Editorial Introduction to the Special Articles in the Spring Issue

Innovative Applications of Artificial Intelligence 2016

Peter Z. Yeh and James Crawford

■ *This issue features expanded versions of articles selected from the 2016 AAAI Conference on Innovative Applications of Artificial Intelligence held in Phoenix, Arizona. We present a selection of three articles that describe deployed applications, two articles that discuss work on emerging applications, and an article based on the 2016 Robert S. Engelmore Memorial Lecture.*

AAAI's Innovative Applications of Artificial Intelligence Conference was founded in 1989 to showcase the successful application of artificial intelligence technology to real-world problems and its deployment into the hands of end users. Since then, we have seen examples of AI applied to domains as varied as medicine, education, manufacturing, transportation, user modeling, military operations, and citizen science. This year, the 2016 conference continued the tradition with a selection of deployed applications describing systems in use by their intended end users, emerging applications describing works in progress, and challenge problem papers that discuss associated research and development challenges in applying artificial intelligence to real-world problems.

Our first article is based on the IAAI-16 Robert S. Engelmore Memorial Lecture given by Reid G. Smith at AAAI/IAAI 2016 in honor of Bob Engelmore's extraordinary service to AAAI and his contributions to applied AI. The article Smith and coauthor Joshua Eckroth later wrote is titled Building AI Applications: Yesterday, Today, and Tomorrow. It focuses on changes in the world of computing over the last three decades that made building AI applications more feasible. The article also examines lessons learned during this time (drawing from experiences that have been documented in IAAI conference proceedings since 1989), and distills these lessons into succinct advice for future application builders.

In the second article, PAWS — A Deployed Game-Theoretic Application to Combat Poaching, Fei Fang, Thanh H. Nguyen, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Brian C. Schwedock, Milind Tambe, and Andrew Lemieux describe a deployed game-theoretic application for optimizing foot patrols to combat poaching in Southeast Asia. The authors report on significant evolution of PAWS from a proposed decision aid to a regularly deployed application over the last two years. Key technical advances that led to PAWS's regular deployment

  ISSN 0738-4602

are outlined. These advances include incorporating complex topographic features, handling uncertainties in species distribution, handling complex patrol scheduling constraints, and more. The authors also report key lessons learned ranging from the importance of firsthand immersion in the security environment of concern to minimizing the need for extra equipment in order to further ease future deployments of PAWS. The benefit of PAWS to its intended end users was demonstrated by the continued deployment of PAWS at existing sites in Malaysia, and steps taken at the time of writing to expand PAWS to additional sites.

In our third article, Deploying nEmesis: Preventing Foodborne Illness by Data Mining Social Media, Adam Sadilek, Henry Kautz, Laren DiPrete, Brian Labus, Eric Portman, Jack Teitel, and Vincent Silenzio describe a deployed application that automatically detects venues likely to pose a public health hazard by applying machine-learning techniques to Twitter data. The authors demonstrate nEmesis's efficacy in the Las Vegas metropolitan area in a double-blind experiment conducted over three months in collaboration with Nevada's health department, and show that the deployed application is 64 percent more effective at identifying problematic venues than the current state of the art. If fully deployed, the nEmesis approach could prevent more than 9000 cases of foodborne illness and 557 hospitalizations annually in Las Vegas alone.

In our fourth article, Ontology Reengineering: A Case Study from the Automotive Industry, Nestor Rychtyckyj, Venkatesh Raman, Baskaran Sankaranarayanan, P. Sreenivasa Kumar, and Deepak Khemani discuss an effort to reengineer an existing ontology deployed at Ford. Ford has been utilizing an AI-based system to manage process planning for vehicle assembly at its assembly plants around the world for more than 25 years. The knowledge about Ford's processes is contained in an ontology originally developed using the KL-ONE representation language and methodology. However, the scope of this AI system has increased over the years to include additional functionality on ergonomics and powertrain assembly. Hence, the existing KL-ONE ontology needs to be reengineered into a semantic web OWL/RDF ontology to satisfy the increased scope of the AI system and to enable other applications within Ford to easily make use of it as well.

In our fifth article, Automated Volumetric Intravascular Plaque Classification Using Optical Coherence Tomography, Ronny Shalev, Daisuke Nakamura, Setsu Nishino, Andrew M. Rollins, Hiram G. Bezerra, David L. Wilson, and Soumya Ray describe an emerging application to identify different plaque types in blood vessel images using machine-learning methods. An estimated 17.5 million people died from a cardiovascular disease in 2012, and most acute coronary events result from rupture of the protective fibrous cap overlying an atherosclerotic plaque. Hence, early identification of plaque types that can potentially rupture is of great importance. The state-of-the-art approach to imaging blood vessels is intravascular optical coherence tomography (IVOCT), but this is an offline approach where the images are first collected and then manually analyzed one image at a time to identify regions at risk. This process is extremely laborious, time consuming, and error prone. Initial empirical results presented by the authors using real OCT data show that the proposed approach can identify different plaque types efficiently and with high accuracy across multiple patients.

In our sixth article, Using Global Constraints to Automate Regression Testing, Arnaud Gotlieb and Dusica Marijan describe an emerging application to automate regression testing. Regression testing is a crucial verification step in the software development and release process, but the selection of test cases is challenging due to several factors such the limited time available for testing. This problem, called test suite reduction (TSR), is usually addressed by validation engineers through manual analysis or by using approximation techniques. To address these limitations, the authors apply AI techniques such as constraint programming and global constraints to automate the process. Moreover, the authors' work is conducted in the context of an industrial application in the communication domain with the goal to eventually deploy the emerging application to test a complete product line of conferencing systems in continuous delivery mode. Initial experimental evaluations show promising results.

In the tradition of previous special issues on innovative applications of artificial intelligence, and consistent with the goals of the IAAI conference, the articles in this issue describe work that is strongly grounded in the needs of end users. We hope that you enjoy the articles, and that they both provide insight into the application development process and help to expand your view of what is possible with AI technology. We also invite you to submit a description of your next AI application to IAAI.

Peter Z. Yeh is a senior manager of AI technology and senior principal scientist at Nuance Communications. His research interests lie at the intersection of semantic technologies, data and web mining, and natural language understanding. At Nuance, Yeh's research focuses on developing the next generation of intelligent virtual assistants and the underlying AI technologies necessary to enable this. Prior to joining Nuance, he was a research lead at Accenture Technology Labs where he was responsible for investigating and applying AI technologies to various enterprise problems ranging from data management to advanced analytics. Yeh is the author of more than 40 papers in peer-reviewed journals and conferences, and holds eight patents. He received his Ph.D. in computer science from the University of Texas at Austin.

James Crawford is the founder and chief executive officer at Orbital Insight, Inc. He has held distinguished positions at NASA, led startups, and in 2009 became engineering director for Google Books in charge of scanning the world's books. Crawford served as vice president of engineering and executive vice president of engineering at Composite Software, Inc. He served at Ames Research Center, NASA's center of excellence for information technology. Prior to joining NASA, he led the optimization team at i2 Technologies, worked at AT&T Bell Laboratories, and cofounded the Computational Intelligence Research Laboratory (CIRL) at the University of Oregon. Crawford is the author of more than 15 papers in referred journals and conferences, and holds five patents. He earned his Ph.D. in artificial intelligence and master's degree in computer science from the University of Texas at Austin.