ISIF Perspectives On Information Fusion

May 2020
Volume 3
Number 1

ISIF INTERVIEW: WHAT IS THE IMPACT OF DATA FUSION ON THE SOCIAL AND POLITICAL LIFE OF CITIES?

AUTONOMY CHALLENGES FOR THE NEXT GENERATION OF MARS ROVERS

BOOK REVIEW: *CAUSAL INFERENCE IN STATISTICS: A PRIMER*

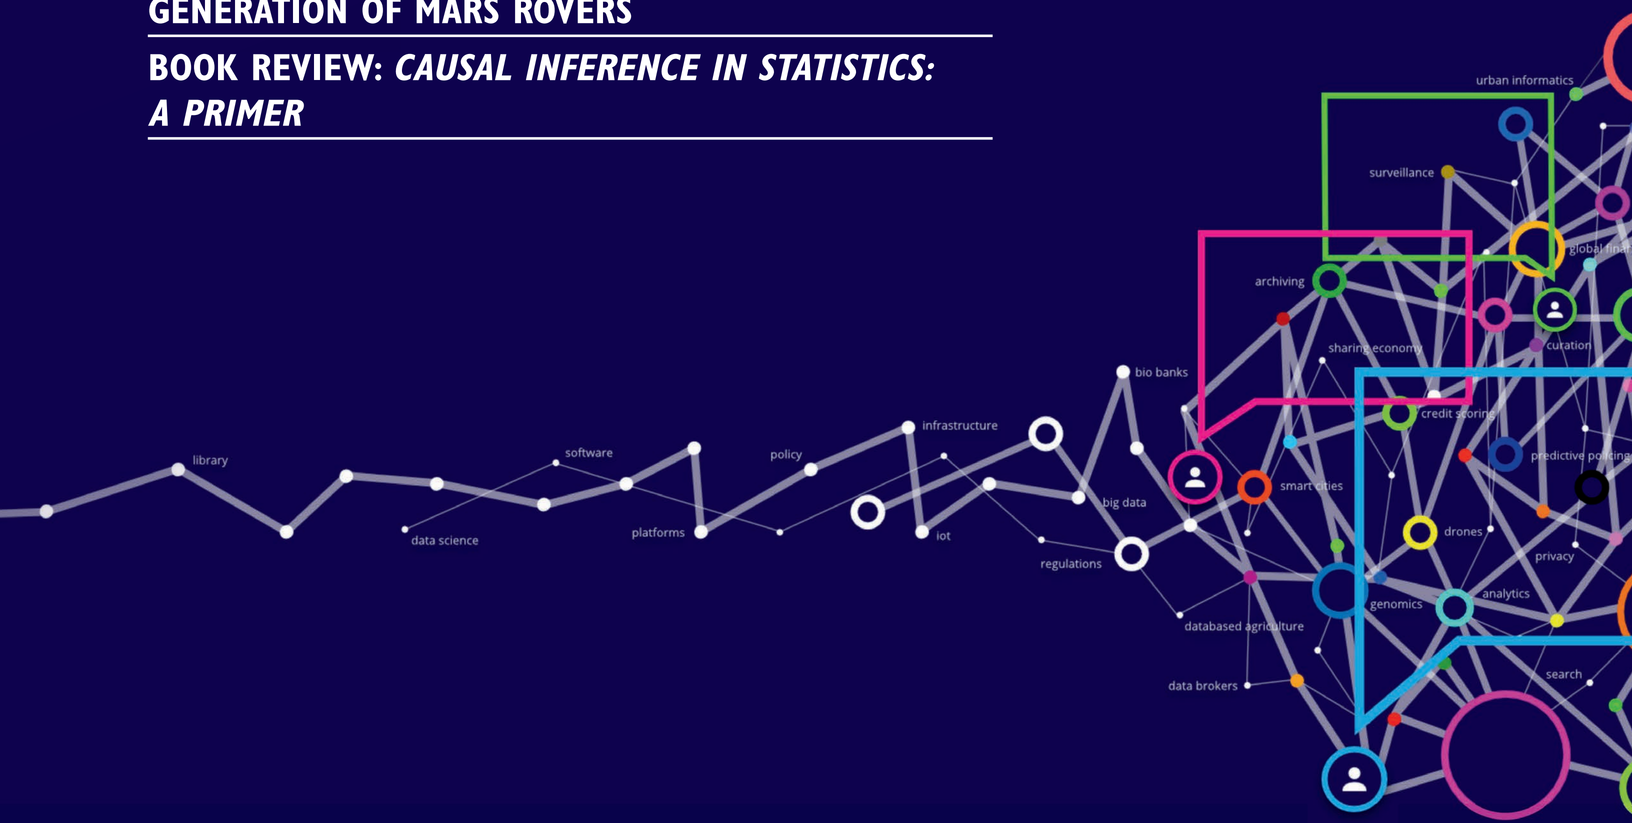

Publication of the
INTERNATIONAL SOCIETY OF INFORMATION FUSION

Perspectives

Editor-in-Chief
Roy Streit
r.streit@ieee.org
Metron, Inc.
Reston, Virginia, United States

Associate Editor-in-Chief
Jason Williams
jlw@ieee.org
CSIRO Data61
Brisbane, Queensland, Australia

Administrative Editor
David W. Krout
dkrout@apl.washington.edu
Applied Physics Laboratory
Seattle, Washington, United States

Production Manager
Kristy Virostek
virostek5@verizon.net

Area/Associate Editors

Higher level fusion
Anne-Laure Jousselme
anne-laure.jousselme@cmre.nato.int
NATO Centre for Maritime Research and Experimentation
La Spezia, Italy

Multitarget tracking and multisensor fusion
Wolfgang Koch
wolfgang.koch@fkie.fraunhofer.de
Fraunhofer Institute for Communication, Information Processing and Ergonomics
Wachtberg, Germany

Bayesian methods, Monte Carlo methods, control, image and video processing, tracking
Lyudmila Mihaylova
l.s.mihaylova@sheffield.ac.uk
University of Sheffield
Sheffield, United Kingdom

Tracking, multistatic radar tracking
Murat Efe
Murat.Efe@eng.ankara.edu.tr
Ankara University
Ankara, Turkey

Statistical signal processing, estimation theory, sensor fusion
Emre Özkan
emreo@metu.edu.tr
Middle East Technical University
Ankara, Turkey

Multiagent systems and distributed sensor networks, machine learning, data mining
Jesus Garcia Herrero
jgherrer@inf.uc3m.es
Universidad Carlos III de Madrid
Colmenarejo, Madrid, Spain

ISIF Perspectives

Perspectives seeks bridging articles, expository papers and tutorials, classroom notes, and announcements on topics of general interest to the ISIF Fusion community. Fresh points of view on established topics are especially welcome, as are articles on topics of interest to the ISIF annual fusion conference. Papers containing new research should be directed to JAIF or other research journal. The standing Call for Papers (CfP) for ***Perspectives*** can be found at http://isif.org/sites/isif.org/files/CfP%20for%20Perspectives%202019_04APRIL2019.pdf. ***Perspectives*** is published annually by ISIF.

More detailed guidelines and submission instructions for authors may be found at http://perspectives.msubmit.net/cgi-bin/main.plex?form_type=display_auth_instructions. The average length for submissions is approximately six (6) pages (in JAIF two-column format). All submissions will be reviewed for content and style, as well as suitability for ***Perspectives***. All papers accepted for publication will be written in a relaxed, colloquial style that facilitates understanding by a wide audience. Articles containing significant original research should be submitted to JAIF.

Cover: Art courtesy of Wataru Wantanabe (https://wataruwatanabe.net). The illustration portrays how data is growing faster and more ubiquitous, complex, and powerful over time. The speech bubbles represent discussions that people and society need and are starting to address, such as privacy, inclusivity, and control. The style is similar to timelines, charts, and maps, which are commonly used to visualize data.

Perspectives (ISSN pending) is published annually by the International Society of Information Fusion (ISIF). The responsibility for the contents rests upon the authors and not upon ISIF, the Society, or its members. ISIF is a California Nonprofit Public Benefit Corporation at P. O. Box 4631, Mountain View, California 94040. Copyright and Reprint Permissions. Abstracting is permitted with credit to the source. For all other copying, reprint, or republication permissions, contact the Administrative Editor. Copyright© 2020 ISIF, Inc.

ISIF Website: http://www.isif.org

Table of Contents

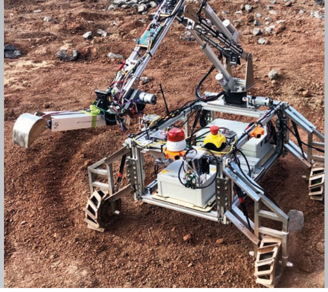

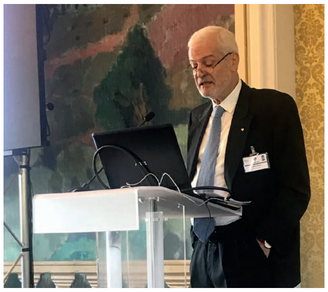

TABLE OF CONTENTS CONTINUED ON NEXT PAGE

OTHER EVENTS AND WORKSHOPS

BOOK REVIEW

BOOK REVIEW REFLECTION

PERSPECTIVES MAGAZINE

Welcome to the third issue of *Perspectives* magazine. We hope that you enjoy the diversity of articles in this issue. The diversity is, of course, intentional! A few highlights…

The issue features an interview with Prof. Tracey Lauriault of Carleton University. It is a follow-on to her plenary presentation at FUSION 2019 in Ottawa, and it was conducted by two *Perspectives* Associate Editors, Wolfgang Koch and Anne-Laure Jousselme. The interview is a first for *Perspectives* and the International Society of Information Fusion (ISIF). It is of special interest because it discusses the increasingly important *societal impact of data fusion technologies* in the context of "Smart Cities". The full transcript can be found on the ISIF website.

The feature article "Autonomy Challenges for the Next Generation of Mars Rovers" is a report on an ambitious prize-winning, student-led UK project for the design of an autonomous scooping mechanism for the European Rover Challenge, currently the biggest space robotics event in Europe. *Perspectives* magazine is always open to reports on exceptional student-led projects such as this one.

Statistics has long taught—quite correctly—that correlation is not to be confused with causality (or cause and effect). Consequently, causality is often deemed inaccessible to statistical methods. The distinguished Prof. Judea Pearl, however, has long argued otherwise in a series of books. His most recent book *Causal Inference in Statistics, A Primer* is reviewed here. The review by Dr. Lawrence Stone is unusual by design, as the request was for him to introduce the subject before actually reviewing the book. The result is a stimulating read. A supplemental response to the review was written by Dr. Alfonso Farina and provides still more perspective.

The issue includes a report on the highly successful FUSION Conference in Ottawa in 2019. The shortlisted papers for Best Paper and Best Student Paper awards are listed, and the abstracts of the winning papers are included here. Reports from several workshops sponsored in part by ISIF that were held throughout the past year are also given. A report on the first ever Maritime Situational Awareness Workshop (MSAW) in the incomparably beautiful Lerici, Italy on the Italian Riviera is included because of its relevance to the data fusion community.

This issue is possible only by the timely and thoughtful contributions of the authors, reviewers, and editors. A special thanks goes to Kristy Virostek, our Production Manager, who somehow knows how to keep us all on schedule. (The EiC needs more help than most.) Thank you all for your generous help with this issue.

Roy Streit
Editor-in-Chief

POSTSCRIPT

I wrote the Introduction in early March. Since then, personal perspectives (pun intended) about what is important in our lives and the best direction of our work are shifting for many of us. It will take time for each of us to sort these things out to find a new normal.

Personal choices and instincts drive many of us working in Information Fusion to seek quantitative models of the world. Consequently, in this New World of the last few months, two areas of research will naturally beckon for our attention in the coming days and months.

One area is the mathematical modeling of the spread of infectious diseases. This kind of work began in the 1920s with differential equation "compartmental" models of the numbers of Susceptible/Infectious/Recovered (SIR) individuals. The models are nowadays both extended (SEIR and SEIRS) and more sophisticated (stochastic models), as well as data adaptive [1]. The senior author (Vincent Poor) was a Plenary Speaker at Fusion 2002 in Annapolis, MD.

The other area concerns spatial-temporal data modeling. With the abundance of multisource data, these models can be both dynamical and stochastic. Several fairly recent books are devoted to this increasingly important area. Both areas fit within the realm of Information Fusion, and both will influence us—individually and as an Information Fusion Society—in the coming days.

My thoughts and best wishes go out to every member of the Fusion community, their families and friends, their countries, and to all.

REFERENCE

1. Eletreby, R., Zhuang, Y., Carley, K. M., Yağan, K. M., and Poor, H. V. The effects of evolutionary adaptations on spreading processes in complex networks. *Proc. National Academy of Sciences*, Vol. 177, 11 (2000), 5664–5670.

What is the Impact of Data Fusion on the Social and Political Life of Cities?

Wolfgang Koch
Fraunhofer FKIE
Wachtberg, Germany
Wolfgang.Koch@fkie.fraunhofer.de

Anne-Laure Jousselme
NATO STO Centre for Maritime Research and Experimentation
La Spezia SP, Italy
Anne-Laure.Jousselme@cmre.nato.int

Tracey Lauriault
Carleton University
Ottawa, Canada
Tracey.Lauriault@carleton.ca

TRACEY LAURIAULT INTERVIEWED BY WOLFGANG KOCH AND ANNE-LAURE JOUSSELME
BONN, GERMANY, SEPTEMBER 14, 2019

Information Fusion has a clear year of birth. Our community was born around 1984 when the Joint Directors of Laboratories (JDL) model of data fusion was created, which is still basic for us.[1] 1984 is George Orwell's very year, Tracey. Are you creating Orwellian surveillance infrastructures for cities that may turn his dystopia into reality?

I was completely unaware of this, because I don't know your community very well. But this is a pretty interesting coincidence, indeed. My whole work, however, is an attempt to counter the more Orwellian notion of the "surveillance city" or the "watching city". Or even the Frankenstein version of the Smart City that we are starting to see more and more often, where the parts do not interact with each other, where the countless machines, sensors, devices, are going to break down and fall apart, as these have short shelf lives in outdoor environments or frequently used on buildings.

The pioneers of data fusion seem to have been unaware of it also. You have delivered a most stimulating keynote at the 22nd International Conference on Information Fusion on July 3, 2019 in Ottawa—"Fusion of Data in an Open Smart City Context". Tell us, Tracey, what are "Smart Cities" all about?

As I am seeing them, and I paraphrase Rob Kitchin's work on networked urbanism, Smart Cities are technologically instrumented and networked systems of systems that are interlinked and integrated. Here, vast troves of big urban data are being generated by sensors and administrative processes that are used to manage and control urban life in real-time. The focus in this kind of Smart City is most often to quantify and manage infrastructure, mobility, business, and online government services. Of course, algorithms for data fusion and resources management play a key role here as strategically placed sensors around the urban landscape monitor the citizens, their behavior in the city, and also city assets, resources, services, and many other factors of urban living.

Who has an interest to realize these complex infrastructures? Building Smart Cities requires big investment.

Administrators and elected officials are investing in Smart City technologies and data analytical systems to inform how to innovatively, economically, efficiently, and objectively run and manage the city. This is a good thing, of course, but it needs to be governed in the public interest. In combination with large data bases, such as Geographic Information Systems, and sensor derived real-time data, local authorities use the insights gained to manage the challenges that the city faces in sectors such as crime prevention, traffic management, energy use, or waste reduction.

Are you aware of any very first lessons learned from Smart City projects in Canada and elsewhere?

Yes, indeed I am. In our case studies, we have identified the reasons for deploying Smart City initiatives, the beneficiaries, underlying governance models and deployment strategies, citizen engagement, Smart City business models, and so on. Evidently, Smart Cities are new and emerging. This means that the citizens themselves do not generally know what is coming and may not be the drivers of the development or even be involved in expressing guidelines to follow to ensure that the Smart City works for them. In general, I feel there is some sort of technological solutionism in this business, over-engineering. Don't you think it is much more healthy to identify real issues to be resolved with technology first, instead of creating technology first and then looking for issues? We definitely need overarching principles that govern Smart City design. Astonishingly enough, very few people consider sociotechnical and ethical principles for these types of complex systems. A more mature "technology-aware citizenship" is required.

[1] The Data Fusion Group of the JDL, a US Department of Defense committee, created the original Data Fusion Model in 1985. This functional model, aimed at facilitating understanding and communication among fusion theoreticians and practitioners, and its "Fusion Levels" have been driving the discussions since then. F. E. White and JDL published "Data Fusion Lexicon" in 1987.

While listening to you, it becomes quite clear that Smart Cities are obviously a huge topic for our information fusion community. Why are Smart Cities of so much interest for you personally? What aspects have attracted you as a researcher?

During the keynote at the Fusion conference in Ottawa, I was asking the question about "How do we govern the data and the technology at the level of the architecture?" Actually, I was talking about "Open Smart Cities" because of my fear of an Orwellian 1984 surveillance Smart City. We therefore have to build Smart Cities with a proper architecture, as Open Smart Cities from the very beginning. I worry that we are going to get into the lock down of information in the context of the Smart City. For quite a while, I started thinking about issues related to open data, open source, open science, open sensors, open architecture, open platforms, in addition to public engagement, public policy, the environment, sustainability, fairness, and accountability. I was thinking about all of those things because of my background as a grassroots activist, many years ago, on environmental issues, social justice issues, and even antinuclear issues in the early '80s. And how would those issues get mapped onto the Smart City.

More precisely speaking, what is the difference between a Smart City and an Open Smart City?

The definition of a Smart City is what we call networked urbanism or connected cities, which I defined earlier, whereas an Open Smart City is defined by the attempt to bring in these other ideas and to understand them in this context. A city is operationalized and managed by a city government. Those who are going to make decisions about Smart Cities are going to be those Smart City officials. Smart City officials manage the blue (the water), the green (the environment), and the gray (the built environment) material of the city to better govern and help people live better in the city. And they may need to rely on data fusion or real-time data to do so.

"Astonishingly enough, very few people consider sociotechnical and ethical principles for these types of complex systems. A more mature 'technology-aware citizenship' is required."

Quite frankly speaking and seen from the perspective of Germany, a country that has had a horrible totalitarian experience, why should the idea of Open Smart Cities architectures prevent oppression?

In my view, "smart" technology is no longer just about operations. It is about governance, governance of the data, the processes, the infrastructure, and the outcomes of processes and decisions, which means the very ways that smart technologies impact citizens, residents, and visitors. Are these technological systems in the public interest and for the public good? This is the key question. And moreover: What are the benefits and the downfalls, not only of each individual part for specific institutions, but for all of those living in a city when these things become interoperable and interconnected?

Why exactly is the concept of an Open Smart City a promising concept in view of the public interest and the common good? What characterizes its "openness" in view of this?

Open Smart Cities, as I see them, are about applying sociotechnological systems-of-systems thinking towards the creation of a city where residents, civil society, academics, and the private sector collaborate with public officials. The key is to mobilize data and technologies when warranted in an ethical, accountable, and transparent way to govern the city as a fair, viable, and livable commons and balance economic development, social progress, and environmental responsibility. That requires integrated social and technological system thinking and doing.

Can you please be a bit more precise? What is the impact of these goals on information communication technologies and information fusion?

Since governance in an Open Smart City is to be ethical, accountable, and transparent, these principles apply to the governance of the social and technical platforms, which include data, algorithms, data fusion processes, skills, infrastructure, and knowledge. The same is true if an Open Smart City is to be participatory, collaborative, and responsive. An informational infrastructure has to be created in such a way that it can technically and organizationally enable meaningful participation of the civil society, the private sector, the media, academia, and residents in the governance of the city and the social and technical processes that operate the city and this involves shared rights and responsibilities. All this has an effect of the design of data technologies to be developed, whereby they are acquired and deployed in such a way that are fit for purpose and can be repaired and queried and governed to mitigate mission creep. Moreover, when it comes to technology, wherever possible, their source codes are open, adhere to open standards, are interoperable, durable, secure, and where possible locally procured and scalable. The information fusion infrastructure of an Open Smart City is used and acquired in such a way as to reduce harm and bias, increase sustainability, and enhance flexibility. It may defer when warranted to automated decision making and therefore the design of these systems makes them legible, responsive, adaptive, and accountable. It is quite clear, of course, that in an Open Smart City, data technologies are not always the solution to many of the systemic issues cities face, nor are there always quick fixes to complex problems, such as homelessness, income inequality, racism, etc. These problems require innovative, sometimes long term, social, organizational, economic, and political processes and solutions.

How can you step out of your roles as technical experts and scientists and also be technological citizens, and collaborate with those on the social side of the equation to ensure that what you are building benefits us all, that you make our world more livable, safe, equal, and fair for all?"

By the way, what is your scholarly background that led you to be interested in Smart Cities?

Oh, I really have done and been interested in many things in my life. But always I had interest in mapping and data. A more direct path to my current interest was cybercartography during my Ph.D. work. It is a multimodal, multisensory, multidimensional, multidisciplinary version of cartography that is interactive and online. I also have an interest in spatial data infrastructures, which are about delivering data, spatial data, over the internet to the Canadian population, that includes Global Positioning Systems, satellites, radar, digital maps, standards, policies, technologies such as sensors, and all kinds of different institutions working together such as agriculture, defense, transportation, or natural resources. It sounds easy, but is very complex to do. If it works well, you don't even know it is happening. But also, on the cybercartography project, I was the student that looked at sensors, in particular the electronic nose and trying to figure out how to do olfactory cartography. Concurrently, outside of the academy, I was becoming very active in the open data movement in Canada.

Another very important step for me was to work with Professor Rob Kitchin in Ireland, an important actor in the critical Smart City area and a very important actor in critical data studies as well. As part of his Programmable City Project, I conducted case studies about city data and technology to better understand their social and material implications.

Obviously, there is an interest of Big Tech to make cities smarter. How do you think we can counter the concentration of ever more data—equaling more power, not only financially—in the hands of very few companies with a financial strength bigger than that of many countries?

They are not investing at all. The companies early on offered these technologies as gifts. The company said, "You can test out", and so Dublin, for example, became a laboratory for technology companies. And in a way, the Smart Cities that are being promised are giant laboratories. I work in a university with smart taps, as we want washrooms where you do not touch things, so there are lots of sensors, but now if we want to repair the tap in your bathroom…well, you need an electrician, you need a plumber, you need a software someone, and a sensor expert. Think about it! When your tap breaks here, you only have to call one person or you fix it yourself! Well, in a smart building, you need whoever installed the entire smart system. But these companies come and go. They don't last forever. So, whom do you call? Who gets the contract for your smart building? Who upgrades the sensors? Who does the maintenance? What happens when you scale that up to a whole city? You suddenly have Frankenstein cities, if you will, with their lack of interoperability, these multiple negotiated experiments, where the companies benefit because they're getting access to people's data.

Would you say that interoperability is the key for the success of Smart Cities?

IBM created the first notion of a Smart City. *Ubiquitous computing*, which you know way more than I do, you could call that some sort of early Smart City thinking. It enables a sensing city, a sentient city, an intelligent city, a wise city, an accessible city, a safe city. A safe city would involve policing, emergency preparedness, military…which is also part of a Smart City. Those highly automated, highly technologically informed environments already exist. They are just not interconnected in an all-seeing way. So, if we have all of these separate systems fused and interoperating, would interoperability be our friend or is interoperability our foe? We have to have interoperability, but actually in a Smart City, a lack of interoperability might save us, because we might mitigate the surveillance city scenario that is being witnessed in some countries.

From your work and your experience, do you think it is possible to prepare procedures in order to create some sort of a certificate that a certain Smart City technology is societally acceptable?

A certificate would be too strong a word, perhaps, but yes, the question is right: Will a certain technology actually lead to the health and well-being of the society? And in addition, will it bring the environment in balance with the economy? So, let us go to the first part of the definition of an Open Smart City: fair, transparent, and accountable and balancing the social, economic, and environmental needs of the city. The economy and the city as a habitat or as a human commons needs to be sustainable, it needs to not damage the environment, and balance economic benefit.

So, we will have different criteria than those we usually use, i.e., measurable, quantifiable criteria. You are telling us to consider criteria which are not as easily quantifiable and which comprise environment and human factors.

That would be really interesting. So, what is the social and technological code of conduct and governance, whatever, the strategies or principles for data fusion experts? And, I mean, you have such a diverse community doing diverse things. Do no harm, how do you benefit all and marginalize none, how do you reduce racial profiling or other forms of bias, and how do you build with equality and justice in mind, but how does that translate in a large social and technical complex systems of systems like a Smart City? What ethical principles do you need to have in play? What is the decision tree? If this then that...if that then this. It is about the whole and it is about the parts, it is about governing the systems of systems.

What about the vulnerability of Smart Cities? We Europeans are afraid of already being a target in hybrid warfare. A cyber-attack on a Smart City should be so easy.

You are asking for a robust Smart City, a defensive Smart City, a city with an electronic bubble around it that makes it impenetrable from a cyber threat? Doesn't this mean we are back to the walled cities of the mediaeval times? Cities of quartz, silicon cities. Is that the direction where we are moving to? Because in a way, information communication technology is perfect for that, isn't it? Actually, I am thinking of something better. I want to know if my city is happy. I want to know from the social media profiles of my citizens if they are feeling blue [depressed]. I want to know from Twitter or whatever is rocking any social media boat at any given time: "How are my people doing?" And how can I measure behavior, sentiment, movement, weather, how they feel during a football match when there is a cool evening summer breeze at the end of a hot day, and how can I understand how they are moving and behaving from their phones? So, should I model the city and the behavior of the citizens so that it's optimal collectively? So, suddenly you're getting into totalitarian kind of thinking…one that is called social physics, or a kind of technological determinist solutionism. There needs to be a balance, fair, accountable, just, etc.

Oh, you are describing the older sister of Orwell's Big Brother then.

No, the scenery is a kind of nudging. But it could be stronger than nudging. But, it's this idea that you can and should mathematically model everything, from sentiment to mobility, and furthermore, that you should live according to the mathematical model that was developed. And it is this false dream of the optimal golden mean, right? There is an optimal equation that should be imposed on the society, and that I could come up with that equation. If only I had all these data! If only I had ubiquitous computing! Should we leave the governing of cities to the big companies and to social physicist who want to design models for us to live by? No, I do not like that kind of Smart City! Smart Cities are vulnerable, not only by an external attack, but also from within, if we do not govern them.

Are community organizations and nonprofit resident organizations thinking like Saul Alinsky's and his "community organizing" in Chicago, or Jane Jacobs, aware of Smart Cities as a tool to "smartly" transform societies into "Smart Tyrannies"?

Wolfgang Koch, Tracey Lauriault, and Anne-Laure Jousselme in Bonn, Germany.

That is a good question! It's different, the issues are different and smart technologies are a different type of city actor. In Barcelona, for example, it is the civic technology community that pushes for better decision making with the use of data analytics for specific problem solving. If you look to the Pirate Party [an elected political group] in Iceland—that too is kind of related. In Germany, you have the movement of the Computer Chaos Club that is involved in what I would call technological citizenship. They know their tech, they know their social science, and they know their politics, and they inform government technology policy. You however cannot solely have these groups of people govern but they should be part of the multisectoral and multidisciplinary teams, or a citizen city committee to advise on Smart City technologies. Many civic technology and open data community organizations are also technologically solutionist in their orientation, and often they are idealists, but they may also be libertarian and will miss social justice, equality, intersectional issues, and so on. As technological solutionists, they may not understand social policy. So, don't let them govern on their own, but we need them to be part the discourse, along with antipoverty organizations, childcare advocacy groups, immigration settlement organizations, policy experts, and of course city administrators, right? We need that sophisticated technical thinking, but not only that.

"Open Smart Cities, as I see them, are about applying sociotechnological systems-of-systems thinking towards the creation of a city where residents, civil society, academics, and the private sector collaborate with public officials."

So, in the end, what is your definition of Smart City?

There are many actors in the running of a city, but just like the Open Smart City V.1.0 Guide is new and just like the definition is new, and Smart Cities are new but are mostly for the moment disconnected smart intelligent systems. There is no one organization or group that is doing all of this, but everybody is doing a piece of it. The intention of the Guide and the Open Smart City definition was to bring all the pieces together in a way that people could understand, and provides something to aim for, a type of road map if you will, to ensure that our cities remain livable and just for all, and that we govern the technology in our best interest instead of having the big technology companies governing us!

How did you get in touch with our information fusion community?

Elisa Shahbazian called me up, as my friend Robert Davidson and colleague whom I serve with on the Multistakeholder Civil Society Advisory Committee to Open Government in Canada recommended me. In the end, I said to her quite frankly, I am a lightweight for you data fusion folks because I don't do the stuff you do. She replied, "We need to hear something different, and I heard your talks and thought you would be a good fit". That's why I chose and wanted to spend time with your community for a couple days, why I stayed as much as I could. I am innately curious and love brand new communities, as each has a culture of its own. I didn't fully comprehend the sessions I went to, but what I got was an understanding of the community. I got an understanding of the problems that the community is generally working on. Again, I don't comprehend the mechanisms and the algorithms, but I understood that this is a community that works in the deep recesses of infrastructure to ensure that things hum along, that they work, that things work together, and that the things are secure. It is really useful for me as a scholar and as a technological citizen to know about the important work that data fusionists do.

If you are using a car, you are not interested in how an automatic transmission actually works. You want to have it and to use it. What do you want from us? From the data fusion community?

Your community is building machines that create the information and that keep systems working. This information is directing autonomous cars or the traffic management system, the airplane, the electrical grid, defense system, weapons, etc. What I wish is that this work gets done with the recognition of some of foundational, ethical, and principles and values. And that that is articulated as part of the system you are involved with, the social and technical systems that are in our shared world, and these affect us. The sense that I get from you Wolfgang and you Anne-Laure, and the group of people that I met, was that you are so efficient at what you do because you are working at another level. You were already working at that principled level. This means if your good work is plugging into a city all the things that a city infrastructure is doing, and is intersecting with your work, and then your work is feeding it back, city officials have to know how to do their work properly to fit into yours; and in such a way that it doesn't become a surveillance city. In a way, what you are doing is you're creating the underlying infrastructure to do it, for good and for bad. So how to do your work to be on Princess Leia's side and not Darth Vader's?

We are worried about 1984, we see some places very much operating like Orwell's science fiction. We are also witnessing regime changes, and more rigid forms of thinking and power concentration. And you, as data fusion people, are building this big fused technical system that is the prerequisite of Orwell's world. How can your community technically support the socially responsible use of the power of the technologies you are building? How can you step out of your roles as technical experts and scientists and also be technological citizens, and collaborate with those on the social side of the equation to ensure that what you are building benefits us all, that you make our world more livable, safe, equal, and fair for all. I cannot give you an actual algorithm on how to do that, but certainly I see, and sense, based on my couple of days with the data fusion community in Ottawa, and after spending a full day with you two here in Bonn, that there is a willingness to do so, and the intellectual flexibility to do so, and dare I say a spiritual inclination or moral compass that would lead you to do so!

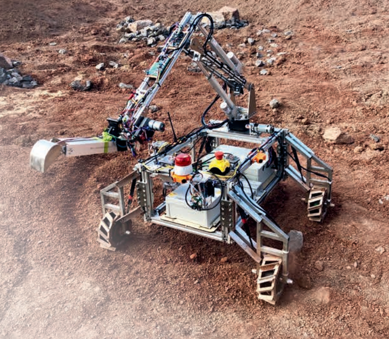

Autonomy Challenges for the Next Generation of Mars Rovers

Victor Covasan
Department of Aerospace Engineering
University of Sheffield, UK
vcovasan1@sheffield.ac.uk

Viktor Fedun
Lyudmila Mihaylova
Marwan Taher
Department of Automatic Control and Systems Engineering
University of Sheffield, UK

Joycelyn Fontanilla
Department of Chemical Engineering
University of Sheffield, UK

Hari Krishna
Department of Computer Science
University of Sheffield, UK

Abstract—Achieving full autonomy for a planetary explorer is the main requirement in rendering feasible missions when the communication time with the ground station does not allow real-time operation and monitoring. The design process involved in building the MarsWorks rover illustrates the challenges to be addressed in a typical surface exploration mission. This paper presents the main stages necessary to achieve rover autonomy in Mars-analogue environments. The focus is on two key areas: rough terrain navigation and autonomous manipulation with a six degree-of-freedom robotic arm. The first topic covers fundamental data fusion and Kalman filtering methods that estimate the current pose, as well as displacements from the starting position by means of visual-inertial odometry. An approach to guidance and control is then presented from the perspective of the dynamic window technique. Subsequently, autonomous grasping with increasing levels of automation is presented: from the low-level proportional–integral–derivative (PID) control to inverse kinematics, motion planning, computer vision, and automatic target recognition. Finally, onboard data handling, fusion of the sensor data used for scientific sample analysis, and communication with the ground station are briefly discussed. Each section presents future ambitions and possible ways of optimising individual subsystems of the MarsWorks rover.

INTRODUCTION

Planetary science missions are arguably the most emblematic accomplishments of the space industry, many of them (Curiosity, Perseverance, Europa Lander) being categorised as flagship-class missions, implying both immense efforts, technically and economically, but also colossal contributions to our understanding of the Universe. Flagship mission proposals are called for based on the Planetary Science Decadal Survey [1] published by the United States National Research Council. Such surveys consider the most relevant scientific questions of the decade and reflect the public interest in the exploration of certain celestial bodies. This combination of public enthusiasm and scientific return often motivate space agencies to organise student rover competitions that simulate innovations towards the next generation of planetary missions. Part of this series of contests and contenders is MarsWorks, an interdisciplinary team of students from the University of Sheffield, UK, a team who dedicated a substantial part of the academic year to designing and building a Mars Rover autonomous vehicle. The vehicle participated in the European Rover Challenge (ERC). Figure 1 shows the MarsWorks rover at ERC2019, which will be the subject of this article.

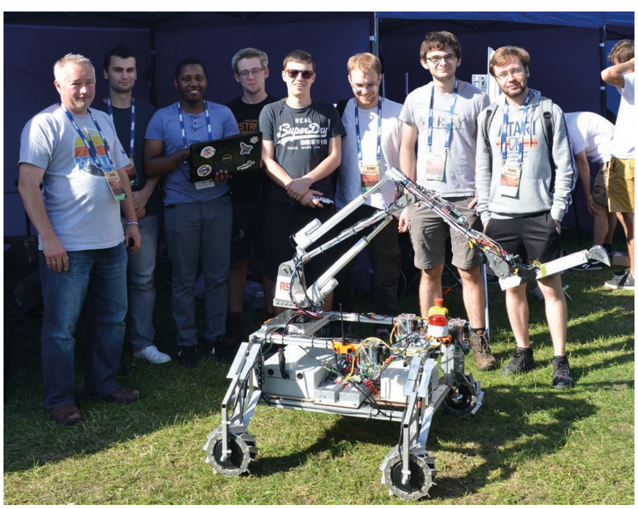

Figure 1
MarsWorks representatives at ERC2019.

The first planetary rovers, Lunakhod [2], were launched in the early 1970s and focused on extreme terrain mobility and small body/microgravity mobility. Navigation and control were difficult since computers were bulky and slow, thus Lunakhod was a teleoperated mission. The first rover on Mars was the National Aeronautics and Space Administration's (NASAs) Sojourner, a small 11.5 kg rover which explored the area within site of the Pathfinder Lander's camera, taking measurements of surface properties, imaging rocks, and obtaining their elemental composition. The Mars Exploration Rovers were the first to use advanced navigation methods such as visual odometry, a technology reaching maturity in the space sector which was also employed on the MarsWorks rover. This progression reveals the evolution of scientific objectives for Mars Rovers, starting with mobility demonstration, to search for water and life and more recently the investigation of in-situ resource utilisation potential or habitability [3].

The European Space Agency (ESA) ExoMars rover Rosalind Franklin is planned for launch in July 2020, with a focus on astrobiology. The mission will revolve around searching for past life on Mars, investigating gases and their sources and, by doing this demonstrating capabilities for a Mars sample-return mission in the future. The results delivered by the ExoMars rover will be complementary with the measurements taken by the Trace Gas Orbiter (TGO), which maps the distributions of hydrogen and methane in the atmosphere of Mars. More importantly, the satellite also serves as a communication link between future landers and the Earth, so rovers only need to uplink their measurements to TGO, which will relay them to the ground station. Sample-return missions have been known to have a great scientific return on investment as analysis is freed from the time, budget, or space constraints of spacecraft sensors; therefore, they are projected to become increasingly relevant over the next decades. Another motivation for sample-return research is the potential for asteroid capture and exploitation of resources from bodies located in the near vicinity of the Earth.

As a result, the tasks in ERC are based on the needs of a sample-return mission as well. The rovers involved shall be capable of autonomous navigation to a desired site and deep surface drilling once at the location, much like Rosalind Franklin [4], but also autonomous detection and collection of predefined targets in the Martian field. Additionally, rovers must be able to conduct onboard scientific analysis of collected soil samples and to operate a control panel in order to aid astronauts with maintenance of a Mars base. The concept of astronaut-aiding robots is motivated by the need to reduce the exposure time of astronauts to high radiation environments. As opposed to Earth, Mars lost its magnetic field with the cooling of its metallic core; therefore, the noxious solar radiation is not deflected around the planet but reaches the surface, which makes it a significant hazard to surface explorers, but especially to future human settlements. This is why it is crucial for most outdoors activities to continue to be performed by rovers and robots even with sustained human presence on Mars.

Achieving the highest degree of autonomy is also a continuing endeavour for current planetary explorers since the communication time between the ground station and spacecraft is much longer than the time available to respond to hazards. Specifically, the communication time to Mars is about 20 minutes; therefore, if a violent sandstorm starts to develop it might be too late for the storm warning to reach the Earth, then the "take shelter" command to be received and executed by the rover, because by that time the storm might have gained threatening proportions. However, the control algorithms employed also need to prove robust and predictable enough to be certified for space applications, this being a reason why inherently black-box approaches such as artificial neural networks are not suitable for safety-critical scenarios. Conversely, advanced model predictive approaches have been widely adopted by the space industry, especially for attitude control of orbiters, since they guarantee the desired performance and stability margins. On the other hand, increased autonomy drives the need for quantitatively more sensor measurements but also adequate accuracy of the estimated states, which in turn prompts the requirement for more computational power as well as more advanced data fusion mechanisms to handle the diverse range of sensors and increased data volume. Current rovers rely on a suite of sensors—infrared cameras, accelerometers, gyroscopes—for trajectory planning and attitude of the vehicle, but also spectrometers (to analyse the composition of sampled materials), atmospheric analysers, and radiation detectors. Often, measurements from all these transducers are requested by the same module of the rover for navigation, guidance, or control tasks, this requirement pointing back to the need for robust onboard data handling and fusion to ensure data compatibility.

With these considerations in mind, the article aims to demonstrate how control and data processing algorithms are implemented in a complex real-world system such as a Mars surface explorer with a view to achieving the highest degree of autonomy. The rest of the paper is organized as follows: the systems architecture is presented, the robotic arm control system is described, followed by the guidance, navigation, and control system. The latter sections focus on the sensor data fusion approaches used for the scientific measurements, and conclusions and the key lessons learned are given.

MarsWorks is a student-led project at the University of Sheffield scoped with designing and building a Mars Rover for the European Rover Challenge (ERC). This project is part of Sheffield Space Initiative (SSI), a highly cross-disciplinary space technology platform that is now developing a real heritage of success for the University of Sheffield and our science, technology, engineering, and mathematics (STEM) students in particular. SSI was founded in 2017 to further engage the University of Sheffield students in the science and engineering challenges involved in space exploration. MarsWorks has its origins in Project MoonWorks, as the first rover was dedicated to the fabrication of a miniature lunar vehicle which could retrieve ice samples from the depths of lunar craters. The team has a broad range of activities. A successful participation in the national competition organised by UK Students for the Exploration and Development of Space in 2018 led to a prize award for the best innovation for the developed advanced scooping mechanism. The project called MarsWorks, moved forward to design and build a fully autonomous Mars Rover to participate in the European Rover Challenge (ERC), the biggest space robotics event in Europe.

THE SYSTEM ARCHITECTURE

In modern space systems engineering, integration is primarily achieved in software, implying that all subsystems and components must communicate with a central computing element. As a result, it is worth starting the discussion with the high-level system breakdown together with the interfaces between

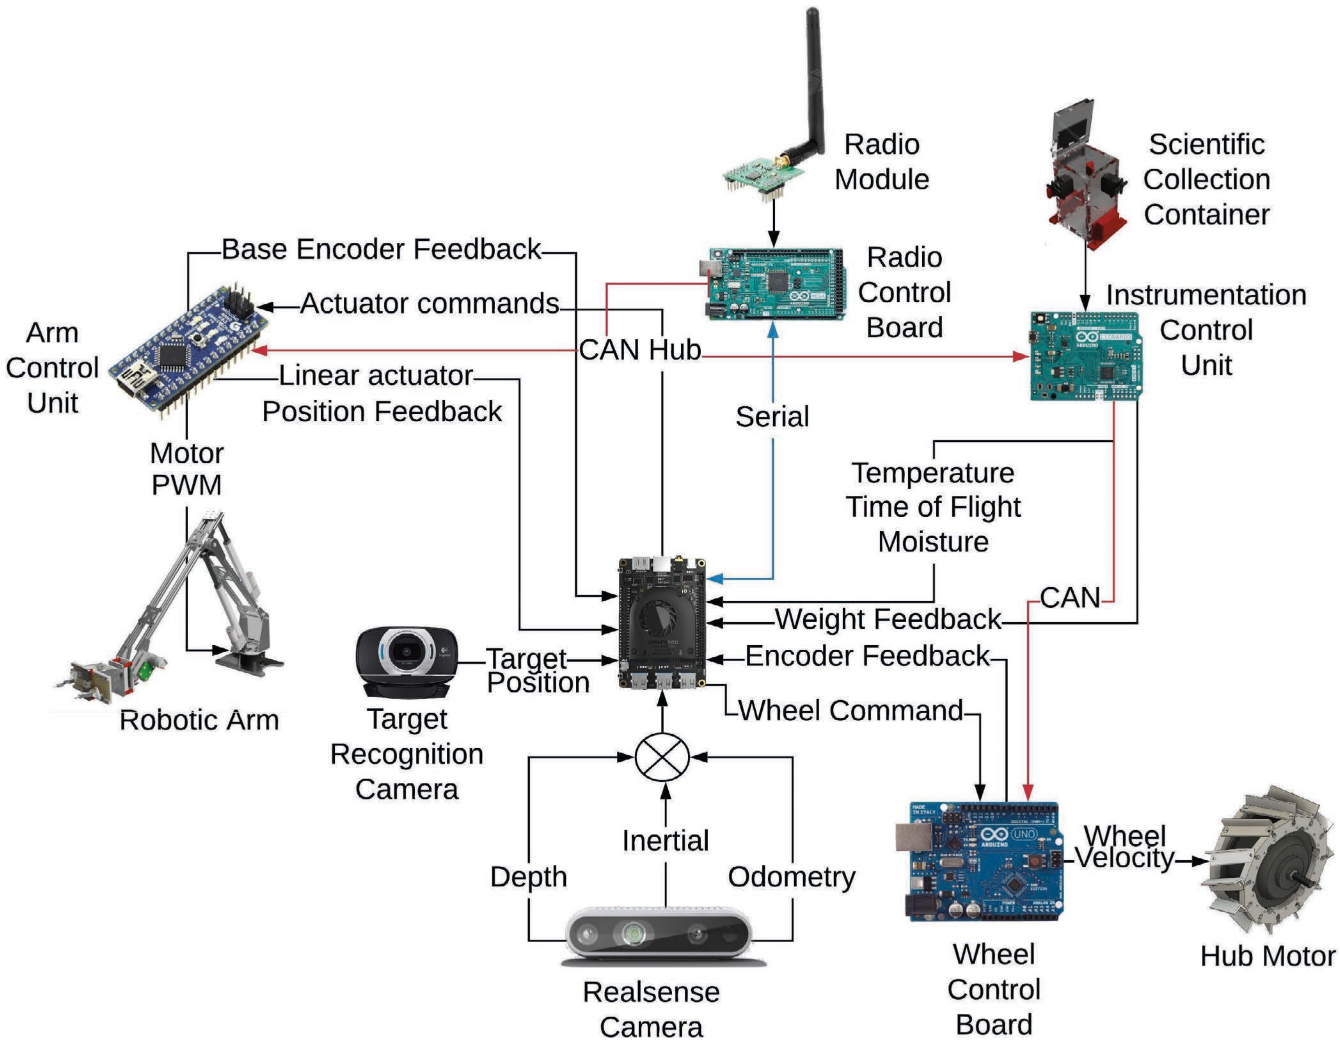

Figure 2
System breakdown structure of the main control instances.

the individual subsystems and the selection of the main onboard computer. Due to the computationally demanding nature of the image processing tasks involved in autonomous navigation and motion planning, the Advanced Reduced Instruction Set Computer (RISC) Machines architectures were considered unable to satisfy the mission requirements. As a result, a LattePanda Alpha development board was found more appropriate. The Alpha serves as the Rover Compute Element (RCE), being the highest authority in the control hierarchy of the system. Figure 2 shows the final control architecture of the MarsWorks rover.

Secondary computing units were designed around popular microchip controllers (ATmega32u4, ATmega328, and ATmega328p [5]) to handle specialized tasks such as motor control, radio communication, and signal processing. This way, the computational load on the RCE is reduced by implementing real-time operating systems together with custom scheduling algorithms on each microcontroller. It also guarantees that no control task will be interrupted by lower priority tasks, hence avoiding instability. These secondary computers are: four Instrumentation Control Units which process the sensor measurements of the collected soil sample (mass, temperature, humidity, and time-of-flight), four Motor Control Boards to run the Proportional–Integral–Derivative (PID) speed control algorithms for each wheel, one Arm Control Board to drive the six arm actuators, one Radio Controller, and one Emergency Stop Controller. The main protocol used to communicate between microcontrollers was Controlled Area Network (CAN), which is the preferred standard at ESA and the other space agencies for most spacecraft data handling applications due to its decentralised architecture.

The framework used to coordinate all the processes running on the craft is the Robot Operating System (ROS) [6], hosted on the RCE. ROS was considered the ideal middleware between low level motion control and autonomous navigation or guidance since it provides built-in capabilities for fusing sensor measurements coming from different subsystems of the rover, state estimation, visual odometry, as well as multiple open-source packages for interfacing with the RealSense cameras or even Arduino. For example, rough terrain traversal is achieved by having the RCE analyse the output of two infrared and depth cameras to detect the next site to be reached as well as determining its own position and pose (navigation). Next the RCE analyses the picture frames for clear paths towards the target, loops through the results to select the shortest one, and computes the trajectories needed to get there (guidance). Finally, these values are transmitted to the individual wheel controllers; they are then converted to Pulse-Width Modulated values which are sent to the direct current motors to drive the wheels (control), all while the same motor controllers read the encoder outputs from the motors to ensure that the desired speed was achieved without offset (PID).

Due to the nature of the competition tasks, the arm and locomotion dynamics are relatively decoupled for most path planning scenarios. Consequently, the detailed analysis will follow each subsystem individually on a component level.

NAVIGATION AND LOCOMOTION CONTROL SYSTEM

The main task to be completed by the navigation stack is autonomous traversal: finding and following the shortest clear path leading to the previously localised target in the odometric frame. To achieve that, depth and stereo data from the RealSense cameras was used to create a two-dimensional occupancy grid (local costmap) of the immediate environment. Then, the costmap is searched for clear paths towards the target using the Dynamic Window Approach algorithm [7], which is traditionally used in robotics for collision avoidance. Then, the algorithm loops through the results found and selects the shortest path. Once the trajectory corresponding to the chosen path is computed, ROS will generate the control actions required to drive the wheels and guide the craft along the path. The control signal is transmitted from the RCE to the Radio Controller through serial, and from there to the individual Wheel Controllers through CAN bus.

For navigation purposes, Visual-Inertial Odometry (VIO) was the preferred approach over Simultaneous Localization and Mapping (SLAM) due to the ease of incorporating wheel-encoder feedback into the visual position data, which together with the displacement estimates obtained from the Inertial Measurement Units (IMUs) provide a very robust system that accounts for slippage and many other uncertainties in hazardous terrains [8]. Furthermore, VIO also proved to be more computationally efficient than SLAM especially in scenarios when the navigation algorithms have to run in parallel with the arm motion planning within ROS. All these contributed to VIO being the preferred navigation methodology for the NASA Mars Exploration Rovers, where it delivered unprecedented performance (97% convergence on Spirit and 95% on Opportunity) [9].

One of the fundamental questions posed by VIO is estimating the distance travelled by the vehicle from the origin of the odometric frame. This is accomplished within ROS by feeding all the measured displacements (encoder ticks, doubly integrated accelerations from the IMU, and stereo camera distances) to a state estimation algorithm (the Extended Kalman Filter function), which weights and fuses the measurements according to the reliability gains specified in the measurement covariance matrix. The output of this function is an optimum distance estimate.

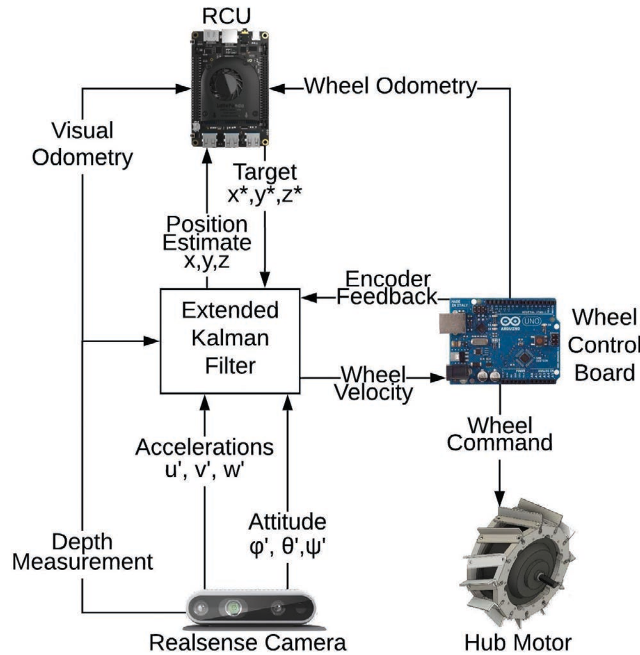

Figure 3
Navigation system—main components.

A major goal for the next design iteration is being able to incorporate a multirotor unmanned aerial vehicle (UAV) to the rover, which would be able to perform extensive mapping of the environment and relay the data in real time to the rover. Using this approach, the rover would become capable of planning the traversal task potentially kilometres ahead and learn about the feasibility of certain routes way advance, therefore making the most out of its time in commission. This would arise fascinating challenges both in terms of translating an aerial perspective to a terrestrial planner but also from the point of view of autonomous docking of the UAV with the rover for recharging, which is a great opportunity to experiment with model predictive control approaches.

ONBOARD SENSOR DATA FUSION

Sensor data fusion [10] as a process of knowledge extraction from multiple sensors plays an important role in the navigation of the Mars Rover. The data from multiple cameras are fused with data from other sensors, such as encoders. As part of the scientific task, basic properties of the collected soil sample have to be measured (temperature, density, humidity), which inevitably involves the fusion of signals from very different transducers, but also filtering, interpolation, or downsampling to enable the variables to be manipulated together in calculations. One basic example is the density determination, which involves combining the readings from a load cell (mass) with the time of flight data (volume occupied); therefore, the sampling rates of the two had to be matched and different low pass filters had to be implemented according to the known transducer dynamics.

For the simulated Martian environment in which the rover was designed to operate, the collected sand was stored in three identical scientific containers equipped with sensors that measure the temperature, humidity, proximity/time-of-flight, and weight of the sample. Each container is equipped with an Instrumentation Control Board, which performs sampling, signal amplification, analog-to-digital conversion (especially for the low voltages coming from the load cells), as well as preliminary filtering before transmitting the data through CAN bus.

Two types of moisture-sensing devices were present onboard, this form of redundancy offering the opportunity to apply more advanced averaging and voting techniques. Namely, the DH11 temperature sensor [11] also provides relative humidity measurements; however, these are not as reliable as the ones returned by the dedicated Grove NE555DR [12] capacitive moisture sensor. As a result, the dual sensor was given a reliability weight of 30% as compared to a weight of 70% for the dedicated moisture sensor. It was also observed that upon collection, the moisture measurement provided by DH11 sensor presented a settling time of about five seconds, hence a moving average filter was included in the loop as a smoothing mechanism for the initialisation period. To further improve the estimation accuracy of the measurements, a simple averaging-and-voting system was implemented. This was based on the assumption that for a given digging site, no radical leap in humidity should occur within a depth of five or less centimetres; therefore, in the event that any of the three containers returns a humidity value 10% or more greater than the other two, the measurement should be excluded from the estimation and the other two are averaged instead.

PID control was also implemented for collection to ensure that the amount of soil poured by the arm into the containers matches the perfectly desired quantity. To achieve this, the weight measurement is used as a feedback signal for the robotic arm upon releasing the sample. It was therefore critical to ensure that no spikes are present in the load cell signal (which often proved to be the case during the first seconds of auto-calibration) as those would lead directly to violent control actions in the robotic arm. To account for that, a low pass filter with a cut-off frequency of 0.5 Hz was introduced in the loop. Low pass filtering is usually resorted to in industrial applications for noise reduction in signal processing or for lowering pixel contrast in image editing, however for the given application it also helped with smoothing the sharp transitions occurring at auto-calibration.

Handling the proximity (time of flight) data was comparatively straightforward. The proximity sensor was used to determine the distance from the lid to the sand surface inside the container, therefore ON-OFF control was implemented to receive distance data only once the lid was closed. Having the distance to the sand surface, one can determine the volume occupied by the sample, and together the measured mass the average density of the collected soil can be estimated.

The decentralised architecture of the data handling system allows for modularity and communicational efficiency, the data being digitised and filtered locally, therefore only relayed through the network once clean and reliable. One immediate advantage is that signal attenuation and contamination are avoided, no analog signal having to traverse a region of the vehicle with potential for electromagnetic interference before analog-to-digital conversion. This approach also reduces the computational load on the RCE and ensures that compact data packets are being transmitted to other subsystems, which is both bandwidth-efficient for transmission and prevents control hazards arising from raw signal contamination. As a result, the scientific data leaves the Instrumentation board already at Data Processing Level 3 [13], before reaching other segments of the rover or the ground station.

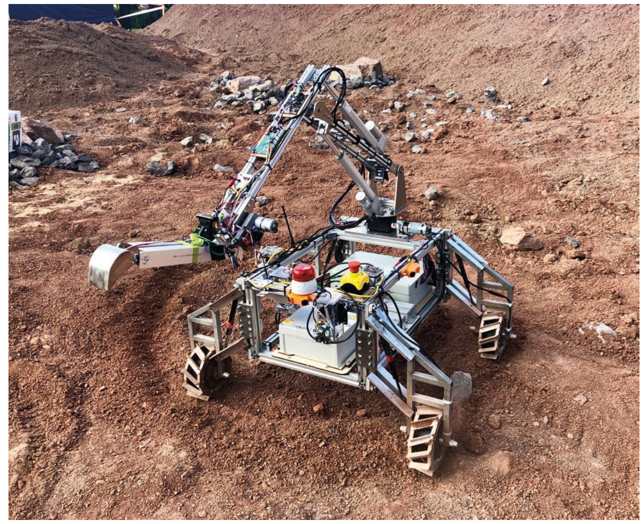

Figure 4
The MarsWorks rover in excavator configuration at ERC2019.

REFLECTIONS AND CONCLUSIONS

As access to space becomes more affordable and launch frequency grows due to space commercialization, ground-in-the-loop guidance of spacecraft will become prohibitively expensive because of scheduling conflicts but also increases in maintenance and labor costs. This will only contribute to an increasing probability of a human error. Automation can prevent such outcomes, enabling greater numbers and types of missions to operate concomitantly while improving robustness, reducing risk, and hence increasing future commercial and scientific return from space.

The ERC prepares young engineers for the challenges posed by the next generation of Mars Rovers, which are expected to have an increasingly high scientific pay-off. The competition itself is akin to a test at an analogue site, achieving striking similarity with the actual martial terrain. The development of autonomous navigation, manipulation, sample collection, and return are the core technology gaps to be addressed for these sample retrieval missions. The benefits of development in autonomous space technologies would also spill into other industries, for example a major potential for these types of rovers would be in farming and agriculture. Most noticeably, agriculture and food production are important sectors for any country, particularly the UK. The UK-Robotics and Autonomous Systems network stated that the Agri-Food industry is the largest manufacturing sector in the UK. This industry has a high manual labour demand which results in low productivity. The UK government recently committed to investing £90 m to boost productivity through automation and process monitoring. A key capability for agricultural robotic vehicles is that they must be able travel on uneven terrain, without damaging crops, which parallels the requirements for Mars Rovers. In the future, more investment is to be expected in this area due to the reduction in the number of manual labour workers and rising wages.

Although building a Mars Rover proved to be an exceedingly rigorous and demanding technical endeavour, some of the most valuable lessons learnt in the process concern the human aspect of the project, revealing that the transfer of knowledge between teams, good communication structures within the organisation, and having dedicated people for systems integration are just as critical to project success as the science and engineering underpinning it. The key lessons generally agreed by the MarsWorks team after ERC2019 are:

- Testing and verification should be carried at every stage as opposed to only happening at the end of the development process.
- Systems integration should happen as soon as possible in order to avoid collisions or incompatibilities close to delivery.
- Independent external reviews from separate University-based teams are a great opportunity to exchange good practices and motivate better documentation.
- Every two subteams should share at least one member. This promotes cohesion between subsystems and makes the management aware of problems at the interface between technical areas or subteams.
- Team members with a background in systems engineering and general knowledge of each technical area tend to be the most effective leaders.

ACKNOWLEDGEMENTS

Funding: MarsWorks was made possible with support from the Automatic Control and Systems Engineering and Mechanical Engineering Departments together with the School of Mathematics and Statistics from the University of Sheffield as well as our Alumni Fund, KE Fund, EEE, AMRC Sheffield, and the UK Mars Society.

Author Contributions: L.M. prompted the motivation for the study and provided invaluable comments on the interpretation of the sensor data. V.F. coordinated the overall project and contributed to the production of the manuscript text. V.C. analyzed the data and wrote the first draft of the manuscript. J.F. presented the management strategies and development process involved in designing and manufacturing the rover. H.K. covered the computer vision methodology applied for cache recognition. M.T. provided documentation on the navigation algorithms developed for the traversal task.

Additional thanks are extended to Dr. Li Sun, Dr. Gary Verth, and Dr. Istvan Ballai for helping establish the scientific foundation on which the system was subsequently designed and engineered.

REFERENCES

1. United States National Research Council. Planetary science decadal survey. Available: https://solarsystem.nasa.gov/science-goals/about/, last access Mar. 2020.
2. National Aeronautics and Space Administration. Solar System Exploration, Missions. Lunakhood 01. Available: https://solarsystem.nasa.gov/missions/lunokhod-01/in-depth/, last access Oct. 2019.
3. National Aeronautics and Space Administration. Mars exploration rover history. Available: https://mars.nasa.gov/mars-exploration/missions/mars-exploration-rovers/, last access Oct. 2019.
4. European Space Agency. ExoMars. Available: https://www.esa.int/Our_Activities/Human_and_Robotic_Exploration/Exploration/ExoMars/ESA_s_Mars_rover_has_a_name_Rosalind_Franklin/, last access Oct. 2019.
5. Microchip. ATmega328, ATmega328P, ATmega328PBp. Available: https://www.microchip.com/wwwproducts/en/ATmega328, last access Oct. 2019.
6. Robot Operating System. ROS kinetic. Available: http://wiki.ros.org/kinetic, last access Oct. 2019.
7. Fox, D. The dynamic window approach to collision avoidance. *IEEE Robotics & Automation Magazine*, Vol. 4, 1 (1997), 23–33.
8. Pavone, M. Spacecraft autonomy challenges for next generation space missions. *Lecture Notes in Control and Information Sciences*, Vol. 460, (2016), 1–34.
9. Maimone, H. *Two Years of Visual Odometry on the Mars Exploration Rovers*. Hoboken, NJ: Wiley, 2007.
10. Fourati, H. (Ed.). *Multisensor Data Fusion: From Algorithms and Architectural Design to Applications*. CRC Press, Taylor & Francis Group, Aug. 26, 2015.
11. Waveshare. DHT11 temperature-humidity sensor. Available: https://www.waveshare.com/wiki/DHT11_Temperature-Humidity_Sensor, last access Oct. 2019.
12. Grove. Capacitive moisture sensor. Available: http://wiki.seeedstudio.com/Grove-Moisture_Sensor/, last access Oct. 2019.
13. National Aeronautics and Space Administration. Data processing levels. Available: https://science.nasa.gov/earth-science/earth-science-data/data-processing-levels-for-eosdis-data-products, last access Oct. 2019.

FUSION CONFERENCE AWARDS

FUSION 2019 BEST PAPER AWARDS

Stefano Coraluppi
Systems & Technology Research LLC
Woburn, Massachusetts, USA
stefano.coraluppi@ieee.org

The 22nd International Conference on Information Fusion (FUSION 2019) was held in July 2019 at the Shaw Centre in Ottawa, Canada. FUSION is the flagship event of the International Society of Information Fusion (ISIF), and the conference is well established as the premiere forum to present and discuss research progress and initiatives in information fusion. This year, there were 403 attendees from around the world, with active participation from industry, government, and academia. The full report is available at http://isif.org/conferences/isif-conference-information.

Since its inception, ISIF has promoted a high-quality technical program at FUSION. One way to encourage this excellence is to promote the paper awards program. Accordingly, each year the conference includes recognition of the best regular papers and the best student papers. Student papers are those for which the lead author is a full-time graduate (or undergraduate) student at an accredited university. As mandated by the ISIF Board of Directors, the best paper receives the ***Jean-Pierre Le Cadre Award***. The best student paper receives the ***Tammy L. Blair Award***. These awards honor the efforts and commitment of both Jean-Pierre and Tammy to the international fusion community over many years.

The FUSION 2019 Awards Cochairs were Erik Blasch, Stefano Coraluppi, Ivan Kadar, and Mahendra Mallick. They began the selection process by examining the reviews of 266 papers by the Technical Program Committee led by the Technical Cochairs Anne-Laure Jousselme, Thia Kirubarajan, Henry Leung, Rakesh Nagi, and Roy Streit. To avoid the possibility of conflicts of interest, all papers coauthored by any FUSION 2019 Organizing Committee member were excluded from further consideration. Based on reviewer scores, the Awards Cochairs selected 15 regular and 15 student papers for detailed assessment. They conducted a thorough review and quantitative

Stefano Coraluppi describes the awards selection process and announces the winners (photo by Josephine Ding).

On behalf of Huajie Shao, Lance Kaplan receives the third-place best student paper award from Erik Blasch (photo by Josephine Ding).

JEAN-PIERRE LE CADRE AWARD

Vikram Krishnamurthy and Muralidhar Rangaswamy, "How to Calibrate Your Enemy's Capabilities? Inverse Filtering for Counter-Autonomous Systems"

Abstract—We consider the following adversarial Bayesian signal processing problem involving "us" and the "enemy": an enemy observes our state in noise, updates its posterior distribution of the state, and then chooses an action based on this posterior. Given knowledge of "our" state and sequence of enemy's actions observed in noise, we consider two problems: (i) How can the enemy's posterior distribution be estimated? Estimating the posterior is an inverse filtering problem involving a random measure. We formulate and solve several versions of this problem in a Bayesian setting. (ii) How can the enemy's observation likelihood be estimated? This tells us how accurate the enemy's sensors are. We compute the maximum likelihood estimator for the enemy's observation likelihood given our measurements of the enemy's actions, where the enemy's actions are in response to estimating our state. The above questions are motivated by the design of counter-autonomous systems: given measurements of the actions of a sophisticated autonomous enemy, how can a counter-autonomous system estimate the underlying belief of the enemy, predict future actions, and therefore guard against these actions.

scoring of these papers, leading to a set of seven regular and seven student papers for further analysis.

Subsequently, the Awards Cochairs defined two committees to examine the papers. Six committee members (Sanjeev Arulampalam, Daniel Clark, Mark Coates, Fredrik Gustafsson, Gustaf Hendeby, and Ruixin Niu) were asked to rank the seven regular papers. Similarly, six committee members (David Crouse, Ondřej Straka, Ángel García-Fernández, Qi Cheng, Pramod Varshney, and Suman Chakravorty) were asked to rank the seven student papers. No committee members were coauthors on any papers that they evaluated, and no conflicts of interest were identified. The sum of scores led to overall rankings that were ratified by the Awards Cochairs.

The best regular papers were the following:

- First place: Vikram Krishnamurthy and Muralidhar Rangaswamy, "How to Calibrate Your Enemy's Capabilities? Inverse Filtering for Counter-Autonomous Systems" (See text box for the paper's abstract.)
- Second place: David Crouse, "Particle Flow Filters: Biases and Bias Avoidance"
- Third place: Ángel García-Fernández and Lennart Svensson, "Spooky Effect in Optimal OSPA Estimation and How GOSPA Solves It"

David Crouse receives the second-place best paper award (photo by Josephine Ding).

Ángel García- Fernández receives the third-place best paper award (photo by Josephine Ding).

The best student papers were the following:

- First place: Benjamin Naujoks, Patrick Burger, and Hans-Joachim Wuensche, "Combining Deep Learning and Model-Based Methods for Robust Real-Time Semantic Landmark Detection" (See text box for the paper's abstract.)
- Second place: Erik Wilthil, Yaakov Bar-Shalom, Peter Willett, and Edmund Brekke, "Estimation of Target Detectability for Maritime Target Tracking in the PDA Framework"
- Third place: Huajie Shao, Shuochao Yao, Yiran Zhao, Lu Su, Zhibo Wang, Dongxin Liu, Shengzhong Liu, Lance Kaplan, and Tarek Abdelzaher, "Unsupervised Fact-Finding with Multi-Modal Data in Social Sensing"

These papers were recognized during the FUSION 2019 banquet dinner. Erik Blasch and Stefano Coraluppi announced the winners and presented award certificates.

The selection process to decide FUSION paper awards is an important stage that complements the larger paper-review process. The awards selection is conducted with great thoroughness, identifying research of significant value that is deserving of the attention of fusion researchers and practitioners. On behalf of ISIF, congratulations to the authors of all six papers for their hard work and impressive achievement!

TAMMY L. BLAIR AWARD

Benjamin Naujoks, Patrick Burger, and Hans-Joachim Wuensche, "Combining Deep Learning and Model-Based Methods for Robust Real-Time Semantic Landmark Detection"

Abstract—Compared to abstract features, significant objects, so-called landmarks, are a more natural means for vehicle localization and navigation, especially in challenging unstructured environments. The major challenge is to recognize landmarks in various lighting conditions and changing environment (growing vegetation) while only having few training samples available. We propose a new method which leverages Deep Learning as well as model-based methods to overcome the need of a large data set. Using red, green, blue (RGB) images and light detection and ranging (LiDAR) point clouds, our approach combines state-of-the-art classification results of Convolutional Neural Networks (CNN) with robust model-based methods by taking prior knowledge of previous time steps into account. Evaluations on a challenging real-world scenario, with trees and bushes as landmarks, show promising results over pure learning-based state-of-the-art three-dimensional (3D) detectors, while being significantly faster.

ISIF-Sponsored Events and Workshops

BELIEF FUNCTIONS AND APPLICATIONS SOCIETY RECENT EVENTS: CONFERENCES AND SCHOOLS

Anne-Laure Jousselme
NATO STO Centre for Maritime Research and Experimentation
La Spezia, Italy
anne-laure.jousselme@cmre.nato.int

Sébastien Destercke
Université de Technologie de Compiègne, and
Centre National de la Recherche Scientifique (CNRS)
sebastien.destercke@hds.utc.fr

Thierry Denoeux
Université de Technologie de Compiègne, and
Institut Universitaire de France
tdenoeux@utc.fr

The Belief Functions and Applications Society (BFAS) organizes biennial events alternating between conferences on even years and schools on odd years. The first edition of the conference was held in Brest, France in 2010, where both Glenn Shafer and Arthur Dempster in their respective keynote speeches were looking back on the origins and foundations of their now famous theory, also referred to as Dempster-Shafer theory, or evidence theory. Since then, the annual BFAS events gather between 30 and 80 researchers and practitioners of belief functions.

The biennial BELIEF conferences are a forum for the confrontation of ideas, the reporting of recent achievements, and the presentation of the wide range of applications of this theory. Since the first edition in Brest in 2010, the conference was held successively in Compiègne, France in 2012, in Oxford, UK in 2014, in Prague, Czech Republic in 2016, and in Compiègne, France in 2018.

BELIEF 2018: Conference attendees prior to the banquet at the Abbaye Royale du Moncel, Compiègne, France.

The biennial BELIEF schools offer a unique opportunity for students and researchers to learn about fundamental and advanced aspects of the theory of belief functions. The school is organized around a set of lectures by prominent researchers who gradually tackle basic to more advanced theoretical concepts and offer complementary tutorial sessions focused on the practical use and implementation of belief functions. The lectures also highlight the links with other uncertainty theories such as imprecise probabilities, random sets, or rough sets, and present some of the belief functions applications in various domains including, notably, information fusion, inference, and machine learning. The schools were successively held in Autrans, France in 2011, in Carthage, Tunisia in 2013, in Stella Plage, France in 2015, and finally in Xi'an, China in 2017. This last location and dates were chosen so that the event would be held in conjunction with the FUSION conference the week prior in a nearby venue to facilitate attendance to both events.

The fifth editions of the conference and school were held respectively in Compiègne, France in 2018 and in Siena, Italy in 2019. Their outcomes are reported below.

BELIEF 2018—5TH INTERNATIONAL CONFERENCE ON BELIEF FUNCTIONS

BELIEF 2018 was held in the Innovation Centre of the Université de technologie de Compiègne (UTC) in France from September 17–21, 2018. This fifth edition of the series of conferences on the theory and application of belief functions was the first edition to be held jointly with the International Conference on Soft Methods in Probability and Statistics (SMPS). This was the occasion to connect with another community and to discuss common problems and differences between the approaches privileged by each of them.

To foster and reinforce this connection between BELIEF and SMPS and truly have a joint event, a unique program committee was set up for all papers, and reviewers of both communities were encouraged to select and bid on papers irrespective of the venue to which they were submitted. This resulted in a program committee of over 55 people. Also, except for the most specific papers, all papers were presented in a single session, with papers from both conferences mixed in the same session. Participants really appreciated this initiative, especially when they could not tell which paper was from which conference.

The BELIEF conference alone received 39 submissions, of which 32 were accepted. Authors who submitted papers to the joint event were from 23 countries, most notably France (36), China (17), Spain (13), Italy (10), Belgium (9), Tunisia (8), and Czech Republic (8). Seventy (70) participants registered for the joint conference.

In the spirit of the joint event, tutorials and plenary talks were chosen to connect the different communities present. The first day of the conference was devoted to invited tutorials: “Statistics with Interval Data” by Scott Ferson, “Imprecise Markov Chains” by Jasper de Bock and Thomas Krak, and “Random Fuzzy Sets and Statistics with Imprecise-valued Data” by Maria Gil.

The plenary talks were the following: “On Statistical Modelling with Imprecise Probabilities” by Thomas Augustin, “Non-Laplacian Uncertainty: Practical Consequences of an Ugly Paradigm Shift about How We Handle Not Knowing” by Scott Ferson, and “Belief Functions and Valid Statistical Inference” by Ryan Martin. This last talk was given by an overseas invited speaker who was sponsored thanks to the ISIF grant. It provided an overview of very recent results showing that belief functions and confidence bands were ideal tools to deliver statistically guaranteed results, providing solutions to some of the issues of Fisher pivotal methods.

Awards for best papers were given during the banquet, which was the occasion for participants to discuss in a relaxed and playful atmosphere, with wooden games available. A shortlist of seven candidate papers, all from young researchers (as this was the first criterion to receive the award), was prepared by the Program Chairs on the basis of the reviewers’ scores and assessments. This list was later submitted to the BFAS Board of Directors for the selection of the best paper awards.

The awards, sponsored by Elsevier and the *International Journal of Approximate Reasoning*, went to the papers “Outer Approximations of Coherent Lower Probabilities Using Belief Functions” by Ignacio Montes, Enrique Miranda, and Paolo Vicig, and “An Evidential K-Nearest Neighbor Classifier Based on Contextual Discounting and Likelihood Maximization” by Orakanya Kanjanatarakul, Siwarat Kuson, and Thierry Denoeux. The first paper studies how sets of probabilities could be approximated by belief functions, while the second one revisits K-Nearest Neighbor classifiers under the light of likelihood estimation.

BELIEF 2018’s proceedings are published by Springer’s Lecture Notes in Artificial Intelligence/Lecture Notes in Computer Science series, Volume 11069.

BFTA 2019—5TH SCHOOL ON BELIEF FUNCTIONS AND THEIR APPLICATIONS

The fifth school on Belief Functions and Their Applications (BFTA) was held at the Certosa di Pontignano, in the vicinity of Sienna, in the province of Tuscany in Italy, from October 27–31, 2019.[1]

Thirteen lectures, accompanied with exercises provided by ten lecturers, succeeded each other over five days. The school opened with the talk of Thierry Denoeux on the “Introduction of Belief Functions”, providing basics of the theory. This was followed by a perspective on “Belief Functions and Boolean Inference” by Sébastien Destercke, addressing the basic combination of propositional logic and belief functions, together with computational challenges illustrated on some applications.

[1] bfasociety.org/BFTA2019

BFTA school 2019: School participants during the social event, after the visit of the Azienda Agricola Losi Querciavalle, Tuscany, Italy.

The second day started with the now famous lecture of Didier Dubois "Positioning Belief Functions among Uncertainty Theories", which illustrates the limited expressiveness of probabilities and provides an overview of alternative uncertainty theories, such as imprecise probabilities, belief functions and random sets, and possibility theory. Davide Ciucci complemented the portrait presenting the links between "Rough Sets and Belief Functions". Thierry Denoeux presented rational approaches for modeling evidence with belief functions in real world applications with "Methods for Building Belief Functions", capturing either expert opinions or statistical information. Frédéric Pichon gave the last lecture of Day 2 on "Information Correction and Fusion Using Belief Functions", tackling the specific problem of the consideration of meta-information on source quality (typically the reliability) in the combination of information from multiple sources. The second day closed with a poster session during which students had the opportunity to present their work and exchange with other school participants and lecturers. The posters were displayed until the end of the school, which stimulated discussions outside specific poster sessions and during breaks.

Day 3 initiated with the third lecture by Thierry Denoeux on "Classification and Clustering Using Belief Functions", which provided an overview of evidential classifiers illustrated with results on several applications. The practical session of Arnaud Martin on the "Implementation of Belief Functions" concluded the lectures of this short day, allowing the tutees to experience basic and advanced calculus with belief functions using the R library. The afternoon was a break and included social events, during which the group first had a guided visit of the vineyard farm close by, followed by a free visit to the city of Siena.

On Day 4, Prakash Shenoy presented "Graphical Models for Belief Functions", looking back at his work on Valuation-Based Systems, a generic framework abstracting several uncertainty *calculi* including belief function theory, probability theory, possibility theory, propositional calculus, and making the fusion process efficiently implemented thanks to local computation. Anne-Laure Jousselme then surveyed "Distances and Conflict between Belief Functions", highlighting the specific properties distinguishing between the two types of measures. Didier Dubois presented his second lecture on "Prejudiced Information Fusion Using Belief Functions", revisiting Philippe Smets' decomposition of belief functions to define an approach based on "diffidence" functions, as opposed to upper and lower probabilities. The last day, Barbara Vantaggi presented the use of "Belief Functions in Decision Theory" and Jean Dezert concluded the school with recent results on "Multi-Criteria Decision-Making Support with Belief Functions".

The BFAS general assembly was held at the end of the 4th day after a panel discussion, during which several ideas on how to render the material on belief functions more accessible were proposed, including a forum for discussion, recorded talks from lecturers, and short videos teaching basic and advanced calculus on belief functions. The next BELIEF conference was announced by Thierry Denoeux and will be held in be Shanghai, China in 2021, co-located with the 1st International Conference on Cognitive Analytics, Granular Computing, and Three-Way Decisions, cosponsored by the International Rough Set Society.[2]

The school gathered 34 attendees including lecturers, students, and senior researchers from 14 universities, research institutes, or companies from eight countries. The success of this event was greatly due to the ISIF, which covered lecturers' travel fees, as well as to the BFAS, who awarded grants covering registration fees, helping seven students travel from India, USA, Tunisia, Morocco, and France to attend the school.

BELIEF 2018 ORGANIZING COMMITTEE

Thierry Denoeux (Chair)
Sébastien Destercke (Chair)
Christine Leheutre
Mylène Masson
Benjamin Quost
David Savourey
Vu-Linh Nguyen
Yonatan Carranza Alarcon

BFTA SCHOOL 2019 ORGANIZING COMMITTEE

Anne-Laure Jousselme (Chair)
Thierry Denoeux
Nadia Ben Abdallah
David Mercier
Frédéric Pichon
Alessandro De Gloria
Roberto Sacile

[2] https://www.lgi2a.univ-artois.fr/events/belief2020/

ISIF-Sponsored Events and Workshops

IMPRESSIONS OF THE 13TH IEEE AESS SYMPOSIUM "SENSOR DATA FUSION—TRENDS, SOLUTIONS, AND APPLICATIONS" (SDF 2019), OCTOBER 15–17, 2019, BONN, GERMANY

Felix Govaers
Wolfgang Koch
Fraunhofer FKIE
Wachtberg, Germany
Felix.Govaers@fkie.fraunhofer.de
Wolfgang.Koch@fkie.fraunhofer.de

Initiated and organized by the Fraunhofer FKIE and with active participation of the international expert community, as well as sponsorship from the International Society of Information Fusion (ISIF), the 13th Symposium *Sensor Data Fusion: Trends, Solutions, and Applications* took place from October 15–17, 2019 at the Uniclub Bonn. Participants from industry, universities, and research institutes from overseas and Europe presented and discussed current developments in modeling, applications for data fusion, as well as theoretical findings.

Sensor data fusion and state estimation at the end of the second decade in this century was clearly influenced by *deep learning* and *artificial intelligence*. Its impact on data science, engineering, and even ethics and philosophy was immense, and the field of *information fusion,* as part of it, got some impulses for new applications, new methods, and better results. This was reflected in Sensor Data Fusion (SDF) 2019 with its 22 presentations in seven sessions, where the trend of hybrid methods that combine model-based tracking with deep learning approaches for classification, big data processing, anomaly detection, and intention prediction continued. In particular, a strong presence of automotive applications could be observed, which have a large variety in the challenges addressed: from novel methods for extended target tracking to

Keynote speech from Prof. Fredrik Gustafsson: "Fusion Theory for Positive Noise".

The coffee breaks were well received for continued discussions and networking.

robust navigation, multisensor fusion, and image segmentation with uncertainties.

Traditionally, the invited Plenary Talk was a special highlight of the conference. This year it was held by Prof. Fredrik Gustafsson, Linköping University of Technology, Sweden, one of the internationally leading protagonists in sensor informatics. He spoke on "Fusion Theory for Positive Noise". Positive noise arises, for instance, in localization problems based on time differences when multipath propagation is present. However, it has not gained much attention until now. Prof. Gustafsson presented multiple approaches to model positive noise and directly gave their optimal estimators to the audience. This sound technical approach with applications to real problems was well received by the participants of the symposium.

The high quality of the numerous submissions to SDF 2019 made it difficult to honor two particular papers with an award, which come with a 100€ book voucher for a popular online store. This time, the committee rewarded novel and exceptional consistent approaches to combine machine learning with statistical model based fusion: The Best Paper Award went to Mahed Javed and Lyudmila Mihaylova for their paper "Leveraging Uncertainty in Adversarial Learning to Improve Deep Learning Based Segmentation", and the Best Student Paper Award was received by Florian Particke, Jiaren Zhou, Markus Hiller, Christian Hofmann, and Jörn Thielecke for "Neural Network Aided Potential Field Approach for Pedestrian Prediction".

SDF 2019 was jointly organized by Wolfgang Koch, Fraunhofer FKIE/University of Bonn and Peter Willett, University of Connecticut, acting as Executive Cochairs. Technical Program Chair was Felix Govaers, Fraunhofer FKIE. There were about 60 participants with a diverse mixture of interested people from industry, academia, and research institutes. The audience was international (SWE, UK, AUS, USA, for instance) but the majority is from Germany due to the local attraction. The audience stays in a single group, which often results in interesting and feisty discussions. We will continue our series irregularly in 2021 with the upcoming SDF Symposium and hope to see many of our participants at the FUSION 2021 conference next year in South Africa.

More information on this and future SDF workshops can be found at http://fkie.fraunhofer.de/sdf2019.

Best Paper Award and Best Student Paper Award for Mahed Javed et al. and Florian Particke et al., respectively.

ISIF-Sponsored Events and Workshops

Carlo Regazzoni
Lucio Marcenaro
University of Genoa
Genoa, Italy
carlo.regazzoni@unige.it
lucio.marcenaro@unige.it

2019 IEEE SPS/EURASIP/ISIF Summer School on Signal Processing (S3P-2019), September 8–13, 2019, Arenzano, Italy

The 2019 Institute of Electrical and Electronics Engineers (IEEE) Signal Processing Society (SPS)/European Association for Signal Processing (EURASIP)/International Society of Information Fusion (ISIF) Summer School on Signal Processing (S3P-2019)/Signal Processing for Autonomous Systems (SP-AS) was the seventh edition of the event, technically cosponsored by the IEEE SPS via the Seasonal Schools in Signal Processing (S3P) initiative and by the SPS Italy Chapter. This Summer School was also the first Summer School related to the IEEE SPS Autonomous Systems Initiative (ASI). The ASI initiative is a volunteer activity of several members of the IEEE SPS who are trying to highlight the signal processing related aspects of autonomous systems. The school took place in Arenzano, Genoa, Italy, from September 8–13, 2019 and involved almost 40 students and 10 lecturers on topics related to signal processing and information fusion for autonomous systems. Several initiatives such as panel discussions, the best poster and demo award, and industrial days were organized to maximize the participation of students, and they have been extremely satisfied with the experience.

The Summer School SP-AS focused on providing an updated state-of-the-art of the most advanced signal processing and information fusion theories and techniques that are relevant for developing autonomous systems. Lectures focused on novel algorithms and technologies but also on in-depth reviewing of state-of-the-art for autonomous systems. Summer school participants had the opportunity to learn and study innovative algorithms and systems for artificial interaction and cognition, to apply these concepts for building real working autonomous systems components, and to cooperate with other students for designing self-aware related modules.

Key Organizers

The Summer School on autonomous systems was directed by Carlo Regazzoni, Full Professor, Senior Member IEEE and Lucio Marcenaro, Assistant Professor, Member IEEE.

The school was held under the framework of the IEEE SPS Italy Chapter. The Summer School was cosponsored by the EURASIP and the ISIF. The school was also supported by the Italian Group for Telecommunication and Information Technology (GTTI).

The Summer School Steering Board was composed of the following professors:

- Prof. Stefano Tubaro, Politecnico di Milano
- Prof. Mauro Barni, Università degli Studi di Siena
- Prof. Francesco G. B. De Natale, Università degli Studi di Trento
- Prof. Alessandro Piva, Università degli Studi di Firenze
- Prof. Giovanni Poggi, Università degli Studi di Napoli Federico II
- Prof. Riccardo Leonardi, Università degli Studi Brescia

The Venue

The program was held at a hotel in Arenzano Punta San Martino, located in a quiet and privileged position overlooking the sea, with beautiful views over the Gulf of Genoa and the Ligurian Riviera di Ponente in Italy.

Activities and Lectures

The Summer School proposed different types of lectures: standard lectures introduced novel signal processing and information fusion algorithms and techniques as enablers for autonomous systems; application lectures presented state-of-the-art self-aware autonomous systems, describing how their concepts can be applied in real scenarios.

Several events were organized for deep involvement of the students in the Summer School activities such as: a full day dedicated to the presentation and exhibition of students' research activities and a best poster and demo award, two panel sessions directly involving lecturers and students with discussions about signal processing for autonomous systems, a full industrial day, with presentations from company representatives, and social events when participants visited Genoa's historical city center and the aquarium.

Program

On Monday, September 9, the first lecture of the Summer School was given by Prof. Walter Kellermann about "Audio Processing for Autonomous Systems". Walter Kellermann has been Professor of Communications at the University of Erlangen-Nuremberg, Germany since 1999.

In the afternoon of the first day, Prof. Ioannis Pitas from the Department of Informatics of the University of Thessaloniki, gave a lecture titled "Deep Learning for Multiple Drone Vision Systems".

On September 10, Prof. Karl Friston discussed "Bayesian Mechanics and the Free Energy Principle". Karl Friston (Scientific Director of the Wellcome Trust Centre for Neuroimaging at the Institute of Neurology, University College London)

The view from the hotel Punta San Martino in Arenzano, Genoa.

is a theoretical neuroscientist and authority on brain imaging. He invented statistical parametric mapping (SPM), voxel-based morphometry (VBM), and dynamic causal modelling (DCM). On the same day, Letizia Marchegiani from Aalborg University gave a lecture about "Sound as an Exteroceptive and Proprioceptive Sensing Modality for Autonomous Driving". A panel discussion about "Representation and Inference in Cognitive Systems: Human and Artificial Autonomous Agents" concluded the second day of the School.

Wednesday, September 11 was the SPS Italy Chapter day and highlighted the IEEE Distinguished Lecturer Prof. Anna Scaglione, a Professor at Ira A. Fulton Schools of Engineering, School of Electrical Computer and Energy Engineering, Arizona State University, Tempe, Arizona (USA). Her lecture was titled "Distributed Learning and Signal Processing Algorithms" and discussed how artificial intelligence today is about developing the capability of a single node to make an inference or to respond to its surroundings with the appropriate action.

On the same day, a poster and demo session was organized in which all the participating students were given the opportunity to present their research. After the evaluation phase, Ali Krayani (University of Genoa), Sara Baldoni (Roma Tre University), and Mohamad Baydoun (University of Genoa) received the third, second, and first prizes, respectively. The first prize of 500€ was fully sponsored by the International Society of Information Fusion. On Wednesday afternoon, the social event was organized, with a guided tour of the aquarium in Genoa and the social dinner at Grand Hotel Savoia.

Prof. K. Friston lecture on "Bayesian Mechanics and the Free Energy Principle".

Thursday, September 12 was the industrial day: three researchers from industry described their research experiences in the autonomous systems field. Gian Luca Mariottini, Principal Member at Draper Laboratories, Cambridge, Massachusetts (USA) gave a lecture titled "Assistive Autonomy". Stefano Coraluppi, Chief Scientist at Systems & Technology Research (STR), Woburn, Massachusetts (USA) discussed "Advances in Multi-Target Tracking". Finally, Alfonso Farina, a consultant from Italy, described "40 years of Cooperation Between Industry and Academia on Radar Tracking Systems".

On the morning of Friday, September 13, there were two "application lectures" by Antoine Deleforge, Research Scientist at Inria Nancy, France and David Martín Gómez, Universidad Carlos III de Madrid, Spain. The titles of the two lectures were "Taking the Best of Physics and Machine Learning in Robot Audition" and "Intelligent Transportation Systems: From Environment Understanding to Autonomous Vehicles", respectively.

IMPACT AND FEEDBACK FROM PARTICIPANTS

The students who attended the Summer School developed the capability to evaluate, understand, design, and develop technologies for self-aware autonomous systems. They were trained in the design and implementation of autonomous systems where users are usually involved in everyday tasks. As in the previous editions of the School, students were informed at the beginning of the event about the initiatives of the IEEE SPS Italy Chapter, IEEE SPS, EURASIP, and ISIF, and how to be active members in these associations. The fees of the school provided significant advantages for SPS, EURASIP, and ISIF registered students and nonstudents. Moreover, IEEE-SPS Italy Chapter Best poster and demo award was granted to the top student presenting research so as to enhance their feeling of being an active part of the Society. All lectures have been recorded and made available on the web.

Feedback from the 38 participants was, in general, extremely positive. The overall satisfaction for the Summer School was 3.5/4, and 93% of the attendees will be interested in attending a follow-up Summer School next year on a related topic. Free comments from the participants were grouped into strengths, weaknesses, and suggestions. In general, comments about the overall organization and venue were very positive. Most of the weaknesses and suggestions were related to the poster and demo session (too short and practically impossible for students to see other participants' contributions) and panel sessions (should have been organized in a more interactive way). For the future editions of the school, a few more "hands-on" activities should be considered as well, to give participants the opportunity to try experimentally the illustrated theoretical approaches.

Other Events and Workshops

MARITIME SITUATIONAL AWARENESS WORKSHOP (MSAW'19), OCTOBER 7–9, 2019, LERICI, ITALY

Anne-Laure Jousselme*
Paolo Braca
Maximilian Zocholl
Leonardo Millefiori
Francesca de Rosa
NATO STO Centre for Maritime Research and Experimentation
La Spezia, Italy
* anne-laure.jousselme@cmre.nato.int

Maritime Situational Awareness (MSA) supports effective and efficient decision-making and enables maritime operations to preemptively identify emerging safety, security, or environmental issues so that a timely intervention is possible. To reach a common and comprehensive understanding of the maritime operational environment, accurate, timely, and standardized information needs to be shared among nations, partners, and civilian agencies, providing the required information superiority to successfully conduct maritime operations. Understanding the maritime situation enables decision makers and emergency responders to focus on relevant events, to prevent malevolent acts, to minimize the impact of a possible threat, and/or to intervene in a timely manner. MSA highly depends on the ability of sensing, collecting, and processing technologies to handle the big data challenges brought by the ever-increasing volume, velocity, and variety of data, which often lack veracity. In this perspective, the achievement of MSA requires a multi- and interdisciplinary approach, spanning several fields including but not limited to sensing technologies, signal processing, data fusion, unmanned systems, machine learning, big data, artificial intelligence, and applied human factors.


Villa Marigola in Lerici.

Between October 7–10, 2019, the North Atlantic Treaty Organization (NATO) Science and Technology Organization-Centre for Maritime Research and Experimentation (STO-CMRE) hosted the Maritime Situational Awareness Workshop 2019 (MSAW'19) at the Villa Marigola in Lerici, La Spezia (Italy). The aim of this workshop was to present and discuss advanced technologies, innovative concepts, and emerging scientific challenges with respect to current and future MSA operational needs. Under the theme *Science and technology meet operational needs*, the MSAW'19 thus aimed at encouraging engagement with operational experts and scientists from national governments, military, academia, and industry to discuss their respective challenges regarding MSA. The objective of MSAW'19 was then to foster the cross-fertilization of ideas from scientific and military domains, toward the design and implementation of future solutions tailored to MSA operational needs.

The MSAW'19 brought together about 170 participants from 23 countries, including 18 NATO nations and 14 nations from the European Union (EU). All five continents were represented and brought together scientists, engineers, researchers from scientific communities with national and international authorities, end users, operators, and industrial representatives: 14% from academia, 39% from applied research institutes, 31% from industry, and 16% from the operational community. The technical program mixed highly technical scientific contributions with implementation perspectives from industry and standpoints from operations, offered as 42 oral presentations, 12 posters, and six technical booths.

MSAW'19 kicked off with both a welcome and opening session chaired by Dr. Catherine Warner, Director of CMRE. Four keynote talks were provided by distinguished experts of radar signal processing, information fusion, target tracking, and data fusion: Dr. Alfonso Farina and Sergio Gallone, Leonardo S.p.A., and Massimo Comparini, e-GEOS Chief Executive Officer, Telespazio Head of Line of Business Geoinformation, shared the floor for opening the workshop with the presentation entitled "Maritime Surveillance: Radar Technologies and Scenario Characteristics". Dr. James Llinas, Emeritus Professor at the State University of New York at Buffalo, in Buffalo, NY, presented remotely in the afternoon of the first day about "Re-Examining Fusion-Sensemaking-Decision-Making Interdependencies Again". Dr. Felix Govaers, Deputy Head of Department Sensor Data and Information Fusion at Fraunhofer In-

Workshop participants at MSAW'19.

stitute FKIE, in Bonn, Germany, keynote speaker of the second day, talked about “Push and Pull in Digitalization: Technology Drivers for Sensor Data Fusion”. Mr. John Waterston, Program Manager in the Strategic Technology Office at the Defense Advance Research Program Agency, US, concluded the keynote talks on the last day, presenting “Ocean of Things”. Four projects funded by the European Commission were presented: “RAdars LoNG Distance Maritime SurveillancE and SaR Operations” (RANGER), cosponsor of this workshop, “MARitime Integrated Surveillance Awareness” (MARISA), “Coordination of Maritime Assets for Persistent and Systematic Surveillance” (COMPASS2020), and Arctic and North Atlantic Security and Preparedness Network (ARCSAR), presented remotely from Iceland.

Alfonso Farina (center), Sergio Gallone, and Massimo Comparini, who presented on “Maritime Surveillance: Radar Technologies and Scenario Characteristics”.

On the second day of the workshop, Fabio Marziani of the NATO Communications and Information Agency (NCIA) chaired a session on NCIA’s 2019 Defence Innovation Challenge that focused on challenges of the High North including improvements to MSA capability. The winning entry was “Dual-use of AIS Data, Combining AIS Tracking with Social Network Analysis for Increased Maritime Network Awareness” submitted by the US Navy Post Graduate School and Norway’s defense research agency Defence Research Establishment (FFI). A brief on the project was presented by Dagfinn Vatne of FFI. The concept presented was to combine historical and live automatic identification system (AIS) data with social network analysis (SNA), in order to identify and geolocate suspicious actors. Other NCIA Defense Innovation Challenge winners were also present at the workshop and presented several posters on their work.

The workshop closed with an expert panel animated by Dr. Sandro Carniel (Head of Research, CMRE) with interventions of Cpt. Jehan-Christophe Charles (Ret., French Navy), Cdr. Jorge Martinez (Spanish Navy, NATO Combined Joint Operations from the Sea Centre of Excellence), Lt. Cdr. Ivo Musulin (HRV Navy, NATO Shipping Centre, MARCOM), and Guy Thomas (DCSA, MBA, Advisor for Maritime Situational Awareness, Multinational Maritime Security Centre of Excellence). Experts provided their views on the most promising technologies to support MSA, the topics where more research effort should be allocated, and the main challenges for MSA.

The MSAW’19 was cosponsored by the EU Horizon 2020 project RANGER and NATO Allied Command Transformation (ACT) as part of the CMRE Data Knowledge Operational Effectiveness (DKOE) project.

Papers and presentations are available at www.cmre.nato.int/msaw-2019-home.

Causal Inference in Statistics: A Primer

Judea Pearl, Madelyn Glymour, and Nicholas P. Jewell
Wiley, 2016
ISBN: 978-1119186847

Lawrence Stone
Metron, Inc.
Reston, Virginia, USA
stone@metsci.com

If you have ever needed to determine causal effect without performing a randomized controlled trial or if you are tired of hearing statisticians say, "correlation is not causation", then I highly recommend that you read *Causal Inference in Statistics: A Primer* [1] by Judea Pearl, Madelyn Glymour, and Nicholas P. Jewel.

Determining the causal effect of an action is something we do instinctively every day. We know from experience that putting a lit match to a piece of paper will cause the paper to burn. If we do not eat, we will get hungry. When the temperature falls below 32 °F outside, ice may form. We determine causes to guide our actions. For example, determining that malaria is caused by a mosquito rather than "bad air", as its name implies, tells us to use mosquito netting to avoid malaria rather than a gas mask. Knowing that smoking causes lung cancer, we can reduce our risk of getting cancer by not smoking.

Determining causal relationships between actions and results allows us to make intelligent decisions about the actions we should take to avoid risks, improve our health, or safeguard the health of our planet. In many cases, such as the match and the burning paper, the causal relationship is clear. In many others, such as smoking and lung cancer or global warming and extreme weather events, the causal relationship is not so clear.

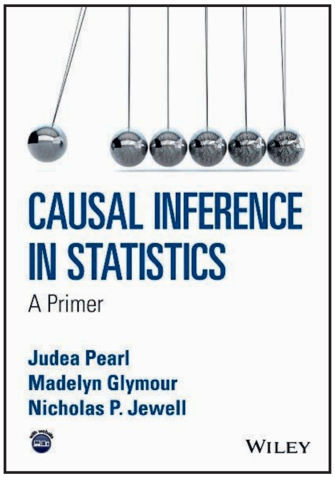

Judea Pearl is one of the leading developers of the theory of causality, along with Donald B. Rubin [2] and James M. Robins [3]. Pearl is also the recipient of the Association for Computing Machinery's Alan Turing Award for fundamental contributions to probabilistic and causal reasoning. *Causal Inference in Statistics: A Primer* is, in effect, a textbook for a first course in causal inference, complete with study questions (problems) whose answers are available from an instructor's companion website. In level of technical difficulty, this book lies between [4], which is described as a comprehensive exposition of the modern analysis of causation and [5], which is a popular science presentation of Pearl's theory of causal inference. We have used the primer for a study group at Metron (Reston, VA) in which we worked our way through most of the chapters and sections, discussing them and presenting solutions to some of the study problems. Although the presentation in this book is elementary and requires little background, we found it requires a lot of effort to understand the definitions of causality and the methods presented for performing causal inference. Nonetheless, this is an important and developing extension of statistical inference which will become an increasingly significant area of statistical analysis. The well-educated statistician or analyst should at minimum understand the concepts of causal inference and ideally be able to perform causal analysis.

Before reviewing the book, I present some background on causality to introduce the reader to this subject and provide some understanding of the long struggle to develop a satisfactory definition of causality and methods for performing causal inference.

BACKGROUND

Neither classical nor Bayesian statistics provide methods for determining or estimating causality. Statistical methods can estimate correlation, but as we are continually reminded, correlation is not causation. A classic example of this is the data Francis Galton collected on the heights of fathers and sons [6]. He determined that every extra inch in height of the father produced (on average) an extra half-inch of height in the son. He called this relationship the *correlation* between the height of the father and that of his son. This type analysis was later formalized by Karl Pearson into a mathematical method for computing the slope of a (properly rescaled) regression line. Pearson called this slope the correlation coefficient. A peculiar feature of this correlation is that it goes both ways. Tall dads tend to have tall sons, and tall sons tend to have tall dads. The correlation coefficient is the same both ways. Which is the cause and which the effect? Correlation and statistics have no way of telling us. However, we know that it is the father's extra height that tends to produce a taller than average son, not the other way around. We know this because we have in our mind a simple causal model of inheritance which says that the father's height is a cause of the son's height, not the other way around. What Galton's analysis does for us is to quantify the causal relationship. The notions of causality, causal model, and statistical estimates of causal effects have been developed by Pearl and his colleagues into a methodology which allows us to estimate, in certain circumstances, the causal effect of an action on an outcome. Before pressing on to explain Pearl's definition of causality, I briefly review the history of attempts to define causality. Kleinberg and Hripcsak [7] provide an excellent overview of causal inference and various definitions of causality from the perspective of bioinformatics. Researchers in econometrics, bioinformatics, and epidemiology have been in the forefront of the development and use of causal inference.

DEFINITIONS OF CAUSALITY

REGULARITY DEFINITIONS

In 1739, David Hume proposed the regularity definition of causation which says, in effect, if one type of object (say a flame) always produces a second type of object (say heat), the first object (flame) is the cause of the second (heat). Of course, the difficulty with this definition is that it is really defining a correlation not causation. In 1748, Hume [8] amended his definition to read:

> *We may define a cause to be an object followed by another, and where all the objects similar to the first, are followed by objects similar to the second. Or in other words, where if the first object had not been, the second never had existed.*

This definition depends on the notion of counterfactuals—"where if the first object had not been". That is, it depends on imagining the result of something that did not happen, i.e., a counterfactual. In this definition, Hume assumes that people intuitively understand counterfactual reasoning and that it does not need to be defined. In fact, people often intuitively perform counterfactual reasoning to determine cause and effect in their everyday lives. In subsequent years, most people who considered the question of causality ignored the second sentence in Hume's 1748 definition and concentrated on the first sentence, the regularity part.

INUS CONDITIONS

In many cases there may be multiple factors that produce an effect. Mackie [9] produced an updated version of Hume's regularity definition that allows for multiple causes. He defined a cause as some condition that is perhaps *Insufficient* by itself to produce the effect but is a *Nonredundant* part of a set of conditions that may be *Unnecessary* but are *Sufficient*. These are termed the INUS conditions.

BRADFORD HILL CRITERIA

In 1965, the English statistician Bradford Hill [10] proposed a set of nine criteria to provide epidemiologic evidence of a causal relationship. These criteria were used to demonstrate the connection between cigarette smoking and lung cancer. At one time, these criteria were widely accepted as useful for identifying causal relationships in epidemiological studies. However, a problem with the use of these criteria is that many of them rely on judgment rather than scientific verification.

PROBABILISTIC CAUSALITY

One of the difficulties with the above causality definitions is that they are deterministic. Specifically, they do not allow us to determine quantitatively what fraction of the effect is due to each cause. Probabilistic theories of causality [11], [12], and [13] have been proposed to deal with this problem. The basic idea of these theories is that a cause raises the probability of and occurs before its effect. The condition that a cause C raises the probability of an effect E is defined using conditional probabilities as follows:

$$P(E \mid C) > P(E). \tag{1}$$

The difficulty with this definition is that the conditions of the cause being prior to the effect and the relationship in (1) being true are neither necessary nor sufficient for a causal relationship. A classic example is a falling barometer and rain. The falling barometer occurs before the rain and may be seen as increasing the probability of rain, but it is actually the decreasing air pressure that causes both.

GRANGER CAUSALITY

This definition of causality is usually applied to time series. The approach attempts to find if one variable (coupled with the appropriate time lag) is informative about another. Specifically, let W_t represent the knowledge that is available at time t. Then the time series X at time t is said to be a Granger-cause of the time series Y at some time $t+s$, where $s > 0$ if

$$P\left(Y_{t+s} \mid W_t\right) \neq P\left(Y_{t+s} \mid W_t - X_t\right) \tag{2}$$

where we use $W_t - X_t$ to mean the information contained in W_t with that in X_t removed. The inequality in (2) indicates that X_t contains some information about Y_{t+s} that is not in the rest of the set W_t. Although Granger causality may be useful for predictions, it is not suitable for causality or explanation. As an example, consider that smoking causes both lung cancer and stained fingers, and that the stained fingers usually occur before the cancer. However, we cannot prevent lung cancer by wearing gloves when smoking. The primary type of error that Granger causality produces is to mistake the correlation between common effects of a cause for a causal relationship.

POTENTIAL OUTCOMES AND COUNTERFACTUALS

Let Y represent an outcome such as that of a drug trial and $Y(u)$ represent the outcome for single individual u. Let X be a variable that may affect the outcome of the trial, such as whether the individual u took a specific drug. For example, we could set $X = 1$ for taking the drug and $X = 0$ for not. The *potential outcome* of the variable $Y(u)$ is the value $Y(u)$ would have taken if $X = x$. This is denoted $Y_{X=x}(u)$. If $X \neq x$ in the trial, then $Y_{X=x}(u)$ is the value $Y(u)$ would have had if (counter to the facts) $X = x$. The crucial assumption here is that such a value exists. Rubin's theory [14] of potential outcomes asserts that this value does exist. The Rubin causal model treats counterfactuals as abstract mathematical objects not derived from a causal model (defined below). In the view of Pearl [5, 280–281], using structural causal models to represent causality relations allows the analyst to clearly visualize and understand their assumptions. By contrast, Pearl claims the purely mathematical assumptions required by Rubin can be difficult to understand and verify.

PEARL'S DEFINITION OF CAUSALITY

In the 1980s, Pearl developed Bayesian Networks and wrote the influential book [15] which decades of analysts have used to model and reason about evidence and uncertainty. Unfortunately, Bayesian networks can estimate only associations, not causality. To estimate causality, Pearl defined three additional notions, causal models, interventions, and counterfactuals.

LADDER OF CAUSATION

In Pearl's ladder of causation [5, 28], he envisions three levels of causal reasoning.

LEVEL 1: ASSOCIATION

On this level, we identify and use regularities in observations. What events are associated with one another? This level allows us to make predictions. For example, what does this poll tell us about the chances of a certain candidate winning an election?

LEVEL 2: INTERVENTION

On this level, one can answer questions such as—*If I give the patient a drug for a certain ailment, how much will that increase his chances of being cured?*

LEVEL 3: COUNTERFACTUALS

On this level, one can answer questions such as—*If I had not taken an aspirin, would my headache have gone away?* Recall that the notion of a counterfactual is crucial to Hume's modified definition of cause: "where if the first object had not been, the second never had existed".

Pearl considers intervention to be a step above association in causal reasoning and counterfactuals a step above that.

We are all familiar with the successes of machine learning algorithms, such as speech recognition and the development of systems such as AlphaGoPlus that can beat the best human Go players. These systems are estimating associations, not causation. For example, speech recognition software recognizes the meaning of ambiguous words by using the meaning associated to that word in the context of the previous words in a sentence. Strategy systems such as AlphaGoPlus are making moves that are associated with positive outcomes in its learning sessions. Thus, machine learning is on the first rung of the causality ladder.

Pearl's insight is that to go up the ladder of causation and estimate the effects of interventions or counterfactuals, we must go beyond classical or Bayesian statistics and beyond systems such as Bayesian networks. To do this, Pearl defines structures called causal models and uses them to give a precise meaning to the terms intervention and counterfactual. A causal model must be added to a standard statistical model to perform causal inference at levels two and three of the ladder of causality. A causal model is an a priori assumption, just like a prior distribution distribution is in Bayesian statistics. Of course, the results of the causal inference will depend on the causal model, just as the results of Bayesian inference will depend on the prior distribution assumed. However, as with Bayesian analysis, the assumptions are explicit for everyone to see, understand, and question if they wish.

CAUSAL MODELS

STRUCTURAL CAUSAL MODEL (SCM)

A structural causal model (SCM) consists of two sets of variables, U and V, and a set F of functions f that assign to each variable in V a value based on the other variables in the model. The variables in U are *exogenous*; they do not depend on any other variables in the model. The variables in V are endogenous and must depend on (be a descendent of) one or more of the exogenous variables. The variables in V can also depend on other endogenous variables.

STRUCTURAL EQUATION MODEL (SEM)

If we know the functions $f \in F$ explicitly, then the SCM becomes a structural equation model (SEM).

GRAPHICAL CAUSAL MODEL (GCM)

Every SCM, M, is associated with a graphical causal model (GCM). The GCM contains one node for each variable in M. If the variable Y in M depends on Z in M, then there is a directed edge from Z to Y. *Causal Inference in Statistics* deals only with SCMs that can be represented by directed acyclic graphs, i.e., directed graphs without loops.

The graphical causal model for ice cream sales given in Figure 1 appears in chapter 3 of [1] and is used to illustrate the concept of intervention and the difference between association and causation. This graph represents an SCM in which

$$U = \{U_X, U_Y, U_Z\}, \quad V = \{X, Y, Z\},$$
$$\text{and } F = \{f_Z(U_Z), f_X(U_X, Z), f_Y(U_Y, Z)\} \tag{3}$$

where f_Z, f_X, f_Y are the functions (possibly unknown) that define Z, X, and Y. For this example, we assume the exogenous variables in U are independent random variables. Note that the temperature Z depends only on the exogenous variable U_Z; X

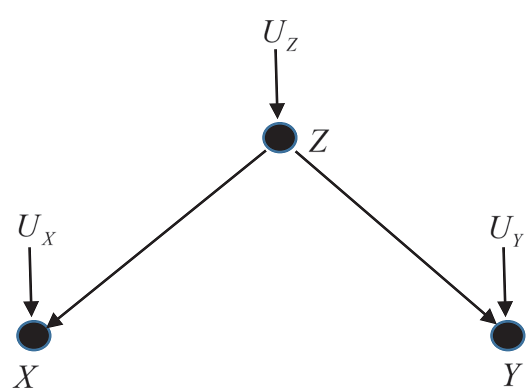

Figure 1
A graphical causal model representing the relationship between temperature Z, ice cream sales X, and crime rates Y. In this graph, X is not independent of Y, but it is conditionally independent of Y given Z.

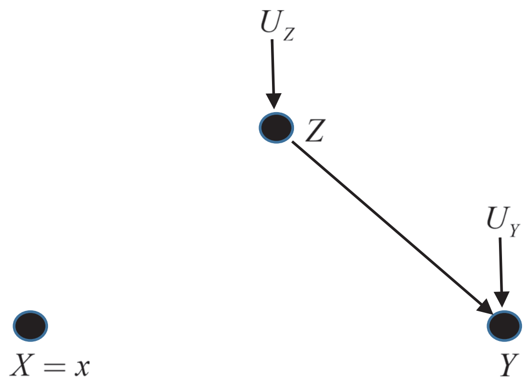

Figure 2
The graphical causal model M_x representing the intervention $X = x$ in the GCM in Figure 1.

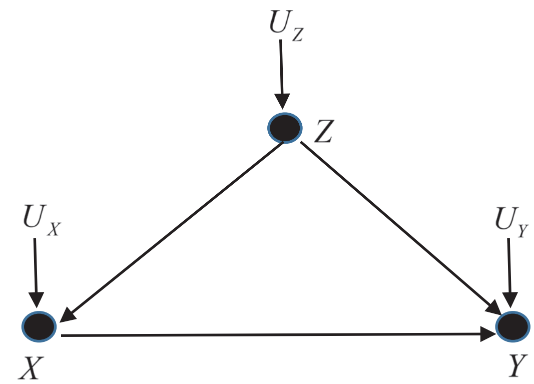

Figure 3
A graphical causal model representing the effects of a new drug, with Z representing gender, X drug usage, and Y recovery.

depends on the endogenous variable Z as well as U_X; and Y depends on Z as well as U_Y. However, there is no dependence of Y, crime rate, on X, ice cream sales. If one performed a statistical analysis of ice cream sales and crime rates, one would likely find a significant correlation between the two. However, from the GCM in Figure 1, we know this is an association not a causal relationship.

If we know the functions f_Z, f_X, f_Y explicitly, then the SCM in (3) becomes an SEM.

INTERVENTIONS AND COUNTERFACTUALS

In order to define interventions and counterfactuals, Pearl first defines the model M_x.

THE MODEL M_x

Suppose we have an SCM M. For defining both an intervention and a counterfactual, we use the SCM model M_x derived from M by setting the endogenous variable X in M to the fixed value x, which we indicate by writing $X = x$.

If M is specified by an SEM, we obtain M_x by replacing X by the fixed value x in all the equations of the model.

If M is specified by a GCM, M_x is obtained by setting the node $X = x$ and removing all the arrows that lead into node X, as shown in Figure 2.

INTERVENTIONS

Pearl uses the notation $P(Y \mid do(X = x))$ to indicate the probability distribution on Y when we intervene to set $X = x$. He points out that this distribution is different than $P(Y \mid X = x)$. He says,

> *In the distributional terminology, $P(Y \mid X = x)$ reflects the population distribution of Y among those individuals whose X value is x. On the other hand, $P(Y \mid do(X = x))$ represents the population distribution of Y if* everyone in the population *had their X value fixed at x.*

To understand what this means, let's consider an SCM in which the variables in U are stochastic and in which we know their joint distribution. We write $P(U = u)$ to denote the probability that the (vector) random variable U has the (vector) value u.

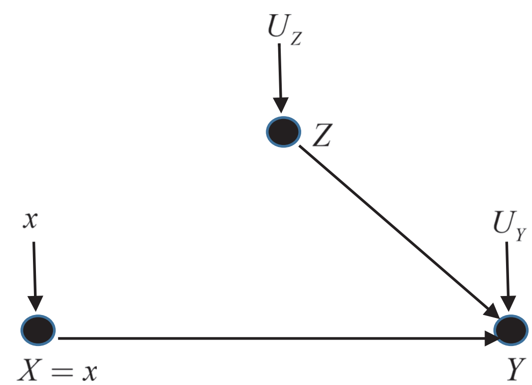

Figure 4
The modified version M_x of the GCM in Figure 3 representing the intervention $X = x$.

To calculate the effect of the intervention $do(X = x)$ on Y, we compute:

$$E\left[Y \mid do(X = x)\right] \equiv E\left[Y_{X=x}\right] \tag{4}$$

where $Y_{X=x}$ represents Y in the modified model M_x, but the expectation is taken over the unmodified (prior) distribution on U. For simplicity of notation we will often write Y_X for $Y_{X=x}$.

As an example, suppose we have an SEM where the functions f_Z, f_X, f_Y are known explicitly and where:

$$\begin{aligned} Z &= f_Z(U_Z) = U_Z \\ X &= f_X(Z, U_X) = g_X(Z) + U_X \\ Y &= f_Y(Z, X, U_Y) = g_Y(Z, X) + U_Y, \end{aligned} \tag{5}$$

i.e., we assume the functions g_X and g_Y are known explicitly. The GCM in Figure 3 shows a graphical version of this model in which Z represents gender, X drug usage, and Y recovery.

Figure 4 shows the graphical model M_x representing the intervention $X = x$.

Using (5), we can calculate the expectation in (4) by

$$E\left[Y \mid do\left(X = x\right)\right] = E\left[Y_x\right] = E\left[g_Y\left(Z, x\right) + U_Y\right] = \int \left[g_Y\left(U_Z, x\right) + U_Y\right] P\left(du\right). \quad (6)$$

In computing the expectation in (6), we have set $X = x$ in the equation for Y and computed the expectation of Y over the prior distribution on U to obtain $E[Y_x]$ and the effect of the intervention $X = x$.

To illustrate the difference between computing a conditional probability and a probability conditioned on an intervention, consider the question of whether a given vaccine tends to protect a person from contracting a disease. If we simply calculate the conditional probability of getting that disease given a person was vaccinated, we are estimating only an association. If the probability is lower for vaccinated than for unvaccinated people, a reason could be that people who tend to get vaccinated are healthier than those who do not or that they tend to have some natural immunity.

A way around this problem is to conduct a randomized controlled trial (RCT). In an RCT, two populations are selected at random. The first receives the vaccine, the other doesn't. The purpose of the randomization is to obtain two populations that are as close to identical as possible so that the only difference between the two is whether people were vaccinated or not. In this case, we can ascribe the decreased chance of contracting the disease to the vaccine. However, in many cases it may be too costly, too difficult, or even unethical to conduct an RCT. For example, you would not want to perform an RCT to determine if smoking causes cancer. For many questions, we are stuck with observational data.

In *Causal Inference in Statistics,* Pearl and his coauthors describe situations and methods by which we can perform causal estimation using observational data. This is the crucial capability provided by Pearl's model of causal inference that is not available from standard statistical techniques.

COUNTERFACTUALS

The counterfactual $P(Y_x = y \mid X = x')$ is the probability of $Y_x = y$ in the model M_x conditioned on the counterfactual $X = x' \neq x$. Recall that Y_x depends on the vector of random variables U. We indicate this by writing $Y_x(u)$ for the value of Y_x when $U=u$.

Let $P(U = u \mid X = x')$ be the probability that $U=u$ given $X = x'$, and let $\mu_{X=x'}$ be the resulting probability distribution on U. Then:

$$P\left(Y_x = y \mid X = x'\right) = E_{\mu_{X=x'}}\left[P\left(Y_x\left(u\right) = y\right)\right] = \int P\left(Y_x\left(u\right) = y\right) P\left(du \mid X = x'\right) \quad (7)$$

where the subscript $\mu_{X=x'}$ on E means that expectation is taken with respect to the measure $\mu_{X=x'}$. More generally:

$$E\left[Y_x \mid X = x'\right] = E_{\mu_{X=x'}}\left[Y_x\right] \quad (8)$$

Furthermore, we can calculate the expectation of Y_x under any counterfactual event $E = e$ by

$$E\left[Y_x \mid E = e\right] = E_{\mu_{E=e}}\left[Y_x\right] \quad (9)$$

COMMENTS ON INTERVENTIONS AND COUNTERFACTUALS

For both interventions and counterfactuals, we compute the expected value of Y_x in model M_x. For interventions, the expectation of Y_x is taken over by the unmodified (prior) distribution on U. For counterfactuals, the expectation of Y_x is taken over by the distribution of U conditioned on the counterfactual, e.g., $X = x' \neq x$.

For counterfactuals, we are estimating the probability of outcomes in an alternate world that does not exist. Surprisingly, there are ways to do this. One way is to have a detailed and accurate model of the system (world) you are analyzing and use that model to compute counterfactual probabilities. Consider the question of whether climate change has increased the probability of extreme weather events. Here we have to consider a counterfactual: if the global average temperature had been 1.5 °F cooler than it is now, how would that have changed the probability of these extreme events occurring? Think about the recent wildfires in Australia at the end of 2019 and beginning of 2020. Is this extreme event now more likely to occur because of climate change?

The special supplement to the Bulletin of the American Meteorological Society of January 2019 [16] includes climate change attribution assessments for seventeen different extreme events from around the world during 2017. For 16 out of those 17 events, the assessments concluded that climate change (global warming) increased the probability of their occurrence. In one of most striking assessments, [17] found that the likelihood of a heat wave at least as hot as the one that happened in the European-Mediterranean region in the summer of 2017 "is at least 3.5 times higher compared to 1950". The probability of this event is now 10% per year.

How was [17] able to make this estimate? The authors developed a statistical model of temperatures in this region. This model depends on the global mean surface temperature (GMST). Using climate models with enough data to analyze the distribution of past and present temperatures, they verified that the GMST influences only the mean of the distribution not the shape. The authors used the temperature distribution resulting from the 2017 value of GMST to determine the probability of experiencing a heat wave in the European-Mediterranean region at least as hot as the one that occurred in 2017. They then compared this to the probability of this event occurring using the (counterfactual) temperature distribution based on the 1950 GMST. The difference or ratio of these two probabilities provides an estimate of the effect of global warming on the probability of occurrence of this extreme event.

As well as providing a quantitative estimate of the effect of global warming on the occurrence of this extreme weather event, this analysis now changes the terms of discussion "from I do not believe in global warming" to "how good is the model that was used to estimate this effect". The latter is a more directed and scientific question and more suitable to quantitative analysis.

Unfortunately, we do not always have a detailed model available to answer our counterfactual questions. Often, we have only observational data. This is where GCMs and Pearl's theory of causality can help us out. If we have a GCM, then in some circumstances (identified in [1]), we can estimate counterfactuals from observational data. Again, the GCM is an assumption, but now the question of the validity of the estimate can be reduced to a more contained and scientific question of whether the GCM is a good model or not. In any case, the assumptions under which the estimate is valid are clearly stated and generally easy to understand.

OUTLINE OF CAUSAL INFLUENCE IN STATISTICS

The book contains four chapters:

1. Preliminaries: Statistical and Causal Models
2. Graphical Models and Their Applications
3. The Effects of Intervention
4. Counterfactuals and Their Applications

CHAPTER 1

Chapter 1 begins with a discussion of why we should study causality and some examples of the fact that standard statistical methods can lead us astray when we use them to estimate causality. The examples are a form of Simpson's paradox, which showed that a statistical association can hold over a whole population but be reversed in every subpopulation. The book presents several examples [1, 3–4] of this paradox, including a hypothetical study of the effects of exercise on cholesterol. If we segregate the data in this study by age, we see that exercise tends to reduce cholesterol levels in every age group; but, if we aggregate the data over all age groups, we find that more exercise tends to produce higher cholesterol levels. The problem here is that older people, who have higher cholesterol levels than younger people, also tend to exercise more. This highlights the question of when to segregate and when to aggregate data when estimating an effect. Statistics by itself has no satisfactory answer to this question. Chapter 1 promises that causal inference will provide an answer to this problem.

The remainder of the chapter covers basic probability, statistics, and graphs. It concludes by defining SCMs and GCMs. (See the section Causal Models.)

CHAPTER 2

Chapter 2 discusses how to use graphs to model dependencies in data. It discusses the notions of chains, forks, and colliders and their importance in causal estimation. It defines the notion of *d-separation* which is an important concept for performing causal estimates. The d-separation property allows one to estimate the effect of an intervention using observational data, i.e., without having to perform an RCT. This chapter also discusses ways in which one can use data to test the validity of a GCM.

CHAPTER 3

Chapter 3 discusses how to estimate the effect of an intervention. Having defined SCMs and GCMs in chapter 2, [1] defines intervention in terms of GCMs as we did in the section Graphical Causal Model (GCM) and defines the "do" operator, e.g., $do(X=x)$ to indicate an intervention. Recall that $P(Y \mid do(X=x))$ is the distribution of the values of Y if every member of the population had their X value set to x. By contrast, $P(Y \mid X=x)$ is the distribution of values of Y among those members of the population whose X value happens to equal x. The latter is an association; the former a causal estimate—what would be the effect on Y of setting $X=x$? The causal estimate is a prediction of the effect produced by the intervention $X=x$. An example is the estimate of the reduction in the probability of contracting a disease if someone is given a vaccination against that disease.

The notion of intervention, although simple to state, is crucial to Pearl's definition of causality. For convenience and emphasis, we repeat it here. First let us recall the definition of an SCM given in the section Structural Causal Model.

An SCM consists of two sets of variables, U and V, and a set F of functions f that assign to each variable in V a value based on the other variables in the model. The variables in U are *exogenous,* they do not depend on any other variables in the model. The variables in V are *endogenous* and must depend on (be a descendent of) one or more of the exogenous variables. The variables in V can also depend on other endogenous variables.

If the equations in F are known explicitly, then the SCM becomes an SEM. Even if we do not know the equations in the SCM explicitly, we can construct a GCM. For example, if we did not have explicit versions of the equations defining the SCM specified by (5), we could represent this SCM by the GCM in Figure 3. The specification of the variables upon which each equation in F depends tells us where to place arrows in the GCM representation of the SCM.

DEFINITION OF CAUSAL EFFECT

Suppose we have a model M specified by an SEM or GCM. (Note any SCM can be represented by a GCM, so we need only consider SEMs and GCMs). Then, M_x is the model obtained by setting the variable $X=x$ in the SEM or the node $X=x$ in the GCM and removing all arrows into X. To calculate the causal effect of the intervention $do(X=x)$ on Y, we compute

$$E\left[Y \mid do\left(X=x\right)\right] \equiv E\left[Y_x\right] \tag{10}$$

Let $X=1$ if a drug is given to a patient and $X=0$ if not. Suppose Y is the outcome where $Y=1$ if cured and $Y=0$ if not. The book defines the average causal effect of the intervention $X=1$ as

$$E\left[Y \mid do\left(X=1\right)\right] - E\left[Y \mid do\left(X=0\right)\right] \tag{11}$$

In this case, we obtain an estimate of the increased probability of cure when a patient takes the drug compared to not taking the drug. More generally, (11) represents the average causal effect on Y of the intervention $do(X = 1)$ whatever the intervention X represents.

Using the definition of intervention in (10), [1] is able to give guidance on when to segregate or adjust data when performing statistical estimations.

As an example, consider the situation represented by Figure 3. To estimate the effectiveness of the drug, we imagine a hypothetical intervention in which the drug is administered uniformly to everyone in the population and compare the recovery rate to the situation where no one takes the drug. That is, we wish to estimate

$$P\left(Y=1 \mid do\left(X=1\right)\right)-P\left(Y=1 \mid do\left(X=0\right)\right) \tag{12}$$

This is the average causal effect defined in (11). However, we cannot simply estimate this effect from observational data because, as we see from Figure 3, our GCM says that gender affects both drug usage and recovery. To estimate the average causal effect, we must change the model M given by Figure 3 to the model M_x given by Figure 4.

Using this model, [1] shows that

$$P\left(Y=y \mid do\left(X=x\right)\right)=\sum_z P\left(Y=y \mid X=x, Z=z\right)P\left(Z=z\right) \tag{13}$$

Note that the right-hand side of (13) contains only probabilities that can be estimated from observational data. This equation is the adjustment equation which says that in this case, we must adjust our estimates of the effectiveness of the drug conditioned on sex and then produce an overall estimate by weighting these estimates by the distribution of sex in the population under consideration. That is, we are adjusting our estimates by the marginal distribution of Z.

BACKDOOR AND FRONT-DOOR CRITERIA

This chapter defines the notions in GCMs of the backdoor and front-door criteria for a GCM which enable us to use observational data to estimate the effect of an intervention. The backdoor criterion is a special case of d-separation. When a GCM satisfies these criteria, [1] gives formulas for calculating the effects of interventions (e.g., the expectation in (10)) using observational data. This is particularly useful for situations where randomized controlled trials are not feasible or their data are not available. Since the backdoor and front-door criteria apply to GCMs, they give us a way to estimate causal effect when we do not have explicit formulas for the functions in F, i.e., when we do not have an SEM. The chapter describes other methods for estimating the effect of an intervention such as inverse probability weighting and mediation.

CAUSAL INFERENCE IN LINEAR SYSTEMS

The chapter finishes by illustrating causal inference in linear systems and discussing the difference between structural (causal) coefficients in linear systems and regression coefficients. If one has an SCM and performs a regression in accordance with that, the resulting regression coefficients are structural coefficients allowing one to compute causal effects. By contrast, regression coefficients without the context of a SCM give estimates of associations only.

CHAPTER 4

Chapter 4, which deals with counterfactuals, is the most challenging and difficult to absorb. One of the reasons that it was difficult for me is that the chapter does not provide a crisp mathematical definition of counterfactual, such as the one given in the section Interventions and Counterfactuals above.

DIFFERENCE BETWEEN COUNTERFACTUALS AND INTERVENTIONS

The crucial difference between the *do* operator and a counterfactual is that the $do(X = x)$ operator captures the behavior of a population under the intervention $X = x$, whereas $Y_{x'}(u)$ describes the behavior of the individual u under the condition $X = x'$. If it happens that $X = x \neq x'$ in our data for u, then we are estimating a counterfactual, e.g., what would a person's salary have been if she had gone to college rather than beginning work right after high school. We stated this difference in mathematical terms, in the section Interventions and the section Counterfactuals.

COUNTERFACTUAL EXAMPLE

A simple example given in [5, 273–279] illustrates the concept of counterfactuals and how they differ from interventions. Table 8.1 in [5] shows data listing employees, their salary S, education ED (= 0, 1, or 2 for high school, college, or graduate degree), and years of work experience EX. From this table one can perform a linear regression to obtain

$$S = \$65{,}000 + 2{,}500 \times EX + 5000 \times ED \tag{14}$$

as the expected salary of a worker as a function of years of experience and education. In the data, Alice has 6 years of experience, a high school education (ED = 0), and a salary of \$81,000. The counterfactual question is, what would Alice's salary be if she had a college degree, i.e., what is the value of $S_{ED=1}(\text{Alice})$? Using the regression in (14), we could answer this question by setting EX = 6 and ED = 1 to obtain \$85,000. However, the regression in (14) does not account for the fact that education and experience are dependent. We know that

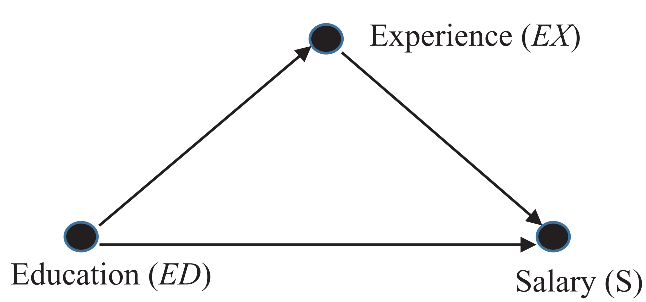

Figure 5
Effect of education (*ED*) and experience (*EX*) on salary *S*.

that the number of years of education tends to lower the number of years of experience. With this in mind, we construct the GCM in Figure 5.

This GCM indicates that education has an effect on experience. It also says that salary does not affect education. We know that causal relationship goes the other way around, namely education affects salary as shown. With this in mind, we rewrite (14) in terms of a SEM, specifically

$$S = \$65{,}000 + \$2{,}500 \times EX + 5000 \times ED + U_S \tag{15}$$

where U_S represent the individual factors that affect a person's salary. Since Figure 5 says that education affects experience, we next perform a linear regression on the data to estimate that effect. By Figure 5, salary does not affect experience, so we set the coefficient of S to 0 in the regression. Suppose we obtain

$$EX = 10 - 4 \times ED + U_{EX} \tag{16}$$

To perform the counterfactual analysis to obtain $S_{ED=1}$(Alice), we first put Alice's experience, education, and salary into (15) and (16) to obtain U_S(Alice) = \$1,000 and U_{EX}(Alice) = −4. We now erase any arrows pointing into ED and set ED=1. In this case there are no arrows pointing into ED, so this step is trivial. In many cases, there are arrows pointing into the counterfactual variable and these must be removed from the model. Finally, we put the values of U_S(Alice) = \$1,000, U_{EX}(Alice) = −4, and ED=1 into the model. First, we compute Alice's experience if she had gone to college. From (16), we see that $EX_{ED=1}$(Alice) = 2. Then from (15) we compute

$$S_{ED=1}(\text{Alice}) = \$65{,}000 + \$2{,}500 \times 2 + \$5{,}000 \times 1 + \$1{,}000 = \$76{,}000$$

as an estimate of Alice's salary if she had gone to college. This is lower than the \$85,000 estimate from the regression. The reason the counterfactual estimate is lower than the regression estimate is that it accounts for the fact that going to college would reduce the number of years of experience that Alice has. In addition, the counterfactual analysis accounts for the terms U_S(Alice) and U_{EX}(Alice) unique to Alice.

As one can see even in this simple deterministic case, counterfactual analysis is not simple.

CONTENTS OF CHAPTER 4

Chapter 4 examines both deterministic and probabilistic counterfactuals. It also provides examples of practical uses of counterfactuals, such as determining the effectiveness of a government program or estimating the effect of sex discrimination in hiring. These examples give us a feeling for the important questions that counterfactual analysis can help us answer. The book finishes with a description of some mathematical tools for estimating probabilities of causation and mediation.

SUMMARY

Modern causal inference, which has developed methods for obtaining quantitative estimates of the effect of interventions and counterfactuals, is an important and relatively new area of analysis. Every analyst should be familiar with the concepts and definitions of causal inference. Causal inference represents a significant extension of standard statistical analysis that should become an increasingly important tool for answering questions about the effectiveness of interventions and for developing artificial intelligence (AI)-like systems that can reason and make decisions. Present AI systems don't reason at the counterfactual level (causality level 3). They make decisions based only on association, unlike humans who can also make decisions based on counterfactual reasoning.

Causal Inference in Statistics is a good introduction to Pearl's version of causal inference. Even though it is a "primer", it requires substantial effort on the reader's part to understand and to apply the concepts and tools presented. The presentation is informal and requires little background beyond basic probability, which is useful for an introduction but frustrating if you are looking for a more mathematical and rigorous approach. For that one has to delve into [4], which can be daunting.

In summary, I highly recommend this book as an introduction to this emerging and important area of analysis. Other introductory texts on causal inference are [2], [3], and [18].

REFERENCES

1. Pearl, J., Glymour, M., and Jewell, N. P. *Causal Inference in Statistics: A Primer*. New York, NY: Wiley, 2016.
2. Imbens, G. W., and Rubin, D. B. *Causal Inference in Statistics, Social and Biomedical Sciences*. Cambridge, UK: Cambridge University Press, 2015.
3. Hernan, M. A., and Robins, J. M. *Causal Inference: What If*. Boca Raton, FL: CRC Press, 2019.
4. Pearl, J. *Causality: Models, Reasoning, and Inference* (2nd ed.). New York, NY: Cambridge University Press, 2009.
5. Pearl, J., and Mackenzie, D. *The Book of Why*. New York, NY: Basic Books, 2018.
6. Galton, F. Regression toward mediocrity in hereditary stature. *Journal of the Anthropological Institute of Great Britain and Ireland*, (1886), 246–263.
7. Kleinberg, S., and Hripcsak, G. A review of causal inference for biomedical informatics. *Journal of Biomedical Informatics*, Vol. 44, 6 (2011), 1102–1112.
8. Hume, D. *An Enquiry Concerning Human Understanding*. 1748.
9. Mackie, J. L. Causes and conditions. *American Philosophical Quarterly*, Vol. 12, (1965), 245–265.
10. Hill, A. B. The environment and disease: Association or causation? *Proceedings of the Royal Society of Medicine*, Vol. 58, 5 (1965), 295–300.
11. Good, I. J. A causal calculus (I). *British Journal for the Philosophy of Science*, Vol. XI, 44 (1961), 305–318.
12. Eells, E. *Probabilistic Theory of Causality*. New York, NY: Cambridge University Press, 1991.
13. Suppes, P. A. *Probabilistic Theory of Causality*. New York, NY: North Holland, 1970.

14. Rubin, D. Estimating causal effects of treatments in randomized and non-randomized studies. *Journal of Educational Psychology*, Vol. 66, (1974), 688–701.
15. Pearl, J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann, 1988.
16. Herring, S. C., Christides, N., Hoell, A., Hoerling M., and Stott, P.A. Explaining extreme events of 2017 from a climate perspective. *Bulletin of the American Meteorological Society*, Vol. 100, I (2019), S1–S117.
17. Kew, S. F., Phillip, S. Y., van Oldenborgh, G. J., Otto, F. E. L., Vautard, R., and van der Schrier, G. The exceptional summer heat wave in southern Europe 2017. *Bulletin of the American Meteorological Society*, Vol. 100, 1 (2019), S49–S53.
18. Morgan, S. L. W. C. *Counterfactuals and Causal Inference*. Cambridge, UK: Cambridge University Press, 2015.

Book Review Reflection

CAUSAL INFERENCE IN STATISTICS: AN ATTEMPT AT SOME REFLECTION

Alfonso Farina
SELEX ES (retired)
Rome, Italy
alfonso.farina@outlook.it

It is a great pleasure and a distinct honor to be invited to read and, possibly, offer some comments to Dr. Lawrence D. Stone's review [1] of the book [2]. I like Dr. Stone's review, which is high quality, highly professional, and carefully written. Determining the causal relationships between actions and results is one of humanity's endless struggling searches. Thus, I very much welcome the books [2] and [3], the second being the one I have browsed in recent months after a suggestion from Dr. R. Streit at the Maritime Situational Awareness Workshop (MSAW 2019, Lerici, Italy) (https://www.cmre.nato.int/msaw-2019-home). The review [1] by Dr. Stone is a precious guide and invitation to read the books [2] and [3]. The background section, the various definitions of causality (including the one proposed by Professor J. Pearl), the graphical causal models, and the definition of interventions and counterfactuals are very helpful and smoothly conduct the reader to the ensuing descriptions of Dr. Stone of the four chapters of book [2].[1,2] The summary and reference list thoroughly complete the book review.

From the summary of [1], I like to quote the "take away" messages: "Modern causal inference, which has developed methods for obtaining quantitative estimates of the effect of interventions and counterfactuals, is an important and relatively new area of analysis. Every analyst should be familiar with the concepts and definitions of causal inference. Causal inference represents a significant extension of standard statistical analysis that should become an increasingly important tool for answering questions about the effectiveness of interventions and for developing artificial intelligence (AI)-like systems that can reason and make decisions. *Present AI systems don't reason at the counterfactual level*. (Italics are mine). They make decisions based only association unlike humans who can also make decisions based on counterfactual reasoning."

Actually, I had already come across the book topic in March 2017 in Singapore. I am referring to the conference "Causality–Reality" organized by Para Limes (https://www.paralimes.org/) and Nanyang Technological University, Singapore [6].[3]

Abridged from the synopsis of the conference: "We seek to manage and control our world by establishing causalities. And we try to use science to help us. However, one of the biggest challenges for science is to untangle or better understand the relationship between causality and reality. This is especially true for complexity science that deals with the real world, or with complex systems like our brains or our immune system. *Causality* is the agency or efficacy that connects one process (the cause) with another (the effect), where the first is understood to be partly responsible for the second. *Reality* is the state of things as they actually exist, rather than as they may appear or might be imagined. Once we have met this challenge, we have the key to finding ways to sustainably manage our lives, our systems, our science, our education, our laws, our healthcare and our cities in a world that is becoming more complex and interconnected than ever before. It is also key to finding new breakthroughs in the sciences that seek to understand 'the human' and its relations" [6].

It would be of interest to have a look at the presentations from the Singapore Conference. In the following, I just picked one statement which struck me from each presentation:

- Bertil Andersson, opening address (see above quoted synopsis).
- Jan W. Vasbinder, welcome remarks ("The ultimate equations in this quote constitute causality. Cataloguing and understanding emergent behavior constitute the relation to reality.") [6]
- George Rzevski, "Managing Organization Complexity: Practical Methods and Tools for Adaptation and Causality Analysis" ("Managing complexity…") [6]
- Stuart Kauffman, Beyond Physics: The Emergence and Evolution of Life ("The emergence and evolution of life is based on physics but it is beyond physics. Evolution is an historical process arising from the non-ergodicity of the universe above the level of atoms… [4] Beyond entailing law, the evolving biosphere literally constructs itself and is the most complex system we know in the universe."[5]) [6]
- De Kai, Translating Reality to Causality ("… the study of cognition…") [6]
- Michael Puett, Rethinking Notions of Causality and Reality: Indigenous Theories from China ("…causality and reality is hotly debated in both the sciences and social sciences") [6]
- James Bailey, Schooling for Life K–12 ("…replacement of K–12 curriculum grounded in network of neurons rather than lines of forces") [6]

[1] Hot topic indeed, Google provides [4].
[2] Intriguing topic, Google provides [5].
[3] Beyond Boundaries, indicating a limitless potential for exploration [6]. "Limes" in Latin is for boundary.

[4] Medaglia, J. D., Ramanathan, D. M., Venkatesan, U. M., and Hillary F. G. The challenge of non-ergodicity in network neuroscience.Available: https://www.ncbi.nlm.nih.gov/pubmed/22149675/.
[5] Kauffman, S. *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. Oxford, UK: University Press, 2001.

- Nick Obolensky, A Military View ("…a non-hierarchical complex adaptive system, instead of as a traditional hierarchical one, can get surprising results") [6]
- Ernst Pöppel, Trust as Basis for the Concept of Causality: A Biological Speculation ("…humans have the tendency to attribute only one cause when trying to understand whatever has happened or whatever is given") [6]
- Stefan Thurner, How Complexity Weakens Causality—Emerging Dangers—and Ways Out ("…we discuss problems that emerge from a world that is getting increasingly complex and seemingly less causal such as the gradual public acceptance of a 'post-factual era'") [6]
- Ilan Chabay, Behavioral Causality—Anthropocene Reality ("The UN Sustainability Development Goals (SDGs) set 17 ambitious targets for moving to sustainable futures. How are we to understand, enable, and foster collective behavior changes that can address these complex changes?") [6]
- Peter Edwards, Technological Myopia ("Innovation is good; disruptive innovation is better! Will explore some of the unforeseen consequences of disruptive technologies and ask why it is that we have such difficulty anticipating them.") [6]
- Mile Gu, Quantum Simplicity: Can quantum mechanics better isolate the causes of natural things? ("Certain observed phenomena may appear to require tracking immense amounts of information to model classically, and yet remarkable little information quantum mechanically.") [6]
- Sydney Brenner, Causality in Evolution ("Complexity... invisible reality") [6]
- Jan W. Vasbinder, closing remarks [6]

To me, the recurrent word of these talks is complexity, which resonates with Pearl's book too.

In the "Epilogue, The Art and Science of Cause and Effect," pp. 401–428 [3], very captivating to read, the pivotal role of Graphical Causal Models (GCM) in identifying causal effects is summarized.[6] The key point is the great complexity of realistic GCM. Detailed models are always difficult, maybe impossible, to achieve together with accurate initial/boundary condition data, otherwise…chaotic models come up with corresponding unpredictability.

I am acquainted with the A. L. Barabasi work (e.g., [7]), the new science of networks, a perspective which seems of interest to consider. In the most basic form, a network is a set of objects and a set of connections between pairs of objects. From a mathematical point of view, a network takes the form of a graph where the interconnected objects are represented by mathematical abstractions called vertices (or nodes), and the connections called edges. A. L. Barabasi's book introduces the science of networks to the general audience. It provides an introduction to the main models and properties of networks and their applications in many areas of real life. The subtitle of the book is also informative: "How Everything is Connected to Everything Else and What It Means for Business, Science, and Everyday Life". Thus, it is evident that network/graph theory plays a key role in modelling physical phenomena and systems. Signal processing over graphs is also becoming an attractive and powerful engineering tool. Emergent behaviors, typical of networks, arise through formation of patterns not reducible to a single agent's behavior.

A. L. Barabasi and coworkers tackled the mathematics of controllability and observability, generalizing the concepts introduced by Prof. R. E. Kalman in 1963, of somehow realistic networks [8] and [9].[7,8] The Nobel Laureate Murray Gell-Mann (who discovered the quark particle) introduced a way to measure the complexity of network by means of its plecticity [10].[9] A brief review of the methods to try to measure complexity is in chapter 10 of [11]. In the same chapter, two study cases have been described to calculate the complexity and controllability of nets of a few hundred nodes and a few thousand connections.

I see potential intersections between GCM and the above quoted theory of nets. I would argue that similar topics related to complexity, controllability, and observability (and more in general the *"ilitis"* of description of dynamic systems) should be explored and taken into account in GCM for realistic applications. Indeed in [3], p. 77, the *identifiability and the causal effect identifiability* are described. Graphical tests of *identifiability* are tackled at p. 89 and the following pages.

Another remarkable topic afforded in [2], and carefully spotted in [1], is the counterfactuality. To me this is an outstanding technical novelty introduced in the GCM.

Speaking with the neuroscientist Ph.D., C. Di Dio (coauthor of [12]), whose kind assistance is warmly acknowledged, I learned something on the neuroscience point of view of counterfactual thinking [12].[10,11] In summary, by means of the counterfactual reasoning, we exploit all our cognitive resources to evaluate which would have been the results we would have achieved if we had acted in an alternative way. This implies comparing the expected result with the one actually obtained. The process serves to modulate the behavior in order to improve it. The emotional component is also involved because the disappointment for not having obtained the desired result translates into the motivation to change. The areas involved are the orbitofrontal area where the comparison between the expected and the real takes place. This is where the signal starts in terms of error, which updates our expectations by means of connec-

[6] A public lecture delivered by Prof. J. Pearl, November 1996, as part of the UCLA Faculty Research Lectureship Program.

[7] "Identify the set of driver nodes with time-dependent control that can guide the system's entire dynamics in a complex directed network." [8]

[8] "A system is called observable if the system's complete internal state can be reconstructed from its outputs." [9]

[9] Simply speaking, it is related to the connectivity of the graph. It is expressed via the "betweeness centrality" of nodes.

[10] Università Cattolica del Sacro Cuore, Milan (I), Member of the Research Unit on Theory of Mind, Department of Psychology.

[11] Additional scientific literature, such as [13] and [14], have been suggested by Ph.D. M. Boccia, kindly acknowledged, University of Rome (I), Sapienza, Department of Psychology.

tions with subcortical (amygdala) and with cortical (frontal). The anterior cingulate cortex and hippocampus play also a key role for modulating behavior. These results extend the possible role of a *mirror-neuron like mechanism* beyond basic emotions. A gambling task and functional-Magnetic-Resonance-Imaging (fMRI) were used to test this hypothesis using *regret*, the negative cognitively-based emotion arising from a counterfactual comparison between the outcome of chosen and discarded options, whereby the discarded option would have produced higher benefits to the individual [12], [15].

The discovering of mirror-neurons effect is one of the most exciting events in neuroscience [16]. Mirror-neurons are related to empathy, imitation, the chameleon effect, and probably language (sing, calls, etc.) development. Researchers in cognitive neuroscience and cognitive psychology consider that this system provides the physiological mechanism for the *perception-action* cycle. The mirror-neurons may be important for understanding the actions and intentions of other people, and for learning new skills by imitation. It is also suggested that mirror-neuron systems may simulate observed actions, and thus contribute to theory of mind skills. It is felt that mirror neurons are the neural basis of the human capacity to feel empathy, and namely to resonate with another's emotional states. Thanks to the visual/audio-motor coupling mediated by the mirror system, some processes, such as understanding others' motor goals and intentions, are faster compared to systems based on mere cognitive, inferential processes. In the 1980s and 1990s, Prof. Giacomo Rizzolatti (coauthor of [12]) was working with G. Di Pellegrino, L. Fadiga, L. Fogassi, and V. Gallese at the University of Parma, Italy, and discovered this phenomenon [17], [18].

This said, I am unable to suggest how this neuroscience view of counterfactual thinking can be embedded in the math of GCM. Counterfactual thinking depends on an integrated network of systems for affective processing, mental simulation, and cognitive control to guide adaptive behavior [13]. This could be the avenue for future research. Teaming with neuroscientists and psychologists would be—in my opinion—most welcome.

Probabilistic counterfactual and probability of causation are powerfully developed in [3] and properly mentioned in [1] from [2]. This could raise the question as to whether probability may continue to be the way to measure randomness, which is the physical phenomenon. Researchers and engineers are exploring—among others—belief functions and imprecise probability theory [19], [20]. I refrain to dwell on this debated topic. I just wish to mention two remarkable books [21] and [22] by N. Taleb. [21] considers the extreme impact of rare and unpredictable outlier events, called the Black Swan theory. [22] introduces the concept of antifragility, which is beyond resilience or robustness. Quoted from N. Taleb: "Antifragility makes us understand fragility better… Anything that has more upside than downside from random events is antifragile; the reverse is fragile." Examples that could be classified as antifragile are: stochastic tinkering, simulated annealing, and stochastic resonance [23], [24]. Many others are presented in the Taleb book. The subtitle of the Taleb book is: "How to live in a world we don't understand." Quoting Taleb's main subject matter: "decision under opacity", that is, a map and a protocol on how we would live in a world we don't understand.

This struggling question appears to resonate also in the close of "Epilogue" of [3]. "But the really challenging problems are still ahead: we still do not have a causal understanding of poverty and cancer and intolerance, and only the accumulation of data and the insight of great minds will eventually lead to such understanding. The data is all over the place, the insight is yours, and now an abacus is at your disposal, too. I hope the combination amplifies each of these components."

Prof. J. Pearl's sense of hope is indeed very welcome.

REFERENCES

1. Stone, L. Book review, *Causal Inference in Statistics: A Primer* by J. Pearl, M. Glymour, and N. P. Jewell. *ISIF Perspectives on Information Fusion*, Vol. 3, (2020).
2. Pearl, J., Glymour, M., and Jewell, N. P. *Causal Inference in Statistics: A Primer*. New York, NY: Wiley, 2016.
3. Pearl, J. *Causality: Models, Reasoning and Inference* (2nd ed.). New York: Cambridge University Press, 2009.
4. Causal inference. [Online]. Available: https://en.wikipedia.org/wiki/Causal_inference.
5. Counterfactual thinking. [Online]. Available: https://en.wikipedia.org/wiki/Counterfactual_thinking.
6. In Proceedings of Conference: Causality—Reality, Para Limes Annual Conference, Singapore, March 6–8 2017. [Online]. Available: https://www.paralimes.org/events/past-events-video-recordings-and-slides/conference-causality-reality/.
7. Barabasi, A. L. *Linked: The New Science of Networks*. New York, NY: Perseus Books Group, 2002.
8. Liu, Y.-Y., Slotine, J.-J., and Barabasi, A. L. Controllability of complex networks. *Nature*, Vol. 473, (2011), 167–173.
9. Liu, Y.-Y., Slotine, J.-J., and Barabasi, A. L. Observability of complex systems. *PNAS*, Vol. 110, 7 (2013), 2460–2465.
10. Gell-Man, M. *The Quark and the Jaguar: Adventures in the Simple and the Complex*. New York, NY: Henry Holt and Company, 1994.
11. Farina, A. A cross-disciplinary overview with potential applications and examples for cognitive radar. In *The Impact of Cognition on Radar Technology*. A. Farina, A. De Maio, and S. Haykin, Eds. London, UK: Scitech Publishing, The Institution of Engineering Technology, 2017, ch. 10, 237–267.
12. Canessa, N., Motterlini, M., Di Dio, C., Perani, D., Scifo, P., Cappa, S. F., and Rizzolatti, G. Understanding others' regret: A fMRI study. *PLoS ONE*, Vol. 4, 10 (2009), 1–19.
13. Van Hoeck, N., Watson, P. D., and Barbey A. K. Cognitive neuroscience of human counterfactual reasoning. *Frontiers in Human Neuroscience*, Vol. 9, (2015).
14. Van Hoeck, N., Ma, N., Ampe, L., Baetens, K., Vandekerckhove, M., and Van Overwalle, F. Counterfactual thinking: An fMRI study on changing the past for a better future. *Social Cognitive Affective Neuroscience,* Vol. 8, 5 (2013), 556–564.
15. Kahneman, D., and Tversky, A. Variants of uncertainty. *Cognition*, Vol. 11, (1982), 143–157.
16. Mirror neuron. [Online]. Available: https://en.wikipedia.org/wiki/Mirror_neuron.

17. Di Pellegrino, G., Fadiga, L., Fogassi, L., Gallese, V., and Rizzolatti, G. Understanding motor events: A neurophysiological study. *Experimental Brain Research*, Vol. 91, (1992), 176–180.
18. Rizzolatti, G., Fadiga, L., Gallese, V., and Fogassi, L. Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, Vol. 3, 2 (1996), 131–141.
19. Benavoli, A., Chisci, L., Ristic, B., Farina, A., and Graziano, A. *Reasoning Under Uncertainty: From Bayes to Valuation Based Systems, Application to Target Classification and Threat Assessment*. Rome, Italy: SELEX Sistemi Integrati, ISBN 978-888665843-0, 2007.
20. Ristic, B., Gillam, C., Byrne, M., and Benavoli, A. A tutorial on uncertainty modelling for machine reasoning. *Information Fusion*, Vol. 55, (2020), 30–44.
21. Taleb, N. N. *The Black Swan: The Impact of the Highly Improbable*. New York, NY: Random House, 2007.
22. Taleb, N. N. *Anti Fragile—How to Live in a World We Don't Understand*. London, UK: Allen Lane, an imprint of Penguin Books, 2012, p. 519.
23. Benzi, R., Sutera, A., and Vulpiani, A. The mechanism of stochastic resonance. *Journal of Physics*, A14, (1981), L453-L457.
24. Benzi, R., Parisi, G., Sutera, A., and Vulpiani, A. Stochastic resonance in climate change. *Tellus*, Vol. 34, (1982), 10–16.

FUSION | 2021

24th International Conference on Information Fusion

Sun City | South Africa | July 12 - 15, 2021

www.fusion2021.isif.org

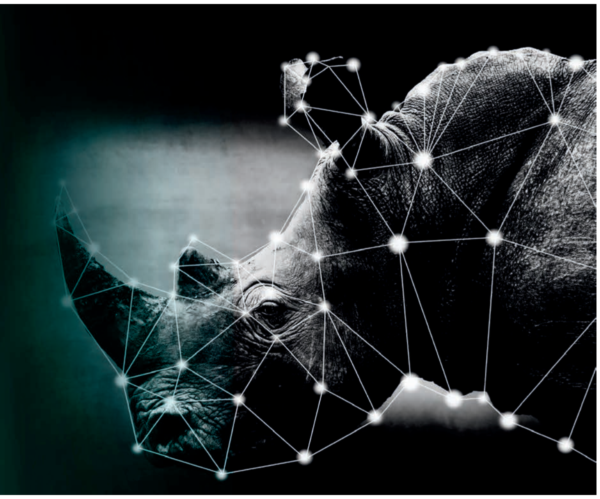

VENUE

The 24th International Conference on Information Fusion will take place at the Sun City Convention Centre, Rustenburg, South Africa. The venue is accustomed to hosting large international events, and being 200km from Johannesburg, it offers easy access via car rental, regular shuttle services and flights to a nearby airport.

PROGRAMME

The conference will take place from 12 to 15 July 2021 and offers the following programme:

- Monday 12 July: Tutorials and conference registration
- Tuesday 13 July: Conference and safari game drive followed by sundowner cocktails
- Wednesday 14 July: Conference and banquet
- Thursday 15 July: Conference and close

The International Conference on Information Fusion is the premier forum for interchange of the latest research in data and information fusion, and its impacts on our society. The conference brings together researchers and practitioners from academia and industry to report on the latest scientific and technical advances. Topics of interest include information fusion theory and representation, algorithms for target tracking, target and sensor modelling and evaluation and applications of information fusion such as IoT, robotics, big data, AI, autonomous systems *etc*.

SOCIAL PROGRAMME

The social programme will include a 3-hour safari game drive on 13 July in the neighbouring Pilanesberg National Park, including a halfway stop to serve sundowner drinks and snacks while listening to sounds of the African bush. Attendees will then proceed to the Botsalanong Boma, a unique open-air African banqueting venue, which with a huge bonfire and African dancers will make it an unforgettable experience under the stars for all. A banquet will take place on the evening of 14 July in the King's Ballroom, where attendees will experience true African abundance!

ACCOMMODATION

Sun City offers four individual on-site hotels. The exclusive Palace of the Lost City is ideal for those looking for a truly luxurious 5-star getaway. The Cascades is an upmarket luxury hotel set in beautiful gardens featuring pools and waterfalls. The Soho was the very first hotel established and offers a modern look and feel. The Cabanas is the ideal affordable accommodation to suit active families and budget travellers. The nearby Kingdom Resort offers neat and affordable accommodation with regular shuttles to Sun City that will be made available for the conference.

ISIF VISION STATEMENT

The International Society of Information Fusion (ISIF) is the premier professional society and global information resource for multidisciplinary approaches for theoretical and applied INFORMATION FUSION technologies. Technical areas of interest include target tracking, detection theory, applications for information fusion methods, image fusion, fusion systems architectures and management issues, classification, learning, data mining, Bayesian and reasoning methods.

ISIF Journal of Advances in Information Fusion (JAIF)

The Journal of Advances in Information Fusion (JAIF) is the flagship journal of ISIF. JAIF is an open-access, peer-reviewed, semi-annual, archival journal published electronically and distributed via the internet. JAIF was founded in July 2006. The journal is indexed at SCOPUS, free for authors, and freely available for readers at http://www.isif.org/journals/all. Authors are invited to submit both regular papers as well as short correspondences describing advances, applications, and new ideas in information fusion, both theory and application. Authors of papers presented at our annual International Conference on Information Fusion are strongly encouraged to consider submitting expanded versions of their papers to JAIF. Manuscripts can be submitted at http://jaif.msubmit.net.

Perspectives

JAIF Editor-in-Chief (2020–2022):

Stefano Coraluppi
Systems & Technology Research LLC, USA

Administrative Editor:

David W. Krout
Applied Physics Laboratory—University of Washington, Washington, USA

ISIF Society Officers

President	Paulo Costa
President Elect	Simon Maskell
Treasurer	Chee-Yee Chong
Secretary	Simon Maskell
VP Communication	David W. Krout
VP Conferences	Lance Kaplan
VP Membership	Anne-Laure Jousselme
VP Publications	W. Dale Blair
VP Working Groups	Darin Dunham
Perspectives EIC (2018–2020)	Roy Streit

ISIF Board of Directors

Term (2018–2020)	Fredrik Gustafsson
Term (2018–2020)	X. Rong Li
Term (2018–2020)	Zhansheng Duan
Term (2019–2021)	Kathryn Laskey
Term (2019–2021)	Felix Govaers
Term (2019–2021)	Simon Maskell
Term (2020–2022)	Pieter De Villiers
Term (2020–2022)	Murat Efe
Term (2020–2022)	Wolfgang Koch

Past ISIF Presidents

2019	Paulo Costa
2017–2018	Lyudmila Mihaylova
2016	Jean Dezert
2014–2015	Darin Dunham
2013	Wolfgang Koch
2012	Roy Streit
2011	Joachim Biermann
2010	Stefano Coraluppi
2009	Elisa Shahbazian
2008	Darko Mušicki
2007	Erik Blasch
2006	Pierre Valin
2005	Dale Blair
2004	Chee-Yee Chong
2003	Xiao-Rong Li
2002	Yaakov Bar-Shalom
2001	Pramod Varshney
2000	Yaakov Bar-Shalom
1998–1999	Jim Llinas

Welcome to the 25th International Conference on Information Fusion in Linköping, the Fusion Capital of Sweden!

Linköping has for the past 100 years been the aerospace capital in Northern Europe, with two airfields, aircraft manufacturing industries, air force training and a whole ecosystem around this, including the C3ISR branch of the Swedish Defence Research Agency (FOI), and the Swedish Air Force Museum. You will find a beautiful city center with walking distance between the conference venue and the hotels.

Linköping University was founded in the 1960s as an engineering university to support Swedish automotive and other industries, but it has sparked the whole city to grow over 150% since then. The Science Park has 350 companies and 7000 employees, with focus areas in automotive safety, telecom, MedTech and IoT. Linköping University with its EE, CS and ME departments, has contributed hundreds of papers to FUSION over the years. It is part of the 500 million USD research program WASP, focusing on autonomous systems (AS) and machine learning (ML). Aside from the traditional scope of the conference, FUSION 2022 will also highlight AS and ML in its plenaries, tutorials, special sessions and exhibitions.

Linköping is an historic place, dating back more than 700 years. It hosts the bishop of one of Sweden's first and most important dioceses (1107) and a beautiful 800 year old cathedral. The cathedral, nearby Vreta Abbey, and the town of Vadstena (home of Saint Bridget) makes it a popular pilgrimage destination. In July you will enjoy nightless days, with 18 hours of daylight and 6 hours of dusk, and an average high of 22.6 degrees Celsius.

Please visit www.isif.org/fusion2022 for more information!

Fredrik Gustafsson and Gustaf Hendeby
General Chairs of Fusion 2022