SA FORUM FOR ARTIFICIAL INTELLIGENCE RESEARCH 2019 | PROGRAMME

WEDNESDAY, 4 DECEMBER 2019

08:00	Registration & Coffee
09:15	Welcoming
09:30	Adnan Darwiche | **What logic can do for AI today**
10:30	Refreshment Break
Theme: Machine learning. Session Chair: Alta De Waal	
11:00	Chris Swanepoel and Katherine Malan | **Towards a visual framework for the incorporation of knowledge in the phases of machine learning**
11:20	Asad Jeewa, Anban Pillay and Edgar Jembere | **Directed curiosity-driven exploration in hard exploration, sparse rewards environments**
11:40	Zach Wolpe and Alta De Waal | **Autoencoding variational Bayes for latent Dirichlet allocation**
12:00	Arnold Pretorius, Marelie Davel and Etienne Barnard | **ReLU and sigmoidal activation functions**
12:20	Tian Theunissen, Marelie Davel and Etienne Barnard | **Insights regarding overfitting on noise in deep learning**
12:40	Marelie Davel | **Activation gap generators in deep neural networks**
13:00	Lunch
Theme: Applications of AI: Speech and Language. Session Chair: Aurona Gerber	
14:00	Lulamile Mzamo, Albert Helberg and Sonja Bosch | **Evaluation of combined bi-directional branching entropy language models for morphological segmentation of isiXhosa**
14:20	Bernardt Duvenhage | **Short text language identification for under resourced language**
14:40	Oluwafemi Oriola and Eduan Kotzé | **Automatic detection of abusive South African tweets using a semi-supervised learning approach**
15:00	Short Break

Mixed Work in Progress oral presentations		
	Session Chair: Anban Pillay	Session Chair: Edgar Jembere
15:10	Dewald Krynauw, Marelie Davel and Stefan Lotz | **Solar flare prediction with temporal convolutional networks**	Julius Stopforth and Deshendran Moodley | **Continuous versus discrete action spaces for deep reinforcement learning**
15:30	Bradley Pillay, Anban Pillay and Edgar Jembere | **Human activity recognition using deep learning approaches**	Cilliers Pretorius and Thomas Meyer | **Improved explanations in the Protégé OWL Ontology Editor**
15:50	Kouame Kouassi and Deshendran Moodley | **Automatic construction of deep neural networks for stock market trend prediction**	Yahlieel Jafta and Louise Leenen | **A semantic tool for the protection of personal information**
16:10	Rest of Poster WIP 1 minute presentations	
16:40	Refreshment Break & Start of Poster Session	
18:00	End of Day 1	

THURSDAY, 5 DECEMBER 2019

08:15	Registration & Coffee
09:00	Andries Engelbrecht | **A hyper-heuristic framework for dynamic optimisation problems**
10:00	Refreshment Break
Theme: Applications of AI: Speech and Language. Session Chair: Marelie Davel	
10:30	Aby Louw | **Neural speech synthesis For resource-scarce languages**
10:50	Ari Holtman, Jan Buys, Leo Du, Maxwell Forbes and Yejin Choi | **The curious case of neural text degeneration**
Theme: Applications of AI. Session Chair: Louise Leenen	
11:10	Wiebke Toussaint and Deshen Moodley | **Comparison of clustering techniques for residential load profiles in South Africa**
11:30	Stefan Lotz, Jacques Beukes and Marelie Davel | **Input parameter ranking for neural networks in a space weather regression problem**

11:50	Stefanie Strachan and Aurona Gerber | **Graph analysis of Twitter feed network maps the detection of network patterns within the South African Twitter community**
12:10	Pablo Fillottrani and C. Maria Keet | **An architecture for efficient knowledge-driven Information and data access**
12:30	Lunch
Theme: Applications of AI. Session Chair: Deshen Moodley	
13:30	Catherine S. Price, Deshendran Moodley and Anban Pillay | **Modelling uncertain adaptive decisions: application to KwaZulu-Natal sugarcane growers**
13:50	Michael White and Patrick Marais | **Supervised learning and image processing for efficient malaria detection**
14:10	Muhammad Shoaib Ikram, Anban Pillay and Edgar Jembere | **Using social networks to enhance a deep learning approach to solve the cold-start problem in recommender systems**
14:30	Mark de Lancey and Inger Fabris-Rotelli | **Effective graph sampling of a nonlinear image transform**
14:50	Alexander Powell, Chad van Wyk, Darren Bates and Darren de Abreu | **A cross-comparison of feature selection algorithms on multiple cyber security data-sets**
15:10	Emmanuel Kabuga, Ian Durbach, Bubacarr Bah and Allan Clark | **Using neural networks to identify individual animals from photographs**
15:30	Refreshment Break
Theme: Knowledge representation. Session Chair: Arina Britz	
16:00	Matthew Morris, Tala Ross and Thomas Meyer | **Defeasible disjunctive datalog**
16:20	Julian Chingoma and Thomas Meyer | **Forrester's paradox using typicality**
16:40	Michael Harrison and Thomas Meyer | **Preferential reasoning for datalog**
18:30	FAIR Dinner **Venue: Life Grand Café** V&A Waterfront 19 Dock Road, Cape Town

FRIDAY, 6 DECEMBER 2019

08:15	Registration & Coffee
09:00	Cobus Bernard | **How machine learning is changing the game (digital transformation in sports)**
10:00	Refreshment Break
Theme: Ethics of AI. Session Chair: Emma Ruttkamp-Bloem	
10:30	Thompson Chengeta | **Is existing law adequate to govern autonomous weapon systems?**
10:50	Wolfhart Totschnig | **Fully autonomous AI**
11:10	Emma Ruttkamp-Bloem | **The ethics of belief in the context of data-driven knowledge**
11:30	Fabio Tollon | **Towards a coherent account of moral agency**
11:50	Karabo Maiyane | **Ethics of Artificial Intelligence: virtue ethics as a solution to artificial moral reasoning in the context of lethal autonomous weapon systems**
12:10	Cindy Friedman | **The application of ethical boundaries for humanoid companion robots: a human perspective**
12:30	Lunch
13:30	Quentin Williams | **AI in industry: has the practitioner surpassed the scientist?**
14:10	Short Break
14:20	FAIR General Meeting
15:20	Closing of FAIR

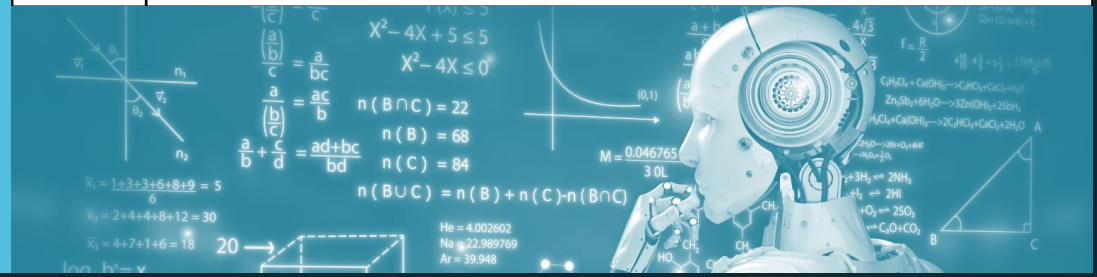

SA FORUM FOR ARTIFICIAL INTELLIGENCE RESEARCH 2019 | KEYNOTE PRESENTATIONS

ADNAN DARWICHE

Adnan Darwiche is a professor and former chairman of the Computer Science department at the University of California. He directs the Automated Reasoning Group, which focuses on symbolic and probabilistic reasoning and their applications, including machine learning. Professor Darwiche is a AAAI and ACM Fellow. He is a former editor-in-chief of the Journal of Artificial Intelligence Research (JAIR) and author of "Modeling and Reasoning with Bayesian Networks," by Cambridge University Press.

What logic can do for AI today

I will discuss a number of roles for logic in AI today, which include probabilistic reasoning, machine learning and explaining AI systems. For probabilistic reasoning, I will show how probabilistic graphical models can be compiled into tractable Boolean circuits, allowing probabilistic reasoning to be conducted efficiently using weighted model counting. For machine learning, I will show how one can learn from a combination of data and knowledge expressed in logical form, where symbolic manipulations end up playing the key role. Finally, I will show how some common machine learning classifiers over discrete features can be compiled intro tractable Boolean circuits that have the same input-output behavior, allowing one to symbolically explain the decisions made by these numeric classifiers.

ANDRIES ENGELBRECHT

Prof Engelbrecht received the Masters and PhD degrees in Computer Science from Stellenbosch University, South Africa, in 1994 and 1999 respectively. He is currently appointed as the Voigt Chair in Data Science in the Department of Industrial Engineering, with a joint appointment as Professor in the Computer Science Division, Stellenbosch University. Prior to his appointment at Stellenbosch University, he has been at the University of Pretoria, Department of Computer Science (1998-2018), where he was appointed as South Africa Research Chair in Artificial Intelligence (2007-2018), the head of the Department of Computer Science (2008-2017), and Director of the Institute for Big Data and Data Science (2017-2018). His research interests include swarm intelligence, evolutionary computation, artificial neural networks, machine learning, data analytics, and optimization. He is the author of two books, "Computational Intelligence: An Introduction" and "Fundamentals of Computational Swarm Intelligence", in addition to a number of articles in journals and the proceedings of international conferences. He is actively involved as a member of the IEEE Computational Intelligence Society, and serves as an associate-editor for a number of journals. He holds an NRF A rating.

A hyper-heuristic framework for dynamic optimisation problems

Dynamic optimisation problems provide a challenge to optimization algorithms in that optima have to be tracked as the search landscape changes. The complexity of a dynamic optimisation problem is determined by the severity and frequency of changes, as well as the behavior of the values and trajectory of optima. While many efficient meta-heuristic algorithms have been developed to solve these types of problems, the choice of the best algorithm is highly dependent on the type of change present in the environment. This presents a significant challenge to practitioners in that it takes some effort and time to understand the nature and characteristics of a dynamic optimisation problem, in order to identify a suitable algorithm for that problem. A "black-box" approach to solving dynamic optimisation problems is needed. Hyper-heuristics offer such a black-box approach. The optimisation process is adapted by continually choosing the best meta-heuristic(s) to apply to track optima over time. This talk presents a hyper-heuristic framework for population-based meta-heuristics for continuous-valued dynamic optimisation problems. Different meta-heuristic selection operators are presented and the performance of these selection operators are analysed and compared to single-instance meta-heuristics and a number of control groups.

SA FORUM FOR ARTIFICIAL INTELLIGENCE RESEARCH 2019 | INDUSTRY PRESENTATIONS

COBUS BERNARD

Cobus is a Senior Technical Evangelist at AWS empowering developers to know and understand how best to use AWS. His primary interests are in security, containers and devops. Prior to joining AWS, he was a customer for 8 years building in the FinTech, HealthCare and Online Gaming space. During his 14 years of development, he has worked on C#/Java backends, done a bit of Android / iOS app development and some game development, one which is available on Steam. In the last few years he has focussed on spreading DevOps best practices and helping companies adopt them. Founder of the Cape Town DevOps Meetup group and co-organiser of DevOpsDays Cape Town.

How machine learning is changing the game (digital transformation in sports)

Machine learning is being built into every aspect of modern life, from recommendations for which products are purchased together to identifying lost dogs after natural disasters to reunite them with their families. In this talk, we will take a look at how machine learning is applied in live sporting events. Using real-time video processing combined with historic data, you can predict the probability of specific set-play pieces of succeeding in American Football (NFL), or help shape your racing strategy for F1 by predicting the probability of overtaking vs making an optimised pitstop. It is even used to predict the probability a player trying to execute a base steal in Major League Baseball (MLB). All of this is done in real-time by processing large datasets as they are received.

QUENTIN WILLIAMS

Dr Williams completed his PhD in 2005 from the University of Oxford in the UK focussing on pattern recognition and classification in the medical imaging domain. He worked at the CSIR researching technologies for persons with disabilities using Artificial Intelligence (AI). He soon became the Manager for Emerging Research and Human Capital Development, a post he held for 5 years, before being appointed as the Strategic Research Manager for the CSIR's Meraka Institute focusing on digitalisation and analytics strategies. During this time he established and led the W3C Southern African Office; the Intelligent Environments for Independent Living research group; and the Data Science for Impact and Decision Enablement programme and was the technical lead for the Cabinet approved ICT RDI Roadmap. He joined Deloitte in February 2019 looking at analytics and AI-driven business.

AI in industry: has the practitioner surpassed the scientist?

Platforms such as DataRobot and H2O.ai have made it all the more easier to develop and deploy analytics and AI models. This has led to the term citizen data scientists and a plethora of data "insights" published from industry. Industry has seen this as the saviour to their data scientist resource problem with extreme valuations for AutoML companies currently in vogue all expressing the promise of opening up the data science "black box". The question therefore remains: what is the true value of the academic data scientist, and how can we translate that value more easily to industry? Should we continue to develop and explore the intricacies of deep learning networks or has AutoML replaced the art and science of the academic?