Towards Software Apprentices that Learn in Dynamic Domains

(FA9550-16-1-0138)

PI: Kenneth D. Forbus (Northwestern University)

Co-PI: Thomas Hinrichs (Northwestern University)

AFOSR Program Review:
Computational Cognition and Machine Intelligence Program
(10/7/20, Arlington, VA)

Towards Software Apprentices

Forbus & Hinrichs, Northwestern University

Objective:

- Discover how to build software apprentices
  - Theory of representations to support learning and reasoning about dynamic environments
  - Theory of language and sketch understanding to learn from demonstrations, advice, and stories

Approach:

- Build on Companion cognitive architecture
  - Human-like analogical reasoning and learning
  - Qualitative representations and reasoning
- Use open-source strategy game as dynamic environment simulation

DoD Benefits:

- Software apprentices would be disruptive breakthrough
  - Enough like us to communicate and be trustable
  - Provide complementary strengths

Progress:

- Analogical reference resolution and unified QP frames to improve communication abilities
- New representations for tactics and spatiotemporal histories for better decision-making

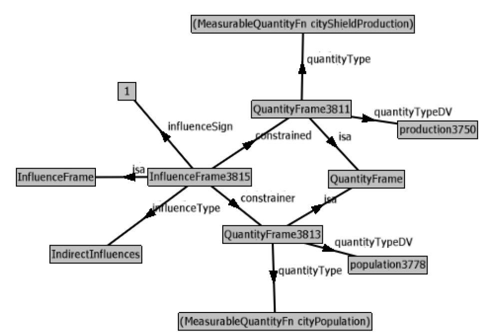

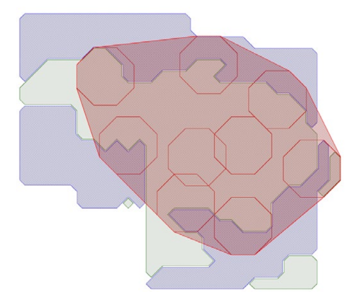

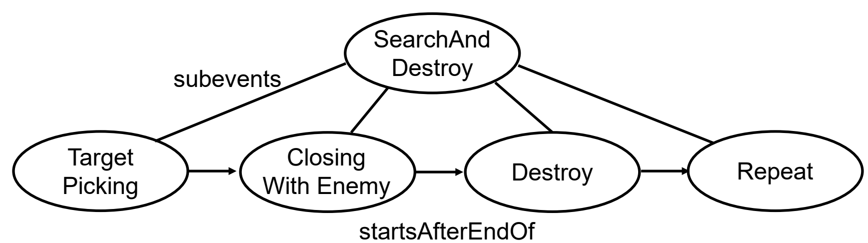

Project Goals

- Develop theories and representations to support learning, reasoning, and communicating about strategies, tactics, and decision-making in dynamic domains
  - Qualitative Representations, Analogical Learning
- Extend Companions to operate as apprentices
  - Learn by self-directed experimentation
  - Learn from human-provided lessons and advice, using natural modalities (language and sketching)

Progress Towards Goals

Better communication:

- Analogical Reference Resolution handles near-miss references and correcting for common ground
- Unified QP frames to support assembling instance-level or type-level qualitative models from language based on context

Better decision-making:

- Developing domain-independent representations of tactics for flexible decision-making and learning
- Developed qualitative spatiotemporal histories for strategic reasoning

Reference Resolution

- To interact with humans via language, agents must be able to interpret *referring expressions*
  - “Pick up *the yellow ball*.”
- *Reference resolution* is the task of matching a referring expression to its intended referent
- People are good at understanding *near misses*
  - A: “Did you hear about *the man who jumped off a bridge*?”
  - B: “He didn’t jump. He was pushed.”
- Exact matching is too brittle
  - Need a notion of semantic similarity
  - Solution: Use analogy!

Analogical Reference Resolution

- Use analogical retrieval to identify the entity that's most similar to a given description
  - Probe case contains semantics of the description
  - Case library contains known information about objects in scene
  - Retrieved case(s) correspond to the most likely referent
- Nakos et al. ACS 2019: Evaluated on TUNA corpus (Gatt et al., 2007)
  - Outperformed baseline due to near misses in dataset
  - Robust to artificial noise (insertions, deletions, substitutions)

the dark-haired man with glasses

the green chair facing left

Modeling Human Two-Stage Strategy

- 1st stage like before
- 2nd stage corrects for what other person can perceive
  - Suppresses objects not in common ground
  - Re-represents as needed (e.g. recomputes adjectives using CogSketch)
- Nakos et al. CogSci 2020: Model matches human behavior in all 16 trials

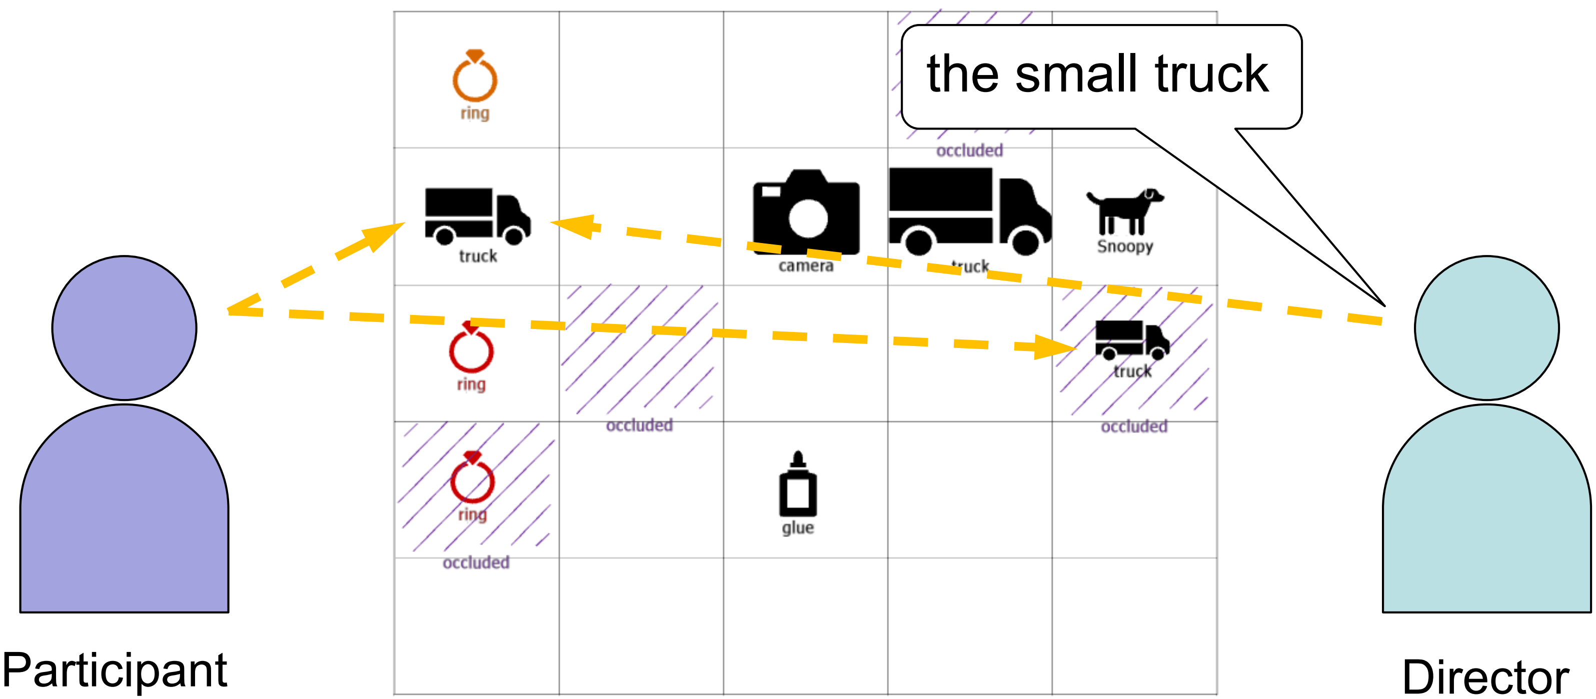

(Epley et al., 2004)

Unified QP Frame System

- Hypothesis: Qualitative representations form an important component of natural language semantics
- Frame systems support the assembly of models from incremental information

Instance-level Frames (Kuehne, 2004)

"Heat flows from the brick to the ground, because the brick is hotter than the ground."

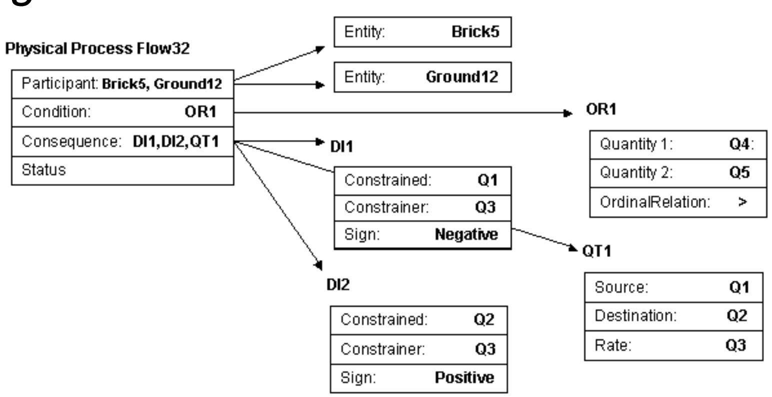

Type-level frames (McFate et al. 2014)

"As the population of the city increases, the food production of the city increases."

```
Qprop9978
  isa: TypeLevelQpropFrame
  sign:+
  constrained: Quantity9706
    isa: TypeLevelQuantityFrame
    entityType: Production-Generic
    quantityType: (RateFn
                    Production-Generic)
  constrainer: Quantity9974
   isa: TypeLevelQuantityFrame
   entityType: Freeciv-City
   quantityType: cityPopulation
```

Unified QP Frames are Level-Neutral

- Language provides information incrementally
  - Example: “Production depends on population.”
  - Sentence provides part of the meaning
  - Context provides the rest
    - “What does the production in Chicago depend on?” “
    - “What does production in a city depend on?”
- Frame system needs to defer level decisions
- Model Assembler converts frames to traditional QP models at either level
  - Enables context to be exploited

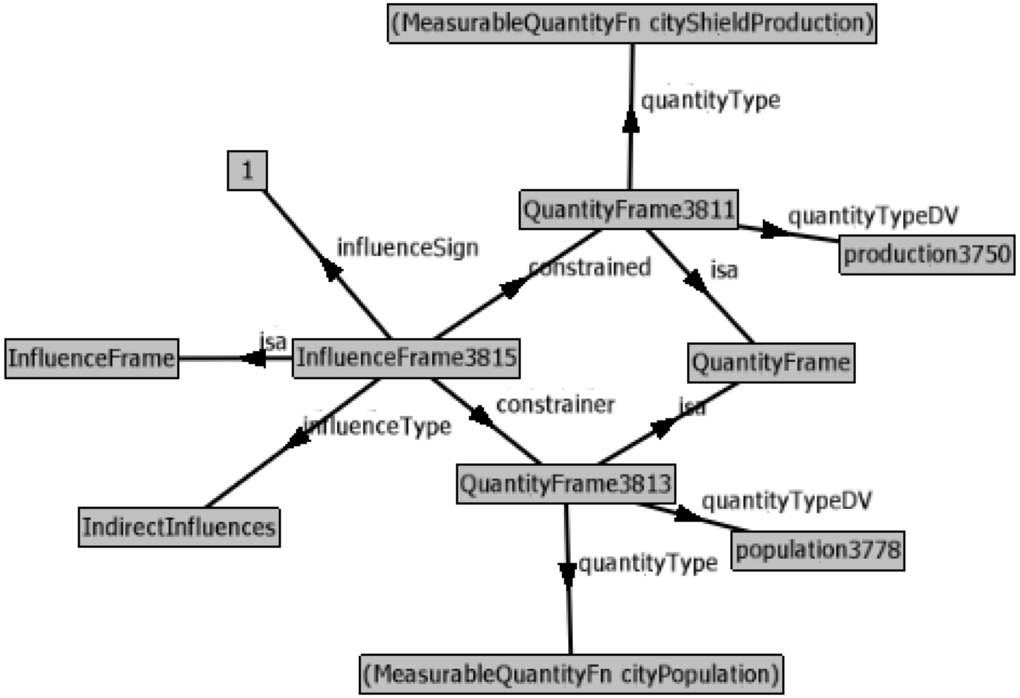

Unified Frame representation for "Production depends on population."

Instance-level model

```
(qprop+ ((MeasurableQuantityFn cityShieldProduction) Chicago)
        ((MeasurableQuantityFn cityPopulation) Chicago))
```

Type-level model

```
(qprop+TypeType (MeasurableQuantityFn cityShieldProduction)
   (MeasurableQuantityFn cityPopulation)
   FreeCiv-City FreeCiv-City same)
```

Forbus & Hinrichs, Proceedings of QR 2020

Goal: Broad Conceptual Coverage

- 89 comparatives added
  - e.g. “The elephant is heavier than the fly” introduces an ordinal quantity frame involving mass
- Lexicalizing quantity types as needed
  - NextKB has > 300 Quantity types
  - e.g. “rough carpet” **`quantityType:SurfaceSmoothness`** **`quantityValue: Rough`** *;;element from qualitative value set linked by ordinals*
- Support domain-independent dialogues about strategies, tactics, and qualitative models

Flexible Representations for Tactical Decision Making

Learning to apply complex tactics one decision at a time

- Tactics represent larger chunks of behavior than primitive actions.
  - Incrementally instantiated in a Course Of Action intent representation
  - Well-suited for learning through experimentation
- Domain-independent representation
  - To support communication and transfer
  - Operationalize for domain in terms of its actions and predicates
- Davidsonian (frame-like) tactic representation
  - Composed of multiple distinct decisions
  - Each decision is an independently learnable action

What is a tactic?

Tactics are compositions of multiple actions and goals that:

- are typically more general than simple macrops or action sequences
- are different from HTNs because they have explicit, reified participants and structure
- achieve strategies (which in turn resolve goal tradeoffs)

Examples:
Business tactics: OrganizationMerger, Divestment, LeverageBuyouts…
Legal tactics: Lawsuit, Incorporating, PlausibleDeniabililty, 5thAmendment
Social tactics: Flattery, Blackmail, Revenge, Anonymity, Competition…
Military tactics: SearchAndDestroy, Seige, PincerAttack, IsolatingSupplyLines…

These are well-known, real-world tactics, not game-specific cliches.

Representation

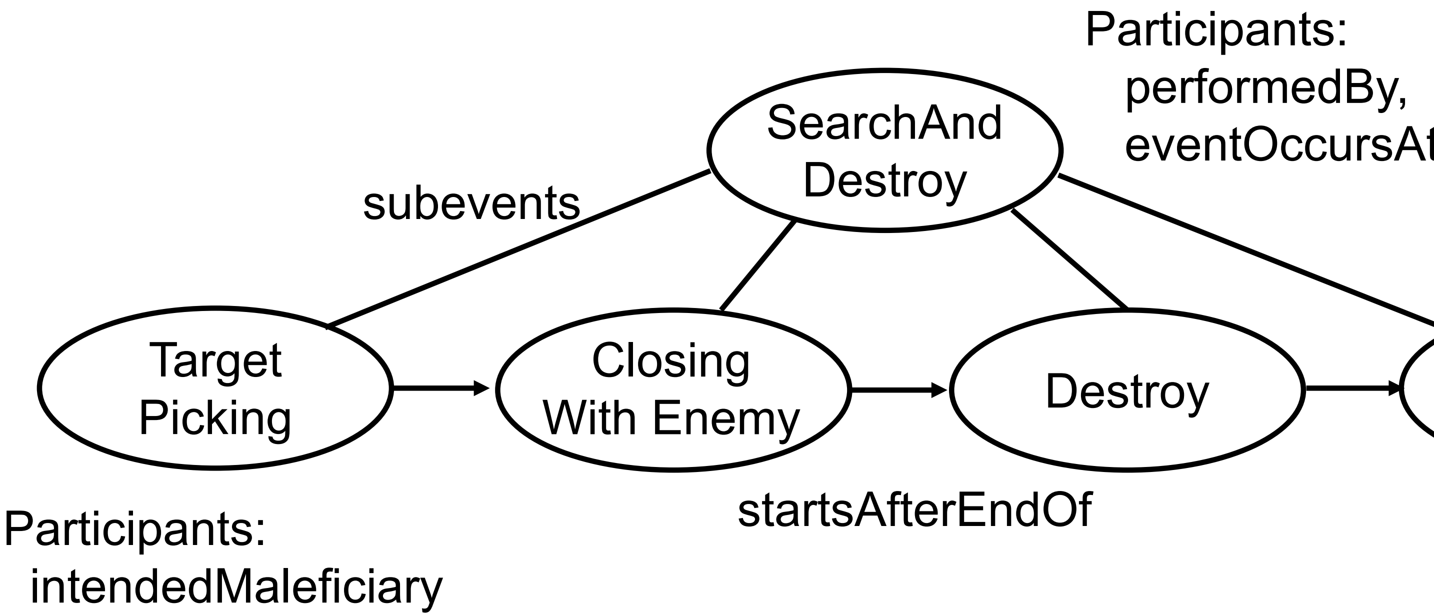

Participants are inherited by sub-events and post-events.
Participants can be selected incrementally, in any order.
Selection of participants is a learned decision.
Credit assignment maps back from game events to goals, and from goals to tactical decisions in order to update selection policies.

Example: Search and Destroy

- Activation of offensive goal posts SearchAndDestroy tactic on COA (1)
  - Decision to assign Catapult to S&D tactic populates a decision case
  - TargetPicking sub-event selects Chariot and also populates a decision case
- ClosingWithEnemyForce continually re-plans to track intended target (2)
- DestroyingAnEnemyForce attacks when target is in range (3)
- Outcome triggers credit assignment to the reified decisions (4)
  - Decision cases added to success or failure libraries
  - Generalized cases become policies for constraining the assignment and target-picking decisions

Histories for Strategic Reasoning

- History = spatiotemporal representation of change
- Use CogSketch for grounding qualitative spatial representations in maps
- Example: In Freeciv, a civilization's footprint consists of its cities and the space between them

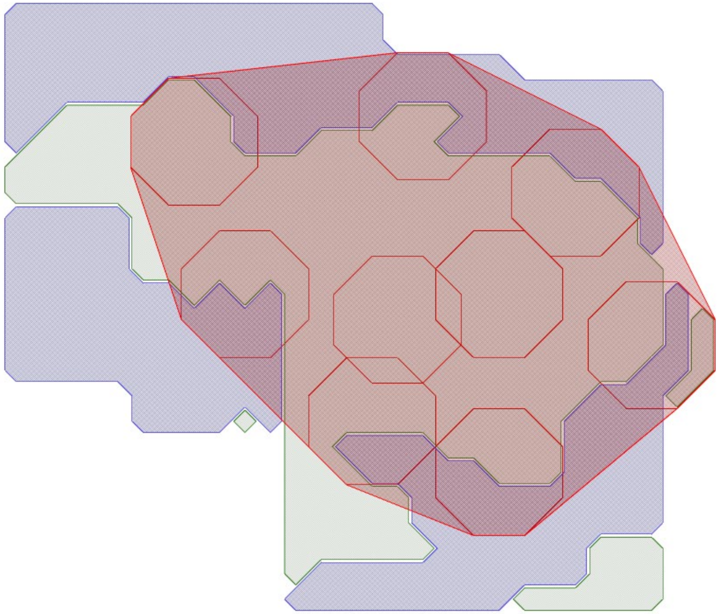

Example of a footprint

Threat Detection

Civilization footprint

Threats detected when entering footprint

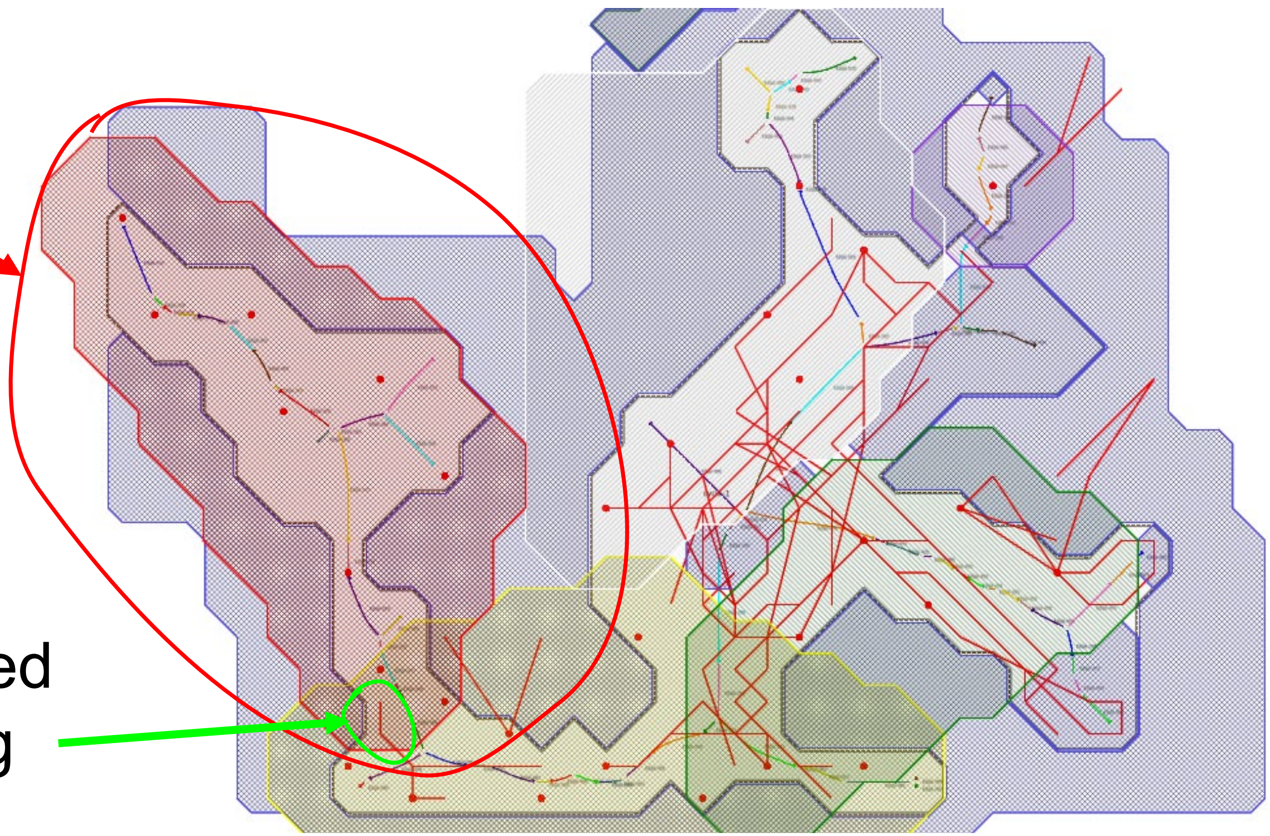

Identify Civilization Expansion Opportunities

Safe expansion area

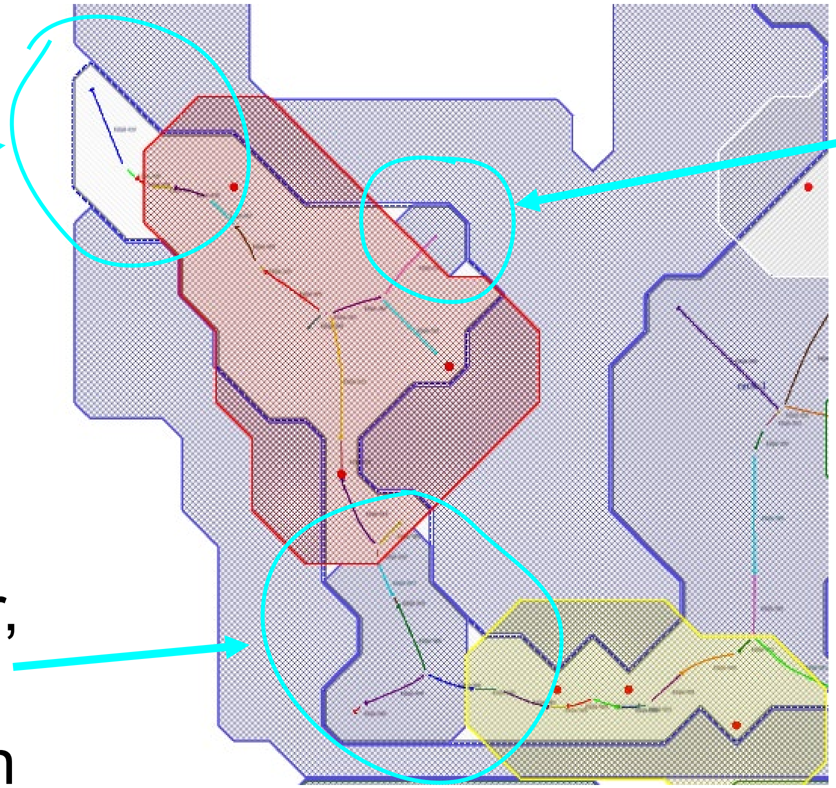

Safe expansion area

Potential danger, build on strong defensive terrain

Suggest City Specialization

Border cities near hostile neighbors need stronger defenses

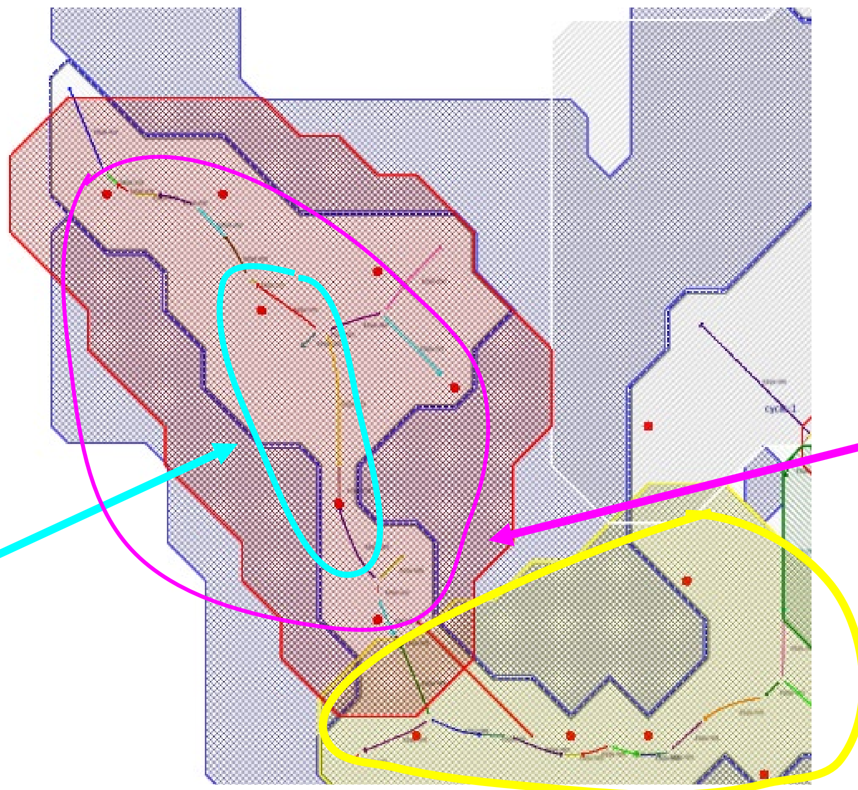

Interior cities can focus on economics, science

Enemy footprint

Experimental Results

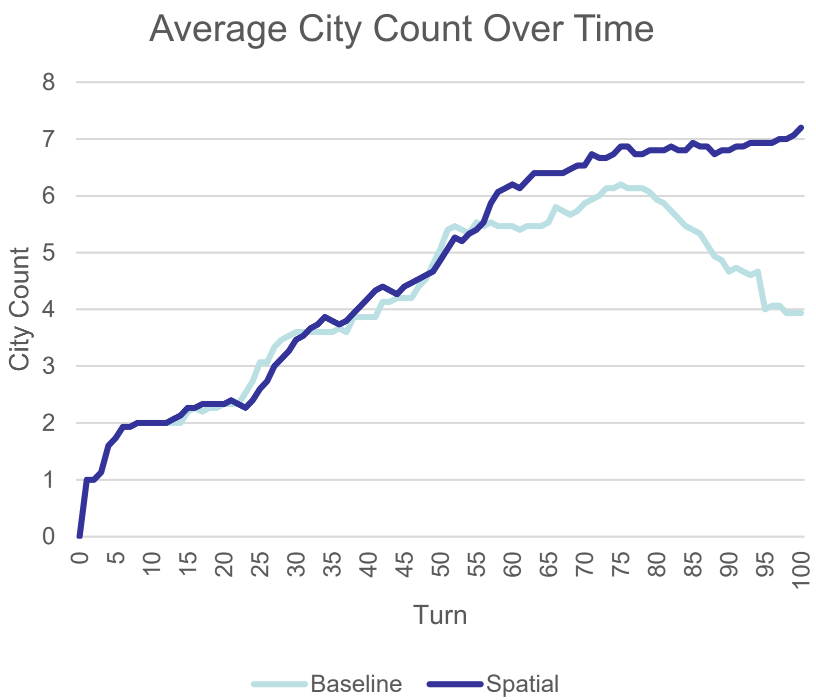

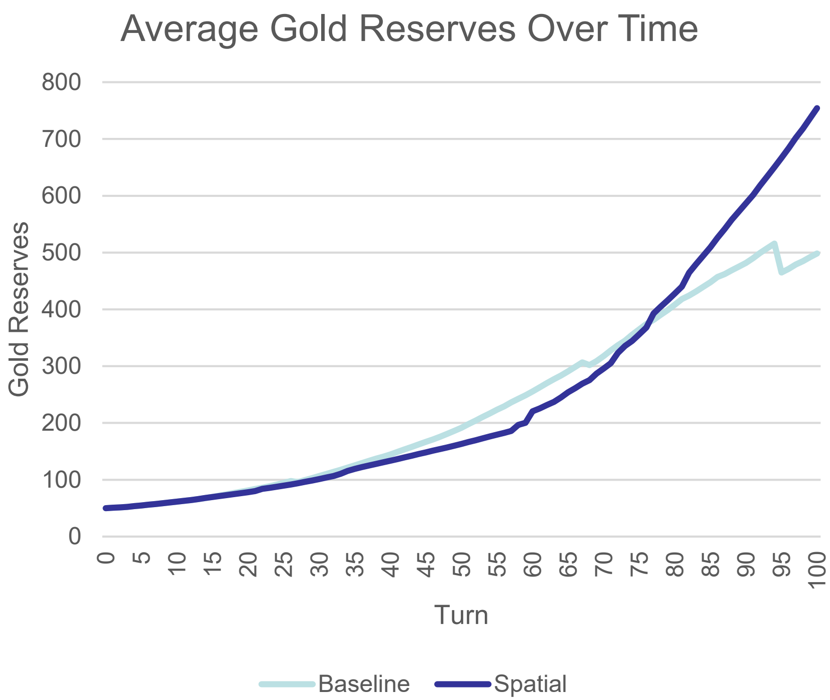

P =.00844 using paired t-test

P=.01128 using paired t-test

Hancock et al. Proceedings of QR 2020

Next Steps

- QP Frames
  - Finish implementation
  - Use in next round of learning from dialogue
- Tactical decision-making
  - Implement more tactics
  - Experiments in learning to operationalize tactics
- Histories
  - Use for episodic memories, for concise representations of temporal behavior
  - Analogical learning from histories

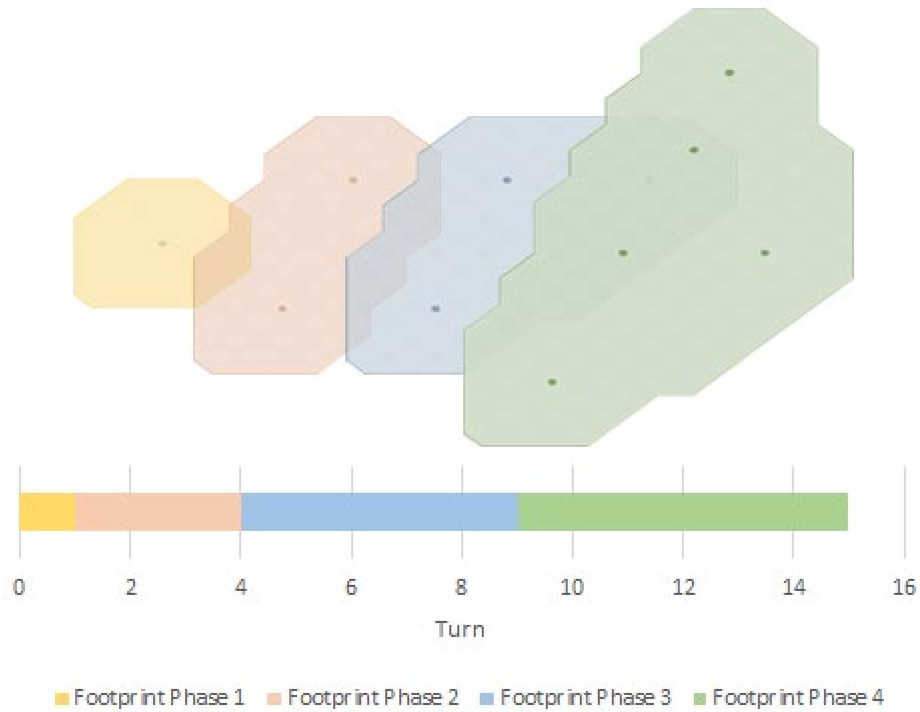

Reasoning for Social Autonomous Agents

Kenneth D. Forbus, Thomas Hinrichs, Northwestern University

<u>Objective</u>: Understand the reasoning capabilities needed to create autonomous software social agents

<u>Approach</u>:

- **Explore *social reasoning*, (e.g. norms, legal concepts)**
- **Explore *cognitive control,* the signals and methods used to guide an agent's reasoning and learning**
- **Explore *broad reasoning*, to handle open-ended questions, including the computational basis for human cognitive illusions**
- **Collaborate with AFRL-ACT3, to transfer Northwestern technologies, and use ACT3 problems and data to help guide the research**

<u>DoD Benefits</u>: Intelligent systems that can predict human and organization reactions to situations and follow social norms would be a disruptive breakthrough

- **Like us enough to communicate, be trustable**
- **Different enough to provide complementary strengths**

<u>Progress</u>:

- **New start**
- **Delivered Companion cognitive architecture executable to ACT-3 for experimentation**
- **Discussions underway to deepen the collaboration**

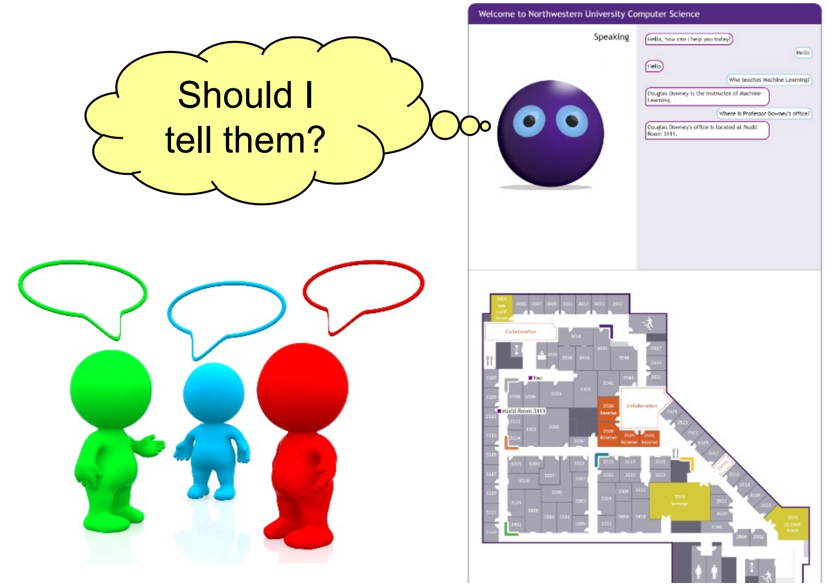

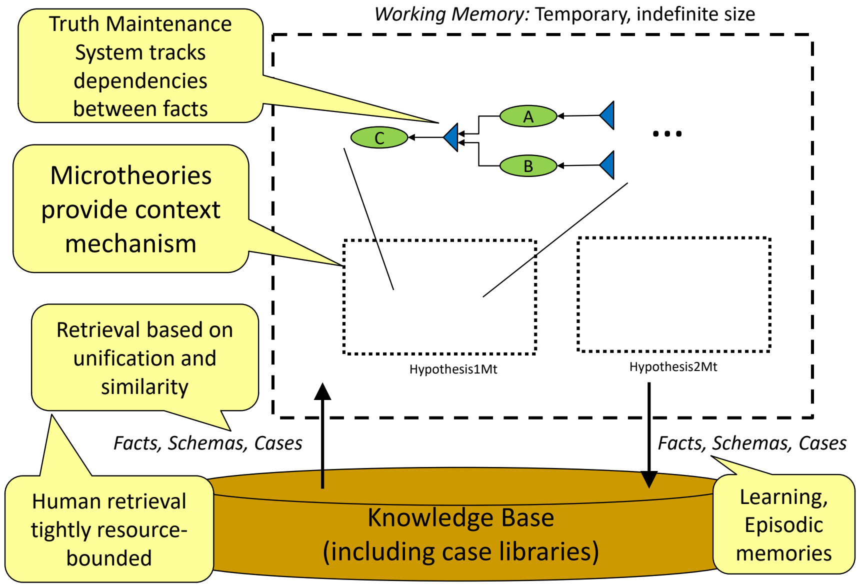

Other Relevant Activities

- Forbus is serving as a technical expert for a UN OECD/National Academies study on AI and the Future of Work
  - Follow-on from initial study in 2017
  - Chapter in progress: “Evaluating Revolutions in AI from a Human Perspective”

Publications

- Forbus, K., & Hinrichs, T. (2019). Qualitative Reasoning about Investment Decisions. *Proceedings of the 32nd International Workshop on Qualitative Reasoning (QR 2019).* Macao, China.
- Forbus, K. & Hinrichs, T. (2020) Unifying Instance-Level and Type-Level QP Frames for Natural Language Understanding. *Proceedings of QR2020.*
- Forbus, K., Hinrichs, T., Crouse, M., & Blass, J. (2020). Analogy versus Rules in Cognitive Architecture. *Proceedings of ACS 2020.*
- Hinrichs, T. and Forbus, K. (2019). Experimentation in a Model-Based Game. In *Proceedings of the Seventh Annual Conference on Advances in Cognitive Systems*. Cambridge, MA.
- Hinrichs, T. and Forbus, K. (2019). How Qualitative Models can Improve Learning by Experimentation. *Proceedings of the 32nd International Workshop on Qualitative Reasoning (QR 2019).* Macao, China.
- Hancock, W., Forbus, K., & Hinrichs, T. (2020). Towards Qualitative Spatiotemporal Representations for Episodic Memory. *Proceedings of QR2020.*
- Nakos, C., Rabkina, I. & Forbus, K. (2019) An Analogical Account of Reference Resolution. *Proceedings of ACS 2019.*
- Nakos, C., Rabkina, I., Hill, S., & Forbus, K. (2020) Corrective Processes in Modeling Reference Resolution. In *Proceedings of CogSci 2020*.