On the Use of Modeling and Simulation in Robotics: A Workshop Report

Stuart Anderson Byron Boots Arunkumar Byravan Evan Drumwright
Christian Duriez Dieter Fox Gregory Hager Jessica Hodgins
Abhinandan Jain Ashish Kapoor Daniel Koditschek Nate Koenig
Edward Lee Chen Li Karen Liu Franziska Meier Dan Negrut
Akshay Rajhans Ludovic Righetti Alberto Rodriguez Stefan Schaal
Jie Tan Yuval Tassa Emanuel Todorov Jeff Trinkle

September 30, 2018

Contents

Backdrop

The last five years marked a manifest surge in interest for and use of "smart" robots that operate in dynamic and unstructured environments and might interact with humans. Yet the principles that govern the design and control of these intelligent agents are anything but clear. Against this backdrop, the premise of this workshop was that computer simulation can provide a fertile testbed instrumental in understanding faster, at lower costs, and more thoroughly how the robots of the future should be designed for safety and performance. The purpose of the one-day workshop was to bring together experts from robotics and the modeling and simulation (M&S) communities in a brainstorming exercise that pinpointed opportunities, identified challenges, and proposed "next steps" vis-à-vis the goal of increasing the role that computer simulation plays in smart robotics.

Robots are no longer stiff/rigid implements operating in the structured environment of assembly lines and performing a limited set of operations. Machine learning and artificial intelligence are poised to endow a new generation of robots with mobility and decision-making skills. These robots might be flexible, reconfigurable, interacting with humans, and operating in environments that are unstructured and changing. The expectation is that in the near future they will be operating on highways as autonomous vehicles, in nursing homes assisting social workers, in schools tutoring young learners, underwater managing oil spills, in adverse and cluttered environments in rescue missions, in operating room assisting surgeons, etc. Yet physically testing such robots to ensure quality and safety is complex, time consuming, and expensive. The power of cloud computing in combination with software-as-a-service facilities can leverage simulation and machine learning techniques to produce a virtual proving ground used to design this new generation of robots in an exercise that can reduce costs, increase testing and operation safety, improve quality and compress time to market. This document attempts to capture the main discussion points that emerged during an engaging one-day workshop on the very topic of using simulation in robotics. We hope that this report will inform both the common research effort and resource allocation planning. These are two key aspects that will shape the trajectory of an industry anticipated to reach $82.7 billion by 2020 and which has registered an annual growth rate of more than 10% since 2014[1].

Workshop invited participants and document signatories: Stuart Anderson, Byron Boots, Arunkumar Byravan, Evan Drumwright, Christian Duriez, Dieter Fox, Gregory Hager, Jessica Hodgins, Abhinandan Jain, Ashish Kapoor, Daniel Koditschek, Nate Koenig, Edward Lee, Chen Li, Karen Liu, Franziska Meier, Dan Negrut, Akshay Rajhans, Ludovic Righetti, Alberto Rodriguez, Stefan Schaal, Jie Tan, Yuval Tassa, Emanuel Todorov, and Jeff Trinkle.

---

[1] Allied Market Research, "Robotics Technology Market by Type." `https://www:alliedmarketresearch:com/robotics-technology-market`. Accessed: 2018-02-09.

Acknowledgments: The signatories would like to acknowledge the National Science Foundation, Army Research Lab, Networking and Information Technology Research and Development Program, and National Institute of Standards and Technologies for providing the financial (NSF NRI Grant 1830129) and logistic support that made this workshop possible.

Disclaimer: The views and opinions expressed in this document are those of the signatories and do not necessarily reflect the position of other workshop participants, the organizations that the signatories belong to, or any US federal agency.

1 Event Logistics

The April 17, 2018 workshop was hosted by the National Institute of Standards and Technologies in Gaithersburg, MD. It was organized by Dan Negrut of the University of Wisconsin-Madison in collaboration with colleagues from the University of Washington: Dieter Fox, Arunkumar Byravan, and Franziska Meier.

The workshop had two groups of participants. The "invited participants and organizers" group included: Stuart Anderson (NVIDIA), Byron Boots (Georgia Institute of Technology), Arunkumar Byravan (University of Washington), Evan Drumright (George Washington University), Christian Duriez (INRIA Lille), Dieter Fox (University of Washington), Gregory Hager (Johns Hopkins University), Jessica Hodgins (Carnegie Mellon University), Abhinandan Jain (Jet Propulsion Lab), Ashish Kapoor (Microsoft), Daniel Koditschek (University of Pennsylvania), Nate Koenig (Open Robotics Foundation), Edward Lee (University of California-Berkeley), Chen Li (John Hopkins University), Karen Liu (Georgia Institute of Technology), Franziska Meier (Max Planck Institute for Intelligent Systems/University of Washington), Dan Negrut (University of Wisconsin-Madison), Akshay Rajhans (MathWorks), Ludovic Righetti (New York University), Alberto Rodriguez (MIT), Stefan Schaal (University of Southern California), Jie Tan (Google Brain Robotics), Yuval Tassa (Google DeepMind), Emanuel Todorov (University of Washington), and Jeff Trinkle (Rensselaer Polytechnic Institute).

The second group of participants comprised federal employees who expressed an interest in attending the workshop in an "observer" capacity. This group included: Jordan Berg (NSF), Hee Sun Choi (NIOSH), Cindy Crump (USAMRMC), Irina Dolinskaya (NSF), Faisal D'Souza (NITRD Program), David Han (ARL), William Harrison (NIST), Stuart Harshbarger (Johns Hopkins University Applied Physics Laboratory), Frank Hearl (NIOSH), Atul Kelkar (NSF), Frederick Leve (US AFSOR), Elena Messina (NIST), Geoff Miller (USAMRMC), Reid Simmons (NSF), Donald Sofge (US Navy), Johnathan Sprinkle (NSF), Ethan Stump (ARL). Abbreviations used: NSF - National Science Foundations, ARL - Army Research Lab, AFSOR - Air Force Office of Scientific Research, USAMRMC - US Army Medical Research and Materiel Command, NIOSH - National Institute for Occupational Safety and Health, NIST - National Institute of Standards and Technologies, NITRD - Networking and Information Technology Research and Development.

Two weeks prior to the workshop each participant was asked to provide four slides with input along the following lines: slide 1 – *Understanding opportunities: how M&S can assist Robotics*; slide 2 – *M&S hurdles: What's stopping us from getting there*; slide 3 – *Pragmatic suggestions for moving forward*; slide 4 – *Additional thoughts/comments you might have*. A word cloud based on these decks of slides is shown in Fig. 1. The four-slide pre-event docs received were minimally curated (removed author name, etc.) and are available online[2]. They were used by the organizers to compile three "lightning presentations" that each seeded one of the three breakout sessions of the workshop:

- **Breakout 1 - How M&S could help/helps Robotics:** The first breakout focused on identifying ways in which M&S can impact Robotics. The main points that emerged in this breakout

---

[2] `https://sbel.wisc.edu/simulation-in-robotics-2018/#pre-workshop-thoughts`

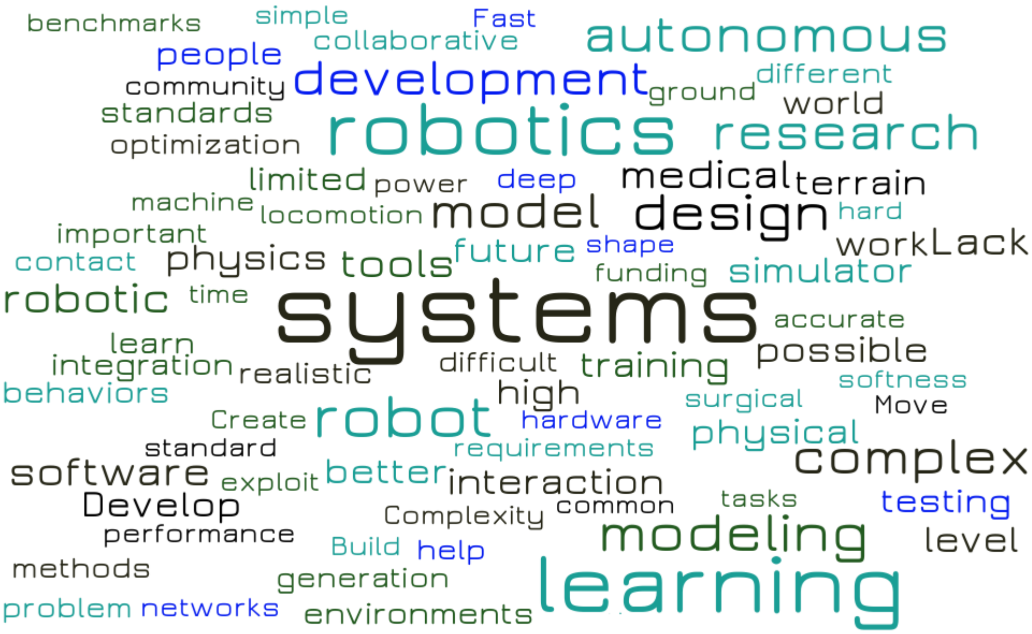

Figure 1: *A word cloud generated based on the pre-workshop slides provided by participants. Word size proportional to use count; "simulation" was excluded.*

are summarized in section §2.

- **Breakout 2 - Issues that continue to hinder M&S in Robotics:** In this breakout the participants sought to compile a list of open problems that presently limit the role that M&S plays in Robotics. Section §3 summarizes the main ideas that coalesced in this breakout.
- **Breakout 3 - Concrete and immediate measures/steps.** There were two themes in the final breakout. First, we sought to understand whether a consensus emerged in relation to concrete short-term and long-term steps towards improving robotics solutions via computer simulation. Second, the experts were asked to express their thoughts on possible breakthroughs and speculate on disruptive technologies that we might face in M&S enabled robotics, augmented reality, virtual reality, etc. The points made in this breakout are discussed in section §4.

The meeting's agenda was as follows:

07:30 - Registration
08:00 - Introduction: goal of the workshop, logistics (Negrut)
08:30 - Participant self-introduction
09:00 - Coffee break/Networking
09:30 - Breakout, Topic 1: *How M&S could help/helps Robotics*. Lightning presentation to "seed" breakout discussion (Dieter Fox). Suggested breakout discussion points:

- Machine learning aspects
- Model based control, path planning, optimal design
- Autonomous vehicles
- Etc.

11:15 - Joint session. Group speaker summarizes each groups thinking on issues related to Topic 1. Follow up, plenary discussion.

12:00 - Lunch (onsite, catered)
12:30 - Breakout, Topic 2: *Issues that continue to hinder M&S in Robotics*. Lightning presentation to "seed" breakout discussion (Arunkumar Byravan). Suggested breakout discussion points:

- Nonsmooth dynamics (friction/contact/impact)
- Soft robotics (nonlinear geometry/nonlinear material/nonlinear boundary conditions)
- Multi-physics (CFD, plasticity)
- Terramechanics
- Central repository of standardized, shared models
- Open-source issues
- Etc.

13:15 - Joint session. Group speaker summarizes each groups thinking on issues related to Topic 2. Follow up, plenary discussion.
14:00 - Coffee break/Networking
14:30 - Breakout, Topic 3: *M&S in Robotics: concrete next steps. Possible breakthroughs. Speculations about disruptive technologies vis-à-vis the issue of M&S in Robotics*. Lightning presentation to "seed" breakout discussion (Franzi Meier). Suggested breakout discussion points:

- Low hanging fruit
- "The long view" issues
- Imminent landscape changes
- Prioritization aspects. Spectacular-return-on-investment opportunities
- What is the role of industry (video-games, VR, AR)? What/How can we leverage?
- Etc.

15:15 - Joint session. Group speaker summarizes each groups thinking on issues related to Topic 3. Follow up, plenary discussion.
16:00 - Coffee break/Networking
16:15 - The big picture, open-floor discussion. Next steps: workshop report and journal paper.
17:00 - Wrap up

The cornerstone of the workshop was the set of three breakouts and the follow-up plenary discussions. In each breakout, six teams of eight or nine participants moved to separate rooms to engage in a 25-minute brainstorming session. Each of the six teams elected a speaker to present the team's perspective on the topic at hand in the follow-up plenary session. The team's thinking was summarized in three slides: a *consensus thinking* slide, an *ideas that were somewhat contentious* slide, and an *"out there" surprising thoughts, comments, odds and ends* slide. The slides for each team and for each breakout are available online[3,4,5]. Each breakout was followed by a 45-minute-long plenary session whose purpose was twofold: allow team speakers to quickly outline the team's thinking on the breakout's topic; and, enable an open discussion whose starting point was the ideas brought forward by each team. In an attempt to create an environment conducive to the emergence of original/diverse ideas, the composition of the teams was different for each breakout. That is, the teams were shuffled – the participants mostly teamed up with different colleagues for each of the three breakouts. Finally, the workshop concluded with a 45 minute-long open-floor discussion where the participants had an opportunity to reflect on the day's discussions.

The planned outcomes of the workshop were this report and a slated co-authored journal submission that would dive into technical details that fell outside the scope of this report. The "raw" information generated in conjunction with this workshop is a set of 25 decks of four slides (pre-workshop input provided by participants), 18 decks of slides containing team opinions (three breakouts, each with six teams), and three lightning talks used to seed the breakout session discussion. All this information is publicly available[6].

---

[3] `https://sbel.wisc.edu/simulation-in-robotics-2018/#breakout-session-1`
[4] `https://sbel.wisc.edu/simulation-in-robotics-2018/#breakout-session-2`
[5] `https://sbel.wisc.edu/simulation-in-robotics-2018/#breakout-session-3`
[6] `https://sbel.wisc.edu/simulation-in-robotics-2018/`

2 M&S in Robotics: Opportunities

Ideas embedded in the pre-workshop slides and further discussions in the first breakout session helped us identify several opportunities (Opps.1 through 8) in which M&S can come into play in Robotics:

1. Generate expeditiously and at low costs large amounts of training data for Machine Learning
2. Accelerated and safe design environment
3. Accelerated, safe, and fully controlled virtual testing and verification environment
4. Opening the door to intelligent robots through simulation-enabled online-inference
5. Controlling solution costs
6. Multi-robot, collaborative scenarios
7. Assessment of adversarial attacks
8. Facilitate the understanding of human-robot interactions (HRI)

Opp.1. **Generate expeditiously and at low costs large amounts of training data for Machine Learning.** In the quest to leverage machine learning and deploy artificial intelligence solutions, the ability to generate through simulation a wealth of data in a short amount of time represents a major selling point of the simulation-in-robotics mantra. A validated virtual environment represents an ideal proving ground for developing systems that can both learn from their mistakes and be verifiable. Simulation can be used to discover new behaviors and to carry out tasks never known to be possible. The advantage of using simulation to generate training data becomes even more compelling when the software used can draw on ubiquitous high-throughput computing resources; i.e., using multiple nodes to carry out batches of simulations in parallel and perhaps in the cloud.

Opp.2. **Accelerated and safe design environment.** Today, the design of robots and associated controls are often time-consuming since this process employs multiple physical prototypes. Physics-based virtual models of the robots could remove a part of the experimental burden. The "morphology" of the robot; i.e., its shape, actuators, sensors, etc., could be optimized through simulation to accomplish given objectives in a process led by the tasks aimed to be executed. In an ideal world, control policies synthesized via simulation would be transferable to actual robots. For instance, simple yet rich predictive models can be used for receding horizon control.

Opp.3. **Accelerated, safe, and fully controlled virtual testing and verification environment.** Approaches to verification of autonomous systems are in their infancy. Approaches to verification of autonomous robotic systems that learn on-line are essentially non-existent. Simulation can help establish the principled and effective means for autonomous system verification and thus be instrumental in developing industry and military standards and guidelines.

Opp.4. **Opening the door to intelligent robots through simulation-enabled online-inference.** By and large, today's robots lack the ability of making decisions on the fly. They are controlled in structured environments and trained to perform a narrow and a priori established catalog of operations. This is bound to change owing to several factors: the types and accuracy of sensing have improved significantly; the actuation and controls are gradually becoming more sophisticated; the compute speeds available on an affordable power budget have increased substantially. The emerging opportunity is that of having robots operate in unstructured environments in which they are faced with decision-making activities. Just like how the concepts of morality, prior experience and ability to predict the consequence of one's activities shape the decision-making process in humans, a set of prior experiences, a set of rules (including ethical) and an ability to predict consequence through simulation can shape the decision-making process of robots.

Opp.5. **Controlling solution costs.** Compressing the robot's design time by limiting the use of physical models translates into reduced solution costs. Cost and time-to-solution savings are expected both in the design and safety/verification stages. Note that during the design stage, one can account not only for function but also durability aspects; i.e., being able to predict the durability of a robotic system, for instance, knowing the mechanical stresses that crop up during various use cycles. Finally, simulation can be instrumental in choosing between competing solutions to identify trade-offs between task efficiency and end product costs.

Opp.6. **Multi-robot, collaborative scenarios.** Collaborative multi-robot systems are composed of robots interacting with each other based on own local decision-making algorithms that factor in sensed and/or shared information. Designing and verifying these systems becomes increasingly complex as the number of interactions between robots goes up. Testing and verification is tedious as the number of scenarios to probe increases dramatically with the number of robots in the system. For instance, it is difficult to systematically test in real conditions multi-robot systems used for environmental monitoring, surveillance or infrastructure management due in part to the challenges posed by operating in these environments. Finally, another opportunity is that of testing multi-robot systems that communicate and collaborate with humans in heterogeneous teams, e.g. in search and rescue scenarios, with a human-in-the-loop component.

Opp.7. **Assessment of adversarial attacks.** Imagine an autonomous vehicle with a security breach that allows a malicious attack on the sensor system. How will the autonomous vehicle react? Having accurate representations of how the information system is built and being able to test its behavior in cases of failure is critical for safety and regulation purposes. Indeed, robots can be bulky and gone awry can become outright dangerous. Simulating the reaction of the robotic system in cases of failure, while it is potentially performing a task, permit the investigation of fail-safe strategies for the system and design of security, "anti-breach" policies employed by robots.

Opp.8. **Facilitate the understanding of human-robot interactions (HRI).** Ability to simulate the interaction between the robot and human presents opportunities in experimenting with tele-surgical robotics in semi-autonomous or autonomous operation; reduction of risk to workers in dangerous occupations; eliminate repetitive motion trauma and musculoskeletal overload; reduced fatalities and injuries from motor vehicle incidents as autonomous vehicles are yet another form in which robots interact with humans. Of particular interest for small teams operating remotely (e.g., a mission to Mars), the human-robot interplay can project specialized medical expertise and care to any point of need, at any time.

3 M&S in Robotics: Challenges

The second breakout concentrated on identifying issues that, to various degrees, restrict the role that simulation plays in smart robotics.

1. The simulation-in-robotics landscape is currently fragmented and only slowly emerging
2. There is no established robotics modeling language that exposes suitable abstractions and hierarchies
3. Model composability aspects not addressed yet
4. Multi-physics simulation prone to taking long run times
5. Model calibration could be tedious
6. Data-driven simulation (surrogate models; replacing simulation with an oracle) is in its infancy
7. No clear understanding and/or consensus on mandatory-level-of-fidelity aspects

Issue 1. **The simulation-in-robotics landscape is currently fragmented and only slowly emerging.** After decades of investments, simulation in *Engineering* is arguably ubiquitous. For instance, one can carry out a highly predictive finite element analysis of an engine crank shaft or of an airplane wing to accurately identify vibration/flutter modes. Likewise, a vehicle ride analysis can be carried out via a dynamics simulation in which a vehicle on a flat, rigid surface and a sequence of driving maneuvers are performed for noise/vibration/harshness purposes. Computer Aided Engineering (CAE) draws on a well understood and "tame" set of boundary conditions, loading functions and/or driving constraints. Compared to CAE, robotics simulation is manifestly less structured and typically more multi-disciplinary. Indeed, in addition to producing models for robot dynamics, which is in the purview of CAE, one has to: bring in sensor simulation required for agent control; simulate in high fidelity the unstructured environment surrounding the agents of interest – a prerequisite for sensing; produce controls strategies; simulate the human dimension in HRI; and, simulate agent-to-agent or agent-to-infrastructure communication. These pre-requisites elicit both research and development work to gradually close the functionality gap between them and the level of sophistication of CAE. They are also at different levels of sophistication in the quality of solution each brings to the table, from reasonably good for simulating agent-to-agent communication to utterly in "embryonic stage" for capturing in simulation the HRI dimension.

Issue 2. **There is no established robotics modeling language that exposes suitable abstractions and hierarchies.** The abstractions defined should be conducive to yielding models with hierarchical structure that accommodate the user's needs vis-à-vis the level of fidelity rendered in simulation. The modeling language should accommodate the manifest lack of consensus in relation to the required level of accuracy in simulation for robotics by allowing one to incrementally build up, if so desired, the level of complexity of a model. The abstraction mechanism should look for inspiration in Electronic Design Automation (EDA), Open System Interconnect (OSI) or the Mead-Conway VLSI Design Revolution. Cues could also be taken from the softwares used to create computer games or 3D computer animated movies (Unreal, Unity, Maya, etc.) These softwares already allow users to create complex virtual worlds with very good visual rendering. They are already able to emulate realistic visual environments for the robots which come into play in camera and LiDAR sensing used for control and decision making. However, these softwares are not programmed to simulate accurately the physics governing the behavior of robots and their interactions with the environment; i.e., their predictive attributes are lacking owing to an emphasis on plausibility rather than accuracy.

Issue 3. **Model composability aspects not addressed yet.** The robotics modeling language adopted should promote a composability perspective on model generation. Simple constructs, like in a Lego game, should be useful by themselves or be composed to produce complex scenarios. The simple constructs enable testing/debugging of actual control algorithms/software when assessing robot control solutions. They are expected to run significantly faster than real time and come into play in control and on-line machine learning. Complex scenarios are investigated via composed models that draw on validated, discipline-oriented, submodels. They might be slower than real time. Reducing the process of model generation calls for integration of existing assets/lessons learned/existing capabilities. For instance, CAD-to-dynamics model pipelines (which leverage, for instance, SolidWorks), as well as "video gaming-to-robotics model" virtual worlds translation can lighten the burden of developing from scratch protocols and principles for composability in robotics simulation. On the up-side, the game community has established sophisticated game authoring tools. On the down side, they are massive chunks of code, cumbersome to work with and, as pointed out in Issue 2 only partially covering the robotics community's needs owing to emphasizing credibility-in-appearance rather than accuracy-in-results.

Issue 4. **Multi-physics simulation prone to taking long run times.** Simulation in robotics can be used for design, on-line control, and machine learning purposes. The latter two applications require fast run times. In general, there is a non-linear increase in simulation run-times relative to the complexity of the model. The rule of thumb is that one robot made up of rigid bodies typically runs in real-time. The real-time attribute is challenged by models that contain flexible/soft/deformable components. An extra layer of complexity that slows down simulations is associated with unstructured virtual worlds in robotics, e.g., deformable and cluttered terrain, fluid-solid interaction, etc. Bringing more physics into the model is important when robots venture outside controlled environments and operate in complex real world scenarios. However, once basic physics principles have been established sufficiently for such complex interactions to provide relatively simple rules, high-fidelity physics model based simulations are perhaps not always necessary; understanding where to stop is tied to the discussion in Issue 7 below. Finally, the "weaponry" available to make a simulation run fast is specific to the goal pursued. The same model can run fast in a use-for-design framework, when one can count on power-hungry hardware such as GPU computing when solving complex systems, while running very slowly in an on-line setup where weight and/or power constraints limit the type of hardware that can be counted on. Another facet of the "simulation time" problem is the need, in on-line control, for a deterministic time attribute of the solution; i.e., the ability to guarantee that an approximation will be produced in a predefined amount of time regardless of the model's state and input.

Issue 5. **Model calibration could be tedious.** Populating a model with correct parameters for robots operating in unstructured environments is time consuming and most often an ad-hoc, case-by-case, process. For instance, friction and contact simulation for rigid-wheel/deformable-terrain interaction requires 14 parameters[7], which are typically determined through a bevameter test carried out in the field. Likewise, when simulating soft robots, or a FEM analysis for that matter, one might have to specify properties for non-linear, inhomogeneous, and anisotropic materials. While a tall order for the specialist, it is difficult for a person building the simulation engine to capture in software the nuances and subtleties of the continuum mechanics problem, and even more so for a user, who needs to yield model input values operating in a model-composition space modulated by the understanding of the software developer. The science of parameter identification is the holy grail, yet it is not expected that simulation-in-robotics will be conducive to out-of-the-box solutions any more so than other fields that pose large identification problems originating in multi-physics applications.

Issue 6. **Data-driven simulation (surrogate models; replacing simulation with an oracle) is in its infancy.** The recent advances in machine learning, together with the standardization of robotic platforms and software open the possibility to collect large amounts of data in real world scenarios that could be leveraged to build predictive models directly from data. The hope is that using statistical learning techniques, one could bypass simulation-specific hurdles such as model generation and calibration by constructing "oracles" able to predict the next system state given its current one. Data-driven simulation could be used to methodically reduce model (and therefore computational) complexity through systematic dimensional reduction and/or model compression. However, such approaches are in their infancy and principled methodologies to design and validate data-driven models remained to be identified.

Issue 7. **No clear understanding and/or consensus on mandatory-level-of-fidelity aspects.** Demands for highly accurate models that capture multi-physics phenomena, have large degree-of-freedom counts, and require a detailed treatment of friction, contact and compliance often lead to long run times. Complex models also require a wealth of input data, which in many cases needs to be approximated leading to situations in which the complex simulation yields less accuracy than the expeditious one owing to bad model parameters. Moreover, debugging complex models is hard since it is difficult to discern whether sub-par results are due to poor model parameters or to weaknesses in the modeling and numerical solution techniques anchoring the simulator. Against this backdrop, there is no consensus regarding when complex models are needed; i.e., at what point sub-par models confine our ability to discover through simulation. Other similar questions that weigh in the mandatory-level-of-fidelity discussion are the need for soft/hard real-time simulation; and, whether machine learning approaches are robust enough to not get tripped by what quality-wise is sub-par training data obtained with fudged parameters or simplified models. "Tripped" in

---

[7] J. Y. Wong and A. R. Reece, "Prediction of rigid wheel performance based on the analysis of soil-wheel stresses part i. performance of driven rigid wheels," Journal of Terramechanics, vol. 4, no. 1, pp. 8198, 1967.

this context refers to solutions that work fine in simulation but fail when deployed in actual robots. Finally, one should be cognizant of the fact that simulation serves multiple end goals – producing data for learning, designing better controls, improving mechanical performance, auditing for safety purposes, etc., which suggests that diversity in level of fidelity might be inevitable.

4 Looking Ahead

The next two subsections summarize, respectively, recommendations and high-vantage-point observations that emerged at the end of the third and last breakout.

4.1 "Nuts-and-bolts" Recommendations

1. Foster adoption of open source simulation platforms
2. Converge towards a small number of community curated and maintained model libraries
3. Establish a small set of "simulation-in-robotics" grand challenges
4. Characterize the human-robot interaction
5. Embrace a compliant-body perspective in robotics and handle friction/contact fallout
6. Embrace and account for uncertainty
7. Capitalize on emerging hardware architectures

Rec. 1. **Foster adoption of open source simulation platforms.** A robust and feature-rich set of simulation tools in the open-source domain is critical to advancing the state of the art in robotics. An open source platform democratizes the "simulation-in-robotics" effort and becomes a source of inspiration for future, more refined open-source or commercial efforts. Owing to the breadth of robotics applications, it is expected that no single platform will emerge as the solution of choice for all targeted simulation scenarios. Producing a catalog of simulation platforms available and compiling a community sanctioned set of recommendations that indicates which tool works well in what scenario would likely improve the adoption rate for simulation in robotics. The presence of several accessible platforms will allow for cross-pollination via algorithm and data recycling; and, foster competition as users will be able to compare and contrast alternatives.

Rec. 2. **Converge towards a small number of community curated and maintained model libraries.** Simulation in robotics calls for the interplay of several types of models: robots, synthetic worlds, sensor models. In some cases, models of the human component can also be required, and, for multi-robot scenarios, one might need to simulate the communication layer. Any one of these models is complex, both to generate and endow with meaningful parameters, see Issue 5. Collaborate with industry to endow each robot with a model library for its simulation. Models at various levels of complexity are desirable for the same problem (goes along with the idea of "composability" and progresses from fast to high accuracy) since different applications call for different modeling choices.

Rec. 3. **Establish a handful of "simulation-in-robotics" grand challenges.** Establishing a collection of grand challenges is expected to pay dividends in several ways. First, a carefully chosen small set of challenges provides a sense of purpose to the community effort. Without direction and clear purpose, the community can choose to focus limited resources along technical thrusts of secondary relevance. Second, a set of challenges would be a catalyst for inter-team collaboration as groups with complementary expertise might come together to address a big problem. The current robotics simulation gaps are best addressed by multi-disciplinary teams, combining model developers and model users. Finally, if continued over a period of three to five years, a "grand challenges" initiative that has national visibility will kindle and foster an ecosystem build up effort to produce the methods, tools, and models of the trade.

Rec. 4. **Characterize the human-robot interaction (HRI).** The issue of establishing human models that capture mechanical attributes of the body and/or psychological and cognitive traits of the human behavior is cross-cutting. Applications in which HRI will come into play include robotic surgery, which is relevant in laparoscopic surgery training and remotely treating patients in combat zones; assisting seniors with tasks such as dressing, personal hygiene, cleaning and cooking; assisting individuals with limited ambulatory ability with transportation needs, etc. The body of work in the area of HRI modeling is meager, which explains the limited knowledge vis-à-vis the issue of human cognitive performance in HRI. In this context, there is a very limited set of science-based requirements and thresholds for safe human-robot interaction.

Rec. 5. **Embrace a compliant-body perspective in robotics and handle friction/contact fallout.** To date, robotics simulation has almost exclusively drawn on rigid body dynamics. Indeed, the underlying modeling is simpler, the software implementation effort is more reasonable as are the typical simulation run-times. Moving to a flexible-body representation of robotics is mandatory, for instance, in HCI. It is anticipated that embracing compliance in our robotics models will elicit new approaches to handling frictional contact with the potential benefit of alleviating numerical artifacts/paradoxes brought to the fore by the rigid body model. Generating through simulation sensory-motor data that matches the multi-resolution dynamics, noise, softness, etc. of sensors and actuators during complex tasks that include frictional contacts would allow for a systematic study of the sensory-motor space for robotic manipulation and locomotion.

Rec. 6. **Embrace and account for uncertainty.** Friction, impact, contact, actuator noise, uncertain external loads, complex and unstructured environments, etc. are sources of uncertainty that mar the physical and virtual robot. Accounting for uncertainty in simulation-for-robotics is difficult both in the modeling and solution phase. In modeling, one has to provide mechanisms to inject uncertainty in the model, e.g., allowing the friction coefficient to change in a deformable and heterogeneous terrain; the geometry of various components to be less than perfect; the actuation forces/torques to display delays, etc. In running simulations, one has to handle discontinuities associated with the solution correctly; produce statistical information in an expeditious way; establish control and optimization algorithms that are stochastic in nature, etc. The perspective that one needs to take in simulation-in-robotics is a statistics one, in which confidence bounds are necessary to gauge the extent to which one can rely on simulation results.

Rec. 7. **Capitalize on emerging hardware architectures.** Recent advances in high-bandwidth memory and large processors counts can partially alleviate increased computational loads associated with terramechanics, soft robotics, fluid-solid interaction and real-time simulation. Looking beyond GPU and multi-core computing, a "federated" approach in which various components of a model are handled separately by different solvers in a co-simulation framework would: facilitate a plug-and-play vision that enables scalability of the solution; open the door to contributions coming from multiple groups; and permit the adoption of best-in-class solvers for subproblems.

4.2 "High-vantage-point" Aspects

Unlike the previous subsection, which contained a list of recommendations, the focus below is on capstone observations.

1. Building a simulation-in-robotics ecosystem is a multi-disciplinary effort that poses numerous open questions in basic research
2. Simulation-in-robotics can provide a grand challenge that would galvanize disparate communities towards a very worthwhile end
3. Simulation-in-robotics has a sizable software development component to it; small and agile groups of domain experts involved in this activity have a role to play

High Level 1. **Building a simulation-in-robotics ecosystem is a multi-disciplinary effort that poses numerous open questions in basic research.** Several types of simulation engines have to be combined to produce a simulation-in-robotics solution platform. A dynamics engine is needed to capture the time evolution of the robots; sensor simulation is required to control the agents; correct sensor simulation is predicated on the ability to simulate the environment and model its unstructured and time-dependent nature; simulating the human dimension in HRI; finally, communication simulation comes into play for agent-to-agent or agent-to-infrastructure communication. Open problems that defy our current level of understanding come up in conjunction with each of these simulation engines. Addressing the open problems calls upon basic research in a variety of areas such as physics-based modeling (friction, contact, impact, soft/compliant bodies, fluid-solid interaction, terramechanics); numerical methods (fast real-time solvers, DAE solution, PDE solution, linear solvers, FEM, CFD, computational geometry); controls techniques (model predictive control); software development (hardware-aware software); data analysis (uncertainty quantification). Serendipitously, several funded initiatives are ongoing, e.g., soft robotics (NSF) and terramechanics (DOD), which are bound to have an impact in this community. Nonetheless, that majority of the simulation-in-robotics multi-disciplinary open problems, e.g., modeling the human component in HRI, sensor modeling, model composability, etc., linger and at best are tangentially addressed in a context that lacks synergy and is non-programmatic. Against this backdrop, we are of the opinion that a broad multi-agency initiative that is two-pronged; i.e., it stimulates fundamental research in relevant areas and at the same time fosters its translation into open source simulation platforms, can markedly accelerate the pace at which simulation impacts smart robotics. A broad initiative of this caliber would bring together, ideally in an international framework, research groups from academia and agile software development outfits from academia, research labs, or industry.

High Level 2. **Simulation-in-robotics can provide a grand challenge that would galvanize disparate communities towards a very worthwhile end.** It is this group's belief that computer simulation can and should play an important role in smart robotics. We take our cue from the enthusiastic adoption and impact of computer simulation in other industries and endeavors – from building cars and airplanes to planetary exploration. Yet compared to the computer aided design (CAD), computer aided manufacturing (CAM) and computer aided engineering (CAE) solutions that anchor both the product life-cycle in industry and a vigorous research enterprise in academia, the simulation in robotics field is in its infancy. Moreover, beyond lack of maturity, simulation in robotics is faced with the task of serving user groups pulling in different directions by virtue of them being engaged in different activities such as *design* (will it brake? is it fast enough? how do we build it?), *controls* (how can we make it reconfigure? how does it climb stairs? how does it work with other robots? will it work under limited sensing?), *machine learning* (how can I generate and label expressive/representative training data sets?) and *artificial intelligence* (does it know what to do in unstructured and alien environments?). Lastly, by comparison with CAD/CAM/CAE, simulation in robotics poses both specific and truly multi-disciplinary challenges that are yet to be addressed, e.g., simulating the human-robot interaction, simulating sensing, controls in unstructured environments, multi-agent dynamics, etc.

While established commercial CAD/CAM/CAE solution providers, e.g., MSC.ADAMS, Siemens, Dassault, RecurDyn, CM Labs, etc., are anticipated to gradually address the needs of the robotics community, our hopes in the immediate future; i.e., 5 to 10 years out, are pinned on nimbler and more focused platforms, e.g., Bullet, Chrono, DART, Drake, MuJoCo, ODE, SOFA, the majority open source, which are plugged deeper into the research community and have the flexibility to rapidly translate modeling, numerical methods, sensor simulation, graphics, and emerging hardware architectures breakthroughs into advances that amplify the impact that simulation plays in robotics. Ultimately, the goal of reaching robust and accurate simulation in robotics is worthwhile. In pursuing this goal, difficult open problems that straddle discipline boundaries remain to be solved since one cannot simulate unless one understands.

High Level 3. **Simulation-in-robotics has a sizable software development component to it; small and agile groups of domain experts involved in this activity have a role to play.** While posing significant challenges, software development is not a research activity insofar as simulation-in-robotics is concerned. Yet the process adopted for developing the necessary software infrastructure and the terms under which the software is released play a role in how soon simulation leaves a mark in robotics. An important aspect is whether or not the software that underlies a simulation-in-robotics initiative should be released as open source. Likewise, there are several licenses, some more permissive than others, vis-à-vis how open-source can be used/modified/distributed. The salient point is that software that solves the problem at hand is desirable in any form. Perhaps, at the onset of a simulation-in-robotics initiative open source released under a permissive license such as MIT is more attractive only for the reason that a component of it, be it a graphical user interface, data input/output facilities, a particular algorithm, a collision detection implementation, etc., might be recycled by another effort, be it open or closed source. In our experience, the argument that an open source code grows faster owing to contributions from volunteers has proven largely insubstantial. Indeed, the level of knowledge required to make meaningful contributions to a project are prohibitive, particularly early on in the trail-blazing phase of the project. In fact, the point in time in the software life cycle when volunteers can help marks the moment when the project would ideally go commercial or be sponsored by such an entity. This brings up the question of who should develop the software that enables the simulation-in-robotics vision. Encouraging domain experts to engage in software development pursuits that have a manifest translation attribute; i.e., demonstrating new algorithms, modeling approaches, etc., would be beneficiary, particularly at the onset of the initiative and when done with a mindset of generating open-source software. Ideally, these domain-expert generated software components would be recycled by more mature simulation platforms. In this ecosystem, as the field matures, one can hope that bigger and perhaps commercial entities would step in to carry the burden of adding the features of convenience that accelerate solution adoption.