2nd Workshop on Closing the Sim2Real Gap

Notes by Jacky Liang

Debates

Debate 1 | “Sim2Real is a waste of time and money” | For: Chris Atkeson, Abhinav Gupta. Against: Ken Goldberg, Peter Welinder

- Opening
  - Gupta
    - simulation is not useful for machine learning on “basic manipulation for a 5 year old”
    - physical domain gap is huge, even for contact dynamics
    - most simulators are too simple
    - simulators require a lot of effort and prior knowledge
    - “real-world robotics is unscalable” is not true - you can collect large data since robots are becoming cheaper
    - “real-world lab data is not diverse” - but simulators are also not diverse?
    - “fast dev/debug time” - but missing richness of real-world data. just b/c something doesn’t work in sim doesn’t mean it won’t work in real
    - “safe exploration” - but the right approach is to have more robust/compliant hardware
  - Welinder
    - sim2real gets bad rep b/c ppl think it’s sim only, and overcoming sim2real gap can be done
    - the fact that training in sim and testing in real works for some applications is promising
    - sim training can bootstrap real training
    - cheap robots are harder to use
    - sim2real lowers barrier to research
    - should treat sim2real as a tool - no need to always use or not use it
  - Atkeson
    - in other industries, premise of using simulator is no sim2real gap
    - it’s cheaper to use real than making sim really accurate
    - “simulation is doomed to succeed”
    - ML sim2real ignores prior works in robust and adaptive controls
    - sim2real might be more useful for meta-learning
  - Goldberg [Opening Statement]
    - define terms: Sim2Real is the use of simulation to study models before they are tested in reality.
    - let’s dispense with the straw man arguments: “simulation isn’t reality” and “sim is the end point” - everyone agrees.
    - calls the latter: “Sim2Null”, showing a simulation and stopping there.
    - sim is necessary (if it doesn’t work in sim it won’t in reality) but never sufficient,
    - iterative system design is faster, cheaper, and safer in sim than real
    - sim allows resets, knowledge of ground truth, and a level playing field for comparing an alg with baselines.
    - the old adage: “Simulation is doomed to succeed.” is false. Simulations fail all the time and that’s helpful.
    - Sim2Real is not a waste of time and money, it’s just the opposite: it saves time and money and we need more research on better simulators and Sim2Real methods.
- Discussion
  - Welinder: But this is the point for sim2real research. current methods are “existence proofs”
  - Gupta: sim is a lower barrier for crappy work. sim2real should be focusing on generalization
  - Goldberg: sim is a necessary step. Analogy with CV has helped robotics research. sim data can achieve much higher scale than real
  - Atkeson: real2real and self2self (generalization across time) are higher priority than sim2real
  - Gupta: need to agree on a definition of “sim2real working”
  - Welinder: real2real can benefit from sim2real
  - Atkeson: sim2real is not democratizing research - sim is also very expensive (see OpenAI)
  - Welinder: cost of first time solving a task is higher than what follows
  - Goldberg: getting real hardware to work is far more challenging than using sim and not easily repeatable.
  - Atkeson: ML uses GD, but RL w/ GD is inefficient, and using sim makes inefficient algorithms popular
  - Gupta: a good learning algorithm should not use resets. democratization is double edged sword - too many papers on sim only

Debate 2 | “Sim2Real is old news” | For: Martha White, Jan Peters. Against: John Leonard, Gregory Dudek

- Opening
  - Peters
    - classical sample-based planning is already “sim2real” - using sim as a forward model
    - DR has been around since 1960s
  - Leonard
    - “simulation is doomed to succeed”
    - sim edge cases are useful, only if sim2real works
    - robots should “close the loop” between sim and real
  - White
    - define sim2real as training policy in sim and deploy w/o adaptation in real
    - same with approximate DP
    - in RL, small model errs -> large value errs
    - learn “selective models” and value-aware models
  - Dudek

        - identifying sim2real as a new subfield is good for research synergy
        - sim2real is emerging in a way similar to other subfields, like CS, controls, and deep learning
        - sim2real in robotics has a set of unique challenges, priorities
- Discussion
    - Peters: learning in sim exploits in sim
    - White: is sim2real a problem or strategy?
    - Leonard: sim2real is more of a strategy, but need to do sim2real2sim, which can mitigate sim artifacts
    - Peters: but real2sim is just adaptive control
    - Dudek: sim2real is dealing w/ problems much harder than adaptive control, which uses an already informative/narrow model
    - White: but this is just Model-based RL
    - Dudek: aggregating tools from different fields under the umbrella of sim2real is useful
    - Leonard: sim2real extends historical contributions and could brings other robotics subdisciplines together
    - White: but focusing on sim2real can be limiting, and ignores “internal-models” which is not a “simulator”
    - Peters: should always be suspicious of models. there’s a need for guarantees on sim bias
    - Leonard: is sim going to be programmed or learned? black box might not be a good idea
    - White: sim2real -> learning with learned models? need to embrace online learning
    - Dudek: hierarchical simulators, safety guarantees, transfer learning, etc are all sim2real

Debate 3 | “Accurate simulators are necessary for sim2real” | For: Karen Liu, Dieter Fox. Against: Anca Dragan, Shuran Song

- Opening
    - Fox
        - realism is measured on input-output relationships
        - sloppy sim can generate behavior that’s “hard to forget”
        - accurate sim lead to better generalization - broad domains
        - should focus on real2sim, keeping the sim2real gap “manageable”
    - Dragan
        - realism/accuracy is about parameters, and DR/diversity can address this
        - focus on accurate sim gives “false sense of security”
        - accuracy is not the end goal, as long as sim is useful
        - many models can induce “best response,” so “accurate models” is a red herring
        - models are task-specific, while simulators are task-agnostic, so models are more useful than simulators
        - especially true for humans - can’t accurately simulate humans
    - Liu
        - this debate is really about investing more on improving sim or improving sim2real
        - realism is about predictions
        - learning policies to compensate for sloppy sim is not worth the cost
        - in the long run, better sim outweighs better sim2real
    - Song
        - accurate physics and params is very hard to achieve, and unnecessary
        - simple and fast sim can give diversity
        - combine simple/inaccurate sim w/ real world learning
        - residual learning can pick up the slack
- Discussion
    - Song: need to identify what’s easy and what’s hard to model
    - Dragan: accuracy at the abstraction level is misleading. inaccurate models can still give good plans
    - Fox: less accurate sim means crazier DR. good sim means more informed DR
    - Dragan: there are going to be errors no matter what. so it’s important to estimate uncertainties. sometimes improving sim acc leads to worse planning perf
    - Liu: (naive) DR is not sufficient for tasks with small feasibility sets. we don’t know what’s important for a task initially
    - Song: humans always use inaccurate models, and it’s more about narrowing uncertainty
    - Fox: models vs. sim - sim needs to be accurate, models not necessarily. DR should be more realistic, which can be facilitated by more accurate sim
    - Liu: babies learn in a perf simulator, but sensing is imperfect
    - Liu: need more unified structure of how to improve sim with data, especially for multi-physics

Panel

- Dragan, Dudek, Atkeson - should focus on task-oriented optimization for sim2real
    - collect data only when it’s useful
    - Dragan - “This is an example of improving accuracy leading to worse inference https://arxiv.org/abs/1903.03877 ; we think it’s because maximizing probability of data given model is different than maximizing probability of correct inference given the model https://arxiv.org/abs/1903.03877”
- Dudek - meta thought - debates on “sim2real” on extreme vs. actual beliefs on the subject
- Differentiable simulators?
    - Dudek - useful but not sufficient
    - Atkeson - only need Jacobian of sim to be useful for task
    - Goldberg - sim shouldn’t be deterministic - multiple derivatives
- Dudek - parameterizing sim might be harder than improving sim perf
- Should sim2sim results be rejected?
    - Atkeson - no but they need to be more convincing
    - Song - sim lowers barrier of entrance, and their contributions can still be useful
    - Dragan - pandemic makes real experiments harder and makes ppl take a more serious look at sim
    - Goldberg - sim gives more reproducibility than real. bias also exists in real exps

    - Welinder - trusting only sim results is hard b/c sim is simplified. focus on what to do to increase trust / likelihood of transfer
    - Dragan - reviewers might overweigh the presence of real experiments
- Leonard - part of intelligence is prediction, and sim can play a big part
- Dragan - model vs sim - different needs/requirements
- Liam Paull - sim2real shouldn't be attached to learning algorithms - that's a separate issue. sim2real is not just about RL
    - Atkeson - agreed. again, self2self and real2real are more important
    - Golemo - but sim design is always related to the task
    - Atkeson - sure but most widely used sim is "general-purpose" and specialized for what the sim designer prioritized, not about the task
    - Liu - might be a way to balance agnostic and oriented sim. also, gradients often don't exist for important events
- Welinder - high chance that sim we have today is all we need for robots to do useful things
    - sim can especially benefit meta-learning
    - Dudek - disagree. sim is still very limited for a lot of tasks.
- Goldberg - robot-specific and task-specific aspects need sim tuning in the loop
- Song - on which level do we adapt sim?
    - Goldberg - not a lot of current sys-id techniques work in stochastic models
    - Song - should we tune small parameters or large parameters?
- Leonard - it's "hopelessly difficult to simulate underwater robotics" - perception, motion, hydrodynamics
- Liu - invest more in sim perf or in control techniques? the community seems to be placing more weight on the latter. need more work on fluids, deformable objs, and other physics
    - Dudek - ppl who work on sim perf are in large organizations
    - Liu - no scientific reward for research in sim perf
    - Atkeson - in computer animation, distillation of complex sim into crude parametric models is common practice. for robotics, there should be many sims at varying levels of details, and they should all be able to improved by data
- Atkeson - should focus on making models useful for specific tasks, and not information gain
    - Song - but this is limiting / not generalizable. information gain can be useful for meta-learning
- Ankur Handa - should we have a separate conference on sim2real?
    - Leonard - maybe. the robotics field has gotten much bigger. sim2real can bring ppl together. maybe another grand challenge on robotics?
    - Dudek - workshops do pretty well on this.
    - Atkeson - sim2real divorced from robotics is not useful
- Juan Higuera - for modular/task-specific sim, how do we set up sim for specific tasks?
    - Liu - sim landscape is too fragmented. advances in CG are not used in robotics, b/c no common abstraction hierarchy. maybe some tools similar to CAD in engineering that's easily shareable. for customizing physics for specific tasks, maybe more data-driven approaches
    - Fox - content creation for tasks is a big problem. robotics still lacks diverse data - generate variety of scenes
    - Leonard - systems need to be self-aware of its limits
- Sebastian Hoefer - what is the ImageNet of robotics?
    - Goldberg - one example is Dex-Net: a dataset of 3d object models labelled with grasps scored by robustness.
    - Fox - needs to be more general - simulation environments with variability that is in line with real world
- Dudek - simulation allows direct intervention to see the effect of changing parameters
    - Goldberg - yes - simulation facilitates sensitivity analysis and grid search to understand and tune hyperparameters. This can increase the potential that physical experiments will succeed.