Summary

AI Roadmap Workshop #1
Integrated Intelligence
Marie desJardins & Ken Forbus

Vignettes: Improve Health Care

- Health insurance advisor/assistant
- Suicide and mental health diagnosis, monitoring, and intervention
- Disease prevention, diagnosis, monitoring, e.g. individualized/customized cancer treatment at scale
- Drug optimizer / esp. in context of opioid crisis, overprescription, addiction
- Holistic medical assistant, social support scenario (medical interview; early diagnosis/triage; convey "bullet" to doctor)
- Eldercare, aging in place with dignity, safety net, occupational therapy & engagement
- Smart support for medical data privacy while leveraging medical data for data mining (voluntary "data donor" program)
- Audience-aware summaries of electronic medical records (for patients with varying backgrounds; for different health care professionals; for use in data mining)

Vignettes: Accelerate Science

- AI-based lab assistant
- Extract knowledge from remote-sensing data, world-scale models
- Climate change
- Context-sensitive personalized education
- Autonomous scientific discovery (in specific domain? Drug discovery?)
- Building rich causal models based on time-series data
- Experiment design and "argumentation testing" (does experiment X actually test hypothesis Y?)
- Fires, flooding, and drought models - prediction & prevention
- Globally available clean water supplies

Vignettes: Education and Training

- Textbook elimination: real-time personalized tailoring of relevant information
- AI simulated patients, mentors students, peers, career counselors
- Depersonalized avatars to reduce bias in interviews & other evaluative contexts
- Support informed citizenry by helping people take filter bubbles, fake news into account

Vignettes: Improve Business Innovation

- Supply chain disruption and dynamics
- Cybersecurity education & advising

Three Big Roadmap Avenues

- Lifelong Personal AIs
  - Social Cognition
  - Creativity and Design
- Science of Integrated AI
  - Modeling Human Intelligence
  - Memory
- Open Knowledge Repository

Lifelong Personal AIs

- Live on the user's machine, owned by them, operating in their interests
- Aligned with human values and goals
- See other slides for details

LPA Area: Social Cognition

- Current state of the art; technical gaps
- Affective computing
- Human-agent teaming
- AI alignment with human values, goals, desires, and norms
- Explainable AI
- Discourse/dialogue/argumentation
- Personalization and modeling context

LPA Area: Creativity and Design

- Current state of the art; technical gaps
- Everyday creativity -- MacGyver
- Design of complex systems
- Novelty production / handling novel situations
- Creativity in the service of goals (to express, understand, explain, solve)

Science of Integrated AI

- Current state of the art; technical gaps
- Deliberative, Guided Autonomy
- Goal Reasoning
- What are the components of intelligence?
  - Current subdisciplines appropriate way to carve phenomena?
  - Explore finer-grained capabilities/alternate formulations
- What are the interfaces and process models between/among components?
- Perceptually grounded semantics and task models
- Metacognition
  - reasoning about their own cognition; "knowing what you don't know"

SIA Area: Memory

- Current state of the art; technical gaps
- How to organize for life-long learning and behavior - learning transfer; small-data/zero-shot learning
- Episodic memory
- Sharing episodic memory across agents to bootstrap/improve-- privacy concern

SIA Area: Modeling Human Intelligence

- Provides clues for how biological systems integrate their capabilities
- Cognitive Science combines AI, Psychology, Linguistics, Anthropology, Neuroscience
- Emotions
- Motivation
- Consciousness
- Evolutionary origins of mind; shared evolutionary structure

Open Knowledge Repository

- Current state of the art; technical gaps, why will this time be different?
- Shared, open, accessible, interoperable repositories of common-sense, domain-independent (or domain-spanning) knowledge
- Rich background knowledge to support transfer and zero-shot learning
- Bridging symbolic representations and statistical data
- Knowledge engineering / elicitation / learning / maintaining

OKN Area: Expressive Knowledge

- Go beyond collections of entities and bare ontologies to support more expressive knowledge
- Representing scientific questions, including open questions
- Causality and causal models from multiple domains (e.g. biology, mechanics, social cognition)
- Metaknowledge and self-knowledge
- Cultural knowledge

OKN Area: Reasoning

- Integrating multiple forms of reasoning to handle more general queries
- Reasoning at scale
- Non-deductive forms of reasoning
  - Analogy
  - Abduction
  - Statistical reasoning
- Automated knowledge refinement processes, to improve repository contents

Big Scientific Questions

- What are minds?
  - What processes constitute them?
  - What knowledge do humans have?
  - What makes humans special?
- How do we represent conceptual structure?
  - How do we integrate information from our senses into coherent models of the world?
- Cognitive Science: Accounts that combine AI, Psychology, and Neuroscience