Benelux AI Newsletter

Would you like to announce your own events on the BNVKI website and newsletter? Send an email to events@bnvki.org! Also reports of past events are welcome. If you have other news or contributions for the BNVKI newsletter, these can be sent to editor@bnvki.org.

BNAIC 2017 Postproceedings Published

The BNAIC 2017 postproceedings were published in the Springer CCIS series. It contains all 11 Type A regular papers accepted for oral presentation (37% of the 30 Type A submissions).
Verheij, B., & Wiering, M. (eds.) (2018). *Artificial Intelligence. 29th Benelux Conference, BNAIC 2017. Groningen, The Netherlands, November 8-9, 2017. Revised Selected Papers (Communications in Computer and Information Science Volume 823).* Cham: Springer. https://doi.org/10.1007/978-3-319-76892-2

EU Communication on AI

On April 25th, 2018, the European Commission has published a Communication outlining a European approach to AI, based on three pillars:

- **Boosting the EU's technological and industrial capacity and AI uptake across the economy**, both by the private and public sectors. This includes investments in research and innovation and better access to data.
- **Preparing for socio-economic changes** brought about by AI by encouraging the modernisation of education and training systems, nurturing talent, anticipating changes in the labour market, supporting labour market transitions and adaptation of social protection systems.

- **Ensuring an appropriate ethical and legal framework**, based on the Union's values and in line with the Charter of Fundamental Rights of the EU. This includes forthcoming guidance on existing product liability rules, a detailed analysis of emerging challenges, and cooperation with stakeholders, through a European AI Alliance, for the development of AI ethics guidelines.

The Communication is available via: https://ec.europa.eu/digital-single-market/en/news/communication-artificial-intelligence-europe.

BNAIC / BeneLearn 2018

BNAIC 2018 will be co-located with BeneLearn 2018! The joint event will take at the Jheronimus Academy of Data Science (JADS) in ‘s-Hertogenbosch on November 8-9, 2018. For more information, see https://bnaic2018.nl.

Call for Nominations: EurAI Artificial Intelligence Dissertation Award 2017

The European Association for Artificial Intelligence has launched a call for nominations for its Dissertation Award (1.500€). BNVKI members are also members of EurAI and may therefore be eligible for this award!

BNAIC 2017: Session report "Machine Learning 2"

The first presentation in this session was given by David Roschewitz from Maastricht University, who presented a joint work with Kurt Driessens and Pieter Collins titled "Simultaneous Ensemble Generation and Hyper parameter Optimization for Regression". The authors presented a method to simultaneously generate ensembles and tune the hyper parameters of the models for regression problems. In addition they also investigated the use of robust loss functions as well as different methodologies for determining the size of the ensemble. For the tested problems, they have observed that the MSE loss function outperforms the robust loss function. Many experimental results

on several challenging problems showed that for models with tunable hyper parameter spaces, the proposed techniques significantly outperform single regressors.

The second presentation in this session was given by Dr. Marieke van Vugt from the University of Groningen. The title of the paper is: "Distracted in a Demanding Task: A Classification Study with Artificial Neural Networks" which is a joint work with Stefan Huijser and Niels Taatgen. The authors analyzed spatial complex working memory tasks by means of Artificial Neural Network (ANN) based classifiers. More precisely, they aim at predicting whether distracted subjects' thoughts were focused on the task based on recorded eye-tracking features and task performance. They also found that Eye-tracking features (e.g., pupil size, blink duration, fixation duration) are much less predictive and that trial-to-trial performance is the strongest predictor of distracted thought.

The third presentation in this session was given by Dimitrios Bountouridis from Utrecht University. The title of the paper is: "Melody Retrieval and Classification Using Biologically-Inspired Techniques" which is a joint work with Dan Brown, Hendrik Vincent Koops, Frans Wiering and Remco Veltkamp. The paper aims at enhancing melody retrieval and classification using bioinformatic based techniques. With regard to efficient classification, the authors employed BLAST (Basic Local Alignment Search Tool) and observed its limitations for complex retrieval tasks. Therefore they also examined the profile hidden Markov models (profile HMMs) which are able to capture salient and robust properties of musical content, called profiles or prototypes. Their experimental results show that BLAST and profile HMMs can be reliable and efficient solutions for large-scale melody classification and retrieval respectively, without the incorporation of musical heuristics.

The last presentation in this session was given by Jaap Kamps from the University of Amsterdam who presented joint work with Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn and W. Bruce Croft titled "Neural Ranking Models with Weak Supervision". They proposed to leverage large amounts of unlabeled data to infer "noisy" or "weak" labels and use that signal for learning supervised models. In particular, the authors used classic unsupervised IR models such as BM25 as a weak supervision signal for training deep neural ranking models. They further studied the impact of weak supervision on various neural ranking models with different ranking architectures and objectives. Other studies that have been conducted in this work are analyzing the behavior of models to understand what they learn, what is the relationship among different models, and how much training data is needed to go beyond the weak supervision signal. They demonstrated that, in the ranking problem, the performance of deep neural networks trained on a limited amount of supervised data significantly

improves when they are initialized from a model pre-trained on weakly labeled data.

BNAIC 2017: Session report "Knowledge Representation & Reasoning"

This session included four fascinating presentations, focussing on different aspects of Knowledge Representation and Reasoning. The first paper, titled 'Constructing Knowledge Graphs of Depression', was presented by Frank van Harmelen, and was co-authored by Zhisheng Huang, Jie Yang, and Qing Hu. Frank presented a collaborative project between the VU and Beijing Anding hospital that focused on constructing a large knowledge graph of depression by integrating various knowledge resources using semantic web technology. Due to its high degree of inter-operability, this approach enables psychiatric doctors to efficiently find answers to queries without having to explore multiple databases.

The second paper, titled 'Using Values and Norms to Model Realistic Social Agents', was presented by Rijk Mercuur, and described joint work with his supervisors at Delft University of Technology, Virginia Dignum and Catholijn Jonker. The presentation addressed an interesting study that aimed to evaluate the benefit of values and norms for developing more 'socially realistic' agents. By simulating a psychological experiment known as the Ultimatum Game with three different types of social agents, the authors found that agents using a theory that combines both values and norms produce behaviour that is most similar to the behaviour of humans.

The third paper, titled 'Evaluating Intelligent Knowledge Systems', was presented by Neil Yorke-Smith of TU Delft. In his engaging talk, Neil shared his experiences from a large seminal project on the development of a user-adaptive personal assistant agent for time management assistance. The presentation focused on the (often underestimated) role of evaluation in such a project, and concluded with several useful 'lessons learned', of which the main message was that researchers and project managers benefit from adoption of best practice in evaluation methodologies from the start of a technology project.

The fourth paper was titled 'On the Problem of Making Autonomous Vehicles Conform to Traffic Law' and was presented by Henry Prakken (Utrecht University & University of Groningen). As the title suggests, the presentation addressed the problem of how we can make the behaviour of self-driving cars conform to traffic laws, although it could also be seen as a special case of the general problem of making intelligent

autonomous systems conform to the relevant laws. Henry discussed various features that make traffic law challenging for AI & Law models (such as exceptions, rule conflicts, and the need for common-sense knowledge) and evaluated three approaches to the design of law-conforming autonomous vehicles in light of these challenges. All in all, this was a very lively session with a lot of interesting discussions.

BNAIC 2017: Session report "Games"

The session on Games consisted of four presentations discussing applications of mathematical game theory in multi agent systems. The first two presentations presented the work of Gleb Polevoy and Mathijs de Weerdt on agents that need to divide their available time over multiple projects, respectively multiple group interactions. An example of the former is the division of time over different research project, and an example of the latter is the active participation in different social media groups. When agents have to divide time over multiple projects, they need to consider the benefits of participating in each of the projects. The benefit of a project depends on whether enough agents investigate time in the project to be successful. When agents participate in multiple group interactions, the benefit of each group interaction depends on the participation of the other agents in the group. The conditions under which (social optimal) Nash equilibria exist, are investigated for agents dividing their available time over multiple projects, respectively multiple group interactions.

The third presentation presented a new gift-giving game developed by Elias Fernández Domingos, Juan Carlos Burguillo and Tom Lenaerts. Gift-giving games can be used to study the emergence of trust, fairness and generosity. Two types of agents were compared, namely reactive and anticipative agents. The anticipative agents showed some of the characteristics of human decision making in the experiments.

The last presentation of the session addressed repeated task allocation problems. This work by Qing Chuan Ye and Yingqian Zhang investigated the influence of past task allocation outcomes on future task allocation outcomes. Two types of agents were investigated, agents that only consider optimality in terms of costs, and agent that consider optimality in terms of primarily fairness and secondarily costs. The experiments demonstrated that the latter agents have an incentive to continue participating, which result in a higher social welfare of the agents.

BNAIC 2017: Session report "Uncertainty in AI"

In the Uncertainty in AI session three papers were presented. Two papers were wholly or partly motivated by legal applications. Remi Wieten from Utrecht university addressed the problem that Bayesian networks, while increasingly popular in forensic science, are poorly understood by legal experts. This problem makes it hard to build reliable Bayesian networks for crime investigation or for evidential reasoning in criminal cases. Wieten further developed the idea that legal experts can express their evidential knowledge in terms of arguments (a style of reasoning with which they are familiar), which are then automatically converted into constraints on the design of a Bayesian network.

Joost Vennekens of the KU Leuven presented a joint paper with Sander Beckers on formalizing the concept of actual causation, where the problem is to define when event X is deemed to have caused event Y in the context of a particular story. This is a problem that often arises in legal cases. After criticizing David Lewis' definition of actual causation as the transitivity of counterfactual dependency, Vennekens presented an alternative approach based on the ideas that counterfactual dependency is a sufficient but not necessary condition for actual causation, and that actual causation is transitive only insofar as it does not violate asymmetry (X cannot have caused Y if we would also have considered not-X as a cause of Y).

Johan Quisthout from the Donders Institute in Nijmegen addressed complexity issues of approximate inference in Bayesian networks. One motivation of this work is the Bayesian Brain hypothesis from cognitive science, according to which the human brain carries out or at least approximates Bayesian updating. This raises the issue of how humans can do this efficiently given that Bayesian inference is intractable in general. Using the formal framework of so-called fixed-error randomized tractability, Kwisthout presented a number of positive and negative tractability results.

BNAIC 2017: Session report "Agent Systems 2"

The Agent Systems session was opened by Bram Wiggers from the University of Groningen. In his paper "The origin of mimicry: Deception or merely coincidence?", he presents an agent model for the evolution of mimicry. Mimicry refers to a situation in which two species of prey animals share a phenotype, while only one of them is dangerous to eat for predators. In this situation, predators avoid eating both species, which allows the harmless prey to benefit from the presence of the dangerous prey. Wiggers showed how mimicry could evolve through deception, where harmless prey

evolve to have the same phenotype as the dangerous prey, but also through coincidence.

Next, Daniel Formolo presented his joint work with Natalie van der Wal on "Simulating Collective Evacuations with Social Elements". He argued that while current simulations for evacuations are helpful, but they do not model human behaviour that corresponds to reality. In particular, current models lack socio-cultural modeling such as trips and falls, fear, social contagion, helping behaviour, or familiarity with the environment. He showed that especially travelling in groups reduces evacuation time. Formolo concluded by saying that the model with socio-cultural modeling has the potential to predict human behaviour in emergency situations more accurately than the current standard.

In the final presentation of the session, Jens Nevens presented his work with Katrien Beuls, titled "The Effect of Tutor Feedback in Language Acquisition Models". Nevens argued that an important aspect of teaching and learning is referent learning, which is learning what physical object is referred to by a given word. In particular, he investigated the role of social feedback, such as pointing or gazing at the object. In his agent model for a color guessing game, Nevens compared cross-situational learning, in which no feedback is given, with interactive learning, in which social feedback is always given. The results show that social feedback influences learning a language in two ways: it allows the learner to understand the referent of a given word more quickly, and also better. In addition, the model showed that the tutor's strategy for selecting topics greatly influences the learner's learning.

BNAIC 2017: Session report "Reinforcement Learning"

In the first presentation, Nico Roos proposed a model for describing a network of distributed services for task executions, and proved that for this model, even in a self-interested setting reinforcement learning can be applied to learn how to cooperate well. Specifically, when agents can decide which agent should perform the next subtask and can assess the quality of service of other agents, and share in rewards (and responsibility) for the completion of the full task, agents will learn to hand off their tasks to the best possible next agents. In this, agents do not have to be fully cooperative to learn how to cooperate - which is important, as they learn to assign blame when the quality of service is low - but have a natural incentive to be truthful, as well as to

cooperate.

Then, Jelle Visser took the floor to introduce us to their approach to learning-from-demonstration as a prequel to independent reinforcement learning in Donkey Kong. This approach mitigates the problem of having to randomly explore - leading to a very long initial period in which the agent does not seem to learn well - when learning from scratch. By playing a hundred games (50 each by two people), they managed to initialise a policy and a critic as input critic for an actor-critic RL approach - ACLA - that resumes learning independently from human demonstration. When the initial policy was already reasonably good, the authors found that ACLA did not improve much anymore from the initial policy, but when the initial policy was not yet good, ACLA did help to improve it. The authors conclude that these preliminary results are inconclusive, and that more research is required to ensure that policies are improved further after initialisation by learning-from-demonstration.

Last but not least, Thomas Moerland introduced us to his approach to variational inference for learning multi-modal transition models in model-based RL. This is highly important, in order to capture the complex stochastic environment dynamics that underlie many real-world problems. They show that using variational autoencoders, they can learn a generative model of a multi-modal distribution for the transition dynamics by introducing unobserved latent variables. This result can possibly lead to many exciting new model-based RL methods based on (deep) neural networks in the future.

BNAIC 2017: Session report "Natural Language Processing"

The first presentation of the Natural Language Processing session was given by Mathijs Pieters from the University of Groningen titled "Comparison of Machine Learning Techniques for Multi-label Genre Classification". This joint work with Marco Wiering won the SNN best paper award and studies classic and novel text classification techniques with the aim to recognize the genre of a movie based on its subtitles. A new dataset was constructed and a novel technique combining a histogram and the Word2vec model was introduced. Experiments were performed with six different machine learning techniques that use the word information as input such as a combination of long short term memory networks with convolutional neural networks. The results showed that the more complex methods did not outperform the simpler bag-of-word technique with a multi-layer perceptron, and that the introduced histogram

performed the best from all methods using Word2vec embeddings.

Verna Dankers from the University of Amsterdam then presented two papers: "Modelling the Generation and Retrieval of Word Associations with Word Embeddings" and "Modelling Word Associations and Interactiveness for Describer Agents in Word-Guessing Games" which are co-authored with Aysenur Bilgin and Raquel Fernández. The first presentation described the use of semantic vector spaces for the Location Taboo Game. In this game, an agent has to guess the name of a city from hints provided by a describer agent. The agent therefore has to learn to associate hints such as "Great pizza" with possibly first a country and then a specific city based on more hints. The results showed that the agent could correctly guess the name of the intended city in around 27% of the games. In the second presentation, Verna Dankers discussed the role of the describer agent for the same game, with the aim to optimize a describer agent in such a way that it gives informative hints so that humans can guess the name of the city. To make this game more complex, a list of Taboo words is given, which are not allowed to be used in the hints. Different approaches were tested in simulation and in an empirical study and the results showed that humans who knew the name of a city could guess it in half of the cases after a limited amount of hints given by the best system.

The fourth talk was presented by Zoltán Szlávik from IBM in Amsterdam and discussed his joint work with Nikita Galinkin, Lora Aroyo and Benjamin Timmermans titled "Catch Them If You Can: Malicious Behavior Simulation in Deep Question Answering". The presentation was about the problem when question-answering systems are developed using human input, when users are malicious and enter incorrect information. Although the talk was focused on the cultural heritage domain, the presenter strongly argued that for building intelligent systems in this way, malicious user actions should be detected as soon as possible.

BNAIC 2017: Session report "Agent Systems"

The first talk, "Get Your Virtual Hands Off Me! - Developing Threatening Agents Using Haptic Feedback" was presented by Tibor Bosse, and described how virtual agents can be used to train people in social situations. For example, virtual agents in a VR setting are used to train tram conductors in Amsterdam on how to best respond to aggressive passengers. A continuous challenge in this research is making the simulation believable. The authors presented evidence that a virtual agent who could raise his voice, and whose pushes could be felt by haptic inputs on a vest, helped to improve

this believability.

In the second talk, Vincent Koopman talked in his presentation "Omniscient Debugging for Cognitive Agent Programs" about how debugging could be optimized in agent-based simulations. A good technique for debugging is to record the execution of a program, but this slows it down dramatically (by about 300x). However, a solution is to capture a higher level of abstraction of the program: not every little detail but only the agent's mental states. This allowed for capture of the program at only 1.1x the runtime speed.

The third talk by Lenin Medeiros, "An Empathic Agent that Alleviates Stress by Providing Support via Social Media" described how virtual agents could potentially function as friends to provide relief in times of stress. The advantage of such a system would be that those agents would always be available, and not suffer from empathy burnout. The virtual agents that were programmed first detected the type of stress of the user, then selected a support strategy and used that to send messages. Users preferred this virtual agent over a random model.

In the final talk "Expectation Management in Child-Robot Interaction" Mike Ligthart talked about how expectation management is crucial for successful child-robot interactions. For example, children may believe that robots actually understand them or that they can help them with delicate tasks such as drawing blood. If the robot is found to be unable to do that, the child then becomes disappointed. They suggested that to prevent this from happening, it is important to identify children's expectations from robots, and maybe eventually the robots themselves could help to manage these expectations by pointing out what they can and cannot do.

BNAIC 2017: Session report "Machine Learning"

After the opening remarks by the conference organizers and an introduction to the scientific context in Groningen, this session featured the first four paper presentations of the conference in the main lecture room. It had an unfortunate start due to a mix of cable and hardware issues, but speakers and chair together solved these cooperatively and very quickly and then the session started with a presentation by Claudio Reggiani titled "Feature selection in high-dimensional dataset using MapReduce", which is joint work with Yann-Aël Le Borgne and Gianluca Bontempi. Claudio presented results of a MapReduce implementation for feature selection (in this case the minimal Redundancy

Maximal Relevance filtering technique). The (open source code) results showed good scalability performance.

The second presentation was given by Dirk van der Hoeven and was titled: "Is Mirror Descent a Special Case of Exponential Weights?" which is joint work with Tim van Erven. The presentation managed to deliver a clear summary of fairly technical work in online convex optimization. One of the goals of the work is to relate online gradient descent, mirror descent, and exponential weights. Possible applications of the main results and future work include efficient sampling in linear bandits, priors for learning rates, and scale free algorithms.

The third presentation was delivered by Siamak Mehrkanoon (joint work with Johan Suykens) and titled: "Regularized Semi-Paired Kernel CCA for Domain Adaptation". The method presented can be placed in the context of transfer learning algorithms and is aimed at generalizing a model trained on a source task to a target domain by taking into account how many labeled instance in both tasks are available. Experimental results showed that a joint representation of the data set across different domains can be learned and utilized.

The final presentation in this session was given by Lynn Houthuys who presented joint work with Zahra Karevan and Johan Suykens on "Multi-View LS-SVM for Temperature Prediction". In this work, applied to a temperature prediction task for cities in Belgium, multiple views of the data (representing each city) are combined in a support vector machine setting. The proposed method enforces alignment of error variables over multiple views, distributed over multiple LS-SVMs, such that when training one SVM on one view, the other views are taken into account. The experiments show that the new approach outperforms both existing weather predictions as well as predictions based on naively concatenating all available features. This session had four speakers who all managed very well to deliver their presentation in time, which resulted in a lot of time for discussion after each talk, which was filled with many questions from the audience and the chair, even on this early hour of the conference.

BNAIC 2017: General report

BNAIC is the annual Benelux Conference on Artificial Intelligence. In 2017, the 29th edition of BNAIC was organized by the Institute of Artificial Intelligence and Cognitive Engineering (ALICE), University of Groningen, under the auspices of the Benelux Association for Artificial Intelligence (BNVKI) and the Dutch Research School for Information and Knowledge Systems (SIKS). BNAIC 2017 took place in Het Kasteel,

Melkweg 1, Groningen, The Netherlands, on Wednesday November 8 and Thursday November 9, 2017. BNAIC 2017 included invited speakers, research presentations, posters, demonstrations, a deep learning workshop (organized by our sponsor NVIDIA) and a research and business session.

The four BNAIC 2017 keynote speakers were:

- Marco Dorigo, Universite Libre de Bruxelles
  Swarm Robotics: Current Research Directions at IRIDIA
- Laurens van der Maaten, Facebook AI Research
  From Visual Recognition to Visual Understanding
- Luc Steels, Institute for Advanced Studies (ICREA), Barcelona
  Digital Replicants and Mind-Uploading
- Rineke Verbrugge, University of Groningen
  Recursive Theory of Mind: Between Logic and Cognition

Furthermore, three FACt talks (FACulty focusing on the FACts of Artificial Intelligence) were scheduled:

- Bert Bredeweg, Universiteit van Amsterdam
  Humanly AI: Creating smart people with AI
- Eric Postma, Tilburg University
  Towards Artificial Human-like Intelligence
- Geraint Wiggins, Queen Mary University of London/Vrije Universiteit Brussel
  Introducing Computational Creativity

Authors were invited to submit papers on all aspects of Artificial Intelligence. We received 68 submissions in total. Of the 30 submitted Type A regular papers, 11 (37%) were accepted for oral presentation, and 14 (47%) for poster presentation. 5 (17%) were rejected. Of the 19 submitted Type B compressed contributions, 17 (89%) were accepted for oral presentation, and 2 (11%) for poster presentation. None were rejected. All 6 submitted Type C demonstration abstracts were accepted. Of the submitted 13 Type D thesis abstracts, 5 (38%) were accepted for oral presentation, and 8 (62%) for poster presentation. None were rejected. The selection was made using peer review. Each submission was assigned to three members of the program committee, and their expert reviews were the basis for our decisions. All submissions accepted for oral or poster presentations and all demonstration abstracts appear in the electronic preproceedings, made available on the conference web site during the conference (http://bnaic2017.ai.rug.nl/). All 11 Type A regular papers accepted for oral presentation will appear in the postproceedings, to be published in the Springer CCIS series.

We are grateful to our sponsors for their generous support of the conference:

- Target Holding
- NVIDIA Deep Learning Institute
- Anchormen
- Quint
- the Netherlands Research School for Information and Knowledge Systems (SIKS)
- SIM-CI
- Textkernel
- LuxAI
- IOS Press
- Stichting Knowledge-Based Systems (SKBS)
- SSN Adaptive Intelligence

Reports on the different sessions of the conference can be found here:

- Machine Learning 1
- Agent Systems 1
- Natural Language Processing
- Reinforcement Learning
- Agent Systems 2
- Uncertainty in AI
- Games
- Knowledge Representation & Reasoning
- Machine Learning 2

Changes to the BNVKI board

As per January 2018, BNVKI has a new board. Koen Hindriks and Marc van Zee have left the board, and are thanked for their excellent contribution over the past years. They have been replaced by Mike Ligthart (TU Delft) and Yingqian Zhang (TU Eindhoven). Mike will take the role of editor of the newsletter, and Yingqian the role of secretary. In addition, Tibor Bosse has been appointed as the new chair, and Kurt Driessens as the new vice-chair of BNVKI. The new composition of the board can be seen on the

website.