Benelux A.I. Newsletter

Summer 2015 Edition (No. 2)

BNVKI
BAIAI

Contents

Interview with... Catholijn Jonker

by Marieke Peeters

Profile. Prof. Dr. Catholijn Jonker is head of the Interactive Intelligence group at Delft University of Technology. Her research interests are the modelling and simulation of sociocognitive processes and concepts such as trust, negotiation, and teamwork dynamics.

What is your academic background and how did you obtain your current position?
My original background lies in theoretical computer science, in logic programming and model theoretic semantics to be more specific. I always loved solving logical puzzles. Later on, I decided to broaden my horizon, starting out with research on agent technology, and then followed by research on cognitive models for agents and the relation between humans and machines, specifically agents.

Catholijn Jonker (photo by Wieke Eefting)

Nowadays my research is focused on questions such as: What knowledge do humans and agents need to have about one another in order for them to collaborate successfully? How can humans and agents understand one another? But also: How can an agent explain its own capacities and limitations? With regard to the latter, it is important that an agent is aware of its own role in a human-agent partnership and realizes how its own performance is related to the performance of other collaborators. My main focus is the modeling of human cognition, interactions, and social processes, such as, for instance, negotiation.

My ideal is that humans and agents become increasingly better attuned to one another, in a way comparable to interpersonal relations between humans. It has been a deliberate choice of mine to do research in such an interdisciplinary area: studying the relation between humans and artificial agents. The artificial intelligence techniques required in this endeavor are still my primary research topic. Yet in order to understand the human component in these relations, and support humans in their collaborations with agents, we also reckon with psychological components in our research.

In 2008 you received a Vici grant from NWO for the project 'The Pocket Negotiator'. Can you tell us about this project?
The Pocket Negotiator (PN) offers users support in negotiating. My goal was for the PN to be the equivalent of the pocket calculator, but for negotiations. The user's first step in accomplishing a negotiation task is to prepare the negotiation with the help of the PN: what is the goal, what are the conditions, what are the upper-and lower limits, what is the ideal outcome? After that, the user and the PN go into the negotiation together and the PN advises the user in accepting or placing bids with the use of various negotiation strategies.

What knowledge was yielded from the Pocket Negotiator project?
Agents clearly outperform humans when it comes to keeping track of the room for negotiation. Moreover, the negotiation outcomes achieved by agents are significantly better than those obtained by humans. As a result of our research, we now possess advanced knowledge about optimal automated bidding strategies between agents.

As a human, you cannot compete with an agent when it comes to negotiating on your behalf. One condition, though, is that the agent is well informed about your needs and wishes. Informing the agent about these things requires thorough preparation: as a user you should take the time to explain the agent what is important to you and what is not. One of the challenges in our research was to extract a person's preference model: as a user you may state that A is more important to you than B, but the question still remains: How much more important? To obtain people's preference models, we asked them to construct pie diagrams in which they revealed their relative (and quantified) preferences. From these pie diagrams we computed a utility function. This input method enables the PN to translate a person's qualitative preference model to a quantitative utility function. The PN then uses this utility function during the actual negotiation to determine whether a specific outcome is desirable or not. The PN is currently able to deal with a person's preferences, and can recommend bids and whether to accept or not, yet without the provision of supportive arguments. Also, as of yet, the PN cannot understand a user's emotions or establish complex relations with the user.

An interesting follow-up question would be how to find out whether a user's wishes are correctly represented by the utility function derived from the pie charts. And yet another question is whether it is a bad thing when the utility function does not perfectly match the user's wishes and preferences, and how big an error margin would be acceptable during negotiation to still obtain a desirable outcome.

What are the next steps in research for the Pocket Negotiator?
I think that the usability of the Pocket Negotiator could be improved if we were to equip the agent with domain knowledge that allows it to understand the world the way humans do. We observed that a generic negotiation system is not sufficient, because it is not very usable for the arbitrary layman. To overcome this problem, the agent should possess knowledge about certain aspects of the object being negotiated. Moreover, the agent should know how those aspects are related to the preference model of the people taking part in the negotiation. Furthermore, the agent must understand the user's motives. If the agent possesses more domain knowledge, it is able to offer the user more specific guidance.

I also think that the Pocket Negotiator could be improved if it had the ability to understand natural language. Currently, users are required to manually enter all necessary information, costing the user an (unnecessarily) large amount of time. Unfortunately, I am not working in natural language processing, so this is not the kind of improvement I would carry out myself.

> “The bond between a supervisor and their PhD student is intense and extraordinary. As a supervisor, you can really place your mark on the work of your student. (...) As a PhD student, you also have a significant influence on (the work of) your supervisor.”
>
> *Catholijn Jonker*

How do you view your role as a full professor?
As a professor, there is a multitude of things you should be able to do well: doing research, teaching, inspiring, organizing, convincing. Moreover, it is important to be sensitive to the wishes and ideas of others. To guide a group, everyone in it should be allowed the space they need in order to fulfil their

potential, but at the same time you want to excite people's shared interests as much as possible in order to further stimulate mutual collaboration.

Personally I also think it is important to train new researchers in my group. I believe it important to send PhD students off with a good start. To accomplish this, we do not merely guide our students in conducting sound research and achieving high quality publications, but also aim to answer questions such as: "What can you do to establish a strong and balanced work resume?" and "What internationally well-respected research groups are worth visiting for a while?"

The bond between a supervisor and their PhD student is intense and extraordinary. As a supervisor, you can really place your mark on the work of your student. Moreover, it gives you the possibility to pass your enthusiasm for a certain topic on to someone else. I only realized years after my own promotion that this influence is completely reciprocal. As a PhD student, you also have a significant influence on (the work of) your supervisor.

The fact that the bond between a supervisor and their PhD student is so important also means that it is important for there to be a mutual click. It is difficult to be a teacher that is a good match for everyone.

Suppose you would have the time and money to do whatever you like for a considerable amount of time. What would you like to do most?
I would continue working on the Pocket Negotiator. It still is a fantastic vehicle for doing research; on the one hand it entails highly technical research on heuristics and optimal strategies, and on the other hand it involves research on extracting a user's value system, and making the application more accessible and understandable by using, for instance, *explainable A.I.* (i.e. Artificial Intelligence able to explain itself – e.g. goals, actions, internal state – to others).

What do you like most about your job?
The mutual inspiration and generation of ideas, for instance through conversations such as this one.

What's hot in...Computational Linguistics?

by Antal van den Bosch, Center for Language Studies, Radboud University

Computational linguistics is a vibrant subfield of artificial intelligence, operating relatively autonomously from other areas in AI since the 1960s. For a grand overview of the field, consult the ACL Anthology, offering all journal articles and proceedings of conferences and workshops that the Association for Computational Linguistics has overseen, or could get their hands on: https://aclweb.org/anthology/. This strong self-organization has, fortunately, not led to a 'splendid isolation' - computational linguistics has been open to and has embraced many developments in artificial intelligence. It has often been in the lead of such developments, or has offered the first applications of new AI algorithms in, for example, neural networks, symbolic machine learning, and statistical machine learning.

The advent of applicable, scalable machine learning algorithms in the early 1990s caused big excitement in the field, which at that time was experiencing its own kind of a winter. Toy problems, grammars and lexicons were replaced by real data. Working with corpora and treebanks with 'millions of words' in the late 1990s evolved in talking about hundreds of millions in the 2000s, and billions in the 2010s. Measured in bytes, language data may not be the biggest data around, but measured in information content, the digital deluge of textual data is arguably the richest treasure trove mankind has ever gathered. While computational linguistics covers mostly text, text itself remains one of the prime media of human information storage and communication. Internet and social media are newly built wells tapping into products of human communication and intelligence that outdo all the aggregated information output of all printed output and archives of the past centuries in each single year.

In short, data-driven algorithms are now the norm in computational linguistics, and their scalability a clear issue, due to the abundance of data. Thanks to some skilled programmers and PhD students, the field has been supplied richly with open-source software. This, along with the usual type of hype cycles one observes in most of the fast-paced subfields of artificial intelligence, has now led to a number of hot areas of which I highlight two: topic modeling and low-dimensional word vector representations.

Low-Dimensional Word Vector Representations

To start with the latter, the story has to begin with *word2vec*, an (as yet) open-source software package from Google's labs (https://code.google.com/p/
word2vec/), see also Mikolov *et al.* (2013). The word2vec development team, led by Tomas Mikolov, has been very successful in kickstarting a large interest in learning low-dimensional vector representations for words. Although this work was already on the agenda in the early 1990s, the relative ease with which software such as word2vec can be applied to very large corpora these days has made the topic hot again. Word2vec exploits the notion that similar words occur in similar contexts. Asked about 'France', word2vec will be able to tell you that the closest vectors are those of 'Spain', 'Belgium', and 'Italy'.

Word	Cosine distance
los_angeles	0.666175
golden_gate	0.571522
oakland	0.557521
california	0.554623
san_diego	0.534939
pasadena	0.519115
seattle	0.512098
taiko	0.507570
houston	0.499762
chicago_illinois	0.491598

Wordvec example output with the closest tokens to 'san_francisco'

What is so appealing about the vector representations learned by word2vec is that you can almost literally compute with words. Subtracting the vector for the word 'man' from the vector for the word 'king', and adding the vector for the word 'woman' to this result, yields a vector very close to that of 'queen'. Applications of these vector representations are known to be useful in any information processing tool in which it is handy to expand words into lists of similar words, and to be able to compute with semantic similarity estimates between words: in information retrieval, question answering, summarization and paraphrasing, and text categorization. The word2vec idea has spawned a great amount of related work that not only finds low-dimensional vectors for words or phrases, but also for sentences, paragraphs, and entire documents.

Topic Modeling

The other hot 'topic' is the concept of 'topic' itself, or rather the unsupervised discovery of latent topics in a collection of documents. In contrast to pre-labeled documents, which are usually limited in size due to the cost of human annotation, many vast document collections exist that are unlabeled. At the same time they obviously exhibit topical structure, across but also within documents. The pivotal algorithm that revived the interest in this topic was the formulation of the Latent Dirichlet Allocation algorithm by David Blei and colleagues (Blei et al., 2003). It is perhaps odd to call a 12-year old algorithm 'hot', but after a slow start the application of this algorithm has soared over the past five years. Most notably, it has found its way well beyond AI and computational linguistics, into diverse fields such as political and management sciences, literary studies, and historical studies. Latent Dirichlet Allocation has two popular implementations (in GenSim, and in Mallet). LDA, not to be confused with the 'other' LDA, the classic Linear Discriminant Analysis dimension reduction method, finds a mixture of N topic models, where N is a number set by the experimenter. Each topic model is a generative model that 'wants' to generate topically coherent text. During training, LDA uses sampling methods such as Gibbs sampling to turn a random initialization into a self-organized mix of topic models. It is mostly the direct readout of topic models, e.g. the top-20 words associated with each model, that induces the 'wows from the crowd', and is already of direct use for further human interpretation in e.g. literary studies. The topic models themselves can be used for further processing in text mining applications such as text categorization and information retrieval.

Topic Modeling of Dreams

As an illustration, we recently analyzed 981 dreams reported by a nearly equal amount of male and female American graduate students. These dream descriptions are part of DreamBank, a curated set of reported dream collections gathered for psychological research purposes. Running LDA on the dream descriptions, we induce topics such as the following:

- Store money buy get pay bill man grocery lot counter bank shopping machine give tickets shop put dollars change bought

- Bathroom toilet hair shower water go room bath floor clean wash see naked sink tub pee bathtub face cut towel
- book paper read books find picture write writing reading letter pictures written name looking look something library letters office computer

These three topics clearly group semantically coherent sets of words and represent typical topics in everyday life (which are the main topics of 75-80% of all dreams according to dream psychologists, see Domhoff, 2003). Since we know whether a dream is written by a man or a woman, we can apply a log-likelihood test to identify topics that are dreamt significantly more often by women or by men. Only with this scientific introduction do I dare to quote the topics that came out. The results are depicted in Table 1, and they are a great start for a dinner table conversation, indeed. The largest danger of topic models, fortunately pointed out by many, is that the 'wow' masks the underlying questions. The method is undeterministic and has sensitive hyperparameters. Outputs can be lucky shots and the further application of LDA models is not trivial. This calls for computational linguists to help out with the experimentation. Please give us a call.

Female students	**Male students**
house mother father baby old boy brother family girl home little sister children years parents child see aunt son kids	car driving drive road truck get going go seat side back cars stop front parked turn driver street station parking
church wedding people married front sit seats seat aisle back table sitting ceremony group place priest getting side room chair	game ball playing team play basketball football field high baseball coach good balls hit players tennis school man other player
wearing white dress black blue red clothes hair dressed shirt put shoes wear pair pants hat green suit see old	gun men man shoot shot people fire kill shooting guns police war run killed enemy escape being soldiers fight get
pool swimming water swim go board end burt dive diving bottom little bathing side deep suit lake watching shallow underwater	building stairs floor go get elevator people going door top hall room roof wall high office find steps way ladder

Table 1: Topics that American female students dream of (left column) vs. topics that American male students dream of (right column)

References

Domhoff, William G. (2003). The Scientific Study of Dreams: Neural Networks, Cognitive Development, and Content Analysis. In *American Psychological Association (APA), 1 edition.*

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. In *the Journal of machine Learning research 3* (2003): 993-1022.

Mikolov, Tomas, Yih, Wen-tau, and Zweig, Geoffrey. Linguistic Regularities in Continuous Space Word Representations. In *Proceedings of NAACL HLT*, 2013.

BNVKI-AIABN Board Seeking a New Member

The board of the BNVKI-AIABN is inviting a new enthusiastic member to strengthen the team and to develop new activities. The term will start at our association's upcoming General Assembly during the BNAIC'15, November 5-6, 2015. The board uses teleconferencing for approx. 9 meetings per year and has one live meeting. For more information, please contact any of the board members. Candidates are requested to send their application to Ann Nowe at ann.nowe@vub.ac.be.

Event Reports

Seminar of Organizational Compliance

Date: **February 20th, 2015**
Venue: **Commissiekamer 3, Aula, TU Delft**
This seminar was a part of the **Graduate School** seminar program.

The seminar of "Organisational Compliance: designing and evaluating organizational interactions" was attended by around 40 attendants from various universities and interest groups. The seminar started at 9:00 with a welcome speech given by Dr Virginia Dignum, in which the speakers and the aim of the seminar were presented.

From 9:15 to 10:00, Prof Leon van der Torre and Dr Silvano Colombo Tosatto from University of Luxembourg gave a talk on the problem of proving regulatory compliance of business process models and its complexity.

From 10.00 to 10.45, Prof Wil van der Aalst from Eindhoven University of Technology gave a talk on how to divide and conquer big event data by means of process mining techniques.

From 10.45 to 11.30, Prof. Munindar Singh from North Carolina State University gave a talk on the multi-agent foundations for social computing that brings together considerations of social interaction and computing.

After each talk, there was a questioning and answering part in which the audience and the speakers discussed about different issues related to the talk.

The seminar continued with the PhD defense of Jie Jiang from Delft University of Technology. The title of the PhD thesis is *Organisational Compliance: designing and evaluating organizational interactions*. The PhD defence started with a layman talk at 12:00 in which Jie Jiang briefly explained her PhD research. The formal defense was from 12:30 to 13:00, in which Jie Jiang defended various questions posed by the committee members.

Summarizing, the seminar brought together many of the researchers that are interested in the problem of organizational compliance, and allowed for exchanging ideas and approaches as well as the latest developments and trends in various research domains.

PhD Abstracts

Studies in Learning Monotonic Models from Data

by Nicola Barile

General Information

2014-01 Nicola Barile (Utrecht University)
Studies in Learning Monotone Models from Data
Promotor: Prof.dr.A.P.J.M. Siebes (Utrecht University)
Copromotor: Dr. A.J. Feelders (Utrecht University)
Promotion: 10 February 2014

Abstract This thesis describes a number of new data mining algorithms which were the result of our research into the enforcement of monotony restrictions when learning (mostly non-parametric) models from data. Not only can judicious use of domain knowledge improve the predictive accuracy of data mining algorithms but also, crucially, models that are consistent with the knowledge of domain experts will be accepted and adopted much earlier than models that are not. Unfortunately, domain knowledge that is most of times available is often informal and poorly structured, which makes its use in practice fraught with difficulty.

099p[;The knowledge of the existence an ascending or descending relationship between predictor variables and the variable to predict represents a notable exception. Moreover, in many applications domain experts can specify such monotonic relationships with relative ease and reliability based on their knowledge and experience. It is known, for instance, that smoking and being overweight increase the risk of cardiovascular disease (an increasing relationship); on the other hand, it is likely that a higher income reduces the probability of default on a loan (a decreasing relationship).

The experiments described in this thesis show that the predictive power of our new data mining algorithms is comparable to, or sometimes even better than, that of their non-monotonic counterparts. This is obtained at a limited additional computational cost. All in all, it is possible to conclude that enforcing monotony restrictions when applicable is practically achievable and has an advantageous effect on the quality of the models produced.

Nicola is currently a data scientist at ING in the International Advanced Analytics Team, located in Amsterdam Area, the Netherlands.

Collaborations in Open Learning Environments

by Howard Spoelstra (OUN)

General Information

2015-04
Howard Spoelstra (Open University of the Netherlands)
Collaborations in Open Learning Environments
Promotor: Prof. dr.P. Sloep (Open University of the Netherlands)
Copromotor: Dr. P. Van Rosmalen (Open University of the Netherlands)
Promotion: 24 April 2015

Abstract This thesis researches automated services for professionals aiming at starting collaborative learning projects in open learning environments, such as MOOCs. It investigates the theoretical backgrounds of team formation for collaborative learning. Based on the outcomes, a model is developed describing the process of 1) project proposal assessment for fit to learning materials, and 2) project team formation based on prior knowledge and personality. Algorithms for the formation of learning, creative, and productive teams are described. A large scale experiment demonstrates the successful use of Latent Semantic Analysis for the modeling of a knowledge domain, the assessment of project fit to the learning environment, and for the assessment of learner prior knowledge. The outcomes contribute to MOOC design, team formation theory, and the LSA knowledge base.

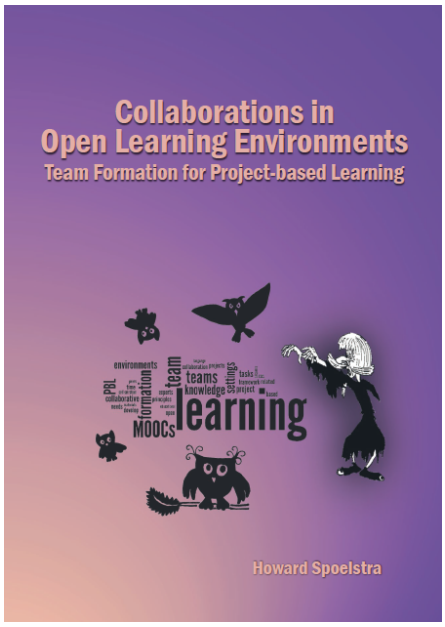

Howard Spoelstra is currently an assistant professor at the Welten Institute Research Centre for Learning, Teaching and Technology at the Open University of the Netherlands

Machine learning for network data

by Twan van Laarhoven (RUN)

General Information

2015-03
Twan van Laarhoven (Radboud University Nijmegen)
Machine learning for network data
Promotor: prof. dr. T. Heskes (Radboud University Nijmegen)
Copromotor: Dr. H. Marchiori (Radboud University Nijmegen)
Promotion: 09 February 2015

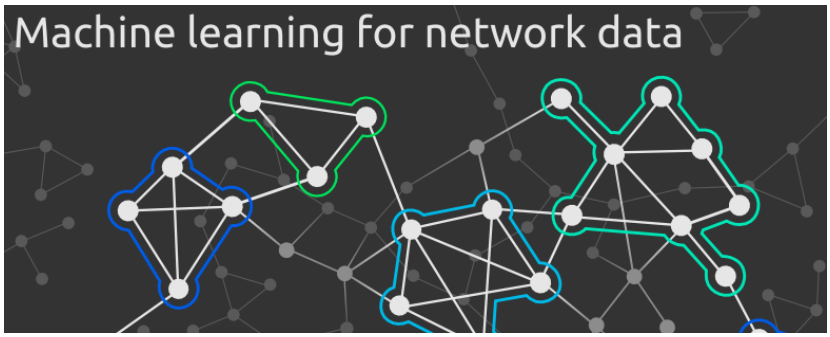

Abstract Networks are a convenient and flexible representation for many datasets. Network data often concerns interactions or relations. These can be relations between people, such as being friend; connections between neurons in the brain; or simply the notion of two things being similar to each other. But also the interactions between proteins, drugs, etc. can be considered as a network.

Given a network, there are several machine learning questions that can be asked. What sets these machine learning problems apart from traditional graph theory is that they involve some kind of inference. The goal is to look beyond the network at hand, and discover something about the underlying structure. Of course, this relies on the assumption that there is a certain underlying structure. If this assumption is wrong, then it may be impossible to say anything about the network. Additionally, we can never get certainty about the results, and we have to settle for very likely ones.

In this thesis we look at two different such machine learning problems on networks.

In part 1 we look at finding clusters in networks. Clustering is a prototypical unsupervised machine learning problem. The goal is to split a dataset into clusters, where the elements of the clusters are similar in some sense. In our setting, the elements of the dataset are the nodes in a network, and similarity is defined by the edges of that network.

In part 2 we will look at predicting links in bipartite networks. A bipartite network is one where the nodes can be split into two sets, and the only edges are between nodes in different sets. Many biological networks take this form, in particular the interaction network of drugs and target proteins which we consider in this part of the thesis.

Twan van Laarhoven successfully defended his thesis with cum laude distinction. He is currently a Post doctoral researcher at the Radboud University in Nijmegen.

Time-Aware Online Reputation Analysis

by Maria-Hendrike Peetz(UvA)

General Information

2015-07
Maria-Hendrike Peetz (University of Amsterdam)
Time-Aware Online Reputation Analysis
Promotor: prof.dr.M.de Rijke (University of Amsterdam)
Copromotor:prof.dr. W.M. van Dolen (University of Amsterdam)
Promotion: 24 March 2015

Abstract Social media has become an integral part of society. Omnipresent mobile devices allow for immediate sharing of experiences. Experiences can be about brands and other entities. For social media analysts a collection of posts mentioning a brand can serve as a magnifying glass on the prevalent opinion towards a brand: The overall estimation of a its reputation is increasingly based on the aggregation of a brand's reputation polarity in social media posts. This polarity of reputation is currently annotated manually. However, with the dramatic increase of social media, this is no longer feasible. This thesis aims to facilitate and automate parts of the process to estimate the reputation of a brand. We motivate this by performing user studies with expert social media analysts. We analyse three resulting datasets: a questionnaire, log data of a manual annotation interface, and videos of annotating experts following the think-aloud protocol. Based on the indicators used for manual annotation, we proceed with the development of algorithms for the automatic estimation of reputation polarity. Unlike earlier, static evaluation scenarios, we follow a dynamic scenario, which mimics the daily workflow of social media analysts. Our algorithms are successful because we distinguish between reputation and sentiment. The second part of this thesis is motivated by the analysts' desire for automation of retrieval and filtering of new media. For information retrieval, we present two improvements to existing algorithms. We conclude that many aspects of the annotation of reputation can be automated - using in particular time series analysis, memory models, and low-impact help from expert social media analysts.

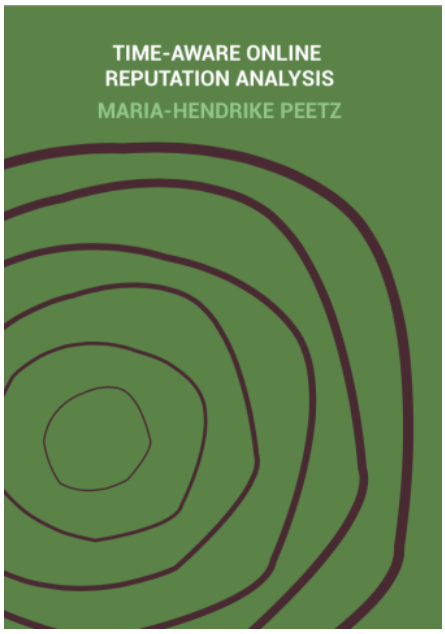

Maria-Hendrike Peetz is currently a Site Reliability Engineer at Google in Zürich, Switzerland.

New ECCAI Fellows

BNVKI is very happy with the recent ECCAI Fellow nomination of 4 of its members. The ECCAI Fellows Program honours only a very small percentage of the total membership of all ECCAI member societies. Our 2015 Fellows are:

- **Hendrik Blockeel** (Belgium - KU Leuven)
  `http://people.cs.kuleuven.be/~hendrik.blockeel/`
- **Catholijn Jonker** (Netherlands - Delft University of Technology)
  `http://ii.tudelft.nl/~catholijn/`
- **Antal van den Bosch** (The Netherlands - Radboud University):
  `http://antalvandenbosch.ruhosting.nl/`
- **Leon van der Torre** (Luxembourg - University of Luxembourg):
  `http://icr.uni.lu/leonvandertorre/`

BNVKI Membership Fees

In the table below you can find the BNVKI membership fees.

	2015
Regular members	€ 20,-
PhD students	€ 10,-
Master students	€ 10,-

Table 2: BNVKI Registration Fees

Becoming a BNVKI member makes you automatically an ECCAI member and allows you register at a reduced registration rate for certain major events, such as ECAI and ACAI. By increasing the number of BNVKI members, our AI community can also nominate more colleagues to become ECCAI fellows, as the maximum number of fellows we are allowed to have is proportional to the number of members. Finally, it might be good to know that ECCAI has decided to sponsor international events through invited speakers and these invited speakers need to be an ECCAI member over the past years.

If you want to know where our members are currently located, check out http://wilma.vub.ac.be/ dvandeun/mapje.html, if your affiliation is not represented, or you would like to see a larger dot, become a member and convince you colleagues to join as well.

BNAIC 2015

The 27th Benelux conference on Artificial Intelligence (BNAIC 2015) will take place on 5-6 November in Hasselt (Belgium). BNAIC 2015 will be held at the city campus of Hasselt University, in the unique setting of the former prison of Hasselt. BNAIC 2015 will include invited speakers, research presentations, posters and demonstrations. Authors are invited to submit papers on all aspects of artificial intelligence.

One of the keynote speakers is Dr. Elpiniki I. Papageorgiou. She is assistant Professor at the Department of Computer Engineering of the Technological Education Institute (TEI) of Central Greece, Lamia, Greece. She has been working for over thirteen years as researcher in several research projects related with the development of novel computational intelligence methodologies for decision support systems, intelligent algorithms for decision making, data analysis and mining and expert systems.

Please visit bnaic2015.org for more information.

Board Members BNVKI

How to Subscribe?

The BNVKI-AIABN Newsletter is a direct benefit of membership of the BAIAI: Benelux Association for Artificial Intelligence. Membership dues are € 20 for regular members and € 10 for students (AIO’s or master). In addition, members will receive access to the electronic version of the European journal AI Communications. The Newsletter appears quarterly. For more information, please visit our website and go to “Membership and Benefits”.

Copy

The editorial board welcomes product announcements, book reviews, product reviews, overviews of AI education, AI research in business, and interviews. Contributions stating controversial opinions or otherwise stimulating discussions are highly encouraged. Please send your submission by E-mail (MS Word or text) to board@bnvki.org.

Advertising

It is possible to have your advertisement included in the BNVKI/AIABN Newsletter. For further information about pricing etc., see our website, section “Sponsoring Rules”.