

CORNELL UNIVERSITY

# Physical AI for Society

---

**Symbiotic computing** across humans, robots, and the cloud.

*[symbiotic-computing@cornell.edu](mailto:symbiotic-computing@cornell.edu)*

*[symbiotic-computing.bowers.cornell.edu](http://symbiotic-computing.bowers.cornell.edu)*

INDUSTRY ENGAGEMENT PROSPECTUS

*An invitation to build the future of Physical AI, with us*

---

# The future of Physical AI is **symbiotic**

## A Symbiotic computational substrate

Humans, robots, and the cloud working in symbiosis to help each other solve some of the most important societal challenges: in homes, farms, warehouses, and classrooms.


## A Symbiotic team

Twelve faculty from seven departments and three campuses working in symbiosis: computer systems, robotics, machine learning, computer vision, sensing, law, and occupational therapy.

## A Symbiotic computing stack

Multiple layers, from sensing to perception to control, working in symbiosis as a unified computing stack: one stack for all applications and all hardware.

## The Symbiotic Computing Stack



*Every component is learned, driven by machine learning algorithms. Compute and hardware is drawn dashed because we do not plan to innovate there, for now. Trust and safety, and evaluation, run the full height because they constrain and measure every layer.*

We are building this stack in the open: everything starts with and returns to the open-source community. Come build with us.

**TECHNICAL WHITE PAPER: HUMAN-ROBOT-CLOUD SYMBIOSIS**

# The Symbiotic Computing Team

Twelve faculty, seven departments, three campuses.



**Rachit Agarwal**

Co-director  
Computer Science  
Cornell



**Tapomayukh Bhattacharjee**

Co-director  
Computer Science  
Cornell



**Mark Campbell**

Mech. & Aerospace Eng.  
Cornell



**Tanzeem Choudhury**

Information Science  
Cornell Tech



**Sarah Dean**

Computer Science  
Cornell



**Chris De Sa**

Computer Science  
Cornell



**Udit Gupta**

Elec. & Computer Eng.  
Cornell Tech



**Bharath Hariharan**

Computer Science  
Cornell



**Manoj Karkee**

Biological & Env. Eng.  
Cornell



**Karen Levy**

Information Science  
Cornell



**Joel Stein**

Rehabilitation  
Medicine  
Weill Cornell



**Mabel Taracena**

Entomology  
Cornell

# The Symbiotic Computing Deployments

Symbiotic evaluation will build upon and extend three systems that our team has taken out of the lab and deployed in the wild: homes, commercial orchards, and public roads.



## Care: mealtime assistance in the home

FEAST lets care recipients personalize the robot themselves. We built it over three years of participatory design with Aimee, who has multiple sclerosis, and Benjamin, who has a spinal cord injury, then ran it in their own homes with minimal researcher intervention.

A formative study with 21 care recipients surfaced 46 distinct personalization needs: bite size, approach angle, when to ask rather than act. The lesson was that deployment success turns not on more autonomy, but on whether the user can re-parameterize the robot without researcher intervention.

TWO PRIVATE HOMES · THREE-YEAR CO-DESIGN



## Agriculture: apple harvesting in commercial orchards

Fifteen years of machine-vision systems, end-effectors, and integrated harvesting robots, rigid and soft, tested in university and commercial orchards in New York and Washington across a wide range of cultivars and canopy architectures.

The field located the boundary. Robots pick accessible fruit very effectively; fruit in clusters, or occluded by branches and foliage, remains the open problem. Commercial deployability runs through multi-arm coordination and human-robot collaboration on canopy perception, planning, and manipulation.

NY & WA ORCHARDS · FIFTEEN YEARS OF FIELD STUDIES



## Mobility: all-weather perception on public roads

An instrumented vehicle drove the same 15 km route through Ithaca hundreds of times over fifteen months, through sun, rain, snow, night, and traffic, capturing high-density lidar, camera, and GPS/inertial data with amodal annotations for partially occluded objects.

Repetition was the instrument: one route driven hundreds of times turns weather, light, and season into controlled variables, which no single-pass collection can do. Released as the Ithaca365 dataset (CVPR 2022), it remains in community use for perception in low light and adverse weather. The rig has since collected in Hanoi.

ITHACA, NY & HANOI · FIFTEEN-MONTH LONGITUDINAL COLLECTION

*Each deployment environment produces unpredictable failure modes that labs do not reproduce on demand: involuntary motion and shifting preference in the home, occlusion and canopy variability in the orchard, low light and degraded sensing under adverse weather. The Symbiotic Computing stack must survive these failure modes; Symbiotic deployments and evaluation measure whether it does.*

# The Symbiotic Computing Testbed

Symbiotic evaluation will build upon and extend a multi-stage evaluation framework, typically including four stages: simulation, hardware validation, controlled studies, and longitudinal deployment.

## 01 Testing robustness in simulation

Sweeping operating conditions at scale—world state, sensing quality, contact conditions, and/or robot embodiment. The outcome is quantifiable: performance boundaries, generalization, safety-margin violations, recovery from perturbations, uncertainty calibration.

## 02 Closing the sim-to-real gap: instrumented hardware

The same tasks as in simulation, now on domain-specific instrumented hardware: force-torque, motion capture, safety monitors, body surrogates, real trees. The outcome is quantifiable: contact quality, controller stability, execution repeatability, and fallback behavior under real sensing and latency.

## 03 Controlled studies in realistic settings

The same tasks, now in realistic environments within the laboratories built in-house. They involve real humans, trees, etc., and mixed-initiative interaction, enabling evaluation before longitudinal deployments. The outcome is quantifiable: trust calibration, perceived safety, workload, legibility.

## 04 Longitudinal deployment in target settings

IRB-approved deployments with target users in target settings, under a conservative escalation policy: fuller autonomy only after safety and usability criteria are met. Each deployment ends with a withdrawal test: the robot is removed for a time, and we measure whether its benefits persist. The outcome is quantifiable: task success, intervention frequency, safety-threshold violations, autonomy preservation, technology acceptance.

Each stage will record the same quantities—task success, safety-threshold violations, intervention frequency, trust calibration—using a unified measurement interface with the same definitions and the same analysis code, so that a safety violation counted in simulation is counted the same way in a home. Our open evaluation standard, described below, will fix those definitions across the four stages and across applications.

## An open evaluation standard

The stages, the measurement interface, the shared measures, and our reference results across various applications are work in progress. We will publish these as the **Symbiotic Computing Evaluation Standard**—a runnable benchmark suite—under a permissive license. Anyone can run it and report the numbers: no fee, no sign-off, no permission from us.

# Engaging with the Symbiotic Computing Initiative

## Benefits of engagement

|                                    |                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>01 Research retreats</b>        | Two per year, two days each. Your engineers engage with our team and hear summary reports on the work under way. Members provide input on industry pain points and problems worth exploring, which the team takes into advisement when setting the benchmarks and the research agenda. This is the center of the program. Retreats are open among members. |
| <b>02 Testbeds</b>                 | Simulation, instrumented hardware, and controlled studies. Your hardware and software platforms are integrated into longitudinal deployments with target users, and into full-scale field deployments, run by our team; the results are incorporated into our publications.*                                                                               |
| <b>03 Benchmarks and artifacts</b> | Open-source systems, benchmark suites, and anonymized datasets—released where participant consent and IRB approval permit—all under a permissive license. Many of our team's datasets are already in public use; there are opportunities available to generate specific research data.*                                                                    |
| <b>04 Students</b>                 | Doctoral students working across systems, robotics, machine learning, computer vision, and applications: experience across the entire stack, by training. A named contact among the faculty nearest your work, and internships.                                                                                                                            |
| <b>05 Boot camp</b>                | An annual two-day, hands-on boot camp: working with the end-to-end stack, including everything partners have integrated into it.                                                                                                                                                                                                                           |

\* There are also opportunities to establish a collaborative research campaign—including co-designed deployment protocols and testing methodologies—or to generate specific research data, under a separate sponsored research agreement.

## What engagement does not include

|                             |                                                                           |
|-----------------------------|---------------------------------------------------------------------------|
| <b>No IP rights</b>         | No licenses, no options, no assignment, no right of first refusal.        |
| <b>No deliverables</b>      | No statement of work, no milestones, no reporting obligations to members. |
| <b>No publication delay</b> | No rights on holding papers, and no rights on reviewing them.             |

# The Symbiotic Computing Engagement Tiers

The initiative runs for five years, through 2031.

| TIER                    | CONTRIBUTION                                                                        | INCLUDES                                                                                                                                                                                                                                                                                                                           |
|-------------------------|-------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Founding partner</b> | \$250,000 / year<br>Three-year term<br>Years 4–5:<br>renewable,<br>\$125,000 / year | <b>Limited to three.</b> Everything below, plus first priority for integrating your hardware and software platforms into our longitudinal deployments—run by our team, with results incorporated into our publications—where appropriate to the faculty-defined research agenda. Also, first priority for boot-camp registration.* |
| <b>Founding member</b>  | \$150,000 / year<br>Three-year term<br>Years 4–5:<br>renewable,<br>\$75,000 / year  | Everything below, plus a seat on the advisory group that provides guidance to the faculty each year; priority, after founding partners, for integrating your hardware and software platforms into our longitudinal deployments and for boot-camp registration.*                                                                    |
| <b>Member</b>           | \$75,000 / year<br>Annually<br>renewable                                            | Retreats, benchmarks and artifacts, facilitated student access, and early access to boot-camp registration.                                                                                                                                                                                                                        |
| <b>In-kind</b>          | By arrangement                                                                      | Robot platforms, sensors, edge compute, cloud credits and instrumentation at a stated cash-equivalent value. In-kind contributions may be accepted by arrangement, subject to Cornell's applicable policies and agreement terms.                                                                                                   |

Founding status, partner or member, closes at the first winter retreat, December 2026, and is recognized for the life of the initiative. Members of every tier are listed on the initiative website.

\* There are also opportunities to establish a collaborative research campaign—including co-designed deployment protocols and testing methodologies—or to generate specific research data, under a separate sponsored research agreement.

## Practical notes

|                  |                                                                                                                                                                                                                                                             |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Form</b>      | A membership under Cornell's standard Corporate Membership Agreement, not a research contract. The membership fee is due within 30 days of signing; benefits are structured consistent with the agreement.                                                  |
| <b>Term</b>      | The agreement begins on the date it is signed. Founding terms run the initiative's first three years; years four and five renew annually at the rates above, and founding recognition persists whether or not you renew. Members renew annually throughout. |
| <b>Licensing</b> | Software and benchmark suites are released under permissive open-source licenses; datasets under permissive data licenses.                                                                                                                                  |
| <b>Data</b>      | All human-subject work is IRB-approved, and all data releases are governed by participant consent.                                                                                                                                                          |

## START A CONVERSATION

## Tell us what you are trying to build.

We are most useful to partners who arrive with something concrete: an open question whose resolution would help you, a system you want to take into the field, a capability you need to depend on under load, a claim you need evidence for.

A first meeting is sixty minutes, in person or remote, with both co-directors and the two or three faculty working closest to what you are building. Nothing is required from you beforehand.

|                     |                                                                                                    |
|---------------------|----------------------------------------------------------------------------------------------------|
| <b>Email</b>        | <a href="mailto:symbiotic-computing@cornell.edu">symbiotic-computing@cornell.edu</a>               |
| <b>Co-directors</b> | Rachit Agarwal and Tapomayukh Bhattacharjee, Department of Computer Science                        |
| <b>Website</b>      | <a href="http://symbiotic-computing.bowers.cornell.edu">symbiotic-computing.bowers.cornell.edu</a> |

# Appendix: how the program runs

## How a year runs

|                       |                                                                                                                             |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <b>Winter retreat</b> | Two days. The faculty set the agenda for the next year here, taking member guidance into advisement.                        |
| <b>Summer retreat</b> | Two days at Cornell, where the robots and testbeds live: mid-year summary reports, benchmark releases, live demonstrations. |
| <b>Late autumn</b>    | The advisory group provides guidance to the faculty. Membership renewals.                                                   |

## The first ninety days

|                       |                                                                                                                                                                                            |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Week 1</b>         | Membership agreement executed. A named faculty contact is assigned to you.                                                                                                                 |
| <b>Week 2</b>         | Introductory session with both co-directors and the faculty nearest your work. You describe what you are trying to build; we tell you candidly where we can help, and where we cannot.     |
| <b>Weeks 3 and 4</b>  | The pain points and problems you bring are taken into advisement by the faculty while deciding benchmarks and the research agenda.                                                         |
| <b>Months 2 and 3</b> | Opportunities for providing input on industry pain points and problems worth exploring over meetings every two weeks; faculty take these into advisement when deciding the research focus. |
| <b>Next retreat</b>   | Your engineers attend.                                                                                                                                                                     |

## The advisory group

|                                            |                                                                                                                                                            |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Structure</b>                           | A representative of each founding partner and each founding member, plus both co-directors. Meets each autumn, with standing contact in between.           |
| <b>What the group provides guidance on</b> | Which problems merit emphasis in the coming year. Which pain points matter most as the faculty decide benchmarks and the research agenda. Retreat agendas. |
| <b>What the group does not decide</b>      | The areas and direction of the research. What gets published, when, or by whom. Which students work on what. Hiring.                                       |
| <b>If members disagree</b>                 | The co-directors decide, and give their reasons.                                                                                                           |

# Appendix: frequently asked questions

**Will you certify our system?**

No. We publish the standard, the benchmark suite and our reference results; anyone can run them and report their numbers. A standard anyone can run is more useful than a mark only we can grant.

**How is this different from a sponsored research agreement?**

Membership carries no statement of work and no IP terms. If you need defined deliverables or specific IP terms, that is a separate sponsored research agreement. The two are complementary, and we are glad to discuss both.

**Who owns what our engineers contribute during a collaborative research campaign via a sponsored research agreement?**

Ownership is set in the sponsored research agreement, which is executed on mutually agreed terms before the campaign starts.

**What does a collaborative research campaign produce?**

Mutually agreed-upon deliverables defined in its sponsored research agreement, such as specific research data, software, or testing.

**What if we need exclusivity?**

We cannot offer it. Everything goes to every member and to the public, which is what allows the work to be published and keeps our results independent of who funds them.

**How do I justify this internally?**

Three points. Longitudinal human-subject evidence of a kind that is difficult to generate internally, through our publications and public releases; a recruiting channel into a scarce engineering profile; and a voice, through guidance, in a layer your products are likely to depend on.

**Where does the contribution go?**

Almost entirely into people and the testbed: faculty, doctoral students and postdocs working across the stack, testbed instrumentation and operations, and the costs of running deployments (recruitment, consent processes, and in-home support). Cornell's standard administration of membership fees applies.

**What happens at the end of five years?**

We move on to the next frontier of open questions. A stated end date means the work is judged against it rather than continuing by default. Artifacts, benchmarks and datasets remain public, and their maintenance transitions to the open-source community, as with any mature open-source project.

**We are a startup and cannot afford this.**

Talk to us anyway. In-kind contributions count, and the right problems and collaborations matter more than the size of the contribution.

# Physical AI for Society

---

**Symbiotic computing** across humans, robots, and the cloud.

*[symbiotic-computing@cornell.edu](mailto:symbiotic-computing@cornell.edu)*

*[symbiotic-computing.bowers.cornell.edu](mailto:symbiotic-computing.bowers.cornell.edu)*

## **INDUSTRY ENGAGEMENT PROSPECTUS**

*An invitation to build the future of Physical AI, with us*

---

*Cornell University is an equal-opportunity employer and educator.*

August 2026