

THE IEEE

Intelligent Informatics

BULLETIN



IEEE Computer Society
Technical Committee
on Intelligent Informatics

August 2017 Vol. 18 No. 1 (ISSN 1727-5997)

Communications..... *Vijay Raghavan* 1

Conference Report

AI as a 30+ Year Florida Statewide Strategic Initiative..... *Pierre Larochelle, Marius Silaghi and Vasile Rus* 2

R&D Profile

Intelligent Informatics and the United Nations: A Window of Opportunity..... *Karen Judd Smith* 5

Short Feature Articles

A Consultant Framework for Natural Language Processing in Integrated Robot Architectures..... *Tom Williams* 10

Process Mining for Trauma Resuscitation
..... *Sen Yang, Jingyuan Li, Xiaoyi Tang, Shuhong Chen, Ivan Marsic and Randall S. Burd* 15

Selected PhD Thesis Abstracts..... *Xin Li* 20

Book Review

Cognitive Computing: Theory and Applications..... *Weiyi Meng* 34

IEEE Computer Society Technical Committee on Intelligent Informatics (TCII)

Executive Committee of the TCII:

Chair: Chengqi Zhang
University of Technology, Sydney,
Australia
Email: chengqi.zhang@uts.edu.au

Vice Chair: Yiu-ming Cheung
(membership, etc.)
Hong Kong Baptist University, HK
Email: ymc@comp.hkbu.edu.hk

Jeffrey M. Bradshaw
(early-career faculty/student mentoring)
Institute for Human and Machine
Cognition, USA
Email: jbradshaw@ihmc.us

Dominik Slezak
(conference sponsorship)
University of Warsaw, Poland.
Email: slezak@mimuw.edu.pl

Gabriella Pasi
(curriculum/training development)
University of Milano Bicocca, Milan, Italy
Email: pasi@disco.unimib.it

Takayuki Ito
(university/industrial relations)
Nagoya Institute of Technology, Japan
Email: ito.takayuki@nitech.ac.jp

Vijay Raghavan
(TCII Bulletin)
University of Louisiana- Lafayette, USA
Email: raghavan@louisiana.edu

Past Chair: Jiming Liu
Hong Kong Baptist University, HK
Email: jiming@comp.hkbu.edu.hk

The Technical Committee on Intelligent Informatics (TCII) of the IEEE Computer Society deals with tools and systems using biologically and linguistically motivated computational paradigms such as artificial neural networks, fuzzy logic, evolutionary optimization, rough sets, data mining, Web intelligence, intelligent agent technology, parallel and distributed information processing, and virtual reality. If you are a member of the IEEE Computer Society, you may join the TCII without cost at <http://computer.org/tcsignup/>.

The IEEE Intelligent Informatics Bulletin

Aims and Scope

The IEEE Intelligent Informatics Bulletin is the official publication of the Technical Committee on Intelligent Informatics (TCII) of the IEEE Computer Society, which is published twice a year in both hardcopies and electronic copies. The contents of the Bulletin include (but may not be limited to):

- 1) Letters and Communications of the TCII Executive Committee
- 2) Feature Articles
- 3) R&D Profiles (R&D organizations, interview profile on individuals, and projects etc.)
- 4) Selected PhD Thesis Abstracts
- 5) Book Reviews
- 6) News, Reports, and Announcements (TCII sponsored or important/related activities)

Materials suitable for publication at the IEEE Intelligent Informatics Bulletin should be sent directly to the Associate Editors of respective sections.

Technical or survey articles are subject to peer reviews, and their scope may include the theories, methods, tools, techniques, systems, and experiences for/in developing and applying biologically and linguistically motivated computational paradigms, such as artificial neural networks, fuzzy logic, evolutionary optimization, rough sets, and self-organization in the research and application domains, such as data mining, Web intelligence, intelligent agent technology, parallel and distributed information processing, and virtual reality.

Editorial Board

Editor-in-Chief:

Vijay Raghavan
University of Louisiana- Lafayette, USA
Email: raghavan@louisiana.edu

Managing Editor:

William K. Cheung
Hong Kong Baptist University, HK
Email: william@comp.hkbu.edu.hk

Assistant Managing Editor:

Xin Li
Beijing Institute of Technology, China
Email: xinli@bit.edu.cn

Associate Editors:

Mike Howard (R & D Profiles)
Information Sciences Laboratory
HRL Laboratories, USA
Email: mhoward@hrl.com

Marius C. Silaghi
(News & Reports on Activities)
Florida Institute of Technology, USA
Email: msilaghi@cs.fit.edu

Ruili Wang (Book Reviews)
Inst. of Info. Sciences and Technology
Massey University, New Zealand
Email: R.Wang@massey.ac.nz

Sanjay Chawla (Feature Articles)
Sydney University, NSW, Australia
Email: chawla@it.usyd.edu.au

Ian Davidson (Feature Articles)
University at Albany, SUNY, USA
Email: davidson@cs.albany.edu

Michel Desmarais (Feature Articles)
Ecole Polytechnique de Montreal, Canada
Email: michel.desmarais@polymtl.ca

Yuefeng Li (Feature Articles)
Queensland University of Technology,
Australia
Email: y2.li@qut.edu.au

Pang-Ning Tan (Feature Articles)
Dept of Computer Science & Engineering
Michigan State University, USA
Email: ptn@cse.msu.edu

Shichao Zhang (Feature Articles)
Guangxi Normal University, China
Email: zhangsc@mailbox.gxnu.edu.cn

Xun Wang (Feature Articles)
Zhejiang Gongshang University, China
Email: wx@zjgsu.edu.cn

Publisher: The IEEE Computer Society Technical Committee on Intelligent Informatics

Address: Department of Computer Science, Hong Kong Baptist University, Kowloon Tong, Hong Kong (Attention: Dr. William K. Cheung; Email: william@comp.hkbu.edu.hk)

ISSN Number: 1727-5997(printed)/1727-6004(on-line)

Abstracting and Indexing: All the published articles will be submitted to the following on-line search engines and bibliographies databases for indexing—Google(www.google.com), The ResearchIndex(citeseer.nj.nec.com), The Collection of Computer Science Bibliographies (liinwww.ira.uka.de/bibliography/index.html), and **DBLP** Computer Science Bibliography (www.informatik.uni-trier.de/~ley/db/index.html).

© 2017 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the **IEEE**.

Message from the Editor-in-Chief

The IEEE Intelligent Informatics Bulletin (IIB) is the official publication of the Technical Committee on Intelligent Informatics (TCII) of the IEEE Computer Society. It aims at promoting the excellent research on the wide spectrum of intelligent informatics, as well as displaying the news/reports on TCII related conferences and other activities. This issue- the first of its kind- of the IEEE Intelligent Informatics Bulletin makes a novel attempt publishing a carefully selected collection of Ph.D. thesis abstracts. Only these defended during the period of January 2016 to July 2017 are targeted considering timeliness. This issue also includes, a Report edited by Marius Silaghi, a profile by Karen J. Smith (edited by Mike Howard), two short feature articles, respectively by Tom Williams and Sen Yang, and a book review by Weiyi Meng (edited by Ruili Wang).

This new attempt of the Intelligent Informatics Bulletin intends to offer an opportunity for young talented scholars at the beginning of their career to introduce their previous and ongoing research work to potentially interested scholars in similar intelligent informatics subdomains. To achieve this goal, it is required that a URL to Ph.D. theses' soft copy should be included with the submitted abstract, and this URL should be accessible and properly maintained continuously.

In total, 43 abstracts of the Ph.D. thesis were successfully submitted. 55 authors come from 21 different countries located in Asia, North America, Europe and Africa. Each submitted abstract was reviewed by at least 2 editorial board members of the bulletin. According to the ranking of reviewing scores, only the top 24 submissions were selected to publish in this Ph.D. thesis abstracts section. The selected submissions cover numerous research and applications under intelligent informatics such as robotics, cyber-physical systems, social intelligence, intelligent agent technology, recommender systems, complex rough sets, multi-view learning, text mining, general-purpose learning systems, Bayesian modeling, knowledge graphs, genetic programming and dynamic adversarial mining.

The key issues of this Ph.D. thesis abstracts section are summarized above. The organizers wish that this special issue can bridge different subdomains of intelligent informatics and encourage further communications and cooperation within the whole scholar community. Special thanks are due to Prof. Xin Li for spearheading this effort.

In addition, we are pleased to announce that the publication of selected PhD abstracts will be continued in future years. Thus, we plan to have two issues of IIB published each year, one in summer and one in winter. We welcome your feedback on this special issue, as well as future issues of IIB.

Vijay Raghavan
(University of Louisiana- Lafayette, USA)
Editor-in-Chief
The IEEE Intelligent Informatics Bulletin

AI as a 30+ Year Florida Statewide Strategic Initiative

30 YEARS ANNIVERSARY OF FCRAR AND FLAIRS

I. A STATE OF FLORIDA INITIATIVE

Approximately 30 year ago, in 1987, an effort of the state of Florida was channeled toward the exploitation of the local intellectual and technical infrastructure for the advancement of local research, in particular through the support and organization of conferences. Groups of scientists from the main Florida universities and representing various key research fields were invited to Tallahassee and encouraged to discuss ways of organizing events, such as local conferences, that would support their technical or scientific activities.

Among these groups, two focused on research directions that are currently converging into topics related to intelligent informatics. These were the group specialized on artificial intelligence (AI) and the group specialized on robotics. Each of the two groups ended up giving rise to a corresponding successful conference, while each of them followed a very different approach towards a distinct vision. The artificial intelligence group decided to create a research society and an international conference, FLAIRS (Florida Artificial Intelligence Research Society), that would bring together researchers from all over the world, exploiting not only the local research base but also the available rich local attractions. The robotics group decided to construct an intimate conference, FCRAR (Florida Conference on Recent Advances in Robotics), coaching local graduate students on research. FLAIRS became a top scientific conference with indexed proceedings while FCRAR is a



FLAIRS: Plenary Session

vibrant conference with free registration and “lightly” reviewed online proceedings.

II. FCRAR 30: THE 30TH FLORIDA CONFERENCE ON RECENT ADVANCES IN ROBOTICS

The FCRAR 2017 conference was organized during May 11-12 at the Florida Atlantic University (FAU) by Dr. Oren Masory and Dr. Zvi Roth. The conference is a traditional venue for presenting ongoing robotic research in a formal environment. It has originally focused on the mechanical part of robotics but artificial intelligence proved to be inseparable from its topic and is now addressed by a significant part of the presentations.

As with previous editions, the 2017 FCRAR did span over two days. In most years FCRAR is a single track conference. However, two parallel sessions were needed this year for the 49



FCRAR Robotic Jellyfish Demo

regular articles accepted. Still, the whole attendance met for the first session and welcome message, as well as for the three valuable invited talks:

A. The first invited talk, by Dr. Andrew Goldenberg, Professor Emeritus from the University of Toronto, provided insights from his rich experience with transferring technology from academia to industry, as well as to the market size and trends for various types of robots. A glimpse was offered into the extent to which robotics is planned to quickly replace significantly more service jobs in the fast food industry. His advice to students stressed the fact that challenges remain mainly with the artificial intelligence component, which persists

as an important constraint and bottleneck of the democratization of the robot. Challenges posed by the development of robots for prostate



FAU Robotic Lab Expo.

surgery under MRI, as well as exploitation of IP in academia made the subject of most questions after the talk.

B. The invited talk of Dr. Deborah Nagle focused on her experience as a robotic surgeon and researcher with surgery robotics. A detailed account was given of the stages of evolution in robotic surgery, as well as of difficulties faced by surgeons in using the equipment. The emphasis was put on the desire of surgeons to get more help in terms of machine learning techniques, to help them during medical procedures. While the surgery robotics market is hard to penetrate, the speaker explained challenges by newcomer Google to the established DaVinci manufacturer.

C. The third invited speaker was University of Florida's Dr. Gloria Wiens, Director of FloridaMakes BRIDG, who described governmental efforts for supporting local innovation with current projects in domains such as smart kitchens and smart agriculture. Most questions from attendees concerned the history of results earned by similar centers.

The conference also contained a robotic expo session where, under the leadership of Dr. Erik Engeberg, participants were shown the robotic jellyfish manufactured at the FAU, as well as their large collection of robotic arms and sensors. Visitors experimented with the control of industrial robots

using gloves, thought, or other techniques to simply capture and replicate human motion or desires.

In regular sessions, presenters addressed various robotic prototypes as well as machine learning techniques to learn patterns and interpret images. New machine learning techniques for detecting seizures in EKG and recognizing fish sounds using Slantlet transforms were described by Ibrahim. Search heuristics described by Medina were part of AI techniques to explore space in NASA Swarmathon. ID3-like entropy reducing decision tree building heuristics were described by Ballard for eye in arm robots exploring workspaces. Heuristics for graph traversal used in remote laser-based soldering of cracks detected visually were described by Findling, while Drada discussed regular soldering systems. A family of path-planning techniques for inspection of nuclear waste tanks were described by Sebastian Zanlongo, extending A* and Disjktra's algorithms with heuristics modeling details concerning travel difficulties and tether cables. Other addressed robotics applications include intelligent robots to carry luggage, the BudeE robot for surveillance of children, vertical farming, construction robots, autonomous buoys. AI related problems addressed span strategy planning, path-planning algorithms, component failure prediction, user trust, and frustration management.

III. FLAIRS 30: THE 30TH FLORIDA ARTIFICIAL INTELLIGENCE RESEARCH SOCIETY CONFERENCE

The 2017 FLAIRS conference was organized by Dr. Ingrid Russell during May 22-24 at the Marco Island Hilton Hotel. As in previous editions, a set of special tracks was integrated with the main conference, each special track having its own topic and program committee. Acceptance criteria were unified by the program committee consisting of Vasile Rus, Zdravko Markov, and Keith Brawner. This edition accommodated 17 special tracks and 192 participants. The conference program was organized in sets of 4 parallel sessions for the 103 accepted

full papers. A poster session was held during the first afternoon session on the first conference day. There were 61 posters, 25 of them having an abstract published in the proceedings that has undergone a separate reviewing process, while the rest corresponding to accepted short papers. Six papers were nominated for the best paper and best student paper awards.

Recent FLAIRS editions include lunches and dinners in the registration package, reducing noon breaks and helping participants to interact and discuss collaborations, one of the appreciated strengths of the conference. There were four special track invited talks, presented in parallel sessions and three main conference invited talks presented in plenary sessions:

A. *The first invited talk was offered by Thomas Dietterich, on "Robust Artificial Intelligence: Why and How". It was suggested to address lack of robustness by modeling randomness as an attacker, using minimax. General AI can take inspiration from the field of Automatic Speech Recognition (ASR) where noisy bands are detected based on unlikely repetition of phonemes. Four ideas emerge:*

- 1) Employ anomaly detectors
- 2) Use dynamic model extensions
- 3) Use causal directions, as they work better than diagnostic models
- 4) Combine techniques

Citing Minsky: "You understand only if you understand in multiple ways".

Public questions addressed: evolution, impact of small errors and wrong datasets, advantage of ensemble methods (portfolios) over thinking too much, and difference between building incomplete models versus building causal models. Science studies causal relations while engineering is satisfied with incomplete models. While ASR was improved with more data, that may not work for adversarial situations. Accordingly DARPA has a project on "Explainable AI".

B. The second plenary invited talk was delivered by Prof. Jiawei Han on "Mining Structures from Massive Text Data: A Data-Driven Approach". The speaker presented his research in using large text data sets to create relatively structured networks, usable as

knowledge. Questions raised the challenge of approximate knowledge needed to answer queries such as: "Who is the President of England?" and representation of nested concepts such as "President of [United States]".

C. The third plenary invited speech was given by Dr. James Allen on broad-coverage deep language understanding. The challenge question was for a definition of deep understanding. It was suggested that solutions include: no word left behind, preservation of details, preservation of ambiguity, and preservation of compositional structure. The TRIPS word organization capturing common sense semantics, proposed by the speaker, was contrasted with WordNet. Comprehensive dictionaries are now parsed to build ontologies. It was stressed that words are used for arguments rather than relations. Disambiguation is achieved by filtering semantics based on annotations and temporal constraints. However, construction & composition require



FLAIRS Awards Ceremony

invention of new semantics as in "The dog barked the cat up the tree". The questions addressed: parsing of emoticons, dots, templates of allowable constructions, different semantics, complex world in long novels, non-English text.

The "Radical-Based Hierarchical Embeddings for Chinese Sentiment Analysis at Sentence Level" by Peng et.al. won the best paper award. Natural language processing is one of the main strength of the FLAIRS community, but most other AI branches are represented.

**31th FLAIRS and 31th FCRAR
will take place during May
2018.**

Contact Information

Pierre Larochelle
South Dakota School of Mines and
Technology

Marius Silaghi
Florida Institute of Technology

Vasile Rus
University of Memphis

Intelligent Informatics and the United Nations: A Window of Opportunity

KAREN JUDD SMITH

Abstract —“Every one of us is responsible for the future of all of us.” Dr. Gro Harlem Brundtland, Former Prime Minister of Norway and former head of the World Health Organization.

K.A. Judd Smith, Co-Chair of the Alliance of NGOs on Crime Prevention and Criminal Justice (www.CPCJAlliance.org) and consultant to the United Nations Office of Information and Communications Technology, outlines the way the flood of new technologies has led to types of criminal activity that are cheaper, more anonymous, and with global reach. She challenges the tech sector to build protections of social good into their technologies, and to uncover novel ways to infuse compassion and empathy into the machinery and so make it safer and human friendly by integrating comprehensive perspectives of how we function socially as human beings.

Index Terms—computational and artificial intelligence, intelligent transportation systems, industry applications, systems, decision support systems, policy, ethics.

I. A STATE OF FLORIDA INITIATIVE

Ten years into the smart phone and mobile device phase of our digital revolution, transnational organized criminals have consistently been agile adopters and adapters of new technologies. Even the once-stable societies are being unsettled by this surge of change, the white noise of abundant information, and a rise in disinformation and “fake news.”

This triple play of an unprecedented pace of technological development, the human ability to integrate tools into our daily lives, and the inbuilt resistance to change of our major social institutions, is creating an ethical Gray Zone where the collective “we” strain toward anomie. Our increasing number of available sets of norms, none of which are clearly binding, blurs the clarity once afforded by social institutions in slower

times. With each day, evolving technologies multiply their impact on society leaving our twentieth century governing institutions unable to dampen the turbulence. Even the relative constancy of the most stable societies is being disrupted with dissent and family-dividing ideological differences. Fast-paced change is the new norm.




Fig. 1. The Gray Zone: This indicates the areas of social life increasingly unregulated by traditional mechanisms.

Both legitimate and criminal innovations outstrip the ability of our national and international institutions to create and implement timely and needed responses. Our key social institutions are caught in old ways of functioning that are in dire need of innovation.

It has now been seventeen years since the 2000 adoption of the United Nations Convention against Transnational Organized Crime (UNTOC). True, these years have seen an increasing number of UN member states that are party to the Convention, but little significant headway has been made.

This problem of transnational organized crime is not some distant, faceless issue. Rather, it affects the daily lives of men, women and children in all our cities as well as rural locations worldwide. Those on both sides of this dark economy are our neighbors. The anonymous agents streaming the rape of children to clients willing to pay with cryptocurrencies often have well known faces in our communities. Drug sales are

shifting from the street to the Dark Web. Zero-day exploits and sophisticated hacking tools can be bought easily in this masked economy.

Few of these types of criminal activities are totally new. What is different, however, are the increasing access, decreasing costs and increased anonymity. In short, we have an explosion of the Dark Web economy. This new territory, including social media and the Internet of Things (IoT), also brings with it, a significant decline in ethical and moral clarity. The twentieth-century forms of regulation and self-regulation still in place are unable to keep up with this increasing onslaught of changing variables.

Privacy issues blended with security concerns are also less and less easy to determine. The values associated with hacking are also complex. Thus, the democratization of information and access to evolving technologies have brought with them a plethora of opportunities both for the greater good and to the detriment of many. Few could have predicted what we face less than thirty years after the open sourcing of the World Wide Web back in 1989 by the British scientist at the Conseil Européen pour la Recherche Nucléaire (CERN), Tim Berners-Lee.

In addition to the freedoms, benefits, convenience, and remarkable progress resulting from this tidal surge of technological development, there are those on the dark side who use these technology tools without regard to human rights or the rule of law. Human trafficking, sexual exploitation, money laundering, hacking-as-a-service, drug and weapons sales, the ability to disseminate radical causes with destructive intent, etc. now fuel violent extremism and contribute to the destabilization of whole regions. Yet no

amount of sophisticated socio-political handwringing will ever be able to put this technological genie back into the bottle.

What we can do, however, is to look for opportunities to leverage institutional innovation from within this growing Gray Zone. To emphasize, we can and should determine to fight this technological fire with an at least equally dedicated technological counter-fire.

Technological advancement now enables us to pursue the social good in ways never before possible. Artificial intelligence (AI), machine language, the emergent fields of material engineering, bio-tech, and the use of biologically and linguistically motivated computational paradigms, place in our hands new ways to contribute to the global social good.

As new challenges are banging hard on our doors, there are reasons to believe we may be on the cusp of significant institutional-level innovation, if we put emerging intelligent technologies to work more strategically in our multicultural, multireligious, multinational, multi-disciplinary, and multipurposed endeavors.

Before turning our attention to one of today's underutilized and strategic opportunities for engagement, we will briefly consider the "Transilience Framework" which highlights key forces at play in any social good endeavor. These provide additional insight for the work at hand whether planning a project, assessing impact, or developing a strategy. This framework is not a strategic planning methodology with milestones and waypoints, but rather a way to assess the sea-conditions within which the endeavor takes place. As a decision-matrix of sorts, it factors in social, rational and neurological forces to our journey in today's Gray Zone.

Finally, we will conclude with the recommendation to pursue a new, if challenging partnership that could open up new levels of thinking, insight, and strategic impact. The timing is right. Just as mobile devices in Africa enabled the continent to bypass the need for building an extensive and expensive wired communications infrastructure, this

confluence of new resources and needs may well enable our global community to leap-frog over some of our most intransigent issues.

II. THE TRANSILIENCE FRAMEWORK

The definition of transilient as "passing quickly from one thing to another" comes from the Latin "transilire—to jump over." In the context of social change, this framework calls attention to our innate human capacity for inventiveness and ways to leverage this potent human capacity. Transilience pays particular attention to the transformative resilience of human beings, the factors affecting it and its role in the midst of turbulence, change, and uncertainty.

The three sets of factors this framework considers are Difference, Drivers, and Domain, each being part of a social good compass of sorts. The concept of True North here indicates pursuits that best serve our metahumanity, an holistic description that includes all aspects of optimal living for human beings as individuals, as social animals with a strong propensity for accounting, and as a species still reliant upon the health of our planet for our own health and well-being.

A. Difference

Difference recognizes the varying forces involved with change. The greater the degree of change sought or needed in a shorter space of time, will require different kinds of effort and leadership than circumstances that value stability, strict adherence to standards and system-wide trustworthiness. The former calls for bold, adaptive leadership to ensure that the culture, talent, and structures operate well in less familiar, riskier environments. The latter, management-oriented leadership focuses on managing the variables in a more closed system to ensure a quality of service output for all stakeholders. Each kind of leadership has its own place and they are not mutually exclusive. Capable leaders or project designers recognize the differing emphasis needed due to the nature of the work at hand.

Difference also takes into

consideration psychological time, that is, whether we are forward-looking or backward-looking. In reality, we cannot go back in time, but we can and do delve into our memories and our efforts. Mindset and vision tend to have an orientation that is forward and exploratory or backward and focused on sustaining a status quo.

In either case, whether the differences are larger or smaller, forward-looking or backward-looking the pathways being chosen by leaders or designers need to be adapted appropriately for the stakeholders and the forces they will encounter.




Fig. 2 Difference: This reminds us of the forces in play when endeavoring to make a difference. Resistance increases with as greater changes are sought. It also indicates that psychologically, we may be forward or backward looking in our focus.

As can be seen in Fig. 2, the greater the distance from the status quo sought, whether forward or backward, efforts are met with increasingly powerful forces or resistance to change. This resistance is compounded by the Drivers, as will be explained in the next section.

The forces at play here can be seen in the slow development of laws on cybercrime or the lack of global initiatives, for example, to globally change router security standards. So it remains possible to hack into a bank or a power plant and leave no trace because there is no extra layer of identification required. This is just one instance where legislators neither adequately understand technology nor include technologists in legislative decision-making in a timely manner.

B. Drivers

Drivers are both the rational and pre-rational dimensions of social

engagement. These are at play in all human interactions. Curiously, while it is known that humans tend to make decisions emotionally (or pre-rationally), our rational sense of self often overlooks this. Paying attention to these Drivers is an effort to acknowledge and factor in the full complexity of human interaction that is being described here as the ongoing interaction of our (evolutionarily speaking) “three brains” and the impact of each of these on our moment-to-moment decision-making.

This is an essential consideration to note because all too often the more analytical, rational actors (such as researchers and engineers) overlook these pre-rational considerations during development. This reverence for rationality can lead to the filtering out of the non-rational as irrelevant. This is especially true when intuitive responses stand in direct opposition of what makes rational, long-term sense. We then end up surprised by voting results, the eruption of violence, and unexpected aspects of social relationships.




Fig. 3. The Drivers: This figure identifies in a symbolic manner, the pre-rational and rational dimensions of human decision-making: our reptilian complex, paleomammalian complex and neomammalian complex. [1] Paying attention to any one alone or ignoring any of these, can lead to poor judgment and misreading of social dynamics.

When we factor in both the logical complexity of changes being sought and the pre-rational drivers that are stimulated by those changes, we can make greater sense out of some of the political dynamics today. Nationalist surges are showing up in many countries in part as a reaction to underlying uncertainties—politically, socially and economically—as technologies hasten the pace of change. The body-language of today’s changing times stimulates our

reptilian complex and our limbic systems which, when unmediated by our more logical neocortex, naturally stirs up the desire to retreat to something more familiar and safe. We see this playing out in the Brexit vote and in the US response to “making America great again,” to mention just two of many phenomena.

The Transilience Framework, therefore, encourages consideration be given to the impact new tech may have on the pre-rational, lizard and limbic aspects of ourselves in addition to the tasks they are being designed to fulfill. We cannot change these hard-wired dimensions of our brains. But by being more fully aware of how much they impact our behavior, we can work to put each to better use. Then perhaps during the next US election, pollsters will find novel ways to measure the temperature all three areas of human concern: survival issues, social concerns and preferred political strategies.

C. Domain or Scope

The final component woven into our social good strategies as indicated by the Transilience Framework addresses the scope of work. Projects and services can target an individual or larger groupings and social institutions. The latter can range from family units, through to national and international structures. As for most systems, moving from the inner individual levels to the outer larger institutions means a shift to significantly more complex structures. Identifying the scope of work facilitates establishing time frames, methodologies, resources, the clarification of which benefits the management expectations.

How does this help identify strategic actions for today’s evolving intelligence informatics community? First, since the focus here is on developing potential partnerships with some aspect of the United Nations, it will help to clarify some notions of the UN. There is a tendency for people to view the UN as an entity that is responsible for peace and security at the global level. This is however, false. While it may be the most global oriented entity we have, it does not have global level authority. It is not

structured as an entity with a mandate to secure world peace. It can advise nations about peace and security and facilitate governmental negotiations on the issues, but the UN itself was never globally purposed. There was, and remains, an aversion to a single “global state” and its accompanying notion of a common political authority for all of humanity.

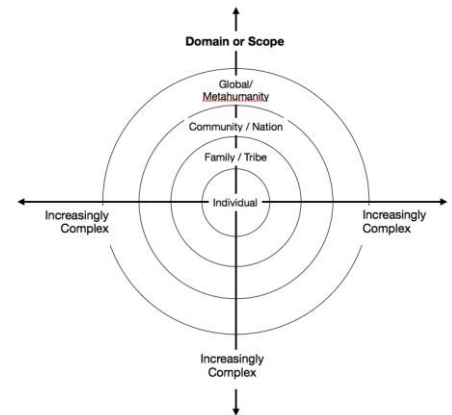


Fig. 4 Domain: This identifies the scope of work at hand ranges from the individual level to that of metahumanity. The degree of complexity of implementing solutions increases exponentially making global issues enormously more complex. This suggests the most critical and perhaps strategic space for the augmentation of human efforts by intelligence informatics is at the global or metahumanity level.

Though the UN is the agreed meeting place where governmental representatives negotiate according to their national agendas and interests. In this relatively transparent environment, Member States often do operate from enlightened self-interest. However there is nothing stopping them from putting national interests above the interests of other nations and peoples. The UN’s Charter makes this clear, that it pursues peace *between* nations; neighborly relations *among* its members; and respect *for* each other’s national sovereignty.

The Charter does encourage the use of “international machinery for the promotion of the economic and social advancement of all people,” [2] but this too is an endeavor of nations. Then there is the 1948 Universal Declaration of Human Rights that has enlightened and strengthened the work of those seeking peaceful solutions. But this universal instrument does not change the fact that

the UN does not have global authority.

When someone asks, “What did the UN do about the UN Convention on Transnational Organized Crime (UNTOC)?” for example, the more telling rephrasing of that question would be something like this series of questions: “What did each nation agree to in that meeting?” “What are they now doing to implement their commitments?” “How do we know?”

While the UN may be the one social institution endowed with the moral authority to act for the greater good, it is not structured to do so. Hence the phrase often heard, “UN resolutions have no teeth.” But none of these and other factors stop people from questing for the betterment of humanity as a whole with or without a global authority established to facilitate this.

As the ubiquity of the Internet and its attendant web-based intelligence, communications and finances evolve, it becomes clearer that significant trans-national organizational issues remain to be addressed. The UN remains the most obvious, yet still limited entity for this work. What is also needed are self-organizing efforts to step up to the plate. That is how the UN International Criminal Court (ICC) and the 1997 Convention on the Prohibition of the Use, Stockpiling, Production and Transfer of Anti-Personnel Mines came into existence. Those with vision, expertise and drive step up and make something happen.

III. GLOBAL GOVERNANCE AND INTELLIGENT INFORMATICS

Many of our global social challenges are already incorporated into the UN’s 2030 Sustainable Development Goals (SDGs) [3], into the work of the United Nations Office of Drugs and Crime, the Security Council, and taken up by its multiple businesses and agencies such as UNICEF, UNESCO, WHO, and dozens of others. With focused effort, the right kinds of leadership in appropriate places and technical savvy, much of this good work could be upgraded and improved further.

The UN already has innovation labs in

field offices around the world bringing together technologists with local social activists in pursuit of the UN’s development mandates. The UN also has significant efforts underway to optimize its technical infrastructure, though there is much more that can be done there too.

The more complex work of innovative uses of existing data and domain experts to provide intelligence for governance in a complex system such as the UN is still missing. But now that the multidisciplinary intelligence and security informatics communities are maturing, and the UN is creating digital platforms to better enable innovative partnerships and engagement with technically oriented civil society and the private sector, these crossing paths provide an unusual opportunity to leap forward and benefit the greater global community.

Bringing the diverse mindsets of technologists and global legislators to work together on some of the stickiest global issues has the potential to be a powerful difference engine fueling needed innovations.

There is yet another gain that could be made. If teams from the think tanks and institutes currently funded by the likes of Steven Hawking, Elon Musk, and Peter Thiel (to name just a few concerned for AI’s existential threats to the future of humanity), would forge partnerships with the UN’s Digital Blue Helmets for example, these small, targeted cross-disciplinary partnerships could work *in situ*. They would have easier access to the data and resources of the UN’s 50-plus agencies and organizations and their domain experts otherwise unavailable to AI technologists. There may be no better work place available for infusing human friendly compassion and empathy into intelligence systems than that alongside those striving to fulfill the UN’s challenging mandates.

The Digital Blue Helmet platform has already been established by forward-looking technologists at the UN for “rapid information exchange and better coordination of protective and defensive measures against information

technology security incidents for the United Nations, its agencies, funds, and programmes.” [4] What is needed now is for those in civil society, the private and philanthropic sectors to step up to these global challenges and infuse the UN’s work with additional innovative energies. Now that Google’s AI has challenged and beaten the world-champion Go player, Ke Jie, it is time to put AI, machine learning, robotics, bioinformatics, big data and the best of the best to work on something even more complex and challenging: world peace.

By bringing web intelligence resources to the challenges such as those the Digital Blue Helmets are contending with: transnational organized crime, cybersecurity, crisis intervention, and the multitude of uses of today’s technologies for the global good, we can all benefit. Of course there are risks and concerns at which those in the C-Suites of the UN, the private sector and academia will need to look. But in principle, it is wise for us all to keep an eye on the clock, and assess the risks both of action and inaction.

Governance intelligence may neither have the same dramatic optics of a moon landing nor the feel of the latest piece of shiny tech; however, bold steps with metahumanity in mind need to be taken today. The urgency of this can be measured in the numbers of people going to the streets in protest, the escalating threats of violence and war, the need for faster, better response to natural and man-made crises, the millions of displaced persons looking for refuge, the economic challenges of rapidly changing job markets, and more. A window of opportunity is opening right now—what better time for the odd couple of technologists and global legislators working for the social good decide to self-organize and step up to tackle the complex problems that affect us all?

IV. REFERENCES

- [1] MacLean, Paul D., and Vojtech A. Kral. A Triune Concept of the Brain and Behaviour.: Including Psychology of Memory [U.a.] :

Papers Presented at Queen's University, Kingston, Ontario, February 1969 / by V.A. Kral. Toronto [u.a.]: Univ. of Toronto Pr, 1973. The triune brain is a model of the evolution of the vertebrate forebrain and behavior.

- [2] United Nations: Department Of Public Information. *Charter of the United Nations and Statute of the International Court of Justice*. [Place of publication not identified]: United Nations, 2015.
- [3] UN's 2030 Sustainable Development Goals, <http://www.undp.org/content/undp/en/home/sustainable-development-goals.html>.
- [4] The UN' s Digital Blue Helmets, <https://unite.un.org/digitalbluehelmets/>

Contact Information

Dr. Karen Judd Smith
Co-Chair, Alliance of NGOs on Crime
Prevention and Criminal Justice
Consultant, United Nations Office of
Information and Communications
Technology
Email: karenjuddsmith@gmail.com

A Consultant Framework for Natural Language Processing in Integrated Robot Architectures

Tom Williams

Abstract—One of the goals of the field of human-robot interaction is to enable robots to interact through *natural language*. This is a particularly challenging problem due to the uncertain and open nature of most application domains. In this paper, we summarize our recent work in developing natural language understanding and generation algorithms. These algorithms are specifically designed to handle the uncertain and open-worlds in which interactive robots must operate, and use a *Consultant Framework* specifically designed to account for the realities of integrated robot architectures.

Index Terms—human-robot interaction; natural language understanding; natural language generation; natural language pragmatics; integrated robot architectures

I. INTRODUCTION

ENGAGING in task-based natural language interactions in realistic situated environments is incredibly challenging [1]. This is especially true for robotic agents for three reasons. First, their knowledge is woefully *incomplete*, both physically (only having knowledge of a small number of objects, people, locations, and so forth) and socially (only having knowledge of a small number of social norms). Second, their knowledge is highly *uncertain* due perceptual and cognitive limitations. Finally, natural language is highly *ambiguous*. As such, language-enabled robotic architectures must be designed to handle uncertainty, ignorance, and ambiguity at each stage of the natural language pipeline.

In this article, we will summarize research performed in the Human-Robot Interaction laboratory at Tufts University in service of this goal¹. Specifically, we will discuss research intending to account for uncertainty, ignorance, and ambiguity in the *referential* and *pragmatic* components of our robot architecture, DIARC [3], as implemented in the Agent Development Environment (ADE) [4]². We will begin by describing our *mnemonic architecture*: a *Consultant Framework* designed to facilitate domain independent memory retrieval in uncertain, open and ambiguous worlds. We will then describe our *linguistic architecture*: the language-processing components of our robot architecture which benefit from those mnemonic design choices.

Tom Williams was with the Department of Computer Science at Tufts University, Medford MA 02145 USA. E-mail: williams@cs.tufts.edu (See also: <http://inside.mines.edu/~twilliams>).

¹Specifically, we summarize work that contributed to the author's doctoral dissertation [2]

²While our laboratory has also examined these topics at other stages of the natural language pipeline [5], that work is beyond the scope of this review.

II. MNEMONIC ARCHITECTURE

A. Distributed Heterogeneous Knowledge Bases

In integrated robot architectures (e.g., [3], [6]), information may be distributed across a variety of architectural components. Information about objects may be stored in one location, information about locations in another, and information about people and social relationships in yet another. Furthermore, each of these stores of knowledge may use a very different representational framework. The set of architectural components capable of providing information about entities that may be referenced in dialogue can thus be viewed as a set of *distributed, heterogeneous knowledge bases* [7]. Our mnemonic architecture makes the following assumptions [2] about such *DHKBs*: (1) Each DHKB has knowledge of some set of *entities*; (2) A subset of knowledge regarding each entity can be *accessed* through introspection and described using positive arity predicate symbols; (3) Any knowledge that can be accessed in this way can also be *assessed* as to strength of belief; and (4) Each sort of knowledge that can be accessed, assessed, and described can also be *imagined* to hold for a given entity.

B. Consultants

These assumptions are exploited using a set of *Consultants*. Each architectural Consultant provides domain-independent access to a particular DHKB, allowing other components to access, assess, and imagine entities without needing to know anything about how such entities are represented (see also Fig. 1). To facilitate this, each Consultant provides four capabilities: (1) providing a set of atomic entities assessable through introspection; (2) advertising a list of predicates that can be assessed with respect to such entities, listed according to a descending *preference ordering*; (3) assessing the extent to which it is believed such properties apply to such entities; and (4) imagining new entities and asserting knowledge regarding them.

III. LINGUISTIC ARCHITECTURE

Now that we have described DIARC's mnemonic architecture, we are ready to describe the linguistic architecture that leverages it (see also Fig. 2). We will begin by discussing natural language understanding components (reference resolution and pragmatic understanding), and then discuss natural language generation components (pragmatic generation and referring expression generation).

Throughout this section, we will use the example utterance "I need the medkit that is on the shelf in the breakroom", and a

mnemonic architecture using three DKHBs (*Mapping, Vision, Social*) and three associated consultants (*locs,objs,ppl*).

A. Reference Resolution

The first architectural component we will discuss is our *Reference Resolution* Component, whose job is to ascertain the identities of any entities referenced through natural language. For example, upon receiving the semantic representation for “I need the medkit that is on the shelf in the breakroom” ($Statement(speaker, self, need(self, X), \{on(X, Y), medkit(X), shelf(Y), breakroom(Z), in(Y, Z)\})$), it is up to the Reference Resolution Component to determine what entities should be associated with variables X, Y and Z . This problem can be broken down into three levels: closed-world reference resolution, open-world reference resolution, and anaphora resolution. In the following subsections we will discuss algorithms for solving each of these increasingly larger problems.

1) *Closed-World Reference Resolution*: Closed-World Reference Resolution is the basic problem of finding the optimal mapping from *references* to *known entities*. Under a simplifying assumption of inter-constraint independence, Closed-World Reference Resolution can be modeled using the following equation, where $\Lambda = \{\lambda_0, \dots, \lambda_n\}$ is a set of semantic constraints, $\Gamma = \{\Gamma_0, \dots, \Gamma_n\}$ is the set of possible bindings to the variables contained in those constraints, and $\Phi = \phi_0, \dots, \phi_n$ is a set of satisfaction variables for which each ϕ_i is True iff formula λ_i holds under a given binding:

$$\Gamma^* = \operatorname{argmax}_{\Gamma \in \bar{\Gamma}} \prod_{i=0}^{|\Lambda|} P(\phi_i \mid \Gamma, \lambda_i)$$

That is, the optimal set of bindings Γ^* is that which maximizes the joint probability of a set of satisfaction variables Φ being satisfied under that binding and the set of provided semantic constraints Λ (under a simplifying assumption of independence between constraints). This model is algorithmically realized using the *DIST-CoWER* algorithm [2], [7], [8], which searches through the space of possible variable-entity assignments, pruning branches whose incrementally computed probability falls below a given threshold. For the example sentence, if the robot knows of a medkit on a shelf in a breakroom, it might return a set of bindings such as $\{X \rightarrow objs_4, Y \rightarrow objs_6, Z \rightarrow locs_9\}$, along with an associated probability value.

2) *Open-World Reference Resolution*: *DIST-CoWER* facilitates reference resolution under uncertainty, but presumes that all entities that could be referenced are known a priori: an assumption which is unwarranted in realistic human-robot interaction scenarios. To continue the previous example, if a robot knows of a breakroom but does not know of a shelf in that breakroom, *DIST-CoWER* will fail to return any bindings for “the medkit that is on the shelf in the breakroom”. Ideally, however, the robot would be able to both resolve the portions of the utterance that it *does* know of a priori, and learn in one shot about other, previously unknown entities referenced in the utterance. We model this problem, which we call *Open-World Reference Resolution*, using the following equation, where Λ is

a set of constraints, Λ^V is an ordering of the variables involved in those constraints, Γ^{i*} is the optimal solution provided by *DIST-CoWER* given the constraints involving only the last $|\Lambda^V| - i$ variables in Λ^V , Γ^{i*P} is the probability associated with this solution, and *complete* is a function which creates new representation to associate with any variables appearing in Λ^V but not in optimal open-world solution Γ^{i*} :

$$\Gamma^* = \operatorname{complete} \left(\operatorname{argmax}_{\Gamma^{i*} \in \{\Gamma^0, \dots, \Gamma^{|\Lambda^V|*}\}} \begin{cases} i, & \text{if } \Gamma^{i*P} > \tau \\ 0, & \text{otherwise} \end{cases} \right).$$

That is, the optimal set of bindings Γ^* is the first sufficiently probable set of bindings returned from a series of calls to *DIST-CoWER* made using different subsets of the full predicate set Λ . This model is algorithmically realized using the *DIST-POWER* algorithm [2], [7], [8], which (1) finds a variable ordering over the variables used in predicate list Λ based on linguistic factors such as prepositional attachment; (2) successively removing variables from this list until calling *DIST-CoWER* on only the predicates involving the remaining variables returns a sufficiently probable solution; (3) for each variable removed in this way, instructing the appropriate consultant to create a mental representation for a new entity; (4) instructing the appropriate Consultants to make any representations created in this way to be consistent with all related predicates in Λ ; (5) returning a unified set of bindings from the set of variables used in Λ to known and/or newly created entity representations.

For example, if in the example sentence the robot does not know of a shelf in a breakroom, it might return the set of bindings $\{X \rightarrow objs_{44}, Y \rightarrow objs_{45}, Z \rightarrow locs_9\}$, where $objs_{44}$ and $objs_{45}$ are references to newly created representations for hypothesized entities, and $locs_9$ is a reference to a previously existing representation for a (grounded or hypothetical) entity.

3) *Anaphora Resolution*: The algorithms discussed in the previous sections allow our robots to resolve references in uncertain and open worlds using *definite noun phrases* (i.e., “the- N ” phrases). Humans, however, have a tendency to use a much wider variety of referring forms, (e.g., “a- N ”, “this”, “it”). In order to handle this multitude of forms, we embed *DIST-POWER* within a larger algorithm called *GH-POWER* [9] for its use of the *Givenness Hierarchy* theory of reference. The *Givenness Hierarchy* divides referential forms into six groups, each of which is associated with a different tier of a hierarchy of six nested cognitive statuses. For example, when one uses “it”, the *Givenness Hierarchy* suggests that the speaker believes their target referent to be at least *in focus* for the listener; when one uses “this”, the *Givenness Hierarchy* suggests the speaker believes their target referent to be at least *activated*, that is, in *short-term memory* for the listener (and possibly also *in focus* because of the nested nature of the six tiers); and so forth. Using this theory, when a particular referential form is used, we infer the set of possible statuses the speaker may believe their target referent to have, and from this infer a sequence of *mnemonic actions* to take: creating a new representation of a referent or searching for an existing one in a particular data structure. For example, when a definite noun phrase is used, *GH-POWER* will search through the set




Fig. 1. **Memory Model:** The Focus of Attention, Short Term Memory, Discourse Context are hierarchically nested, and contain references to entities stored in the Distributed, Heterogeneous Knowledge Bases which comprise long-term memory (*Mapping, Vision, Social*), access to which is enabled and controlled using a set of Consultants (*locs, objs, ppl*).

of activated entities, the focus of attention, the set of familiar entities, and, if none of these steps are successful, will use *DIST-POWER* to search through all of the robot’s long term memory, and hypothesizing a new representation if that search fails. Fig. 1 shows the set of Givenness Hierarchy-theoretic data structures we use (Focus of Attention, Short Term Memory, Discourse Context), their hierarchical relationship with each other, and how the representations within each other can be viewed as references to entities stored in the DHKBs which comprise long-term memory (*Mapping, Vision, Social*), access to which is enabled and controlled using a set of Consultants (*locs, objs, ppl*).

For example, in the example sentence, all three entities are described using “the”; as such, *GH-POWER* will use the sequence of mnemonic actions Search STM, Search FoA, Search DC, Search LTM, Hypothesize, where these last two steps are achieved (if necessary) by performing a *DIST-POWER* query. If sufficiently probable candidate referents can be found in the upper level GH-theoretic data structures, this may not be necessary.

4) *Discussion:* To summarize thus far, given the unbound semantic interpretation of an incoming utterance, our referential components will first identify sequences of mnemonic actions to take to resolve the references found in that utterance; those sequences will then be used to initiate searches through various data structures for entities that, according to the appropriate Consultants, are likely to satisfy the predicates that comprise the semantic interpretation; in the worst case this will require the use of *DIST-POWER* to search all of long-term memory. Once the reference resolution process is

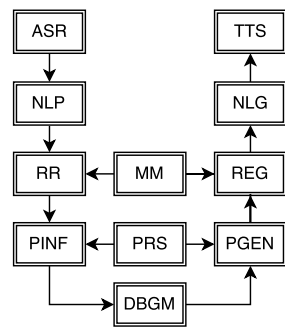


Fig. 2. **Architectural Diagram:** Information flows through the following components: Automatic Speech Recognition (ASR), Natural Language Processing (NLP), Reference Resolution (RR), Pragmatic Understanding (PUND), the Dialogue Belief and Goal Manager (DBGM), Pragmatic Generation (PGEN), Referring Expression Generation (REG), Natural Language Generation (NLG), Text to Speech (TTS). Utilized by these components are the consultant framework and associated Memory Model (MM) and the Pragmatic Rule Set (PRS)

completed, the end result is a set of hypotheses, each of which associate the variables found in the incoming semantic interpretation with a different set of entities, and each of which has a particular likelihood. As described in [10], these hypotheses are then used to create a set of *bound utterances*, which are combined into a single Dempster-Shafer theoretic *body of evidence* which is passed to our *pragmatic reasoning* component.

Our approach significantly differs from most other recent language understanding work in robotics. Most other work has focused on tackling the full *language grounding* problem of mapping from natural language to continuous perceptual representations [11]–[20]. In contrast, we separate this problem into two halves: *reference resolution*, in which natural language is mapped to discrete symbols representing unique entities, and *symbol grounding*, in which those symbols are associated with continuous perceptual representations, and focus on the development of reference resolution algorithms. This division has allowed us to develop a framework for resolving references both to grounded, observed entities, as well as to heretofore unknown or even hypothetical and imaginary entities, thus providing us the means to tackle *open-world* reference resolution.

B. Pragmatic Reasoning

In this section we will discuss the Pragmatic Reasoning components of our architecture: Pragmatic Understanding and its counterpart, Pragmatic Generation.

1) *Pragmatic Understanding:* The pragmatic reasoning component’s primary purpose is to infer the intentions behind incoming utterances. Specifically, this component attempts to infer the intentions behind conventionally indirect utterance forms, also known as *Indirect Speech Acts*, which recent work has shown to be prevalent throughout human-robot dialogue [21]. Provided with the Dempster-Shafer theoretic set of candidate utterances Θ_u produced by *GH-POWER*, a Dempster-Shafer theoretic set of contextual information Θ_c and a set of Dempster-Shafer theoretic rules of the form $u \wedge c \Rightarrow i$, the pragmatic reasoning component produces a Dempster-Shafer theoretic set of intentions Θ_i by computing $m_i(\cdot) = ((m_u \otimes m_c) \odot m_{uc \rightarrow i})(\cdot)$, where \otimes is a Dempster-Shafer theoretic *And* operator, and \odot is a Dempster-Shafer theoretic *Modus Ponens* operator [22].

For example, when processing the example sentence, a robot might have a pragmatic rule that says (with some probability between, say, 0.7 and 0.8) that when someone says to someone they believe to be their subordinate that they need something, they likely want their subordinate to bring them that something:

$$\text{Stmt}(A, B, \text{need}(A, C)) \wedge \text{bel}(A, \text{subordinate}(B, A)) \\ \xrightarrow{[0.7, 0.8]} \text{want}(A, \text{goal}(B, \text{bring}(B, C, A)))$$

Because the example utterance matches this rule’s utterance form, the uncertainty interval reflecting the confidence that that utterance was what was actually said is combined with the uncertainty interval reflecting the robot’s confidence that the speaker believes the robot to be their subordinate, as well as with the uncertainty interval associated with the rule itself,

producing an uncertainty interval reflecting how confident the robot is that the speaker wants it to bring them the described object.

Each candidate intention I inferred in this way is thus augmented with a Dempster-Shafer theoretic interval $[\alpha, \beta]$ within which the probability that I is true can be said to lie. If an interval is determined to reflect sufficient uncertainty by Nunez' uncertainty measure [23]

$$\lambda = 1 + \frac{\beta}{1 + \beta - \alpha} \log_2 \frac{\beta}{1 + \beta - \alpha} + \frac{1 - \alpha}{1 + \beta - \alpha} \log_2 \frac{1 - \alpha}{1 + \beta - \alpha},$$

the robot generates it's own intention – an intention to know whether or not it should actually infer that candidate intention. This allows the robot to identify sources of both pragmatic and referential uncertainty and ignorance. If such an intention is generated, it will be satisfied by generating a *clarification request* [10]. This highlights the other capability of the pragmatic reasoning component: to perform *pragmatic generation*.

2) *Pragmatic Generation*: Due to our use of a Dempster-Shafer-theoretic approach, the same rules used to infer the intentions behind utterances (pragmatic understanding) can also be used to abduce utterances that can be used to communicate intentions (pragmatic generation). Pragmatic generation is performed using the same set of Dempster-Shafer theoretic rules and logical operators used during pragmatic understanding [22], with the addition of one pre-processing step and one post-processing step.

Before pragmatic generation is performed, it is determined whether the robot is generating a clarification request, and if so, whether there are more than two choices must be arbitrated between [10]. If so (e.g., if, in the example, the robot determines it knows of two medkits on a shelf in a breakroom), those choices are unified into a single predicate which will allow a generic WH-question (e.g., “Which medkit would you like?”) rather than a many-item YN-question (e.g., “Would you like the red medkit or the blue medkit or the white medkit or the green medkit?”), which are only when there are two or fewer options. Note, however, that at this stage of processing, only the bare utterance form has been generated (e.g., QuestionYN(self,speaker, or(would(speaker,like(speaker,obj₁)), would(speaker,like(speaker,obj₂))))), and not the properties used to describe each referenced entity.

After pragmatic generation has yielded a set of utterance forms which could be communicated, each is passed *forwards* through the pragmatic reasoning module, in order to simulate the utterance understanding process. This creates a set of intentions the robot may believe its interlocutor will infer if it chooses to use that particular utterance form. This allows the robot to detect unintended side-effects of different candidate utterance forms so that it can choose the best possible utterance to communicate its intentions. The best candidate utterance form is then selected and sent to our *Referring Expression Generation* module.

3) *Discussion*: Our work on pragmatic understanding directly builds off of previous work from Briggs et al. [24]. Like that work, our own understands utterances based on its beliefs *about the speaker's beliefs*; but we improve on that

work by handling uncertainty (and ignorance) and allowing for adaptation: capabilities also largely lacking from previous computational approaches to ISA understanding (e.g., [25]–[27]).

Our techniques for generating clarification requests compare favorably to previous work due to our accounting for human preferences (cf. [28], [29]) and our ability to handle uncertainty (cf. [12]). Similarly, while there has been some previous work on generating indirect language (e.g., [24], [30]), we believe that our work is the first to enable robots' generation of conventionalized indirect speech acts under uncertainty.

C. Referring Expression Generation

Once a robot has chosen an utterance form to communicate, it must decide what properties to use to describe the things it wishes to communicate about. This is a problem known as Referring Expression Generation. To solve this problem, we use *DIST-PIA* [31], a version of the classic *Incremental Algorithm* [32], modified to use our consultant framework. When crafting a referring expression for a given entity, our algorithm proceeds through the ordered list of properties provided by the consultant responsible for that entity: each property is added to the list of properties to be used in the description if it is sufficiently probable that it applies to the target referent, and if it is not sufficiently probable that it applies to one or more *distractors*, thus allowing those distractors to be ruled out. This algorithm improves on the classic Incremental Algorithm in its use of our Mnemonic Architecture (to enable use in integrated robot architectures) and in its ability to operate under uncertainty.

For example, suppose the *objs* consultant advertises that it can handle the following properties: $\{shelf(X - objs), medkit(X - objs), green(X - objs), red(X - objs), on(X - objs, Y - objs), in(X - objs, Y - locs)\}$, and that the *DIST-PIA* algorithm is used to generate a description of $objs_4$. Suppose $objs_4$ is not likely to be a shelf – that property will be ignored. Suppose $objs_4$ is likely to be a medkit, as are two other objects – $medkit(objs_4)$ will be added to the set of properties to use. Suppose $objs_4$ is not likely to be green – that property will be ignored. Suppose $objs_4$ is likely to be red, and that neither of the two *distractor* medkits are likely to be so – $red(objs_4)$ will be added to the set of properties to be used, and since all distractors have been eliminated, the set $\{medkit(objs_4), red(objs_4)\}$ will be returned, allowing a description such as “the red medkit” to be generated.

Of equal contribution in the work that presented this algorithm [31] is our development of a novel evaluation framework that solicits certainty estimates from humans in order to craft probability distributions that can then be used by uncertainty-handling REG algorithms. This allows such algorithms to be evaluated with respect to both other algorithms and human beings, without committing to any particular set of visual classifiers (cf. [33]–[35]).

IV. CONCLUSION

In this article, we have described a natural language understanding and generation pipeline designed specifically for use

in integrated robot architectures and for operation in uncertain and open worlds. It is our hope that the general frameworks presented in this work will allow researchers to more easily integrate together disparate approaches – and that this work will draw researchers’ attention to the under-studied area of open-world language processing. For a complete treatment of the work described in this paper, we direct the interested reader to the authors’ recent dissertation, which describes the work cited herein in much more detail [2].

ACKNOWLEDGMENT

This work was funded in part by ONR grants #N00014-11-1-0289, #N00014-11-1-0493, #N00014-10-1-0140 #N00014-14-1-0144, #N00014-14-1-0149, #N00014-14-1-0751, and NSF grants #1111323, #1038257.

REFERENCES

- [1] N. Mavridis, “A review of verbal and non-verbal human–robot interactive communication,” *Robotics and Autonomous Systems*, vol. 63, pp. 22–35, 2015.
- [2] T. Williams, “Situated natural language interaction in uncertain and open worlds,” Ph.D. dissertation, Tufts University, 2017.
- [3] P. W. Schermerhorn, J. F. Kramer, C. Middendorff, and M. Scheutz, “DIARC: A testbed for natural human-robot interaction,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2006, pp. 1972–1973.
- [4] M. Scheutz, G. Briggs, R. Cantrell, E. Krause, T. Williams, and R. Veale, “Novel mechanisms for natural human-robot interactions in the diarc architecture,” in *Proceedings of AAAI Workshop on Intelligent Robotic Systems*, 2013.
- [5] M. Scheutz, E. Krause, B. Oosterveld, T. Frasca, and R. Platt, “Spoken instruction-based one-shot object and action learning in a cognitive robotic architecture,” in *Proceedings of the Sixteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2017.
- [6] M. Quigley, J. Faust, T. Foote, and J. Leibs, “ROS: an open-source robot operating system,” in *ICRA Workshop on Open Source Software*, 2009.
- [7] T. Williams and M. Scheutz, “A framework for resolving open-world referential expressions in distributed heterogeneous knowledge bases,” in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [8] —, “POWER: A domain-independent algorithm for probabilistic, open-world entity resolution,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015.
- [9] T. Williams, S. Acharya, S. Schreitter, and M. Scheutz, “Situated open world reference resolution for human-robot dialogue,” in *Proceedings of the Eleventh ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2016.
- [10] T. Williams and M. Scheutz, “Resolution of referential ambiguity in human-robot dialogue using dempster-shafer theoretic pragmatics,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2017.
- [11] P. Gorniak and D. Roy, “Grounded semantic composition for visual scenes,” *Journal of Artificial Intelligence Research*, vol. 21, pp. 429–470, 2004.
- [12] G.-J. M. Kruijff, P. Lison, T. Benjamin, H. Jacobsson, and N. Hawes, “Incremental, multi-level processing for comprehending situated dialogue in human-robot interaction,” in *Symposium on Language and Robots*, 2007.
- [13] S. Lemaignan, R. Ros, R. Alami, and M. Beetz, “What are you talking about? grounding dialogue in a perspective-aware robotic architecture,” in *The Eighteenth IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2011, pp. 107–112.
- [14] F. Meyer, “Grounding words to objects: A joint model for co-reference and entity resolution using markov logic for robot instruction processing,” Ph.D. dissertation, TUHH.
- [15] J. Y. Chai, L. She, R. Fang, S. Ottarson, C. Littley, C. Liu, and K. Hanson, “Collaborative effort towards common ground in situated human-robot dialogue,” in *Proceedings of the ACM/IEEE International Conference on Human-robot Interaction (HRI)*, 2014, pp. 33–40.
- [16] J. Fasola and M. J. Matrić, “Interpreting instruction sequences in spatial language discourse with pragmatics towards natural human-robot interaction,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 2720–2727.
- [17] C. Kennington and D. Schlangen, “A simple generative model of incremental reference resolution for situated dialogue,” *Computer Speech & Language*, vol. 41, pp. 43–67, 2017.
- [18] S. Tellex, T. Kollar, S. Dickerson, M. R. Walter, A. G. Banerjee, S. Teller, and N. Roy, “Approaching the symbol grounding problem with probabilistic graphical models,” *AI Magazine*, 2011.
- [19] I. Chung, O. Propp, M. R. Walter, and T. M. Howard, “On the performance of hierarchical distributed correspondence graphs for efficient symbol grounding of robot instructions,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 5247–5252.
- [20] C. Matuszek, E. Herbst, L. Zettlemoyer, and D. Fox, “Learning to parse natural language commands to a robot control system,” in *Proceedings of the Thirteenth International Symposium on Experimental Robotics (ISER)*, 2012.
- [21] G. Briggs, T. Williams, and M. Scheutz, “Enabling robots to understand indirect speech acts in task-based interactions,” *Journal of Human-Robot Interaction*, 2017.
- [22] T. Williams, G. Briggs, B. Oosterveld, and M. Scheutz, “Going beyond command-based instructions: Extending robotic natural language interaction capabilities,” in *Proceedings of 29th AAAI Conference on Artificial Intelligence*, 2015.
- [23] R. C. Núñez, R. Dabarera, M. Scheutz, G. Briggs, O. Bueno, K. Premaratne, and M. N. Murthi, “DS-Based Uncertain Implication Rules for Inference and Fusion Applications,” in *Sixteenth International Conference on Information Fusion*, July 2013.
- [24] G. Briggs and M. Scheutz, “A hybrid architectural approach to understanding and appropriately generating indirect speech acts,” in *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- [25] D. J. Litman and J. F. Allen, “A plan recognition model for subdialogues in conversations,” *Cognitive science*, vol. 11, no. 2, pp. 163–200, 1987.
- [26] E. A. Hinkelman and J. F. Allen, “Two constraints on speech act ambiguity,” in *Proceedings of the Twenty-Seventh annual meeting of the Association for Computational Linguistics (ACL)*, 1989, pp. 212–219.
- [27] S. Wilske and G.-J. Kruijff, “Service robots dealing with indirect speech acts,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2006, pp. 4698–4703.
- [28] R. Deits, S. Tellex, T. Kollar, and N. Roy, “Clarifying commands with information-theoretic human-robot dialog,” *Journal of Human-Robot Interaction*, 2013.
- [29] S. Hemachandra, M. R. Walter, S. Tellex, and S. Teller, “Learning spatial-semantic representations from natural language descriptions and scene classifications,” in *International Conference on Robotics and Automation (ICRA)*, 2014.
- [30] R. A. Knepper, C. I. Mavrogiannis, J. Proft, and C. Liang, “Implicit communication in a joint action,” in *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. ACM, 2017, pp. 283–292.
- [31] T. Williams and M. Scheutz, “Referring expression generation under uncertainty: Algorithm and evaluation framework,” in *Proceedings of the Tenth International Conference on Natural Language Generation (INLG)*, 2017.
- [32] R. Dale and E. Reiter, “Computational interpretations of the gricean maxims in the generation of referring expressions,” *Cognitive science*, vol. 19, no. 2, pp. 233–263, 1995.
- [33] H. Horacek, “Generating referential descriptions under conditions of uncertainty,” in *Proceedings of the Tenth European Workshop on Natural Language Generation (ENLG)*, 2005, pp. 58–67.
- [34] A. Sadovnik, A. Gallagher, and T. Chen, “Not everybody’s special: Using neighbors in referring expressions with uncertain attributes,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2013, pp. 269–276.
- [35] R. Fang, C. Liu, L. She, and J. Y. Chai, “Towards situated dialogue: Revisiting referring expression generation,” in *Proceedings of the Conference on Empirical Methods for Natural Language Processing (EMNLP)*, 2013, pp. 392–402.

Process Mining for Trauma Resuscitation

Sen Yang, Jingyuan Li, Xiaoyi Tang, Shuhong Chen, Ivan Marsic, Randall S. Burd

Abstract—We present our process mining system for analyzing the trauma resuscitation process to improve medical team performance and patient outcomes. Our system has four main parts: trauma resuscitation process model discovery, process model enhancement (or repair), process deviation analysis, and process recommendation. We developed novel algorithms to address the technical challenges for each problem. We validated our system on real-world trauma resuscitation data from the Children’s National Medical Center (CNMC), a level 1 trauma center. Our results show our system’s capability of supporting complex medical processes. Our approaches were also implemented in an interactive visual analytic tool.

Index Terms—Process Mining, Trauma Resuscitation, Medical Process Diagnosis

I. INTRODUCTION

TRAUMA is the leading cause of death and acquired disability among children and young adults. Because early trauma evaluation and management strongly impact the injury’s outcome, it is critical that severely injured patients receive efficient and error-free treatment in the first several hours of injury. During the trauma resuscitation, multidisciplinary teams rapidly identify and treat potential life-threatening injuries, then develop and execute a short-term management plan. The Advanced Trauma Life Support (ATLS) [1] protocol has been widely adopted as the initial evaluation and management strategy for injured patients worldwide. Although its implementation has been associated with improved outcomes, the application of this protocol has been shown to vary considerably, even with experienced teams. Many deviations from the ATLS protocol, e.g. the omission or delaying of steps, may have minimal impact on the outcome, but have been shown to increase the likelihood of a major uncorrected error that may lead to an adverse outcome.

The objective of this project is to develop a computerized decision support system that can automatically identify deviations during trauma resuscitations and provide real-time alerts of risk conditions to the medical team. We are approaching this goal using four process mining techniques [2] (Figure 1). (1) Process model discovery, extracting workflow models from data. (2) Process model enhancement, repairing the workflow model to mitigate the divergence between data-driven workflow models and expert hand-made models. (3) Process deviation analysis, discovering and analyzing the medical team er-

rors. (4) Process recommendation, building a recommender system that can give treatment suggestions to medical teams.

Key challenges for this study include limited data size, permissible deviations, variable patients, and concurrent teamwork. (1) Limited data: the trauma resuscitation workflow data needs to be coded manually in a labor-intensive way. To our best knowledge, there is no reliable system that can automatically capture trauma resuscitation workflow data. (2) Permissible deviations: it is necessary to distinguish the acceptable deviations (false alarms) from unexpected deviations (true alarms). (3) Variable patients: patients come to the trauma bay with different injuries that need different treatment plans. (4) Concurrent teamwork: trauma resuscitations need concurrent collaborative work within a medical team comprised of examining providers (surveying physicians), bedside nurses (left nurse, right nurse and charge nurse), and other team roles (e.g., surgical coordinator, anesthetist).

In this study, we contributed a comprehensive process mining framework with related techniques. We showed what problem each process mining technique solved and how these techniques worked together to achieve our project goals. We evaluated our process mining framework and techniques on a complex real-world medical case, the trauma resuscitation process. Our results showed the effectiveness of our techniques over existing process mining methods. Note that this paper is an overview of our previous work. The technical techniques can be found in our previous work.

II. RELATED WORK

Many medical processes are performed based on domain knowledge and standard protocols. For trauma resuscitations, the ATLS [1] protocol suggests a medical examination flow based on treatment priorities: Airway → Breathing → Circulation → Neurological Disability. Despite the use of standardized evaluation and management protocols, deviations are observed in up to 85% of trauma resuscitations [3]. Although most deviations are variations that result from the flexibility or adaptability needed for managing patients with different injuries, other deviations represent “errors” that can contribute to significant adverse patient outcomes, including death [4][5].

To discover and analyze the process errors, previous research used two different approaches: data-driven and expert-model-based. Data-driven methods rely on process models or patterns discovered from historic data, while expert-model methods rely on process models or rules designed by domain experts. Data-driven methods work by comparing individual process enactments to the discovered average process representation, such as the average process trace [6],

Sen Yang, Jingyuan Li, Xiaoyi Tang, Shuhong Chen and Ivan Marsic are with the Department of Electrical and Computer Engineering at Rutgers University, NJ, USA. Emails: sy358@rutgers.edu, jl2056@rutgers.edu, xt59@rutgers.edu, sc1624@rutgers.edu, marsic@jove.rutgers.edu

Randall S. Burd is with Children’s National Medical Center, Washington DC, USA. Email: rburd@childrensnational.org

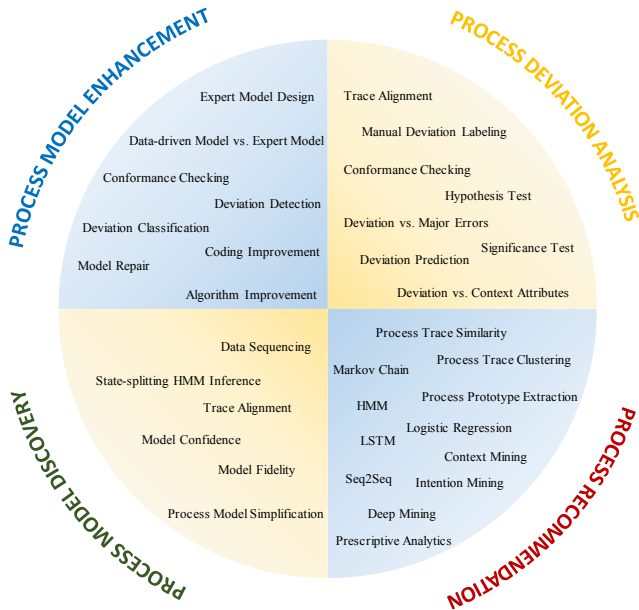


Figure 1. Process mining framework with related techniques we used in trauma resuscitation process analysis.

the data-driven model [7], or frequently occurring patterns [8]. Expert-model-based approaches locate the deviations by checking the conformance between particular enactments to the expert model [9], or constraints or rules specified by medical experts [10]. In our study, we combined both methods. First, we discovered the data-driven model and designed the expert model based on medical domain knowledge. Then, we compared the data-driven and expert-model to uncover the discrepancies between practice and expectation. The discrepancies were evaluated by medical experts to determine if model enhancement is needed. Lastly, we used the enhanced expert model to discover and analyze more process deviations.

To reduce process errors, Clarke et al. [11] and Fitzgerald et al. [5] developed computer-aided decision support systems that recommend treatment steps. These systems, however, rely on hand-made rules specified by medical experts, lack generalizability, and are subject to human bias. We developed an automatic, data-driven, label-free framework for process analysis and recommendation.

III. DATA DESCRIPTION AND DEFINITIONS

Ninety-five resuscitations were coded by medical experts from video recordings collected at trauma bay of CNMC between August 2014 and October 2016. Collection and use of the data was approved by the Institutional Review Board.

The coded **trauma resuscitation cases** $\mathbf{c} = [c^{(1)}, \dots, c^{(l)}]^T$ is a vector of elements $c^{(i)}$. Each $c^{(i)} = \{id^{(i)}, \mathbf{x}^{(i)}, \mathbf{T}^{(i)}\}$ denotes a **resuscitation case** (TABLE II), which is indexed with a unique case id, contains the resuscitation trace $\mathbf{T}^{(i)}$, and has a vector $\mathbf{x}^{(i)}$ of context attributes. A **resuscitation trace** $\mathbf{T}^{(i)} = [a_1^{(i)}, \dots, a_k^{(i)}]^T$ includes k activities that are ordered based on activity occurrence time. **Context attributes**

TABLE II

SAMPLE TRAUMA RESUSCITATION DATA

Case ID	Activity	Start Time	End Time	Case ID	xxx1	xxx2
xx1	Pt arrival	0:00:00	0:00:01	Age category	24-96	24-96
xx1	Visual assessment-AA	0:00:45	0:00:52	Sex	Male	Female
xx1	Chest Auscultation-BA	0:00:55	0:00:58	Night Shift	0	1
xx1	R DP/PT-PC	0:01:04	0:01:05	Weekend	0	0
xx1	Total Verbalized-GCS	0:01:29	0:01:30	Pre-arrival/Notification	1	0
xx1	Total Verbalized-GCS	0:01:50	0:01:51	Trauma Activation Level	Transfer	Attending
xx1	Right pupil-PU	0:02:12	0:02:18	Intubation	0	0
xx1	Left pupil-PU	0:02:19	0:02:24	Glasgow Coma Score >13	1	0
xx1	Right pupil-PU	0:02:24	0:02:25	Injury Type	Blunt	Penetrating
xx1	Visual inspection-H	0:02:33	0:02:34	Injury Severity Score	5	12
xx1	Palpation-H	0:02:33	0:02:37	Neck Injury Severity Score	3	5

(a) Trauma resuscitation trace

Properties	Stats
Num. Cases (or Patients)	95
Num. Total Activities	10851
Num. Activity Types	132
Num. External Attributes	17
Time Period	2014.08 – 2016.12
Size of Medical Team	[7, 12]
Longest Trace (Num. Acts.)	196
Shortest Trace (Num. Acts.)	60
Avg. Num. Acts. in Traces	114.2

(c) Data statistics

(b) Context attributes

ID	Ext. Attributes	Resus. Traces
$id^{(1)}$	$[x_1^{(1)}, \dots, x_g^{(1)}]$	$[a_1^{(1)}, \dots, a_k^{(1)}]$
$id^{(2)}$	$[x_1^{(2)}, \dots, x_g^{(2)}]$	$[a_1^{(2)}, \dots, a_k^{(2)}]$
\vdots	\vdots	\vdots
$id^{(n)}$	$[x_1^{(n)}, \dots, x_g^{(n)}]$	$[a_1^{(n)}, \dots, a_k^{(n)}]$

(d) Data formalization

$\mathbf{x}^{(i)} = [x_1^{(i)}, \dots, x_g^{(i)}]^T$ is a vector of g recorded patient attributes (e.g., patient age, injury type) and hospital factors (e.g., day vs. night shift, prehospital triage of injury severity).

TABLE I

COMPARISON OF AGSS WITH EXISTING STATE-SPLITTING ALGORITHMS ON PRIMARY SURVEY AND SECONDARY SURVEY IN TRAUMA RESUSCITATION PROCESSES. MODEL FIDELITY (M_f) AND MODEL CONFIDENCE (M_c) ARE SCALED BY THE NUMBER OF PROCESS TRACES IN EACH PROCESS.

	Primary Survey		Secondary Survey	
	M_f	M_c	M_f	M_c
Markov Chain	-10.00	-10.38	-47.59	-44.09
AGSS	-9.98	-10.16	-45.22	-44.13
ML-SSS (0.01)	-9.05	-10.66	-48.79	-60.52
Heuristic	-8.50	-10.95	-47.59	-44.10
MDL	-12.37	-12.97	-64.44	-65.37
STACT	-9.32	-10.19	-49.22	-58.08

IV. PROCESS MINING METHODS AND RESULTS

Process model discovery, process model enhancement, and process deviation analysis are descriptive analytics, aiming to extract insights and knowledge from historic data. Process recommendation is predictive or prescriptive analytics, aiming to determine the best treatment procedure given observed context attributes (e.g., patient demographics). For each subsection, we discuss the methods first, followed by our achieved results.

A. Process Discovery of Resuscitation Workflow

Existing workflow discovery algorithms (e.g., heuristic miner, genetic miner, alpha miner [2]) have problems in handling duplicate activities. For example, during trauma resuscitation, the medical team checks patient's eyes twice for different reasons. First, during the primary survey, they assess the patient's pupillary response for neurological disability. Second, during the secondary survey, they examine patient's eyes in more detail, looking for injuries to the cornea, sclera and eyelids. An accurate workflow model should be able to distinguish the first eye check from the second one, despite their identical labels. Existing workflow mining algorithms

assume that duplicate activities in a process trace are equivalent. We solved this problem using Hidden Markov Models (HMM) to represent the workflow. To induce an HMM, we proposed a novel inference algorithm guided by trace alignment (AGSS) [12]. AGSS (Figure 2 and Problem 1) first initializes a general Markov chain λ_0 . After initialization, AGSS determines two factors: which states to split and when to stop splitting. AGSS determines the splitting candidates from the alignment matrix and orders them by column frequency. After calculating the splitting candidates, AGSS performs iterative splitting.

Problem 1: AGSS Medical Process Model Discovery:

Given: A set of resuscitation traces $T = [T^{(1)}, \dots, T^{(n)}]$.

Objective: Successively splitting state $s_j \in S$ to find an HMM topology λ_s that maximizes probability $P(T|\lambda)$:

$$\lambda_s(T) = \operatorname{argmax}_{\lambda} P(T|\lambda) = \operatorname{argmax}_{\lambda} \prod_i P(T^{(i)}|\lambda) \quad (1)$$


where $P(T^{(i)}|\lambda)$ is observation sequence probability, solved with the Forward algorithm [13]. S , the set of states to be split, is calculated using the trace alignment algorithm [14].

We tested AGSS on trauma resuscitation workflow data and compared it to existing state-splitting HMM inference algorithms (e.g., ML-SSS, STACT) [12]. The performance was measured based on (1) model's quality and (2) the algorithm's computational efficiency. The model quality is measured by *model fidelity* and *model confidence*. Model fidelity (or accuracy) measures the agreement between a given workflow model and the observed process traces (i.e., the log likelihood of generating the observed traces using the given model). Model confidence measures how well a workflow model represents the underlying process that generates the observed process traces. High model confidence means the model describes not only the observed traces, but also the unobserved realizations of the underlying process. Our results show that AGSS is not only more computationally efficient (by a factor of $\mathcal{O}(n)$, where n is the number of hidden states), but also produces HMMs of higher model fidelity (M_f) and model confidence (M_c) (TABLE I). Our results also show the workflow model discovered by AGSS is more readable and more representative than models discovered by existing HMM inference and process mining algorithms [12].

B. Process Model Enhancement of Resuscitation Workflow

As required for process deviation analysis, an initial expert model was developed by medical experts. The initial model underwent several revisions until the domain experts reached a consensus about which medical activities to include and in which order. We used two different approaches to enhance the expert model. First, we compared expert model to the data-driven model discovered using our AGSS algorithm. The model induced by AGSS helped discover three discrepancies between the initial expert model and actual practice. These discrepancies were analyzed for repairing the expert model

Trace Alignment



State-Splitting




Figure 2. Our Alignment-guided State-splitting algorithm for discovering a more representative data-driven workflow model [12]. (a) Trace alignment algorithm to find splitting candidates. (b)(c)(d) State-splitting HMM inference.

[12]. Afterwards, we performed a more detailed and comprehensive model diagnosis using conformance checking [15] with twenty-four trauma resuscitation cases (Problem 2). We discovered more discrepancies between the model and practice, and came to the conclusion that our initial model could not fully represent the resuscitation process. In our preliminary analysis, we identified 57.3% (630 out of 1099) activities as deviations based on the initial expert model, of which only 24.6% (155 out of 630) were true deviations (i.e., process errors), while the remaining 75.4% deviations were false alarms due to process model incompleteness, coding errors, or inadequate algorithms. We then applied different strategies to address the false alarms, e.g., repairing the expert model, improving the coding procedure, and updating the algorithm. We tested the repaired model on ten unseen resuscitations, achieving 93.1% deviation discovery accuracy.

Problem 2: Process Enhancement:

Given: Historic resuscitation traces $T = [T^{(1)}, \dots, T^{(n)}]$ and a hand-made expert model λ_e .

Objective #1: Discover process deviations ϵ from T :

$$\epsilon = \mathcal{C}(T, \lambda_e) \quad (2)$$

where $\mathcal{C}(T, \lambda_e)$ is the conformance checking algorithm [15], taking process traces and expert model as input and outputting the process deviations.

Objective #2: Classify ϵ as true alarms (i.e., process errors) and acceptable variations (i.e., false alarms ϵ'). Repair the expert model λ_e to eliminate ϵ' .

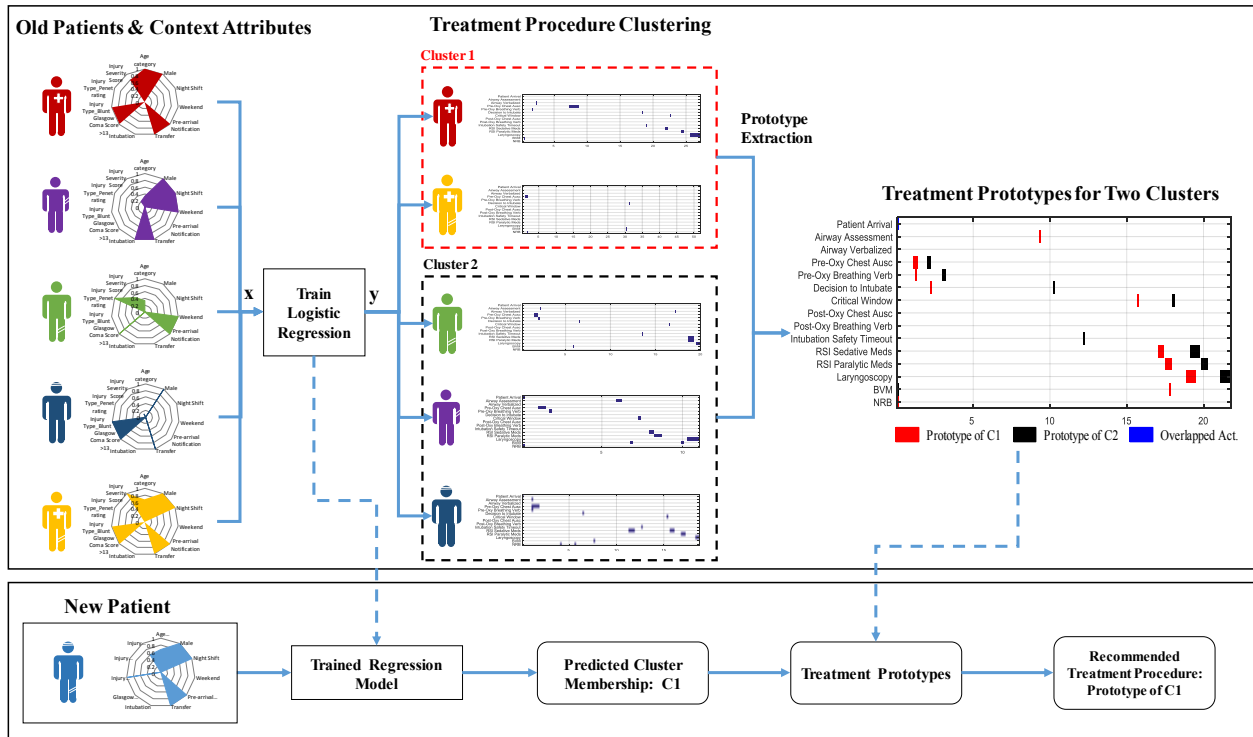


Figure 3. Medical process recommender framework [16]. Treatment procedures were clustered based on similarity (horizontal axis of the matrices represents time and vertical axis represents activity type). A logistic regression model is trained between context attributes and cluster membership. Treatment prototypes were calculated from each cluster. When a new patient comes to the trauma bay (bottom box), trained regression model takes context attributes as input and outputs the cluster membership (e.g., C1). A treatment prototype is recommended based on the cluster membership (e.g., Prototype of C1).

C. Process Deviation Analysis of Resuscitation Workflow

The main goal of process deviation analysis is to test our hypotheses, proving the adverse effects of accumulated deviations on trauma patient outcomes and team's ability to compensate for major errors. The goal can be achieved in two steps. First, identify the process deviations (manually or based on conformance checking). Second, contact statistical analysis to test the association between the process deviations and patient outcomes (or patient attributes).

Our previous deviation analyses were performed both manually by medical experts and automatically based on the conformance checking algorithms. For manual deviation analysis, we analyzed thirty-nine resuscitations and discovered the number of errors and the number of high-risk errors per resuscitation increased with the number of non-error process deviations per resuscitation (correlation-coefficient = 0.42, $p = 0.01$ and correlation-coefficient = 0.62, $p < 0.001$, respectively) [18]. For automated deviation analysis, we performed conformance checking on the enhanced process model using ninety-five trauma resuscitations. We detected 1,059 process deviations in 5,659 activities of 42 commonly performed assessment activity types (11.1 deviations per case). Our results also show that the resuscitations of patients with no pre-arrival notification ($p = 0.037$) and blunt injury (blunt vs. non-blunt $p = 0.013$) were significantly correlated with more deviations.

D. Process Recommendation of Resuscitation Workflow

To help reduce medical team errors, we developed a process

recommender system [16][17] (Figure 3 and Problem 3) that provides data-driven step-by-step treatment recommendations. Our system was built based on the associations between similar historic process performances and contextual information (e.g., patient demographics). We introduced a novel similarity metric (named TwS-TP [16]) that incorporates temporal information about activity performance, and handles concurrent activities. Our recommender system selects the appropriate prototype performance of the process based on user-provided context attributes. Our approach for determining the prototypes discovers the commonly performed activities and their temporal relationships.

Problem 3: Process Recommendation:

Given: A new patient with context attributes \mathbf{x}' .

Objective: Recommend a treatment plan $\mathbf{T}_{x'} = [a_1, \dots, a_k]^T$ to the medical team that maximizes:

$$\mathbf{T}_{x'} = \underset{T}{\operatorname{argmax}} S(\mathbf{T}, \mathbf{T}_g) \quad (3)$$

where \mathbf{T}_g is the ground-truth treatment procedure and $S(\mathbf{T}_{x'}, \mathbf{T}_g)$ is the similarity measure of two process traces [16].

We implemented our recommender system with several different similarity metrics: edit distance (ED), sequential pattern based (SP) and TwS-TP, and different clustering algorithms: hierarchical clustering (HC), density peak clustering (DPC) and affinity propagation clustering (APC). We clustered the process traces using different combinations of similarity

TABLE III

RECOMMENDATION EVALUATION ON TRAUMA RESUSCITATION PROCESS DATASET. THE FORMAT α (τ) REPRESENTS THE REGRESSION MODEL RESULT A AND THE BASELINE (ZERO R) RESULT (τ). REC NC STANDS FOR RECOMMENDED NUMBER OF CLUSTERS.

Trauma Resuscitation Data		
Rec NC	ED(2), SP (2), Time-warping (2)	
Metrics	F-Score	G-means
ED + HC	0.634 (0.654)	0.448 (0.428)
ED + DPC	0.692 (0.686)	0.436 (0.413)
ED + APC	0.346 (0.353)	0.392 (0.500)
SP + HC	0.637 (0.533)	0.603 (0.471)
SP + DPC	0.637 (0.533)	0.603 (0.471)
SP + APC	0.645 (0.519)	0.591 (0.475)
TwS-PT + HC	0.526 (0.392)	0.520 (0.497)
TwS-PT + DPC	0.713 (0.670)	0.556 (0.421)
TwS-PT + APC	0.767 (0.366)	0.683 (0.499)

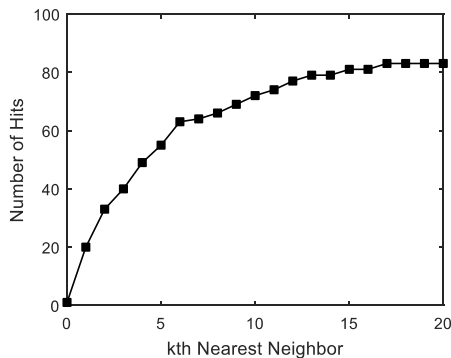


Figure 4. Number of hits of recommended process enactment within k nearest neighbors of the actual enactment.

and clustering algorithms. We used ZeroR as the baseline and selected F-score and G-means as the recommendation accuracy metrics because they are suitable for evaluating multi-class imbalance learning problem. Our results show our system on 87 trauma resuscitation cases achieved recommendation accuracy of up to 0.77 F1 score (TwS-PT + APC in TABLE III) (compared to 0.37 F1 score using ZeroR). In addition, in 55 of 87 cases (63.2%), our recommended prototype was among the 5 nearest neighbors of the actual historic procedures in a set of eighty-seven cases (Figure 4).

V. CONCLUSION

We presented the framework and associated techniques used in our trauma resuscitation process analysis and diagnosis. Our framework includes four parts: process model discovery, process model enhancement, process deviation analysis, and process recommendation. For each part, we showed the core problems, methods, and results we achieved so far. This paper shows a substantial implementation of process mining techniques on a significant real-world problem.

ACKNOWLEDGMENT

This paper is based on research supported by National Institutes of Health under grant number 1R01LM011834-01A1.

REFERENCES

- [1] American College of Surgeons. Committee on Trauma. Advanced trauma life support: ATLS: student course manual. American College of Surgeons, 2012.
- [2] Van Der Aalst, Wil. Process mining: discovery, conformance and enhancement of business processes. Springer Science & Business Media, 2011.
- [3] Carter, Elizabeth A., et al. "Adherence to ATLS primary and secondary surveys during pediatric trauma resuscitation." *Resuscitation* 84.1 (2013): 66-71.
- [4] Clarke, John R., Spejewski, B., Gertner, A.S., Webber, B.L., Hayward, C.Z., Santora, T.A., et al. "An objective analysis of process errors in trauma resuscitations." *Academic Emergency Medicine*, 7.11 (2000): 1303-1310.
- [5] Fitzgerald, Mark, et al. "Trauma resuscitation errors and computer-assisted decision support." *Archives of Surgery* 146.2 (2011): 218-225.
- [6] Bouarfa, Loubna, and Jenny Dankelman. "Workflow mining and outlier detection from clinical activity logs." *Journal of Biomedical Informatics* 45.6 (2012): 1185-1190.
- [7] Caron, Filip, et al. "Monitoring care processes in the gynecologic oncology department." *Computers in biology and medicine* 44 (2014): 88-96.
- [8] Christov, Stefan C., George S. Avrunin, and Lori A. Clarke. "Online Deviation Detection for Medical Processes." *AMIA Annual Symposium Proceedings*. American Medical Informatics Association, 2014.
- [9] Kelleher, Deirdre C., et al. "Effect of a checklist on advanced trauma life support workflow deviations during trauma resuscitations without pre-arrival notification." *Journal of the American College of Surgeons* 218.3 (2014): 459-466.
- [10] Rovani, Marcella, et al. "Declarative process mining in healthcare." *Expert Systems with Applications* 42.23 (2015): 9236-9251.
- [11] Clarke, John R., et al. "Computer-generated trauma management plans: comparison with actual care." *World journal of surgery* 26.5 (2002): 536-538.
- [12] Yang, Sen, et al. "Medical Workflow Modeling Using Alignment-Guided State-Splitting HMM." 2016 IEEE International Conference on Healthcare Informatics (ICHI), 2016.
- [13] Rabiner LR. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE*, vol.77, no.2, pp. 257-286, February 1989.
- [14] Yang, Sen, et al. "Duration-Aware Alignment of Process Traces." *Industrial Conference on Data Mining*. Springer, 2016.
- [15] Adriansyah, Arya, Natalia Sidorova, and Boudewijn F. van Dongen. "Cost-based fitness in conformance checking." the 11th International Conference on Application of Concurrency to System Design, 2011.
- [16] Yang, Sen, et al. "A Data-driven Process Recommender Framework." *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2017.
- [17] Yang, Sen, et al. "VIT-PLA: Visual Interactive Tool for Process Log Analysis." *KDD 2016 Workshop on Interactive Data Exploration and Analytics*.
- [18] Webman, Rachel, et al. "Classification and team response to non-routine events occurring during pediatric trauma resuscitation." *The journal of trauma and acute care surgery* 81.4 (2016): 666.

ROBOTICS has the potential to be one of the most revolutionary technologies in human history. The impact of cheap and potentially limitless manpower could have a profound influence on our everyday life and overall onto our society. As envisioned by Iain M. Banks, Asimov and many other science fictions writers, the effects of robotics on our society might lead to the disappearance of physical labor and a generalized increase of the quality of life. However, the large-scale deployment of robots in our society is still far from reality, except perhaps in a few niche markets such as manufacturing. One reason for this limited deployment of robots is that, despite the tremendous advances in the capabilities of the robotic hardware, a similar advance on the control software is still lacking. The use of robots in our everyday life is still hindered by the necessary complexity to manually design and tune the controllers used to execute tasks. As a result, the deployment of robots often requires lengthy and extensive validations based on human expert knowledge, which limit their adaptation capabilities and their widespread diffusion. In the future, in order to truly achieve an ubiquitous robotization of our society, it is necessary to reduce the complexity of deploying new robots in new environments and tasks.

The goal of this dissertation is to provide automatic tools based on Machine Learning techniques to simplify and streamline the design of controllers for new tasks. In particular, we here argue that Bayesian modeling is an important tool for automatically learning models from raw data and properly capture the uncertainty of the such models. Automatically learning models however requires the definition of appropriate features used as input for the model. Hence, we present an approach that extend traditional Gaussian process models by jointly learning an appropriate feature representation and the subsequent model. By doing so, we can strongly guide the features representation to be useful for the subsequent prediction task.

A first robotics application where the use of Bayesian modeling is beneficial is the accurate learning of complex dynamics models. For highly non-linear robotic systems, such as in presence of contacts, the use of analytical system identification techniques can be challenging and time-consuming, or even intractable. We introduce a new approach for learning inverse dynamics models exploiting artificial tactile sensors. This approach allows to recognize and compensate for the presence of unknown contacts, without requiring a spatial calibration of the tactile sensors. We demonstrate on the humanoid robot iCub that our approach outperforms state-of-the-art analytical models, and when employed in control tasks significantly improves the tracking accuracy.

A second robotics application of Bayesian modeling is automatic black-box optimization of the parameters of a controller. When the dynamics of a system cannot be modeled (either out of complexity or due to the lack of a full state representation), it is still possible to solve a task by adapting an existing controller. The approach used in this thesis is Bayesian optimization, which allows to automatically optimize the parameters of the controller for a specific task. We evaluate and compare the performance of Bayesian optimization on a

gait optimization task on the dynamic bipedal walker Fox. Our experiments highlight the benefit of this approach by reducing the parameters tuning time from weeks to a single day.

In many robotic application, it is however not possible to always define a single straightforward desired objective. More often, multiple conflicting objectives are desirable at the same time, and thus the designer needs to take a decision about the desired trade-off between such objectives (e.g., velocity vs. energy consumption). One framework that is useful to assist in this decision making is the multi-objective optimization framework, and in particular the definition of Pareto optimality. We propose a novel framework that leverages the use of Bayesian modeling to improve the quality of traditional multi-objective optimization approaches, even in low-data regimes. By removing the misleading effects of stochastic noise, the designer is presented with an accurate and continuous Pareto front from which to choose the desired trade-off. Additionally, our framework allows the seamless introduction of multiple robustness metrics which can be considered during the design phase. These contributions allow an unprecedented support to the design process of complex robotic systems in presence of multiple objective, and in particular with regards to robustness.

The overall work in this thesis successfully demonstrates on real robots that the complexity of deploying robots to solve new tasks can be greatly reduced trough automatic learning techniques. We believe this is a first step towards a future where robots can be used outside of closely supervised environments, and where a newly deployed robot could quickly and automatically adapt to accomplish the desired tasks (http://tuprints.ulb.tu-darmstadt.de/5878/13/thesis_roberto_calandra.pdf).

APPLICATION OF AGENT TECHNOLOGY FOR FAULT DIAGNOSIS OF TELECOMMUNICATION NETWORKS

Alvaro Carrera
a.carrera@upm.es

Universidad Politecnica de Madrid, Spain

THIS PhD thesis contributes to the problem of autonomous fault diagnosis of telecommunication networks. Nowadays, in telecommunication networks, operators perform manual diagnosis tasks. Those operations must be carried out by high skilled network engineers which have increasing difficulties to properly manage the growing of those networks, both in size, complexity and heterogeneity. Moreover, the advent of the Future Internet makes the demand of solutions which simplifies and automates the telecommunication network management has been increased in recent years. To collect the domain knowledge required to developed the proposed solutions and to simplify its adoption by the operators, an agile testing methodology is defined for multi-agent systems. This methodology is focused on the communication gap between the different work groups involved in any software development project, stakeholders and developers. To contribute to overcoming the problem of autonomous fault diagnosis, an agent architecture for fault diagnosis of telecommunication networks is defined. That architecture extends the Belief-Desire-Intention (BDI) agent model with different diagnostic

models which handle the different subtasks of the process. The proposed architecture combines different reasoning techniques to achieve its objective using a structural model of the network, which uses ontology-based reasoning, and a causal model, which uses Bayesian reasoning to properly handle the uncertainty of the diagnosis process. To ensure the suitability of the proposed architecture in complex and heterogeneous environments, an argumentation framework is defined. This framework allows agents to perform fault diagnosis in federated domains. To apply this framework in a multi-agent system, a coordination protocol is defined. This protocol is used by agents to dialogue until a reliable conclusion for a specific diagnosis case is reached. Future work comprises the further extension of the agent architecture to approach other managements problems, such as self-discovery or self-optimisation; the application of reputation techniques in the argumentation framework to improve the extensibility of the diagnostic system in federated domains; and the application of the proposed agent architecture in emergent networking architectures, such as SDN, which offers new capabilities of control for the network (<http://oa.upm.es/39170/>).

MULTI-AGENT SYSTEM FOR COORDINATION OF DISTRIBUTED ENERGY RESOURCES IN VIRTUAL POWER PLANTS

Anders Clausen
ancla@mmmi.sdu.dk

Centre for Smart Energy Solution, University of Southern Denmark, Denmark

THE electricity grid is facing challenges as a result of an increase in the share of renewable energy in electricity production. In this context, Demand Response (DR) is considered an inexpensive and CO_2 -friendly approach to handle the resulting fluctuations in the electricity production. The concept of DR refers to changes in consumption patterns of Distributed Energy Resources (DER), in response to incentive payments or changes in the price of electricity over time. Existing DR programs have capacity requirements, which individual DER entities are often unable to meet. As a consequence, the concept of a Virtual Power Plant (VPP) has emerged. A VPP aggregates multiple DER, and exposes them as a single, controllable entity.

The contribution of this thesis is a general method for integrating DER in VPPs. The approach constitutes a meta-model for VPPs, which describes DER and VPPs as entities. Each entity is constituted by a group of software agents. The meta-model describes the interaction between groups and contains two negotiation models: the intradomain- and the interdomain negotiation models. The intradomain negotiation model describes agent decision logic and communication between agents in a group. The model contains a mediator-based negotiation protocol, where agents negotiate over a set of issues, allowing each entity to pursue several objectives and decide upon several issues.

The interdomain negotiation model describes negotiation between groups of agents. In practice this means that the interdomain negotiation model ties instances of intradomain

negotiation together. A key aspect of the interdomain negotiation model is that it ensures group autonomy.

Both negotiation models have been implemented in Controleum, a framework for multi-objective optimization. In this context, Controleum has been refactored to allow for the abstractions of the negotiation models to be implemented in the framework. Furthermore, a Domain Specific Language (DSL) has been developed and implemented to allow for easy configuration of negotiations.

Experiments with simulations of different VPP scenarios have been conducted. These experiments indicate that the proposed approach is capable of integrating complex and heterogeneous DER in VPPs, while preserving the autonomous nature of the DER. Experiments are also conducted on instances of the 0/1 Knapsack problem. These experiments serve to illustrate the general applicability of the proposed solution. Future experiments will test the solution in real scenarios (<http://aclausen.dk/documents/thesis.pdf>).

EFFICIENT KNOWLEDGE MANAGEMENT FOR NAMED ENTITIES FROM TEXT

Sourav Dutta
sourav.dutta@nokia.com
 Saarland University, Germany

THE evolution of search from keywords to entities has necessitated the efficient harvesting and management of entity-centric information for constructing knowledge bases catering to various applications such as semantic search, question answering, and information retrieval. The vast amounts of natural language texts available across diverse domains on the Web provide rich sources for discovering facts about named entities such as people, places, and organizations.

A key challenge, in this regard, entails the need for precise identification and disambiguation of entities across documents for extraction of attributes/relations and their proper representation in knowledge bases. Additionally, the applicability of such repositories not only involves the quality and accuracy of the stored information, but also storage management and query processing efficiency. This dissertation aims to tackle the above problems by presenting efficient approaches for entity-centric knowledge acquisition from texts and its representation in knowledge repositories.

This dissertation presents a robust approach for identifying text phrases pertaining to the same named entity across huge corpora, and their disambiguation to canonical entities present in a knowledge base, by using enriched semantic contexts and link validation encapsulated in a hierarchical clustering framework. This work further presents language and consistency features for classification models to compute the credibility of obtained textual facts, ensuring quality of the extracted information. Finally, an encoding algorithm, using frequent term detection and improved data locality, to represent entities for enhanced knowledge base storage and query performance is presented (<http://scidok.sulb.uni-saarland.de/volltexte/2017/6792/>).

BAYESIAN MODELS OF CATEGORY ACQUISITION AND MEANING DEVELOPMENT

Lea Frermann
l.frermann@sms.ed.ac.uk
University of Edinburgh, UK.

THE ability to organize concepts (e.g., dog, chair) into efficient mental representations, i.e., categories (e.g., animal, furniture) is a fundamental mechanism which allows humans to perceive, organize, and adapt to their world. Much research has been dedicated to the questions of how categories emerge and how they are represented. Experimental evidence suggests that (i) concepts and categories are represented through sets of features (e.g., dogs bark, chairs are made of wood) which are structured into different types (e.g., behavior, material); (ii) categories and their featural representations are learnt jointly and incrementally; and (iii) categories are dynamic and their representations adapt to changing environments.

This thesis investigates the mechanisms underlying the incremental and dynamic formation of categories and their featural representations through cognitively motivated Bayesian computational models. Models of category acquisition have been extensively studied in cognitive science and primarily tested on perceptual abstractions or artificial stimuli. In this thesis, we focus on categories acquired from natural language stimuli, using nouns as a stand-in for their reference concepts, and their linguistic contexts as a representation of the concepts features. The use of text corpora allows us to (i) develop large-scale unsupervised models thus simulating human learning, and (ii) model child category acquisition, leveraging the linguistic input available to children in the form of transcribed child-directed language.

In the first part of this thesis we investigate the incremental process of category acquisition. We present a Bayesian model and an incremental learning algorithm which sequentially integrates newly observed data. We evaluate our model output against gold standard categories (elicited experimentally from human participants), and show that high-quality categories are learnt both from child-directed data and from large, thematically unrestricted text corpora. We find that the model performs well even under constrained memory resources, resembling human cognitive limitations. While lists of representative features for categories emerge from this model, they are neither structured nor jointly optimized with the categories.

We address these shortcomings in the second part of the thesis, and present a Bayesian model which jointly learns categories and structured featural representations. We present both batch and incremental learning algorithms, and demonstrate the models effectiveness on both encyclopedic and child-directed data. We show that high-quality categories and features emerge in the joint learning process, and that the structured features are intuitively interpretable through human plausibility judgment evaluation.

In the third part of the thesis we turn to the dynamic nature of meaning: categories and their featural representations change over time, e.g., children distinguish some types of features (such as size and shade) less clearly than adults, and word meanings adapt to our ever changing environment and its

structure. We present a dynamic Bayesian model of meaning change, which infers time-specific concept representations as a set of feature types and their prevalence, and captures their development as a smooth process. We analyze the development of concept representations in their complexity over time from child-directed data, and show that our model captures established patterns of child concept learning. We also apply our model to diachronic change of word meaning, modeling how word senses change internally and in prevalence over centuries.

The contributions of this thesis are threefold. Firstly, we show that a variety of experimental results on the acquisition and representation of categories can be captured with computational models within the framework of Bayesian modeling. Secondly, we show that natural language text is an appropriate source of information for modeling categorization-related phenomena suggesting that the environmental structure that drives category formation is encoded in this data. Thirdly, we show that the experimental findings hold on a larger scale. Our models are trained and tested on a larger set of concepts and categories than is common in behavioral experiments and the categories and featural representations they can learn from linguistic text are in principle unrestricted (http://frermann.de/dataFiles/phd_thesis_leafermann.pdf).

LATENT FACTOR MODELS FOR COLLABORATIVE FILTERING

Anupriya Gogna
anupriyag@iiitd.ac.in
Indraprastha Institute of Information Technology - Delhi,
India

THE enormous growth in online availability of information content has made Recommender Systems (RS) an integral part of most online portals and e-commerce sites. Most websites and service portals, be it movie rental services, online shopping or travel package providers, offer some form of recommendations to users. These recommendations provide the users more clarity, that too expeditiously and accurately in limiting (shortlisting) the items/information they need to search through, thereby improving the customer's experience. The direct link between customer's satisfaction and revenue of e-commerce sites induce widespread interest of both, academia and industry, in the design of efficient recommender systems.

The current de-facto approach for RS design is Collaborative Filtering (CF). CF techniques use the ratings provided by users, to a subset of the items in the repository, to make future recommendations. However, the rating information is hard to acquire; often a user has rated less than 5% of the items. Thus, the biggest challenge in recommender system design is to infer users preference from this extremely limited predilection information. The lack of adequate (explicit) preference information has motivated several works to augment the rating data with auxiliary information such as users demographics, trust networks, and item tags. Further, the scale of the problem, i.e. the amount of the data to be processed (selecting few items out of hundreds and thousands of items for an equally large number of users) adds another dimension to the concerns

surrounding the design of a good RS. There have been several developments in the field of RS design over the past decades. However, the difficulty in achieving the desired accuracy and effectiveness in recommendations leaves considerable scope for improvement.

In this work, we model effective recommendation strategies, using optimization centric frameworks, by exploiting reliable and readily available information, to address several pertinent issues concerning RS design. Our proposed recommendation strategies are built on the principals of latent factor models (LFM). LFM are constructed on the belief that a users choice for an item is governed by a handful of factors C the latent factors. For example, in the case of movies, these factors may be genre, director, language while for hotels it can be price and location.

Our first contribution targets improvement in prediction accuracy as well the speed of processing by suggesting modifications to the standard LFM frameworks. We develop a more intuitive model, supported by effective algorithm design, which better captures the underlying structure of the rating database while ensuring a reduction in run time compared to standard CF techniques. In the next step, we build upon these proposed frameworks to address the problem of lack of collaborative data, especially for cold start (new) users and items, by making use of readily available user and item metadata - item category and user demographics. Our suggested frameworks make use of available metadata to add additional constraints in the standard models; thereby presenting a comprehensive strategy to improve prediction accuracy in both warm (existing users/items for which rating data is available) and cold start scenario.

Although, high recommendation accuracy is the hallmark of a good RS, over-emphasis on accuracy compromises on variety and leads to monotony. Our next set of models aims to address this concern and promote diversity and novelty in recommendations. Most existing works, targeting diversity, build ad-hoc exploratory models relying heavily on heuristic formulations. In the proposed work, we modify the latent factor model to formulate a joint optimization strategy to establish accuracy-diversity balance; our models yield superior results than existing works.

The last contribution of this work is to explore the use of another representation learning tool for collaborative filtering C Autoencoder (AE). Conventional AE based designs, use only the rating information; lack of adequate data hampers the performance of these structures, thus, they do not perform as well as conventional LFM based designs. In this work, we propose a modification of the standard autoencoder C the Supervised Autoencoder C which can jointly accommodate information from multiple sources resulting in better performance than existing architectures (<https://repository.iiitd.edu.in/xmlui/handle/123456789/501>).

MACHINE LEARNING THROUGH EXPLORATION FOR PERCEPTION-DRIVEN ROBOTICS

Herke van Hoof
herke.vanhoof@mail.mcgill.ca

McGill University, Canada

THE ability of robots to perform tasks in human environments, such as our homes, has largely been limited to rather simple tasks, such as lawn mowing and vacuum cleaning. One reason for this limitation is that every home has a different layout with different objects and furniture. Thus, it is impossible for a human designer to anticipate all challenges a robot might face, and equip the robot a priori with all necessary perceptual and manipulation skills.

Instead, robots could use machine learning techniques to adapt to new environments. Many current learning techniques, however, rely on human supervisors to provide data in the form of annotations, demonstrations, and parameter settings. As such, making a robot perform a task in a novel environment can still require a significant time investment. In this thesis, instead, multiple techniques are studied to let robots collect their own training data through autonomous exploration.

The first study concerns an unsupervised robot that learns from sensory feedback obtained through interactive exploration of its environment. In a novel bottom-up, probabilistic approach, the robot tries to segment the objects in its environment through clustering with minimal prior knowledge. This clustering is based on cues elicited through the robots actions. Evaluations on a real robot system show that the proposed method handles noisy inputs better than previous methods. Furthermore, a proposed scheme for action selection criterion according to the expected information gain criterion is shown to increase the learning speed.

Often, however, the goal of a robot is not just to learn the structure of the environment, but to learn how to perform a task encoded by a reward signal. In a second study, a novel robot reinforcement learning algorithm is proposed that uses learned non-parametric models, value functions, and policies that can deal with high-dimensional sensory representations. To avoid that the robot converges prematurely to a sub-optimal solution, the information loss of policy updates is limited. The experiments show that the proposed algorithm performs well relative to prior methods. Furthermore, the method is validated on a real-robot setup with high-dimensional camera image inputs.

One problem with typical exploration strategies is that the behavior is perturbed independently in each time step. The resulting incoherent exploration behavior can result in inefficient random walk behavior and wear and tear on the robot. Perturbing policy parameters for an entire episode yields coherent exploration, but tends to increase the number of episodes needed. In a third study, a strategy is introduced that makes a balanced trade-off between these two approaches. The experiments show that such trade-offs are beneficial across different tasks and learning algorithms.

This thesis thus addresses how robots can learn autonomously by exploring the world. Throughout the thesis, new approaches and algorithms are introduced: a probabilistic interactive segmentation approach, the non-parametric relative entropy policy search algorithm, and a framework for generalized exploration. These approaches and algorithms contribute towards the capability of robots to autonomously

learn useful skills in human environments in a practical manner (<http://tuprints.ulb.tu-darmstadt.de/5749/>).

A SERVICE-ORIENTED FRAMEWORK FOR THE SPECIFICATION, DEPLOYMENT, EXECUTION, BENCHMARKING, AND PREDICTION OF PERFORMANCE OF SCALABLE PRIVACY-PRESERVING RECORD LINKAGE TECHNIQUES

Dimitrios Karapiperis
dkarapiperis@eap.gr
Hellenic Open University, Greece

AT the dawn of a new era of computing and the growth of big data, information integration is more important than ever before. Large organizations, such as corporations, health providers, public sector agencies, or research institutes, integrate their data in order to generate insightful data analytics. This data integration and analysis enables these organizations to make certain decisions toward deriving better business outcomes.

Record linkage, also known as entity resolution or data matching, is the process of resolving whether two records that belong to disparate data sets, refer to the same real-world entity. The lack of common identifiers, as well as typos and inconsistencies in the data, render the process of record linkage very challenging and mandatory for organizations which need to integrate their records. When data is deemed as private, then specialized techniques are employed that perform Privacy-Preserving Record Linkage (PPRL) in a secure manner, by respecting the privacy of the individuals who are represented by those records. For instance, in the public sector, there are databases which contain records that refer to the same citizen holding outdated information. Although, there is an urgent need of integration, the lack of common identifiers poses significant impediments in the linkage process.

Due to the large volumes of records contained in the data sets, core component of PPRL is the blocking phase, in which records are inserted into overlapping blocks and, then, are compared with one another. The purpose of the blocking phase is to formulate as many as possible matching record pairs. The blocking methods proposed thus far in the literature apply empirical techniques, which, given the particularities and technical characteristics of the data sets at hand, produce arbitrary results. This dissertation is the first to provide theoretical guarantees of completeness in the generated result set of the PPRL process, by introducing a randomized framework that allows for easy tuning of its configuration. Its flexibility lies in the fact that one can specify the level of its performance, with respect to the completeness of the results, taking into account multiple factors, such as: the urgency of the problem being solved, the desired response time, or the criticalness of the results completeness. Additionally, we enhance its main functionality, by providing certain extensions, and illustrate its applicability to both offline and online settings. The framework has been materialized by a prototype that is freely available so that it can be used by practitioners and researchers in their tasks.

This dissertation is divided into several chapters; we first introduce the core of our framework and its capabilities, and, then, we present its several extensions, such as the integration with the map/reduce paradigm for scaling up large volumes of records, or the add-on for performing PPRL using numeric values. In each of these chapters, we report on an extensive evaluation of the application of the constituent methods with real data sets, which illustrates that they outperform existing approaches (<https://drive.google.com/file/d/0B2tBkOmLy8WZc0Z4OU0zQ2dOMVE/view?usp=sharing>).

LEARNING WITH MULTIPLE VIEWS OF DATA: SOME THEORETICAL AND PRACTICAL CONTRIBUTIONS

Chao Lan
pete.chaolan@gmail.com
University of Kansas, USA.

IN machine learning applications, instances are often describable by multiple feature sets, each somewhat interpretable and sufficient for the learning task. In this case, each feature set is called a view, the instances are called multi-view data, and the study of learning with multi-view data is called multi-view learning. In this dissertation, we investigated several issues of multi-view learning in two settings: one is called statistical setting, where one aims to learn predictive models based on random multi-view sample; the other one is called matrix setting, where one aims to recover missing values in the feature matrices of multiple views. In statistical setting, we first theoretically investigated the possibility of training an accurate predictive model using as few unlabeled multi-view data as possible, and concluded such possibility by improving the state-of-the-art unlabeled sample complexity of semi-supervised multi-view learning by a logarithm factor. We then investigated the application of multi-view clustering methods in social circle detection on ego social networks. We finally proposed a simple multi-view multi-class learning algorithm that consistently outperforms the state-of-the-art algorithm. In matrix setting, we focused on the negative transfer problem in Collective Matrix Factorization (CMF), which is a popular method to recover missing values in multi-view feature matrices. We first theoretically characterized negative transfer in a CMF estimator, as the decrease of its ideal minimax learning rate by a root function. We then showed our presented ideal rate is tight (up to a constant factor), by employing the statistical PAC theory to derive a matching upper bound for it; our employment of the PAC theory improved the state-of-the-art one, by relaxing its strong i.i.d. assumption of matrix recovery errors. We finally proposed a simple variant of CMF that outperforms a current variant in small sample case. At the conceptual level, we have been bridging gaps between research in statistical setting and matrix setting. Specifically, we employed statistical PAC theory to analyze matrix recovery error in both active and passive learning settings; we employed statistical multi-view learning framework to develop a variant of CMF for matrix recovery; we employed statistical mini-max theory to analyze CMF performance. Finally, our research in multi-view learning has motivated two

other studies. One study challenged a common assumption in cheminformatics that unreported substance-compound binding profiles are all negative. The other study proposed a first active matrix recovery method with PAC guarantee. (<https://www.dropbox.com/s/1zf3qksqoc8oqnr/dissertation.pdf?dl=0>)

TRANSFORMATION-BASED COMMUNITY DETECTION FROM SOCIAL NETWORKS

Sungsu Lim

ssungssu@kaist.ac.kr

Korea Advanced Institute of Science and Technology, Korea

IN recent decades social network analysis has become one of the most attractive issues in data mining. In particular, community detection is a fundamental problem in social network analysis. Many theories, models, and methods have been developed for this purpose. However, owing to a wide variety of network structures, there remain challenges to determining community structures from social networks. Among them, we focus on solving four important problems for community detection with different underlying structures. These are identifying the community structure of a graph when it consists of (i) overlapping community structure, (ii) highly mixed community structure, (iii) complex sub-structure, and (iv) highly mixed overlapping community structure.

We address these problems by developing transformation techniques. Our key motivation is that the transformation of a given network can provide us an improved structure to identify the community structure of an original network, as a kernel trick does. We propose a transformation-based algorithm that converts an original graph to a transformed graph that reflects the structure of the original network and has a superior structure for each problem. We identify the community structure using the transformed graph and then the membership on the transformed graph is translated back to that of the original graph.

For the first problem, we present a notion of the linkspace transformation that enables us to combine the advantages of both the original graph and the line graph, thereby conveniently achieving overlapping communities. Based on this notion, we develop an overlapping community detection algorithm LinkSCAN* [1]. The experimental results demonstrate that the proposed algorithm outperforms existing algorithms, especially for networks with many highly overlapping nodes.

For the second problem, we propose a community detection algorithm BlackHole [2]. The proposed graph embedding model attempts to place the vertices of the same community at a single position, similar to how a black hole attracts everything nearby. Then, these positions are easily grouped by a conventional clustering algorithm. The proposed algorithm is proven to achieve extremely high performance regardless of the difficulty of community detection in terms of the mixing and the clustering coefficient.

For the third problem, we propose a motif-based embedding method for graph clustering by modeling higher-order relationships among vertices in a graph [3]. The proposed method considers motif-based attractive forces to enhance the

clustering tendency of points in the output space of graph embedding. We prove the relationship with graph clustering and verify the performance and applicability.

For the fourth problem, we develop an algorithm to address the highly mixed overlapping community detection problem. The transformation of the proposed algorithm LinkBlackHole consists of a sequence of two different transformations. By first applying the link-space transformation and then the BlackHole transformation, we can detect highly mixed link communities.

The strength of this dissertation is based on a wide variety of community detection problems from social networks. We believe that this work will enhance the quality of community discovery for social network analysis (http://library.kaist.ac.kr/thesis02/2016/2016D020115245_S1.pdf).

INCREMENTAL AND DEVELOPMENTAL PERSPECTIVES FOR GENERAL-PURPOSE LEARNING SYSTEMS

Fernando Martínez-Plumed

fmartinez@dsic.upv.es

Universitat Politècnica de Valencia, Spain

THE stupefying success of Artificial Intelligence (AI) for specific problems, from recommender systems to self-driving cars, has not yet been matched with a similar progress in general AI systems, coping with a variety of (different) problems. This dissertation deals with the long-standing problem of creating more general AI systems, through the analysis of their development and the evaluation of their cognitive abilities.

Firstly, this thesis contributes with a general-purpose declarative learning system gErl [2],[3] that meets several desirable characteristics in terms of expressiveness, comprehensibility and versatility. The system works with approaches that are inherently general: inductive programming and reinforcement learning. The system does not rely on a fixed library of learning operators, but can be endowed with new ones, so being able to operate in a wide variety of contexts. This flexibility, jointly with its declarative character, makes it possible to use the system as an instrument for better understanding the role (and difficulty) of the constructs that each task requires.

Secondly, the learning process is also overhauled with a new developmental and lifelong approach for knowledge acquisition, consolidation and forgetting, which is necessary when bounded resources (memory and time) are considered. In this sense we present a parametrisable (hierarchical) approach [4] for structuring knowledge (based on coverage) which is able to check whether the new learnt knowledge can be considered redundant, irrelevant or inconsistent with the old one, and whether it may be built upon previously acquired knowledge.

Thirdly, this thesis analyses whether the use of more ability-oriented evaluation techniques for AI (such as intelligence tests) is a much better alternative to most task-oriented evaluation approaches in AI. Accordingly, we make a review of what has been done when AI systems have been confronted against tasks taken from intelligence tests [5]. In this regard, we scrutinise what intelligence tests measure in machines, whether they are useful to evaluate AI systems, whether they

are really challenging problems, and whether they are useful to understand (human) intelligence. Our aim here is to contribute to a more widespread realisation that more general classes of problems are needed when constructing benchmarks for AI evaluation.

As a final contribution, we show that intelligence tests can also be useful to examine concept dependencies (mental operational constructs) in the cognitive development of artificial systems, therefore supporting the assumption that, even for fluid intelligence tests, the difficult items require a more advanced cognitive development than the simpler ones. In this sense, in [3] we show how several fluid intelligence test problems are addressed by our general-purpose learning system gErl, which, although it is not particularly designed on purpose to solve intelligence tests, is able to perform relatively well for this kind of tests. gErl makes it explicit how complex each pattern is and what operators are used for each problem, thus providing useful insight into the characteristics and usefulness of these tests when assessing the abilities and cognitive development of AI systems.

Summing up, this dissertation represents one step forward in the hard and long pursuit of making more general AI systems and fostering less customary (and challenging) ability-oriented evaluation approach.

NOVEL METHODS OF MEASURING THE SIMILARITY AND DISTANCE BETWEEN COMPLEX FUZZY SETS

Josie McCulloch
josie.mcculloch@nottingham.ac.uk
University of Nottingham, UK.

THIS thesis develops measures that enable comparisons of subjective information that is represented through fuzzy sets. Many applications rely on information that is subjective and imprecise due to varying contexts and so fuzzy sets were developed as a method of modelling uncertain data. However, making relative comparisons between data-driven fuzzy sets can be challenging. For example, when data sets are ambiguous or contradictory, then the fuzzy set models often become non-normal or non-convex, making them difficult to compare.

This thesis presents methods of comparing data that may be represented by such (complex) non-normal or non-convex fuzzy sets. The developed approaches for calculating relative comparisons also enable fusing methods of measuring similarity and distance between fuzzy sets. By using multiple methods, more meaningful comparisons of fuzzy sets are possible. Whereas if only a single type of measure is used, ambiguous results are more likely to occur.

This thesis provides a series of advances around the measuring of similarity and distance. Based on them, novel applications are possible, such as personalised and crowd-driven product recommendations. To demonstrate the value of the proposed methods, a recommendation system is developed that enables a person to describe their desired product in relation to one or more other known products. Relative comparisons are then used to find and recommend something that matches a person's subjective preferences.

Demonstrations illustrate that the proposed method is useful for comparing complex, nonnormal and non-convex fuzzy sets. In addition, the recommendation system is effective at using this approach to find products that match a given query (<http://eprints.nottingham.ac.uk/id/eprint/33401>).

MULTI-DOCUMENT SUMMARIZATION BASED ON PATTERNS WITH WILDCARDS AND PROBABILISTIC TOPIC MODELING

Jipeng Qiang
qjp2100@163.com
Yangzhou University, China

WITH the rapid development of information technology, a huge amount of electronic documents are available online, such as Web news, scientific literature, digital books, email, microblogging, and etc. How to effectively organize and manage such vast amount of text data, and make the system facilitate and show the information to users, have become challenges in the field of intelligent information processing. Therefore, now more than ever, users need access to robust text summarization systems, which can effectively condense information found in a large amount of documents into a short, readable synopsis, or summary. In recent years, with the rapid development of e-commerce and social networks, we can obtain a large amount of short texts, e.g., book reviews, movie reviews, online chatting, and product introductions. A short text probably contains a lot of useful information that can help to learn hidden topics among texts. Meanwhile, only very limited word co-occurrence information is available in short texts compared with long texts, so traditional multi-document summarization algorithms cannot work very well on these texts. Thus, how to generate a summary from multiple documents has important research and practical values.

In this thesis, we study multi-document summarization (MDS) on long texts and multi-document summarization on short texts, and propose several multi-document summarization algorithms based on patterns with wildcards and probability topic modeling. Our main contributions are as follows.

(1) A novel pattern-based model for generic multi-document summarization is proposed. There are two main categories of multi-document summarization: term-based and ontology-based methods. A term-based method cannot deal with the problems of polysemy and synonymy. An ontology-based approach addresses such problems by taking into account of the semantic information of document content, but the construction of ontology requires lots of manpower. To overcome these open problems, this paper presents a pattern-based model for generic multi-document summarization, which exploits closed patterns to extract the most salient sentences from a document collection and reduce redundancy in the summary. Our method calculates the weight of each sentence of a document collection by accumulating the weights of its covering closed patterns with respect to this sentence, and iteratively selects one sentence that owns the highest weight and less similarity to the previously selected sentences, until reaching the length limitation. Our method combines the advantages of the term-based and ontology-based models

while avoiding their weaknesses. Empirical studies on the benchmark DUC2004 datasets demonstrate that our pattern-based method significantly outperforms the state-of-the-art methods.

(2) A new MDS paradigm called user-aware multi-document summarization is proposed. The aim of MDS meets the demands of users, and the comments contain implicit information of their care. Therefore, the generated summaries from the reports for an event should be salient according to not only the reports but also the comments. Recently, Bayesian models have successfully been applied to multi-document summarization showing state-of-the-art results in summarization competitions. News articles are often long. Tweets and news comments can be short texts. In this thesis, the corpus includes both short texts and long texts, referred to as heterogeneous text. Long text topic modeling views texts as a mixture of probabilistic topics, and short text topic modeling adopts simple assumption that each text is sampled from only one latent topic. For heterogeneous texts, in this case neither method developed for only long texts nor methods for only short texts can generate satisfying results. In this thesis, we present an innovative method to discover latent topics from a heterogeneous corpus including both long and short texts. Then, we apply the learned topics to the generation of summarizations. Experiments on real-world datasets validate the effectiveness of the proposed model in comparison with other state-of-the-art models.

(3) A new short text topic model based on word embeddings is proposed. Existing methods such as probabilistic latent semantic analysis (PLSA) and latent Dirichlet allocation (LDA) cannot solve this problem very well since only very limited word co-occurrence information is available in short texts. Based on recent results in word embeddings that learn semantical representations for words from a large corpus, we introduce a novel method, Embedding-based Topic Modeling (ETM), to learn latent topics from short texts. ETM not only solves the problem of very limited word co-occurrence information by aggregating short texts into long pseudo-texts, but also utilizes a Markov Random Field regularized model that gives correlated words a better chance to be put into the same topic. Experiments on real-world datasets validate the effectiveness of our model comparing with the state-of-the-art models (<https://drive.google.com/file/d/0B3BkztuizBiyeTltMmFlci10dUU/view?usp=sharing>).

PROBABILISTIC MODELS FOR LEARNING FROM CROWDSOURCED DATA

Filipe Rodrigues
rodr@dtu.dk

Denmark Technical University, Denmark

THIS thesis leverages the general framework of probabilistic graphical models to develop probabilistic approaches for learning from crowdsourced data. This type of data is rapidly changing the way we approach many machine learning problems in different areas such as natural language processing, computer vision and music. By exploiting the wisdom of crowds, machine learning researchers and practitioners

are able to develop approaches to perform complex tasks in a much more scalable manner. For instance, crowdsourcing platforms like Amazon mechanical turk provide users with an inexpensive and accessible resource for labeling large datasets efficiently. However, the different biases and levels of expertise that are commonly found among different annotators in these platforms deem the development of targeted approaches necessary.

With the issue of annotator heterogeneity in mind, we start by introducing a class of latent expertise models which are able to discern reliable annotators from random ones without access to the ground truth, while jointly learning a logistic regression classifier or a conditional random field. Then, a generalization of Gaussian process classifiers to multiple-annotator settings is developed, which makes it possible to learn non-linear decision boundaries between classes and to develop an active learning methodology that is able to increase the efficiency of crowdsourcing while reducing its cost. Lastly, since the majority of the tasks for which crowdsourced data is commonly used involves complex high-dimensional data such as images or text, two supervised topic models are also proposed, one for classification and another for regression problems. Using real crowdsourced data from Mechanical Turk, we empirically demonstrate the superiority of the aforementioned models over state-of-the-art approaches in many different tasks such as classifying posts, news stories, images and music, or even predicting the sentiment of a text, the number of stars of a review or the rating of movie.

But the concept of crowdsourcing is not limited to dedicated platforms such as Mechanical Turk. For example, if we consider the social aspects of the modern Web, we begin to perceive the true ubiquitous nature of crowdsourcing. This opened up an exciting new world of possibilities in artificial intelligence. For instance, from the perspective of intelligent transportation systems, the information shared online by crowds provides the context that allows us to better understand how people move in urban environments. In the second part of this thesis, we explore the use of data generated by crowds as additional inputs in order to improve machine learning models. Namely, the problem of understanding public transport demand in the presence of special events such as concerts, sports games or festivals, is considered. First, a probabilistic model is developed for explaining non-habitual overcrowding using crowd-generated information mined from the Web. Then, a Bayesian additive model with Gaussian process components is proposed. Using real data from Singapore's transport system and crowd-generated data regarding special events, this model is empirically shown to be able to outperform state-of-the-art approaches for predicting public transport demand. Furthermore, due to its additive formulation, the proposed model is able to breakdown an observed time-series of transport demand into a routine component corresponding to commuting and the contributions of individual special events.

Overall, the models proposed in this thesis for learning from crowdsourced data are of wide applicability and can be of great value to a broad range of research communities (http://www.fprodriues.com/thesis_phd_fmpr.pdf).

EXPLORING MIXED REALITY IN DISTRIBUTED COLLABORATIVE LEARNING ENVIRONMENTS

Anasol Pena Rios
acpena@essex.ac.uk
University of Essex, UK.

SOCIETY is moving rapidly towards a world, where technology enables people to exist in a blend of physical and virtual realities. In education, this vision involves technologies ranging from smart classrooms to e-learning, creating greater opportunities for distance learners, bringing the potential to change the fundamental nature of universities. However, to date, most online educational platforms have focused on conveying information rather than supporting collaborative physical activities which are common in university science and engineering laboratories. Moreover, even when online laboratory support is considered, such systems tend to be confined to the use of simulations or pre-recorded videos. The lack of support for online collaborative physical laboratory activities, is a serious shortcoming for distance learners and a significant challenge to educators and researchers.

In working towards a solution to this challenge, this thesis presents an innovative mixed reality framework (computational model, conceptual architecture and proof-of-concept implementation) that enables geographically dispersed learners to perform co-creative teamwork using a computer-based prototype comprising hardware and software components.

Contributions from this work include a novel distributed computational model for synchronising physical objects and their 3D virtual representations, expanding the dual-reality paradigm from single linked pairs to complex groupings, addressing the challenge of interconnecting geographically dispersed environments; and the creation of a computational paradigm that blends a model of distributed learning objects with a constructionist pedagogical model, to produce a solution for distributed mixed reality laboratories.

By way of evidence to support the research findings, this thesis reports on evaluations performed with students from eight different universities in six countries, namely China, Malaysia, Mexico, UAE, USA and UK; providing an important insight to the role of social interactions in distance learning, and demonstrating that the inclusion of a physical component made a positive difference to students learning experience, supporting the use of mixed reality objects in educational activities (<http://repository.essex.ac.uk/16172/>).

DYNAMIC ADVERSARIAL MINING - EFFECTIVELY APPLYING MACHINE LEARNING IN ADVERSARIAL NON-STATIONARY ENVIRONMENTS

Tegiyot Singh Sethi
t0seth01@louisville.edu
University of Louisville, USA.

WHILE understanding of machine learning and data mining is still in its budding stages, the engineering applications of the same has found immense acceptance and success. Cybersecurity applications such as intrusion detection systems, spam filtering, and CAPTCHA authentication, have

all begun adopting machine learning as a viable technique to deal with large scale adversarial activity. However, the naive usage of machine learning in an adversarial setting is prone to reverse engineering and evasion attacks, as most of these techniques were designed primarily for a static setting. The security domain is a dynamic landscape, with an ongoing never ending arms race between the system designer and the attackers. Any solution designed for such a domain needs to take into account an active adversary and needs to evolve over time, in the face of emerging threats. We term this as the Dynamic Adversarial Mining problem, and the presented work provides the foundation for this new interdisciplinary area of research, at the crossroads of Machine Learning, Cybersecurity, and Streaming Data Mining.

We start with a white hat analysis of the vulnerabilities of classification systems to exploratory attack. The proposed Seed-Explore-Exploit framework provides characterization and modeling of attacks, ranging from simple random evasion attacks to sophisticated reverse engineering. It is observed that, even systems having prediction accuracy close to 100%, can be easily evaded with more than 90% precision. This evasion can be performed without any information about the underlying classifier, training dataset, or the domain of application.

Attacks on machine learning systems cause the data to exhibit non stationarity (i.e., the training and the testing data have different distributions). It is necessary to detect these changes in distribution, called concept drift, as they could cause the prediction performance of the model to degrade over time. However, the detection cannot overly rely on labeled data to compute performance explicitly and monitor a drop, as labeling is expensive and time consuming, and at times may not be a possibility altogether. As such, we propose the Margin Density Drift Detection (MD3) algorithm, which can reliably detect concept drift from unlabeled data only. MD3 provides high detection accuracy with a low false alarm rate, making it suitable for cybersecurity applications; where excessive false alarms are expensive and can lead to loss of trust in the warning system. Additionally, MD3 is designed as a classifier independent and streaming algorithm for usage in a variety of continuous never-ending learning systems.

We then propose a Dynamic Adversarial Mining based learning framework, for learning in non stationary and adversarial environments, which provides security by design. The proposed Predict-Detect classifier framework, aims to provide: robustness against attacks, ease of attack detection using unlabeled data, and swift recovery from attacks. Ideas of feature hiding and obfuscation of feature importance are proposed as strategies to enhance the learning frameworks security. Metrics for evaluating the dynamic security of a system and recover-ability after an attack are introduced to provide a practical way of measuring efficacy of dynamic security strategies. The framework is developed as a streaming data methodology, capable of continually functioning with limited supervision and effectively responding to adversarial dynamics.

The developed ideas, methodology, algorithms, and experimental analysis, aim to provide a foundation for future work

in the area of Dynamic Adversarial Mining, wherein a holistic approach to machine learning based security is motivated.

**LEVERAGING MULTIMODAL INFORMATION IN
SEMANTICS AND SENTICS ANALYSIS OF
USER-GENERATED CONTENT**

Rajiv Shah

rajivshah@smu.edu.sg

Singapore Management University, Singapore

THE amount of user-generated multimedia content (UGC) has increased rapidly in recent years due to the ubiquitous availability of smartphones, digital cameras, and affordable network infrastructures. To benefit people from an automatic semantics and sentics understanding of UGC, this thesis^{1,2} focuses on developing effective algorithms for several significant multimedia analytics problems. Sentics are common affective patterns associated with natural language concepts exploited for tasks such as emotion recognition from text/speech or sentiment analysis. Knowledge structures derived from UGC are beneficial in an efficient multimedia search, retrieval, and recommendation. However, real-world UGC is complex, and extracting the semantics and sentics from only multimedia content is very difficult because suitable concepts may be exhibited in different representations. Advancements in technology enable users to collect a significant amount of contextual information (e.g., spatial, temporal, and preferential information). Thus, it necessitates analyzing UGC from multiple modalities to facilitate problems such as multimedia summarization, tag ranking and recommendation, preference-aware multimedia recommendation, and multimedia-based e-learning.

We focus on the semantics and sentics understanding of UGC leveraging both content and contextual information. First, for a better semantics understanding of an event from a large collection of photos, we present the EventBuilder system [2]. It enables people to automatically generate a summary for the event in real-time by visualizing different social media such as Wikipedia and Flickr. In particular, we exploit Wikipedia as the event background knowledge to obtain more contextual information about the event. This information is very useful in effective event detection. Next, we solve an optimization problem to produce text summaries for the event. Subsequently, we present the EventSensor system [6] that aims to address sentics understanding and produces a multimedia summary for a given mood. It extracts concepts and mood tags from the visual content and textual metadata of photos and exploits them in a sentics-based multimedia summary. Moreover, we focus on computing tag relevance for UGIs. Specifically, we leverage personal and social contexts of UGIs and follow a neighbor voting scheme to predict and rank tags [1, 5]. Furthermore, we focus on semantics and sentics understanding from videos (<http://dl.acm.org/citation.cfm?id=2912032>).

**BIG DATA FOR SOCIAL SCIENCES: MEASURING
PATTERNS OF HUMAN BEHAVIOR THROUGH
LARGE-SCALE MOBILE PHONE DATA**

Pal Sundsoy

sundsoy@gmail.com
University of OSLO, Norway

ANALYSIS of large amounts of data, so called Big Data, is changing the way we think about science and society. One of the most promising rich Big Data sources is mobile phone data, which has the potential to deliver near real-time information of human behaviour on an individual and societal scale. Several challenges in society can be tackled in a more efficient way if such information is applied in a useful manner. Through seven publications this dissertation shows how anonymized mobile phone data can contribute to the social good and provide insights into human behaviour on a largescale.

The size of the datasets analysed ranges from 500 million to 300 billion phone records, covering millions of people. The key contributions are two-fold: Big Data for Social Good: Through prediction algorithms the results show how mobile phone data can be useful to predict important socio-economic indicators, such as income, illiteracy and poverty in developing countries. Such knowledge can be used to identify where vulnerable groups in society are, improve allocation of resources for poverty alleviation programs, reduce economic shocks, and is a critical component for monitoring poverty rates over time. Further, the dissertation demonstrates how mobile phone data can be used to better understand human behaviour during large shocks and disasters in society, exemplified by an analysis of data from the terror attack 22nd July 2011 in Norway and a natural disaster on the south-coast in Bangladesh. This work leads to an increased understanding of how information spreads, and how millions of people move around. The intention is to identify displaced people faster, cheaper and more accurately than existing survey-based methods.

Big Data for efficient marketing: Finally, the dissertation offers an insight into how anonymised mobile phone data can be used to map out large social networks, covering millions of people, to understand how products spread inside these networks. Results show that by including social patterns and machine learning techniques in a large-scale marketing experiment in Asia, the adoption rate is increased by 13 times compared to the approach used by experienced marketers. A data-driven and scientific approach to marketing, through more tailored campaigns, contributes to less irrelevant offers for the customers, and better cost efficiency for the companies (<https://www.duo.uio.no/handle/10852/55139>).

**TOWARD INTELLIGENT CYBER-PHYSICAL
SYSTEMS: ALGORITHMS, ARCHITECTURES,
AND APPLICATIONS**

Bo Tang

btang@ele.uri.edu

University of Rhode Island, USA.

CYBER-physical systems (CPS) are the new generation of engineered systems integrated with computation and physical processes. The integration of computation, communication and control adds new capabilities to the systems being able to interact with physical world. The uncertainty in

physical environment makes future CPS to be more reliant on machine learning algorithms which can learn and accumulate knowledge from historical data to support intelligent decision making. Such CPS with the incorporation of intelligence or smartness are termed as intelligent CPS which are safer, more reliable and more efficient.

This thesis studies fundamental machine learning algorithms in supervised and unsupervised manners and examines new computing architecture for the development of next generation CPS. Two important applications of CPS, including smart pipeline and smart grid, are also studied in this thesis. Particularly, regarding supervised machine learning algorithms, several generative learning and discriminative learning methods are proposed to improve learning performance. For the generative learning, we build novel classification methods based on exponentially embedded families (EEF), a new probability density function (PDF) estimation method, when some of the sufficient statistics are known. For the discriminative learning, we develop an extended nearest neighbor (ENN) method to predict patterns according to the maximum gain of intra-class coherence. The new method makes a prediction in a “two-way communication” style: it considers not only who are the nearest neighbors of the test sample, but also who consider the test sample as their nearest neighbors. By exploiting the generalized class-wise statistics from all training data, the proposed ENN is able to learn from the global distribution, therefore improving pattern recognition performance and providing a powerful technique for a wide range of data analysis applications. Based on the concept of ENN, an anomaly detection method is also developed in an unsupervised manner.

CPS usually have high-dimensional data, such as text, video, and other multi-modal sensor data. It is necessary to reduce feature dimensions to facilitate the learning. We propose an optimal feature selection framework which aims to select feature subsets with maximum discrimination capacity. To further address the information loss issue in feature reduction, we develop a novel learning method, termed generalized PDF projection theorem (GPPT), to reconstruct the distribution in high-dimensional raw data space from the low-dimensional feature subspace.

To support the distributed computations throughout the CPS, it needs a novel computing architecture to offer high-performance computing over multiple spatial and temporal scales and to support Internet of Things for machine-to-machine communications. We develop a hierarchical distributed Fog computing architecture for the next generation CPS. A prototype of such architecture for smart pipeline monitoring is implemented to verify its feasibility in real world applications.

Regarding the applications, we examine false data injection detection in smart grid. False data injection is a type of malicious attack which can threaten the security of energy systems. We examine the observability of false data injection and develop statistical models to estimate underlying system states and detect false data injection attacks under different scenarios to enhance the security of power systems (http://digitalcommons.uri.edu/cgi/viewcontent.cgi?article=1508&context=oa_diss).

GENETIC PROGRAMMING TECHNIQUES FOR REGULAR EXPRESSION INFERENCE FROM EXAMPLES

Fabiano Tarlao
ftarlao@gmail.com

Department of Engineering and Architecture, University of Trieste, Italy

MACHINE Learning (ML) techniques have proven their effectiveness for obtaining solutions to a given problem automatically, from observations of problem instances and from examples of the desired solution behaviour. In this thesis we describe the work carried out at the Machine Learning Lab at University of Trieste on several real world problems of practical interest.

The main contribution is the design and implementation of a tool, based on Genetic Programming (GP), capable of constructing regular expressions for text extraction automatically, based only on examples of the text to be extracted as well as of the text not to be extracted. Regular expressions are used in a number of different application domains but writing a regular expression for solving a specific task is usually quite difficult, requiring significant technical skills and creativity. The results demonstrate that our tool not only generates text extractors with much higher effectiveness on more difficult tasks than previously proposed algorithms, it is also human-competitive in terms of both accuracy of the regular expressions and time required for their construction. We base this claim on a large-scale experiment involving more than 1700 users on 10 text extraction tasks of realistic complexity. Thanks to these achievements, our proposal has been awarded the Silver Medal at the 13th Annual “Humies” Award, an international competition that establishes the state of the art in genetic and evolutionary computation.

Moreover, in this thesis we consider two variants of the proposed regular expressions generator, tailored to different application domains(i) an automatic Regex Golf game player, i.e., a tool for constructing a regular expression that matches a given list of strings and that does not match another given list of strings; and, (ii) the identification of Genic Interactions in sentences extracted from scientific papers.

This thesis also encompasses contributions beyond the field of Genetic Programming, including: a methodology for predicting the accuracy of text extractors; a novel learning algorithm able to generate a string similarity function tailored to problems of syntax-based entity extraction from unstructured text; a system for continuous reauthentication of web users based on the observed mouse dynamics; a method for the authentication of an unknown document, given a set of verified documents from the same author; a method for user profiling based on a set of his/her tweets; automatic text generators capable of generating fake reviews for a given scientific paper and fake consumer reviews for a restaurant. The proposed algorithms employ several ML techniques ranging from Grammatical Evolution to Support Vector Machines, from Random Forests to Recurrent Neural Networks (<https://drive.google.com/file/d/0B67gF86BZtPLdmNyYzZUNF8wTDg/view?usp=sharing>).

DATA-DRIVEN STUDY ON TWO DYNAMIC EVOLUTION PHENOMENA OF SOCIAL NETWORKS: RUMOR DIFFUSION IN ONLINE SOCIAL MEDIA AND BANKRUPTCY EVOLUTION AMONG FIRMS

Shihan Wang
ShihanW@trn.dis.titech.ac.jp
Tokyo Institute of Technology, Japan

THE fast growth of computational technologies and unprecedented volume of data have revolutionized the way we understand our society. While the social network structure is commonly used to conceptualize and describe individuals and collectives in the highly connected world, social network analysis becomes an important means of exploring insights behind this social structure. Social networks usually keep evolving slowly over time. This evolution can become very dramatic when facing external influences, which raised new challenges for scholars in understanding the complex social phenomenon.

In the thesis, we concentrate on the dynamic evolution phenomena of social networks caused by external factors from both interpersonal and inter-organizational perspectives: rumor diffusion in online social media (i.e. interpersonal network) and bankruptcy evolution among the firms (i.e. inter-organizational network). Driven by real big data resources, we applied various computational technologies to explore the behavioral patterns in dynamic social networks and provide implications for solving these social problems.

From the individual perspective, we explore the rumor diffusion phenomenon in online social media (i.e. Twitter in particular). With the extremely fast and wide spread of information, online trending rumors cause devastating socioeconomic damage before being effectively identified and corrected. To fix the gap in real-time situation, we propose an early detection mechanism to monitor and identify rumors in the online streaming social media as early as possible. The rumor-related patterns (combining features of users attitude and network structure in information propagation) are first defined, as well as a pattern matching algorithm for tracking the patterns in streaming data. Then, we analyze the snapshots of data stream and alarm matched patterns automatically based on the sliding window mechanism. The experiments in two different real Twitter datasets show that our approach captures early signal patterns of trending rumors and have a good potential to be used in real-time rumor discovery.

From the organizational perspective, we understand the dynamic evolution phenomenon of inter-firm network emerging from bankruptcy. When the bankruptcy transfers as a chain among trade partners (i.e. firms), it causes serious socioeconomic concerns. Beyond previous studies in statistical analysis and propagation modeling, we focus on one underlying human-related factor, the social network among senior executives of firms, and investigate its effects on this social phenomenon. Based on empirical analysis of real Japanese firms data in ten years, an agent-based model is particularly proposed to understand the role of this human factor in two perspectives: the number of social partners and the local interaction mechanism among firms (i.e. triangle structure in

inter-firm social network). Using real and artificial datasets, the beneficial effects of a number of social partners are well examined and validated in various simulated scenarios from both micro and macro levels. Our results also indicate the influential strategies to keep firms resilience when facing the bankrupt emergency.

Besides the contributions we made in each research field respectively, our study in this thesis enhances the understanding of dynamic independent and interactive behaviors in complex social phenomena, and provides a good perspective to seek solutions in other computational social problems.

(https://drive.google.com/file/d/0B_X4oFwLeg7udDhCempyN3JFbEk/view?usp=sharing)

EXPLORING THE KNOWLEDGE CURATION PROCESS OF ONLINE HEALTH COMMUNITIES

Wanli Xing
wanli.xing@ttu.edu
Texas Tech University, USA.

MORE and more people turn to online health communities for social support to satisfy their health-related needs. Previous studies on social support and online health communities in general have focused on the content of social support and the relationship of social support with other entities using traditional social science methods. Little is known about how social support facilitates the knowledge curation process in an online health community. Moreover, the presence of misinformation in online health communities also calls for research into the knowledge curation process in order to reduce the risk of misinformation. This study uses data mining technologies to analyze around one million posts across 23 online health communities along with 900 post information accuracy data. It aims to reveal how information, through social support, flows between the community users working as a whole to dynamically curate knowledge and further interacts with information accuracy.

Text mining methods was used to analyze the 1 million posts data to characterize information flow among the three user positional roles from a quantitative perspective. The results showed that (i) xperiphery users instead of core users dominate the quantitative information flow to request and receive informational support for knowledge curation in online health communities and (iii) the xperiphery users showed the best potential for generating new information. Granger causality was then employed to analyze the data mining results to characterize information flow between the three user positional roles from content perspective. The results demonstrated that (i) it was the xperiphery users who played a central role in directing the content information flow and development; (ii) however, the xperiphery users were still the least active user group in responding to other user positional roles.

Information flow was further quantified from temporal perspective using Directed Acyclic Graphs. K-means clustering and negative binomial regression were further employed to identify three distinct information flow patterns for users to curate knowledge in online health communities

and each with distinct characteristics. Further logistic regression models were built to examine the interaction between the identified information flow patterns with information accuracy. The results showed that (i) information patterns and time had a statistically significant influence on information accuracy, and (ii) information accuracy also showed distinct variation trends between information flow patterns. These findings not only have important implications for social support use, delivery and social support research methodologies but also can inform future online health platform design.

(https://www.researchgate.net/publication/317224472_EXPLORING_THE_COLLECTIVE_KNOWLEDGE_CURATION_PROCESS_OF_ONLINE_HEALTH_COMMUNITIES)

SEMANTIC SIMILARITY ANALYSIS AND APPLICATION IN KNOWLEDGE GRAPHS

Ganggao Zhu

gzhu@dit.upm.es


Universidad Politecnica de Madrid, Spain

THE advanced information extraction techniques and increasing availability of linked data have given birth to the notion of large-scale Knowledge Graph (KG). With the increasing popularity of KGs containing millions of concepts and entities, the research of fundamental tools studying semantic features of KGs is critical for the development of KG-based applications, apart from the study of KG population techniques. With such focus, this thesis exploits semantic similarity in KGs taking into consideration of concept taxonomy, concept distribution, entity descriptions, and categories. Semantic similarity captures the closeness of meanings. Through studying the semantic network of concepts and entities with meaningful relations in KGs, we proposed a novel WPath semantic similarity metric and new graph-based Information Content (IC) computation method. With the WPath and graph-based IC, semantic similarity of concepts can be computed directly and only based on the structural and statistical knowledge contained in KG. The word similarity experiments have shown that the improvement of the proposed methods is statistically significant comparing to conventional methods. Moreover, observing that concepts are usually collocated with textual descriptions, we propose a novel embedding approach to train concept and word embedding jointly. The shared vector space of concepts and words has provided convenient similarity computation between concepts and words through vector similarity. Furthermore, the applications of knowledge-based, corpus-based and embedding-based similarity methods are shown and compared to the task of semantic disambiguation and classification, in order to demonstrate the capability and suitability of different similarity methods in the specific application. Finally, semantic entity search is used as an illustrative showcase to demonstrate the higher level of the application consisting of text matching, disambiguation and query expansion. To implement the complete demonstration of entity-centric information querying, we also propose a rule-based approach for constructing and executing SPARQL queries automatically. In summary, the thesis exploits

various similarity methods and illustrates their corresponding applications for KGs. The proposed similarity methods and presented similarity based applications would help in facilitating the research and development of applications in KGs (<http://oa.upm.es/47031/>).

Cognitive Computing: Theory and Applications

BY V. N. GUDIVADA, V.V. RAGHAVAN, V. GOVINDARAJU, C.R. RAO (EDITORS) - ISBN: 978-0-4446-3744-4



REVIEWED BY WEIYI MENG

This is a timely book about an emergent discipline – cognitive computing. It actually belongs to an interdisciplinary domain encompassing cognitive science, neuroscience, data science, and high performance computing. The book provides an excellent coverage of two primary lines of research in this discipline. One is cognitive science which is a discipline that studies human mind and human cognition. This line of research covers neuroscience, psychology, linguistics, among others. The other line is largely based on computer science, especially the following subdisciplines: high-performance and cloud computing, machine learning, NLP, computer vision, information retrieval, data management and data science.

The book is comprised of 11 chapters written by 20 leading researchers in various relevant areas of cognitive computing. The first two chapters provide an excellent

introduction to cognitive computing and set the stage for reading the rest of the book. They introduce key concepts, architectures and systems, principles and theorems, as well as recent advances in cognitive computing. The next five chapters present various concrete methods that can be applied to tackling cognitive computing problems. These methods include graph-based visual analytics, machine learning algorithms for cognitive analytics, random forest model based big data classification, and Bayesian additive regression tree. The final four chapters present in-depth case studies of four application areas that can benefit from research in cognitive computing: food-water-energy, education and learning, spoken language processing, and Internet of Things (IoT). The book strikes a good balance in breadth and depth in introducing the state-of-the-art of cognitive computing to the readers. The chapters I read are all very well written – rich in content, informative and up-to-date, and clear in writing.

This book can be easily adopted as a textbook for an advanced undergraduate or graduate course on cognitive computing. It will also be an excellent reference book for a number of computer science courses in the areas of data science, big data analytics, and machine learning. Researchers and practitioners who are interested in these areas can also benefit greatly from reading this book. I personally learned a lot from reading this book and I strongly recommend it.

THE BOOK:

GUDIVADA, V.N.,
RAGHAVAN, V.V., GOVINDARAJU,
V., RAO, C.R. (EDS) (2016),
COGNITIVE COMPUTING:
THEORY AND APPLICATIONS, 404
P.
ELSEVIER
PRINT BOOK ISBN : 9780444637444

ABOUT THE REVIEWER:

WEIYI MENG
Chair and Professor, Department of
Computer Science, Binghamton
University, Binghamton, NY 13902
Contact him at: meng@binghamton.edu

IEEE Computer Society
1730 Massachusetts Ave, NW
Washington, D.C. 20036-1903

Non-profit Org.
U.S. Postage
PAID
Silver Spring, MD
Permit 1398