

Story-Enabled Intelligence

Communication, Narration, Explanation

Organizers: Leilani H. Gilpin (lgilpin@mit.edu), Dylan Holmes (dxh@mit.edu),
and Jamie C. Macbeth (jamie.macbeth@gmail.com)

Systems that can describe their own behavior exhibit intelligence of a higher order. Storytelling-like capabilities empower systems to explain their decisions, describe their activities, align their present situation against precedent, consider hypothetical alternatives, diagnose their mistakes, learn from stories, and generalize their experiences. A number of research efforts have independently explored aspects of such machine-generated descriptions, but interaction between specialized subareas of AI remains sparse in the literature. These specialized areas include:

- **Explainable AI (XAI)**, focusing on augmenting machine learning frameworks with compositional, humanly-interpretable representations.
- **Case-based reasoning**, focusing on deploying previous stories to understand new phenomena.
- **Hypothetical reasoning**, focusing on understanding which is augmented by speculation about variations in past, present, and future circumstances.
- **Narratology**, focusing on modeling the human aspects of producing and interpreting narratives and explanations.
- **Cognitive architectures**, focusing on the design of large-scale symbolic systems.
- **Natural language generation and understanding**, focusing on modeling human facility with language, including description, conversation, persuasion, explanation, and teaching.
- **Planning**, focusing on generating problem-solving action sequences, often involving skillful incorporation of contingency, constraints, other agents, and incomplete information.

In pursuit of a unified approach, four common challenges emerge: (1) **Architecture**—How can we augment existing opaque systems so as to add a layer of explainability? Alternatively, how can we design systems so as to incorporate explainability from the bottom up? (2) **Representation**—How can we design systems out of composable, explainable parts? Which representations capture application-specific information and expose constraint? (3) **Procedure**—How do we effectively identify and deploy sources of constraint? How do we develop anytime explanations, and how can these explanations integrate new information? What capabilities, such as matching against precedent or reasoning about hypotheticals, are enabled by explainable parts? (4) **Cognition**—What kinds of explanations are most effective for human users? How do humans produce and consume stories? How do we tailor explanations for different audiences and purposes?

This symposium will introduce a shared forum for researchers to discuss and debate intersectional issues involved in developing intelligent systems that use and manipulate story-like representations. To encourage diverse approaches, we invite contributions not only from artificial intelligence but also

directly relevant work from nearby fields such as computer science, narratology, cognitive science, anthropology, and linguistics. With this symposium, we will promote discussion and exchange of ideas among the following topics, all of which bear directly on the general problem of machine-generated narratives:

- Computational models of human storytelling and understanding
- Systems that explain when they fail
- Large-scale architectures built from communicating parts
- Problem-solving programs that tell their own story
- Systems that combine information from many sources
- Tools that summarize and interpret
- Tools that support planning and exploration of alternative scenarios
- Self-explaining software engineering tools
- Systems that provide compositional explanations of what they're doing.

Building from our experiences with workshops in related areas, we offer a multitude of ways to engage. Given that this unified approach to story-enabled intelligence is relatively new, we furthermore aim to encourage promising preliminary work and empower early-stage researchers by showcasing student work, preliminary papers, and open discussion. In detail, we anticipate including:

1. Both shorter talks (20-30 minutes) for papers accepted for presentation, as well as long-form invited talks (1 hour).
2. Brief intermittent poster sessions, scheduled throughout, showcasing more preliminary work.
3. Break-out sessions (1-2 hrs) between talk sessions. These moderated panel discussions elaborate on ideas from the talk sessions, heavily focusing on enabling audience interaction.

Finally, we invite contributions in a variety of forms: paper submissions (long and short), panel discussants, and poster abstracts.

Organizing Committee

Leilani H. Gilpin

77 Massachusetts Ave

Bldg 32-G506

Cambridge, MA 02139

lgilpin@mit.edu

Dylan Holmes
77 Massachusetts Ave
Bldg 32-G258
Cambridge, MA 02139
dxh@mit.edu

Jamie C. Macbeth
Smith College Department of Computer Science
Northampton, MA
jamie.macbeth@gmail.com

About the organizers. Professor Jamie Macbeth brings expertise in the fields of natural language understanding and story processing, as well as established collaborators across multiple institutions. Leilani Gilpin and Dylan Holmes are senior PhD students operating as chiefs of staff for their respective research groups at MIT. Their duties include one-on-one mentoring of all students in lab and managing the groups' meetings and publication timetables. Additionally, Leilani has organized and run MIT EECS PhD visit days for over 150 admitted students, and Dylan coordinates all aspects of the yearly AI course involving 300+ students and 12 TAs. All three organizers are involved in a number of related communities, including Advances in Cognitive Systems, AAI, CHI, and IJCAI.

Potential participants

The following people have all expressed keen interest:

- Andrew Gordon (USC)
- Mark Finlayson (FIU)
- Doug Riecken (AFOSR)
- Dagmar Gromann (Technical University Dresden)
- Patrick Winston (MIT)
- Gerald Jay Sussman (MIT)
- Henry Lieberman (MIT)
- Margaret Minsky (MIT Media Lab)
- Jason Wilson (Northwestern)
- Irina Rabkina (Northwestern)
- Tristan Thrush (MIT)
- Gaurang Gawai (Palo Alto Research Center)

Kyle Dent (Palo Alto Research Center)

David Aha (NRL)

Mark Roberts (NRL)

We have also recently reached out to the following people in our network, but have not yet heard back:

Rafael Perez y Perez (UAM, Unidad Cuajimalpa)

Pat Langley (ISLE)

Robert Berwick (MIT)

Michael Witbrock (IBM)

David Ferrucci (Elemental Cognition)

Larry Birnbaum (Northwestern and Narrative Science)

Kris Hammond (Northwestern and Narrative Science)

Danny Hillis

Ben Shapiro (University of Colorado - Boulder)

Doug Lenat (Cycorp)

Carl Hewett

Catherine Havasi (Luminoso)

Tom Knight (Ginkgo Bioworks and MIT CSAIL)