



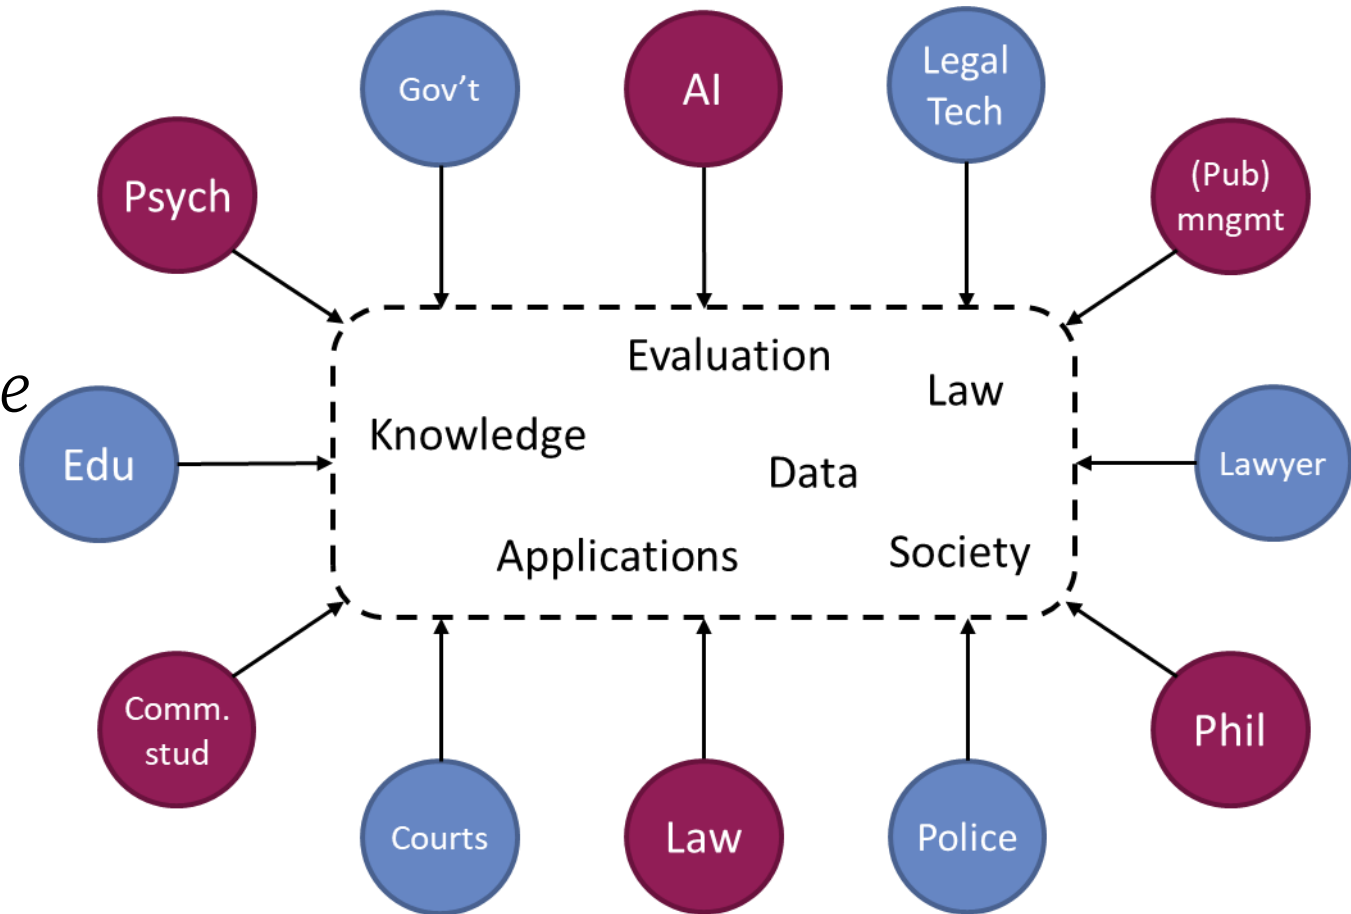
Utrecht University

AI, Law and beyond

A transdisciplinary ecosystem for the future of AI & Law

Floris Bex

Scientific Director National Police Lab AI (Utrecht)
Associate Professor Artificial Intelligence (Utrecht University)
Full Professor Data Science and the Judiciary (Tilburg University)



35 years of AI and Law

The Future of AI & Law

L. Karl Branting

20 Years of ICAIL – Reflections on the Field of AI and Law

Thomas F. Gordon

25 Years of AI & Law ICAIL 1987 - 2013

Radboud Winkels

AI and Law in 2017: Turning the hype into real world solutions

Katie Atkinson



Artificial Intelligence & Law: Through the Lens of IAAIL – Past, Present and Future

Jack G. Conrad, IAAIL President

The Winter, the Summer and the Summer Dream of AI in Law

The journey of ICAIL conference series from my perspective

Enrico Francesconi




AI & Law is doing great in 2023



Natural Legal Language Processing - NLLP

Society for Empirical Legal
Studies (SELS)



CompuLaw

GOVERNANCE OF COMPUTATIONAL ENTITIES THROUGH
AN INTEGRATED LEGAL AND TECHNICAL FRAMEWORK




LEDs4XAIL



Hybrid
Intelligence



algorithms Justice



Scimago Journal Rank (Q1-Green,..., Q4-Red)

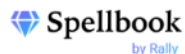
AI & Law is doing great in 2023

Pakistani judge uses ChatGPT to make court decision

Colombian judge says he used ChatGPT in ruling

China's court AI reaches every corner of justice system, advising judges and streamlining punishment

South China Morning Post, July 2022



Draft contracts 10x faster with AI

Sequoia and OpenAI Back Harvey to Redefine Professional Services, Starting with Legal

Developing AI for the legal sector: innovation and collaboration at the University of Liverpool


Artificial Intelligence for Law Enforcement and Community Safety (AiLECS) Lab



Utrecht University



National Policelab AI



AI & Law is doing great in 2023 (right??)

FEATURE-As Malaysia tests AI court sentencing, some lawyers fear for justice

Reuters, April 2022

Automating Fairness?

Artificial Intelligence in the Chinese Courts

RACHEL E. STERN,* BENJAMIN L. LIEBMAN,**
MARGARET E. ROBERTS,*** AND ALICE Z. WANG****

Colombian judge in ruling

Pakistani judge uses ChatGPT to

A lawyer used ChatGPT to cite bogus

cases. What are the ethics?

Reuters, May 2023

OpenAI's GPT-4 Is Closed Source and Shrouded in Secrecy

Vice, 16-3-2023

WE SENSE TROUBLE

AUTOMATED DISCRIMINATION AND MASS SURVEILLANCE IN PREDICTIVE POLICING IN THE NETHERLANDS

Amnesty, September 2020

ChatGPT Lands OpenAI in Legal Trouble, Globally

AnalyticsIndia, 2023

'The Godfather of A.I.' Leaves Google and Warns of Danger Ahead

NY Times, May 2023

An “algorithmic drama”

Futurama - Fox



Robocop - MGM




Ziewitz (2016). Governing algorithms: Myth, mess, and methods. *Science, Technology, & Human Values*

people.dsv.su.se/~jpalme/reports/right-wrong.html



Robocar Poli – Roi Visual





An “algorithmic drama” in AI (& Law)

- Data-driven deep learning vs. knowledge-driven reasoning
- Techno-optimism vs. techno scepticism
- Building AI vs. Regulating AI



Human-centered Artificial Intelligence



Governing the Digital Society

Stepping away from the drama – A way forward for AI & Law

1. Combine knowledge & data
 - Use new techniques without forgetting about the old ones
2. Evaluate how AI & Law is being used in practice
 - Develop and broadly evaluate AI & Law applications
3. Combine multiple disciplines
 - Law, AI, and beyond



- Collaboration between police and universities
 - Research, develop & evaluate AI for real police problems, in actual police context
- 20 PhDs
 - Majority also works at police
 - 2/3 CS/AI background, 1/3 other (information systems, public management, law, communication studies)

Example 1: AI for citizen complaint/report intake

- Trade fraud: false webshops, malicious traders on Ebay
 - 40,000+ reports of alleged online fraud per year
 - Not all fraud: wrong product, not paid
- Automatically recommend to file report or not
 - Citizen fills in a form w. details & free text story
 - Possible fraud or not?

The screenshot shows the Dutch Police website's 'Aangifte internetoplichting' (Internet Scam Report) form. The header includes the police logo and contact information: 'Bij spoed: 112' and 'Geen spoed: 0900-8844'. The navigation bar has links for 'Home', 'Aangifte of melding doen', 'Mijn buurt', 'Nieuws', 'Gezocht & Vermist', 'Thema's', and 'Over de'. The main heading is 'Aangifte internetoplichting'. Below it, a note states: 'Vul onderstaande velden zo volledig mogelijk in. Wij wijzen u erop dat alles naar waarheid ingevuld moet worden.' Another note says: 'De velden met een sterretje (*) moet u in elk geval invullen.' The form is divided into four steps: '1. Aangever', '2. Wederpartij', '3. Conflict', and '4. Overzicht'. The '1. Aangever' step is active. The form fields are organized into two main sections: 'Advertentie gegevens' (Advertisement details) and 'Transactiegegevens' (Transaction details). The 'Advertentie gegevens' section includes: 'Waar bent u opgelicht?' (Where were you deceived?), 'Advertentietitel' (Advertisement title), 'Advertentie nummer' (Advertisement number), 'Uw accountnaam' (Your account name), 'Accountnaam wederpartij' (Counterparty account name), and 'Wat is er gebeurd?' (What happened?). The 'Transactiegegevens' section includes: 'Uw bankrekeningnummer' (Your bank account number), 'Datum betaling' (Payment date), and 'Tijdstip betaling' (Payment time).

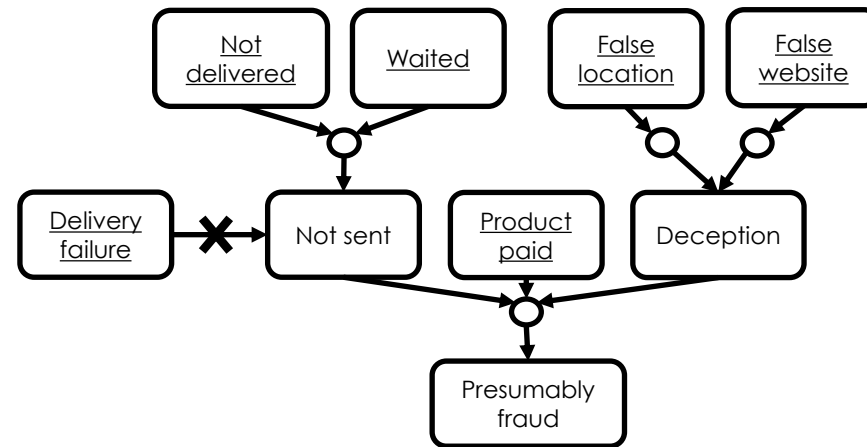
AI for intake – data & knowledge

- Combine data- and knowledge-driven AI
 - Relevant legal rules are known, bounded domain
 - Free-text interpretation needs data-driven AI

The screenshot shows the Dutch Police (Politie) website interface for reporting internet fraud. At the top, there's a header with the contact number 112 and 0900-8844, and the Politie logo. Below the header is a navigation bar with links: Home, Aangifte of melding doen, Mijn buurt, Nieuws, Gezocht & Vermist, Thema's, and Over de Politie. The main heading is 'Aangifte internetoplichting'. Below this, there's a note: 'Vul onderstaande velden zo volledig mogelijk in. Wij wijzen u erop dat alles naar waarheid ingevuld moet worden.' and a reminder: 'De velden met een sterretje (*) moet u in elk geval invullen.' There are four steps: 1. Aangever, 2. Wederpartij, 3. Conflict, and 4. Overzicht. The 'Aangifte gegevens' section includes fields for 'Waar bent u opgelicht?' (dropdown), 'Advertentietitel' (text), 'Advertentie nummer' (text), 'Uw accountnaam' (text), 'Accountnaam wederpartij' (text), and 'Wat is er gebeurd?' (large text area). The 'Transactiegegevens' section includes 'Uw bankrekeningnummer' (text), 'Datum betaling' (dd-mm-jjjj), and 'Tijdstip betaling' (uu:mm).

AI for intake – legal model

Legal model



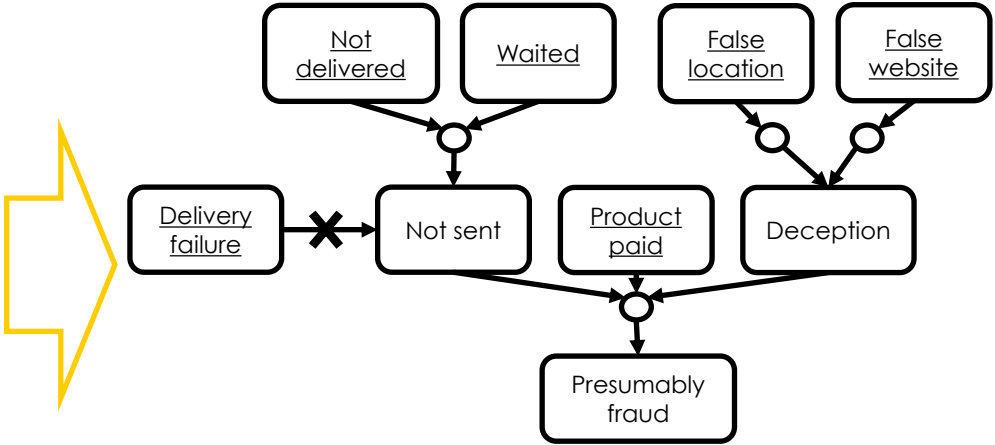
Computational argumentation
*Rules w. exceptions based on
DCC & police policy rules*

AI for intake – free text

Complaint form

Fictitious example report 1
I would like to report fraud. I recently saw a bicycle for sale on eBay and contacted the advertiser. He said he lived far away, so he would send me the bike. I paid him in good faith, but have still not received anything. I saw on Facebook he lives nearby.

Legal model



Computational argumentation
*Rules w. exceptions based on
DCC & police policy rules*

AI for intake – combining IR and argumentation

Extracting observations
from complaint form

Paid

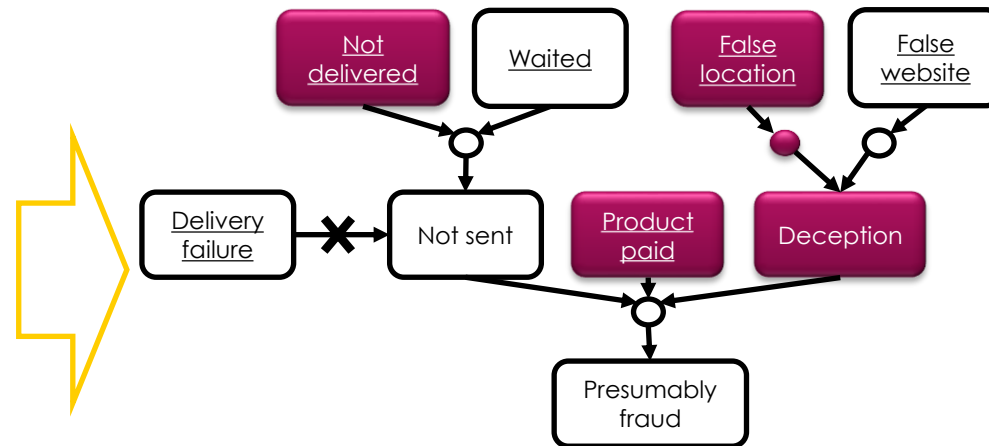
Fictitious example report 1
I would like to report fraud. I recently saw a bicycle for sale on eBay and contacted the advertiser. He said he lived far away, so he would send me the bike. I paid him in good faith, but have still not received anything. I saw on Facebook he lives nearby.

False location

Not delivered

Basic information extraction

Inferring possible fraud (or not)



Computational argumentation
*Rules w. exceptions based on
DCC & police policy rules*

AI for intake – asking the right questions

Extracting observations
from complaint form

Paid

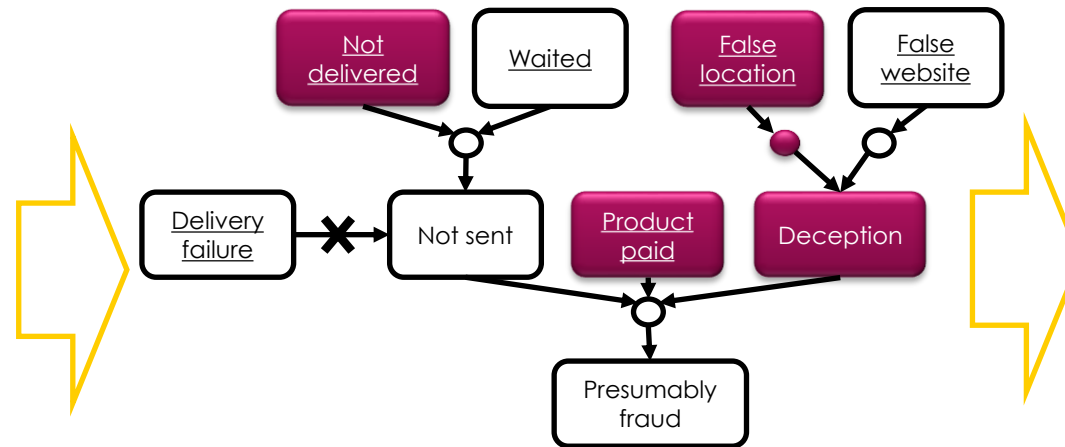
Fictitious example report 1
I would like to report fraud. I recently saw a bicycle for sale on eBay and contacted the advertiser. He said he lived far away, so he would send me the bike. I paid him in good faith, but have still not received anything. I saw on Facebook he lives nearby.

False location

Not delivered

Basic information extraction

Inferring possible fraud (or not)



Computational argumentation
*Rules w. exceptions based on
DCC & police policy rules*

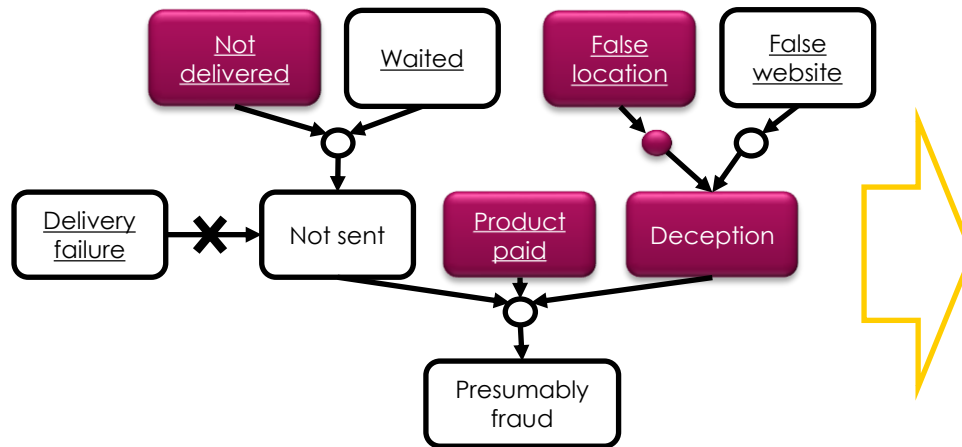
Asking for missing
observations



Approximation algorithms
*Can new info still change the
conclusion (and if so which)?*

AI for intake - explanations

Inferring possible fraud (or not)



Computational argumentation
*Rules w. exceptions based on
DCC & police policy rules*

Response

Thank you for your complaint. In your case, the system has concluded that it is not a case of fraud, since you did not wait for at least 5 days. We recommend you do not file an official report at this point.

Explanations

*Explaining (non-)acceptance in terms
of arguments and counterarguments*

AI for intake - evaluation

- Evaluate accuracy, user satisfaction
- Investigate citizen trust in automatic recommendations
 - How do users perceive recommendations by the system?
 - Do explanations matter?

The screenshot shows the Dutch Police (Politie) website's online intake form for internet reporting. The header includes the contact number 112 (for emergencies) and 0900-8844 (for non-emergencies), along with the 'POLITIE' logo. A navigation bar contains links: Home, Aangifte of melding doen (selected), Mijn buurt, Nieuws, Gezocht & Vermist, Thema's, and Over de... Below the navigation bar, the page title is 'Aangifte internetoplichting'. A note states: 'Vul onderstaande velden zo volledig mogelijk in. Wij wijzen u erop dat alles naar waarheid ingevuld moet worden.' Another note says: 'De velden met een sterretje (*) moet u in elk geval invullen.' A progress bar shows four steps: 1. Aangever (selected), 2. Wederpartij, 3. Conflict, and 4. Overzicht. The form is divided into two main sections: 'Advertentie gegevens' and 'Transactiegegevens'. The 'Advertentie gegevens' section includes fields for 'Waar bent u opgelicht?' (a dropdown menu), 'Advertentietitel', 'Advertentie nummer', 'Uw accountnaam', 'Accountnaam wederpartij', and 'Wat is er gebeurd?' (a large text area). Each of these fields has a red asterisk (*) indicating it is required, and the last four fields have a blue question mark icon to the right. The 'Transactiegegevens' section includes fields for 'Uw bankrekeningnummer', 'Datum betaling' (with a date picker showing dd-mm-jjjj), and 'Tijdstip betaling' (with a time picker showing ww:mm).

AI for intake – citizen trust & explanations

- Do citizens trust the system with and without an explanation?
 - Controlled experiments 1700+ participants
- Not fraud – still file an official report? (trusting behaviour)?
 - No explanation (control): 40-60% still filed report
 - With explanation: only 20-35% still filed report

AI, transparency and citizen trust

- Transparency
 - About the system/decision: XAI
 - About the use of systems in the organisation
 - About AI and regulation
- How do citizens react when AI is more contentious?
 - Predictive policing
- What's the influence of basic trust in police?
 - US vs. Netherlands

Example 2: AI for (explainable) text classification

- Police generate, use and analyse lots of text data
 - Citizen reports, Incident reports, Lab reports, Social Media, Seized Data Carriers
- Text classification for search, for use in AI systems

Fictitious example report 1

I recently saw a bicycle for sale online and contacted the advertiser. He said he lived far away, so he would send me the bike. I paid him in good faith, but have still not received anything. I saw on Facebook he lives nearby.

Paid



Not paid

Fictitious example report 2

I wanted to buy champagne from John Doe via Ebay. Up to now, I have not received anything, and he does not respond to my e-mails, so I haven't transferred the money yet.

Threat



No Threat



AI for explainable text classification

- Being able to explain why something was classified
 - Model testing and improvement
 - Transparency and accountability
 - Use in legal (criminal) cases

Fictitious example report 1

I recently saw a bicycle for sale online and contacted the advertiser. He said he lived far away, so he would send me the bike. I paid him in good faith, but have still not received anything. I saw on Facebook he lives nearby.

Paid



Not paid

Fictitious example report 2

I wanted to buy champagne from John Doe via Ebay. Up to now, I have not received anything, and he does not respond to my e-mails, so I haven't transferred the money yet.

Threat



No Threat



Explainable text classification - Rationales

- Explaining text classification
 - Using *machine generated rationales* (highlighted sentences)

Fictitious example report 1

I recently saw a bicycle for sale online and contacted the advertiser. He said he lived far away, so he would send me the bike. I paid him in good faith, but have still not received anything. I saw on Facebook he lives nearby.

Paid



Not paid

Fictitious example report 1

I recently saw a bicycle for sale online and contacted the advertiser. He said he lived far away, so he would send me the bike. **I paid him in good faith**, but have still not received anything. I saw on Facebook he lives nearby.

Fictitious example report 2

I wanted to buy champagne from John Doe via Ebay. Up to now, I have not received anything, and he does not respond to my e-mails, **so I haven't transferred the money yet.**

Fictitious example report 2

I wanted to buy champagne from John Doe via Ebay. Up to now, I have not received anything, and he does not respond to my e-mails, so I haven't transferred the money yet.

Explainable text classification - Counterfactuals

- Explaining text classification
 - By generating *counterfactual example text*

Fictitious example report 1

I recently saw a bicycle for sale online and contacted the advertiser. He said he lived far away, so he would send me the bike. **I paid him** in good faith, but have still not received anything. I saw on Facebook he lives nearby.

Paid



(not paid)

Fictitious example report 1

I recently saw a bicycle for sale online and contacted the advertiser. He said he lived far away, so he would send me the bike. **I did not pay him** in good faith, but have still not received anything. I saw on Facebook he lives nearby.

Fictitious example report 2

I wanted to buy champagne from John Doe via Ebay. Up to now, I have not received anything, and he does not respond to my e-mails, **so I haven't transferred the money yet.**

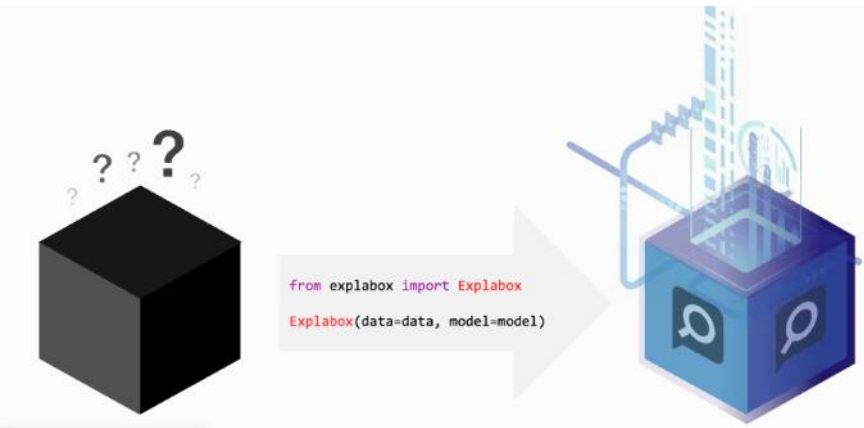
Not paid

(paid)

Fictitious example report 2

I wanted to buy champagne from John Doe via Ebay. Up to now, I have not received anything, and he does not respond to my e-mails, **so I already transferred the money.**

Explainable AI for legal decisions



<https://explabox.readthedocs.io/>

- Open-source libraries & toolkit for AI model inspection
 - Data statistics
 - XAI: rationales, counterfactuals, LIME/SHAP
 - Robustness: spelling mistakes, typo's
 - Biases: names, gender, etc.
- A holistic view on the AI system
 - What kind of data? How (good) does the system perform? Why does the system do what it does?

Rules, tools, and metrics

Brussels, 21.4.2021
COM(2021) 206 final

2021/0106 (COD)



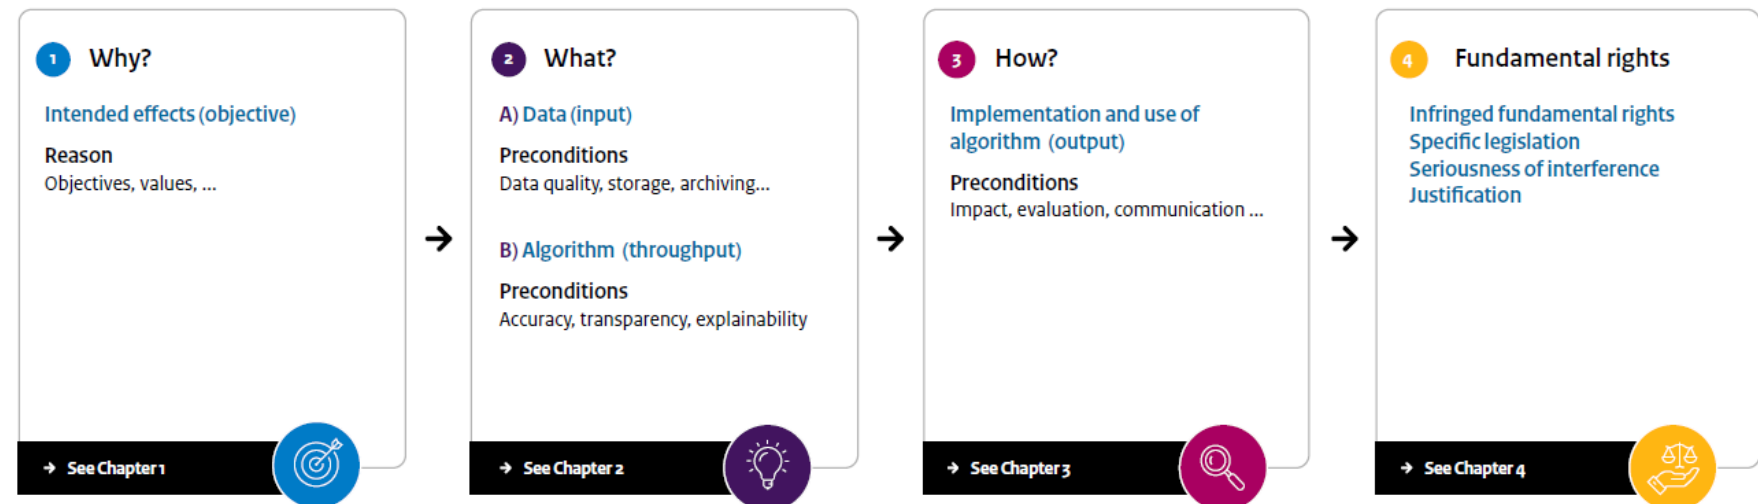
Proposal for a

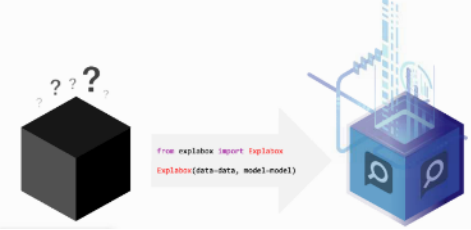
REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL

**LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE
(ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION
LEGISLATIVE ACTS**



Impact Assessment
Fundamental rights and algorithms





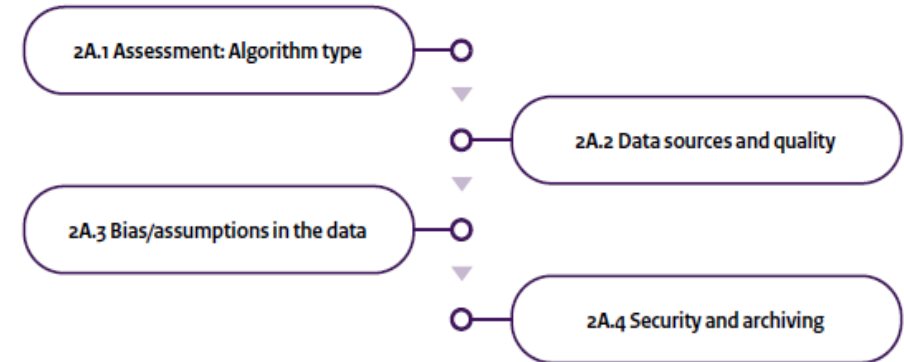
Explabox as assessment aid

- Use information from Explabox for assessment
 - What kind of data? How (good) does the system perform? Why does the system do what it does?

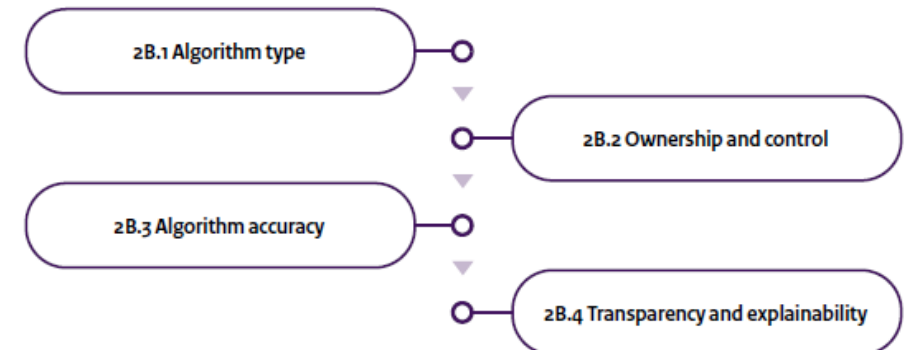


Part 2A: What? Data – input

This section covers the following topics:




Impact Assessment | Fundamental rights and algorithms



Rules, tools, and metrics

- Tools & metrics
 - What use are they? Intended and actual effects?
- New roles and responsibilities in organisations



Explainable AI for legal decisions

- Rules: Operationalising transparency and contestability in the law
 - Equality of arms

Proposal for a

REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL

**LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE
(ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION
LEGISLATIVE ACTS**

Explainable AI for legal decisions

- Rules: Operationalising transparency and contestability in the law
 - Equality of arms
 - Evaluating evidence and motivating decisions

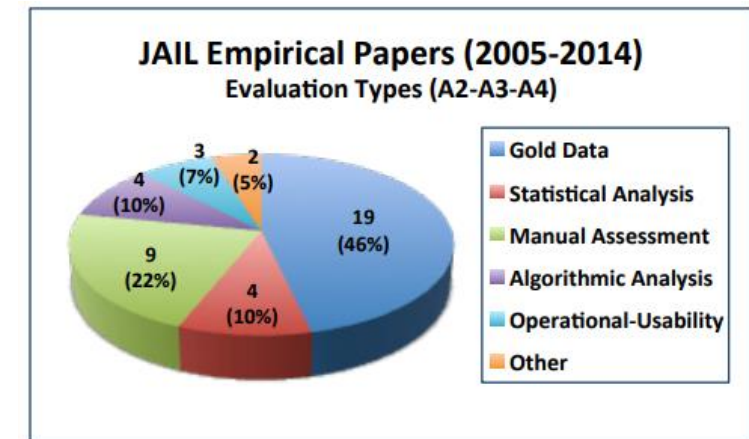
Proposal for a

REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL

LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE
(ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION
LEGISLATIVE ACTS

Evaluating AI in practice

- AI & Law Tech development & evaluation
 - Argumentation
 - Natural Language Processing
- How to evaluate?
 - Real systems in real user context



Atkinson, Collette, Bench-Capon, Dzehtsiarou (2021) Practical tools from formal models: the ECHR as a case study. *ICAIL 2021*.

Odekerken & Bex (2020) Towards transparent human-in-the-loop classification of fraudulent web shops. *JURIX 2020*

Van Binsbergen, Liu, Van Doesburg & Van Engers (2020) eFLINT: a domain-specific language for executable norm specifications. *ACM SIGPLAN Conference on Generative Programming*.

Conrad & Zeleznikow (2015) The role of evaluation in AI and law: an examination of its different forms in the AI and law journal. *ICAIL 2015*.

AI for detecting mobile phone usage while driving

- AI does initial filtering for pictures of cars/drivers who seem to be holding a phone, officer then checks the picture.

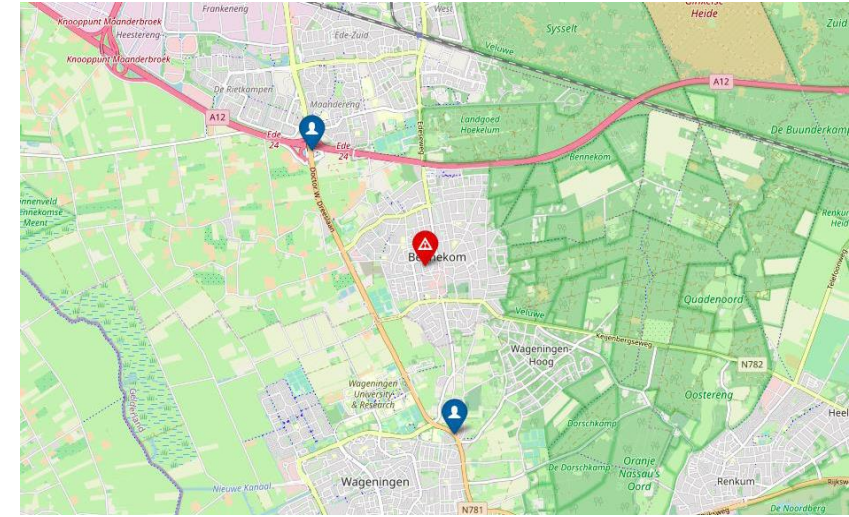


AI for detecting mobile phone usage while driving

- Best-practice in value-sensitive design
 - Data protection and anonymization
 - Training models with representative datasets
 - Develop and retain control in-house
- In practice:
 - New windscreen foil on cars
 - Officers share photographs with other officers to get second opinion
- Continuous training of both AI and human!



AI for police interception



- Notification of crime (e.g. robbery, smash & grab) and fleeing suspects
- Using knowledge about suspect behaviour, roads, etc., predict the suspect's route
- “Just like I thought”
 - Expert dispatchers only followed the recommendations of the system if they coincided with their own intuitions
 - Explanations hardly influence whether they trust/follow the recommendation

AI, Law and beyond

1. Combine knowledge & data
 - Trade fraud complaint intake
2. Evaluate how AI & Law is being used in practice
 - Citizen interaction & trust
 - XAI techniques in practice
 - Police officers using systems
3. Combine multiple disciplines
 - Public management: empirically investigating AI audits
 - Law: XAI, decision motivation and equality of arms

Combining knowledge and data

- Legal information extraction
- Non-statistical models to reason with data
- Combining machine learning and knowledge representation approaches
 - ML to extract information, KR to reason
 - Both ML and KR for one (complex) task
 - Solving KR problems with ML models
 - Constraining ML models using KR models

Argumentation Structure Prediction in CJEU Decisions on Fiscal State Aid – Santin et al.

Automatic Identification and Empirical Analysis of Legally Relevant Factors – Gray et al.

Argument Mining with Graph Representation Learning – Zhang et al.

Improving Translation of Case Descriptions into Logical Fact Formulas using LegalCaseNER – Zin et al.

Computable Contracts by Extracting Obligation Logic Graphs – Servantez et al.

Justification, stability and relevance for case-based reasoning with incomplete focus cases – Odekerken et al.

Model- and data-agnostic justifications with a fortiori case-based argumentation - Peters et al.

Analogical Reasoning, Generalization, and Rule Learning for Common Law Reasoning – Blass & Forbus

Combining a Legal Knowledge Model with Machine Learning for Reasoning with Legal Cases – Mumford et al.

Beyond Readability with RateMyPDF: A Combined Rule-based and Machine Learning Approach to Improving Court Forms – Steenhuis et al.

Evaluate AI & Law in practice

- Innovative applications
 - Evaluate operational-usability by “disinterested domain experts”
- Work together with stakeholders from practice
- Evaluate with proxy users
- Work with easily accessible user groups
 - Legal education
 - Academics

sustain.AI: a Recommender System to analyze Sustainability Reports – Hillebrand et al.

Image Analysis Approach to Trademark Congestion and Depletion – Haim & Kesari

“What’s wrong with this product?” - Detection of product safety issues based on information consumers share online – Fuchs et al.

A Methodology for Building Augmented Intelligence Tools for Laypeople to Increase Access to Justice – Westermann & Benyekhlef

Beyond Readability with RateMyPDF: A Combined Rule-based and Machine Learning Approach to Improving Court Forms – Steenhuis et al.

Combine multiple disciplines

- AI for Law
 - Apply AI for law, studying law with AI
- Law for AI
 - Legal-by-design
 - Legal aspects of AI for Law
- Include other disciplines
 - Broaden the AI & Law ecosystem

Using Agent-Based Simulations to Evaluate Bayesian Networks for Criminal Scenarios - van Leeuwen et al.

Do agents dream of abiding by the rules? Learning norms via behavioral exploration and sparse human supervision – Fratrič et al.

Uncovering Trauma in Genocide Tribunals: An NLP Approach Using the Genocide Transcript Corpus – Schirmer et al.

On predicting and explaining asylum adjudication – Katsikouli et al.

Gender Disparities in Child Custody Sentencing in Spain: a Data Driven Analysis – Riera et al.

The Perfect Victim: Computational Analysis of Judicial Attitudes towards Victims of Sexual Violence – Habba et al.

Computational Accountability - Hulstijn

Effects of XAI on Legal Process – Nielsen et al.

Rebuilding ‘ethics’ to govern ‘AI’: How to re-set the boundaries for the legal sector? - Unver

Araujo et al. (2020) In AI we trust? Perceptions about automated decision-making by artificial intelligence. *AI & Society*

Yalcin et al. (2023) Perceptions of Justice By Algorithms. *AI & Law Journal*

Barysé & Sarel (2023) Algorithms in the court: does it matter which part of the judicial decision-making is automated? *AI & Law Journal*

AI, law and beyond: building a transdisciplinary ecosystem

