magnitude more energy to walk than an equivalently sized person. This drastically limits the range that robots can travel on a limited energy supply, pushing legged roboticists to find ways to do more with less energy.

One efficiency-driven approach, underactuation, involves building bipedal robots with fewer motors, forcing the controller to do the same task using less power. Underactuation can also reduce robot weight as well as stiffness due to highly geared drive transmissions. These combined improvements often lead to increased efficiency: A simple underactuated walker called the Cornell Ranger was able to walk 40 miles on a single battery charge.¹

Another approach to increased efficiency comes in the form of compliance, or elasticity in the robot. Robots traditionally are built with highly rigid bodies, in part to make control algorithms simpler. Humans, however, have elastic tendons in their legs that can store and return energy otherwise lost to heat while walking. Following this example, robots like DURUS (figure 2) have spring-legged feet that reduce the energy cost of locomotion by 70 percent over previous humanoid robots (Hereid et al. 2018).

For the field of manufacturing, assembly-line manipulators may save on energy costs by omitting actuators and including flexible linkages.

ROBUSTNESS TO UNKNOWN TERRAIN


The terrain of the outside world is messy. Not only can it be uneven or rocky, it can be soft like soil, sand, or snow. This means that vision alone cannot always reveal all the necessary properties of oncoming terrain to inform control. Real-world environments call for robots that are robust to unknown terrain. Figure 1

¹ Reported by PI Andy Ruina (http://ruina.tam.cornell.edu/research/topics/locomotion_and_robotics/ranger/Ranger2011/).



Robustness to Environments

FIGURE 1 The Atlas humanoid from Boston Dynamics (www.bostondynamics.com) demonstrates remarkable robustness to challenging terrain.



Physical Compliance and Passive Dynamics

FIGURE 2 The DURUS humanoid from SRI International is an example of built-in flexibility (compliance).

shows the Atlas humanoid from Boston Dynamics walking over a snow bank, an effective example of terrain robustness.

This need for terrain robustness led researchers to develop force control techniques for locomotion: if a user controls a robot leg to produce a specified *force*, it has a vastly different behavior in response to disturbances than if its *position* is controlled. If a force-controlled leg steps in a soft patch of earth, the force controller automatically pushes the leg harder into the ground to hold up its weight, instead of stumbling from the unexpected terrain. This allows legged robots to travel on all kinds of terrain without falling, including the Cassie biped (Agility Robotics) and the MIT Cheetah, both of which can walk up slopes and stairs that they can't even "see."

In the context of manufacturing, force control is a useful mode for sensitive manipulation using robotic arms. A force-controlled end effector can grip and move an object of unknown geometry or softness, thereby decreasing sensitivity to unknowns in manufacturing.

SELF-STABLE AND EMERGENT CONTROL BEHAVIORS

Roboticians have given a lot of thought to a key question about effective control: What quantities need to be controlled—the position of a robot's legs, the orientation of its torso, the forces at its feet? All of these options have been useful in different applications of the field of legged robotics. The right choice of control target can lead to self-stable and emergent behaviors that are far more capable than their programming was predicted to accomplish.


The bipedal robot ATRIAS (figure 3) was programmed with a self-stabilizing walking controller that systematically cycles its feet and controls its forward speed without needing elaborate computation to maintain balance. Further, when commanded to speed up using this relatively simple controller, the robot did something unexpected—it started running without being explicitly commanded to do so (Hubicki et al. 2018). Imagine how this emergent behavior might manifest in an automated factory: In a task where one robotic arm must transfer a part to another robotic arm, if commanded to transfer faster, it might toss the part to the other. With emergent control behavior, a factory can be inherently more clever and effective than initially imagined by its engineers.

TASK-FLEXIBLE CONTROL ALGORITHMS

The locomotive robotics field has pushed for not only more stable, faster, and more efficient locomotion but also a variety of walking and running behaviors. How does a robot jump over an obstacle, run around a corner, or walk with a fragile object? Preprogramming each task with its own human-derived controller quickly becomes impractical. Consequently, achieving a task-flexible control framework has become a critical push for practical viability.


The push for task flexibility has led to a proliferation of real-time optimization methods to generate stable controllers on the fly for a given task. One method is model-predictive control (MPC), where easy-to-compute optimizations are solved extremely quickly and the solutions form a plan for the controller to complete its task in a specified timeframe. The key benefit of MPC is the ability to appropriately react to disturbances, which are myriad out in the field. If an obstacle appears or the goal suddenly changes, fast optimization enables responsive replanning.

This task flexibility made MPC and similar optimization-based methods popular approaches for the vaunted DARPA Robotics Challenge in 2015. Teams at this challenge were tasked with building a robot that would respond to a simulated industrial disaster. They needed to control a robot to drive a car to the site, exit the vehicle, open a building door, walk over rubble, climb a staircase, shut off a valve, and drill a hole in a wall—all with minimal human supervision on site. The teams succeeded in formulating optimizations that could handle the complexity of robots with dozens of degrees of freedom, solved hundreds of times per second.



Self-stable Control and Emergent Behaviors

FIGURE 3 The ATRIAS biped demonstrates that stable dynamic behaviors (like running) can emerge that are not explicitly preprogrammed.



Task-flexible Control Algorithms

FIGURE 4 The Digit robot from Agility Robotics (www.agilityrobotics.com) is capable of a variety of dynamic maneuvers including expeditiously conveying packages.

The emergence of these planning algorithms enabled the next generation of versatile legged robots, including prototypes for package delivery with Agility Robotics' Digit robot (figure 4). Future manufacturing may need such real-time task flexibility as well: If a robot performing one specialized task malfunctions, another robot not designed explicitly for the role could be adapted to replace it.

CONCLUSION

Legged robots have had to make a number of advances in order to move out in the real world. Many of these approaches have straightforward extensions to the needs of improved manufacturing. At the same time, legged robots would benefit from the reliability of modern manufacturing processes. And there are lessons that human legs can take from factory robots, too. Ideally, manufacturing assembly lines of the future will be as robust and adaptable as human walking,

and conversely, the reliability of walking robots will be measured in “sigmas,” like manufactured products.

REFERENCES

- Hereid A, Hubicki CM, Cousineau EA, Ames AD. 2018. Dynamic humanoid locomotion: A scalable formulation for HZD gait optimization. *IEEE Transactions on Robotics* 34(2):370–87.
- Hubicki C, Abate A, Clary P, Rezazadeh S, Jones M, Peekema A, Van Why J, Domres R, Wu A, Martin W, and 2 others. 2018. Walking and running with passive compliance: Lessons from engineering—A live demonstration of the ATRIAS biped. *IEEE Robotics & Automation* 25(3):23–39.

The Digital Twin Concept

PAMELA A. KOBRYN

Air Force Research Laboratory at Wright-Patterson Air Force Base

The digital twin concept involves simulating the future performance of a specific product or system based on current knowledge about the system and how it is operated. Key aspects of the concept include near- and long-term performance predictions individualized to both the particular product/system (e.g., by serial number) and its use; delivery of results in an intuitive, interactive, and affordable manner in a timeframe suitable for use; and timely and automated updating of results.

While the original concept focused on the health management of engineered systems with stringent reliability and safety requirements (e.g., airplanes), the scope of digital twin applications is rapidly expanding across the entire product/system lifecycle. Advanced computing, information system, network, and device technologies are being joined with advanced analytical methods to unlock new digital twin applications that bring value to enterprises, communities, and individuals across a broad spectrum of uses—in transportation and logistics, mining and construction, manufacturing and production, power generation and distribution, communication and computing networks, and medicine.

MOTIVATING FACTORS AND SIGNIFICANCE

An early motivation for the digital twin concept was the structural health management of military aircraft, which led to investment in Airframe Digital Twin (ADT) technologies by the Air Force Research Laboratory (AFRL) beginning in about 2009 (Tuegel et al. 2011). At the time, the US Air Force's desire to reduce the impact of maintenance on aircraft availability and operating costs inspired AFRL engineers to devise new concepts for predicting structural maintenance

needs. These engineers sought to develop methods to increase the fidelity and timeliness of the analyses used to decide when to perform such maintenance (Tuegel and Babish 2014).

Newer applications have a similar motivation: delivering and sustaining the predictable, safe, reliable, and affordable operational capability of engineered products and systems to achieve the outcome desired or required by the end user. Recent advances—in high-performance computing capability; modeling, simulation, and analysis methods; data analytics and information technology; metrology and sensor technology; the internet of things/industrial internet of things (IoT/IIoT); and mobile and cloud computing—support the integration of these capabilities in a digital twin simulation framework for a variety of applications across the product lifecycle, including design, development, testing, and manufacturing, in addition to the original applications in system sustainment.

Several aspects of US Air Force aviation led the AFRL engineers to use the term “digital twin” for their new concept. The facts that every flight of each Air Force airplane is unique and that the types of missions for which the aircraft are used change periodically led to the idea of using flight simulation to predict airplane performance over time.

While the use of flight simulators is not new, the idea of using flight simulation to predict the engineering performance of an individual aircraft over time is novel. Because the physical configuration of different aircraft of the same make and model is unique and changes periodically,¹ tail-number-specific configuration data are used for the flight simulations.

A digital twin simulates the performance of its physical twin using current, periodically updated knowledge about the state and use of its physical twin. This is in contrast to typical engineering-level analyses that use a nominal physical configuration with average or worst-case initial conditions and boundary conditions and are updated only when major changes in configuration or usage occur.

The significance of the digital twin concept is derived from its key elements:

- Digital twins are designed to provide timely and actionable information about an asset to a decision maker.
- The output is tailored for the asset operator(s), based on both the known physical characteristics of the assets and the details of past, current, and planned use.
- The output of digital twin simulations is updated based on new information about the physical characteristics of the system and/or its past, current, and future use.

¹ Configuration changes are due to (i) repairs associated with wear and other damage that can occur during manufacturing, operation, and maintenance; (ii) design updates to address newly identified performance deficiencies; (iii) design updates to add new capability to the aircraft; and (iv) the installation of missionized equipment (e.g., weapons, external fuel tanks, sensor pods).

AN EXAMPLE: AFRL'S AIRFRAME DIGITAL TWIN PROGRAM

For AFRL's ADT program, the decision maker is the airframe structures engineer and the decision is when to require safety- and maintenance-critical structural inspections for each aircraft in the fleet. Because the engineer is not the owner or operator of the fleet, the decision of when to require inspections must include operational considerations such as an adequate planning horizon and minimized downtime and cost.

Better Maintenance and Planning

No operator wants an engineer to call for maintenance on short notice, particularly if that maintenance takes the aircraft out of service and/or is expensive. The operator typically wants to defer or eliminate maintenance actions as much as possible! ADT aims to provide information about operational and economic risks as a function of flight hours and/or calendar time for each aircraft to help the engineer justify inspection requirements to the operator. Furthermore, ADT enables the engineer to provide these requirements to the operator early enough for the inspections to be incorporated in the operator's plans.

Because of the pressure to reduce maintenance requirements without compromising safety, engineers are always looking for ways to improve their ability to forecast system degradation. In the case of airframe structures, the primary degradation mechanism is fatigue cracking of metallic parts. Such cracking is very difficult to predict because it is driven by factors that are challenging or impossible to know *a priori*. Current engineering methods for forecasting fatigue cracking employ various safety factors that are uniformly applied to an entire fleet for the duration of its service life.

ADT aims to reduce or eliminate the use of uniform, fleetwide safety factors in favor of aircraft-specific probabilistic analyses. Individualized analyses can reduce some uncertainties in the factors of safety, making the analysis results more precise and reducing the likelihood of over- or underinspecting. ADT analyses of the physical characteristics of a given plane will account for differences induced by manufacturing, assembly, operation, and maintenance that influence fatigue cracking behavior based on data gathered throughout its life (figure 1).

Analysis Based on Actual Use

The other way ADT aims to refine analyses is by accounting for how an individual operator uses its aircraft. An operator at a training site flies differently than an operator at a forward operating site, and ADT accounts for such systematic differences to reduce analysis uncertainty. Hence, to forecast fatigue cracking, one must first forecast operations. For ADT, these forecasts are in the form of simulated future flights based on synthesized data from previous

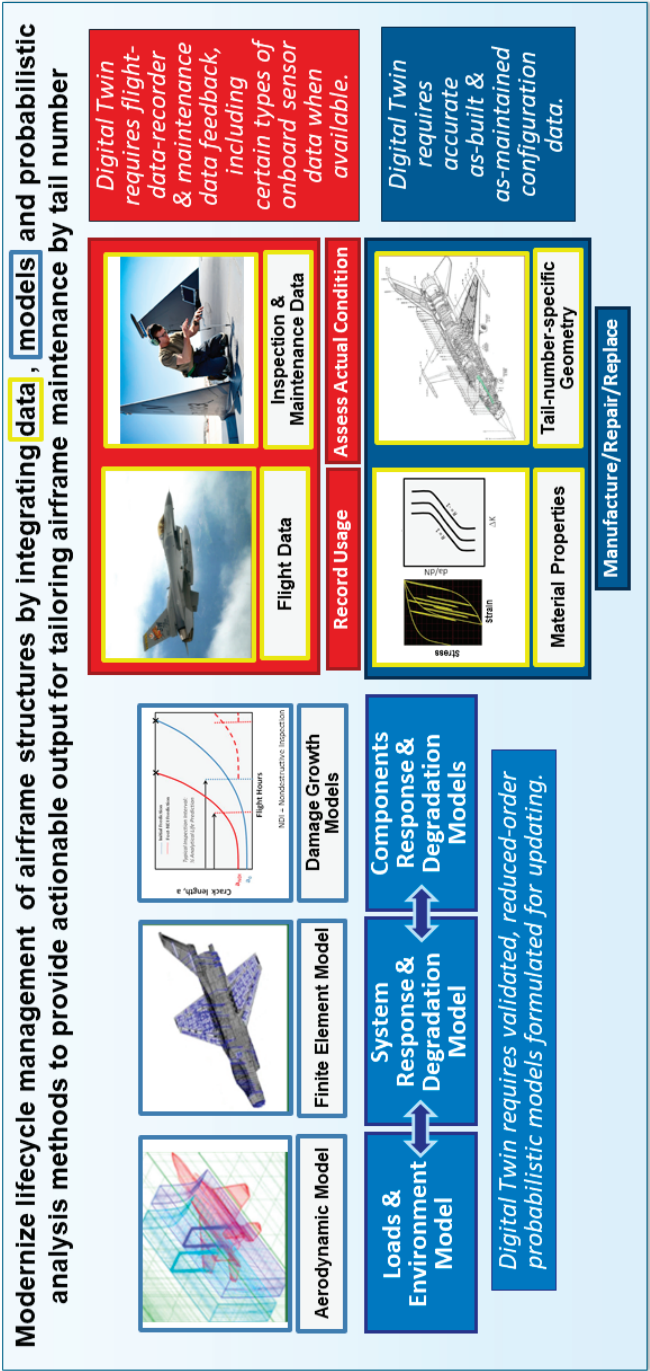


FIGURE 1 Schema of digital twin components for aircraft lifecycle management.

use and assumptions about how the operator intends to fly its aircraft (Asher et al. 2017).

Finally, ADT aims to automatically update when the aircraft configuration is changed and when new flight or maintenance records become available. In this manner, ADT further reduces uncertainty, providing additional opportunity to tailor maintenance requirements.

Proof of Concept

While simulating the engineering performance of a physical aircraft and updating it over its lifetime is a simple concept, in practice it involves the synchronization of numerous models, analyses, and data elements. The time and cost of developing and validating a digital twin are not trivial, so proving that the concept is viable and can benefit both engineers and operators is a necessary early step. However, given the time scale of fatigue cracking in operations—typically thousands of flight hours—proving the concept using operational aircraft would simply take too long.

AFRL engineers therefore focused on developing a laboratory-based method for proof of concept. This effort resulted in a one-of-a-kind full-scale structural experiment in which the external aerodynamic loads from individual flights are applied to full-scale aircraft wings in a laboratory environment at the rate of 200 simulated flights per workweek. This experiment is currently running in AFRL's Structures Validation Facility at Wright-Patterson Air Force Base.

EXCITING FRONTIERS

Since the early days of AFRL's ADT program, the digital twin concept has become increasingly common and enabling technology has advanced. One exciting recent example comes from the US Food and Drug Administration's Office of Science and Engineering Laboratories (FDA 2019), which solicited information on "the capability to perform whole human heart computations with a medically implanted device" and "to create 'virtual patients' and a 'virtual population' such that the FDA can conduct an *in silico* clinical trial with data that can be used to support a proposal for a real clinical trial." While this project doesn't use the term digital twin, many similarities exist, including decision support, tailoring for the individual, uncertainty quantification, and statistical model updating.

CURRENT LIMITATIONS AND CHALLENGES

Digital twin simulations have many potential applications, but significant technical, economic, and social limitations and challenges remain, including the need to

- determine what information to present to the decision maker and how often to update it
- determine the proper level of fidelity for the simulations
- develop methods to reduce the order of the underlying models to reduce computation time
- decide how much to tailor simulations to the individual asset/operator
- develop affordable, reliable means of collecting state and usage data
- develop computationally efficient methods of updating probabilistic simulations
- develop methods to validate probabilistic simulations
- develop methods to synthesize usage and state data
- protect personal privacy and intellectual property
- secure data and models
- address liability for operational failures.

SUMMARY

The concept of simulating engineering performance of physical assets and updating the simulations with state and usage data over time is a powerful idea that is becoming increasingly feasible as enabling technologies mature. Though challenges remain, engineers are envisioning new applications and finding ways to bring them to fruition.

REFERENCES

- Asher I, Wang L, Khan G, Ling Y, Viana F. 2017. Developing a probabilistic load spectrum for fatigue modeling. AIAA SciTech Forum, Jan 9–13, Grapevine TX.
- FDA [US Food and Drug Administration]. 2019. Computational human heart modeling software and services. Solicitation FDA-RFI-1215586, Aug 1. Rockville MD.
- Tuegel EJ, Babish CA IV. 2014. Continuing airworthiness and the Airframe Digital Twin. Proceedings, NATO STO Workshop on Continuing Airworthiness of Ageing Aircraft Systems, STO-MP-AVT-222. Brussels: North Atlantic Treaty Organization.
- Tuegel EJ, Ingraffea AR, Eason TG, Spottswood SM. 2011. Reengineering aircraft structural life prediction using a digital twin. *International Journal of Aerospace Engineering* 2011:154798.

ENGINEERING THE GENOME

Engineering the Genome

CHARLES GERSBACH
Duke University

RENEE WEGRZYN
DARPA

Genome engineering is an ever-growing part of the news cycle as understanding and capabilities in DNA sequencing, synthesis, and modification continue to advance. However, as the use of genome engineering tools—including zinc fingers (ZFNs), transcription activator-like effector nucleases (TALENs), and clustered regularly interspaced short palindromic repeats-associated protein 9 (CRISPR/Cas9)—continues to grow in research and clinical settings, genome cleavage specificity and alternative methods to precisely control gene expression become increasingly important. If specificity and other challenges can be addressed, genome engineering has the potential to alter any DNA or RNA sequence, whether in a bacterium, plant, animal, or human being, and could result in an almost limitless range of possible applications in living things. Development of programmable nucleases could eventually enable the broad application of these or other programmable nucleases to treat human genetic diseases, develop new industrial biotechnological products, improve crop and livestock productivity, and address conservation and invasive species challenges.

The first speaker, Krishanu Saha (University of Wisconsin), introduced genome engineering and the rise of CRISPR-Cas9, and explained the wide-ranging application areas available through this technology, including research and human therapeutics. Subsequent speakers provided specific applications of this technology across species and industries. Omar Akbari (University of California, San Diego) discussed the impact of genome engineering on ecosystems, through the example of mosquito transmission of human diseases, including a discussion of ethical, legal, and social implications considered by researchers in the field. Patrick Boyle of Ginkgo Bioworks described industrial scale-up and wide-ranging applications for manufacturing molecules through

iterative design and development. Finally, Samantha Maragh (National Institute of Standards and Technology) illuminated the need for standards and data sharing for this rapidly evolving field, and showcased the unique perspective and contributions of NIST to help advance the field.

Genome Editing with Precision and Accuracy

KRISHANU SAHA
University of Wisconsin–Madison

Editing the genetic code of living organisms with word-processing-like capabilities has been a goal of life scientists and engineers for decades. For biomedical applications, changing as little as a single base in the 3 billion bases of the human genome might cure disorders such as muscular dystrophy and cystic fibrosis. New genome editing tools are now capable of “one in a billion” specificity, leading to the prospect of new classes of gene and cell therapies. To illustrate both the excitement and the risks, I begin with a cautionary tale that made global headlines, before describing opportunities and challenges in efforts to develop genome editors for biomedical applications.

A CAUTIONARY TALE

In November 2018 Chinese biophysics researcher He Jiankui surprised the world by announcing the birth of two genome-edited babies, the so-called “CRISPR¹ twins” (Cyranoski 2019). It was a momentous step: the intentional alteration of the human germline to produce changes transferred to the next generation.

But the editing outcomes were not as Dr. He expected (Ryder 2018). The Chinese team had injected CRISPR-associated protein 9 (Cas9) genome editing RNA and proteins into human embryos to modify the *CCR5* gene, which encodes a receptor on the surface of immune cells for human immunodeficiency virus (HIV). CRISPR-Cas9 is a nuclease that can cut DNA, and the location of the cut can be programmed by the sequence of a single guide RNA (sgRNA).

¹ CRISPR (clustered regularly interspaced short palindromic repeats) is a family of DNA sequences in the genomes of prokaryotic organisms such as bacteria.

The sgRNAs used by Dr. He targeted two locations in the *CCR5* gene to delete 32 amino acids of the translated CCR5 protein. The deletion of the *CCR5* gene was meant to destroy the ability of HIV to infect T cells and thus make the treated embryos resistant to HIV. But genomic analysis of the CRISPR twins indicated that the intended mutation was not achieved; instead, different insertions and deletions of DNA bases (termed indels) were generated in the *CCR5* gene. These insertions and deletions are anticipated to make the twins more susceptible to influenza, with unknown effects on their susceptibility to HIV.

While there are many lessons to be learned from this experiment (Barrangou 2019; Jasanoff et al. 2019), it is a cautionary tale about using genome editing tools with poor precision. Critiques of the approach indicate suboptimal use of both Cas9 nucleases (there are protein-engineered, higher-fidelity variants) and sgRNAs (other target sequences could have been used) as well as questionable timing of the intervention (the proteins and RNA could have been introduced at a different stage of embryonic development).

The challenges of getting all the parameters right are daunting, and many people, including leaders in the scientific community (Lander et al. 2019), argue that human embryo editing should not even proceed. There is intense activity, however, to attempt to tackle the challenges of editing the human genome after birth, so-called “somatic genome editing,” to mitigate certain diseases. I briefly describe efforts to overcome four types of challenges through improved precision and accuracy in genome editing (figure 1).

CHALLENGE 1: ON-TARGET NUCLEASE ACTIVITY

The most common interpretation of “precision” in genome editing is the ability to edit the genome at the intended target site while limiting edits elsewhere in the genome, commonly called “off-target effects.”

For the CRISPR-Cas9 genome editing systems, Cas9 has been observed to create excessive undesirable mutations (Cradick et al. 2013; Duan et al. 2014; Pattanayak et al. 2013). New methods have therefore been developed to controllably introduce genome editing components, such as ribonucleoproteins (RNPs), and to regulate when and where Cas9 is expressed (Chen et al. 2016; Davis et al. 2015; Hemphill et al. 2015). For example, modified Cas9 nucleases can be selectively activated with small molecules to decrease the gene editing time window (Davis et al. 2015).

Several groups have engineered nickase Cas9 proteins with only one active nuclease domain. A single nickase cannot create a full DNA double-strand break, but when two nicks are paired, the break can be repaired via nonhomologous end joining (NHEJ) in mammalian cells (Ran et al. 2013). While this method lowers off-target effects, the efficiency of the genome editing is greatly decreased as two nicks and two sgRNAs need to be delivered to the nucleus to perform simultaneous cuts.

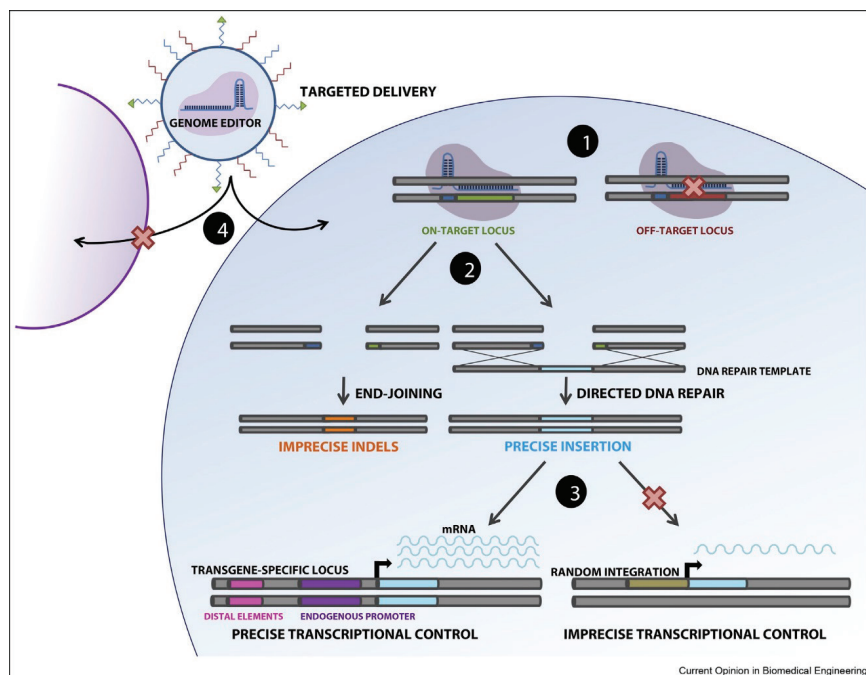


FIGURE 1 Four ways to achieve precision and accuracy for genome editing therapeutics: (1) Binding of genome editing machinery to the intended target genomic locus. The protoadjacent motif in the host genome used by Cas proteins for target recognition are shaded in blue; the DNA sequence in the host genome bound by the genome editor is in green at the intended (on-target) locus and in red at the unintended (off-target) locus. (2) Incorporation of the correct sequence in the edited locus after DNA double-strand break formation or after base editing (not shown). Indels = insertions or deletions of DNA. (3) Precise regulation of integrated transgenes by endogenous promoters and distal elements in comparison to random integration. mRNA = messenger RNA, which represents the expression level of the inserted transgene; perturbed expression levels in cells after genome editing can lead to poor efficacy or adverse events. (4) Delivery to specific cell types by engineered nanomaterials or viral capsids. Reprinted from Mueller et al. (2018) with permission from Elsevier.

Rational protein engineering approaches to modify the nuclease have also generated high-fidelity variants of Cas9: eSpCas9 (Slaymaker et al. 2016), Cas9-HF1 (Kleinstiver et al. 2016), and xCas9 (Hu et al. 2018). Cas9 variants decrease the binding time of the sgRNA to the target sites in the genome, resulting in a decrease in off-target binding and cutting. These high-fidelity Cas9 variants may represent a quick path to clinical relevance as they can greatly reduce off-target events.

CHALLENGE 2: “SCARLESS” INCORPORATION OF NEW SEQUENCES

While the intended on-target gene can be modified using a targeted DNA break, attempting to insert specific bases (i.e., to “write in” sequences) into the genomic target site adds a layer of complexity.

DNA repair pathways in the cell dictate whether new nucleic acids can be inserted, and a major pathway used is the homology-directed repair (HDR) pathway. HDR can generate perfect incorporation of the desired sequence without modifying any other bases in the genome and so is called “scarless editing.” Researchers have attempted to modify DNA repair pathways (Chu et al. 2015; Yu et al. 2015) to increase both the efficiency of HDR and the ratio of precise edits to imprecise mutations. These methods are most applicable for in vitro cell culture applications where potential toxicity is less limiting.

For short insertions, single-stranded oligodeoxynucleotide (ssODN) templates hold significant promise for treating disease because of their ease of synthesis. However, sequence changes encoded by the ssODN are infrequently incorporated after editing (<10 percent), and desired edits are typically outnumbered by other sequence outcomes (presumably from NHEJ). Strategies to link the ssODN to Cas9 help increase HDR (Carlson-Stevermer et al. 2017; Lee et al. 2017). Other methods avoid the use of HDR and leverage other DNA repair pathways (Suzuki et al. 2016).

Base editors are particularly attractive for clinical translation as they avoid DNA double-strand breaks entirely. They employ a catalytically dead version of Cas9 fused to a DNA deaminase to modify base pairs proximal to the sgRNA target, deaminating cytidine bases to form uridine. The modified bases are recognized by the cell as mismatched and corrected to thymidine (Komor et al. 2016). Current work in this area mostly focuses on cytosine (C) > thymine (T) (or analogous guanine (G) > adenine (A)) base conversions, although future versions will aim to allow modifications of any single base (Gaudelli et al. 2017). Additionally, a new strategy, called prime editing, also avoids DNA double-strand breaks and has been demonstrated in vitro with mammalian cells for scarless incorporation of new sequences (Anzalone et al. 2019).

CHALLENGE 3: PRECISE TRANSCRIPTIONAL CONTROL

Even if challenges 1 and 2 are met with perfect accuracy and precision, expression of edited genes to generate RNA transcripts can vary over time as well as across cell differentiation and behavior patterns. Misregulation of the edited transcript can compromise therapeutic efficacy or lead to adverse events. Therefore, it is critical to consider strategies to maximize transcriptional control, especially when inserting new bases.

A striking discovery about the necessity of precise transgene expression recently emerged in the field of chimeric antigen receptor (CAR) T cell therapy. In the CAR T paradigm, a synthetic CAR transgene targeting a cancer-enriched antigen is inserted *ex vivo* into a patient's T cells, which are then expanded and reinfused, thereby engineering the immune system to recognize and target cells bearing the antigen (Piscopo et al. 2018). CRISPR-Cas9 was recently used to generate CAR T cells featuring a transgene at the T cell receptor alpha (*TRAC*) locus, which ensured that CAR expression was regulated by the endogenous *TRAC* promoter (Eyquem et al. 2017). These CAR T cells demonstrated promising results in a leukemic mouse model and also displayed fewer biomarkers of dysfunctional CAR T cells, suggesting that precise transgene control may yield a more potent cell product.

CHALLENGE 4: PRECISE EDITING IN SPECIFIC CELLS AND TISSUES

Precise delivery of editing components to the right cells and tissues remains a challenge as many delivery agents suffer from low efficiency, high toxicity, and immunogenicity, but viral and nonviral delivery agents have been engineered to achieve cell and tissue specificity.

Viral constructs can be engineered to harbor cell- and tissue-specific promoters that drive expression of the gene editing system (Ran et al. 2015; Swiech et al. 2014) so that editing machinery is not expressed in undesired cell types. And several nonviral designs have demonstrated high gene editing efficiencies when used with RNPs, ranging from 30 to 40 percent in cell lines, and up to 90 percent delivery efficiency (Chen et al. 2019; Mout et al. 2017; Sun et al. 2015; Zuris et al. 2015).

Custom biomaterials can also be engineered to direct genetic payloads to specific tissue types to allow gene editing *in situ*, thereby bypassing many of the biomanufacturing challenges associated with *ex vivo* cell therapy. Researchers recently developed DNA nanocarriers with the capacity to deliver CAR transgenes to T cells in a leukemic mouse model by coupling anti-CD3 ligands to polyglutamic acid (Smith et al. 2017). These nanocarriers demonstrated specificity to circulating T cells over other blood cell types shortly after delivery, causing tumor regression.

OUTLOOK

It is likely that strategies to meet these four challenges will be complementary, ultimately enabling more precise genomic surgery in patients' cells. For *in vitro* applications, drug discovery will probably be accelerated by enhanced tools for disease modeling, target validation, and toxicological studies. For *ex vivo* uses, precision-engineered cell and tissue therapies may incorporate more functionality

from synthetic circuits (Weinberg et al. 2017). Finally, for in vivo somatic gene editing applications, injectable viral and nanoparticle strategies could specifically edit stem cells to regenerate tissues and correct disease-causing mutations.

Successful strategies to overcome the challenges described above may pave the way for a new wave of transformative therapeutics.

REFERENCES

- Anzalone AV, Randolph PB, Davis JR, Sousa AA, Koblan LW, Levy JM, Chen PJ, Wilson C, Newby GA, Raguram A, Liu DR. 2019. Search-and-replace genome editing without double-strand breaks or donor DNA. *Nature*, doi:10.1038/s41586-019-1711-4.
- Barrangou R. 2019. Thinking about CRISPR: The ethics of human genome editing. *CRISPR Journal* 2(5).
- Carlson-Stevermer J, Abdeen AA, Kohlenberg L, Goedland M, Molugu K, Lou M, Saha K. 2017. Assembly of CRISPR ribonucleoproteins with biotinylated oligonucleotides via an RNA aptamer for precise gene editing. *Nature Communications* 8:1711.
- Chen Y, Liu X, Zhang Y, Wang H. 2016. A self-restricted CRISPR system to reduce off-target effects. *Molecular Therapy* 24(9):P1508–10.
- Chen G, Abdeen AA, Wang, Y, Shahi PK, Robertson S, Xie R, Suzuki M, Pattnaik BR, Saha K, Gong S. 2019. A biodegradable nanocapsule delivers a Cas9 ribonucleoprotein complex for in vivo genome editing. *Nature Nanotechnology* 14:974–80.
- Chu VT, Weber T, Wefers B, Wurst W, Sander S, Rajewsky K, Kühn R. 2015. Increasing the efficiency of homology-directed repair for CRISPR-Cas9-induced precise gene editing in mammalian cells. *Nature Biotechnology* 33:543–48.
- Cradick TJ, Fine EJ, Antico CJ, Bao G. 2013. CRISPR/Cas9 systems targeting β -globin and *CCR5* genes have substantial off-target activity. *Nucleic Acids Research* 41(20):9584–92.
- Cyranoski D. 2019. The CRISPR-baby scandal: What's next for human gene-editing. *Nature* 566:440–42.
- Davis KM, Pattanayak V, Thompson DB, Zuris JA, Liu DR. 2015. Small molecule-triggered Cas9 protein with improved genome-editing specificity. *Nature Chemical Biology* 11(5):316–18.
- Duan J, Lu G, Xie Z, Lou M, Luo J, Guo L, Zhang Y. 2014. Genome-wide identification of CRISPR/Cas9 off-targets in human genome. *Cell Research* 24:1009–12.
- Eyquem J, Mansilla-Soto J, Giavridis T, van der Stegen SJC, Hamieh M, Cunanan KM, Odak A, Gönen M, Sadelain M. 2017. Targeting a CAR to the *TRAC* locus with CRISPR/Cas9 enhances tumour rejection. *Nature* 543(7643):113–17.
- Gaudelli NM, Komor AC, Rees HA, Packer MS, Badran AH, Bryson DI, Liu DR. 2017. Programmable base editing of A•T to G•C in genomic DNA without DNA cleavage. *Nature* 551:464–71.
- Hemphill J, Borchardt EK, Brown K, Asokan A, Deiters A. 2015. Optical control of CRISPR/Cas9 gene editing. *Journal of the American Chemical Society* 137(17):5642–45.
- Hu JH, Miller SM, Geurts MH, Tang W, Chen L, Sun N, Zeina CM, Gao X, Rees HA, Lin Z, Liu DR. 2018. Evolved Cas9 variants with broad PAM compatibility and high DNA specificity. *Nature* 556:57–63.
- Jasanoff S, Hurlbut JB, Saha K. 2019. Democratic governance of human germline genome editing. *CRISPR Journal* 2(5).
- Kleinstiver BP, Pattanayak V, Prew MS, Tsai SQ, Nguyen NT, Zheng Z, Joung JK. 2016. High-fidelity CRISPR-Cas9 nucleases with no detectable genome-wide off-target effects. *Nature* 529(7587):490–95.
- Komor AC, Kim YB, Packer MS, Zuris JA, Liu DR. 2016. Programmable editing of a target base in genomic DNA without double-stranded DNA cleavage. *Nature* 533:420–24.

- Lander E, Baylis F, Zheng F, Charpentier E, Berg P, Bourgain C, Friedrich B, Joung JK, Li J, Liu D, and 8 others. Comment: Adopt a moratorium on heritable genome editing. *Nature* 567:165–68.
- Lee K, Mackley VA, Rao A, Chong AT, Dewitt MA, Corn JE, Murthy N. 2017. Synthetically modified guide RNA and donor DNA are a versatile platform for CRISPR-Cas9 engineering. *eLife* 6:e25312.
- Mout R, Ray M, Tonga GY, Lee Y-W, Tay T, Sasaki K, Rotello VM. 2017. Direct cytosolic delivery of CRISPR/Cas9-ribonucleoprotein for efficient gene editing. *ACS Nano* 11(3):2452–58.
- Mueller K, Carlson-Stevermer J, Saha K. 2018. Increasing the precision of gene editing in vitro, ex vivo, and in vivo. *Current Opinion in Biomedical Engineering* 7:83–90.
- Pattanayak V, Lin S, Guilinger JP, Ma E, Doudna JA, Liu DR. 2013. High-throughput profiling of off-target DNA cleavage reveals RNA-programmed Cas9 nuclease specificity. *Nature Biotechnology* 31(9):839–43.
- Piscopo NJ, Mueller KP, Das A, Hematti P, Murphy WL, Palecek SP, Capitini CM, Saha K. 2018. Bioengineering solutions for manufacturing challenges in CAR T cells. *Biotechnology* 13(2).
- Ran FA, Hsu PD, Lin C-Y, Gootenberg JS, Konermann S, Trevino AE, Scott DA, Inoue A, Matoba S, Zhang Y, Zhang F. 2013. Double nicking by RNA-guided CRISPR Cas9 for enhanced genome editing specificity. *Cell* 154(6):1380–89.
- Ran FA, Cong L, Yan WX, Scott DA, Gootenberg JS, Kriz AJ, Zetsche B, Shalem O, Wu X, Makarova KS, and 3 others. 2015. In vivo genome editing using *Staphylococcus aureus* Cas9. *Nature* 520(7546):186–91.
- Ryder SP. 2018. #CRISPRbabies: Notes on a scandal. *CRISPR Journal* 1(6):355–57.
- Slaymaker IM, Gao L, Zetsche B, Scott DA, Yan WX, Zhang F. 2016. Rationally engineered Cas9 nucleases with improved specificity. *Science* 351(6268):84–88.
- Smith TT, Stephan SB, Moffett HF, Mcknight LE, Ji W, Reiman D, Bonagofski E, Wohlfahrt ME, Pillai SPS, Stephan MT. 2017. In situ programming of leukaemia-specific T cells using synthetic DNA nanocarriers. *Nature Nanotechnology* 12:813–20.
- Sun W, Ji W, Hall JM, Hu Q, Wang C, Beisel CL, Gu Z. 2015. Self-assembled DNA nanoclews for the efficient delivery of CRISPR-Cas9 for genome editing. *Angewandte Chemie* 54(41):12029–33.
- Suzuki K, Tsunekawa, Y, Hernandez-Benitez R, Wu J, Zhu J, Kim EJ, Hatanaka F, Yamamoto M, Araoka T, Li Z, and 25 others. 2016. In vivo genome editing via CRISPR/Cas9 mediated homology-independent targeted integration. *Nature* 540:144–49.
- Swiech L, Heidenreich M, Banerjee A, Habib N, Li Y, Trombetta J, Sur M, Zhang F. 2014. In vivo interrogation of gene function in the mammalian brain using CRISPR-Cas9. *Nature Biotechnology* 33(1):102–106.
- Weinberg BH, Hang Pham NT, Caraballo LD, Lozanoski T, Engel A, Bhatia S, Wong WW. 2017. Large-scale design of robust genetic circuits with multiple inputs and outputs for mammalian cells. *Nature Biotechnology* 35:453–62.
- Yu C, Liu Y, Ma T, Liu K, Xu S, Zhang Y, Liu H, La Russa M, Xie M, Ding S, Qi LS. 2015. Small molecules enhance CRISPR genome editing in pluripotent stem cells. *Cell Stem Cell* 16(2):142–47.
- Zuris JA, Thompson DB, Shu Y, Guilinger JP, Bessen JL, Hu JH, Maeder ML, Joung JK, Chen Z-Y, Liu DR. 2015. Cationic lipid-mediated delivery of proteins enables efficient protein-based genome editing in vitro and in vivo. *Nature Biotechnology* 33:73–80.

Using CRISPR to Combat Human Disease Vectors

OMAR S. AKBARI

University of California, San Diego

The annual incidence of vector-borne disease exceeds 1 billion globally—roughly half of the world’s population is at risk of infection.¹ Mosquito-borne diseases account for the majority of cases (WHO 2014), but there are no vaccines for most of them, so prevention, mainly through inefficient vector control of limited effectiveness, is the primary method to reduce disease burden. Furthermore, treatments for most mosquito-borne pathogens are also limited, and those that are effective are under threat from increasing pathogen drug resistance.

The severity of the problem is best exemplified by the repeated development of antimalarial resistance in Southeast Asia. In the 1990s parasite resistance to first- and second-line malaria drugs necessitated the development of combination therapies for treatment (Nosten et al. 1987, 1994). However, high resistance to these combination drugs and their later derivatives resulted in an increase in malaria-related deaths in this region (Dondorp et al. 2009; Ménard et al. 2016; Phyo et al. 2016). Therefore, in most cases, vector control is the best approach for reducing the burden of vector-borne diseases.

VECTOR CONTROL TOOLS

Chemical insecticides have historically been an important tool for mosquito control, but they have limitations, most notably their limited efficacy due to increasing vector insecticide resistance and their limited species specificity and duration. While insecticide-driven approaches have been successful in some dis-

¹ World Health Organization, “Vector-borne diseases,” October 31, 2017 (<https://www.who.int/en/news-room/fact-sheets/detail/vector-borne-diseases>).

ease prevention programs (Pluess et al. 2010), for a myriad of reasons they have mixed results overall (Esu et al. 2010; George et al. 2015; Maciel-de-Freitas et al. 2014). Even in areas where sustained vector control has been achieved in the past, insecticide resistance has greatly reduced or eliminated the impact of vector control on disease transmission (Hemingway et al. 2002; Liu 2015; Maciel-de-Freitas et al. 2014).

Given the widespread use of insecticides and limited number of insecticide families available for vector control programs, insecticide resistance will continue to be a barrier to insecticide-based vector control. New control techniques are therefore being evaluated to complement vector control programs.

STERILE INSECT TECHNIQUE FOR INSECT CONTROL

Sterile insect technique (SIT) is the gold standard for genetics-based insect population control. In classic SIT, insects are treated with ionizing radiation to induce male sterility and then released in high frequency to mate with wild females, resulting in nonviable progeny. Over time, repeated mass releases of sterile males suppress and can even eliminate target populations. This approach was used to eradicate the screwworm fly (*Cochliomyia hominivorax*; Krafur et al. 1986), the Mexican fruit fly (*Anastrepha ludens*), and the Mediterranean fruit fly (*Ceratitis capitata*) from regions of North America (Hendrichs et al. 2002).

But in mosquitoes irradiation-based SIT causes high male mortality and exceedingly high fitness costs. For example, field studies show that the release of irradiated, sterile male *Aedes albopictus* led to very limited population reduction (Bellini et al. 2013) likely for these reasons.

So although irradiation-based SIT presents an environmentally friendly method of local population suppression, it is not feasible or scalable in its current form for large-scale control of mosquito populations.

NOVEL VECTOR CONTROL METHODS

In recent years innovative genetic vector control methods, such as the release of insects carrying a dominant lethal (RIDL) (Thomas et al. 2000), have demonstrated large reductions in wild vector populations (Carvalho et al. 2015; Harris et al. 2012). Other novel disease or vector control methods, such as Dengue and Zika virus transmission-blocking *Wolbachia*-infected *Aedes aegypti* and the *Wolbachia* incompatible insect technique (IIT), respectively, are being evaluated in the field (Schmidt et al. 2017). While effective, these methods require large numbers of mosquitoes to be raised, manually sexed, and released as adults in the field near target sites.

Building mosquito mass rearing factories in local disease endemic areas is costly and labor intensive and current procedures are error prone (Gilles et al. 2014; Papathanos et al. 2009). Female release, even in small numbers, is particu-

larly problematic to the *Wolbachia* IIT technology as the release will immunize the target population to the incompatible *Wolbachia* strain and ultimately lead to the failure of the approach. Some studies even indicate that in some contexts, *Wolbachia* actually enhances pathogen infection (Dodson et al. 2014; Hughes et al. 2014) or can have large vector fitness costs, which can be problematic (Joshi et al. 2014).

Additionally, the antibiotic drugs required during rearing of RIDL mosquitoes have high male fitness costs (about 5 percent that of wild-type male fitness) based on RIDL field trials in the Cayman Islands (Harris et al. 2011) and Brazil (Carvalho et al. 2015), due to the loss or alteration of gut microbiome or symbiotic bacteria as well as toxicity to mitochondrial cell functions (Chatzisprourou et al. 2015; Moullan et al. 2015). Therefore, there is still an urgent need for new vector control technologies for the suppression of wild vector populations.

USING CRISPR

The advent of CRISPR² technology has excited the potential to engineer new game-changing technologies and innovative systems that can be used to control wild populations of mosquitoes. Two developments of particular interest are a self-limiting system termed *precision-guided sterile insect technique* (pgSIT) (Kandul et al. 2019) and a *homing-based gene drive* (HGD) (Champer et al. 2016; Esvelt et al. 2014). The unique features of these systems can make them valuable in the future to control mosquitoes, as elaborated below.

pgSIT

The novel CRISPR-based pgSIT mechanistically relies on a dominant genetic technology that enables simultaneous sexing and sterilization, facilitating the release of eggs into the environment and ensuring that only sterile adult males emerge. Importantly, for field applications, the release of eggs will eliminate burdens of manually sexing and sterilizing males, reducing the time and effort involved and increasing scalability. Moreover, the release of eggs should reduce the need to build factories near release sites as eggs could be shipped to release locations from a centralized facility and hatched directly in the environment.

This system was recently systematically engineered in an insect fly model system and was shown to be extremely efficient at generating 100 percent sterile males that could suppress populations. The system functions by mass producing two strains, one expressing the CRISPR-associated protein 9 (Cas9) endonuclease and the other expressing two guide RNAs (gRNAs), one targeting a gene important for female viability and the other a gene important for male fertility. When

² CRISPR (clustered regularly interspaced short palindromic repeats) is a family of DNA sequences in the genomes of prokaryotic organisms such as bacteria.

the two separate strains are crossed the only surviving progeny are sterile males, which can be directly deployed (figure 1A).

Efforts are underway to transfer this technology to mosquitoes, and in the coming years it may be deployed in the field.

Homing-Based Gene Drives

Replacement of wild insect populations with genetically modified individuals unable to transmit disease provides an environmentally friendly, sustainable, and self-perpetuating method of disease prevention. However, transgenes that mediate disease resistance to treatment (refractoriness) may inadvertently compromise the fitness of insects that carry them. Furthermore, wild populations are large, partially reproductively isolated, and dispersed over wide areas.

Population replacement therefore requires a gene drive mechanism to spread linked genes that mediate disease refractoriness through wild populations at greater than Mendelian frequencies. In an effort to achieve this, CRISPR methods have been used to accelerate the development of HGDs in model systems in addition to mosquitoes and even mammals (Champer et al. 2017, 2018; DiCarlo et al. 2015; Gantz and Bier 2015; Gantz et al. 2015; Grunwald et al. 2019; Hammond et al. 2016, 2018; KaramiNejadRanjbar et al. 2018; Kyrou et al. 2018; Li et al. 2019; Windbichler et al. 2011; Yan and Finnigan 2018).

HGDs function by encoding the Cas9 endonuclease and an independently expressed gRNA responsible for mediating DNA base pairing directing Cas9-mediated cleavage at a predetermined site (Champer et al. 2016; Esvelt et al. 2014; Gantz and Bier 2016; Marshall and Akbari 2018). When the HGD is positioned in its target site in a heterozygote, double-stranded DNA breakage of the opposite chromosome can cause the drive allele to be used as a template (i.e., donor chromosome) for DNA repair mediated by homologous recombination. This can result in copying, or “homing,” of the HGD into the broken (receiver) chromosome, thereby converting heterozygotes to homozygotes in the germline, which can bias Mendelian inheritance ratios and lead to an increase in HGD frequency in a population (figure 1B,C).

Given recent progress toward developing HGDs in pest species such as mosquitoes (Gantz et al. 2015; Hammond et al. 2016, 2018; Kyrou et al. 2018; Li et al. 2019), there is significant enthusiasm for their potential use to control wild populations. For example, release of HGDs linked with effector genes that inhibit mosquito pathogen transmission (Buchman et al. 2019a,b; Isaacs et al. 2011; Jupatanakul et al. 2017) may lead to replacement of disease-susceptible mosquitoes with disease-resistant counterparts, thereby reducing pathogen transmission (i.e., population modification drive). Alternatively, HGDs targeting genes that affect the fitness of female mosquitoes could also lead to gradual population declines and potentially even elimination (i.e., population suppression drive) (Kyrou et al. 2018; Windbichler et al. 2008, 2011).

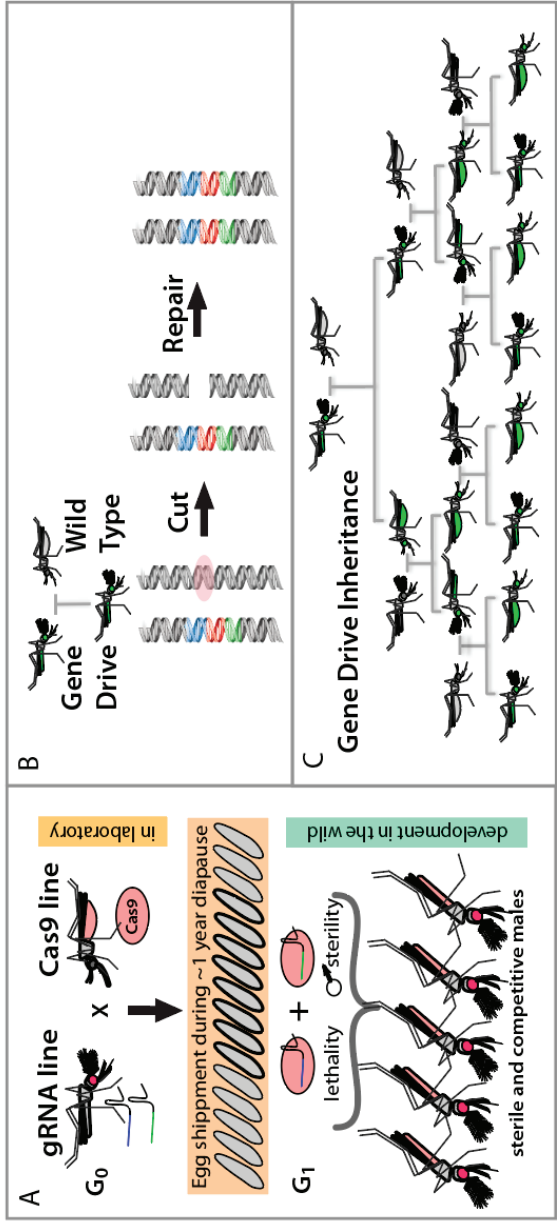


FIGURE 1 Precision-guided sterile insect technique (pgSIT) and homing-based gene drives (HGDs). pgSIT relies on mass rearing two separate strains: the first expresses two guide RNAs (gRNAs) designed to target female viability and male fertility genes, the second expresses the CRISPR-associated protein 9 (Cas9) endonuclease. When crossed, the only surviving progeny are sterile males, which can be repeatedly released as eggs into the environment, resulting in population suppression as they compete with wild males for females (A). HGDs convert heterozygotes to homozygotes using a cut/repair process (B) resulting in biased inheritance and rapid spread into a population (C; green denotes individuals with the gene drive, grey denotes wild-type mosquitoes).

CONCLUSION

Both genetic SIT systems and modification and suppression drives have the potential to transform mosquito population control measures (Burt 2003; Champer et al. 2016; Esvelt et al. 2014), and therefore have excited discussions about their potential use, regulation, safety, ethics, and governance (Adelman et al. 2017; Akbari et al. 2015; NASEM 2016; Oye et al. 2014). Field testing of these systems over the next 5 to 10 years will help illuminate the efficacy and safety concerns of these systems.

REFERENCES

- Adelman Z, Akbari O, Bauer J, Bier E, Bloss C, Carter SR, Callender C, Costero-Saint Denis A, Cowhey P, Dass B, and 23 others. 2017. Rules of the road for insect gene drive research and testing. *Nature Biotechnology* 35(8):716–18.
- Akbari OS, Bellen HJ, Bier E, Bullock SL, Burt A, Church GM, Cook KR, Duchek P, Edwards OR, Esvelt KM, and 17 others. 2015. Safeguarding gene drive experiments in the laboratory. *Science* 349(6251):927–29.
- Bellini R, Medici A, Puggioli A, Balestrino F, Carrieri M. 2013. Pilot field trials with *Aedes albopictus* irradiated sterile males in Italian urban areas. *Journal of Medical Entomology* 50(2):317–25.
- Buchman A, Gamez S, Li M, Antoshechkin I, Lee S-H, Wang S-W, Chen C-H, Klein MJ, Duchemin J-B, Crowe JE Jr, and 2 others. 2019a. Broad dengue neutralization in mosquitoes expressing an engineered antibody. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3398490>.
- Buchman A, Gamez S, Li M, Antoshechkin I, Li H-H, Wang H-W, Chen C-H, Klein, MJ, Duchemin J-B, Paradkar PN, Akbari OS. 2019b. Engineered resistance to Zika virus in transgenic *Aedes aegypti* expressing a polycistronic cluster of synthetic small RNAs. *Proceedings, National Academy of Sciences* 116(9):3656–61.
- Burt A. 2003. Site-specific selfish genes as tools for the control and genetic engineering of natural populations. *Proceedings, Royal Society B: Biological Sciences* 270(1518):921–28.
- Carvalho DO, McKemey AR, Garziera L, Lacroix R, Donnelly CA, Alphey L, Malavasi A, Capurro ML. 2015. Suppression of a field population of *Aedes aegypti* in Brazil by sustained release of transgenic male mosquitoes. *PLoS Neglected Tropical Diseases* 9(7):e0003864.
- Champer J, Buchman A, Akbari OS. 2016. Cheating evolution: Engineering gene drives to manipulate the fate of wild populations. *Nature Reviews Genetics* 17(3):146–59.
- Champer J, Reeves R, Oh SY, Liu C, Liu J, Clark AG, Messer PW. 2017. Novel CRISPR/Cas9 gene drive constructs reveal insights into mechanisms of resistance allele formation and drive efficiency in genetically diverse populations. *PLoS Genetics* 13(7):e1006796.
- Champer J, Liu J, Oh SY, Reeves R, Luthra A, Oakes N, Clark AG, Messer PW. 2018. Reducing resistance allele formation in CRISPR gene drive. *Proceedings, National Academy of Sciences* 115(21):5522–27.
- Chatzisprou IA, Held NM, Mouchiroud L, Auwerx J, Houtkooper RH. 2015. Tetracycline antibiotics impair mitochondrial function and its experimental use confounds research. *Cancer Research* 75(21):4446–49.
- DiCarlo JE, Chavez A, Dietz SL, Esvelt KM, Church GM. 2015. Safeguarding CRISPR-Cas9 gene drives in yeast. *Nature Biotechnology* 33(12):1250–55.
- Dodson BL, Hughes GL, Paul O, Maticchiero AC, Kramer LD, Rasgon JL. 2014. Wolbachia enhances West Nile Virus (WNV) infection in the mosquito *Culex tarsalis*. *PLoS Neglected Tropical Diseases* 8(7):e2965.

- Dondorp AM, Nosten F, Yi P, Das D, Phylo AP, Tarning J, Lwin KM, Arie F, Hanpithakpong W, Lee SJ, and 12 others. 2009. Artemisinin resistance in *Plasmodium falciparum* malaria. *New England Journal of Medicine* 361(5):455–67.
- Esu E, Lenhart A, Smith L, Horstick O. 2010. Effectiveness of peridomestic space spraying with insecticide on dengue transmission: Systematic review. *Tropical Medicine & International Health* 15(5):619–31.
- Esvelt KM, Smidler AL, Catteruccia F, Church GM. 2014. Emerging technology: Concerning RNA-guided gene drives for the alteration of wild populations. *eLife* 3:e03401.
- Gantz VM, Bier E. 2015. The mutagenic chain reaction: A method for converting heterozygous to homozygous mutations. *Science* 348(6233):442–44.
- Gantz VM, Bier E. 2016. The dawn of active genetics. *BioEssays* 38(1):50–63.
- Gantz VM, Jasinskiene N, Tatarenkova O, Fazekas A, Macias VM, Bier E, James AA. 2015. Highly efficient Cas9-mediated gene drive for population modification of the malaria vector mosquito *Anopheles stephensi*. *Proceedings, National Academy of Sciences* 112(49):E6736–43.
- George L, Lenhart A, Toledo J, Lazaro A, Han WW, Velayudhan R, Ranzinger SR, Horstick O. 2015. Community-effectiveness of temephos for dengue vector control: A systematic literature review. *PLoS Neglected Tropical Diseases* 9(9):e0004006.
- Gilles JRL, Schetelig MF, Scolari F, Marec F, Capurro ML, Franz G, Bourtzis K. 2014. Towards mosquito sterile insect technique programmes: Exploring genetic, molecular, mechanical and behavioural methods of sex separation in mosquitoes. *Acta Tropica* 132 (Suppl):S178–87.
- Grunwald HA, Gantz VM, Poplawski G, Xu X-RS, Bier E, Cooper KL. 2019. Super-Mendelian inheritance mediated by CRISPR-Cas9 in the female mouse germline. *Nature* 566(7742):105–109.
- Hammond A, Galizi R, Kyrou K, Simoni A, Siniscalchi C, Katsanos D, Gribble M, Baker D, Marois E, Russell S, and 4 others. 2016. A CRISPR-Cas9 gene drive system targeting female reproduction in the malaria mosquito vector *Anopheles gambiae*. *Nature Biotechnology* 34(1):78–83.
- Hammond AM, Kyrou K, Gribble M, Karlsson X, Morianou I, Galizi R, Beaghton A, Crisanti A, Nolan T. 2018. Improved CRISPR-based suppression gene drives mitigate resistance and impose a large reproductive load on laboratory-contained mosquito populations. *bioRxiv*. <https://doi.org/10.1101/360339>.
- Harris AF, Nimmo D, McKemey AR, Kelly N, Scaife S, Donnelly CA, Beech C, Petrie WD, Alphey L. 2011. Field performance of engineered male mosquitoes. *Nature Biotechnology* 29(11):1034–37.
- Harris AF, McKemey AR, Nimmo D, Curtis Z, Black I, Morgan SA, Oviedo MN, Lacroix R, Naish N, Morrison NI, and 13 others. 2012. Successful suppression of a field mosquito population by sustained release of engineered male mosquitoes. *Nature Biotechnology* 30:828–30.
- Hemingway J, Field L, Vontas J. 2002. An overview of insecticide resistance. *Science* 298(5591):96–97.
- Hendrichs J, Robinson AS, Cayol JP, Enkerlin W. 2002. Medfly areawide sterile insect technique programmes for prevention, suppression or eradication: The importance of mating behavior studies. *Florida Entomologist* 85(1):1–13.
- Hughes GL, Rivero A, Rasgon JL. 2014. Wolbachia can enhance plasmodium infection in mosquitoes: Implications for malaria control? *PLoS Pathogens* 10(9):e1004182.
- Isaacs AT, Li F, Jasinskiene N, Chen X, Nirmala X, Marinotti O, Vinetz JM, James AA. 2011. Engineered resistance to plasmodium falciparum development in transgenic *Anopheles stephensi*. *PLoS Pathogens* 7(4):e1002017.
- Joshi D, McFadden MJ, Bevins D, Zhang F, Xi Z. 2014. Wolbachia strain wAlbB confers both fitness costs and benefit on *Anopheles stephensi*. *Parasites and Vectors* 7:336.
- Jupatanakul N, Sim S, Angleró-Rodríguez YI, Souza-Neto J, Das S, Poti KE, Rossi SL, Bergren N, Vasilakis N, Dimopoulos G. 2017. Engineered *Aedes aegypti* JAK/STAT pathway-mediated immunity to dengue virus. *PLoS Neglected Tropical Diseases* 11(1):e0005187.
- Kandul NP, Liu J, Sánchez Castellanos HM, Wu SL, Marshall JM, Akbari OS. 2019. Transforming insect population control with precision guided sterile males with demonstration in flies. *Nature Communications* 10(1):84.

- KaramiNejadRanjbar M, Eckermann KN, Ahmed HMM, Sánchez Castellanos HM, Dippel S, Marshall JM, Wimmer EA. 2018. Consequences of resistance evolution in a Cas9-based sex conversion-suppression gene drive for insect pest management. *Proceedings, National Academy of Sciences* 115(24):6189–94.
- Krafsur ES, Townson H, Davidson G, Curtis CF. 1986. Screwworm eradication is what it seems. *Nature* 323(6088):495–96.
- Kyrou K, Hammond AM, Galizi R, Kranjc N, Burt A, Beaghton AK, Nolan T, Crisanti A. 2018. A CRISPR-Cas9 gene drive targeting doublesex causes complete population suppression in caged *Anopheles gambiae* mosquitoes. *Nature Biotechnology* 36:1062–66.
- Li M, Yang T, Kandul NP, Bui M, Gamez S, Raban R, Bennett J, Sanchez C HM, Lanzaro GC, Schmidt H, and 3 others. 2019. Development of a confinable gene-drive system in the human disease vector, *Aedes aegypti*. *bioRxiv*. <https://doi.org/10.1101/645440>.
- Liu N. 2015. Insecticide resistance in mosquitoes: Impact, mechanisms, and research directions. *Annual Review of Entomology* 60(1):537–59.
- Maciel-de-Freitas R, Avendanho FC, Santos R, Sylvestre G, Araújo SC, Lima JBP, Martins AJ, Coelho GE, Valle D. 2014. Undesirable consequences of insecticide resistance following *Aedes aegypti* control activities due to a dengue outbreak. *PloS One* 9(3):e92424.
- Marshall JM, Akbari OS. 2018. Can CRISPR-based gene drive be confined in the wild? A question for molecular and population biology. *ACS Chemical Biology* 13(2):424–30.
- Ménard D, Khim N, Beghain J, Adegnika AA, Shafiul-Alam M, Amodu O, Rahim-Awab G, Barnadas C, Berry A, Boum Y, and 80 others. 2016. A worldwide map of *Plasmodium falciparum* K13-propeller polymorphisms. *New England Journal of Medicine* 374(25):2453–64.
- Moullan N, Mouchiroud L, Wang X, Ryu D, Williams EG, Mottis A, Jovaisaite V, Frochoux MV, Quiros PM, Deplancke B, and 2 others. 2015. Tetracyclines disturb mitochondrial function across eukaryotic models: A call for caution in biomedical research. *Cell Reports* 10(10):1681–91.
- NASEM [National Academies of Sciences, Engineering, and Medicine]. 2016. *Gene Drives on the Horizon: Advancing Science, Navigating Uncertainty, and Aligning Research with Public Values*. Washington: National Academies Press.
- Nosten F, Imvithaya S, Vincenti M, Delmas G, Lebihan G, Hausler B, White N. 1987. Malaria on the Thai-Burmese border: Treatment of 5192 patients with mefloquine-sulfadoxine-pyrimethamine. *Bulletin of the World Health Organization* 65(6):891–96.
- Nosten F, Luxemburger C, ter Kuile FO, Woodrow C, Eh JP, Chongsuphajaisiddhi T, White NJ. 1994. Treatment of multidrug-resistant *Plasmodium falciparum* malaria with 3-day artesunate-mefloquine combination. *Journal of Infectious Diseases* 170(4):971–77.
- Oye KA, Esvelt K, Appleton E, Catteruccia F, Church G, Kuiken T, Lightfoot SBY, McNamara J, Smidler A, Collins JP. 2014. Regulating gene drives. *Science* 345(6197):626–28.
- Papathanos PA, Bossin HC, Benedict MQ, Catteruccia F, Malcolm CA, Alphey L, Crisanti A. 2009. Sex separation strategies: Past experience and new approaches. *Malaria Journal* 8(Suppl 2):S5.
- Phyo AP, Ashley EA, Anderson, TJC, Bozdech, Z, Carrara VI, Sripawat K, Nair S, White MM, Dziekan J, Ling C, and 10 others. 2016. Declining efficacy of artemisinin combination therapy against *P. falciparum* malaria on the Thai-Myanmar border (2003–2013): The role of parasite genetic factors. *Clinical Infectious Diseases* 63(6):784–91.
- Pluess B, Tanser FC, Lengeler C, Sharp BL. 2010. Indoor residual spraying for preventing malaria. *Cochrane Database of Systematic Reviews* 4:CD006657.
- Schmidt TL, Barton NH, Rašić G, Turley AP, Montgomery BL, Iturbe-Ormaetxe I, Cook PE, Ryan PA, Ritchie SA, Hoffmann AA, and 2 others. 2017. Local introduction and heterogeneous spatial spread of dengue-suppressing Wolbachia through an urban population of *Aedes aegypti*. *PLoS Biology* 15(5):e2001894.
- Thomas DD, Donnelly CA, Wood RJ, Alphey LS. 2000. Insect population control using a dominant, repressible, lethal genetic system. *Science* 287(5462):2474–76.
- WHO [World Health Organization]. 2014. *A Global Brief on Vector-Borne Diseases*. Geneva.

- Windbichler N, Papathanos PA, Crisanti A. 2008. Targeting the X chromosome during spermatogenesis induces Y chromosome transmission ratio distortion and early dominant embryo lethality in *Anopheles gambiae*. *PLoS Genetics* 4(12):e1000291.
- Windbichler N, Menichelli M, Papathanos PA, Thyme SB, Li H, Ulge UY, Hovde BT, Baker D, Monnat RJ, Burt A, and 1 other. 2011. A synthetic homing endonuclease-based gene drive system in the human malaria mosquito. *Nature* 473(7346):212–15.
- Yan Y, Finnigan GC. 2018. Development of a multi-locus CRISPR gene drive system in budding yeast. *Scientific Reports* 8(1):17277.

Microbes and Manufacturing: Moore's Law Meets Biology

PATRICK BOYLE
Ginkgo Bioworks

Biology is the most powerful known manufacturing “technology.” Proof of this is all around: at the continental scale, the Earth’s land surface is defined by plant life, much of which has been harnessed with agriculture. At the nanoscale, biological systems routinely self-organize with a precision that can’t be matched by the most advanced silicon chip fabrication methods. Even nonbiological technology like petrochemistry uses building blocks that were once biological: petrochemicals, the defining building block of 20th century manufacturing, are derived from the decomposition of prehistoric biomass.

BACKGROUND

For more than 4 billion years, biology has been evolving solutions that scientists and engineers are only now beginning to understand and adapt. For example,

- Antibiotics, aspirin, and many other drugs were isolated from nature. Today, it is possible to further engineer microbes to produce new drug variants.
- Spider silk has been prized for its high strength-to-weight ratio and its promise as a next-generation material. Multiple companies are producing spider silk via engineered microbes.
- Many petrochemicals can now be produced from sustainable carbon sources via engineered microbes.
- Traditional petrochemical products are being enhanced with biological components. For example, modern laundry detergent contains enzymes

(again from engineered microbes) that function in cold water and save heating energy.

All of these applications are advantaged by the fact that biological systems self-assemble, self-repair, and self-replicate. In effect, a microbrewery can serve as a common manufacturing platform for any number of products, simply by engineering the microbe grown in the fermenter.

These advances are possible because the tools are finally available to read (sequence) and write (synthesize) DNA. Both of these technologies have been improving at a rate faster than Moore's law for nearly 20 years. This exponential improvement in the ability to program DNA is driving a technological revolution that rivals the computer revolution of the 20th century—and is impacting manufacturing at a scale not seen since the industrial revolution of the 19th century.

HISTORY OF SYNTHETIC BIOLOGY

Manufacturing with biology far predates the ability to genetically engineer biology. The domestication and breeding of plants and animals for food, clothing, and other materials is synonymous with the emergence of civilization, as these biotechnologies allowed humans to settle in towns and cities with access to cultivated bio-based products.

The earliest domestication efforts were considerable engineering feats in their own right: modern corn bears little resemblance to the teosinte grass that served as the starting point for domestication (Doebley et al. 2006). Similarly, many distinct vegetables such as mustard, broccoli, cauliflower, and even kohlrabi are human-crafted variants of common ancestor species (Dixon 2017). Dogs, cattle, and other animals were similarly differentiated from their wild ancestors via selective breeding over thousands of years.

In the 20th century, the advent of genetic tools and the ability to read and write DNA allowed biologists to consider directly engineering biological organisms for the first time. Many of the early examples of genetic engineering have been extraordinarily successful: human insulin produced in microbes, developed by Genentech in the 1980s, allowed a transition away from the use of animal insulins isolated from pig and cow pancreases (Fraser 2016). In agriculture, genetically modified crops entered use in the United States in the 1990s, and today more than 90 percent of US-grown soybean, cotton, and corn is genetically modified (USDA 2019).

Simply put, biology appeared to be the only technology capable of coordinating atoms with nanometer precision into complex three-dimensional structures. For example, bacterial flagella (tail-like features that propel many bacteria) are self-assembling rotary motors, with a diameter of approximately 25 nm, that rotate at greater than 100 Hz. A typical *Escherichia coli* cell is about 1 μm in length and has several flagella (van den Heuvel and Dekker 2007).

It is hard to imagine how to design machines at the nanometer scale of comparable complexity without biology. Inspired by this, in the mid-1990s a group of electrical engineers, computer scientists, and biologists began to meet regularly to discuss the application of engineering principles to biology. DARPA worked with this group to convene an Information Science and Technology (ISAT) study in 1996 on “cellular computing” that laid the groundwork for the field: seeking to develop methods to understand and program DNA for the purposes of engineering biological organisms to produce new products (Knight and Matsudaira 2016).

Synthetic biology combines efforts from many fields: computer science and electrical engineering abstractions to describe cellular circuitry, metabolic engineering to engineer the metabolic pathways of cells, genetics to understand the control elements of gene expression, and systems biology to measure and simulate cellular systems, among others. Many of the principles developed in the 1996 ISAT study remain relevant to understanding the approaches and applications of synthetic biology today. In particular, a technology development roadmap from that study predicted the development of progressively better tools and modeling capabilities that underlie much of today’s rapidly developing synthetic biology “stack” (figures 1 and 2; Canine 2018).

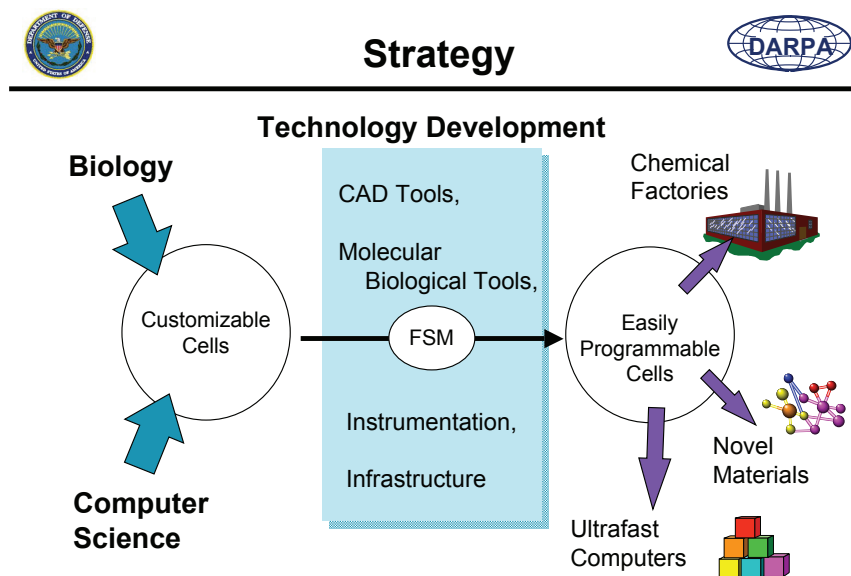


FIGURE 1 Notional strategy for learning how to program biological systems, developed via the 1996 DARPA Information Science and Technology (ISAT) study. CAD = computer-aided design; FSM = finite-state machine. Reprinted from Knight and Matsudaira (2016); image shared under Creative Commons license CC BY-NC-ND 3.0 US.

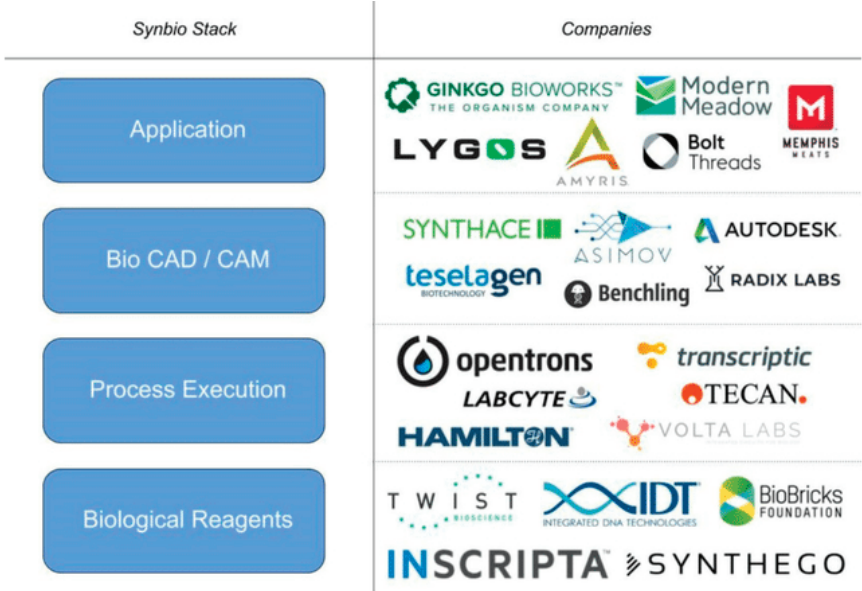


FIGURE 2 Overview of the synthetic biology technology stack. The left-hand column denotes tiers of services, from most application independent at the bottom to purely application focused at the top. Each layer generally draws on resources and services from the layer below it. Companies primarily developing platform technologies in each layer are shown on the right. CAD = computer-aided design; CAM = computer-aided manufacturing. Figure by Will Canine (Opentrons), reused with permission from SynBioBeta (Canine 2018).

DESIGN PRINCIPLES FOR SYNTHETIC BIOLOGY

Two founding design principles of synthetic biology remain especially relevant: the concept of reusable parts and the engineering design cycle.

Synthetic biologists seek to identify and take advantage of modular subunits of biology as reusable parts, to allow the design of more complex systems. For example, genetic control switches such as promoters (the DNA elements that control transcription of a gene into messenger RNA), ribosome binding sites (RNA elements that control the translation of messenger RNA into protein), and other genetic parts have been repurposed to construct oscillators, logic gates, and memory circuits in cells (Boyle and Silver 2009).

The engineering design cycle breaks down the process of engineering into three stages: design, build, and test (DBT).

Design

In contrast to other engineering disciplines, synthetic biologists engineer organisms shaped by evolution, not design. As such, the DBT process in biology requires many more iteration loops than is typical for more mature fields such as mechanical engineering (Petzold et al. 2015). Many successes in synthetic biology follow hundreds or even thousands of failed designs, and often function in only a narrow range of conditions, such as a tightly controlled fermentation tank. These challenges have led to a worldwide effort to develop better “foundries,” facilities that leverage automation to enable rapid prototyping of biological designs, often by conducting many experiments in parallel (Hillson et al. 2019). But trends in DBT technologies have accelerated progress in synthetic biology.

Systems biology, modeling of cellular systems, and data science have enabled synthetic biologists to develop better design algorithms. As in many other fields, machine and deep learning methods are being applied to large biological datasets to refine biological designs (Camacho et al. 2018).

Build

Build technologies in biology have centered around the ability to read and write DNA, the core programming substrate for biology. Here improvement has been defined by two technologies advancing faster than Moore's law: DNA sequencing and synthesis.

Over the past 20 years, the cost to sequence a human genome has fallen more than a millionfold, to less than \$1,000 per genome (NIH 2019). This revolution in sequencing technology has led to exponential growth in the number of sequenced genomes across the tree of life, yielding novel functional parts for synthetic biologists.

Similarly, the cost of DNA synthesis has steadily decreased to today's price of pennies per base pair (Carlson 2017). DNA synthesis is impressively cheap considering the chemistry involved, but still represents a key bottleneck to progress: imagine paying \$0.07 per bit when writing a software program.

Test

Finally, test approaches in biology often make use of cheap DNA sequencing as readouts, and new high-throughput methods for mass spectrometry are allowing researchers to measure the majority of metabolites and proteins in engineered cells (Petzold et al. 2015).

SECURITY FOR BIOLOGY

Synthetic biology is the only engineering discipline where the engineers are made of the same substrate that they are engineering. Since the advent of DNA engineering technologies in the 1970s, researchers and the broader community have raised concerns about the potential misuse of engineered biology to cause harm. As early as 1975, researchers convened to consider the hazards of engineering DNA (Berg et al. 1975). In Cambridge, Massachusetts, public hearings were held in 1976 to develop guidelines for using DNA editing technology as a research tool (Lindsay 1976). These hearings and resulting regulations (such as standard biosafety ratings) have been credited for the emergence of Cambridge and Boston as leading biotech hubs, as the regulations allowed universities and companies to perform this research in a sanctioned environment.

The rapid progress of biological research has led to continual reassessments of biosecurity (NASEM 2017, 2018; NRC 2004). Given the lessons learned in other fields of engineering—particularly in computing and constantly evolving challenges to cybersecurity—safety and security standards, methods of forensics and attribution, and design of biological safety mechanisms must be continually anticipated and addressed. Some examples of these approaches include the biosafety level (BSL) standard, screening protocols to prevent the synthesis of known harmful sequences (DHHS 2015), and deep learning research to identify engineered DNA in sequencing experiments.¹

APPLICATIONS OF ENGINEERED BIOLOGY

Many of the current applications of engineered biology are products of engineered microbes. Microbes have a number of properties that make them useful to engineers: they exhibit fast growth rates, have many genetic tools, and can produce products at commercial scale via fermentation. Many of the early applications for synthetic biology sought to engineer microbes to produce sustainable drop-in replacements for products typically derived from petrochemicals, such as 1,3-propanediol (used in specialty polymers like DuPont's Sorona product), 1,4-butanediol (used in compostable plastics), lactic acid (used to produce polylactic acid polymers), and farnesene (both a fuel and bio-rubber monomer) (Gustavsson and Lee 2016). But the commercial viability of commodity petrochemical replacements was challenged by the falling price of oil in the 2000s, leading to a pivot to higher-value products.

Today, most companies in the synthetic biology space are focusing on products such as fragrances, higher-value materials, and drugs (Schmidt 2017). Because fragrances typically command a high price but are produced in low

¹ The Finding Engineering-Linked Indicators (FELIX) program (<https://www.iarpa.gov/index.php/research-programs/felix>) is an initiative of the Intelligence Advanced Research Projects Activity (IARPA).

volume, they were a natural starting point for companies seeking to develop and commercialize new biomanufactured products.

This focus has parallels with the development of synthetic chemistry as a field, which initially focused on the production of high-price low-volume synthetic dyes before expanding to other products (Yeh and Lim 2007). The approach to transfer the production of volume-limited high-value products to more scalable microbial platforms may be best exemplified by the current competition to produce cannabinoids via fermentation, with hundreds of millions of dollars invested in just the past 2 years (Costa et al. 2019).

Beyond drop-in chemical replacements, many new applications are emerging that are unique to biology. More energy-efficient laundry detergents are effective in cold water in part because they contain enzymes that improve stain removal (Reed 2018). And several companies, such as Indigo Ag, Pivot Bio, and (Ginkgo-affiliated) Joyn Bio, are developing microbial treatments that enhance plant growth or lower the need for conventional fertilizer (Molteni 2018).

Next-generation materials like fermented spider silk may revolutionize textiles, with both Bolt Threads in California and Spiber in Japan developing clothing made of the product (Feldman 2018). Prized for its high strength-to-weight ratio, spider silk is also being explored as a product for aerospace use via a partnership between the German company AMSilk and Airbus (Hyde 2018). Moving the production of silk to microbes means that the proteins that make up silk fibers can be rapidly customized to fit new applications.

Similarly, there has been a growing interest in the production of animal proteins in microbes, allowing vegan production of meat and other animal products without harm to animals. Products such as the Impossible Burger by Impossible Foods in California use microbially produced leghemoglobin protein as a replacement for the hemoglobin proteins that contribute to meat flavor (Wolf 2019). Other companies (including Ginkgo spinout Motif Foodworks) are pursuing the production of a variety of animal proteins to produce vegan dairy foods and other animal-derived products like leather. This approach is also seen as a means to provide high-protein diets more sustainably, given the high-energy requirements for animal-based meat production (Sheikh 2019).

CONCLUSION

It is impossible to predict which of the many applications of synthetic biology will come to define the field as it matures. Unlike all other fields of physical engineering, biology is unique in that it depends on a programmable substrate in the form of DNA. As such, rapid progress has been made on the basis of exponentially improving tools for reading, writing, and debugging biological systems.

It is too soon to know just where and how synthetic biology will evolve, but the stunning diversity of the natural world provides a compelling example of what can be achieved with biology.

REFERENCES

- Berg P, Baltimore D, Brenner S, Roblin RO, Singer MF. 1975. Summary statement of the Asilomar conference on recombinant DNA molecules. *Proceedings, National Academy of Sciences* 72(6):1981–84.
- Boyle P, Silver P. 2009. Harnessing nature's toolbox: Regulatory elements for synthetic biology. *Journal of the Royal Society: Interface* 6(Suppl 4):S535–46.
- Camacho D, Collins K, Powers R, Costello J, Collins J. 2018. Next-generation machine learning for biological networks. *Cell* 173(7):1581–92.
- Canine W. 2018. The synbio stack, part 1. Pleasant Hill CA: SynBioBeta. Available at <https://synbiobeta.com/the-synbio-stack-part-1>.
- Carlson R. 2017. Guesstimating the size of the global array synthesis market. *Synthesis*, Aug 30.
- Costa KA, Vavitsas K, Limas M, Joseph-Nelson B, Cumbers J. 2019. Cannabinoid Fermentation: Scalability, Purity, and Sustainability. Pleasant Hill CA: SynBioBeta.
- DHHS [US Department of Health and Human Services]. 2015. Screening Framework Guidance for Providers of Synthetic Double-Stranded DNA. Washington.
- Dixon G. 2017. The origins of edible brassicas. *Plantsman* 16(3):180–85.
- Doebley J, Gaut B, Smith B. 2006. The molecular genetics of crop domestication. *Cell* 127(7):1309–21.
- Feldman A. 2018. Clothes from a Petri dish: \$700 million bolt threads may have cracked the code on spider silk. *Forbes*, Aug 15.
- Fraser L. 2016. Cloning insulin. Genentech. Available at <https://www.gene.com/stories/cloning-insulin>.
- Gustavsson M, Lee SY. 2016. Prospects of microbial cell factories developed through systems metabolic engineering. *Microbial Biotechnology* 9(5):610–17.
- Hillson N, Caddick M, Cai Y, Carrasco JA, Chang MW, Curach NC, Bell DJ, Le Feuvre R, Friedman DC, Fu X, and 4 others. 2019. Building a global alliance of biofoundries. *Nature Communications* 10(1).
- Hyde E. 2018. Why is global aeronautics leader Airbus partnering with synthetic biology startup AMSilk? Pleasant Hill CA: SynBioBeta. Available at <https://synbiobeta.com/why-is-airbus-partnering-with-synthetic-biology-startup-amsilk>.
- Knight T, Matsudaira P. 2016. Cellular computing, ISAT summer study, August 1996. Available at <https://dspace.mit.edu/handle/1721.1/101009>.
- Lindsay LL. 1976. Recombinant DNA and the Cambridge City Council. Boston: State House Science Resource Office. Available at <https://www.biodiversitylibrary.org/item/129417#page/61mode/1up>.
- Molteni M. 2018. Farmers can now buy designer microbes to replace fertilizer. *Wired*, Oct 2.
- NASEM [National Academies of Sciences, Engineering, and Medicine]. 2017. Dual Use Research of Concern in the Life Sciences: Current Issues and Controversies. Washington: National Academies Press.
- NASEM. 2018. Biodefense in the Age of Synthetic Biology. Washington: National Academies Press.
- NIH [National Institutes of Health]. 2019. The cost of sequencing a human genome. Bethesda MD: NIH National Human Genome Research Institute. Available at <https://www.genome.gov/about-genomics/fact-sheets/Sequencing-Human-Genome-cost>.
- NRC [National Research Council]. 2004. Biotechnology Research in an Age of Terrorism. Washington: National Academies Press.
- Petzold CJ, Chan LJG, Nhan M, Adams PD. 2015. Analytics for metabolic engineering. *Frontiers in Bioengineering and Biotechnology* 3(135).
- Reed S. 2018. Fighting climate change, one laundry load at a time. *New York Times*, Jan 1.
- Schmidt C. 2017. Takeaways from SynBioBeta SF 2017. Pleasant Hill CA: SynBioBeta. Available at <https://synbiobeta.com/takeaways-from-synbiobeta-sf-2017>.
- Sheikh K. 2019. Got impossible milk? The quest for lab-made dairy. *New York Times*, Aug 2.

- USDA [US Department of Agriculture]. 2019. Recent trends in GE adoption. Washington: USDA Economic Research Service. Available at <https://www.ers.usda.gov/dataproducts/adoption-of-genetically-engineered-crops-in-the-us/recent-trends-in-ge-adoption.aspx>.
- van den Heuvel M, Dekker C. 2007. Motor proteins at work for nanotechnology. *Science* 317(5836):333–36.
- Wolf J. 2019. The microbial reasons why the Impossible Burger tastes so good. *ASM Articles*, May 16.
- Yeh B, Lim W. 2007. Synthetic biology: Lessons from the history of synthetic organic chemistry. *Nature Chemical Biology* 3(9):521–25.

Empowering Genome Editing Through Standards

SAMANTHA MARAGH

National Institute of Standards and Technology

A revolution is underway to reengineer the blueprint for life: the genetic code, whose sequence determines identity and function for every living organism. The genome (expressed in DNA base pairs) is the entire complement of an organism's genetic code and is housed in the basic functional unit of life, the cell.

Genome engineering involves tools and techniques to target a specific sequence in a genome and alter the genetic code (genome editing) or to alter the chemical signatures associated with the genetic code (epigenetic engineering). The technology operates by biochemical principles generally applicable to every kind of cell (Carroll 2014; Kim and Kim 2014).

WHAT IS GENOME EDITING?

Genome editing aims to generate edited cells that have permanently changed genetic codes and are functional (Doudna and Charpentier 2014; Gaj et al. 2016). The process typically involves the following:

1. Determine a target location in the genome.
2. Design the editing system to bind to the target location, which may be zinc finger nucleases (ZFNs), transcription activator-like effector nucleases (TALENs), or clustered regularly interspaced short palindromic repeats (CRISPR) with CRISPR-associated proteins (Cas) (Sander and Joung 2014; Urnov 2018).
3. Formulate and deliver the editing system into live cells, where the system finds and binds to the target location in the genome and damages the DNA. This DNA damage gets recognized and repaired by the cell either perfectly,

restoring the original genomic sequence, or imperfectly, as one or more bases of sequence are changed, deleted, or inserted (Komor et al. 2017). Genome editing is achieved as a result of imperfect DNA break repair.

4. Confirm whether and where the genome sequence was changed, what sequence change resulted, and the percentage of the DNA that was changed.
5. Determine whether the engineered cells are fit for the intended purpose. Most often, genome editing is performed on cells of interest in a controlled laboratory setting (*ex vivo* editing), which allows the opportunity to thoroughly analyze edited cells before using them. For complex organisms like humans and tissues like the heart or nervous system, editing systems may need to be delivered directly into the body (*in vivo* editing), in which case the cells edited are already in the organs or tissues, leaving no opportunity to assess whether the cells are fit for the intended purpose before use (Maeder and Gersbach 2016; Yin et al. 2017a).

APPLICATIONS OF GENOME EDITING

Genome editing is being pursued globally by government, academic, and private sectors to transform medicine and bioscience to enable previously impossible advances in areas such as basic biology research, gene therapy, synthetic biology, novel antimicrobials and antivirals, biomanufacturing, agriculture, and food production (Barrangou and Doudna 2016). In some human diseases, just a handful of incorrect letters in the genetic code or as little as a single incorrect letter at a specific position (e.g., sickle cell disease) of the approximately 6.6 billion-letter human genetic code can cause a serious and/or deadly disease.

Genome editing ushers in the first era of technology where the medical field isn't limited to solely managing symptoms and treating illness episodes, but where the cells of a patient may be edited to "fix" a disease at the genetic code level. The biomedical field is daring to think and even speak the word "cure" for diseases that once had few or no treatment options (Fellmann et al. 2017; Porteus 2015; Salsman et al. 2017).

The agriculture industry has also begun targeted editing of crop genomes to have more favorable traits, higher yield, and better disease resistance (Mao et al. 2013; Yin et al. 2017b). And there are large efforts in the environmental science community to engineer microbes with new abilities to produce biofuels and act as biosensors (Ng et al. 2017).

THE EDIT'S IN THE DETAILS: CHALLENGES AND OPPORTUNITIES FOR STANDARDS

Concurrent with the global pursuit to leverage existing genome editing systems, there is significant technological innovation taking place to expand

capabilities for the high-precision targeting of any genomic sequence to make any intended change in any cell. But crucial measurement challenges must be addressed to facilitate the transfer of these technologies into trusted data and products.

Challenges

An editing process is carried out on many cells at a time—often thousands to millions—but there is little technical control over the efficiency of editing and which sequence changes result. Each cell is a closed system where the resulting edits are independent events.

Technical Limitations

Because of technical limitations on the ability to measure the sequence of individual cells at high throughput, edited sequence confirmation involves a bulk measurement sampling of the genomes from a heterogeneously edited cell pool that may contain both intended and unintended edits even at the target site. In addition, because of biological limitations, particularly for human therapeutics, this heterogeneous pool of cells may be the final product.

It is also technically challenging to accurately parse sequencing data from bulk analyses because the number of edits detected can range into the tens or even hundreds for a single genomic location. Bioinformatic pipelines to parse these data were benchmarked on what has been observed in nature: only a handful (if that many) of variants are expected to occur at any one location. It is therefore unclear how accurately sequencing analysis pipelines report what is biologically present in a sample.

The type of edit also contributes to detection difficulty. In general, small edits (e.g., one to tens of bases) are less challenging to sequence verify; large edits (e.g., hundreds of bases or more) are technically challenging to reliably detect and sequence verify.

Off-Target Editing

A prominent measurement challenge for human therapeutic genome editing is the occurrence of unintended, or off-target, editing, when a sequence change occurs at a site or sites other than the target site (Fu et al. 2013; Tsai et al. 2015). Even if the intended edit was effected at the target site, off-target editing may change the cell in a way that makes the therapy or product unsuitable or unsafe. To assess off-target editing, the entire genome, or at least sites where there is reason to think off-target edits could occur, is sequenced after an editing process.

Reliability, Accuracy

There are technical limitations to sequencing the entirety of large genomes like the human genome with sufficient sensitivity to report any off-target edit at any location in the heterogeneous sample. Means of limiting where to sequence for off-target edits are being developed, but there is very limited understanding of their accuracy. Even with reliable sequence data in hand, there's still the challenge of interpreting the significance of unintended edits and understanding whether edited cells are fit for the intended purpose.

Finally, a genome-edited product must be manufactured. Traditional manufacturing approaches don't directly translate to this field where the input is live cells, which must be manipulated precisely and still be live and functional at the end of the process (Harrison et al. 2017, 2018). Moreover, the product may be a personalized therapy for a patient, involving a short storage life and requiring small batch manufacturing with rapid distribution and use.

Opportunities

Standards play an essential role in the translation and durable adoption of technology (figure 1) (Plant et al. 2014, 2018). Standards for genome editing will support and enhance innovation and technology adoption as well as evidence that new biological understanding is based on sound data and that products generated with these technologies are suitable and safe.



FIGURE 1 Types of standards for the translation and durable adoption of technology.

Standards in the form of traceable materials can help increase confidence that a process reliably reports where editing occurred, what edit(s) resulted, and the relative abundance of each edit. For this purpose, traceable material or control samples (a well-qualified series of cells or genomes containing a variety of edits at known relative abundance across the genome) can serve as “ground truth samples” for assessing sequencing methods. These control materials would enable comparability among operators at the same site, operators at different sites or organizations, and sequencing methods.

Standards for datasets and metadata can help enhance the accuracy of sequencing data analysis pipelines, transfer and reproducibility of data, analysis, and data interpretation within and between labs. Shared standard datasets, along with associated metadata detailing the editing and data handling process, will provide a means to benchmark data analysis pipelines, to compare the performance between iterations of both a single pipeline and different pipelines (figure 2). Supporting these technical standards is the need for standardized definitions of key terms in genome editing to enable clear communication of results both within the field and to regulatory agencies.

The US National Institute of Standards and Technology (NIST) has launched the NIST Genome Editing Consortium to work across the genome editing community to develop standards and norms toward filling the needs stated above. Standards for the safe and efficient manufacture of engineered cell products will likely require a paradigm shift and disruptive technologies to address the particular challenges of manufacturing living cells or manufacturing genome editing delivery systems for in vivo editing at large scale.

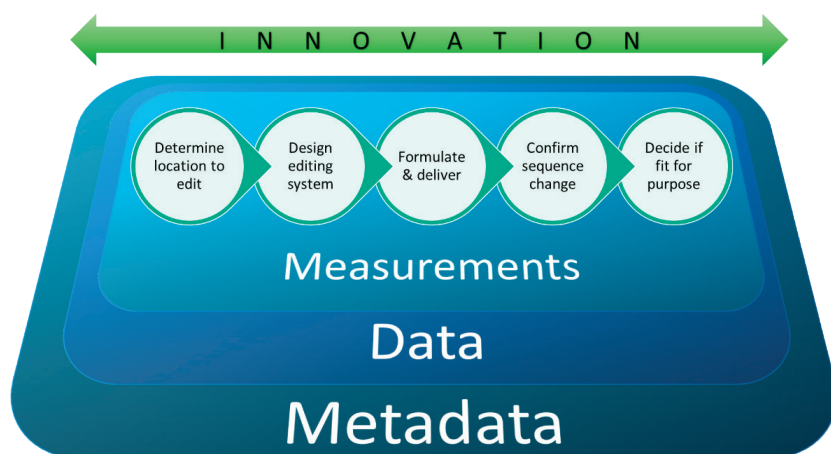


FIGURE 2 Genome editing process overview in context of opportunities for standardization.

CLOSING THOUGHTS

Translating the promise of genome editing into production and medical practice requires robust quantitative assays, accurate data tools, and associated standards and benchmarks to enable high confidence in the characterization of engineered genomes and cells. Steps are being taken to address some of these needs through standards such as physical controls, standard datasets, and a standard lexicon for this field.

Genome engineering technology that is more controlled or not reliant on damaging the genome or changing the sequence at all (e.g., a genome's chemical signature might be changed) is rapidly progressing (Anzalone et al. 2019; Liao et al. 2017; Thakore et al. 2016). Further progress can be made through the development of a suite of tools and technology employing a multidisciplinary approach to address unmet measurement needs such as single cell editing detection and in vivo tracking and monitoring of edited cells once introduced into the environment or a live organism (e.g., in a human for therapeutic treatment).

As genome engineering matures, there will be a need for continuous evaluation of new standards and norms that can support rapid innovation and expansion of this field.

REFERENCES

- Anzalone AV, Randolph PB, Davis JR, Sousa AA, Koblan LW, Levy JM, Chen PJ, Wilson C, Newby GA, Raguram A, Liu DR. 2019. Search-and-replace genome editing without double-strand breaks or donor DNA. *Nature*, doi:10.1038/s41586-019-1711-4.
- Barrangou R, Doudna JA. 2016. Applications of CRISPR technologies in research and beyond. *Nature Biotechnology* 34(9):933–41.
- Carroll D. 2014. Genome engineering with targetable nucleases. *Annual Review of Biochemistry* 83:409–39.
- Doudna JA, Charpentier E. 2014. Genome editing: The new frontier of genome engineering with CRISPR-Cas9. *Science* 346(6213):1258096.
- Fellmann C, Gowen BG, Lin PC, Doudna JA, Corn JE. 2017. Cornerstones of CRISPR-Cas in drug discovery and therapy. *Nature Reviews Drug Discovery* 16(2):89–100.
- Fu Y, Foden JA, Khayter C, Maeder ML, Reyon D, Joung JK, Sander JD. 2013. High-frequency off-target mutagenesis induced by CRISPR-Cas nucleases in human cells. *Nature Biotechnology* 31(9):822–26.
- Gaj T, Sirk SJ, Shui SL, Liu J. 2016. Genome-editing technologies: Principles and applications. *Cold Spring Harbor Perspectives in Biology* 8(12).
- Harrison RP, Ruck S, Medcalf N, Rafiq QA. 2017. Decentralized manufacturing of cell and gene therapies: Overcoming challenges and identifying opportunities. *Cytherapy* 19(10):1140–51.
- Harrison RP, Ruck S, Rafiq QA, Medcalf N. 2018. Decentralised manufacturing of cell and gene therapy products: Learning from other healthcare sectors. *Biotechnology Advances* 36(2):345–57.
- Kim H, Kim JS. 2014. A guide to genome engineering with programmable nucleases. *Nature Reviews Genetics* 15(5):321–34.
- Komor AC, Badran AH, Liu DR. 2017. CRISPR-based technologies for the manipulation of eukaryotic genomes. *Cell* 168(1-2):20–36.

- Liao HK, Hatanaka F, Araoka T, Reddy P, Wu MZ, Sui Y, Yamauchi T, Sakurai M, O'Keefe DD, Núñez-Delicado E, and 6 others. 2017. In vivo target gene activation via CRISPR/Cas9-mediated trans-epigenetic modulation. *Cell* 171(7):1495–507.
- Maeder ML, Gersbach CA. 2016. Genome-editing technologies for gene and cell therapy. *Molecular Therapy* 24(3):430–46.
- Mao Y, Zhang H, Xu N, Zhang B, Gou F, Zhu JK. 2013. Application of the CRISPR-Cas system for efficient genome engineering in plants. *Molecular Plant* 6(6):2008–11.
- Ng IS, Tan SI, Kao PH, Chang YK, Chang JS. 2017. Recent developments on genetic engineering of microalgae for biofuels and bio-based chemicals. *Biotechnology* 12(10).
- Plant AL, Locascio LE, May WE, Gallagher PD. 2014. Improved reproducibility by assuring confidence in measurements in biomedical research. *Nature Methods* 11(9):895–98.
- Plant AL, Becker CA, Hanisch RJ, Boisvert RF, Possolo AM, Elliott JT. 2018. How measurement science can improve confidence in research results. *PLoS Biology* 16(4):e2004299.
- Porteus MH. 2015. Towards a new era in medicine: Therapeutic genome editing. *Genome Biology* 16:286.
- Salsman J, Masson JY, Orthwein A, Dellaire G. 2017. CRISPR/Cas9 gene editing: From basic mechanisms to improved strategies for enhanced genome engineering in vivo. *Current Gene Therapy* 17(4):263–74.
- Sander JD, Joung JK. 2014. CRISPR-Cas systems for editing, regulating and targeting genomes. *Nature Biotechnology* 32(4):347–55.
- Thakore PI, Black JB, Hilton IB, Gersbach CA. 2016. Editing the epigenome: Technologies for programmable transcription and epigenetic modulation. *Nature Methods* 13(2):127–37.
- Tsai SQ, Zheng Z, Nguyen NT, Liebers M, Topkar VV, Thapar V, Wyvekens N, Khayter C, Iafrate AJ, Le LP, and 2 others. 2015. GUIDE-seq enables genome-wide profiling of off-target cleavage by CRISPR-Cas nucleases. *Nature Biotechnology* 33(2):187–97.
- Urnov FD. 2018. Genome editing BC (before CRISPR): Lasting lessons from the Old Testament. *CRISPR Journal* 1:34–46.
- Yin H, Kauffman KJ, Anderson DG. 2017a. Delivery technologies for genome editing. *Nature Reviews Drug Discovery* 16(6):387–99.
- Yin K, Gao C, Qiu JL. 2017b. Progress and prospects in plant genome editing. *Nature Plants* 3:17107.

SELF-DRIVING CARS: TECHNOLOGY AND ETHICS

Self-Driving Cars: Technology and Ethics

CHRISTOFFER HECKMAN
University of Colorado Boulder

HAE-JONG SEO
NVIDIA

Self-driving cars are already taking to the roads, but the technology behind them is still far from fully understood or developed. The challenges ahead for researchers and industry are complex. To overcome them, scientific and engineering breakthroughs combined with fundamental changes in how the public perceives transportation as a service are required. The potential effects on infrastructure, the economy, and society are challenging to quantify due to the number of ways transportation factors into daily life. In response to recent high-profile disasters involving self-driving vehicles, including some that have led to loss of life, there has been an intensified focus on these vehicles' capabilities. Simultaneously, the ethical implications of handing off safety-critical applications to increasingly sophisticated autonomy algorithms have become the subject of intense debate.

The first speaker, Christoffer Heckman (University of Colorado Boulder) provided an overview of the opportunities and challenges presented by self-driving cars.¹ Next, Tae Eun Choe (Baidu) introduced the concept of self-driving vehicles and how they are being developed at scale. Then, John Basl (Northeastern University) explored conversations in philosophy and ethics related to the development of self-driving cars and the technologies and programming that support them. The session closed with a talk by Dorsa Sadigh (Stanford University) about human drivers, their interactions with autonomous and intelligent systems in other vehicles on the shared road, and the societal implications of those interactions.

¹ Paper not included in this volume.

Perceptions of Low-Cost Autonomous Driving

TAE EUN CHOE, XIAOSHU LIU, GUANG CHEN, WEIDE ZHANG,
YULIANG GUO, AND KA WAI TSOI
Baidu

The paper introduces perception algorithms for low-cost autonomous driving in a platform with a full stack of hardware and software developed by the autonomous driving community. We review pros and cons of each sensor and discuss what functionality and level of autonomy can be achieved with them. We also discuss perception modules for dynamic and stationary object detection, sensor fusion (using Dempster–Shafer theory), and virtual lane line and camera calibration.

INTRODUCTION

The development of longer-range and higher-resolution lidar enabled Level 4 autonomous driving with more accurate perception and localization. However, lidar is a less reliable sensor under extreme weather conditions such as heavy rain or snow. Furthermore, its high cost prevents its use in consumer-targeted autonomous cars. In contrast, a camera is more cost-effective and more robust to weather and is a key sensor for traffic light recognition and lane line detection. We present algorithms to achieve autonomous driving using economical sensors such as a camera and a radar.

There are four main pillars in camera- and radar-based perception: pre-processing, deep network, postprocessing, and fusion (figure 1).

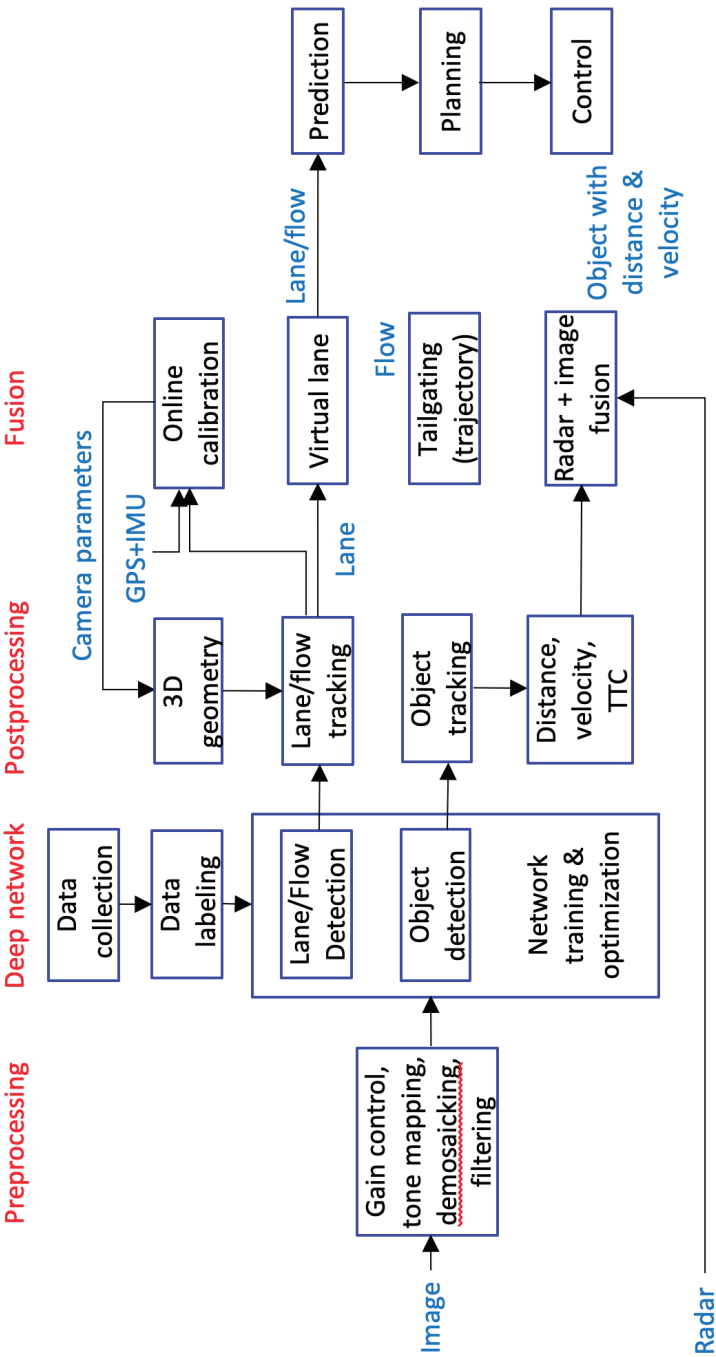


FIGURE 1 Flow diagram of vision-based autonomous driving. (GPS = Global Positioning System; IMU = Inertial Measurement Unit; TTC = Time-To-Collision). Reprinted courtesy of Baidu.

DEEP NETWORK

Data Collection and Labeling

Balanced Data Collection

For deep learning, data collection and labeling are important tasks. Labeled data should be well balanced over time, weather, and road conditions, and should cover night, dawn, sunrise, strong shadow, and sunset on one axis. Another axis is for weather (sunny, rainy, snowy, foggy). The third axis is for road conditions such as straight, curved, fork, merge, or intersection. All the data should be evenly distributed on each axis. Figure 2(a) illustrates the actual data distribution for different environments and 2(b) shows the distribution of data after balanced data collection.

Data Labeling

Before driving, the routine to log vehicle information (yaw rate, speed), GPS, wipers (rain, snow), low/high beam, timestamp, location, and all sensor data should be implemented. The more such data are saved, the easier the labeling process. The car should be driven in the center of a lane as much as possible to imitate autonomous driving. In addition, there should be a simple button to save the past 30 seconds of data when the car experiences specific or rare events. After driving, any duplicated or similar scenes should be removed, especially when the car stops. Pedestrian faces and license plates of other vehicles should also be removed to protect privacy. After such processes are completed, data can be labeled.

Autolabeling

Because manual labeling is costly and subject to human error, automatically labeled data should be included in the training dataset. A good candidate for autolabeling is a stationary object such as a lane, traffic light, traffic sign, or any road landmark. First, near-view objects are detected by an existing detector while driving. After driving some distance (e.g., 200 meters), the previous scenes are reviewed. Assuming near-view object detection and motion estimation are accurate, the accumulated set of detected near-view objects can be labeled far away.

For smart recording, we designed multiple events such as deceleration, curves, cut-in of a neighboring vehicle, cut-out of a closest in-path vehicle, and bumps in the road. When such an event occurs, the data before and after it are saved automatically.

For autolabeling, speed and yaw rate from the vehicle's Controller Area Network (CAN) bus or inertial measurement unit (IMU) data should be recorded for accurate motion estimation. IMU is a useful sensor to measure a vehicle pose.

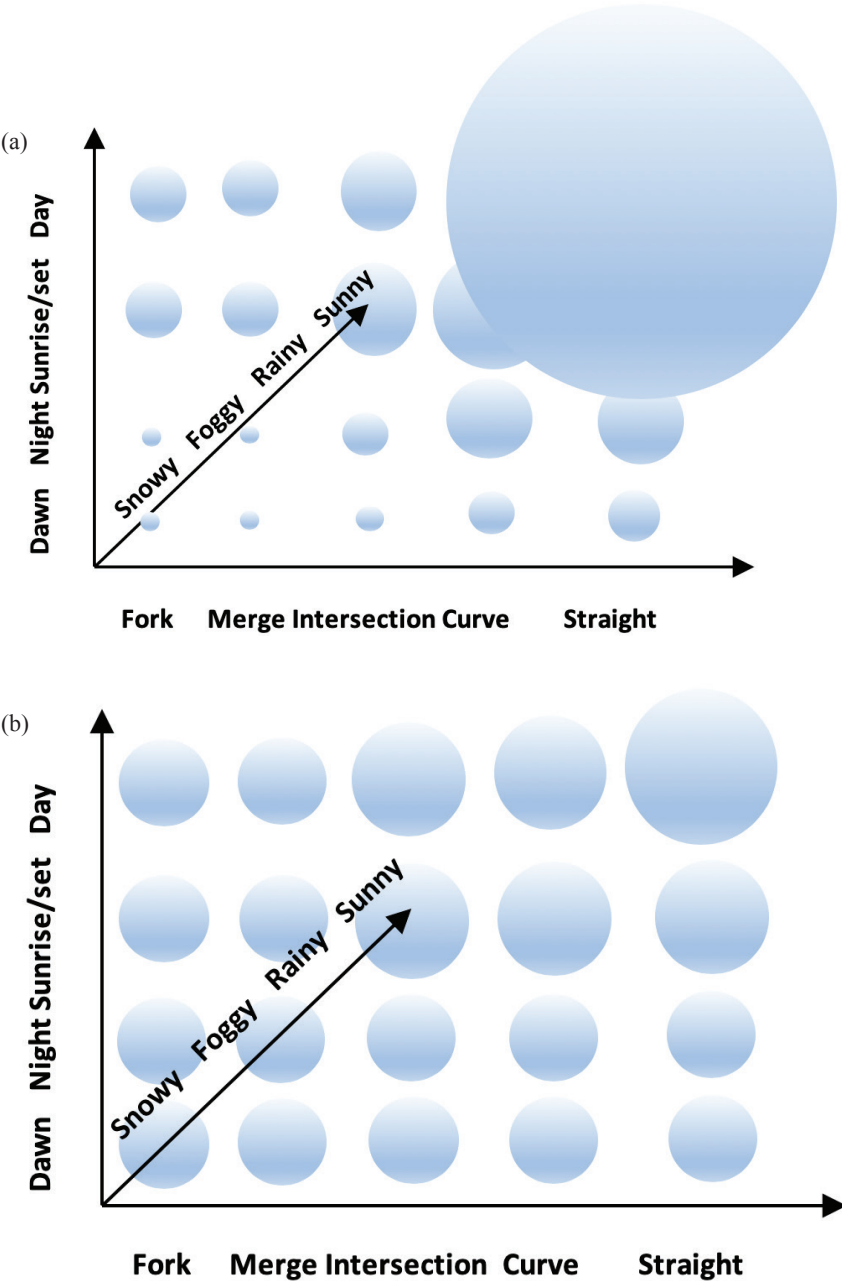


FIGURE 2 Data distribution. (a) Unbalanced real data. (b) Balanced data. Reprinted courtesy of Baidu.

With IMU data accumulated over time, motion from time t to time $t-n$ can be easily estimated.

When the ego-vehicle's pose at time t is $M_{t@t-1}$, where M is a 4×4 matrix with rotation and translation elements, the motion from time t to $t-n$ will be

$$M_{t@t-n} = M_{t@t-1} M_{t-2@t-3} M_{t-3@t-4} \dots M_{t-n+1@t-n}$$

The simple motion matrix $M_{t@t-n}$ directly converts the current vehicle pose to n -th previous pose. For autolabeling, multiple near-view detections are accumulated over time with the estimated motion. The accumulated detections are later projected to the image 200 meters ago for autolabeling. Unlike manual labeling, autolabeling captures objects in three-dimensional form (3D); this includes the road surface, which is reconstructed in 3D to label hill crests, bumps, and even clover leaves. An example of automated labeling is illustrated in figure 3. Autolabeling can capture invisible lane lines (figure 3a) and 3D lane lines (figure 3b).

Network Training and Optimization

Preprocessed images are transferred to a deep neural network for object detection and tracking, lane line and landmark detection, and other computer vision problems. For real-time processing of high-framerate and high-resolution imagery data, network compression is required. In the literature, there are two main network compression approaches:

- Lower-bit approximation: Rather than using the conventional 32-bit float as a weight representation, float32 is quantized into INT8 to achieve real-time implementation (Dettmers 2016).
- Network layer reorganization: When there are multiple tasks, the network structure can be reorganized by sharing common layers and removing unnecessary layers.

Object Detection

In a traffic scene, there are two kinds of objects, stationary and dynamic. The former includes the lane, traffic lights, streetlamps, barriers, bridges over the road, and the skyline; dynamic objects are pedestrians, cars, trucks, bicycles, motorcycles, and animals, among others. For object detection, YOLOv3 (darknet) is used as a base network (Redmon and Farhadi 2018); it accounts for additional object attributes such as 3D size, 3D position, orientation, and type. Detected multiple objects are tracked across multiple frames using a cascade-based multiple hypothesis object tracker.



FIGURE 3 Example of autolabeling of lane lines. (a) Hidden lane lines and (b) 3D lane lines are autolabeled. Reprinted courtesy of Baidu.

Lane Detection

Among stationary objects, a lane is a key stationary object for both longitudinal and lateral control. An “ego-lane”¹ monitor guides lateral control, and any dynamic object in the lane determines longitudinal control. We use the same YOLO (darknet) as a base network and add extra lane tasks to detect the relative positions (left, right, next left, next right, curb lines) and types (white/yellow, solid/broken, fork/split) of lanes.

¹ The “ego-lane” is that of the autonomous vehicle.

POSTPROCESSING AND FUSION

Sensor Fusion

Installation of multiple sensors around a vehicle facilitates full coverage of the environment and redundancy for safety. Each sensor has different capacities: the range of lidar is short but its 3D measurement is accurate; radar provides longitudinally accurate but laterally inaccurate distance and velocity measurements; a camera is accurate for lateral measurement but less so for longitudinal measurement. We learn a *prior* and belief function of each sensor and fuse all sensor output using Dempster–Shafer theory (Wu et al. 2002).

CIPV Detection and Tailgating

The trajectories of all vehicles are captured with respect to the autonomous vehicle. Among them, the closest in-path vehicle (CIPV) is chosen for longitudinal control and for tailgating a CIPV when, for example, there is no lane, such as at an intersection.

Camera Calibration

Camera calibration is challenging but the most important procedure. There are three categories of camera calibration:

- Factory (initial) calibration: At the factory, we estimate intrinsic and extrinsic camera parameters using fixed targets. However, the camera position changes over time and therefore the parameters need to be updated frequently.
- Online calibration: The long-term position of the camera must be estimated with respect to the car body. An online camera calibration module calibrates the camera position in every frame. A change in pitch angle of even 0.3 degrees can result in seriously incorrect vehicle control. For calibration, any object on the road can be used, such as parallel lane lines, vertical landmarks, or any known size of cars or optical flow.
- Instant pose estimation: The pose of a car changes in every frame. When the vehicle passes over a bump, the pose changes a lot. The pose can be estimated by IMU but these data are too noisy to use directly. It can instead be estimated by visual features: by tracking stationary objects, it is possible to estimate motion and provide a much more accurate 3D perception of the scene.

Virtual Lane

When there is no lane line, all lane detection results and tailgating flow are combined spatially and temporally to determine a virtual lane. The virtual lane output is fed to planning and control modules for actuation of the self-driving vehicle.

CONCLUSION

We have shown perception algorithms for low-cost autonomous driving using a camera and radar. As deep neural networks are the key tool for solving perception issues, data collection and labeling became more important tasks. For sustainable data labeling, autolabeling is introduced. For autonomous driving, dynamic object detection/tracking and stationary object detection algorithms are discussed. A Dempster–Shafer-based sensor fusion algorithm is used to handle multiple sensor fusion. Additionally CIPV, tailgating, and camera calibration algorithms are introduced.

REFERENCES

- Dettmers T. 2016. 8-bit approximations for parallelism in deep learning. International Conference on Learning Representations, May 2–4, San Juan, Puerto Rico. Available at arXiv:1511.04561.
- Redmon J, Farhadi A. 2018. YOLOv3: An incremental improvement. <https://arxiv.org/abs/1804.02767>.
- Wu H, Siegel M, Stiefelhagen R, Yang J. 2002. Sensor fusion using Dempster–Shafer theory. Proceedings, IEEE Instrumentation Technology Conference, May 20–22, Vail.

Why Everyone Has It Wrong About the Ethics of Autonomous Vehicles

JOHN BASL
Northeastern University

JEFF BEHREND
Harvard University

Autonomous vehicles (AVs) raise a host of ethical challenges, including determining how they should interact with human drivers in mixed-traffic environments, assigning responsibility when an AV crashes or causes a crash, and how to manage the social and economic impacts of AVs that displace human workers, among others. However, public and academic discussion of the ethics of AVs has been dominated by the question of how to program AVs to manage accident scenarios, and in particular whether and how to draw on so-called “trolley cases” to help resolve this issue. Some in the debate are optimistic that trolley cases are especially useful when addressing accident scenarios, while others are pessimistic, insisting that such cases are of little to no value.

We summarize the debate between the optimists and pessimists, articulate why both sides have failed to recognize the appropriate relationship between trolley cases and AV design, and explain how to better draw on the resources of philosophy to resolve issues in the ethics of AV design and development.

AV ACCIDENT SCENARIOS AND TROLLEY CASES

Autonomous vehicles will inevitably be in accident scenarios in which an accident that causes harm (to pedestrians, passengers, etc.) is unavoidable. Whereas human drivers in these circumstances have very limited ability to navigate them with any sort of control, AVs might be in a position to “decide” how to distribute the harms. It has seemed to many that because with AVs there is some ability to exercise control over how harms are distributed, it is essential to think carefully about how to program AVs for accident scenarios. The question is how to do so?

It has not escaped notice that some accident scenarios bear a resemblance to what are known in philosophy as “trolley cases.” These are imagined scenarios in which a runaway trolley will result in the death of some number of individuals unless a choice is made to divert or otherwise alter the trolley’s course, resulting in some other number of deaths.

In the classic trolley case, the trolley is headed down a track and will kill five people who cannot escape. A bystander has the ability to pull a switch and divert the trolley onto another track. However, on this track there is one person who cannot escape and will die if the trolley is diverted. A similar-seeming scenario involves an AV that is traveling down a street when suddenly a group of pedestrians runs into the street. The only way to avoid hitting them is to take a turn that will result in the death of a pedestrian on the sidewalk.

In another version of the trolley case, a trolley cannot stop and will kill five people unless an object of sufficient weight is pushed in front of it. A bystander has the option of pushing a large person off a bridge and onto the tracks in a way that would stop the train before it kills the five. Again, a case involving an AV might have a similar structure: Perhaps an empty AV has gone out of control and will hit five pedestrians unless another AV with a single passenger drives itself into the first AV.

TROLLEY OPTIMISM

Trolley Optimism is the view that trolley cases can and should inform how AVs are programmed to behave in these sorts of accident scenarios. The general proposal is that various kinds of trolley cases can be constructed, a verdict is reached about what action or behavior is appropriate in that case, and then that verdict is applied in the case of AVs, programming them to behave in a way that mirrors the correct decision in the analogous trolley case (Hübner and White 2018; Lin 2013; Wallach and Allen 2009).

While trolley cases may be born of philosophy, Trolley Optimism is not confined to philosophy departments (see Achenbach 2015; Doctorow 2015; Hao 2018; Marshall 2018; Worstall 2014). Consider the Massachusetts Institute of Technology’s (MIT’s) Moral Machine project, which has a variety of components. One is a website that presents visitors with different accident scenarios and asks how the visitor thinks the car ought to behave in that scenario. The scenarios involve many variables, testing visitors’ judgments about, for example, how to trade off people and animals, men and women, the elderly and children, and those who obey walk signals and those who don’t.

While some might see the Moral Machine project as simply a tool for collecting sociological data, others think that the data, in aggregate, should be used to decide how AVs should be programmed to behave in accident scenarios (Noothigattu et al. 2017). Whereas philosophers might endorse a type of Trolley Optimism that aims to determine *the* correct thing to do in a trolley case and

program AV behavior in accident scenarios accordingly, the Moral Machine's democratic variant leaves it up to the people.

SOME QUESTIONABLE GROUNDS FOR PESSIMISM

Trolley Pessimism is the view that it is a mistake to draw on trolley cases to think about how to program AVs to behave in specific accident scenarios. Different forms of Trolley Pessimism can be distinguished on the basis of what mistake they identify.

Challenges to the Validity of Thought Experiments

One basis for Trolley Pessimism is a distaste for using thought experiments to arrive at conclusions. Sometimes, this is grounded in the idea that thought experiments that philosophers deploy are so idealized and unrealistic that they are useless for navigating the real world.

We think these sorts of objections rest on a mistaken view of the function and value of thought experiments; we set that aside except to note that a key motivation for Trolley Optimism is that accident scenarios seem to closely resemble trolley cases. If trolley cases are useless for thinking about accident scenarios it isn't because the cases are too unrealistic to be of any use. At the very least, a plausible basis for pessimism must articulate the differences between trolley cases and AV accident scenarios that prevent reasonable conclusions about what to do in the latter based on judgments about the former.

Disanalogy

Another basis for Trolley Pessimism tries to show that there is indeed some point of difference between trolley cases and the behavior of AVs in accident scenarios that makes verdicts in the former inapplicable to decisions about what to do about the latter. What are the differences between trolley cases and AV accident scenarios that justify this form of pessimism?

Nyholm and Smids (2016) point to several points of disanalogy. For example, trolley cases set aside questions of the moral and legal liability of those who are deciding how to act. The person who will decide whether to divert the trolley, it is assumed, will not be held responsible or liable for whichever choice they make. But these considerations should inform deliberations about how AVs should behave in accident scenarios.

Another point of disanalogy is that, in trolley cases, the outcomes of various decisions are stipulated to be known with certainty, whereas in the case of AV accident scenarios, despite what one may want or intend a vehicle to do, there is some uncertainty about whether the vehicle's behavior will generate the desired outcome.

Again, we think this is not a plausible basis for Trolley Pessimism, as explained below.

Ways to Address Trolley Pessimism

While it is true that traditional trolley cases do stipulate away issues of legal and moral liability and stipulate outcomes with certainty, there is in principle no reason why thought experiments can't take these variables into account.

It is possible to develop a case that asks what should be done assuming some particular legal liability regime, enumerating the costs to the agent making the decision. Similarly, a case could be constructed in which pulling a switch has an 80 percent chance of altering the course of a trolley, incorporating deliberations about whether this alters one's moral obligations. The creators of the Moral Machine might even be invited to build these variables into their cases, collect data about what people think should be done in those circumstances, and then aggregate the data to dictate the behavior of AVs in accident scenarios.

THE TECHNOLOGICAL BASIS FOR TROLLEY PESSIMISM: LESSONS FROM MACHINE LEARNING

There is a better basis for pessimism, in the very nature of AV enabling technology: machine learning (ML) algorithms.

What Is Machine Learning?

For those unfamiliar with ML algorithms, we contrast them with what we call traditional algorithms. An algorithm is a set of instructions for executing a task or series of tasks to generate some output given some input. In a traditional algorithm the instructions are laid out by hand, each step specified by a programmer or designer. In contrast, ML algorithms themselves generate algorithms, which do not have the steps used to carry out some task specified by a programmer.

A good analogy for some forms of machine learning, namely *supervised* and *reinforcement* learning, is dog training. It is not possible to just program a dog to respond to the words "sit," "stay," "come," and "heel" by wiring its brain by hand. Instead, when training a dog, it is common to arrange for situations where the dog will engage in some desired behavior and then reward the dog. For example, a trainer might hold a treat in front of a dog's nose and then lift the treat into the air, causing the dog naturally to raise its head and drop its back legs. The dog is then rewarded. After many repetitions, the word "sit" is said right before the treat is lifted. Eventually the dog sits on command, having learned an output for the input "sit."

For machine learning, a programmer provides an ML algorithm with a training set, a dataset that includes information about which outputs are desirable and

which are not. The learner then generates an algorithm that is meant to not only yield appropriate input–output pairs when it is fed inputs that match those in the test set, but to extrapolate beyond the test set, yielding, the programmer hopes, desirable outputs for new input data.

Machine Learning and Autonomous Vehicles

Machine learning is a powerful tool. It allows programmers to develop algorithms to solve problems that would otherwise be extremely tedious or impossible.

The AVs likely to be on the road in the foreseeable future will rely on ML technologies; at the very least, machine learning is at the heart of the detection systems used in AVs. Those systems take in data from various sensors (radar, lidar, cameras) and translate the data to some output that other AV systems use to drive the car, to maintain its position within driving lanes, to slow when there is a car in front of it but not when there is merely a piece of litter.

The fact that AVs depend so heavily on ML algorithms grounds a case for Trolley Pessimism. To see why, first note that how an AV behaves in any given accident scenario is mediated by how the algorithm that governs its behavior is trained. The ML training set must be organized to achieve the AV's behavior in a particular accident scenario. For example, to produce an AV that suddenly confronts a scenario where it must swerve and risk harm to its passenger or maintain course and hit a number of pedestrians, such scenarios must be included in the training set and a particular input–output pair marked as desirable.

This is not the only way to achieve the desired behavior; the point is that behavior in particular scenarios is influenced by choices that programmers and designers make about how to train the ML algorithms. These choices involve ethical choices.

Ethical Choices in Machine Learning

AV programmers will have to make choices about, for example, what proportion of the training data is dedicated to accident scenarios at all. One programmer might focus on nonaccident or typical driving scenarios, including no data about how a car should behave in accident scenarios. Another might dedicate half the training data to everyday driving scenarios and half to accident scenarios. Let's imagine these two programmers are on the same team and arguing about what proportion of the training set should be dedicated to scenarios where the car detects itself to be in an accident where harms can't be avoided. The first programmer argues that the car will very rarely be in those kinds of situations and instead should be trained for the most likely scenarios. The second argues that even if the accident scenarios are rare, it's extremely important to make sure the car does the right thing! The first programmer counters that if they dedicate enough of the training set to getting certain behaviors in accident scenarios, it could make the car

less safe in typical driving scenarios or even put the car into accident scenarios more often! Clearly this argument over how to train the algorithm that will help govern AV behavior is an ethical one: it invokes various value judgments and judgments about how those values are implicated in potential outcomes.

It follows from the facts that decisions about how to organize the training regime for AV behavior are ethical decisions and that they mediate questions about how AVs should behave in particular driving situations. Trolley cases do not provide direct guidance about how AVs should behave in accident scenarios, despite any superficial similarities. There are several ways to see why.

Let's suppose that in the imagined argument between the programmers above, the first programmer is correct, that the algorithms that generate AV behavior should not be based on any data about accident scenarios. That is, after engaging in careful deliberation about relevant values, no accident scenarios or anyone's verdicts about how an AV should behave in them should inform the training of ML algorithms used in the AV. The resulting algorithm will still generate behaviors in such scenarios, but the training set won't have been designed to generate any particular behaviors in those scenarios. In this case, the answer to the question "should programmers try to model the behaviors of AVs on the verdicts of trolley cases?" is clearly "no!" because the programmers have accepted that they shouldn't be trying to train for accident scenarios *at all*.

A Thought Experiment

Another way to illustrate the point is to recognize the way trolley cases—thought experiments, imagined scenarios used to help test more general principles—typically function in ethical theorizing.

Let's imagine we are wondering whether we should accept a principle that we should act in such a way as to maximize the total number of lives saved (holding fixed things like whether the people whose lives are saved are good people, how large their families are, etc.). Someone asks us to consider the standard trolley case. We imagine a train hurtling down the tracks and must decide whether diverting the trolley onto a track that results in fewer deaths is the right thing to do.

Let's assume we come to see this trolley case as supporting the principle that we should maximize total lives saved. If we think that principle is true, programmers and designers should abide it by when deciding how to train AVs. The Trolley Optimist might think that the above case justifies efforts to ensure that an AV in an accident scenario will not drive into a larger crowd to spare a smaller. However, it could very well turn out that abiding by the principle we've settled on has the implication that we are not justified in doing so.

To see why, imagine that in the programmers' debate above both are committed to maximizing lives saved. The first programmer argues that this can be done by avoiding accident scenarios as much as possible and to do that they should not train the algorithm for accident scenarios at all but for how to stay out of them.

This might have the result that when an AV is in an accident scenario it does veer into a larger crowd to save a smaller, but given that the programmers' decision is to program for the whole range of behaviors the car will encounter, they haven't failed to take into account the lesson of the trolley case; they've taken it into account in just the right way. The other programmer might come to agree, seeing that if they were to emphasize a training regime that included more accident scenarios that looked like trolley cases, the AV would end up in those scenarios more often or perform poorly in other driving scenarios, causing additional fatalities. This programmer might see it as regrettable that the best way to maximize lives saved overall, given the decision the design team faces, will produce an AV that veers to kill the five instead of the one in a very narrow range of cases, while still acknowledging that this is the approach that conforms with the principle.

CONCLUSION

To be clear, we are not endorsing any particular view of how AVs should be trained or a particular principle as governing that decision. Our point is that the Trolley Optimist makes a mistake in thinking that the lesson from trolley cases is a lesson for how an AV should behave in a superficially similar case.

The ethical question that designers face is not about the right thing to do in a specific scenario but about how to design for the wide range of scenarios that AVs will find themselves in. Choices about how to design for one scenario are not isolated from design choices for others.

The upshot of this is not pessimism about the need for ethics in AV design, nor that trolley cases are useless for the task. The upshot is that designers and ethicists must be much more careful evaluating the appropriate decision and consider how the technologies at issue relate to the ethical principles and reasoning to be deployed.

We hope this paper motivates a closer working relationship between ethicists and designers of AVs to ensure that the right problems are solved in the right way.

REFERENCES

- Achenbach J. 2015. Driverless cars are colliding with the creepy trolley problem. *Washington Post*, Dec 29.
- Doctorow C. 2015. The problem with self-driving cars: Who controls the code? *The Guardian*, Dec 23.
- Hao K. 2018. Should a self-driving car kill the baby or the grandma? Depends where you're from. *MIT Technology Review*, Oct 24.
- Hübner D, White L. 2018. Crash algorithms for autonomous cars: How the trolley problem can move us beyond harm minimisation. *Ethical Theory and Moral Practice* 21(3):685–98.
- Lin P. 2013. The ethics of autonomous cars. *The Atlantic*, Oct 8.
- Marshall A. 2018. What can the trolley problem teach self-driving car engineers? *Wired*, Oct 24.
- Noothigattu R, Gaikwad SS, Awad E, Dsouza S, Rahwan I, Ravikumar P, Procaccia AD. 2017. A voting-based system for ethical decision making. Available at <https://arxiv.org/abs/1709.06692>.

- Nyholm S, Smids J. 2016. The ethics of accident-algorithms for self-driving cars: An applied trolley problem? *Ethical Theory and Moral Practice* 19(5):1275–89.
- Wallach W, Allen C. 2009. *Moral Machines: Teaching Robots Right from Wrong*. Oxford: Oxford University Press.
- Worstell T. 2014. When should your driverless car from Google be allowed to kill you? *Forbes*, Jun 18.

Influencing Interactions Between Human Drivers and Autonomous Vehicles

DORSA SADIGH
Stanford University

Society is rapidly advancing toward autonomous systems that interact and collaborate with humans—semiautonomous vehicles interacting with drivers and pedestrians, medical robots used in collaboration with doctors, or service robots interacting with their users in smart homes.

A key aspect of safe and seamless interaction between autonomous systems and humans is the ways robots such as autonomous cars can *influence* humans' actions in one-on-one or group settings. This is usually overlooked by the autonomous driving industry, where the common assumption is that humans act as external disturbances like moving obstacles, or that automation can always help societies without actually considering how humans may be impacted.

Humans are not simply a disturbance to be avoided, and they do not always easily adapt to the proliferation of automation in their lives. Humans are intelligent agents with approximately rational strategies who can be influenced and act in novel ways when interacting with other autonomous and intelligent agents.

In this paper I discuss a unifying framework for influencing interactions in autonomous driving—actions of autonomous vehicles (AVs) that can positively influence human-driven vehicles in large-scale or vehicle-to-vehicle (V2V) interactions. Influencing such interactions can be a significant contributor to the safe and reliable integration of AVs.

INFLUENCING INTERACTIONS AT VEHICLE LEVEL

We have designed a novel framework for understanding the interaction between autonomous and human-driven vehicles. We model this interaction as a

dynamical system, where the state of the environment evolves based on the actions of the two vehicles at each time step:

$$x^{t+1} = f(x^t, u_A^t, u_H^t)$$

Here, x^t denotes the state of the environment computed based on the sensor values at each time step including the coordinates, velocity, and heading of each vehicle in the interaction, and the road and lane boundaries. The set of actions of each vehicle u_A^t for the autonomous car and u_H^t for the human-driven car includes steering angle and acceleration.

Our key insight is that the actions of autonomous cars can influence the behavior of human-driven cars on the same road. This can be seen when, for example, a car tries to change lanes: it starts nudging into the destination lane, influencing the cars in that lane to slow down. Similarly, the actions of an autonomous car can result in a human driver changing lanes, slowing down, or speeding up.

Our approach to planning for AV influencing interactions has a few fundamental components. We developed imitation learning techniques¹ to build predictive models of human driving behavior, and designed interaction-aware controllers that model the interaction between a human and a robot as a two-player leader–follower game. Leveraging optimization-based and game theoretic techniques, our work produces robot policies that influence human behavior toward safer outcomes in V2V interaction with autonomous cars (Sadigh et al. 2016a,b, 2018).

Human Driver Models

Imitation learning attempts to learn models of humans by imitating a human expert's demonstrations to enable robots to act in similar ways. Here, we leverage similar techniques, modeling each human driver as an agent who approximately optimizes his or her own objective, referred to as a reward function (e.g., a driver's preferences about avoiding collisions or keeping distance from road boundaries):

$$u_H^* = \arg \max_{u_H} R_H(x, u_H, u_R).$$

We assume that $R_H(x, u_H, u_R) = w \cdot \varphi(x, u_H, u_R)$ represents the human's underlying reward function and that this reward function is a linear combination of a set of hand-coded features $\varphi(x, u_H, u_R)$. These features in the setting of driving can include distances between the AV and the edges of the road, lane boundaries, or other cars, including their velocity and direction.

¹ Imitation learning is a set of algorithms that involves training a robot policy to make decisions based on a collected set of expert demonstrations.

We collect training data in a driving simulator and use them in the form of demonstrations or preferences to learn the parameters w of the reward function using techniques such as maximum entropy inverse reinforcement learning or active preference-based learning of reward functions (Basu et al. 2019; Bıyık and Sadigh 2018; Palan et al. 2019; Sadigh et al. 2016a,b, 2017, 2018).

Planning for Interaction-Aware Controllers

Once we have a predictive human driving model, we can plan for autonomous cars that better interact with humans by being “mindful” of how their actions influence humans. We consider a setting where the autonomous car optimizes for its own reward function:

$$u_R^* = \arg \max_{u_R} R_R(x, u_R, u_H^*).$$

Here, the robot’s reward function directly depends on and influences u_H^* , the learned and predicted human behavior (called human policy).

In game theory, this interaction modeling results in a two-player game between a human-driven and an autonomous car. The actions of the autonomous car influence those of the human-driven car, and vice versa. To efficiently solve this interaction game and plan for AVs, we approximately solve the game as a Stackelberg (leader-follower) game. Our work results in influencing actions by the AV that are more assertive, more efficient, and in many settings safer. Some of these trajectories are shown in figure 1. Our user studies suggest that autonomous cars that are programmed to be aware of their interactions with humans can achieve tasks such as lane changing or coordinating at intersections safely and efficiently (Sadigh et al. 2016a,b, 2018).

INFLUENCING INTERACTIONS AT THE GLOBAL LEVEL

Influencing interactions at the vehicle level can be observed in many driving settings—such as changing lanes, merging, or exiting from a highway—and has substantial effects on the larger traffic system (Bıyık et al. 2018, 2019; Fisac et al. 2019; Lazar et al. 2018; Stefansson et al. 2019). For instance, the presence of a large number of AVs on roads can influence the state of traffic—such as congestion, delay, or flow—and hence human drivers’ routing choices.

We now discuss the challenges arising in mixed-autonomy traffic settings where a large number of autonomous and human-driven vehicles interact.

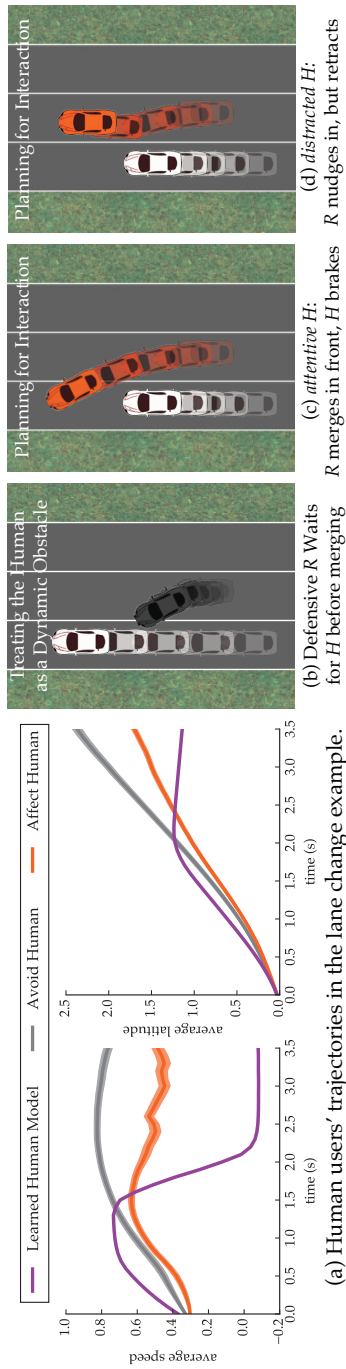


FIGURE 1 Planning for an autonomous car (R) that influences a human-driven car (H) in a lane change maneuver. The graphs in (a) demonstrate the trajectories from our users: orange represents the human trajectories where the autonomous car uses a model of the human, grey represents the human trajectories where the autonomous car does not use a model of the human, and purple represents the simulated learned human model. Figures (b–d) demonstrate the autonomous vehicle's trajectory when changing lanes under the assumption of not having a human model (grey in (b)) or of having a human model (orange in (c–d)). If the autonomous car senses that the human-driven (white) car is distracted, it avoids the merge behavior (d).

Equilibria in Mixed-Autonomy Traffic: Altruistic Autonomy

Traffic congestion has large economic and social costs. The introduction of AVs may reduce congestion both by increasing network throughput and by enabling a social planner to incentivize AV users to take longer routes that can alleviate congestion on more direct roads.²

To formalize the effects of *altruistic autonomy* on roads shared by human drivers and autonomous vehicles we developed a model of road congestion based on a fundamental diagram of traffic (showing the relation between traffic flux [vehicles/hr] and traffic density [vehicles/km]). We considered a network of parallel roads and created algorithms that compute optimal equilibria that are robust to additional unforeseen demand.

Our results show that even with arbitrarily small altruism, total latency can be unboundedly better than without altruism, and that the best selfish equilibrium can be similarly better than the worst selfish equilibrium. We validate our theoretical results through microscopic traffic simulations and show average latency decrease of a factor of 4 from worst-case selfish equilibrium to the optimal equilibrium when AVs are altruistic (Bıyık et al. 2018).

Humans' Routing Choice Models

When users of a road network choose their routes selfishly, the resulting traffic configuration may become very inefficient. Because of this, we consider how to influence human routing decisions so as to decrease congestion on these roads.

We consider a network of parallel roads with two modes of transportation: (i) human drivers who will choose the quickest route available to them, and (ii) a ride hailing service that provides users with an array of AV ride options, each with different prices.

We designed a pricing scheme for the AVs such that when autonomous service users choose from their options and human drivers selfishly choose their routes, road use is optimized and transit delay minimized. To do so, we formalized a model of how autonomous service users make choices between routes with different prices versus delay values.

We developed a preference-based algorithm (similar to our work in learning reward functions discussed above) to learn users' preferences and used a vehicle flow model related to the fundamental diagram of traffic. Based on these, we formulated a planning optimization to support the objective of reduced congestion and demonstrate the benefit of the proposed routing and learning scheme (Bıyık et al. 2019).

² This can be done through pricing schemes or latency management. If, for example, there are two highways to the same destination and one is shorter than the other, drivers will likely select the shorter one, increasing traffic on that route. If a few autonomous cars choose the longer highway, their latency will be lower than that of the congested route—and will also help the latency of the shorter road.

Dynamic Routing in Mixed-Autonomy Traffic

We are developing a social planner by studying a dynamic routing game in which the route choices of autonomous vehicles can be controlled and the human drivers react selfishly and dynamically to the AV actions. As the problem is prohibitively large, we use deep reinforcement learning to develop a policy for controlling the AVs. This policy influences human drivers to route themselves in such a way that minimizes congestion on the network (figure 2).

To gauge the effectiveness of our learned policies, we established theoretical results characterizing equilibria on a network of parallel roads and empirically compared the learned policy results with best possible equilibria. We found that, in the absence of these policies, high demands and network perturbations result in large congestion, whereas using the policy greatly decreases travel times by minimizing congestion.

SUMMARY

We have described our work in planning for influencing interactions in autonomous driving at two levels: (i) vehicle-to-vehicle interaction, in which an autonomous car influences human-driven cars for safer and more efficient driving behavior; and (ii) global-level interaction, in which a large number of autonomous and human-driven vehicles interact in the same traffic network. We design routing decisions for AVs that influence humans' routing choices in order to decrease the total delay of the traffic network for a more desirable societal objective.

Autonomous systems are weaving their way into daily life as robots and the internet of things move into homes and smart cities become a reality. Our long-term goal is to develop a theory for modeling and designing the effects of automation and robotics on human decision making, and this work is a first step toward developing efficient robotics algorithms that lead to safe and transparent autonomous systems as they interact with and influence humans and society.

ACKNOWLEDGMENT

The work discussed is a summary of technical work done in collaboration with Anca Dragan, Erdem Bıyık, Daniel Lazar, Ramtin Pedarsani, S. Shankar Sastry, and Sanjit Seshia.




FIGURE 2 Using deep reinforcement learning (RL) to dynamically route autonomous cars.

REFERENCES

- Basu C, Biyik E, He Z, Singhal M, Sadigh D. 2019. Active learning of reward dynamics from hierarchical queries. *Proceedings, IEEE/RSJ Internatl Conf Intelligent Robots and Systems*, Nov 4–8, Macau, China.
- Biyik E, Sadigh D. 2018. Batch active preference-based learning of reward functions. *Proceedings, 2nd Conf Robot Learning*, Oct 29–31, Zürich.
- Biyik E, Lazar DA, Pedarsani R, Sadigh D. 2018. Altruistic autonomy: Beating congestion on shared roads. *Proceedings, 13th Internatl Workshop on Algorithmic Foundations of Robotics*, Dec 9–11, Mérida.
- Biyik E, Lazar DA, Sadigh D, Pedarsani R. 2019. The green choice: Learning and influencing human decisions on shared roads. *arXiv:1904.02209v2*.
- Fisac JF, Bronstein E, Stefansson E, Sadigh D, Sastry SS, Dragan AD. 2019. Hierarchical game-theoretic planning for autonomous vehicles. *Internatl Conf Robotics and Automation*, May 20–24, Montréal.
- Lazar DA, Chandrasekher K, Pedarsani R, Sadigh D. 2018. Maximizing road capacity using cars that influence people. *Proceedings, 57th IEEE Conf Decision and Control*, Dec 17–19, Miami Beach.
- Palan M, Landolfi NC, Shevchuk G, Sadigh D. 2019. Learning reward functions by integrating human demonstrations and preferences. *Proceedings, Robotics: Science and Systems*, Jun 12–17, Corvallis.
- Sadigh D, Sastry S, Seshia SA, Dragan AD. 2016a. Planning for autonomous cars that leverage effects on human actions. *Proceedings, Robotics: Science and Systems*, Jun 18–22, Ann Arbor.
- Sadigh D, Sastry SS, Seshia SA, Dragan A. 2016b. Information gathering actions over human internal state. *Proceedings, IEEE/RSJ Internatl Conf Intelligent Robots and Systems*, Oct 9–14, Daejeon, Korea.
- Sadigh D, Dragan AD, Sastry S, Seshia SA. 2017. Active preference-based learning of reward functions. *Proceedings, Robotics: Science and Systems*, Jul 12–16, Cambridge MA.
- Sadigh D, Landolfi N, Sastry SS, Seshia SA, Dragan AD. 2018. Planning for cars that coordinate with people: Leveraging effects on human actions for planning and active information gathering over human internal state. *Autonomous Robots* 42:1405–26.
- Stefansson E, Fisac JF, Sadigh D, Sastry SS, Johansson KH. 2019. Human-robot interaction for truck platooning using hierarchical dynamic games. *European Control Conf*, Jun 25–28, Naples.

BLOCKCHAIN TECHNOLOGY

Blockchain Technology

PETR NOVOTNY
IBM T.J. Watson Research Center

ELAINE (RUNTING) SHI
Cornell University

Since its conception by Satoshi Nakamoto in 2008 and first implementation in Bitcoin, the popularity of blockchain and interest in it are rapidly accelerating. As the technology that underlies implementation of Bitcoin and other applications, blockchain is designed as a decentralized network of peers who collectively share and manage a distributed ledger structured as a series of ordered and cryptographically connected blocks, each containing a list of transactions. New transactions are created by smart contracts that embody an application-specific logic (e.g., the transfer of money between accounts or creation of a shipping record). Before inclusion in the ledger, new transactions and blocks must be approved by the network participants through a consensus mechanism. Once included in the distributed ledger, the transactions are immutable, time-stamped, and retrospectively verifiable by any network participant. With the use of these mechanisms, blockchain provides data availability, transparency, and digital trust unparalleled by other systems.

The initial Bitcoin network inspired many new blockchain applications with extended capabilities and uses in a range of areas. However, despite advances in recent years, many fundamental challenges remain unresolved and are the subject of intense scientific research and technology development. Additionally, the trust that blockchain provides opens new opportunities for applications in governance and economy as well as social, healthcare, and other sectors that are being actively investigated.

In this session, Elaine Shi introduced the history and key concepts of blockchain and provided an overview of the major platforms and applications, including

Bitcoin, Ethereum, and Hyperledger.¹ Next, Hong Wan (North Carolina State University) discussed the domain of private and permissioned blockchain platforms, such as Hyperledger Fabric and Corda, designed as building blocks of networks among consortiums of enterprises, as well as the advantages, threads, and weaknesses of these platforms. This talk was followed by Jacob Leshno (University of Chicago), who discussed the use of blockchain technologies in cryptocurrencies.

¹ Paper not included in this volume.

Blockchain Beyond Cryptocurrency: An Overview

HONG WAN
North Carolina State University

In the past 2½ years, blockchain has evolved from a computer science term to a buzzword in the ongoing hype about cryptocurrency, for which the blockchain is the foundation. Articles have even praised the blockchain as “the most important invention after the Internet” (Metry 2017).

One topic of discussion is how to apply blockchains beyond cryptocurrency. Like all emerging technology, there has been much confusion about and misunderstanding of the blockchain concept. One popular misconception is that “blockchains are all like the ones used in Bitcoin.” This is wrong. In this paper I describe various kinds of blockchains and explain why more centralized blockchain structures are more appropriate for business use.

WHAT IS A BLOCKCHAIN?

A blockchain is a digital, append-only, time-stamped ledger. It is a consensus-based, peer-to-peer distributed network with “a growing list (chain) of records, called blocks, which are linked using cryptography. Each block contains a unique hash value of the previous block, a timestamp, and transaction data.”¹ Here, consensus refers to a set of rules that users follow to agree on the states of the system; it makes the blockchain a self-auditing ecosystem.² Hash is a cryptographic function that converts a string into a nonmeaningful, fixed-length output. It is nonreversible since the hash value is highly sensitive to the input—a small change in input leads to a completely different hash value.

¹ <https://en.wikipedia.org/wiki/Blockchain>.

² <https://lisk.io/academy/blockchain-basics/how-does-blockchain-work/consensus-protocols>.

The mathematical structure of the blockchain implies two essential properties. First, the data (in block) are immutable. Specifically, if a block is changed, all blocks before it become invalid since all hash values in these blocks become invalid. Second, a distributed network with consensus allows users to communicate directly with each other to broadcast a new block and synchronize the blockchain status. All users can download a copy of the current ledger and add blocks, which means that there is redundancy of the data in the network. Therefore, the blockchain is more tolerant of node failures. With these two properties, the longer the chain and more users (nodes) in the network, the harder to hack into the chain and change blocks without detection, making the blockchain more reliable (Nakamoto 2008).

The first work on building a cryptographically secured chain of blocks was proposed almost 30 years ago by Stuart Haber and W. Scott Stornetta (1991). The concept was formalized in 2008 by Satoshi Nakamoto,³ who proposed “A purely peer-to-peer version of electronic cash that would allow online payments to be sent directly from one party to another without going through a financial institution...but *through a system that [is] based on cryptographic proof instead of trust*” (emphasis added; Nakamoto 2008). In January 2009, Nakamoto mined the first *bitcoin* and started the era of cryptocurrency.

In recent years, big tech companies like IBM and Microsoft as well as many startups have put significant effort into extending blockchain systems to various industries, many of which use blockchains with a significantly different structure compared to cryptocurrency, as elaborated below.

PUBLIC OR PRIVATE, PERMISSIONED OR PERMISSIONLESS BLOCKCHAINS

From a governance point of view, the blockchain can be public (open) or private (closed). In a public chain, anyone can initiate transactions, generate and broadcast blocks, and download a copy of the whole ledger. In a private chain, only authorized users can access the network.

The blockchain can also be permissioned or permissionless. The permissioned chain means the rights/authorization of users can be different; some have more authorities (e.g., to validate the block) than others (read-only access). The permissionless chain means that all users have the same rights: anyone in the network can download the ledger, generate blocks, and validate transactions.

While many consider permissioned and private the same, they are, in fact, two concepts. A blockchain that is public or private is independent of whether it is permissioned or not. As explained in testimony to Congress (Jaikaran 2017), “Discussing a blockchain as public or private refers to the level of freedom users have to create identities and read data on that blockchain. Discussing a blockchain

³ The identity of Satoshi Nakamoto remains a mystery.

as permissioned or permissionless refers to the level of access the user would have on that blockchain.” Specifically, the public and permissionless chain has the most decentralized structure and assumes no trust among users. The private and permissioned chain, on the other hand, has the most centralized architecture and the highest level of trust, and usually does not need tokens/coins. The other two combinations are hybrid. Each kind of blockchain has unique properties that are suitable for different applications, as elaborated below.

Public and Permissionless Chains

A majority of cryptocurrencies use public and permissionless chains, which are the most well-known blockchain type. To understand how this type of chain works, I compare how financial transactions are handled by a traditional bank and by a blockchain. Suppose Alice wants to transfer \$10 to Bob. After she initiates the transfer, someone needs to check that her account has enough money, deduct \$10 from it, and deposit the \$10 to Bob’s account. In traditional banking, this is done in a centralized database controlled by the bank. In a cryptocurrency wallet (e.g., a bitcoin wallet), there is no trusted third party. The transaction instead is broadcasted to the whole network in the following format (the address in this sample is generic; Kadiyala 2018):

*15N3yGu3UFHeyUNdzQ5sS3aRFRzu5Ae7EZ sent 0.00086 bitcoin to
1JHG2qjdk5KHiq7X5xQrr1wfigepJEK3t on August 8th, 2019, between
11:10 and 11:20 a.m.*

Users compete to validate the transaction for rewards⁴ through the mining process, grouping transactions in a specific period to form a block. Whoever wants to post the block needs to solve a computationally intensive problem. For each period (e.g., 10 minutes for Bitcoin), only the first one to solve the problem can publish the block and claim the rewards.

Proof of Work Consensus

There are variations of these rules in other chains, but the concept is similar. This *proof of work* consensus causes most of the confusion and criticism of the blockchain. Because of the difficulty of the nonmeaningful problems solved, the mining process both significantly slows the transaction speed and consumes enormous computational power and energy. The Bitcoin blockchain can currently guarantee only 4.6 transactions per second (10 minutes per block), compared to Visa at 1,736 transactions per second, making the bitcoin process too slow for

⁴ The reward is usually the cryptocurrency of the network. For example, Bitcoin system offers 12.5 bitcoin for each block mined.

everyday use. This is also called “poor scalability.” And it is estimated that “the global Bitcoin network is consuming more [electricity] ... than the country of Switzerland uses over the same time period” (Vincent 2019).

Why is proof of work necessary for a public and permissionless chain? Because the assumption is that malicious users will try to spam the network with fraudulent blocks and modify the existing chain to add self-benefiting transactions. The proof of work consensus makes sure that it is expensive to create blocks (in terms of both time and energy), so it is difficult and expensive to generate many blocks to flood the system. Also, given the 10-minute waiting period for each block, anyone who generates fake blocks is likely to be caught and rejected by other users. Users are thus incentivized to spend their computational power on legitimate blocks so their work can be compensated. Besides, transactions within blocks are kept secure by the amount of energy spent on mining blocks before them⁵: the longer the chain, the more expensive for hackers to replace it with a fake one.

The proof of work consensus allows a truly distributed network, maximum number of pseudonymous users, fairness among nodes, and chain nontamperability, with efficiency and reasonable cost in computational power.

Proof of State Consensus

Another popular, more recent consensus for the public and permissionless chain is *proof of state*, which attributes mining power to the proportion of coins held by a miner (Young 2016). This consensus significantly reduces the computational power and time required to add blocks. The logic is clear: those with more coins have less motivation to sabotage the chain. On the other hand, a system in which the major stakeholder enjoys extensive control and authority over both technical and economic aspects of the network creates a monopoly problem. In addition, block generators lose nothing by voting for multiple versions of chains (the “nothing-at-stake” problem). Because of this, “some cryptocurrencies are vulnerable to Fake Stake attacks, where an attacker uses no or minimal stake to crash an affected node.”⁶

Even with recent improvements, transaction validation speeds for cryptocurrencies based on point of sale are still not on par with traditional systems.

⁵ Public versus private (permissioned) blockchain comparison. DevTeam.Space (<https://www.devteam.space/blog/public-vs-private-permissioned-blockchain-comparison/>).

⁶ “Fake Stake” attacks on chain-based Proof-of-Stake cryptocurrencies. Medium/Cryptocurrency. (https://medium.com/@dsl_uiuc/fake-stake-attacks-on-chain-based-proof-of-stake-cryptocurrencies-b8b05723f806).

Private and Permissioned Chains

For many businesses and organizations, blockchain is attractive as a record-keeping and sharing system. However, the slow transaction validating speed, poor scalability, and lack of data privacy make the public and permissionless chain a poor choice. For healthcare records, for example, the network should be neither public (it must be Health Insurance Portability and Accountability Act protected) nor permissionless (only authorized people should be allowed to add records). In these cases, the private and permissioned chain is a good option as it allows only authorized users with different levels of authority in the chain.

The IBM Hyperledger Fabric (Mamun 2018) is a blockchain network set by various collaborative organizations called “members.” Each member selects the peers from his or her organization to involve in the blockchain. These peers have different authorizations: *endorsers* validate transactions and decide to approve or disapprove transactions, *anchors* are in charge of broadcasting updates, and *orderers* are in charge of creating and delivering blocks to all the peers. In this network there is intrinsic trust both among peers (since they are from the same organization) and among members since they are collaborating and only dedicated peers who organizations trust can validate transactions and add blocks. There is usually no incentive for malicious behaviors. This consensus is called *proof of authority*, as a number of nodes are “authorities” in charge of validating transactions.

The other popular consensus are a *round-robin scheme*, where users on the network take turns adding new blocks, and *proof of elapsed time*, where each node is assigned a random waiting period and the first node to reach the elapsed time gets to create the next node.

The main criticism of the private and permissioned chain is that it is not a real blockchain. Its multicentralized control system (members) defeats the decentralization property of the blockchain. Many of these chains also do not need a token/coin to incentivize participants, which means that there is no mining. In these cases, the question becomes, Why not just use a distributed database? This is a legitimate question. If a business has critical data that it wants to share internally, a combination of the current database, cloud, and identity management technologies will likely be adequate for its needs. However, if several organizations seek to collaborate and want the data to be immutable and auditable to avoid data discrepancies, a blockchain is probably more appropriate and convenient (e.g., for a supplier–vendor relationship or an insurance company–clinic–patient relationship).

Private and Permissionless, Public and Permissioned Chains

At the two ends of the spectrum there are private and permissionless chains and public and permissioned chains. Neither type of chain has been widely implemented, and both are mainly at conceptual levels.

The private and permissionless chain would allow anybody in the network to submit and process transactions but would control who can be involved in the chain. It could be used to, for example, handle government records, compile research results from different teams, or maintain data privacy.

The public and permissioned chain would allow anyone to access the ledger, but only verified parties to submit, process, and validate transactions. Public and permissioned chains emphasize who can write in/regulate data and can be applied to, for example, real estate registries, diploma checking system, or other scenarios involving regulation or protection from forgery. The categorization can be expanded with partial permissions and a combination of public and private schemes. Why is this interesting? I discuss it in the next session.

CHALLENGES OF SCALABILITY, SECURITY, AND DECENTRALIZATION

From the system design point of view, there are three major characteristics of blockchains: scalability, security, and decentralization. Scalability refers to the ledger's ability to handle growth, security denotes attack resistance, and decentralization refers to the network's transparency, synchronization, and fairness. The trade-off among these three is called the *scalability trilemma*: it is hard to maximize two without sacrificing the third. A useful figure for illustrating this concept can be found at <https://steemit.com/blockchain/@reverseacid/the-scalability-trilemma>.

Blockchain mechanism design must find a specific (usually hybrid) blockchain structure based on the requirements of the application. This goal motivates my laboratory to seek a better understanding of the blockchain as a complex system, quantifying and modeling its key features and performance, and simulating various structures for a structural mechanism design paradigm.

CONCLUSION

Blockchain is not for everything. In many (if not most) cases, as discussed above, a distributed database with access control is more than adequate, and is much faster and cheaper. Caution is appropriate amid the blockchain hype. For example, is “blockchaining the system” necessary? What problems does it solve? Are there options?

On the other hand, the blockchain is a tool for big data and artificial intelligence (AI). It may serve as the *neural network* that connects data and AI by validating, auditing, and sharing data safely and, if wanted, anonymously. It is especially promising when combined with internet of things (IoT) and wearable sensors to collect and distribute data automatically, without individuals needing to worry about misuse of their data by big companies. More solid business and social applications of blockchains are in the future.

REFERENCES

- Haber S, Stornetta WS. 1991. How to time-stamp a digital document. *Journal of Cryptology* 3(2):99–111.
- Jaikaran C. 2017. Consumer data security and the credit bureaus. Testimony before the Senate Committee on Banking, Housing, and Urban Affairs, Oct 17. Available at <https://www.banking.senate.gov/download/jaikaran-testimony-10-17-17pdf>.
- Kadiyala A. 2018. Nuances between permissionless and permissioned blockchains. Medium, Feb 17. Available at <https://medium.com/@akadiyala/nuances-between-permissionless-and-permissioned-blockchains-f5b566f5d483>.
- Mamun M. 2018. How does Hyperledger Fabric work? Medium, Apr 17. Available at <https://medium.com/coinmonks/how-does-hyperledger-fabric-works-cdb68e6066f5>.
- Metry M. 2017. Blockchain technology is the most significant invention since the Internet and electricity. Medium, Sep 2.
- Nakamoto S. 2008. Bitcoin: A peer-to-peer electronic cash system. Bitcoin.org. Available at <https://bitcoin.org/bitcoin.pdf>.
- Steemit. 2018. The scalability trilemma. Available at <https://steemit.com/blockchain/@reverseacid/the-scalability-trilemma>.
- Vincent J. 2019. Bitcoin consumes more energy than Switzerland, according to new estimate. The Verge, Jul 4.
- Young J. 2016. Proof-of-work vs. proof-of-stake: Merits and disadvantages. Coinfox, Sep 14. Available at <http://www.coinfox.info/news/reviews/6417-proof-of-work-vs-proof-of-stake-merits-and-disadvantages>.

Cryptocurrencies as Marketplaces

JACOB LESHNO
University of Chicago

Bitcoin was introduced in 2008 as a computer protocol establishing a decentralized system that allows users to hold balances and make transfers to one another (Nakamoto 2008). Computer systems that provided similar services have existed for decades, but required a trusted party to control and operate them. For example, PayPal Holdings Inc. maintains the required computer infrastructure and charges usage fees to fund its activities and make a profit.

BITCOIN INFRASTRUCTURE AND PROTOCOL

Bitcoin is a decentralized system. Instead of having a company that is responsible for maintaining the system's infrastructure, it is operated by a decentralized network of computers called miners (the term also refers to the people who operate these computers). Much like Uber and Lyft, which allow anyone with a car to provide transportation services in return for compensation, Bitcoin allows anyone with a computer to provide the payment processing infrastructure in return for compensation. It eliminates the need for a centralized infrastructure by creating an open marketplace.

But Bitcoin is unlike Uber and Lyft in that no entity is in control of the marketplace. Uber can change the price paid to drivers, add or remove the option to tip, and charge fees from the participants in its market. In contrast, Bitcoin is governed by its protocol, which no single entity can change (making changes to Bitcoin is akin to changing communication protocols such as TCP/IP). The computer protocol dictates the rules that govern the system and its implied marketplace, determining how miners are compensated and the fees users pay.

The viability and success of Bitcoin, and other cryptocurrencies that followed (e.g., Litecoin, Ethereum, Dogecoin), require that the protocol establish a functioning marketplace. But cryptocurrencies cannot control the miners who provide the infrastructure, and incentives are required to get miners to follow a desired behavior (Carlsten et al. 2016; Eyal and Sirer 2018). Miners provide their services at will and can withdraw from the system at any time, or try to exploit the system for profit and jeopardize its security (Auer 2019; Budish 2018). Game theory provides tools to understand how miners and users will behave in such an environment and to determine whether the system is secure.

Since Bitcoin is not integrated with (and does not wish to rely on) other financial services, payments to miners can be made only with the system's native coin, bitcoin, whose value is determined by financial markets, raising questions from monetary theory (Schilling and Uhlig 2019). These elements and others differentiate cryptocurrencies from traditional computer systems and make them economic objects, akin to marketplaces.

MONOPOLY WITHOUT A MONOPOLIST

In my work with Gur Huberman and Ciamac Moallemi (Huberman et al. 2019) we study the properties of this marketplace for transaction processing and ask who pays for the costs of operating the platform, how, and how much. We compare the Bitcoin payment system (BPS) with a traditional payment system (e.g., PayPal) and ask whether the decentralized design offers new benefits. (While we focus on Bitcoin's design, our analysis also applies to other cryptocurrencies with similar design features.)

The system processes transactions in batches called blocks. To ensure that a block is propagated throughout the network before the next one is issued, the protocol limits block size and frequency, limiting the system's transaction processing capacity. Because of stochastic elements in the system, the system can periodically get congested and transactions can be delayed.

We observe that the blockchain design of the BPS has the following features, which are key elements of its economics:

- Miners can enter or leave the system as they see fit.
- Miners can select the transactions they process and are rewarded with protocol-determined block rewards and transaction fees offered by the users (because Bitcoin's protocol specifies that block rewards are halved approximately every 4 years, transaction fees will become more important over time, as they will eventually be the only form of payment to miners).
- The system's transaction processing capacity (that is, average number of transactions that can be processed per unit time) is independent of the number of miners.

- Users choose the transaction fees they pay when their transaction is processed; they may even choose to pay only a minor minimal transaction fee (but when the system is sufficiently congested, the delay can be so long that payment is required to get through before a transaction is timed out).

We offer a simplified economic model of the BPS that allows analysis of the implied marketplace based on the following: (i) some users are willing to pay to expedite the processing of their transactions, (ii) miners are profit maximizers, and (iii) miners can freely enter or exit the system.

TRANSACTION FEES ARE DETERMINED IN EQUILIBRIUM

We find that the BPS is well described by an equilibrium in which users choose a transaction fee to gain processing priority over other users; miners process the transactions that offer the highest fee, up to capacity. Nobody dictates the equilibrium fee schedule. Transaction fees are set in an implicit auction without any explicit auctioneer.

We offer closed-form expressions for the equilibrium fees and waiting times. We find that total transaction fees depend on three parameters: maximal block size, congestion or load (transaction arrival rate divided by system's capacity), and the distribution of user willingness to pay higher fees to reduce transaction processing delay.

When the system is not congested, the fees are low and essentially insensitive to its use—the expected processing delay is similar across transactions. As the system's use approaches capacity, fees and cross-transaction variation in processing delays rise rapidly. The fee schedule satisfies the Vickrey–Clarke–Groves property: each transaction fee is equal to the externality it imposes by increasing the delay for transactions that offer lower fees.

COMPARISON WITH A PROFIT-MAXIMIZING FIRM

Pricing under the BPS is structurally different from the pricing of a profit-maximizing firm. A firm sets a price and denies service to users who are unwilling to pay that price. When the BPS has sufficient capacity, the system can raise revenue without denying service to anybody; users who are willing to bear delays can have their transaction processed even without paying transaction fees.

Because the miners who collectively operate the system compete with each other, they cannot profitably affect the level of fees paid by users. This provides users protection from price increases: even if the system becomes a monopolist (in the sense that users have no alternative payment methods) users will still pay a low competitive transaction fee. In that way, the decentralized nature of the system may provide economic benefits to users.

However, the design has several weaknesses.

- Transaction processing delays are essential to fee generation and therefore to the BPS's long-run revenue model.
- The amount of infrastructure consumed by the system is determined in equilibrium, and there is no mechanism that ensures an efficient level of infrastructure.
- The amount of energy consumed by Bitcoin has received much media attention. It varies depending on the demand for transactions and the bitcoin-to-USD exchange rate, and the system design does not indicate either a desired level or a way of reaching such a level.

DESIGN SUGGESTIONS

We provide a design that can partly address these concerns. It modifies a component of the protocol so that instead of maintaining a constant capacity, the protocol scales capacity according to demand (within a feasible region) to maintain congestion at a moderate level. This ensures that total transaction fees and the level of infrastructure are kept at a constant level. Our analysis also indicates that smaller block sizes allow the system to raise revenue more efficiently: a smaller block size allows the system to raise the same amount of revenue with shorter transaction processing delays.

GOVERNING A DECENTRALIZED SYSTEM

The limitations of the Bitcoin protocol have motivated much research and the development of other decentralized systems (e.g., Bentov et al. 2016; Chen and Micali 2016; Poon and Buterin 2017). To update such systems, agreement is needed on a new protocol, but without an entity that controls the system such agreement can be difficult to achieve.

The implied rigidity of the system can be advantageous to users, who are guaranteed continuation of service at the same terms (with no ratcheting of fees), but it also reduces the system's ability to react to new circumstances, which is especially important given the early stage of the technology. Game theoretic analysis can shed light on governance issues and help in the design of systems accordingly (Barrera and Hurder 2018).

CONCLUSION

Through a combination of cryptographic tools and economic incentives, Bitcoin and its followers have shown that it is feasible to create a global decentralized system controlled by no one. Services that previously could be provided only by a trusted firm can now be provided by a community coordinated only by

a protocol. This allows for new economic models for the operation and funding of such services. The interdisciplinary nature of these systems calls for exciting future collaboration.

REFERENCES

- Auer R. 2019. Beyond the Doomsday Economics of “Proof-of-Work” in Cryptocurrencies. BIS Working Paper No. 765. Basel: Bank for International Settlements.
- Barrera C, Hurder S. 2018. Blockchain upgrade as a coordination game. Prysm Group.
- Bentov I, Pass R, Shi E. 2016. Snow White: Provably secure proofs of stake. IACR Cryptology ePrint Archive 2016:919.
- Budish E. 2018. The Economic Limits of Bitcoin and the Blockchain. NBER Working Paper No. 74717. Cambridge MA: National Bureau of Economic Research.
- Carlsten M, Kalodner H, Weinberg SM, Narayanan A. 2016. On the instability of bitcoin without the block reward. Proceedings, 2016 ACM SIGSAC Conf Computer and Communications Security, Oct 24–28, Vienna.
- Chen J, Micali S. 2016. Algorand. arXiv:1607.01341.
- Eyal I, Sirer EG. 2018. Majority is not enough: Bitcoin mining is vulnerable. Communications of the ACM 61(7):95–102.
- Huberman G, Leshno JD, Moallemi CC. 2019. An economic analysis of the Bitcoin payment system. Columbia Business School Research Paper No. 17-92.
- Nakamoto S. 2008. Bitcoin: A peer-to-peer electronic cash system. Available at bitcoin.org.
- Poon J, Buterin V. 2017. Plasma: Scalable autonomous smart contracts. Available at <https://www.plasma.io/plasma.pdf>.
- Schilling L, Uhlig H. 2019. Some Simple Bitcoin Economics. NBER Working Paper No. 24483. Cambridge MA: National Bureau of Economic Research.

APPENDIXES

Contributors

Omar Akbari is an associate professor of biological sciences at the University of California, San Diego, where he develops animal population control technologies.

John Basl is an associate professor of philosophy and religion at Northeastern University and works primarily in applied or practical ethics, including the ethics of artificial intelligence, the ethics of emerging technologies, and environmental ethics.

Patrick Boyle is the head of codebase at Ginkgo Bioworks, where he leads a growing portfolio of reusable biological code that builds highly automated foundries to rapidly design, build, and test engineered organisms.

Gabriel Burnett is an engineering manager for Boeing Research and Technology in Seattle, where he develops and applies machine learning, simulation, optimization, and model-based engineering to improve the performance of platforms, production systems, and aftermarket services.

Li Chang is a technical fellow for The Boeing Company, specializing in high-intensity radiated fields, lightning strike protection, advanced manufacturing, and innovation systems architecture.

Tae Eun Choe is the principal engineer at NVIDIA, where he focuses on computer vision, machine learning, and autonomous driving.

Tarik Dickens is an assistant professor of industrial and manufacturing engineering at Florida A&M University and Florida State University, where he researches composite structure design, intelligent processing, additive and automated processing, and mechanical testing with a focus on failure prognosis with interconnected sensors.

Charles Gersbach is an associate professor of biomedical engineering at Duke University, where he researches genome and epigenome editing, gene therapy, regenerative medicine, biomolecular and cellular engineering, synthetic biology, optogenetics, genomics, and epigenomics. His laboratory applies molecular and cellular engineering to develop new methods to genetically modify genome sequences and cellular gene networks in a precise and targeted manner.

Christoffer Heckman is an assistant professor of computer science at the University of Colorado Boulder, where he focuses on autonomy, robotic perception, field robotics, self-driving cars, machine learning, and artificial intelligence. He develops practical and explainable techniques in probabilistic artificial intelligence. His work in computer vision, machine learning, and sensor fusion has applications to intelligence, defense, and environmental monitoring.

Christian Hubicki is an assistant professor of mechanical engineering in the Aero-propulsion, Mechatronics, and Energy Center at Florida State University, where he specializes in bipedal locomotion, specifically optimization methods that apply to both legged robotics and biomechanics.

Pamela Kobryn is a principal aerospace engineer for the structures technology branch of the Air Force Research Laboratory, where she researches hybrid physics-/data-informed modeling, simulation, and analysis of aircraft performance with a focus on structural integrity.

Christopher Lang is an aerospace engineer in structural mechanics and concepts at the National Aeronautics and Space Administration Langley Research Center, where he focuses on computational modeling for the design and certification of additively manufactured components for aerospace applications.

Jacob Leshno is an assistant professor of economics at the University of Chicago, where he works on market design, matching markets, and economic theory. He researches the design of school choice procedures that guide the allocation of students to schools, the assignment of patients to nursing homes, and the design of decentralized cryptocurrencies.

Samantha Maragh serves as leader of the genome editing program at the National Institute of Standards and Technology, where she focuses on mea-

surement assurance and technology development for genome editing, bioassay and biomarker validation, with applications in engineering biology, precision/regenerative medicine, and cancer biology.

Petr Novotny is a research staff member of enterprise solutions at IBM's Thomas J. Watson Research Center, where he leads development of blockchain-based analytical, audit, and compliance methods and tools and seminal blockchain-based industry use cases in finance, trading, media, and other areas.

Joan Robinson-Berry is vice president and chief engineer for Boeing Global Services. As Global Services' engineering leader, Robinson-Berry is responsible for strengthening Boeing's enterprise engineering expertise and encouraging a companywide approach to meeting customer, business, and operational priorities, with an emphasis on safety, quality, and integrity.

Dorsa Sadigh is an assistant professor of computer science and electrical engineering at Stanford University, where she designs algorithms for robots to enable safe and reliable interaction with humans.

Krishanu Saha is an associate professor of biomedical engineering at the University of Wisconsin–Madison and the Wisconsin Institute for Discovery, where he researches cell and genome engineering with patient-derived cells. His laboratory brings together genomics, stem cell biology, immunology, materials science, and bioethics for the development of novel therapeutics targeting the eye, blood, and brain.

Hae-Jong Seo researches autonomous driving, artificial intelligence, and computer vision for NVIDIA's autonomous driving research and development group.

Elaine Shi is an associate professor of computer science at Cornell University, where she designs and builds systems that are efficient, provably secure, and friendly to users and programmers. She takes an interdisciplinary approach that blends cryptography, systems security, language-based security, distributed systems, randomized algorithms, programming languages, and blockchains.

Hong Wan is an associate professor of industrial and systems engineering at North Carolina State University, where she focuses on the design and analysis of complex simulation experiments, statistical learning methods (feature selection), blockchain simulation, data analysis, mechanism design, healthcare engineering, and quality management.

Renee Wegrzyn is a program manager in the biological technologies office of the Defense Advanced Research Projects Agency, where she applies the tools of syn-

thetic biology and genome engineering to support biosecurity, outpace infectious disease, support and grow the bio-economy, and enable human survival in extreme environments, including space.

Jennifer West is the Fitzpatrick Family University Professor of Biomedical Engineering at Duke University, where she focuses on the design of novel bio-functional materials for medical applications ranging from coronary artery bypass grafting to cancer therapy. This includes research in tissue engineering, through which synthetic polymer design can mimic functions of the natural extracellular matrix, thus providing cues to support and guide the formation of new tissues in the laboratory to use for transplantation.

Participants

Omar Akbari**
Assistant Professor
Division of Biological Sciences
University of California, San Diego

Brenna Argall
Associate Professor
Departments of Mechanical
Engineering, Computer Science,
and Physical Medicine and
Rehabilitation
Northwestern University

Oznur Arslan
R&D Manager
Air Separation Cryogenics Technology
Air Products

John Basl**
Associate Professor
Department of Philosophy and Religion
Northeastern University

Ryan Beams
Staff Scientist
Division of Imaging, Diagnostics, and
Software Reliability
US Food and Drug Administration

Agnes Blom-Schieber
Associate Technical Fellow - Principal
Structural Analysis Engineer
Product Strategy and Future Airplane
Development
Boeing Commercial Airplanes

Paul Bogdan
Assistant Professor
Department of Electrical and
Computer Engineering
University of Southern California

Darryl Boyd
Research Chemist
Optical Sciences Division
Naval Research Laboratory

*Organizing Committee Member

**Speaker

Patrick Boyle**
Head of Codebase
Ginkgo Bioworks

Emma Buneci
Principal Technical Program Manager
Robotics AI
Amazon

Gabriel Burnett**
Engineering Manager
Boeing Research and Technology
Boeing South Carolina

Steven Callender
Research Scientist
Wireless Components Research
Intel Corporation

Roman Caudillo
Senior Research Scientist
Tech Manufacturing Group,
Components Research, and
Quantum Computing
Intel Corporation

Li Chang*
Technical Fellow
Boeing Commercial Airplanes
The Boeing Company

Theodora Chaspari
Assistant Professor
Department of Computer Science and
Engineering
Texas A&M University

Lily Cheung
Assistant Professor
School of Chemical and Biomolecular
Engineering
Georgia Institute of Technology

Tae Eun Choe**
Chief Architect
Intelligent Driving Group
Baidu USA

Srabanti Chowdhury
Associate Professor
Department of Electrical Engineering
Stanford University

Eun Ji Chung
Assistant Professor
Department of Biomedical
Engineering
University of Southern California

Benjamin Cohen
Principal Scientist
Chemical and Synthetic Development
Engineering
Bristol-Myers Squibb

Nicolle Correa
Senior Manager, Applied Science
Core AI
Amazon

Ayse Coskun
Associate Professor
Department of Electrical and
Computer Engineering
Boston University

Danielle Cote
Assistant Professor
Department of Mechanical
Engineering
Worcester Polytechnic Institute

Andrew Coughlin
Formulation Team Lead
Formulation Technology Group
Syngenta

Laura Coyle
Senior Optical Engineer
Optical and Payload Systems
Engineering
Ball Aerospace

Simone D'Amico
Assistant Professor
Department of Aeronautics and
Astronautics
Stanford University

Canan Dagdeviren
Assistant Professor
Media Lab
Massachusetts Institute of Technology

Kathryn Daltorio
Assistant Professor
Department of Mechanical and
Aerospace Engineering
Case Western Reserve University

Emily Day
Assistant Professor
Department of Biomedical
Engineering
University of Delaware

Cole DeForest
Dan Evans Career Development
Assistant Professor
Department of Chemical Engineering
and Bioengineering
University of Washington

Tarik Dickens*
Assistant Professor
Department of Industrial and
Manufacturing Engineering
Florida A&M University and Florida
State University

Marine Durel
Avionics Application Software
Engineer
Engineering Test and Technology
Boeing

Rose Faghih
Assistant Professor
Department of Electrical and
Computer Engineering
University of Houston

Karen Feigh
Associate Professor
Department of Aerospace Engineering
Georgia Institute of Technology

Catherine Fromen
Assistant Professor
Department of Chemical and
Biomolecular Engineering
University of Delaware

Charles Gersbach*
Associate Professor
Department of Biomedical Engineering
Duke University

Krishnanjan Gubba Ravikumar
Principal Research Engineer
Research and Development
Schweitzer Engineering Laboratories

A. Kate Gurnon
Lead Engineer
Characterization Laboratory
GE Research

Amberlee Haselhuhn
Senior Researcher
Research and Development
General Motors

Marta Hatzell
Assistant Professor
Woodruff School of Mechanical
Engineering
Georgia Institute of Technology

Yeganeh Hayeri
Assistant Professor
School of Systems and Enterprises
Stevens Institute of Technology

Christoffer Heckman*
Assistant Professor
Department of Computer Science
University of Colorado Boulder

Moneer Helu
Life Cycle Engineering Group Leader
Engineering Laboratory
National Institute of Standards and
Technology

Cathleen Hoel
Senior Scientist
Structural Materials
GE Research

Yongjie Hu
Assistant Professor
Department of Mechanical and
Aerospace Engineering
University of California, Los Angeles

Christian Hubicki**
Assistant Professor
Aero-Propulsion, Mechatronics, and
Energy Center
Florida State University

Heather Hussain
Flight Guidance, Navigation, and
Controls Engineer
Controls and Autonomy/Engineering
Test and Technology
Boeing Corporation

Shirin Jalali
Research Scientist
Mathematics of Communications
Research Department
Nokia Bell Labs

Deep Jariwala
Assistant Professor
Department of Electrical and Systems
Engineering
University of Pennsylvania

Edward Jimenez
R&D S&E, Computer Science
Mission Algorithms Research and
Solutions
Sandia National Laboratories

Ece Kamar
Principal Researcher, AI
Microsoft Research

Sertac Karaman
Associate Professor
Department of Aeronautics and
Astronautics
Massachusetts Institute of Technology

Yasaman Khazaeni
Research Staff Member and Manager
AI Research Group
IBM Research

Pamela Kobryn**
Principal Aerospace Engineer
Structures Technology Branch
Air Force Research Laboratory

Christopher Lang**
Aerospace Engineer
Structural Mechanics and Concepts
NASA Langley

Jacob Leshno**
Assistant Professor
Department of Economics
University of Chicago

Jamal Lewis
Assistant Professor
Department of Biomedical
Engineering
University of California, Davis

Richard Linares
Charles Stark Draper Assistant
Professor
Department of Aeronautics and
Astronautics
Massachusetts Institute of Technology

David Liu
Research Scientist
Corporate Research
Arkema Inc.

Dougal Maclaurin
Research Scientist, Google Brain
Google

Sasikanth Manipatruni
Chief Technology Officer
Computing and Memory
Kepler Computing

Samantha Maragh**
Leader
Genome Editing Program
National Institute of Standards and
Technology

Alexandria Marchi
R&D Scientist
Materials Synthesis and Integrated
Devices
Los Alamos National Laboratory

Johanna Mathieu
Assistant Professor
Department of Electrical Engineering
and Computer Science
University of Michigan

Maria Milina
Senior Research Engineer
Chemicals Research
ExxonMobil Chemical Company

Pramita Mitra
Research Supervisor
Vehicle Research Technologies
Ford Motor Company

Teejay Momoh
Global Systems Engineering Solutions
Architect
Corporate Systems Engineering
Cummins

Marissa Morales-Rodriguez
Research and Development Associate
Electrical and Electronics Systems
Research Division
Oak Ridge National Laboratory

Rebecca Muenich
Assistant Professor
School of Sustainable Engineering and
the Built Environment
Arizona State University

Harish Nagarajan
Senior R&D Manager
Systems Biology and Bioinformatics
R&D
Genomatica, Inc.

Petr Novotny*
Research Staff Member, Enterprise
Solutions
Thomas J. Watson Research Center
IBM

Saad Omar
Senior Scientist
Applied Math and Data Analytics
Schlumberger-Doll Research

Masahiro Ono
Research Technologist
Mobility and Robotic Systems Section
Jet Propulsion Laboratory

Amit Patel
Senior Research Staff Physicist
Sensors and Electronics Laboratory
HRL Laboratories

Alessandro Pinto
Technical Fellow, Embedded AI
Autonomous and Intelligent Systems
United Technologies Research Center

Kalai Ramea
Member of Research Staff
Machine Learning and Data Science
PARC, a Xerox Company

Mark Rentschler
Associate Professor
Department of Mechanical Engineering
University of Colorado Boulder

Reza Rock
Group Leader, Insight Team
Applied Coatings Science
PPG Industries, Inc.

Lisa Rueschhoff
Materials Research Engineer
Composites Branch, Materials and
Manufacturing Directorate
Air Force Research Laboratory

Eric Sacia
Reaction Engineering Consultant
DuPont

Dorsa Sadigh**
Assistant Professor
Department of Computer Science and
Electrical Engineering
Stanford University

Krishanu Saha**
Assistant Professor
Department of Biomedical Engineering
University of Wisconsin–Madison

Michele Sarazen
Assistant Professor
Department of Chemical and
Biological Engineering
Princeton University

Jonathan Sauder
Mechatronics Engineer
Technology Infusion
Jet Propulsion Laboratory

Hae-Jong Seo*
Autonomous Driving R&D
NVIDIA

Kathleen Shafer
Advanced Research Specialist
Corporate Research Materials
Laboratory
3M

Danelle Shah
Assistant Group Leader
Department of Artificial Intelligence
Technology and Systems
MIT Lincoln Laboratory

Elaine Shi*
Associate Professor
Department of Computer Science
Cornell University

Honghui Shi
Research Staff Member, AI
IBM Research

Anshumali Shrivastava
Assistant Professor
Department of Computer Science
Rice University

Katherine Smyth
Research Scientist
Facebook Reality Labs
Facebook

Cynthia Sung
Assistant Professor
Department of Mechanical
Engineering and Applied
Mechanics
University of Pennsylvania

Micaela Tabora Claire
Senior Researcher
Process Technology
ExxonMobil Research and Engineering

Fatima Toor
Assistant Professor
Department of Electrical and
Computer Engineering
University of Iowa

Alejandro Toriello
Associate Professor
Department of Industrial and Systems
Engineering
Georgia Institute of Technology

Lauren Toth
Senior Principal Engineer
Technology Services Sector
Northrop Grumman Corporation

Sonia Tulyani
Senior Director, Materials and
Processes
Materials and Processes Engineering
Pratt and Whitney

Pavan Valavala
Research Scientist
Packaging and Specialty Plastics R&D
Dow, Inc.

Antonios Vytiniotis
Managing Engineer
Civil Engineering
Exponent, Inc.

Alan Wagner
Hartz Family Early Career Assistant
Professor
Department of Aerospace Engineering
Pennsylvania State University

Hong Wan**
Associate Professor
Department of Industrial and Systems
Engineering
North Carolina State University

Qingyu Wang
Research Scientist
Agilent Research Laboratories
Agilent Technologies

Renee Wegrzyn*
Program Manager
Biological Technologies Office
DARPA

Jennifer West*
Fitzpatrick Family University Professor
Department of Biomedical
Engineering
Duke University

Kate Whitefoot
Assistant Professor
Department of Mechanical Engineering
and Engineering and Public Policy
Carnegie Mellon University

Cranos Williams
Associate Professor
Department of Electrical and
Computer Engineering
North Carolina State University

Gregory Wilson
Principal Engineer – Materials
Analytics and Modeling
Engineering Material Systems
GE Aviation

Va'Juanna Wilson
Software Engineer
HorizonX Engineering
Boeing

Brent Winslow
Chief Scientist and Division Head
Biosignature Analytics Division
Design Interactive, Inc.

Sheng Xu
Assistant Professor
Department of NanoEngineering
University of California, San Diego

Dinner Speaker

Joan Robinson-Berry
Vice President and Chief Engineer
Boeing Global Services

Guests

Wesley L. Harris
Charles Stark Draper Professor of
Aeronautics and Astronautics
Massachusetts Institute of Technology

William (Bill) Hayden
Vice President
The Grainger Foundation

Chung-Yan Koh
Systems Engineering and Technical
Advisor
Biological Technologies Office
Defense Advanced Research Projects
Agency

Raheem Lawal
Principal Aerospace Engineer
Air Force Office of Scientific Research

Timothy Lieuwen
Regents Professor and David S. Lewis
Jr. Chair
Daniel Guggenheim School of
Aerospace Engineering
Georgia Institute of Technology

Sohi Rastegar
Senior Advisor
Emerging Frontiers and
Multidisciplinary Activities
Directorate for Engineering
National Science Foundation

National Academy of Engineering

John L. Anderson
President

Alton D. Romig, Jr.
Executive Officer

Guru Madhavan
Director of Programs

Janet Hunziker
Director, Frontiers of Engineering
Program

Sherri Hunter
Meetings Coordinator

Program

NATIONAL ACADEMY OF ENGINEERING

2019 US Frontiers of Engineering

September 25–27, 2019

Chair: Jennifer West, Duke University

ADVANCED MANUFACTURING IN THE AGE OF DIGITAL TRANSFORMATION

Organizers:

Li Chang, Boeing, and Tarik Dickens,
FAMU–FSU College of Engineering

21st Century Engineering Systems:

Boeing's Digital Transformation

Gabriel Burnett, Boeing

Computational Materials for the Design and Qualification of

Additively Manufactured Components

Christopher Lang, NASA Langley

Robots That Walk: What the Challenge of Locomotion

Says About Next-Generation Manufacturing

Christian Hubicki, Florida State University

Digital Twin

Pamela Kobryn, Air Force Research Laboratory

ENGINEERING THE GENOME

Organizers:

Renee Wegrzyn, DARPA, and
Charles Gersbach, Duke University

Genome Editing with Precision and Accuracy

Kris Saha, University of Wisconsin

Using CRISPR to Combat Human Disease Vectors

Omar Akbari, University of California, San Diego

Microbes and Manufacturing: Moore's Law Meets Biology

Patrick Boyle, Ginkgo Bioworks

Empowering Genome Editing Through Standards

Samantha Maragh, National Institute of Standards and Technology

SELF-DRIVING CARS: TECHNOLOGY AND ETHICS

Organizers:

Christoffer Heckman, University of Colorado Boulder, and
Hae-Jong Seo, NVIDIA

Autonomous Vehicles: Opportunities and Challenges

Christoffer Heckman, University of Colorado Boulder

Perception of Low-Cost Autonomous Driving

Tae Eun Choe, Baidu

Why Everyone Has It Wrong About the Ethics of Autonomous Vehicles

John Basl, Northeastern University

Influencing Interactions in Autonomous Driving

Dorsa Sadigh, Stanford University

BLOCKCHAIN TECHNOLOGY

Organizers:

Petr Novotny, IBM, and
Elaine Shi, Cornell University

Blockchains – An Introduction

Elaine Shi, Cornell University

Blockchain Beyond Cryptocurrency: An Overview

Hong Wan, North Carolina State University

Cryptocurrencies as Marketplaces

Jacob Leshno, University of Chicago

