Research Area 1: Distributed Intelligence

Nick Roy, MIT
&
Ethan Stump, ARL

Challenges

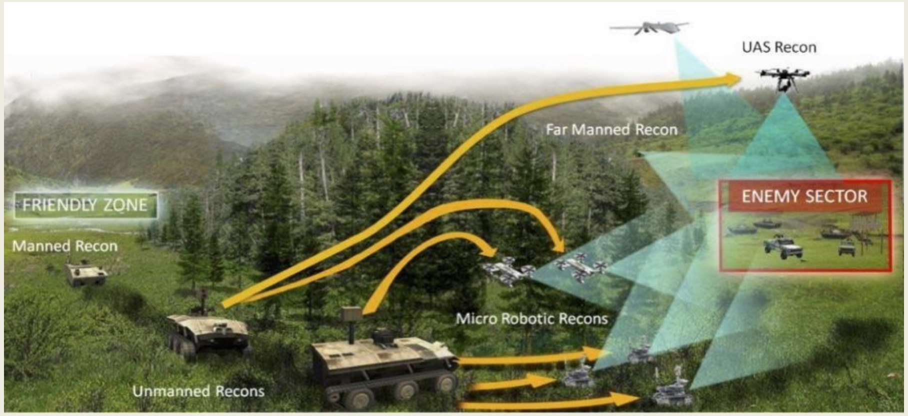

Force Multiplication

- High op-tempo
- Complex, dynamic missions with adversaries
- Distributed agents with heterogeneous sensors
- Large scale needs abstractions
- Intermittent networks with limited bandwidth

Force Protection

- High op-tempo
- Unstructured, dynamic environments with human teammates
- Novel concepts that must be acquired and stored
- Tactical scale demands details
- Information must be human interpretable

How do diverse, embodied agents collectively sense, infer, reason, and plan to support these broad needs?

Key Ideas

- Contextual Abstractions
  - Representations of the world that can be flexibly composed and recomposed efficiently and robustly.
- The Swarm as the Tactical Cloud
  - Ability to share information, representation and capabilities as the problem demands and the data allows
- Perception-Action-Communication Loops for Robust, Responsive and Resilient Swarms
  - Models of sensing, action and learning that are cognizant of constraints in communication resources.

RA1.A Abstractions

- Conventional (deployed) view assumes fixed sensing and planning hierarchy, with fixed inference, planning and learning loops at each level
  - Task planner often hand-coded
- State-of-the-art view acknowledges that inference and planning must reason across levels of abstraction
  - Still fixed hierarchy
- ARCHES view is that abstraction and hierarchy are dynamic and determined by the task at hand.
  - Agents should decide in real-time to modify the world model and the hierarchical model of abstraction

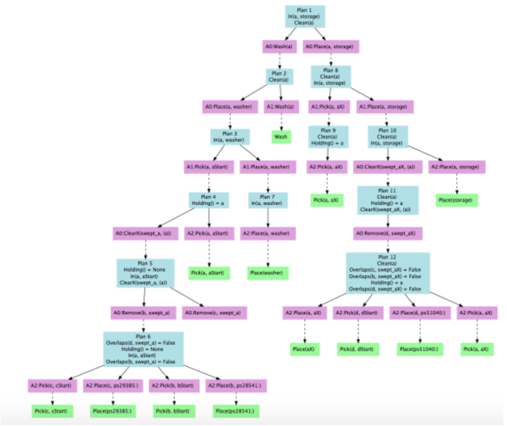

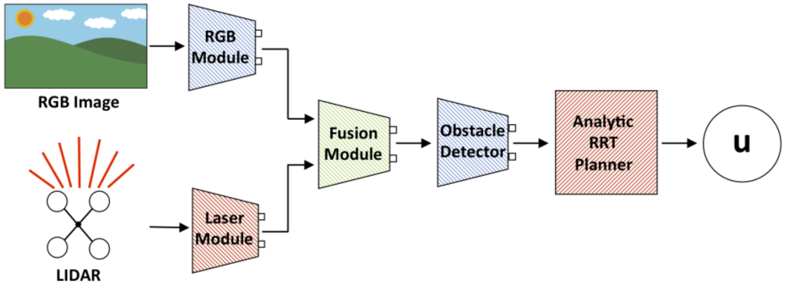

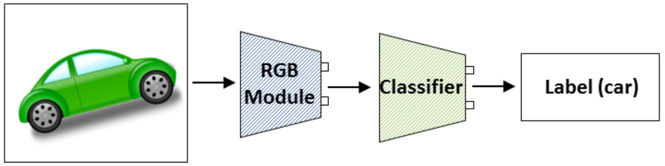

Hierarchical & Composable Models & Contextual Abstractions

- How to infer or learn the structure of these models in the context of the specific environment, mission and distributed resources.
  - RA1.A1 Contextual Perceptual Representations *(Atanasov, Carlone, How, Rogers, Wigness)*
  - **Year 1 Goal:** Multi-modal (metric/semantic/temporal) environment representations
- How to learn with data-efficiency by combining learned components and by combining learned components with conventional models.
  - RA1.A2 Hierarchical, Composable and Adaptable learning *(Levine, Ribeiro, Koppel)*
  - **Year 1 Goal:** hybrid model-based + model-free reinforcement learning for vision-based robotic manipulation and small-scale ground robot
- How to plan in structured models while preserving completeness and optimality
  - RA1.A3 Hierarchical & Composable Planning for Sufficient Optimality *(Karaman, Roy, Tsiotras, Ribeiro, Stump, Hayes)*
  - **Near term Goal:** Hierarchical and composable representations that enable planning with performance and completeness guarantees.

Complex, Collaborative, Distributed Inference & Decision-Making

- **How and when to share and fuse data, in the context of the specific environment, mission and distributed resources?**
  - RA1.B1 Distributed Learning, Inference & Planning *(How, Atanasov, Carlone, Christensen, Rogers, Koppel)*
  - **Near term Goal:** Demonstrate resource-aware inter-robot loop closure detection
- **How to learn across a team when agents do not have identical representations?**
  - RA1.B2 Collaborative Learning in Multi-Agent Networks *(Daniilidis, Atanasov, Levine, Sadler, Wigness*
  - **Year 1 Goal:** Demonstration for functionally relevant metric learning in simulated environment
- **Which abstractions should be used to share information with human partners?**
  - RA1.B3 Interaction with Human Teammates *(Shah, Christensen, Chernova, Loianno, Roy, Bassett, Rogers, Fink, Reardon, Warnell, Cummings, Holder)*
  - **Year 1 Goal:** Human-interpretable induction of dynamic system behavior

RA1.C PAC Loops

- Conventional (deployed) view assumes fixed models and little-to-no interaction between agents
- State-of-the-art view acknowledges a priori training data used to build models, with fixed (or maximal) communication
- ARCHES view is that adaptation and communication are part of the control loop.
  - Agents should decide in real-time to adapt (or not) to new data and to transmit/receive (or not) new data

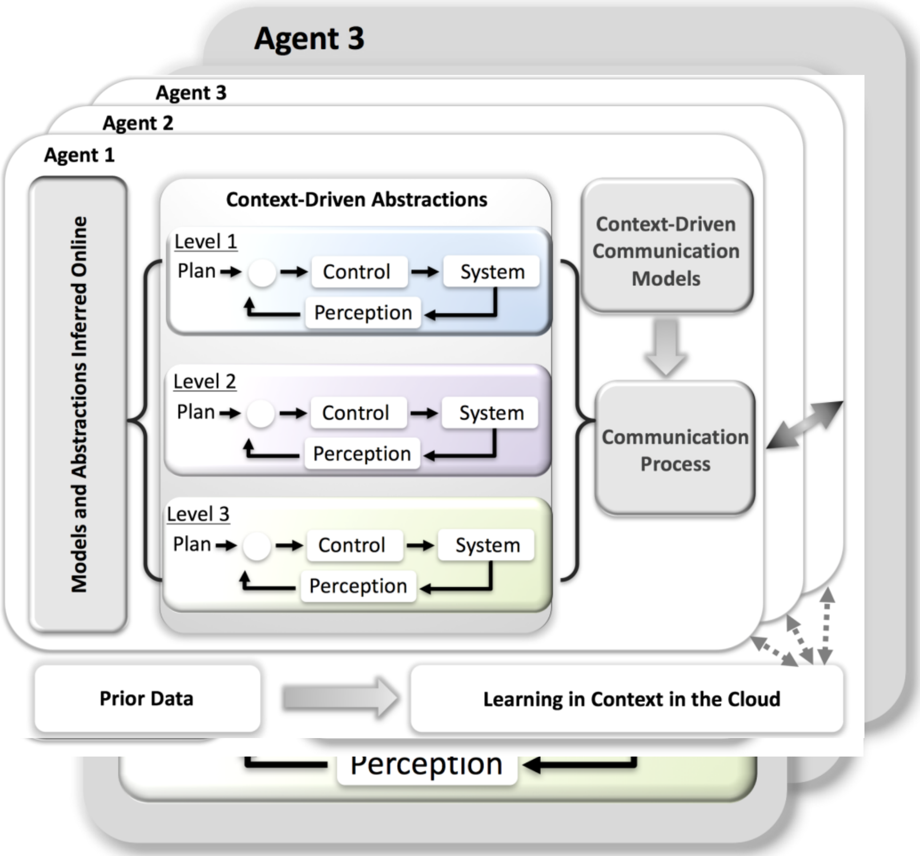

Distributed Perception-Action-Communication Loops

- **How to design these loops in the presence of finite communication and computational resources?**
  - RA1.C1 Joint Resource Allocation in Perception-Action-Communication Loops *(Ribeiro, Pappas, Fink)*
  - **Year 1 Goal:** Design opportunistic communication and resource allocation policies for closing multiple PAC control loops
- **How to use the PAC to adapt the perception system in the context of the specific environment, mission and distributed resources?**
  - RA1.C2 Resource-Aware Perception-Action-Communication Loops *(Tsiotras, Carlone, Roy, Conroy, Dasari, Stump)*
  - **Near-term Goal:** Develop a framework for co-design of perception-action-communication (PAC) loops
- **How to use learning to adapt the communication networks themselves?**
  - RA1.C3 Adaptation and Learning in Wireless Autonomous System *(Romberg, Ribeiro)*
  - **Year 1 Goal:** Understand the effect of imperfect communication on the convergence properties of distributed optimization algorithms

Expected Outcomes & Impact

- ***New theories of representation***
  - How to decompose complex problems for efficient planning, inference and learning, to preserve guarantees of completeness, correctness and (bounded) optimality
  - How to represent complex problems in a human-interpretable manner
- ***New theories of inference, perception and learning***
  - How to learn with new (actually useful) bounds on data efficiency
  - How to act efficiently to acquire new concepts
- ***New theories of communication***
  - How to incorporate communications (and especially limitations) as part of the action loop
  - How to learn new models of communication layers

Deep Learning *cut loose from* ***Big Data***

Distributed World Modeling

Learned Communication Coding

Communication-Aware Planning