Lecture Notes in Computer Science 13000

More information about this subseries at http://www.springer.com/series/7409

William F. Lawless · James Llinas ·
Donald A. Sofge · Ranjeev Mittu (Eds.)

Engineering Artificially Intelligent Systems

A Systems Engineering Approach to Realizing Synergistic Capabilities

Springer

Editors
William F. Lawless
Paine College
Augusta, GA, USA

James Llinas
University at Buffalo, Director Emeritus,
Center for Multi-source Information Fusion
Buffalo, NY, USA

Donald A. Sofge
U.S. Naval Research Laboratory
Washington, DC, USA

Ranjeev Mittu
U.S. Naval Research Laboratory
Washington, DC, USA

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-89384-2 ISBN 978-3-030-89385-9 (eBook)
https://doi.org/10.1007/978-3-030-89385-9

LNCS Sublibrary: SL3 – Information Systems and Applications, incl. Internet/Web, and HCI

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

Jay Forrester, creator of system dynamics, remarked we "live in a complex world of nested feedback loops," involving cascading interdependencies across these loops that vary in complexity, space, and time. Many, if not most, of the current AI/machine learning algorithms, data and information fusion processes, and related methods are attempting, in software, to estimate situations in this complex world. Hence, these algorithms and processes must gracefully and efficiently adapt to the technical challenges (e.g., data quality) induced by nested feedback loops and interdependencies.

To realize effective and efficient designs of these computational systems, a systems engineering perspective may provide a framework for identifying the interrelationships and patterns of change between components, rather than static snapshots, as systems engineering methods are about seeing "wholes." The increasing prevalence of intelligent systems within society will reveal yet more interdependencies, in that AI-enabled intelligent systems have both upstream processes impacting them and downstream processes they affect. The benefits of studying cascading interdependencies through a systems engineering perspective are essential to understanding their behavior and for the adoption of complex system-of-systems in society. Knowledge about the world is needed for multiple applications.

Models of complex situations (e.g., patterns of life) are typically attacked with reductionism. However, in the absence of modeling dynamics and evolving interrelationships, these strategies fall short. If causal models of feedback loops and interrelations can be developed, then system modeling, and the ability to estimate system dynamics, will be as correct as possible. However, real-world situations of Forrester's "nested feedback loops" involve uncertainty and incompleteness, impeding the calculation of accurate estimates. These topics continue to be researched by the data and information fusion communities, whose processes form snapshots and estimates about component entities and situational states. However, the data and information fusion communities should also seek opportunities (where appropriate) for employing AI/machine learning beyond these component-level estimates, toward modeling synergies and dependencies within and across estimation loops to achieve maximum situational awareness.

The genesis of this book was the scholarly presentations given at the Association for the Advancement of Artificial Intelligence (AAAI) Spring Symposium, held during March 22–24, 2021, titled "Leveraging Systems Engineering to Realize Synergistic AI/Machine Learning Capabilities". However, we decided to refocus the title of the book to Engineering Artificially Intelligent Systems to emphasize pragmatic aspects. More information about the AAAI 2021 Spring Symposium series can be found online (https://aaai.org/Symposia/Spring/sss21.php).

Due to the worldwide COVID-19 pandemic the symposium was held virtually. We invited leading experts from around the world to explore the effects of cascading interdependencies, real-world dynamics, and subsequent opportunities to leverage systems engineering principles to design and develop AI/machine learning and data and information fusion estimation processes that accurately represent complex world states. These relationships are exacerbated by the unpredictability of human decision-makers, uncertainty in raw and fused data, and the large trade space of algorithm permutations orchestrated to solve given problems. The field of systems engineering brings opportunities to address and model the full range of complex, synergistic feedback loops in modern complex systems, toward the realization of cost-effective designs.

The chapters of this book are extensions of abstracts and presentations from the AAAI 2021 Spring Symposium on Leveraging Systems Engineering to Realize Synergistic AI/Machine Learning Capabilities. The topics included AI, machine learning and reasoning, data and information fusion, systems engineering, interdependence, human systems, human biases and limitations, trust and complex AI systems, and ethics. In particular, the call for papers sought contributions on the following:

- Artificial Intelligence, Machine Learning, and Reasoning
- Data and Information Fusion
- Systems Engineering of Complex Interdependent Subsystems
- Systems Engineering for Artificial Intelligence
- Artificial Intelligence for Systems Engineering
- Systems Engineering and Federated Learning
- Human Systems Integration and Visualizations
- Human Biases, Limitations, and Capabilities with Large Scale AI-based Systems
- Trust and Acceptance of Complex AI-based Systems
- Emergence in Complex System-of-Systems
- Ethics in Deploying Complex AI-based systems
- Societal Consequences of Interacting Complex AI-based systems
- Explainability and Interpretability of Artificial Intelligence
- Uncertainty Propagation in AI-based Complex Systems
- Modeling User Interactions with Large Scale AI-based Systems

More details about the symposium, including the full program, can be found on our webpage (https://sites.google.com/view/systems-engineering-ai-ml/home). In addition, we invited additional authors, who did not participate in our symposium, to join us by contributing chapters to this book.

In our deliberations and preparations for the symposium and the LNCS book, we were assisted by our Program Committee, some of whom became presenters at our symposium and contributors to this LNCS book. We thank them, we thank our invited speakers, and we thank our regular speakers.

August 2021

William F. Lawless
James Llinas
Donald A. Sofge
Ranjeev Mittu

Organization

Organizing Committee

W. F. Lawless	Paine College, USA
Ranjeev Mittu	U.S. Naval Research Laboratory, USA
Donald A. Sofge	U.S. Naval Research Laboratory, USA
Thomas Shortell	Lockheed Martin Company, USA
Thomas McDermott	Stevens Institute of Technology, USA
James Llinas (Chair)	University at Buffalo, USA
Julie L. Marble	Johns Hopkins University, USA

Publication Chair

W. F. Lawless	Paine College, USA

Program Committee

Raj Dasgupta	U.S. Naval Research Laboratory, USA
Ali K. Raz	Purdue University, USA
Erik Blasch	Air Force Office of Scientific Research, USA
Edward Waltz	Naval Postgraduate School, USA
Will Williamson	Naval Postgraduate School, USA
Ken Hintz	George Mason University/Perquire Research Co., USA
Noelle Brown	U.S. Naval Research Laboratory, USA
Michael Floyd	Knexus Research Corporation, USA
Rino Falcone	Institute of Cognitive Sciences and Technologies, Italy
Paul Robinette	University of Massachusetts Lowell, USA
Ciara Sibley	U.S. Naval Research Laboratory, USA

Addendum

The Association for the Advancement of Artificial Intelligence (AAAI) Symposium in the spring of 2021, originally scheduled to be held at Stanford University, USA, was instead converted into a virtual symposium conducted over Zoom as a result of the COVID-19 pandemic. Its reduced two-day agenda consisted of eight invited speakers who gave hour-long presentations and eight regular speakers who gave half hour presentations, along with the two speakers who gave opening comments on both days. All of the symposium papers were submitted to the lead organizer via EasyChair, who conducted a review of each paper that was accepted for presentation, assisted by the co-organizers (reducing potential conflicts of interest, the lead organizer did not give a talk at the symposium). In a normal year not affected by COVID-19, we would have had more speakers participating on the regularly scheduled third day of the symposium (usually a half-day).

Of the 18 speakers at the AAAI Spring Symposium, 10 papers were expanded into chapters and submitted to the lead editor of this LNCS book. An 11th chapter contributed by the lead organizer gives an introduction, an overview of the science, and a review of the chapters contributed, which was only reviewed by the lead editor of the book. Of the 10 content chapters, four were written by invited speakers at the symposium. Each of these 10 chapters was thus triple-reviewed: for the symposium, following editing for inclusion in this LNCS book, and for review comments as part of the introductions to each chapter.

Finally, having collected only 10 content chapters, a shortfall that we attributed partly to the COVID-19 pandemic, we invited five additional authors to contribute to the LNCS book. These invited chapters were reviewed twice, once by the LNCS lead editor for editorial content, and once by the symposium's lead organizer for the introduction to each chapter. All of the chapter contributors were asked to address all of the editorial comments, which they did.

The editorial process was arduous, but it assured that the contents of the book were uniformly edited, that conflicts of interest were prevented or minimized, and that the state of the science was advanced.

Contents

Motivations for and Initiatives on AI Engineering 1
James Llinas

Architecting Information Acquisition to Satisfy Competing Goals 19
Kenneth J. Hintz

Trusted Entropy-Based Information Maneuverability for AI Information Systems Engineering 34
Erik Blasch, Tod Schuck, and Oliver B. Gagne

BioSecure Digital Twin: Manufacturing Innovation and Cybersecurity Resilience 53
Michael Mylrea, Charles Fracchia, Howard Grimes, Wayne Austad, Gregory Shannon, Bill Reid, and Nathan Case

Finding the Path Toward Design of Synergistic Human-Centric Complex Systems 73
Hesham Y. Fouad, Ali K. Raz, James Llinas, William F. Lawless, and Ranjeev Mittu

Agent Team Action, Brownian Motion and Gambler's Ruin 90
Ira S. Moskowitz

How Deep Learning Model Architecture and Software Stack Impacts Training Performance in the Cloud 109
Egor Bykov, Evgeny Protasenko, and Vladimir Kobzev

How Interdependence Explains the World of Teamwork 122
Matthew Johnson and Jeffrey M. Bradshaw

Designing Interactive Machine Learning Systems for GIS Applications 147
Jaelle Scheuerman, Chris J. Michael, Brad Landreneau, Dina M. Acklin, and Jason L. Harman

Faithful Post-hoc Explanation of Recommendation Using Optimally Selected Features 159
Shun Morisawa and Hayato Yamana

Risk Reduction for Autonomous Systems 174
Tony Gillespie

Agile Systems Engineering in Building Complex AI Systems 192
Subrata Das, Zainab Ali, Sai-Nishant Bandi, Ankush Bhagat, Nithya Chandrakumar, Pritesh Kucheria, Mariella Pariente, Anand Singh, and Blake Tipping

Platforms for Assessing Relationships: Trust with Near Ecologically-Valid Risk, and Team Interaction 209
Julie L. Marble, Ariel M. Greenberg, Justin W. Bonny, Sean M. Kain, Brandon J. Scott, Ian M. Hughes, and Mary E. Luongo

Design Principles for AI-Assisted Attention Aware Systems in Human-in-the-Loop Safety Critical Applications 230
Max Nicosia and Per Ola Kristensson

Interdependence and Vulnerability in Systems: A Review of Theory for Autonomous Human-Machine Teams 247
W. F. Lawless and Donald A. Sofge

Principles of an Accurate Decision and Sense-Making for Virtual Minds 260
Olivier Bartheye and Laurent Chaudron

Author Index 281