Fourteenth International Conference on

Artificial Intelligence and Law

Rome, June 10 - 14, 2013
Consiglio Nazionale delle Ricerche
(National Research Council of Italy)

Proceedings of the Conference

Sponsored by:
International Association for Artificial Intelligence and Law
ITTIG Consiglio Nazionale delle Ricerche
Artificial Intelligence, Elsevier
e-Codex
OASIS
Springer

In cooperation with:
Association for the Advancement of Artificial Intelligence (AAAI)
ACM SIGART

Under the auspices of:
the Senate of the Italian Republic
CNR - Dipartimento Scienze Umane e Sociali - Patrimonio Culturale

The Association for Computing Machinery
2 Penn Plaza, Suite 701
New York New York 10121-0701

ACM ISBN: 978-1-4503-2080-1

Conference Organization

Program Chair
Bart Verheij
Institute of Artificial Intelligence and Cognitive Engineering, University of Groningen
Groningen, The Netherlands
b.verheij@ai.rug.nl

Conference Chair
Enrico Francesconi
Institute of Legal Information Theory and Techniques - ITTIG-CNR
Florence, Italy
francesconi@ittig.cnr.it

Secretary / Treasurer
Anne v.d.L. Gardner
Atherton, California, USA
gardner@cs.stanford.edu

Organizing Committee
Daniela Tiscornia, Tommaso Agnoloni, Ginevra Peruginelli, Maria Teresa Sagri,
Institute of Legal Information Theory and Techniques - ITTIG-CNR
Giovanni Sartor, European University Institute, Florence & CIRSFID, University
of Bologna
Monica Palmirani, CIRSFID, University of Bologna
Aldo Gangemi, Institute of Cognitive Sciences and Technologies - ISTC-CNR

Programme Committee
Thomas Ågotnes, University of Bergen, Norway
Francisco Andrade, Universidade do Minho Law School, Portugal
Kevin D. Ashley, University of Pittsburgh, USA
Katie Atkinson, University of Liverpool, United Kingdom
Trevor Bench-Capon, University of Liverpool, United Kingdom
Floris Bex, University of Dundee, United Kingdom
Guido Boella, University of Turin, Italy
Alexander Boer, University of Amsterdam, The Netherlands
Daniele Bourcier, Université de Paris 2, France
Karl Branting, MITRE Corporation, USA
Joost Breuker, University of Amsterdam, The Netherlands
Scott Brewer, Harvard Law School, USA
Jan Broersen, Utrecht University, The Netherlands
Thomas Bruce, Cornell Law School, USA
Pompeu Casanovas, Autonomous University of Barcelona (UAB), Spain

Jack G. Conrad, Thomson Reuters, Switzerland
Phan Minh Dung, Asian Institute of Technology, Thailand
Tom van Engers, University of Amsterdam, The Netherlands
Norman Fenton, Queen Mary University of London, United Kingdom
Enrico Francesconi, CNR - ITTIG, Italy
Aldo Gangemi, Institute of Cognitive Sciences and Technologies - ISTC-CNR, Italy
Anne Gardner, Atherton, California, USA
Thomas Gordon, Fraunhofer FOKUS, Germany
Guido Governatori, NICTA, Australia
Davide Grossi, University of Liverpool, United Kingdom
Nicola Guarino, Consiglio Nazionale delle Ricerche, Italy
Ben Hachey, Thomson Reuters R&D, USA
Carole Hafner, Northeastern University, USA
Jaap C. Hage, University of Maastricht, The Netherlands
Bruce Hedin, H5, USA
Hans Henseler, Amsterdam University of Applied Sciences, The Netherlands
Rinke Hoekstra, VU Amsterdam/University of Amsterdam, The Netherlands
John F. Horty, University of Maryland, USA
Andrew J. I. Jones, King's College London, United Kingdom
Daniel Katz, Michigan State University, USA
Jeroen Keppens, King's College London, United Kingdom
Steven Kimbrough, University of Pennsylvania, USA
Patries Kordelaar, University of Amsterdam, The Netherlands
Marc Lauritsen, Capstone Practice Systems, USA
Arno R. Lodder, Vrije Universiteit, The Netherlands
L. Thorne McCarty, Rutgers University, USA
Raquel Mochales, Nuance, Belgium
Marie-Francine Moens, Katholieke Universiteit Leuven, Belgium
Katsumi Nitta, Tokyo Institute of Technology, Japan
Anja Oskamp, Open Universiteit, The Netherlands
Monica Palmirani, University of Bologna, Italy
Jeremy Pitt, Imperial College London, United Kingdom
Henry Prakken, Utrecht University & University of Groningen, The Netherlands
Edwina Rissland, University of Massachusetts/Amherst, USA
Antonino Rotolo, CIRSFID - University of Bologna, Italy
Giovanni Sartor, European University Institute, Florence & CIRSFID, University of Bologna
Ken Satoh, National Institute of Informatics, Japan
Burkhard Schafer, University of Edinburgh, United Kingdom
Uri J. Schild, Bar Ilan University, Israel
Erich Schweighofer, University of Vienna, Austria
Daniela Tiscornia, CNR - ITTIG, Italy
Leon van der Torre, University of Luxembourg, Luxembourg
Vern R. Walker, Maurice A. Deane School of Law at Hofstra University, USA

Douglas Walton, University of Windsor, Canada
Radboud Winkels, University of Amsterdam, The Netherlands
Adam Wyner, University of Liverpool, United Kingdom
Hajime Yoshino, Meiji Gakuin University, Japan
John Zeleznikow, Victoria University, Australia

Additional reviewers
Mi-Young Kim, University of Alberta, Canada
Yusuke Miyao, National Institute of Informatics, Japan
Kazuko Takahashi, Kwansei Gakuin University, Japan

Preface

It is not a coincidence that the research fields of Artificial Intelligence and of Law have met, and merged into the interdisciplinary research field of Artificial Intelligence and Law. Both fields use formal methods, with their strengths and limitations; in AI there are software, logic and statistics, in Law there are statutes, procedures and institutions. Both fields are creative; in AI systems are built, experiments designed and paradigms replaced, in Law regulations are passed by lawmakers, precedents are set and ideologies balanced. Both fields struggle with the inevitable complexity of modeling human behavior; in AI with the goal to reconstruct human behavior, in Law with the goal to steer human behavior. These and other similarities are driving the active and dedicated community of AI and Law. Researchers are taking their inspiration from the Law with its insights developed over millennia combining them with AI's half a century of lessons.

The beginnings of AI and Law are marked by the first International Conference on Artificial Intelligence and Law in Boston in 1987, a quarter of a century ago. Ever since that first conference, the biennial ICAIL conference series has been a primary forum for the exchange and discussion of the latest research insights in the interdisciplinary field of Artificial Intelligence and Law. The Fourteenth International Conference on Artificial Intelligence and Law (Rome, Italy, June 10-14, 2013) continues from these 25 year old achievements, and provides a program that consists of invited lectures, full papers, research abstracts, system demonstrations, workshops and tutorials.

The invited speakers of the conference are Rosaria Conte, who discusses the emergence and change of norms of different types using agent simulations as an experimental tool; Paul Thagard, who presents a neural process theory of intentions, connecting to free will and legal responsibility; and Radboud Winkels, who speaks about 25 years of AI & Law and the difficulties of turning data into knowledge. Peter van Koppen discusses the handling of evidence in law, and what can (and cannot) be expected from modeling tools.

In response to our call for contributions, we received 53 submissions, 17 of which were accepted as full papers (10 pages), and 13 of which as research abstracts (5 pages). In order to emphasize the importance of implemented systems for the field, we also called for system demonstrations; 7 were accepted for the conference, 1 of them associated with a research abstract, 6 of them described in a demonstration extended abstract (2 pages).

The papers, research abstracts and demonstrations have been selected on the basis of the scholarly reviews and discussions by the members of the Program Committee and additional reviewers. We thank them all for their hard work.

Thanks go also to the members of the organization committee: Daniela Tiscornia, Giovanni Sartor, Monica Palmirani and Aldo Gangemi, and to the other people who worked hard for the realization of ICAIL 2013: Tommaso Agnoloni, Ginevra Peruginelli, Maria Teresa Sagri, Giuseppina Sabato and Simona Binazzi. We thank the staff of the elegant conference venue Casa dell'Aviatore who helped providing the best circumstances possible for a pleasant and productive conference.

We gratefully acknowledge the support by our sponsors: the International Association for Artificial Intelligence and Law, ITTIG Consiglio Nazionale delle Ricerche, Artificial Intelligence, Elsevier, e-Codex, OASIS and Springer. The conference was organized in cooperation with the Association for the Advancement of Artificial Intelligence (AAAI) and ACM SIGART and under the auspices of the Senate of the Italian Republic and CNR - Dipartimento Scienze Umane e Sociali - Patrimonio Culturale.

Groningen/Florence/Atherton, May 2013

Bart Verheij, program chair
Enrico Francesconi, conference chair
Anne Gardner, secretary/treasurer

ICAIL 2013 Program and Schedule of Events

Monday, June 10

8:30-9:00	**Registration**
9:00-11:00	**Workshops and Tutorials**
	T1 -- Practical Web Ontology and Data Design for Legal Knowledge (full-day tutorial) *Chairs*: Aldo Gangemi, Valentina Presutti
	T2 -- An Introduction to Artificial Intelligence and Law (morning tutorial) *Chairs*: Kevin Ashley, Matthias Grabmair
	T3 -- LegalRuleML (morning tutorial) *Chair*: Monica Palmirani
	W1 -- Coherence 2013 - Artificial Intelligence, Coherence and Legal Reasoning (full-day workshop) *Chair*: Michał Araszkiewicz
	W2 -- Argumentation in AI and Law: what do we know and where should we go? (full-day workshop) *Chair*: Trevor Bench-Capon
11:00-11:30	**Coffee Break**
11:30-13:00	**Full Day Workshops continued; Morning Tutorials concluded**
13:00-14:30	**Lunch**
14:30-15:30	**Full day Workshops continued; Afternoon Workshops; Afternoon Tutorials**
	T1 -- Practical Web Ontology and Data Design for Legal Knowledge (continued) *Chairs*: Aldo Gangemi, Valentina Presutti
	T4 -- Textual Information Extraction from Legal Resources (afternoon tutorial) *Chair*: Adam Wyner
	W1 -- Coherence 2013 - Artificial Intelligence, Coherence and Legal Reasoning (continued) *Chair*: Michał Araszkiewicz
	W2 -- Argumentation in AI and Law: what do we know and where should we go? (continued) *Chair*: Trevor Bench-Capon
	W3 -- Legal Open Data: from Institutions to Crowd-sourcing (afternoon workshop) *Chair*: Monica Palmirani
15:30-16:00	**Coffee Break**
16:00-17:30	**Workshops and Tutorials concluded**
17:45-19:30	**Welcome Reception** Presentation of special issue *Artificial Intelligence and Law* on 25 years of ICAIL

Tuesday, June 11

8:30-9:15	**Registration**		
9:15-9:30	**Opening remarks**		
9:30-11:00	**Paper session**		
9:30-10:00	Davide Gianfelice, Leonardo Lesmo, Monica Palmirani, Daniele Perlo, Daniele P. Radicioni *Modificatory Provisions Detection: a Hybrid NLP Approach*	full paper	43
10:00-10:30	Guido Governatori, Francesco Olivieri, Antonino Rotolo, Simone Scannapieco *Legal Contractions: A Logical Analysis*	full paper	63
10:30-11:00	Oanh Thi Tran, Minh Le Nguyen, Akira Shimazu *Reference Resolution in Legal Texts*	full paper	101
11:00-11:30	**Break**		
11:30-13:00	**Paper session**		
11:30-12:00	Zaher Salah, Frans Coenen, Davide Grossi *Extracting Debate Graphs from Parliamentary Transcripts: A Study Directed at UK House of Commons Debates*	full paper	121
12:00-12:30	Marc van Opijnen *A Model for Automated Rating of Case Law*	full paper	140
12:30-13:00	Guido Boella, Marijn Janssen, Joris Hulstijn, Llio Humphreys, Leendert van der Torre *Managing Legal Interpretation in Regulatory Compliance*	full paper	23
13:00-14:30	**Lunch**		
14:30-15:30	**Invited Address** Rosaria Conte *Minding Norms*		xvii
15:30-16:00	**Break**		
16:00-17:15	**Demonstration session**		
	Guido Boella, Luigi Di Caro, Daniele Rispoli, Livio Robaldo *A System for Classifying Multi-label Text into EuroVoc*		239
	Luc Ferrand, Isabelle Pesquié-Geday *Hammurabi, the Legal Expert Assistant Platform for the French Judge*		241
	Thomas F. Gordon *Introducing the Carneades Web Application*		243
	Guido Governatori, Sidney Shek *Regorous: A Business Process Compliance Checker*		245

	Jop Hofste, Hans Henseler, Maurice van Keulen *Computer Assisted Extraction, Merging and Correlation of Identities with Tracks Inspector*		247
	Anna Ronkainen *TrademarkNow NameRank and NameCheck*		227
	Maya Wardeh, Adam Wyner, Katie Atkinson, Trevor Bench-Capon *Argumentation Based Tools for Policy-Making*		249
17:15-18:15	**Paper session**		
17:15-17:45	Matthias Grabmair, Kevin D. Ashley *Using Event Progression to Enhance Purposive Argumentation in the Value Judgment Formalism*	full paper	73
17:45-18:00	Radboud Winkels, Jochem Douw, Sara Veldhoen *Experiments in Automated Support for Argument Reconstruction*	research abstract	232
18:00-18:15	Shumpei Kubosawa, Kei Nishina, Masaki Sugimoto, Shogo Okada, Katsumi Nitta *A Discussion Training Support System and Its Evaluation*	research abstract	197
18:15	**Close of the day**		

Wednesday, June 12

9:15-10:45	**Paper session**		
9:15-9:45	Tara Athan, Harold Boley, Guido Governatori, Monica Palmirani, Adrian Paschke, Adam Wyner *OASIS LegalRuleML*	full paper	3
9:45-10:00	Tingting Li, Tina Balke, Marina De Vos, Julian Padget, Ken Satoh *A Model-based Approach to the Automatic Revision of Secondary Legislation*	research abstract	202
10:00-10:15	Régis Riveret, Dídac Busquets, Jeremy Pitt, Giuseppe Contissa, Antonino Rotolo, Giovanni Sartor *Vicarious Reinforcement and Ex Ante Law Enforcement: A Study in Norm-Governed Learning Agents*	research abstract	222
10:15-10:30	Anna Ronkainen *Intelligent Trademark Analysis: Experiments in Large-Scale Evaluation of Real-World Legal AI*	research abstract	227
10:30-10:45	Jack G. Conrad, John Zeleznikow *The Significance of Evaluation in AI and Law: A Case Study Re-examining ICAIL Proceedings*	research abstract	186

10:45-11:15	**Break**		
11:15-13:00	**Paper session**		
11:15-11:45	Marc Lauritsen *On Balance*	full paper	83
11:45-12:15	Tomasz Zurek, Michał Araszkiewicz *Modeling Teleological Interpretation*	full paper	160
12:15-12:30	Doris Liebwald *Vagueness in Law: A Stimulus for 'Artificial Intelligence & Law'*	research abstract	207
12:30-12:45	Michał Araszkiewicz, Agata Łopatkiewicz, Adam Zienkiewicz *Factor-Based Parent Plan Support System*	research abstract	171
12:45-13:00	Gordon J. Pace, Fernando Schapachnik *Synthesising Implicit Contracts*	research abstract	217
13:00-14:30	**Lunch**		
14:30-15:30	**Invited Address** Radboud Winkels *25 Years of AI & Law*		xxi
15:30-16:00	**Break**		
16:00-17:00	**Paper session**		
16:00-16:30	Charlotte Vlek, Henry Prakken, Silja Renooij, Bart Verheij *Modeling Crime Scenarios in a Bayesian Network*	full paper	150
16:30-16:45	Kevin D. Ashley, Vern R. Walker *Toward Constructing Evidence-Based Legal Arguments Using Legal Decision Documents and Machine Learning*	research abstract	176
16:45-17:00	Michal Chalamish, Moshe Hazoom, Uri Schild *Semi-Automatic Creation of Wigmore Diagrams*	research abstract	181
17:15-18:15	**IAAIL membership meeting** Tribute to Peter Jackson		
20:00	**Conference dinner**		
	Dinner speech Peter van Koppen *Spoiled Judges and the Fluidity of Evidence in Criminal Cases*		xxii

Thursday, June 13

10:00-11:00	**Paper session**		
10:00-10:30	Laura Giordano, Alberto Martelli, Daniele Theseider Dupré *Temporal Deontic Action Logic for the Verification of Compliance to Norms in ASP*	full paper	53
10:30-11:00	Antonino Rotolo, Serena Villata, Fabien Gandon *A Deontic Logic Semantics for Licenses Composition in the Web of Data*	full paper	111
11:00-11:30	**Break**		
11:30-13:00	**Paper session**		
11:30-12:00	Isabella Distinto, Nicola Guarino, Claudio Masolo *A Well-founded Ontological Framework for Modeling Personal Income Tax*	full paper	33
12:00-12:30	Antonio Mastropaolo, Francesco Pallante, Daniele P. Radicioni *Legal Documents Categorization by Compression*	full paper	92
12:30-12:45	Michael Curtotti, Eric McCreath, Srinivas Sridharan *Software Tools for the Visualization of Definition Networks in Legal Contracts*	research abstract	192
12:45-13:00	Nada Mimouni, Meritxell Fernàndez, Adeline Nazarenko, Danièle Bourcier, Sylvie Salotti *A Relational Approach for Information Retrieval on XML Legal Sources*	research abstract	212
13:00-14:30	**Lunch**		
14:30-15:30	**Invited Address** Paul Thagard *Intention, Action, and Legal Responsibility*		xix
15:30-16:00	**Break**		
16:00-17:00	**Paper session**		
16:00-16:30	Trevor Bench-Capon, Henry Prakken, Adam Wyner, Katie Atkinson *Argument Schemes for Reasoning with Legal Cases Using Values*	full paper	13
16:30-17:00	Mihai Surdeanu, Sara Jeruss *Identifying Patent Monetization Entities*	full paper	131
17:00	**Closing Remarks**		

Friday, June 14

8:30-9:00	**Registration**
9:00-11:00	**Workshops**
	W4 -- 13th International Workshop on Computational Models of Natural Argument (CMNA XIII) (full-day workshop) *Chairs*: Floriana Grasso, Chris Reed, Nancy Green
	W6 -- Workshop on Formal Argument and Evidential Inference (full-day workshop) *Chairs*: Giovanni Sartor, Scott Brewer, Gustavo Ribeiro
	W7 -- Discovery of Electronically Stored Information Workshop (DESI V) (full-day workshop) *Chair*: Jason Baron
	W8 -- Cross-border e-justice and e-Codex (morning workshop) *Chair*: Marco Fabri
	W9 -- Network Analysis in Law (morning workshop) *Chair*: Radboud Winkels
11:00-11:30	**Coffee Break**
11:30-13:00	**Full Day Workshops continued; Morning Workshops concluded**
13:00-14:30	**Lunch**
14:30-15:30	**Full Day Workshops continued; Afternoon Workshops**
	W4 -- 13th International Workshop on Computational Models of Natural Argument (CMNA XIII) (continued) *Chairs*: Floriana Grasso, Chris Reed, Nancy Green
	W6 -- Workshop on Formal Argument and Evidential Inference (continued) *Chairs*: Giovanni Sartor, Scott Brewer, Gustavo Ribeiro
	W5 -- SmartData: the New Face of AI, and the Law (afternoon workshop) *Chairs*: Ann Cavoukian, Stefano Nolfi
	W7 -- Discovery of Electronically Stored Information Workshop (DESI V) (continued) *Chair*: Jason Baron
15:30-16:00	**Coffee Break**
16:00-17:30	**Workshops concluded**

Author Index

INVITED SPEAKERS

Minding Norms[1]

Rosaria Conte
LABSS (Laboratory of Agent Based Social Simulation)
at the ISTC-CNR (Institute for Cognitive Science and Technology of the CNR)

Under the multiform treatment that norms have received so far in the scientific literature, it is possible to envisage a dichotomy. On one hand, social scientists view norms as *regular behaviours*, possibly enforced by social expectations and sanctions, and seeing none or poor reason for a specific, norm-related form of cognition. On the other hand, philosophers of law and logicians conceptualise norms as specific mental representations. The former see norms as behavioural regularities, the latter as *obligations issued* by definite authorities. Hence, the first set of questions that will be addressed: *what are the differences and commonalities among various types of norms*, social, moral, and legal? Is it possible to aim at an integrated view of norms, at the same time keeping them distinct from social conformism on one hand and acquiescence under menace on the other? In the approach presented in the talk, norms will be conceptualised as hybrid - social and cognitive - phenomena undergoing a complex dynamics, in which the social process of emergence and the mental process of immergence are intertwined in a circular fashion.

A second set of questions will be examined, i.e. *how do norms emerge and change*? The hybrid view of norms characterising our approach allows this question to be addressed in a rather innovative way: the interplay between the mental and the social dynamics allows norms to emerge and change. The main thesis is that observable conformity is only the tip of the normative iceberg. Norms cannot emerge in society if they do not previously immerge in the mind, i.e. if they are not first converted into mental representations of some sort. The crucial dynamics lies in the minds of the agents, beneath the line of observation. The mental dynamics of norms, called immergence, also allows a subset of the initial questions to be answered. Norm immergence is common to different types of norms, but at the same time allows norm-based behaviour to be kept distinct: agents abiding with norms, or violating them, act on a set of specific, norm-related, mental representations.

The mental dynamics of norms brings about a third set of questions that will be addressed: *how should we characterize the agents from among which norms emerge*? Why and how do people represent, reason upon, abide with or violate norms? We intend to tackle this set of questions by means of a computational, agent-based, approach, at the same time presenting an integrated and highly dynamic view of norms. We propose a modular normative architecture, accounting for the mental representation and dynamics of norms.

[1] This abstract reports on research performed in collaboration with Giulia Andrighetto.

After a review of different approaches, we will present a dynamic model of norms, a normative agent architecture, a simulation platform and artificial experiments testing the view of norms and the architecture proposed against a number of more or less realistic social scenarios.

Intention, Action, and Legal Responsibility

Paul Thagard
University of Waterloo

The concept of intention is important in many disciplines, including philosophy, psychology, artificial intelligence, cognitive neuroscience, and law. Criminal law treats cases where one person intends to kill another very differently from cases where death results unintentionally from negligence. Despite decades of discussions, however, there is no received theory of intention within any of these disciplines, let alone a theory that accounts for all the phenomena identified across all of the disciplines.

Schröder, Stewart, and Thagard (in press) have proposed a unified theory of intentions as neural processes that integrate representations of states of affairs, actions, and emotional evaluation. Our theory of intention ties together biologically plausible mechanisms for belief, planning, and motor control. The feasibility of these mechanisms is shown by a computational model that successfully simulates psychologically important cases of intention. These simulations support the plausibility of the claim that human intentions are neurocomputational processes operating in the brains of individuals.

We defend the following theoretical claims:

1. Intentions are semantic pointers, which are patterns of activity in populations of spiking neurons that function as compressed representations by binding together other patterns.
2. Specifically, intentions bind representations of situations, emotional evaluations of situations, the doing of actions, and sometimes the self.
3. Intentions can cause actions because of neural processes that connect semantic pointers with motor instructions.
4. Intentions can fail to cause actions because of various kinds of disruptions affecting any of:

    (a) Evaluation of the situation and doing.
    (b) Binding of the evaluation, situation, and doing.
    (c) Processes that connect the intention semantic pointer with motor processes.

According to Eliasmith (2013), semantic pointers are neural processes that (1) provide *shallow* meanings through symbol-like relations to the world and other representations, (2) expand to provide *deeper meanings* with relations to perceptual, motor, and emotional information, (3) support complex syntactic operations, and (4) help to control the flow of information through a cognitive system to accomplish its goals. A semantic pointer consists of spiking patterns in a large population of neurons that provide a kind

of compressed representation, analogous to JPEG picture files or iTunes audio files. The term “pointer” comes from computer science where it refers to a kind of data structure that gets its value from a machine address to which it points. A semantic pointer is a neural process that compresses information in other neural processes to which it points and into which it can be expanded when needed.

My talk at ICAIL: (1) describes how semantic pointers provide a new synthesis of symbolic and neural network approaches to cognition; (2) shows how the semantic pointer theory of intention and action can explain important psychological phenomena such as intention-action gaps; (3) argues that this account of intention and action undermines assumptions about free will built into most legal systems; and (4) discusses alternative understandings of legal responsibility more consistent with how brains make decisions.

References

Eliasmith, C. (2013). *How to build a brain*. Oxford: Oxford University Press.
Schröder, T., Stewart, T. C., & Thagard, P. (in press). Intention, emotion, and action: A neural theory based on semantic pointers. *Cognitive Science*.

25 Years of AI & Law
Presidential address at ICAIL 2013

Radboud Winkels
Leibniz Center for Law, University of Amsterdam, the Netherlands
winkels@uva.nl

In 1988 I went to my first conference on AI & Law, also the first JURIX conference in Amsterdam. At the time I was not working in the field, but in that of intelligent tutoring systems and the overlap between the two was — and is — small. It interested me however and a few years later I was completely involved.

In this presentation I will try to give *my* overview of the field, how it evolved in the last 25 years, what its strengths and weaknesses are, and where I think it should (also) be headed. I will illustrate my ideas with data from past ICAIL and JURIX conferences and some research by myself and others. I do not pretend to give a complete or perhaps even accurate overview (cf. the excellent special issue of AI & Law by Bench-Capon, editor, 2012[1]), but hope to stimulate discussion and future research.

I will discuss three 'trends' I see over the past decades in our research:

1. A move from 'deep modelling' of (legal) knowledge to more 'shallow' data and document analysis, perhaps from more 'subjective' interpretation to more 'objective' methods;
2. Less emphasis on or interest in 'building systems';
3. The rise of 'argumentation'.

Although 'law' is pervasive in society, our multidisciplinary community is small. This is not intrinsically bad, but at times it seems too closed and internally oriented. Our research and community could benefit more from neighbouring disciplines (like AI and computer science, linguistics, economics or even social sciences) and contribute more to these fields at the same time. Even more surprisingly perhaps, is the lack of interaction with the research field of law. Not many of our researchers appear at law conferences, nor do many legal scholars attend ours.

This is in contrast with an apparent growing interest in public administrations, governments and large companies in rule management solutions, e-discovery and data mining in (para-) legal contexts and smart government.

How to resolve this conflict? One way out as I see it, is to refocus part of our work on solving practical problems, validating our theories and methods in real life situations and a rigorous comparison to competing approaches, also human ones. We would benefit from a 'Deep Blue' or 'Watson' like challenge!

[1] A History of AI and Law in 50 papers: 25 Years of the International Conference on AI and Law.

Spoiled Judges and the Fluidity of Evidence in Criminal Cases

Peter J. van Koppen
Maastricht University and Free University Amsterdam

The introduction of DNA changed much in the perception of evidence in criminal cases. Judges, attorneys, prosecutors and police officers were used to having witnesses deliver most of the evidence. Okay, sometimes there was technical forensic evidence, such as fingerprint evidence. The latter was considered infallible; which it, in fact, is not. Apart from fingerprints, evidence was always discussed in terms of 'strong evidence' or 'weak evidence' or 'circumstantial evidence'.

DNA-evidence has changed all this. Now, experts dare to give numerical values to their evidence. They speak of random match probabilities: the chance that by sheer accident an individual has the same DNA-profile as the sample from the crime scene. They come up with enticingly small values like 1 in a billion, in cases where a full DNA-profile is available from the crime scene that matches the profile of the suspect.

The numerical values of DNA-evidence have led to three dramatic changes in the field of criminal law. First, judges and prosecutors and attorneys now expect expert witnesses to produce numbers. I particularly notice such if I am teaching judges and prosecutors together with colleagues from the forensic laboratory. After their class on DNA-evidence, judges start to whine about number for all sorts of evidence. They suddenly expect me to give an exact diagnostic value or likelihood ratio for, for instance, a positive identification in a line-up. I have a hard time explaining to them how silly that is. Of course we could produce some number for a particular line-up as we study them in the psychological laboratory. Line-ups in police practise tend to differ markedly from those in our labs. When the police line-ups differ dramatically from what is desirable, calculations might become of interest. But also: if the police make all kinds of errors, these fine laboratory studies are not very good models for what is happening and are useless for estimating a diagnostic value.

A second dramatic recent change is that judges have become spoiled by these small numbers like 1 in a billion, meaning that the evidence is extremely strong. Since they are lawyers who are not used to calculations, they misperceive numbers. The following fine case can illustrate this. In a place near Rotterdam two men robbed an Aldi Supermarket. Something rare happened: shots were fired and the supermarket owner died. The shooter fled the scene of the crime and his accomplice was apprehended shortly after. He told the police his robbery buddy was Barry and that Barry was the shooter. Barry was arrested but denied all allegations. Apart from the witness statement of the other robber, proof against him was a shoe the shooting robber had left at the scene of the crime. The shoe was Barry's size. But, amazingly, the laboratory did not succeed in securing a DNA-profile from the shoe. There was something else: the

supermarket owner had started a fight such that blood of the robber was found on his nails. The laboratory found a partial DNA-profile in the blood that matched the DNA of Barry. The DNA-expert estimated the rarity of the DNA-profile with a random match probability of 1 in 100,000. Now, assume that evidence with this property would be the only evidence against the defendant in 100,000 different cases. And assume that a judge would use such a single piece of evidence to convict each of these 100,000 robbers. That would mean that the judge would commit a miscarriage of justice in roughly 1 out of 100,000 cases. To me, that seems quite an acceptable error rate.

Yes, I know some of you would accuse me that I now commit the so-called prosecutor's fallacy. And indeed, I do, but that only is a problem if you are a Bayesian. Bayesianism, however, is a religion that, if applied to evidence in criminal cases may cause more problems than it solves. That is another discussion. But, even Bayesians should admit that although DNA evidence with a random match probability of 1 to 100,000 is not as strong as DNA evidence can be, it is still quite strong, nevertheless. But in the Aldi case it was not strong enough for the judges: the defendant was wrongly acquitted both by the trial court in Rotterdam and the appellate court in The Hague. It is fair to say that both courts misunderstood the strength of the evidence, mainly because they are so spoiled by DNA experts presenting them with silly small random match probabilities. They are so used to evidence that is presented as very strong that, if such evidence is not present, they become unsure and tend to acquit, rather than taking the chance of making an error, a miscarriage of justice.

In reality all these numbers are nonsense. At best a number given to the strength of DNA evidence is a shot in the dark, but most of the time just rubbish. There are so many assumptions that must hold before such a number may be true, that these numbers are almost always meaningless. One must assume that the database that is used to estimate the strength is in all respects representative for the specific DNA-comparison at hand. One must assume that no mistakes have been made in the whole chain from the crime scene to the final report by the DNA-expert. One must also assume that the DNA-comparison is simple, between a full profile from the defendant to a full profile found in a crime scene sample. In practice, however, the samples found at the crime scene almost always are mixtures from two or more individuals. With these mixtures calculations soon become quite difficult, often too difficult to do anyway.

A third consequence of the apparent calculability of evidence is that some people think that it is possible to model evidence in a manner that would lead to a meaningful analysis of evidence that in turn can be used to decide the case. I am sorry to say that is a misconception. I should admit that I secretly know that such models in the form of evidence charts are much older that DNA profiles. But admitting that would spoil my argument.

My argument begins with the following. The evidence structure of criminal cases is very simple. In most cases that come to court it is evident that the defendant committed the crime. There is abundant evidence and the defendant confessed to the crime or both. Whatever system one uses to evaluate the evidence and to decide the case, it does not matter. The outcome is always the same: guilty. There may be all kinds of legal problems and the punishment may be a point of discussion, but the question of

guilt does not pose a problem. One could even leave the decision to laymen like a lay jury.

Only a relatively small proportion of the criminal cases are difficult from an evidentiary point of view. For these cases it may matter what method is used to evaluate the evidence. Would then modelling like evidence charts be the right option? Let us follow for a moment the work of the New Evidence Scholarship, represented by people like Terry Anderson, David Schum and William Twining, work that is rooted in the work by Wigmore, and that has recently inspired researchers in the field of AI & Law. They propose a plan by which it can be decided whether the ultimate probandum has been proven. They start out by making a key list, a list of all propositions in a case, all elements that are relevant. Apart from the simplest case where, we just ascertained, no special decision method is necessary, such a key list is almost always very extensive, too extensive to be viable. At that very point, a lot of choices have to be made to limit the extent of the discussion on the case. That limiting is done at every stage of the development of the case. Suppose, for instance, that in the room where you are now reading this piece, a murdered person is found and suppose you are the technical detective investigating the crime scene. Where would you take samples? Would you, for instance, dust all the pages of all books in the bookcase for fingerprints? Would you take DNA-samples from all surfaces in the room? And what would you do with the other rooms in the house? From the very beginning of the case choices are made and the choices are based on scenarios the people involved have, implicitly of explicitly. If a key list is made, the key list is guided and limited by what we think might be relevant; and what we think that might be relevant again depends on what we expect to be relevant scenarios.

After the key list, charts are made. A simple problem will produce a simple chart. A more complicated problem will not just produce a more complicated chart, but very soon a chart that is prohibitively complicated. Especially cases where an analysis using charts might be useful will produce charts that are so complicated that they are not useful any more. And, again, a chart represents a particular scenario of a limited set of scenarios for the case. In reality there are often all types of scenarios that may be more or less relevant for an analysis of the case. Also, charts are not suitable for complicated pieces of evidence, as for instance pieces of evidence that can mean different things, depending on the context in which they are found.

In short, charts are useful for an analysis of local problems in evidence evaluation. But they will never succeed in modelling the evidence in a complete case in such a way that it is essential for making a good decision.

Contents

FULL PAPERS

RESEARCH ABSTRACTS

DEMONSTRATION EXTENDED ABSTRACTS