Operations Research Society of South Africa

At the Forefront of Analytics in Africa

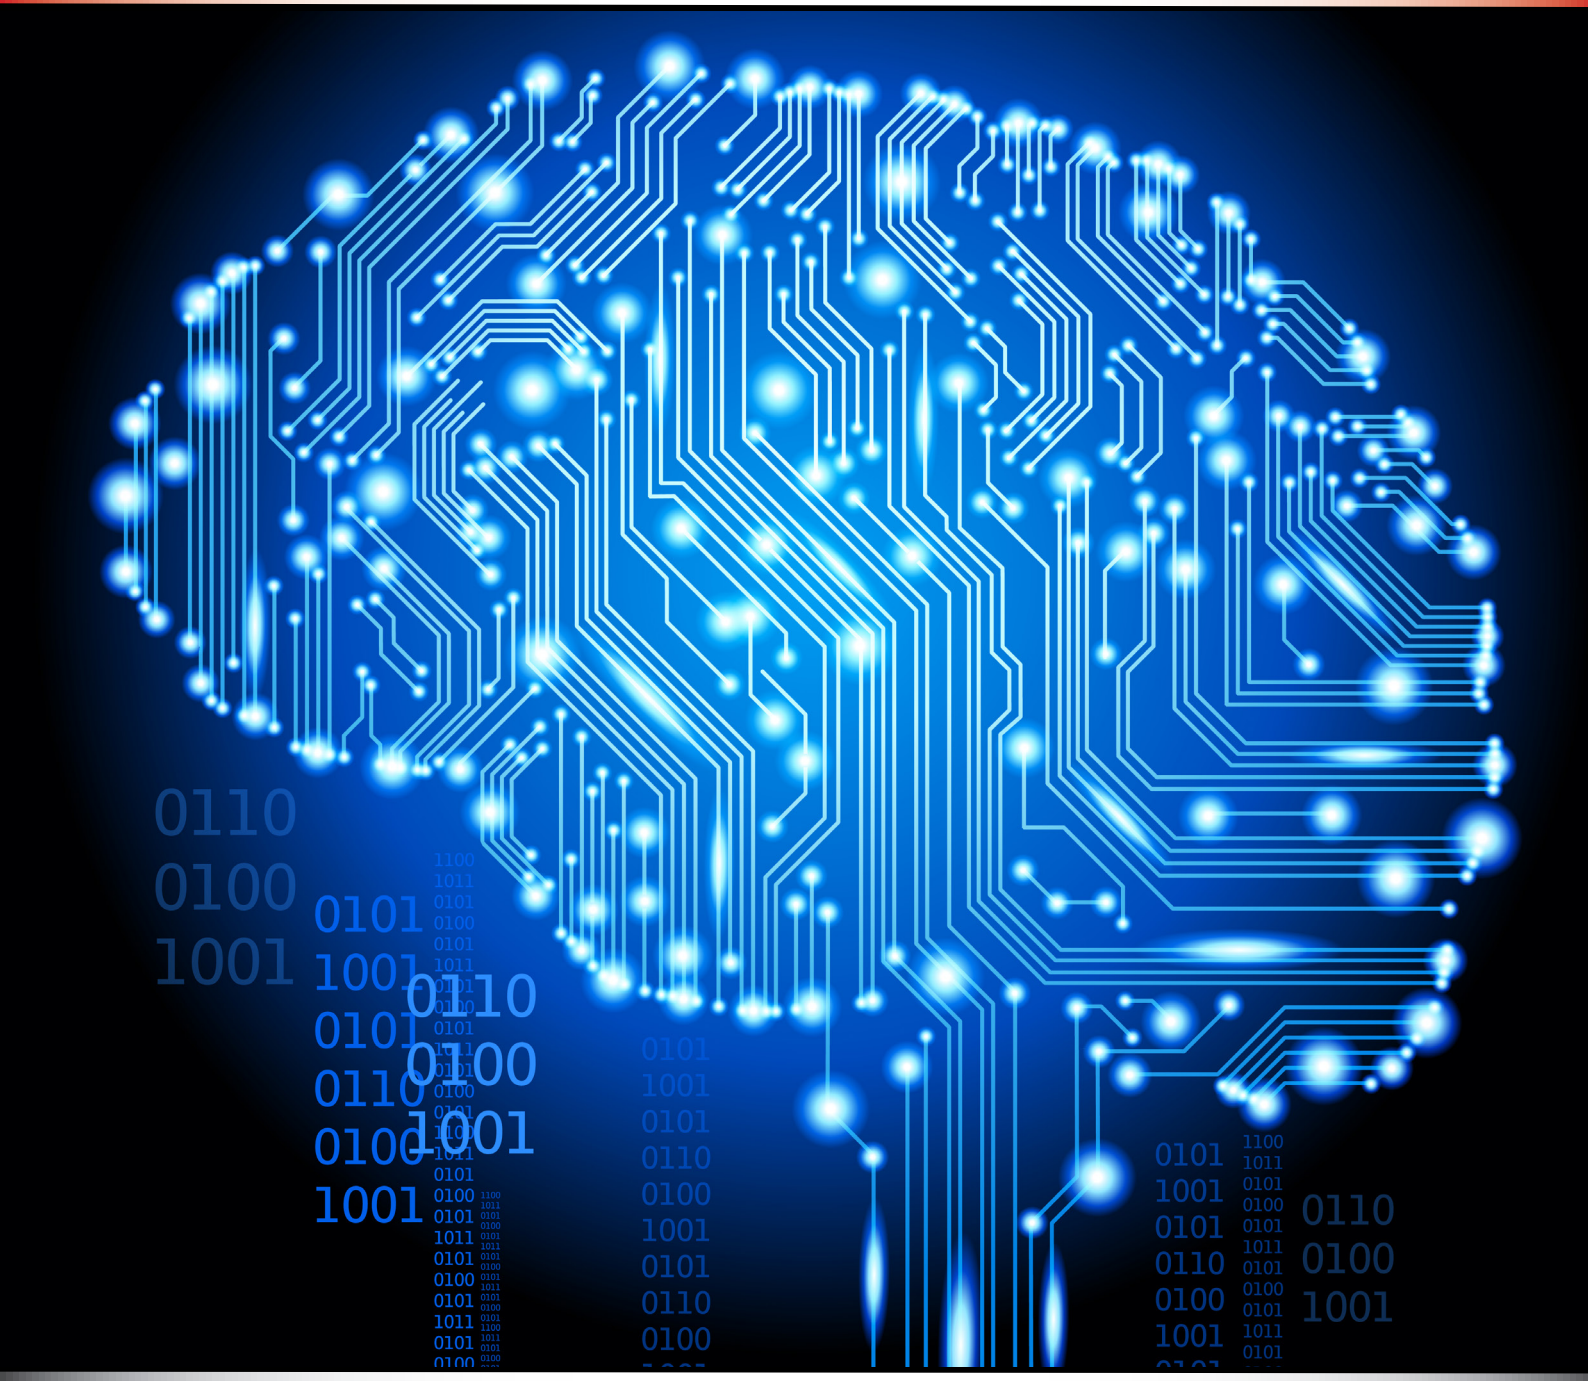

ORSSA Newsletter March 2016

www.orssa.org.za

45th ORSSA Annual Conference

11–14 September 2016

An advanced warm welcome to the 45th Annual Conference of the *Operations Research Society of South Africa* (ORSSA). The conference will be hosted by the Western Cape Chapter of ORSSA, and held at Lanzerac Wine Estate in Stellenbosch from 11-14 September 2016. The theme of the conference will be announced in due course.

The conference will open with an optional tutorial on Sunday afternoon and a welcome reception on Sunday evening, September 11th and will close at lunchtime on Wednesday, September 14th. Participation over the full spectrum of Operations Research is encouraged, including papers of a more fundamental nature, those on the application of Operations Research techniques in business and industry, about topical issues in Operations Research, and about the philosophy, teaching and marketing of Operations Research.

Following the successful introduction of published conference proceedings in 2011, authors will again have the choice of either (a) only presenting papers orally at the conference, or (b) submitting full papers (which will also be presented orally at the conference) for inclusion in the peer-reviewed conference proceedings. Registration, and submissions of abstracts and full papers open on the **14th of March 2016.**

Delegates are responsible for their own travel and accommodation arrangements. Lanzerac Wine Estate is recommended, as the Society has arranged competitive rates for delegates at the venue.

Lanzerac Wine Estate – http://www.lanzerac.co.za/

The conference keynote speakers will be Elena Fernandez, Erwin Pesch and Paul Fatti.

Elana Fernandez (opening keynote)

Erwin Pesch (mid-conference keynote)

Paul Fatti (closing keynote)

Lanzerac Manor House

Important Dates

14 March 2016	Early bird registration & abstract/paper submission opens
11 April 2016	Abstract submission closes for reviewed papers
18 April 2016	Notification of acceptance of abstracts of reviewed papers and go-ahead to submit full papers for peer-review
23 May 2016	Submission of full papers for inclusion in the conference proceedings closes
15 July 2016	Abstract submission closes for oral presentation of all papers
23 July 2016	Notification of acceptance of reviewed papers for proceedings
24 July 2016	Notification of abstract acceptance for non-reviewed papers
29 July 2016	Early bird registration closes
14 August 2016	Cut-off for qualification of reduced room rates at the hotel
21 August 2016	Registration closes

Please visit the ORSSA website and click on the link *ORSSA 2016* for more information:

www.orssa.org.za

From The Editor

By Bernie Lindner (berndtlindner@gmail.com)

Bernie Lindner

Brian van Vuuren

Welcome to 2016, a bit late, but this is our first letter of *news* for the year. *News* from the *new* President, Winnie Pelser, can be found on page 2. The *new* Executive Committee for 2016 can be found on page 3.

This edition is focussed on Artificial Intelligence or AI. Apparently the term "AI" is much more popular according to Google Trends (see graph on page 10).

The main article (page 3) is on something that we should probably start thinking of sooner rather than later it seems, Machine Ethics. David Clark describes a field I did not know was already so well founded/researched. The piece also includes a fictitious story by Dave.

Marc Hatton discusses "What exactly is AI" on page 10. Lastly on the AI front, Robert Bennetto describes his experience (page 8) of the CIS conference he attended in Cape Town last year December. It seems it was almost as good as last year's ORSSA conference organised by the Johannesburg chapter chair.

Some pictures and information on the Western Cape Chapter's latest event is on page 2, more information and photos are available on our Facebook and Twitter pages, please visit, like and follow these pages. The newsletter ends on page 11 with a review by Hans Ittmann on the book "Handling Societal Complexity".

Thank you to Brian van Vuuren for helping edit this edition. Please feel free to send me any articles to be included in the newsletter. Also, any questions and suggestions are welcome.

Until June all the best

Bernie

Features Page

SOCIAL MEDIA

- **Facebook:** Please visit (and like!) our page at *www.facebook.com/ORSocietySA*
- **Twitter**: @_ORSSA
- **LinkedIn:** Please visit our page at *www.linkedin.com/company/the-operations-research-society-of-south-africa*

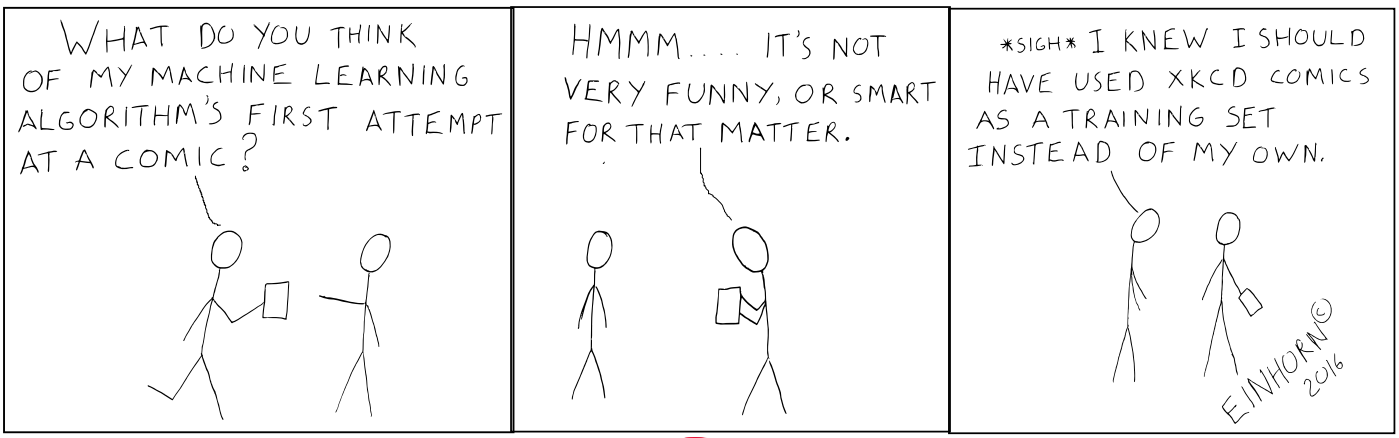

From The President's Desk

By Winnie Pelser
(winnie.pelser@gmail.com)
ORSSA President

Winnie Pelser

Dear ORSSA members, the year 2016 is already well under way and the festive season and holidays are but a memory. I would like to wish every member of the Society the very best for the remainder of 2016. The prospects for a year full of interesting activities and opportunities for ORSSA and *Operations Research* (OR) lies ahead.

I want to thank you for the privilege of entrusting me with the responsibility of president of our Society. I am humbled by your trust in me and will do everything in my power to live up to the expectations of the position. I particularly want to thank Hennie Kruger, the outgoing president, for his very significant contribution during the past two years. He will fortunately continue to serve as Vice-President for another year. Hennie, I shall certainly lean on you for guidance.

On behalf of the Society, I would like to thank every member of the *Executive Committee* (EC) of 2015 for their time, contribution and dedication in each portfolio. Each member committed their time and attention in a most professional manner despite work and other pressures. For those that are experiencing the end of the Financial Year pressure, good luck.

Some members of the EC had to resign at the end of 2015. They are Tiny du Toit (Treasurer), Margarete Bester and Louzanne Oosthuizen (additional members). I want to thank each of the members for their dedicated service to the Society. New members joining the EC are Isabelle Nieuwoudt (Treasurer), Brian van Vuuren (additional member) and Thorsten Schmidt-Dumont (Western Cape Chapter Chair). Thorsten takes over from Danie Lötter, who remains on the EC as Newsletter Business Manager. I want to thank these new members for their willingness to serve the society. The entire Executive Committee of 2016, with their respective portfolios, can be found on the webpage.

The year 2016 will certainly be full of opportunities. A highlight on the calendar is our annual conference during September, which will be organised by the Western Cape chapter (see the conference advertisement). In the trying economic and climatic circumstances we are currently in, operations researchers can contribute to alleviate the situation. This year, we hope to see more members who will endeavour to write the Certified Analytics Professional (CAP) international examinations.

I am looking forward to working closely with the Executive Committee during 2016, ensuring that ORSSA is and remains the professional and vibrant home of all the operational researchers in South Africa. I am also looking forward to join forces with all the ORSSA members to realise our operational research ideals and I want to invite everyone to contact me or any other member of the Executive Committee with suggestions and ideas to improve ORSSA's services and activities even further.

With best wishes. Alles van die beste.
Winnie Pelser

Western Cape Chapter Event

by Berndt Lindner at Stellenobsch University, Industrial Engineering (berndtlindner@gmail.com).

On 10 February 2016, the Western Cape Chapter of ORSSA held its first event of 2016 with a visit from Professor Emeritus Paulo Toth from the University of Bologna, Italy. Professor Toth gave a presentation on *Rolling horizon approaches to the aircraft sequencing problem.* Visit the ORSSA Facebook page for more about this event.

ORSSA WC Chapter Chair, Thorsten Schmidt-Dumont and Prof. Toth

CHAPTER EVENTS

Is your chapter hosting exciting events? Please send a short event summary and some pictures to the editor at *berndtlindner@gmail.com* to be featured in the newsletter.

The 2016 ORSSA Executive Committee

President:
Winnie Pelser

Vice-President:
Hennie Kruger

Secretary:
Lieschen Venter

Treasurer:
Isabelle Niewoudt

Database Manager:
Dave Evans

Marketing Manager:
Mark Einhorn

ORiON Editor-in-Chief:
Stephan Visagie

ORiON Journal Manager:
Martin Kidd

Newsletter Editor:
Bernie Lindner

Newsletter Business Manager:
Daniel Lötter

Webmaster:
Jacques du Toit

Additional Elected Members

Brian van Vuuren

Jan van Vuuren

Linke Potgieter

Angela Rademeyer

Chapter Chairs

Johannesburg:
Robert Bennetto

Pretoria:
Quintin van Heerden

Vaal Triangle:
Fanie Terblanche

Western Cape:
Thorsten Schmidt-Dumont

Kwazulu-Natal:
Aderemi Adewumi

Co-opted External Liaison Representatives

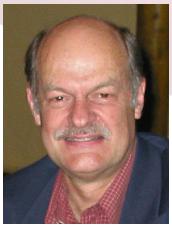

IFORS Representative:
Hans Ittmann

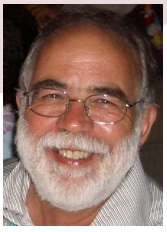

EURO Representative:
Theo Stewart

Ethics, Machines and Decisions

by David Clark at OPSI systems, Johannesburg (david.clark@opsi.co.za)

It was the start of another rainy overcast work day. Autonomous Transporter DB5 was two minutes behind schedule. Unware of the early morning malaise that most human drivers felt sitting in semi-autos on either side of its dedicated lane. To say 'drivers' would be to assign them more responsibility than what was usual. Most people rarely needed to intervene these days, even in vehicles not licenced as fully automated. In any event they would have struggled to see out of the water coated windscreens. DB5's vision system cut through the hazy weather with ease, augmented with feeds from sensors both on the roadway, on other ATs in front, and the public video broadcasts streamed by traffic drones. It formed a seamless and precise view of its surroundings that was impossible for humans to comprehend. Its real-time verification, monitoring and reflection system showed no anomalies, apart from having to override the on-time priority of the drop-off with the dynamic legal speed limit set on the current stretch of road. Still, its routing algorithms actively searched for ways to make up the time, and one option had just shown a possible improvement within an acceptable confidence interval. It was a lane shared by manually driven vehicles, but it looked clear enough. It took the next turn right and headed down a road less travelled.

Hello, World

Developers, when faced with learning a new language, often write a very simple application called the 'hello, world' program[1]. The exercise of creating the program not only requires some basic understanding of the language's syntax and libraries for writing to an output device, but also tests that the development environment itself is working. In Linear Programming one might argue that the Diet problem is the equivalent[2] 'hello, world' exercise, used to test the simplex method by Dantzig in 1947. In ethics, there is the trolley problem.

The Trolley problem was first proposed by Phillippa Foot in 1967 as a thought experiment in ethical decision making. Numerous versions of the problem have been explored and compared since, with equally interesting names, such as 'the fat man', and 'the transplant'. The variant that follows is by Judith Thomson[3], called 'the bystander at the switch'. Crossing a bridge over railway tracks you notice a switch that allows a trolley to be guided from the main line to a side branch. On the main branch you notice five men working, and on the side branch is one worker. Suddenly a trolley thunders around the bend out of control heading for the five workers on the main line. There is no chance for them to escape being hit and there is no doubt that they will all be killed. However, you have the opportunity to throw the switch, diverting the trolley onto the side branch which will surely kill the one unsuspecting worker, but save the five up ahead. What do you do?

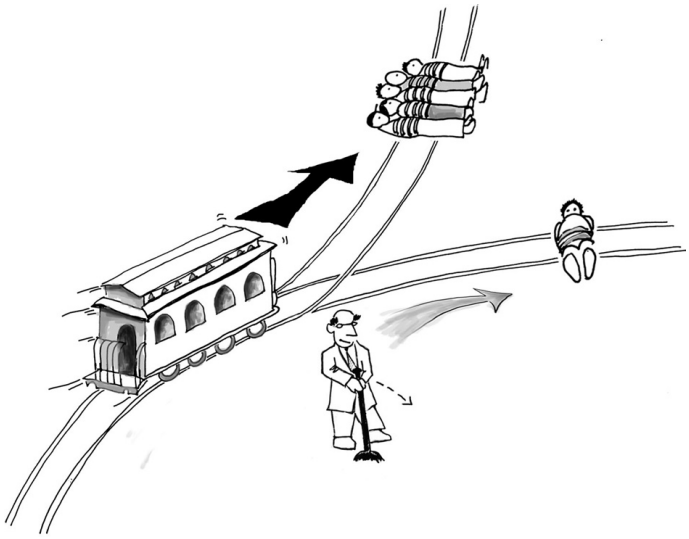

The Trolley Problem

The sensor array at the intersection was reporting erroneous data again. John, the team diagnostics lead had already confirmed that the health assessment unit had been breached by sensor hackers. This was allowing a steady stream of carefully crafted misinformation to flood the network. Information which thousands of agents relied on to perform everything from estimating the time to receive a parcel to opening channels in the traffic to allow emergency vehicles through. Most of the cybersecurity issues were simple acts of vandalism and opportunism. In this case it appeared the sensors where reporting a much lighter load on the roads than was the case, usually someone trying to influence the routing algorithms to divert traffic into the area. These days the more sophisticated units read data from so many sources that is was fairly easy for them to spot dirty data. In fact, that was how they had been made aware of the problem. The whole unit was scheduled for replacement this time. The sensor sat at the intersection of one of the few roads on which manually driven cars were allowed to operate. As the unit went down, the autonomous cars lost virtual sight of the old manuals that had no sensors on them. They would have to rely on their own vision tracking systems to see them now.

1 Supposedly the first use of this example is by Brain Kernighan for the BCPL programming language. It appears again in a Bell Labs 1974 internal memorandum by Kernighan titled "Programming in C: A tutorial."

2 Although in 1947 it certainly wouldn't have been classified as a simple problem, taking 9 clerks 120 man days to obtain the solution.

3 Judith Jarvis Thomson, *The Yale Law Journal,* Vol 94. No 6, May 1985, pp 1395-1415, The Trolley Problem

Machine ethics

What if, instead of you standing by the switch, an autonomous artificial agent controls how and when the switch is thrown. How do we make such agents choose the safer option, if such an option exists? Let's assume there is no one on the side track, but a sensor connected to the switching agent determines that there is possibly something on the track up head. Would we expect that it should throw the switch, should it also try and take into account the damage that may occur if people inside the trolley are injured as it takes the branch at its current speed? What if instead of one agent making such a significant decision, there are thousands of autonomous agents making what appear to be far less life and death decisions but that ultimately crash the stock market. These are the sorts of questions being studied as a subfield of AI known as Artificial Moral Agents. But before we can discuss these further, let's take a quick diversion and look at intelligence and rationality in the field of AI.

The Strong, the Weak and the Bounded

It does feel, at first glance, that a solid understanding of intelligence is required before we discuss the implications of ethical decision making. As humans we identify closely with being a species capable of intelligent behaviour, and have no intention of bestowing this sought after title on anything else[4] (least it turns out to have more of it than we would like). But, we don't seem to mind if things act intelligently. The ability to appear intelligent is known as the weak AI hypothesis, that is, machines are only capable of simulating intelligent behaviour. Strong AI, on the other hand, makes the claim that there is essentially no difference between human minds and intelligent machines. The term Strong AI was put forward by John Searle. And he went on to provide us with the famous Chinese Room Argument in "Minds, Brains and Programs" in the journal The Behavioural and Brain Sciences in 1980 in an attempt to show why the Strong AI position is fundamentally flawed. The debate, however, rages on, with Daniel Dennet as one of the foremost supporters of the Strong AI hypothesis.

The applied artificial research community, however, is not overly concerned with such deep issues. For most, if it looks intelligent, then it is intelligent, is good enough for them. And let's face it, they have bigger problems in just trying to get something to appear intelligent in the first place. Allan Turing in the 1950s cut through much of the confusion, with what many consider still to be the best practical answer, when it comes to trying to navigate through philosophical, psychological and neuroscience definitions of mind and intelligence. Until we have terms we can agree on, let's just test if we can distinguish humans from machines.

Current conversations with virtual assistants like Siri, Cortana and Google Now are unlikely to convince anyone that we are in danger of intelligent machines (although they do surprise you from time to time). But the film 'her' in 2013 is a good example of how easily we could end up changing our position from 'artificial' to 'real'. Not only does the main character, Theodore, become emotionally attached to his talking operating system, which names itself Samantha during configuration. But even as viewers there seems to be a point where we cross a boundary of thinking of Samantha as an algorithmic process to being more of an intelligent equal. Obviously the film is meant to take you on this journey, but it does raise interesting questions regarding our future relationship with ever more sophisticated processes, specifically at the end of the movie in which Samantha begins to exceed human intelligence.

It is also important to note that the AI community continues to refine the notion of intelligence so that the definition is precise enough to allow results to be compared between theory and practice. Stuart Russell describes four definitions of rationality that have been put forward in the history of AI and related fields.

Very briefly, to just give you a taste of the nature of the formal definitions, the first is P1: Perfect Rationality. Such an agent is always able to exhibit the best possible behaviour for a given environment. Theoretically interesting, but practically speaking, such an agent can never exist. Next is P2: Calculative Rationality. Here the formulism explicitly acknowledges that the agent is a machine M, running a program p. Such an agent would produce the same results as P1 if the machine it is running on is infinitely fast. As such, the definition is not useful in real-time worlds. Any

4 Some of us are still smarting from the being displaced from the centre of the universe and then discovering that we are simply one branch of a very large tree of life. Contemplating that our tree may be one of many is still a bit of a sensitive topic.

system that selects suboptimal actions, even in the presence of time constraints falls outside of calculative rationality.

More realistic definitions of rationality include, as defined by I J Good in 1971, "the maximization of expected utility taking into account deliberation costs". Good called this Type II rationality, and it falls under P3: Metalevel Rationality. Here, the cost of acquiring information and the cost of computation is taken into account, but it is still very seldom possible in real environments as it requires that computations and information are controlled optimally, something rarely possible in practice.

Finally we have P4: Bounded Optimality, and the notion of Asymptotic Bounded Optimality which comes the closest to bridging theory and practice. This work came out of investigations in resource-bounded rationality, where the feasible set of all programs that can be run on machine M is taking into account when trying to calculate an answer. According to Russell, the formulation of Bounded Optimality is consistent with "the importance of resource limitation faced by relatively tiny minds in large worlds", and that it also "takes into account the limitations on computational resources that are presumably responsible for most of the regrettable deviation from perfect rationality exhibited by humans."

The old tour bus was a hit with visitors. A rare chance to be driven in a vintage vehicle by an actual person. The driver, on the other hand, was not feeling terribly well and lack of coffee and a very late night didn't help matters. He noticed he had a green on the traffic light up head crossing one of the bustling fully autonomous lanes. He put his foot down a bit harder to make sure he would make it. The bus lurched forward as it picked up speed and a child squealed in delight. The sound of an old combustion vehicle approaching the intersection made John look up from his tablet. People, he thought, shaking his head, always in a rush to beat the light.

Moral Agents

A quick recap. For practical purposes we can ignore the concept of Strong AI, and work within the framework of bounded optimality for rational decision making. And given the example of the trolley problem it seems clear that some of these decisions look like those that we would call ethical decisions if made by a human. As with the question of intelligence, Artificial Moral Agents, come in two distinct philosophical flavours. Implicit Ethical Agents, behave in an ethically responsible manner (from society's viewpoint), while full ethical agents are able to extend moral reasoning and justify new rules of moral reasoning. It seems reasonable to assume that until we have agents that exhibit high levels of intelligence we are unlikely to have to worry about full ethical agents for some time. But before we move on too quickly we should note that many prominent researchers believe that even if the probability of a super intelligent agent arising in the next few decades is low, that the impact it would have warrants concern now. Good, in 1965, had already coined the term 'intelligence explosion' as a warning of the consequences of a positive intelligence feedback loop. A superintelligence would clearly be capable of reasoning and extending any ethical framework it was presented, but whether this reasoning would be understandable by humans, or even in our interest is questionable.

Some solutions that have been suggested are *Boxing the system*, sealing it off from interacting with other environments or creating *Oracle AIs* that are limited to pure question-answering systems. But these solutions assume that a superintelligence could be created within such constraints in the first place. And it seems unlikely that such constraints could be imposed post the event of a singularity. Another more practical solution has been to use *inverse reinforcement learning* (IRL) to extract a reward function from observed optimal behaviour. The idea being to let an agent learn the correct ethical behaviours by watching humans solve ethical problems. The problem here should be obvious, humans as a rule are notorious for making irrational and inconsistent decisions. Learning from us may not be in our best interests.

The massive double tanker held its speed as it approached the intersection. It based its velocity on the timing of the lights which it could access. The lights themselves were unnecessary anyway, it received the intersection status data directly from the traffic controllers. A warning flag switched on in the unit's decision system, it was receiving no additional sensor data around the crossing and would be blind to vehicles approaching from its right. It would have clear right of way though, still as a precaution it dropped its speed slightly.

The Sorcerer's Apprentice and Paperclips

Many have pointed out that it is not Superintelligences we should be worried about just yet; it's the somewhat less intelligent agents that have already been tasked with decision making in areas that affect us already, or will affect us in the near future. Implicit ethical agents can probably be best thought of being obedient to the rules given to them. In this sense, the current focus on issues of ethics in AI presents us with few new dilemmas which aren't already present in the age of modern technology. AI and automated systems can, however, greatly amplify the risks that for now may be more isolated simply due to lack of resources that the algorithms have access to.

Anyone who has seen Fantasia knows how quickly things can go wrong when access to powerful forces coupled

with a limited understanding of how to control them can quickly spiral out of hand. The classic Disney film is based on a poem by Goethe called "The Sorcerer's Apprentice." Nick Bostrom's paperclip example shows that even benign sounding goals may have disastrous consequences when a super intelligent agent is given the goal of making paperclips and then proceeds to transform the entire earth into paperclip manufacturing plants. What happens though when someone tells a self-driving car to get them to their destination as quickly[5] as possible? Or tells the system it's an emergency. In practice this ability to game the system would be quickly closed down, one could almost view such a request as a hacking attempt, which would be covered by the security, integrity and governance issues of such systems.

And we are all familiar with recommendation engines, eagerly showing us what we should be buying next. Is there a point where those recommendations should be vetted for harm? How would we view a recommendation for someone who searches for alcoholics anonymous being bombarded by adverts for whisky and vodka? Are recommendation engines for buying products simply fuelling our weaknesses and addictive behaviours for the benefit of sellers? And what happens when recommendation engines start advising doctors about medical treatments and drug interactions, donor matching and medical triage decisions. What assurances should society ask for in order to be protected from unethical behaviour in how the results are generated?

I'm not going to make it, the driver of the bus realized too late. Instinctively he pressed down even harder on the accelerator, hoping to at least make it through on caution. It was a bad call. The light flicked to red a full second before the bus entered the intersection. He heard the shuddering sounds of the pulsating ABS on the tanker before he saw it loom large out the corner of his eye.

Into Battle

While the classic image of an army of robots descending on us will likely remain the realm of science fiction, there is growing concern over the design and deployment of l*ethal autonomous weapons systems* (LAWS). The UN, in an April 2013 report called for an immediate moratorium on the development and deployment of "lethal autonomous robots". On the semi-autonomous front, however, everything from inter-continental warplanes, to robotic machine guns, intelligent mines and mine-clearing systems are being developed in a fascinating if not disconcerting arms race.

John Aquilla, who developed the concept of netwar or swarm-tactics, notes that while there has been significant reluctance to embrace fully autonomous military action by machines, that soldiers do become attached to their so-called machine buddies. With instances of drones being awarded medals, and ceremonial burials being performed when 'killed in action'. Increased autonomy of weapons can also lead to a decrease in accountability of actions that occur on the battle field.

And back again

Max Tegmark in a recent book called 'The mathematical universe', pointed out that over the centuries we have consistently underestimated the size of our reality. We underestimated the size of the earth, the distance to the moon, the distance to the sun, the distance to the nearest stars, and the size of the galaxy. In his book he goes on to show, how in his opinion, we have vastly underestimated the size of the universe and the reality in which it is embedded.

Russell in a talk subtitled A "New Dawn" for Artificial Intelligence, reminds us that in the 1950s AI, Control Theory, Operations Research and Statistics failed to find common ground as they had no common mathematical formalism. However, since the 1990s work in Bayesian networks and statistical machine learning has started to reintegrate some of these fields. The goal of unifying these disciplines of learning, reasoning, planning, perception, optimization and language could lead to the light at the end of the tunnel. What have we underestimated in our quest for intelligent machines? What decisions will we need to make as we head down the tunnel? And will we make the right ones?

To say there was a tug of war going on in DB5s control units would be an understatement given the fact that every computational unit was fully engaged. The AT had already calculated a possible impact between the bus and the tanker as it approached the intersection. The bus heading towards it from the other side had been flagged as a manual passenger vehicle. DB5 could certainly stop well before the potential accident site; there would be no harm to its payload, vehicles behind it, or the workers it registered opposite a virtual barrier. However, there was an alternate option. With a high degree of confidence it could skip the red light and hit the tanker with enough force to deflect it sufficiently from a full side impact with the bus. Payload protection and on-time arrival objectives could be suppressed in emergencies, but overriding the red light traffic sensor was proving to be difficult. Its prediction algorithms noted that all other autonomous traffic would immediately take evasive action or stop safely, and so the only uncertainty was whether the workman would be accidently struck. It didn't have the time or computational resources to play out the simulation far enough ahead. But it could save the bus from a fatal impact. The decision window closed. DB5 accelerated.

5 Uber drivers are often asked this, at least they have the option of rating their passenger with 1 star.

Some interesting further reading and sites:

- One Hundred Year Study on Artificial Intelligence - https://ai100.stanford.edu
- Rationality and Intelligence: A Brief Update by Stuart Russell - https://www.cs.berkeley.edu/~russell/papers/ptai13-intelligence.pdf
- Stanford Encyclopaedia of Philosophy http://plato.stanford.edu/entries/computing-responsibility/
- Ethical Issues in Advanced Artificial Intelligence - Nick Bostrom http://www.nickbostrom.com/ethics/ai.html
- Artificial Intelligence: A Morden Approach, 3rd Edition – Peter Norvig and Stuart Russell.
- Ethical Machines, Philosophies of Man and Machine, I.J. Good - http://aitopics.org/sites/default/files/classic/Machine_Intelligence_10/MI10-Ch29-Good.pdf
- What to think about Machines, Edited by John Brockman, 2015

CIS 2015 Conference

by Robert Bennetto at OPSI systems, Johannesburg (robert.bennetto@opsi.co.za)

Robert Bennetto

Cape Town was the host city to the IEEE Symposium Series on Computational Intelligence (SSCI) in December 2015 and I was fortunate enough to attend. The event was hosted by the Computation Intelligence department at the University of Pretoria and straight off the bat I can commend them for a job well done as the Local Organising Committee. There was a great turnout from the South African university students and they had loads of experts in the field attending from abroad. The opening act of the show was Marco Dorigo, famous for his *Ant Colony Optimisation* (ACO) metaheuristic that gained huge popularity in the early '00's. Prof Dorigo was receiving the equivalent of a lifetime achievement award at the conference and gave an excellent plenary on his new research which has to do with swarms of robotic robots, working together in differing capacities, to solve problems. He mentioned that he had moved his research away from ACO, but the new research looks to be an undeniable physical manifestation of ACO. The emphasis of his new research is on scalability, parallelism and fault tolerance.

It's tricky to really summarise what a Computational Intelligence conference is all about. The attendees are multi-disciplinary, spanning computer scientists, engineers and statisticians. One could almost mistake the composition of attendees at the conference for a Operations Research conference. The key difference is the emphasis on *learning* techniques — a feature that originally attracted me to the conference. Machine Learning takes on many flavours and depending who you're talking to, and the context they have, often reveals what they perceive machine learning to be. The treatment of the the topic of learning, how it's defined and contextualised in different problem domains was really impressive. These kinds of events are very useful if you're looking to get a temperature reading in a field to see if it's worth going deeper or not.

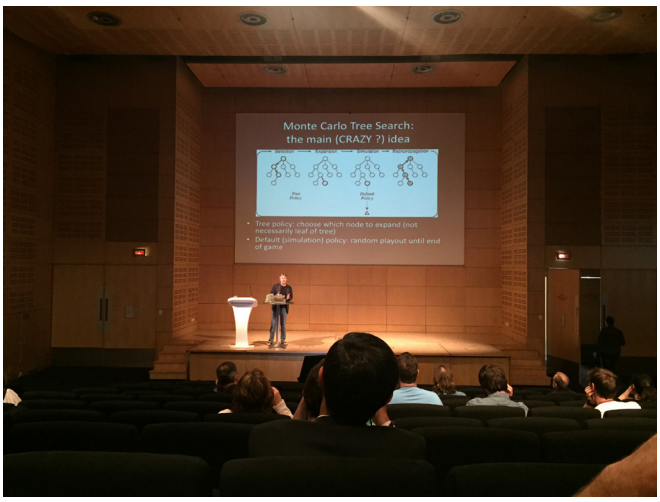

Simon Lucas presenting on Monte Carlo tree search in reinforcement learning for computer games.

There were a few talks that stood out for me on the learning side. Many of them involved applications in robotics, which probably ties back to their roots in control theory. A group of post-docs in Sweden have *taught* a robotic arm to play ping pong. It sounds silly but it's really entertaining to watch the footage of them "teaching" the robot to play. The concept is quite simple; hit the ping pong ball over the net towards the robotic arm and the robot gets scored on its ability to connect with the ball and then return it to the student doing the serving. The research for this particular group was focussed on being able to teach the robot in as few iterations as possible - maximising the amount of information squeezed out of each trial. They were able to get the robot playing as well as the student in around 50 trials which is a surprisingly low number considering the complexity of the task. It was mentioned that the student wasn't particularly good at the game, so perhaps not the best teacher, but they were also in the process of learning.

It didn't come as a shock that the student paper that won the coveted best-paper award was particularly excellent. The topic was related to a field close to my heart (as a statistics-centric person), what to do with non i.i.d data. Independently and identically distributed data can be hard to find in a relational database (because rows are often not independent due to the structure of the data) and this paper addresses how to adjust your predictive models to

compensate for this heterogeneous structure[1].

Barbara Hammer's group presented great work on determining feature relevance (for $p > n$). The statistics framework binds us to only estimate as many parameters as the data will allow, the degrees of freedom. Here $p > n$ refers to having more parameters to estimate (p) than you have data points observed (n). Now you might say: *c'mon Rob, that's such a 90's problem, I've got big data and more observations that I know what do to do with - I'll happily share some with you if you're a bit short*. And you'd be allowed to think that you're safe from this problem, but I suspect that's because you haven't been shown the whole picture yet. At the heart of every big data set lies a fundamental issue. If you have a lot of data, you can probably get more. More data about stuff you don't know about — co-variates, or other variables that may correlate. The joy of big data is not only that you have a massive number of rows but that for any data set you can add more columns than you have rows. This is largely because of the amount of things being measured and the number of open data sources now available (the Internet of Things is one of the contributors). More columns than rows is really the statisticians' worst nightmare; how do we now decide which columns are important (or significant) and which are just noise (or subject to over fitting - the dreaded bias in the model)? Hammers' group are heavily focussed on these problems; see *Inferring Feature Relevance from Metric Learning* by Hammer, Biehl, Mokbel and Schulz.

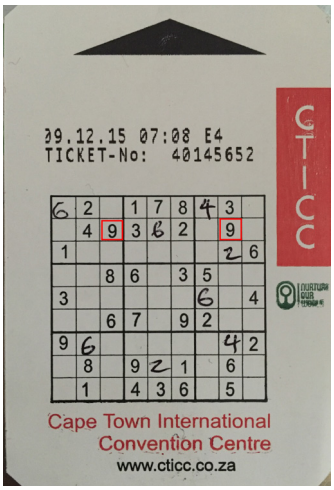

Welcome to the CTICC parking lot. Where Sudoku puzzles are not taken seriously one bit.

The other highlight of the conference for me was enjoying the keynote talks on reinforcement learning. The technique is exploding in popularity largely due to its abstraction from the problem domain. At the time of the conference, a Reinforcement Learning algorithm had just beaten the first *Go*[2] professional (the European champion) in five games to zero. This may sound like a flashback to the 90's where DeepBlue beat Kasperov in Chess, but there's a subtle difference. DeepBlue was programmed with specific strategies to play chess - the machine was shown the environment, explained the rules and given tools to evaluate the quality of the current game state. *AlphaGo,* the equivalent of DeepBlue for Go, was not programmed specifically to play Go - it had to learn to play the game through trial and error (with some help from historic games to get started). It turns out that computers can learn to do tasks quicker than we can program them to do the task, which may lead to a few questions about our longevity as professional programmers, but that's another discussion. The seminal talks at the conference covered a wide range of application areas in Reinforcement Learning ranging from playing computer games (in virtual worlds) to learning to swing a ball on a string into a cup (physical world, robotics). One of the most common questions asked at the conference was: "does it scale?" Practitioners were interested in whether techniques being showcased were application ready and what the computational complexity was of the methods.

Xin Yao (the current president of the IEEE CIS) is one of the foremost experts on co-evolutionary techniques to solve large, non-linear, optimisation problems. The theme of his plenary talk was around how to identify the possible (approximately) independent partitions in large problems so that they can be solved in a distributed fashion. The intimate nature of the conference gave me an opportunity to really talk his ear off about some of the pro's and con's of the approach on combinatorial problems. Fortunately for me, Prof Yao is passionate about his research and was eager to share his insights and own concerns.

Prof Malik gave a talk on research that had spanned 25 years. It transpires that control systems for power station turbines are pretty important. The more efficiently a stabiliser can react to instability in the turbine, the more efficient the turbine can be in generating power. The problem was originally solved using heuristics that were designed to give quick, predictable responses. What made the talk interesting is that all these years later, they've reviewed their original approach and found that with modern hardware, they are now able to take a data-driven approach to solving the control problem. Using a Neural Network as the underlying learning model they were not only able to match the original heuristics, but were now also able to significantly outperform them (in terms of effective operations to correct for turbine instability as rapidly as possible). It was interesting to see how the engineers toolbox has continually matured to take advantage of the state of the art in the field.

I had a positive experience at the conference and would recommend that if you're ever on the fence about going to a CI IEEE Big Data Conference, that you take the plunge. The people were knowledgable and the topics relevant, especially so if you're looking at doing interesting things in the Machine Learning or Optimisation space. The Symposium Series on Computational Intelligence is being held in Greece in December 2016.

1 Oliver Schulte and Fatemeh Riahi; Model-based Outlier Detection for Object-Relational Data, 2015.

2 Go is the game of Chinese chess - played on a *19* x *19* grid with white and black stones.

What exactly is AI?

by Marc Hatton at Adapt IT Holdings, Cape Town (hatton.mn@gmail.com)

Marc Hatton

Is AI a buzzword to score higher hit-rates? Including Artificial Intelligence as a key word in your publication (or even – hold your breath – the title of a published work) would increase views, but what qualifies work as being in the field of AI? On a side note, the graph shown below (sourced from Google Trends) illustrates the relative search phrase popularity over time between AI and Artificial Intelligence. A word of advice: If you decide to reference Artificial Intelligence / AI, use the sexier form – AI

Artificial can be defined as something that does not occur in nature or a copy of what can be found in nature.

Intelligence. How to put a yardstick against the ability to learn, accumulate and apply knowledge and skills? Furthermore, why is artificial intelligence likened to human-like intelligence? As *Homo Sapiens* rules the planet due to its unparalleled complex brain structure (relative to what life forms that have been discovered), perhaps it's fair to measure artificial-intelligent features/behaviour against those of humans – comparing the created to the creator.

That still, however, does not answer the question of how to measure intelligence. In general, psychology, aptitude tests and achievement tests are used to measure potential development and mastery of a field. However, there are many outspoken critics who claim, amongst others, 1) they tend to be culturally bias, 2) some people perform well in test environments whilst others do not, 3) they do not test all areas of intelligence of which there are plenty (pattern recognition, creative, musical, linguistic, kinetic, problem-solving and arithmetic – to name a few). Besides, more practical and sophisticated tests are required (and exist) to test AI.

Artificial-intelligent features include autonomy, self-awareness, visual perception, empathy, speech recognition, and decision-making. Which combinations of these features, summed up, equate to artificial intelligence? The first machines that began to really catch the attention of the general public made use of large databases to reference information. However, reference and decision-making computers generally fall short when asked contextual based questions. In 1959, Alan Turing defined what he perceived to be a threshold to true intelligence – Natural Language Processing.

For those of you not familiar with the Turing Test, one or many judge(s)/interrogator(s) is(are) given the task of determining which player (player 1 or 2) is human and which is a super computer, through reading written responses to the judge(s) questioning. The test seems sensible, yet a heavy bias exists based on the questions asked.

Variations of the Turing Test exist such as: 1) the Feigenbaum Test whereby a computer is required to beat a subject matter expert in her/his field of expertise. 2) Total Turing Test which entails testing perceptual and motion-based aptitudes (requires Robotics). 3) The Reverse Turing Test where the roles are swopped and the judge is a computer, with the best example being CAPTCHA (frequently used on websites to distinguish between bots and humans). Futurist Ray Kurzweil predicts that machines will be indistinguishable from humans and pass the Turing Test by the year 2029.

The world has seen several computer achievements in recent years, such as: i) A computer beat the previous champions of Jeopardy! ii) 'Intelligent' personal assistants, for example Siri, are a form of fairly sophisticated speech recognition (someday these assistants might evolve to be companions – have you watched the movie Her?). iii) Deep Blue defeated world Chess champion Gary Kasparov in 1997 (credit goes

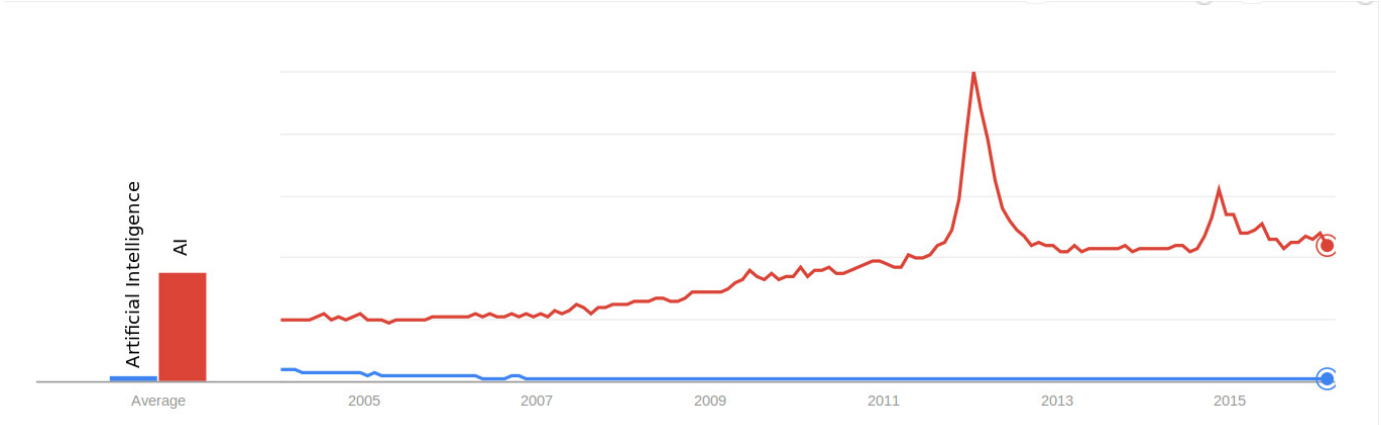

Google Trends' "interest over time" comparison.

to Gary Kasparov for winning the previous year – perhaps the machine was 'trained' in 1996). iv) Lamus (a computer) explored the creative side of the brain, through composing music that captures the attention of professional musicians.

The morning of writing this article I asked my true love Siri the following question: "Well Arsenal win the English Premier League this season?" to which Siri corrected her interpretation of my ascent to "Will Arsenal win the English Premier League this season?" I was hoping she could predict the future, but for one of the few times in my life she let me down. (Maybe the day of 'intelligent' personal assistants being companions has already arrived).

Stephen Hawking, Elon Musk and Bill Gates are openly against AI. "Success in creating AI would be the biggest event in human history," 'said' Stephen Hawking. "Unfortunately, it might also be the last, unless we learn how to avoid the risks. In the near term, world militaries are considering autonomous-weapon systems that can choose and eliminate targets." Professor Hawking added in a 2014 interview with the BBC, "humans, limited by slow biological evolution, couldn't compete and would be superseded by AI"

From a Machine Ethics point of view, a precise definition of AI is paramount to 1) being able to define laws and policies and 2) practicing them to enforce sound regulation and governance (we don't want evil robots to take over the world – we don't have a real world Del Spooner).

Nevertheless, as practitioners of Operations Research, I feel that we are lucky in that we get to venture into and explore the space of what I refer to as "the first couple of stepping-stones on a long flight of stairs which leads to eventual Artificial Intelligence", namely decision-making tools and optimisation algorithms.

Decision-making is one of the most important flags of intelligence and therefore I find OR tools exciting, amongst others: neural networks, mixed-integer programming, machine learning, simulation-optimisation techniques, and evolutionary algorithms (apologies for any awesome tools not listed – I'm only human after all).

Defining AI is not trivial. Even if/when a definition were decided upon, it would nevertheless be challenging to change the public's perception of AI because the phrase has been misappropriated by advertising agencies and re-

Book Review: Handling Societal Complexity

by Hans Ittmann, University of Johannesburg (hittmann01@gmail.com)

(This book review also appeared in the December 2105 edition of the IFORS newsletter)

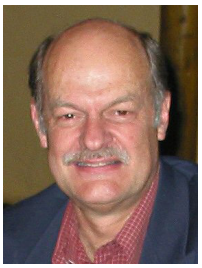

Hans Ittmann

Society worldwide is almost continuously confronted with disasters, complex issues and problems. Terrorist attacks, refugees from Syria and Africa, global warming, HIV/Aids or the world financial crises are just a few examples of complex societal problems. The recent terrorist attacks in Paris have had devastating effects not only on society but also on humanity as a whole. All of these are clearly very difficult problems. The framework for handling such societal complexity falls within the ambit of Soft OR. Complex societal problems are by nature almost impossible to solve and therefore the emphasis is on handling or addressing them. This book outlines the theory and methodology of societal complexity and proposes the COMPRAM Methodology (COMplex PRoblem HAndling Methodology).

The book consists of two parts, with the first part laying the groundwork for the eventual introduction of the methodology. The second part takes the reader through several examples where the methodology is applied. Within the first seven chapters, the theoretical foundation is outlined and described in great detail. The COMPRAM methodology is then introduced and explained. The focus in the second part of the book, comprising the last seven chapters, is on examples of methodology applications in health care, sustainable development, and terrorism, to name but a few. In each of these diverse problem areas, different aspects of the methodology are illustrated and emphasized.

In the first part, boundaries of complex problems addressed using the problem handling process is set. In chapters three and four, the methodological theory of societal complexity is developed almost from first principles in a logical and extensive manner. For example, a general definition of a

problem is developed: "something is called a problem when there are discrepancies between the actual or (near) future situation and the desired future situation and/or there is a lack of knowledge, and/or a lack of relevant data and/ or a lack of data". Every aspect related to a problem is discussed and this logically leads to problem handling with all the related issues. From this, a definition is provided for a complex societal problem, which is characterised as: a real-life problem, having a large and often different impact on different groups of society, is undefined or ill-defined, lacking in data and knowledge, concerns many different domains, with many actors/players involved, provokes emotions, is unique and has never been handled before. Human problem handling is the topic of chapter four. Aspects addressed here, among others, are the problem handling cycle, problem development and model conceptualisation, a wide variety of problem solving methods, possible interventions emanating from the results, and scenarios with their limitations.

Chapter five illustrates how computers, with a focus on *Group Decision Support Systems* (GDSS), have and could assist in the process. A real-life example using GDSS is presented, showing the importance of a methodology for handing complex problems. The next chapter goes into the different aspects required from the methodology, culminating in an exhaustive list of 26 conditions which should be considered, included or discussed in supporting the problem handling process. Each of these is briefly stated with a short accompanying outline of what it entails. Chapter seven is devoted to the COMPRAM methodology which is a framework giving guidelines, suggestions and heuristics on how to approach complex societal problems. It consists of six steps, namely, the analysis, problem description by neutral content experts, analysis and description of the problem by different teams of actors, identification of interventions by experts and actors, anticipation of the societal reactions, implementation of the interventions and evaluation of the changes.

In most of the second part of the book, chapters eight to fourteen, examples of the use of the COMPRAM methodology are discussed in what the author calls "the domain of global safety." Examples include healthcare, economics, climate change, terrorism, large city problems, large technological projects and floods. The credit crisis is likewise presented as a hypothetical application. The time period and geographical area is chosen, demarcating boundaries of the problem. Issues considered include: the unequal distribution of wealth and power, capitalism and democracy, corruption, illegal activities and tolerance, the credit crisis of 2008 and the various actors involved, the role of business banks, the role of private equity funds and hedge funds, and worldwide financial systems. Given all the background, information and aspects of the problem, the methodology is illustrated.

The issue of implementation of interventions as well as ethical aspects, validation and testing of the methodology are addressed and illustrated in chapter thirteen. The final chapter shows how the outcomes could be used for policy formulation and policy making. A final summary is also provided.

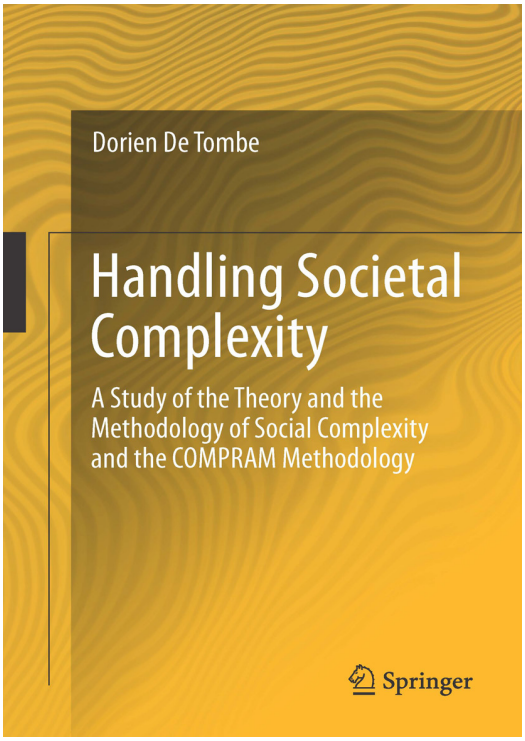

This book is thoroughly researched and provides an in-depth understanding of the multi-disciplinary, multi-actor, multi-level and often also multi-continental approach presented and how it takes into account emotional aspects of a complex societal problem. The author makes it clear that these problems require different approaches and that, in the majority of cases, these are impossible to solve. Nevertheless OR problem solvers, including researchers, lecturers and students, can benefit immensely from an exposure to a problem handling methodology that, if applied properly, has the potential to reduce conflicts, save money and ultimately, even save lives.

Handling Societal Complexity – A Study of the Theory and the Methodology of Societal Complexity and the COMPRAM Methodology by Dorien DeTombe, 2015. Springer-Verlag, Berlin. pp. 551, ISBN: 978- 3-662-43917-6, EURO 129.99 (Hardcover) and ISBN: 978-3-662-43917-3, EURO 107.09 (e-book).

QUERIES & CONTRIBUTIONS

Any queries or contributions to the Newsletter are most welcome, especially article submissions. For any queries or contributions, please contact the Newsletter editor at *berndtlindner@gmail.com.*

anylogic® 7

THE ONLY MULTIMETHOD SIMULATION SOFTWARE

Distributed by

BLUESTALLION technologies

BLUE STALLION TECHNOLOGIES cc

Technical, Scientific and Business Software Solutions
Visit us at www.bluestallion.co.za

Tel.: +27 (0) 11 447 9916
E-mail: info@bluestallion.co.za

Dowload the free trial at **www.anylogic.com**

AnyLogic simulation modeling software is an essential productivity and decision-making tool, which supports all existing modeling methods and therefore covers all your company's possible needs in simulation. Industrial engineers, supply chain analysts, HR managers, market research analysts and strategic planners are using AnyLogic to leverage the power of dynamic simulation throughout the entire business lifecycle.

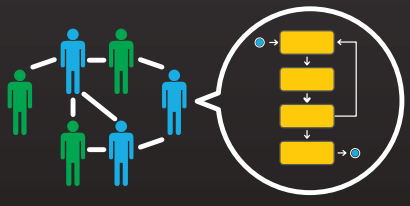

AGENT BASED MODELING

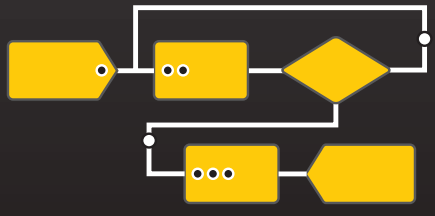

PROCESS-CENTRIC (DISCRETE EVENT)

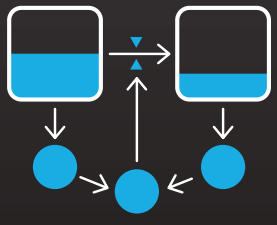

SYSTEM DYNAMICS

- Health Care/Pharma
- Military/Defense
- Airports/Stations/Malls
- Transportation/Rail
- Manufacturing
- Supply Chains
- Logistics/3PL
- Insurance/Finance
- Marketing/Social Media
- Business Processes

- DECISION SUPPORT
- RISK REDUCTION
- CHANGE SIMULATION
- OPTIMIZATION
- VISUALIZATION
- "WHAT-IF" ANALISIS

Optimise your career through research collaboration or by obtaining a degree through the Industrial Engineering Department at Stellenbosch University.

BEng (Industrial), MEng (Research) or MEng (Structured) (Engineering Management), MEng (Research) or MEng (Structured) (Industrial Engineering), PhD, Postgraduate Diploma in Engineering Management or Industrial Engineering

Industrial Engineering @ Stellenbosch University

www.ie.sun.ac.za