Getting Robots Deployed at Scale; Robot Demo Videos are Not Sufficient

Rodney Brooks, September 17th, 2025

robust AI

MIT CSAIL

rodneybrooks.com/blog -- @rodneyabrooks.bsky.social

It Is A Golden Age For Robotics Deployments

- Amazing amounts of compute available onboard robots
  - GPUs galore, deep learning models, SLAM is mature, real-time people&pose detection, trainable image labeling classes
- Commodity hardware has been driven by other product classes
  - Smart phones:
    - Cameras with and without depth sensing hardware overlaid
    - Tiny GPS chips
    - Tiny IMUs
  - Electric scooters
    - Hi torque brushless DC motors, including hub motors
    - Hi capacity battery cells
  - Smart appliances
    - $3 Embedded 32bit processors for motor control and dynamic models that outclass the main processors of 15yr ago robots.
    - Plentiful touch screen options
- Seventy years of AI tools can be harnessed (newer ones with more care)

But, robotics hype promises robot Nirvana any moment now.

Nirvana is not actually imminent.

My History

- 1977 finished MS in Machine Learning in Australia, joined Stanford AI Lab for a 1981 PhD in Computer Science
- 1981 Post-doc at CMU and then MIT AI Lab
- 1983 CS Faculty at Stanford
- 1984 Co-Founded Lucid Inc., Palo Alto, AI Software
- 1984 CS Faculty at MIT, AI Lab Director 1997, Founding CSAIL Director 2003
- 1990 Co-Founded IS Robotics, planetary rovers
- 1992 Co-Founded Artificial Creatures, sentient toys
- 1997 IS Robotics and Artificial Creatures merged to become iRobot
- 1999 Co-Founded Robotics Ventures, a VC firm
- 2008 Co-Founded Rethink Robotics, humanoid-ish factory robots
- 2019 Co-Founded Robust.AI, 4th gen AMRs for logistics and manufacturing

My Major Commercial Robot Families

8 digits

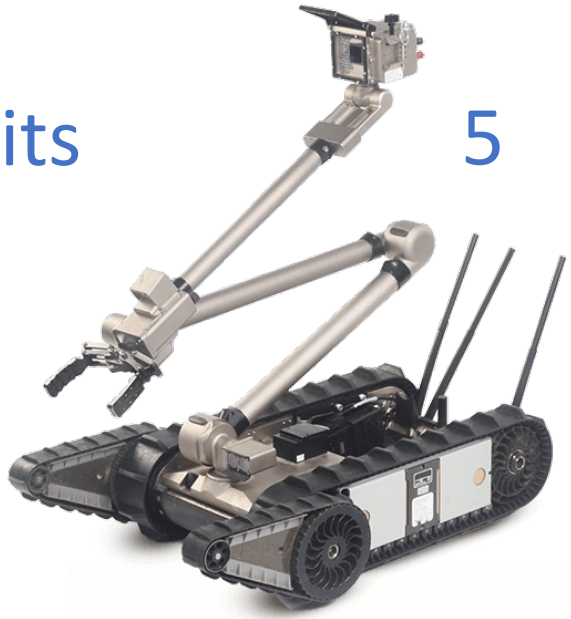

5

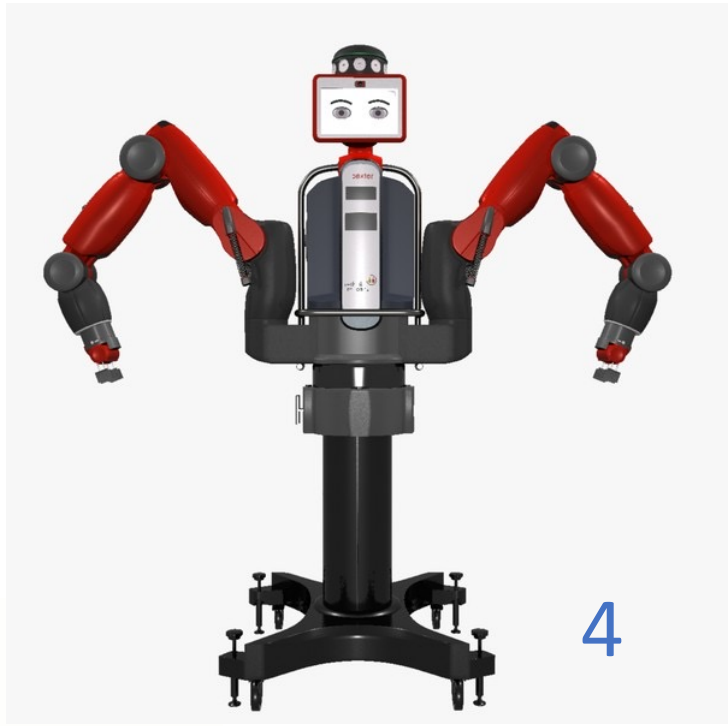

4

4

My Three Laws of Robotics…

1. The visual appearance of a robot makes a promise about what it can do and how smart it is. It needs to deliver or slightly over deliver on that promise or it will not be accepted.

My Three Laws of Robotics…

1. The visual appearance of a robot makes a promise about what it can do and how smart it is. It needs to deliver or slightly over deliver on that promise or it will not be accepted.
2. When robots and people coexist in the same spaces, the robots must not take away from people's agency, particularly when the robots are failing, as inevitably they will at times.

My Three Laws of Robotics…

1. The visual appearance of a robot makes a promise about what it can do and how smart it is. It needs to deliver or slightly over deliver on that promise or it will not be accepted.
2. When robots and people coexist in the same spaces, the robots must not take away from people's agency, particularly when the robots are failing, as inevitably they will at times.
3. Technologies for robots need 10+ years of steady improvement beyond lab demos of the target tasks to mature to low cost and to have their limitations characterized well enough that they can deliver 99.9% of the time. Every 10 more years gets another 9 in reliability.

My Major Commercial Robot Families

8 digits

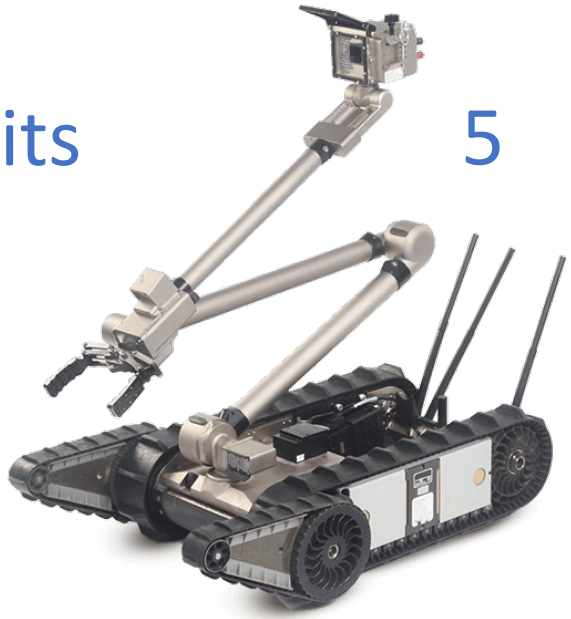

5

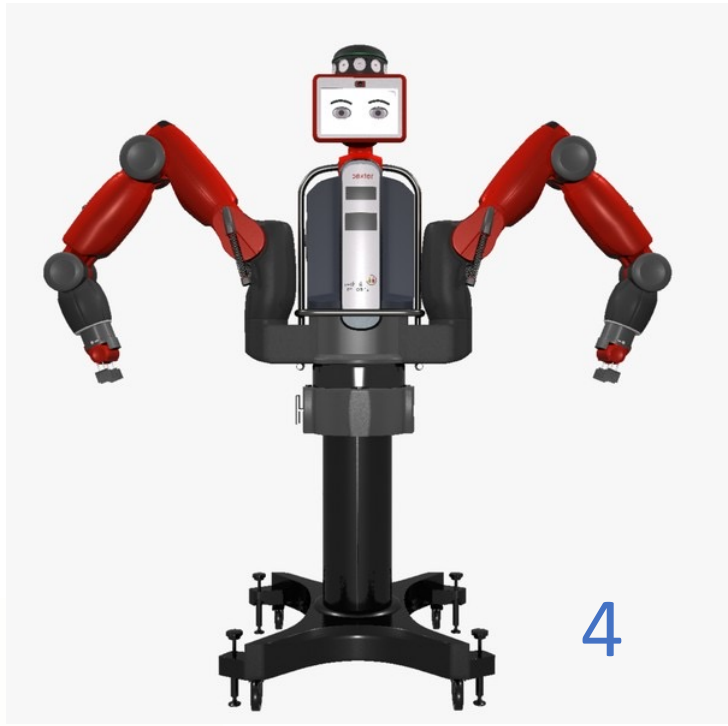

4

4

Carter is safe and approachable

leveraging familiar forms used in warehouses

DHL WAREHOUSE 6:30 AM

The Hat: *Integrated unit that houses*

- *360° vision system*
- *Main compute - NVIDIA Jetson*
- *High-visibility matrix displays to communicate status and next steps*

Custom Shelving:

- *Flexible and configurable*
- *Standard 48” long x 24” wide*
- *Field adjustable*
- *Put to light LED strip*

The Base:

- *Motor controllers & electronics*
- *Main controller board*
- *Batteries and power management*
- *Whole-robot force sensing*
- *Patented mechanical and control holonomic drive design*

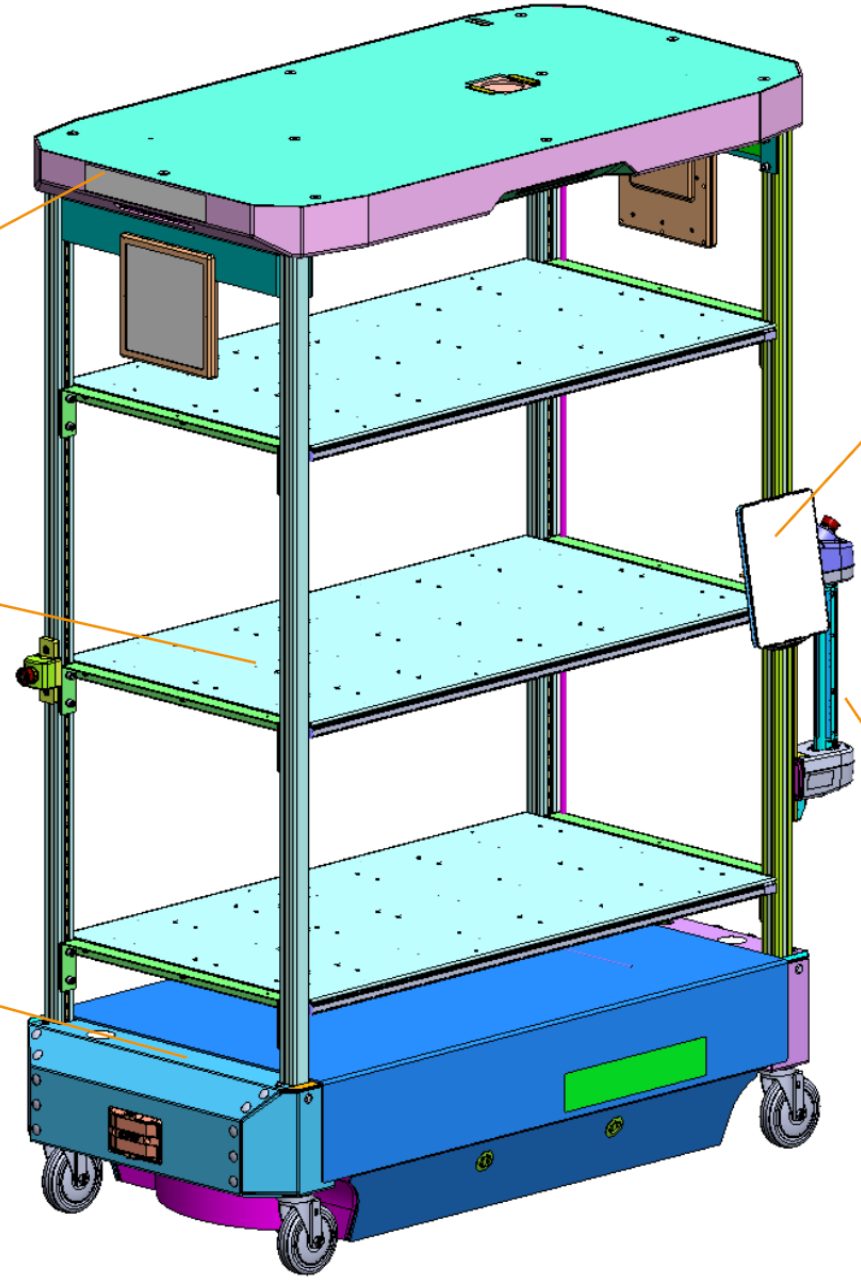

The Interaction Zone:

- *Instructive touch screen*
- *Built-in barcode scanner*

The Force Sensing Handlebar:

- *Light weight robot control*
- *Effortless, single-handed steerability*
- *Ergonomic for shorter and taller workers*
- *Patented design*

robust AI

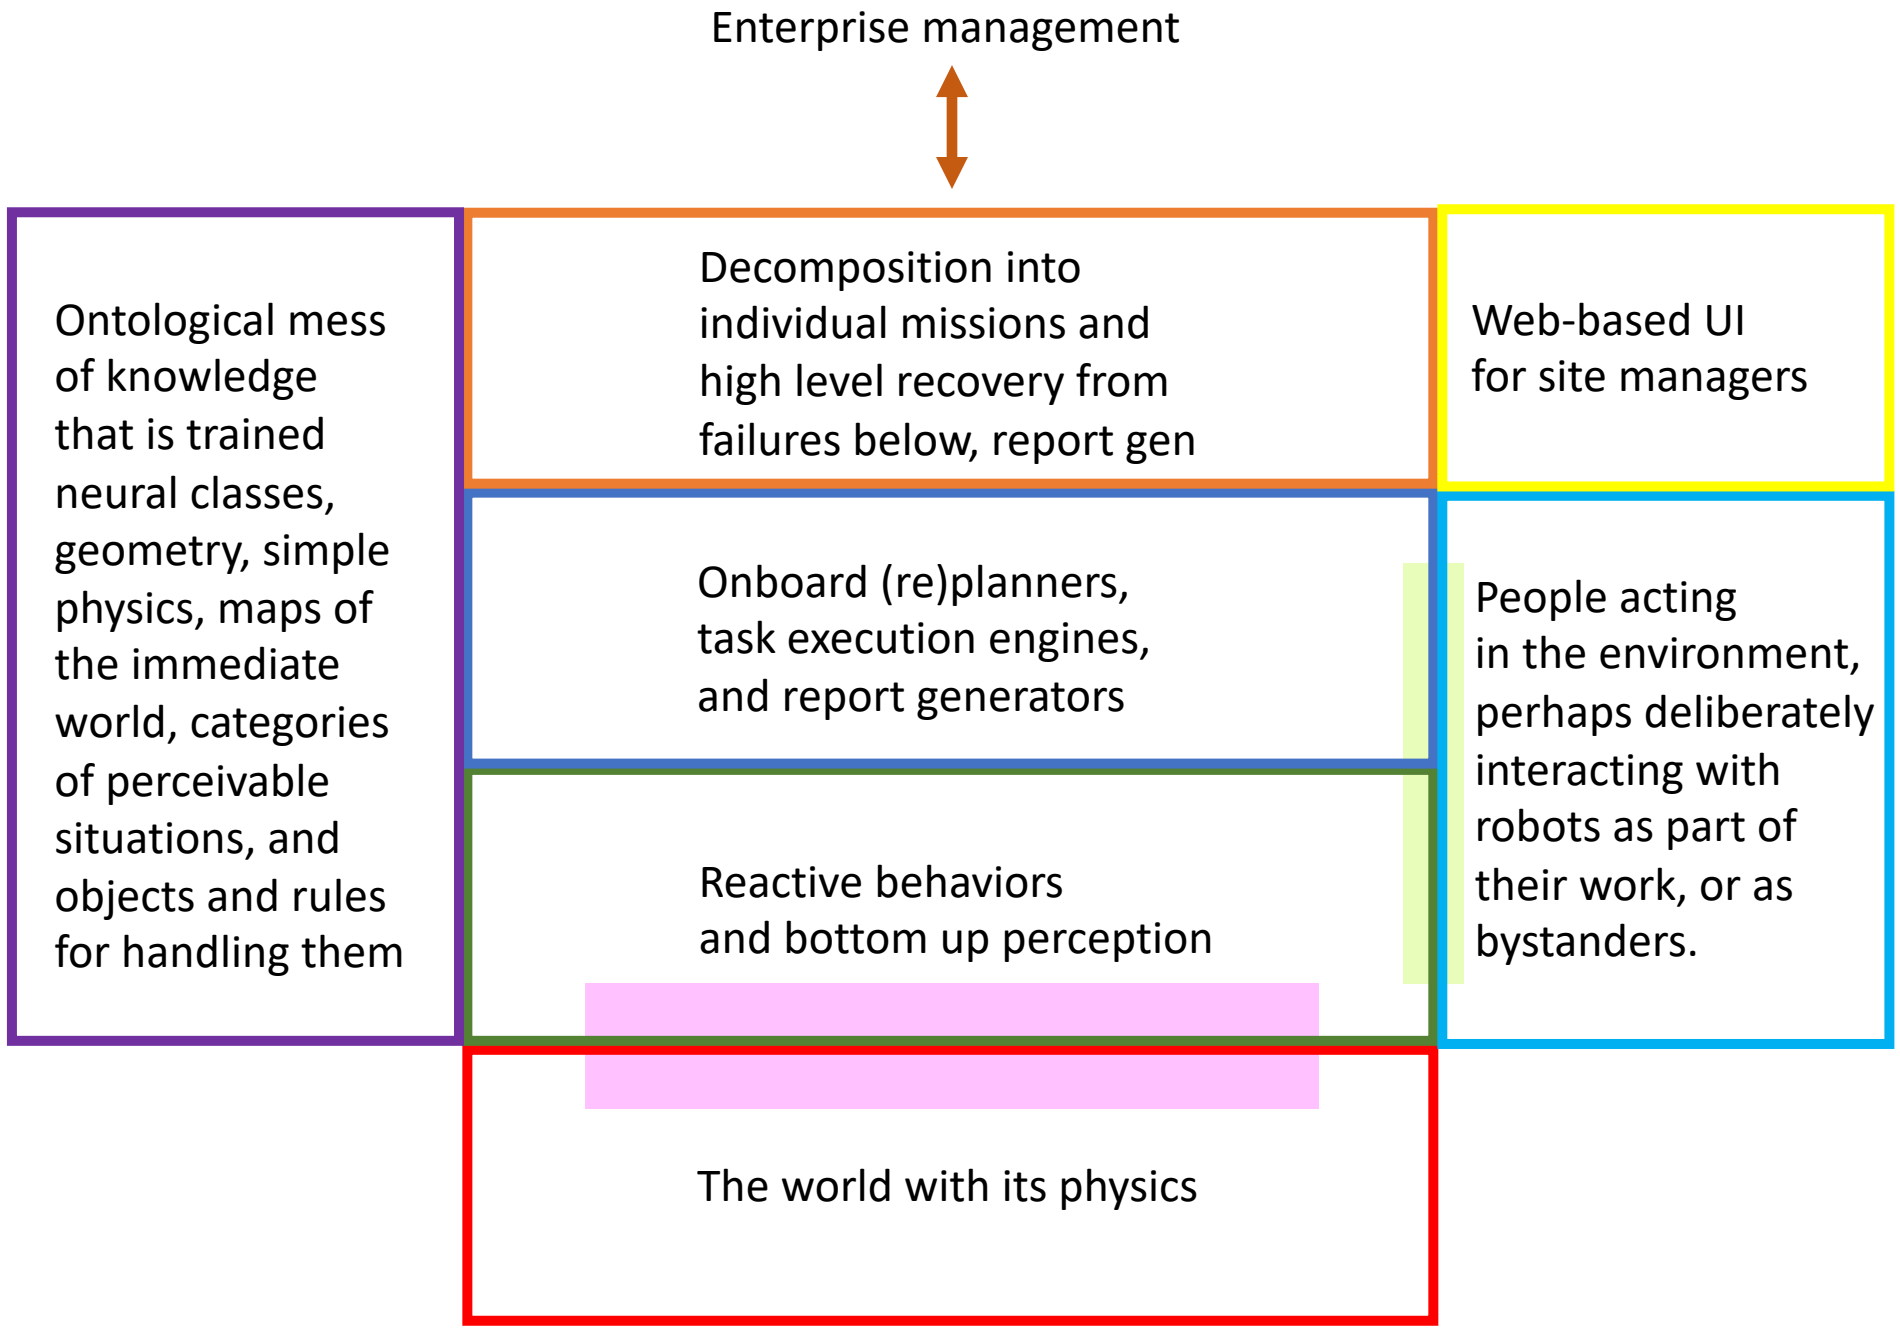
Enterprise management
Ontological mess of knowledge that is trained neural classes, geometry, simple physics, maps of the immediate world, categories of perceivable situations, and objects and rules for handling them
Decomposition into individual missions and high level recovery from failures below, report gen
Onboard (re)planners, task execution engines, and report generators
Reactive behaviors and bottom up perception
The world with its physics
Web-based UI for site managers
People acting in the environment, perhaps deliberately interacting with robots as part of their work, or as bystanders.

Using Ready-For-Primetime Capabilities

- Building vision-based maps automatically (fully **matured** last 3 years)
- Planning geometric 2D paths in fuzzy maps (20 years ago)
- Reactive navigation in dynamic environments (**matured** last 10 years)
- Wheeled (including holonomic) locomotion (recent holonomic++)
- Physical interaction with people (last 15 years++)
- Fleet level management and optimization (**matured** in last 10 years)
- Labeling images with actionable semantic categories (last 10 years++)
- Tracking and categorizing human whole body motion (last 5 years++

Using Ready-For-Primetime Capabilities

- Building vision-based maps automatically (fully **matured** last 3 years)
- Planning geometric 2D paths in fuzzy maps (20 years ago)
- Reactive navigation in dynamic environments (**matured** last 10 years)
- Wheeled (including holonomic) locomotion (recent holonomic++)
- Physical interaction with people (last 15 years++)
- Fleet level management and optimization (**matured** in last 10 years)
- Labeling images with actionable semantic categories (last 10 years++)
- Tracking and categorizing human whole body motion (last 5 years++

Only Very Recently, Practical vSLAM for 4th Generation AMRs

- The GPU bounty that has come to edge computation for running neural inference is the key to this.
- Computation weighs only about a pound.
- Uses less power than the motors and the displays.
- Lets us run multiple stereo flows at 10 frames a second
  - get depth and unique abstract visual features
- “AI Enhanced” perception, e.g., inside stereo algs, floor models, etc.
- Can have fleets all send updates to the central map in the cloud
- Initial mapping is fast
- And when the robot is not localizing those same GPUs can still be used for running neural inference recognizing objects and people behaviors of interest

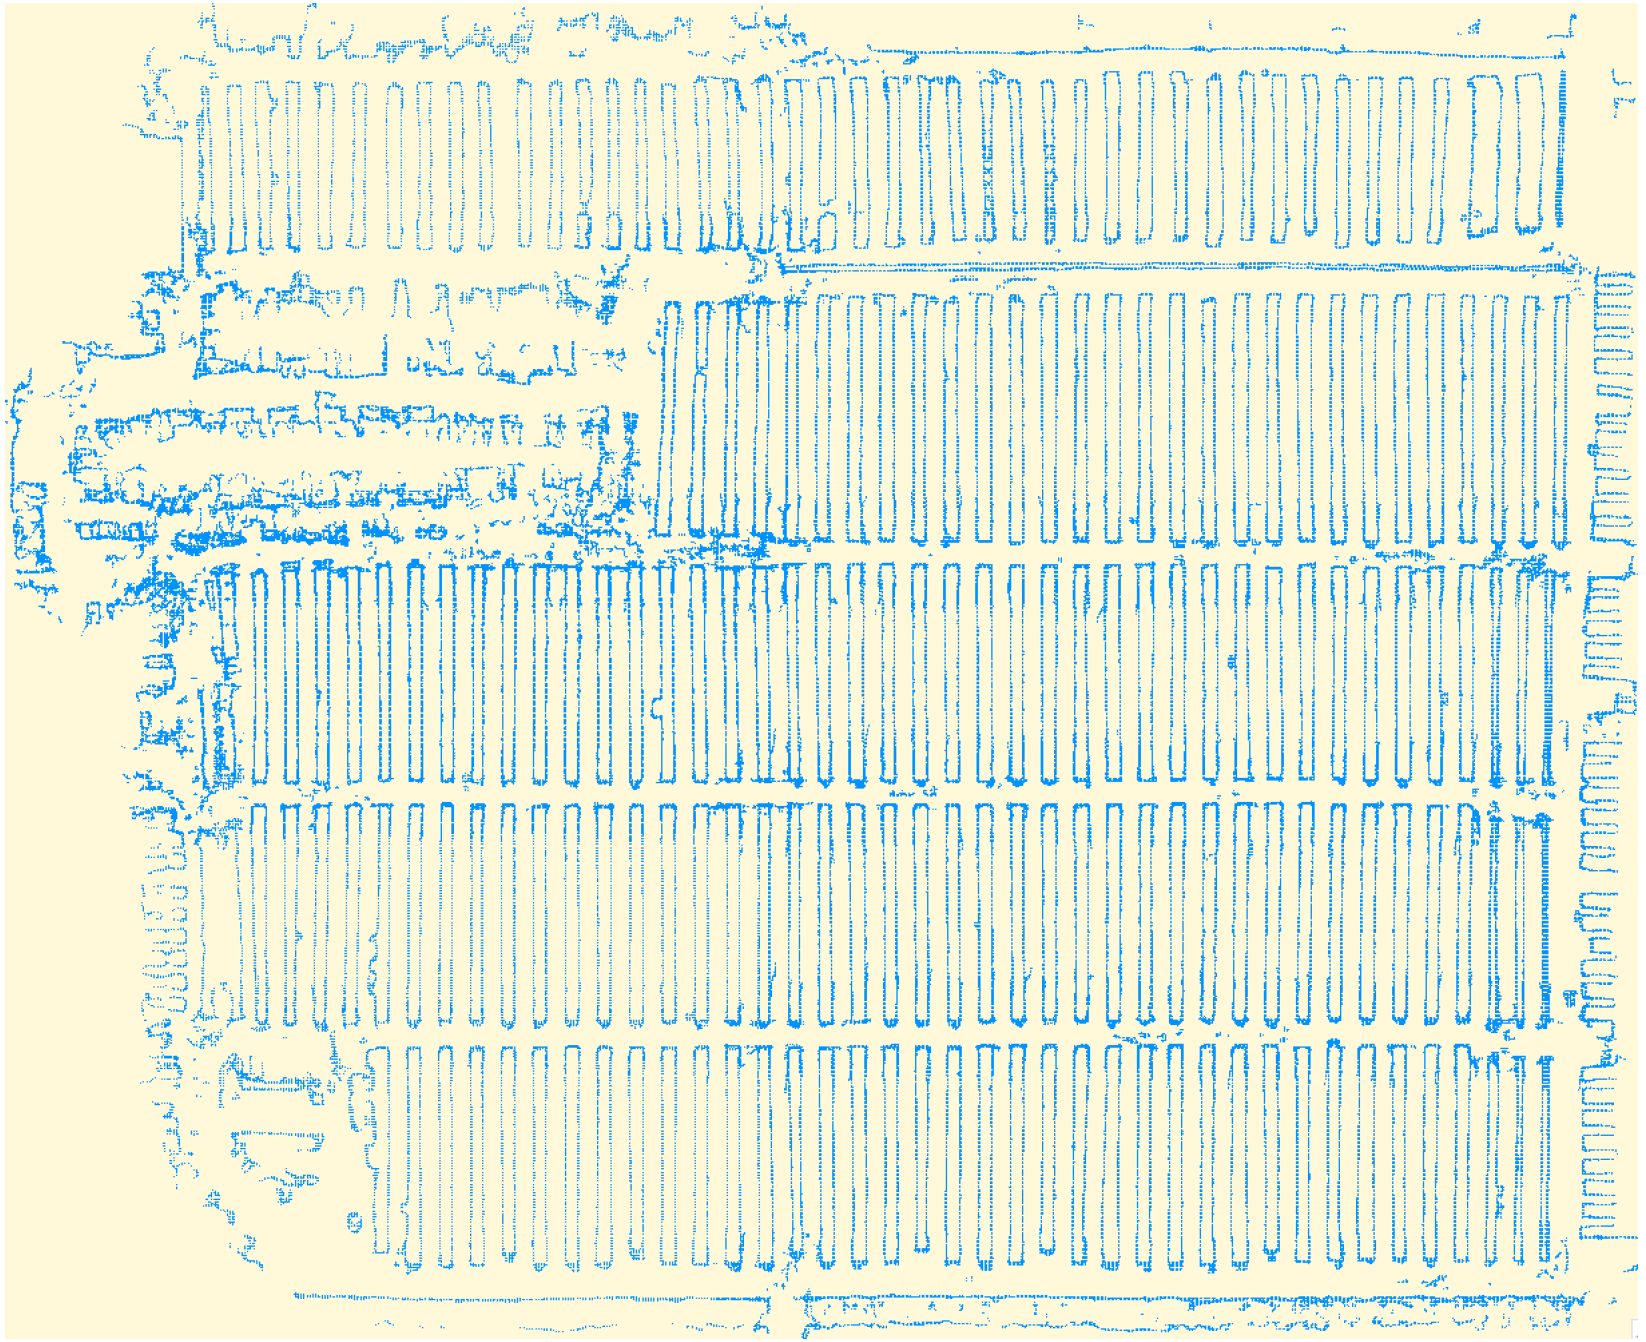

2024: vSLAM

Uses lots of GPUs to do quad directional stereo vision from floor to 7 feet high.

Robots localize to within 5cm at 10 frames per second

258,000 square feet mapped in three hours

Using Ready-For-Primetime Capabilities

- Building vision-based maps automatically (fully **matured** last 3 years)
- Planning geometric 2D paths in fuzzy maps (20 years ago)
- Reactive navigation in dynamic environments (**matured** last 10 years)
- Wheeled (including holonomic) locomotion (recent holonomic++)
- Physical interaction with people (last 15 years++)
- Fleet level management and optimization (**matured** in last 10 years)
- Labeling images with actionable semantic categories (last 10 years++)
- Tracking and categorizing human whole body motion (last 5 years++

Using Ready-For-Primetime Capabilities

- Building vision-based maps automatically (fully **matured** last 3 years)
- Planning geometric 2D paths in fuzzy maps (20 years ago)
- Reactive navigation in dynamic environments (**matured** last 10 years)
- Wheeled (including holonomic) locomotion (recent holonomic++)
- Physical interaction with people (last 15 years++)
- Fleet level management and optimization (**matured** in last 10 years)
- Labeling images with actionable semantic categories (last 10 years++)
- Tracking and categorizing human whole body motion (last 5 years++


robust AI
robust AI


CARTER
TO
IDLE
XAP001
robust


robust AI
robust AI

AI tool chest we use at Robust.AI

- Control theory
- Kalman filters
- Systems dynamics
- Geometric path planning
- Behavior-based systems
- AI-based optimization toolkits

- Visual SLAM
- Neural enhanced stereo
- Neural inference for image labelling
- Synthetic data generation
- Humungous #s of GPUs
- Full body tracking
- Human behavior recognition

AI tool chest we use at Robust.AI

- Control theory
- Kalman filters
- Systems dynamics
- Geometric path planning
- Behavior-based systems
- AI-based optimization toolkits

- Visual SLAM
- Neural enhanced stereo
- Neural inference for image labelling
- Synthetic data generation
- Humungous #s of GPUs
- Full body tracking
- Human behavior recognition

My History

- 1977 finished MS in Machine Learning in Australia, joined Stanford AI Lab for a 1981 PhD in Computer Science
- 1981 Post-doc at CMU and then MIT AI Lab
- 1983 CS Faculty at Stanford
- 1984 Co-Founded Lucid Inc., Palo Alto, AI Software
- 1984 CS Faculty at MIT, AI Lab Director 1997, Founding CSAIL Director 2003
- 1990 Co-Founded IS Robotics, planetary rovers
- 1992 Co-Founded Artificial Creatures, sentient toys
- 1997 IS Robotics and Artificial Creatures merged to become iRobot
- 1999 Co-Founded Robotics Ventures, a VC firm
- 2008 Co-Founded Rethink Robotics, humanoid-ish factory robots
- 2019 Co-Founded Robust.AI, 4th gen AMRs for logistics and manufacturing

Some of the AI Hype Cycles that I Have Seen

- Pre-me:
  - Neural networks (vol 1), 1943, McCulloch & Pitts
  - Turing, 1950 computers fast enough for AGI
  - Neural networks (vol 2) 1950s, Minsky & Rosenblatt
  - McCarthy, 1956 computers fast enough for AGI
  - Pandemonium (AI Agents!), Selfridge, 1959
- Michie, 1959, reinforcement learning (1)
- Minsky, Michie, Nilsson, 1960: search
- It is all first order logic & resolution theorem proving
- Neural networks (vol 3) 1968 Minsky again
- MacHack (search, chess 1)
- Fuzzy logic, STRIPS, knowledge-based systems (revolution in medicine 1)
- Neural networks (vol 4) back propagation
- Primal sketch, 1979, Marr
- Self driving cars 1, 1987, Dickmanns
- 1980s reinforcement learning (2, & 3)
- SOAR (AI Agents!), qualitative reasoning, support vector machines
- Humanoid Robots I, 1992 to 2018
- Self driving cars 2, 1997, Kanade et al
- Deep Blue (search, chess 2)
- Self driving cars 3, 2007, Thrun
- Bayesian inference
- Watson (Jeopardy & revolution in medicine 2)
- Deep learning, neural networks (vol 5)
- Self driving cars 4, 2010+, commercial
- Alpha Go, reinforcement learning (4)
- Generative AI, 2023, neural networks (vol 6)
- Humanoid Robots II, 2024
- AI Agents!, 2024

AI: great expectations

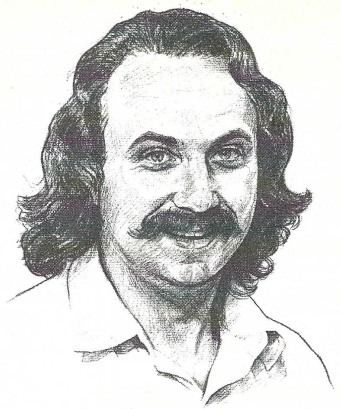

Artificial intelligence (AI) has repeatedly inspired great expectations in people who see the possibilities of applying its techniques. Sometimes it delivers. Sometimes expectations are dashed.

The idea that machines can be rendered intelligent has always been seductive, and demonstrations of limited scope tend to raise greater expectations than hindsight analysis shows were warranted. In his 1949 book *GIANT BRAINS or Machines That Think*, Edmund Berkeley ponders the amazing ability of machines such as ENIAC carrying out 500 multiplications of two 10-digit numbers per second, and envisions machines that would act as automatic stenographers, translators, and psychiatrists.

This pattern is still evident. A few years ago there were high hopes that robots would revolutionize factories. In a way they have, but not in the grand manner predicted in the business plans of start-up companies six or eight years ago. Programming problems, combined with a lack of flexibility, made it impossible to overcome the systems-level problems of integrating assembly robots into the world of manufacturing. Great expectations raised by demonstrations of such robots glossed over other critical aspects of a complete operational enterprise—not the least of which is capital. Robots, indeed, have proven to be useful, but not as useful as was first predicted.

We recently entered the bust side of another set of boom expectations: expert systems. Though mildly successful in industrial applications, once again the expectations of the expert systems industry have not been borne out over time. The too simple representations of the problem domains of expert systems make them extremely brittle when the world diverges from the narrow range of applicability of their knowledge rules. Broadening those representations has been frustratingly difficult. Expert systems are here to stay—though not yet ready to solve all of the world's problems.

With every bust there is a new boom, and in the past year just what the new fashion will be has become clear—neural networks. These networks incorporate an appealing idea in that instead of having to work out all of the details of a task, we'll simply let some randomly organized network of neuron models "learn" through trial and error how to do the right thing. Although neural networks have rarely accomplished anything beyond a computer simulation, business plans are being cranked out for new start-up companies to apply the technology.

But the current neural networks phenomenon is more than just another set of high expectations. This is the second time around for neural networks. It happened in the early '60s. In 1962, a distinguished Stanford professor predicted that computer programming would be obsolete by 1963 because, by then, users would simply converse in English with the front-end neural networks. Since then, there have been a few technical improvements, and computers are much faster, broadening the scope of the applicability and likely successes of neural networks. But, again, they really can't be expected to solve the world's problems. The old-timers, "immunized" the first time around, seem less enamored than the new converts.

I recently worked with a group from industry, offering a detailed explanation of a technical AI method. After some time, the lead technical member of the group—who had no previous exposure to AI—exclaimed, "But that's not intelligence! All you're doing is writing a computer program to solve the problem." Well folks, I'm sorry—but that's all there is. There is no magic in AI. All we do is tackle areas and tasks that people previously were unable to write computer programs to handle. Because we have developed sets of tools and methodologies throughout the years to accomplish this, AI approaches have had a number of good successes. But there is no universal set of magic ideas.

Every so often a new AI development comes along and great excitement ensues as people stumble over themselves, convinced that the key to intelligence has been unlocked. Maybe it *will* happen someday, but I rather doubt it. I don't think there is a single key to intelligence but rather that, unfortunately for both the philosophers and dreamers, intelligence is a vast, complex collection of simpler processes. To develop truly intelligent computers and robots, we're going to have to unlock those processes one by one—causing flurries of great expectations, followed by more modest real successes. This may sound boring and unimaginative, but I find it exciting. Intelligence really is a complex interaction of many things. As we unlock its secrets in the next few years and decades, we will see a constant flow of ideas that have real and immediate practical applications. Finally, when we truly understand AI, it won't seem like just a computer program but will appear as a wondrous testament to the creative genius of evolution. ■

Rodney A Brooks

Dr. Rodney A. Brooks
Associate Professor
Electrical Engineering and Computer Science Dept.
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA

Artificial intelligence (AI) has repeatedly inspired great expectations in people who see the possibilities of applying its techniques. Sometimes it delivers. Sometimes expectations are dashed.

Every so often a new AI development comes along and great excitement ensues as people stumble over themselves, convinced that the key to intelligence has been unlocked.

"Those who don't study history are doomed to repeat it. Yet those who *do* study history are doomed to stand by helplessly while everyone else repeats it."

The mistake that robotics newcomers repeat…

They think robotics is about specifying motions for the robot (and its components)

In fact, it is about deciding what forces to apply when, and managing the total kinetic energy of all parts of the robot.

The motions are the results of those forces and the physics of the robot interacting with its environment.

It is gated by physical contact with other stuff.

Robotics is hard

but

Progress is being made

Capabilities That Are Coming Along…

Engineering

- Robotification of existing indoor wheeled mobility platforms
- Picking stacked boxes from a heavy statically stable robot base
- LLM-based speech interaction with indoor mobile platforms
- Auto-docking mobile robot for recharge or machine interaction

Some innovation

- No-touchscreen nonverbal communication between robots and people
- Robotified farm implements for low crop planting, weeding, harvesting
- Human-safe four legged walking platforms
- Body reshaping to maintain stability when reaching and manipulating

Requires invention

- Whole arm and whole body manipulation
- Human-safe adult sized two legged walking platforms
- Dexterous manipulation at 10% of human capability
- ** Automated last 10 meters of delivery to consumers' homes

The Humanoid Promise

- tall, to get sensors up high and see over human clutter
- skinny to fit in space that humans fit in but with a long reach
- can navigate steps and stairs
- change posture to keep center of mass in support zone
- use their hands like a person can
- walk like a person does
- whole arm/whole body manipulation
- radical change in height of both manipulators and sensor packages

Humanoid companies are punting on the things in dark blue right now and have not delivered on light blue

Living Human(oid)s Also:

- somewhere between omnidirectional and holonomic
- bracing against things in the environment to increase reach
- touching things and surfaces
- touching people
- letting people brace against them
- being back drivable by people
- being able to fall and get up

Outside of current thinking, though humans can do all these and they are critical to our successes as humans

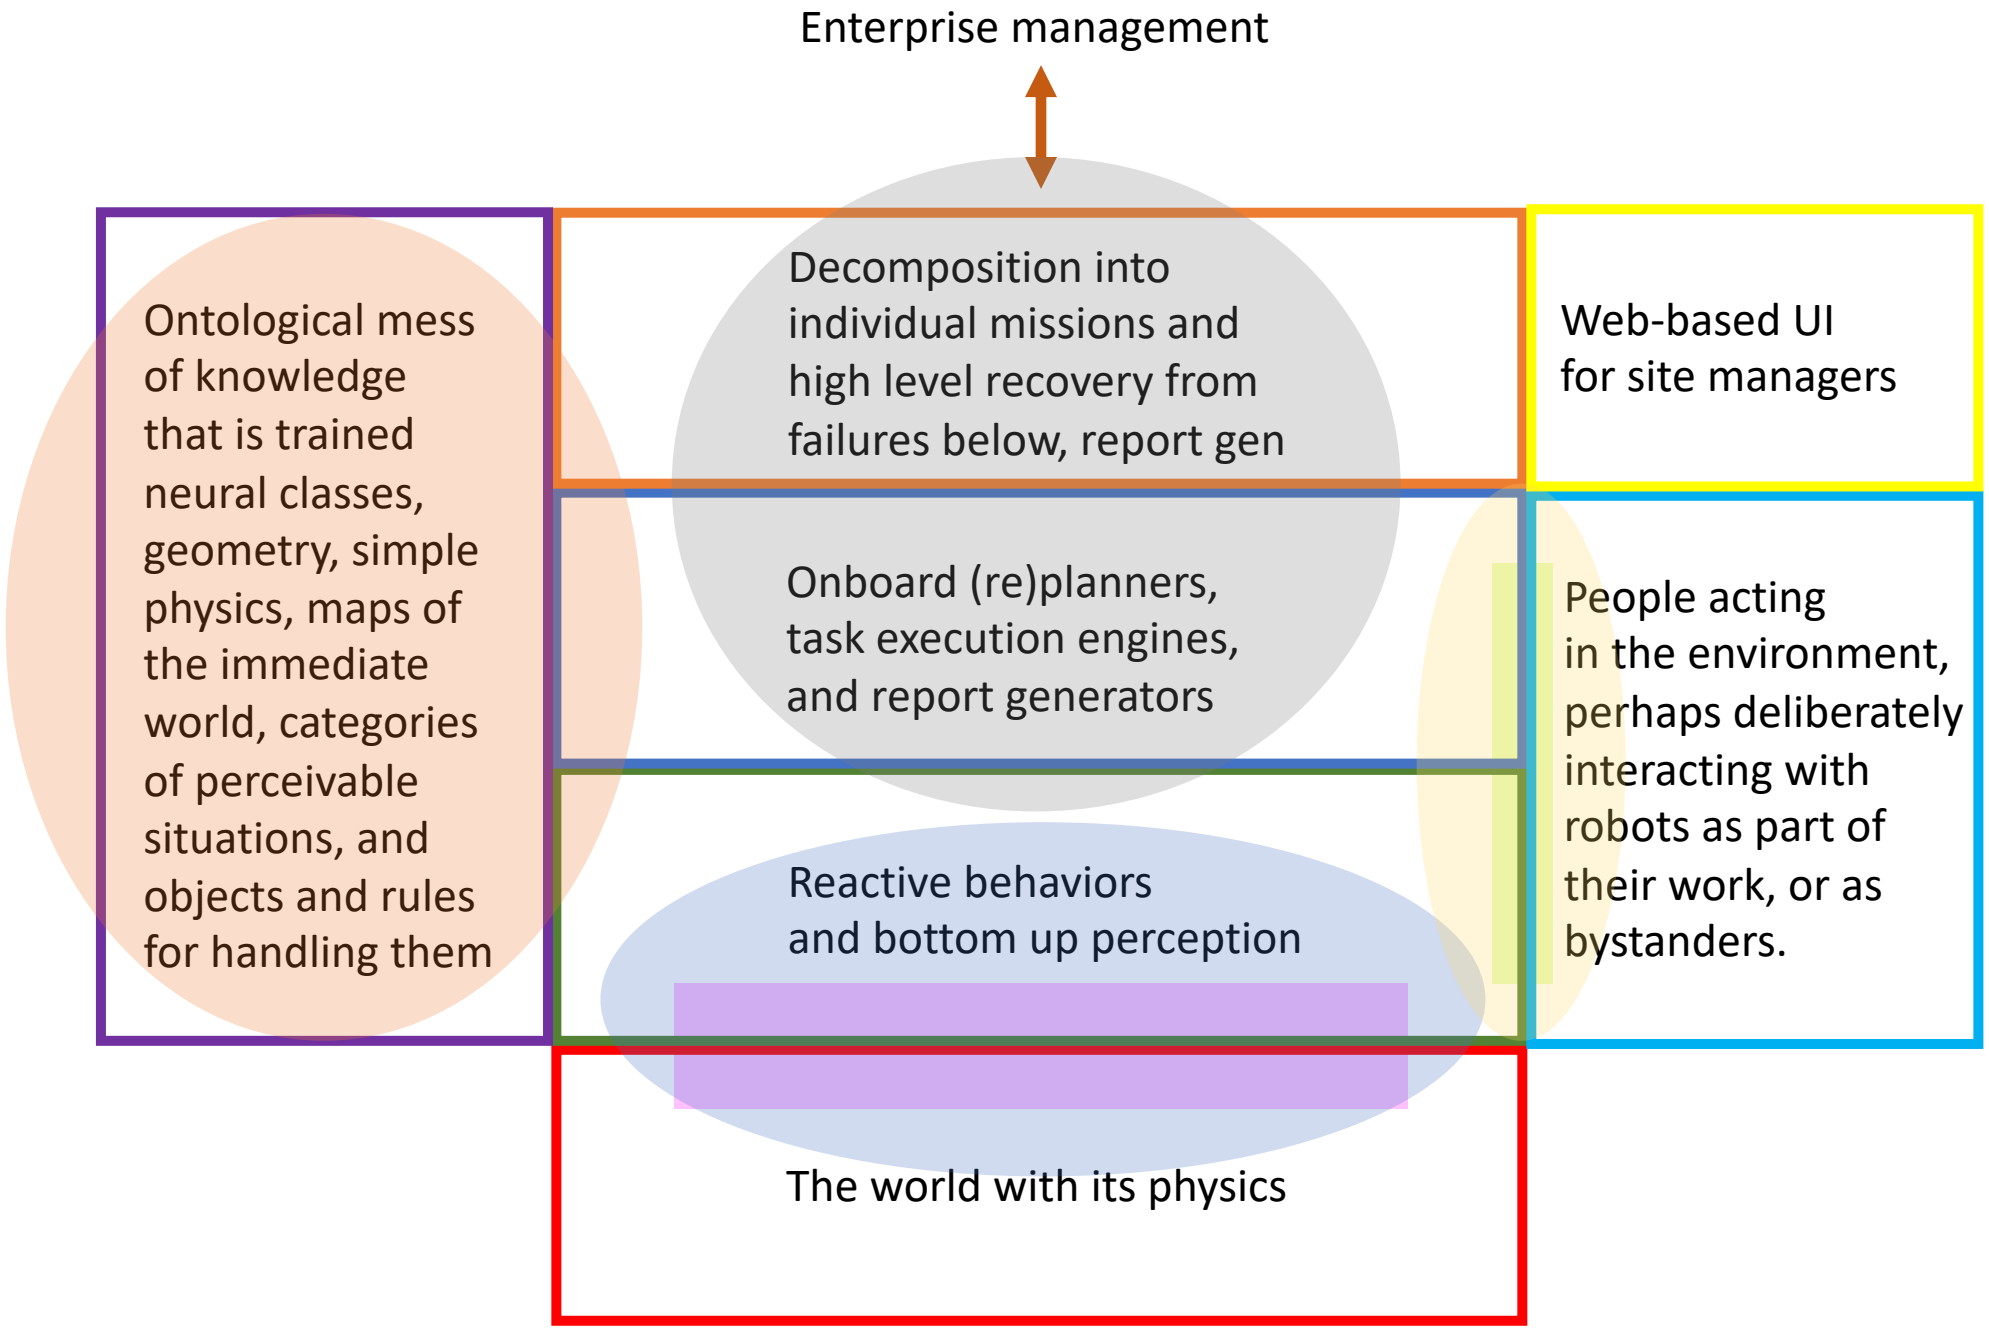
Enterprise management
Ontological mess of knowledge that is trained neural classes, geometry, simple physics, maps of the immediate world, categories of perceivable situations, and objects and rules for handling them
Decomposition into individual missions and high level recovery from failures below, report gen
Onboard (re)planners, task execution engines, and report generators
Reactive behaviors and bottom up perception
The world with its physics
Web-based UI for site managers
People acting in the environment, perhaps deliberately interacting with robots as part of their work, or as bystanders.

Progress is being made

It just takes longer than everyone thinks to get to deployment at scale