Communication and Competitions

Matthew Stone

Department of Computer Science
Center for Cognitive Science
Rutgers University

Communication

Actions that potentially allow agents to give evidence about their preferences, knowledge or intentions.

Outline

Research Questions
Frameworks for Communication
Two Reasoning Paradigms
Looking Ahead

Q1: Social Reasoning

How can agents recognize others' strategies and coordinate their activities

Arises in simple situations (drive on the right, pick a beach stand)

Depends on models of strategy, integrating idealizations, practical logic, empirical facts

Q2: Ontology of meaning

How do real-world actions acquire interpretation as signals?

Conceptual, scientific, engineering problems
- foundational analysis (Lewis on convention)
- explanations of particular properties of human meaning (Kripke on names)
- design of useful robot languages (Steels)

Q3: Problem solving

How can specialized agents combine forces to do things they couldn't do on their own?

Focuses on rich environments, practical tasks
- multi-agent planning
- information gathering and exchange
- realizing teamwork skills

Q4: Interacting with people

How can NL technology, answers to Q1-Q3, lead to better interfaces or interactions?

Focus on design

- leading people to effective and simple interactions
- working around limitations in technology (and people)

Outline

Research Questions
Frameworks for Communication
Two Reasoning Paradigms
Looking Ahead

Costly signals

Real-world actions with natural effects

- give information about agents' preferences to observers

Present (implicitly) in lots of interactions

- Think of betting games: ante, bid, etc.

Costly signals

Advantages:

- Continuous with planning, simple social competence (Bayesian receiver infers sender's type from natural meaning of action)
- Generally trustworthy

Disadvantages:

- No problem-solving, negotiation

Cheap Talk

Signals with no costs or effects

- most often: assigned meanings through agents' strategies in using them

Representative case (Steels)

- referential communication task where agents share payoffs for coordinating on objects
- learn code mapping signals to properties

Cheap Talk

Advantages:

- Applicable in simple settings
- Amenable to formal, algorithmic techniques

Cheap Talk

Downsides:

- Not human-like meanings (problems with reference, compositionality, speech acts)
- Somewhat unconstrained (slow to converge)
- Depends on aligned interests ("pooling equilibria" where communication is not in some speakers' interests)

Negotiation languages

Agents exchange messages in a formal language with a specified semantics

General AI approaches

- Cohen & Levesque (KQML), Sidner

Specific experiments

- Color trails (Gal, Grosz, Pfeffer, Shieber)

Negotiation languages

Advantages:

- Messages can be binding (focus on problem solving, preferences; avoid problems of trust)
- Simplest way to handle expressive content

Drawbacks:

- Design challenges for good mechanisms

Constrained natural language

Talk that's naturally limited to specific problem-solving domain

Examples:

- Coconut (Moore, Thomason, Di Eugenio) Problem solving task (décor arrangement) with reference and negotiation
- GIVE (Koller) – generating instructions in virutal environments

Constrained natural language

Advantages:

- Easiest for real human players
- Closest to applications

Drawbacks:

- Methodology involves harsh tradeoff of coverage/performance against task complexity

Outline

Research Questions

Frameworks for Communication

Two Reasoning Paradigms

Looking Ahead

P1: Collaborative Agency

Communication is a kind of teamwork

- in the tradition of Allen, Cohen, Grosz, etc.

Understanding real-world action

attributing mental state,
intention or ***commitment,***
linking action to context and goals.
[e.g., Pollack 90]

Same for language use

Real-world teamwork

Hey, I should help that guy open the door.

engages with agents' inferred commitments.
[e.g., Cohen & Levesque 90]

Same for language use

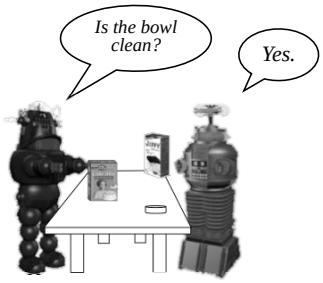

Same for language use

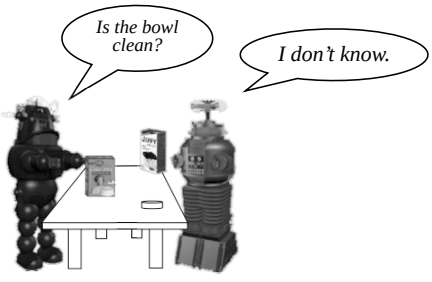

Real-world responses

must be ***presented*** as public contributions.
(don't just sneak and open the door unnoticed)
[e.g., Sengers 99]

Same for language use

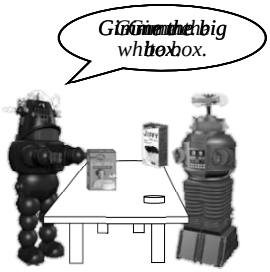

P2: Behavioral Economics

Understanding and producing utterances reflects people's heuristic and biased decision making

- recent work by Pinker and colleagues

Strategic inference

A new twist:

- Observe utterance
- Infer speaker's information
- Infer speaker's values
- From what speaker said
- *And your empirical knowledge about how the speaker would decide what to say*

Three case studies

Plausible deniability

- Human decision makers can be naive

Calculated ambiguity

- We pursue paradoxical social purposes

Obvious indirection

- We are irrationally sensitive to framing

Gee officer, is there any way I could take care of this right now?

Explanation

Sophisticated speaker implicates offer of bribe

- But naïve speaker might use utterance (irrationally) for its literal meaning

Ambiguity of interpretation gives sophisticated speaker reason to use the utterance

Why indirection?

Asymmetries of strategy

- Sophisticated speaker
  second-guessing honest officer,
  second-guessing corrupt officer,
  second-guessing honest but naïve driver.

Like "beauty contest" results in economics

Beauty contest

Keynes's metaphor of market speculation

- speculators guess how assets will look to naïve buyers

Example from behavioral economics

- guess a number 0-100
- person who guesses 2/3 average wins

Beauty contest

Keynes's metaphor of market speculation

- speculators guess how assets will look to naïve buyers

Example from behavioral economics

- guess a number 0-100
- person who guesses 2/3 average wins
- human winners in range 15-25

We really do appreciate the good service we get here.

Why indirection?

Social preferences

- Speaker doesn't really know what he means
- "Almost" implicates request
- Figures maitre d will catch on if he's on the same wavelength
- Prefers misunderstanding to conflict if maitre d is not on the same wavelength
- Feelings not calculations drive social reasoning

Like "ultimatum game" results in economics

Ultimatum game

Experimenter gives player 1 $20.
Player 1 gives $X to player 2, keeping rest.
Player 2 has 2 choices:

- accept, and everyone keeps their money
- reject, and nobody gets anything

Ultimatum game

Experimenter gives player 1 $20.
Player 1 gives $X to player 2, keeping rest.
Player 2 has 2 choices:

- accept, and everyone keeps their money
- reject, and nobody gets anything

Player 2 consistently rejects low offers

- (irrational) feeling of unfairness

Do you want to come up and see my etchings?

Why indirection?

Framing

- Indirection is always off the record (however obvious it is)
- This way of presenting information prompts different feelings and decisions
- These (irrational) feelings and decisions are part of the point or meaning of the utterance

Like “reference point” results in economics

Reference point results

Two coffee shops, A & B

At A:

- Coffee costs $2.90, but cup costs $.10

At B:

- Coffee costs $3.00, but $.10 discount if you bring your mug

Outline

Research Questions
Frameworks for Communication
Two Reasoning Paradigms
Looking Ahead

Understanding people

Goals for competitions

- Avoid solipsism
  competition agents that model each other (and nothing else)
- Don't require expertise
  novices see task failure or unfair outcomes (particularly: naïve human players)
- Make playing fun and rewarding
- Have hooks to the real world