AI, Law and Data

Floris Bex

Department of Information and Computing Sciences

Tilburg Institute for Law, Society and Technology

Universiteit Utrecht

Tilburg University

What is AI?

The AI in question, machine learning, is a technique for recognising patterns in relevant and preferably as complete as possible data files with the aim of discovering patterns in reality.

Minister of Justice to Parliament of the Netherlands

What is AI?

Systems that exhibit intelligent behaviour by analysing their environment and - with a certain degree of autonomy - taking action to achieve specific objectives.

European Commission
Coordinated strategy on AI

The possibilities of AI

- Expectations and hype exceeds reality
  - Big successes come from big companies (Google, Baidu)
  - AI is hard work!

- China is becoming world leader in AI
  - Computer vision, machine learning, medical AI

- But: AI for legal applications is different
  - Transparency, privacy, legal rules and regulations

    vs.

  - Statistical machine learning, Big Data & Deep Neural Networks

Police Lab AI

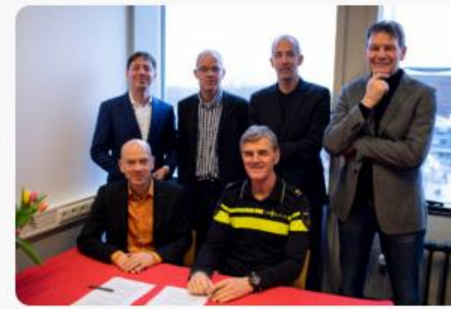

At the front of the developments in AI

AI in practice: handling citizen reports on cybercrime

- System can:
  - Read reports filed by citizens online
  - Monitor incoming reports
  - Build structured case files
  - Reason and ask questions based on reports

IA system architecture

- Different types of AI
  - Text classification (machine learning)
  - Reasoning (symbolic AI)
  - Search algorithms (symbolic AI)
  - Learning which actions to perform (reinforcement machine learning)

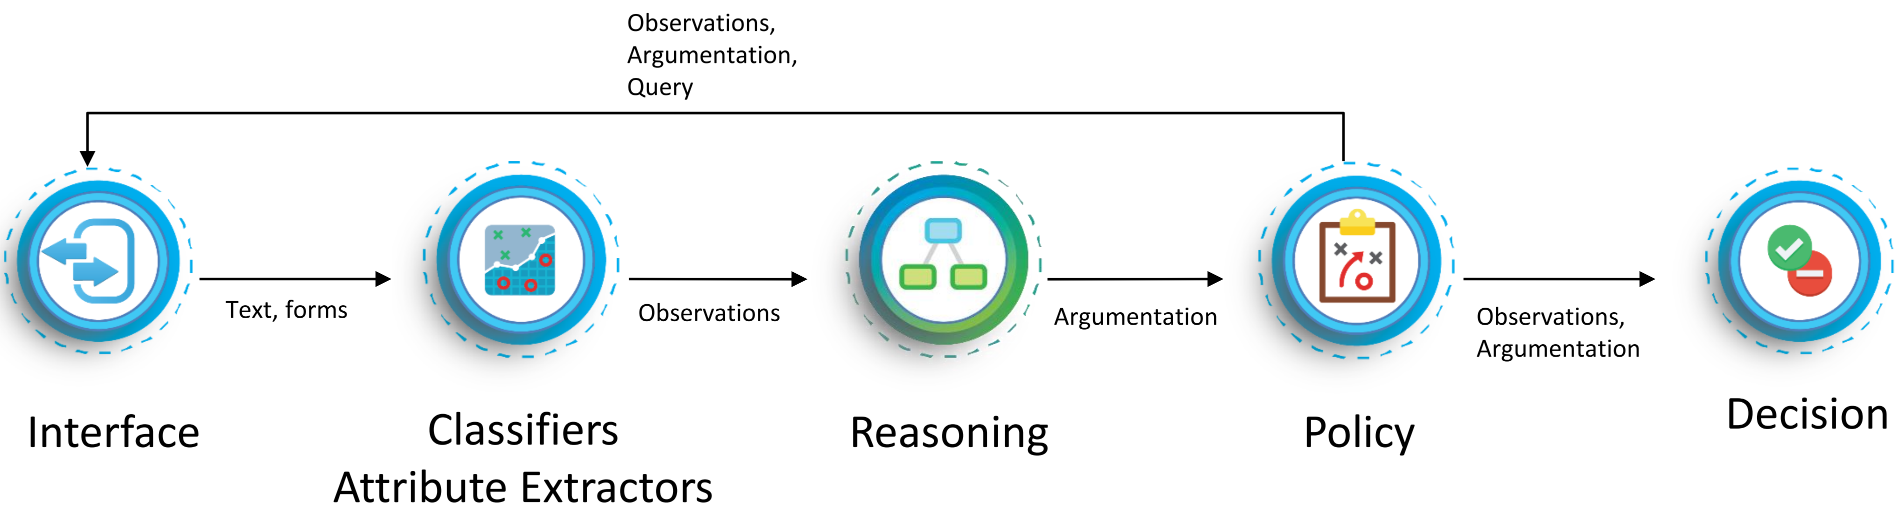

From text to observations

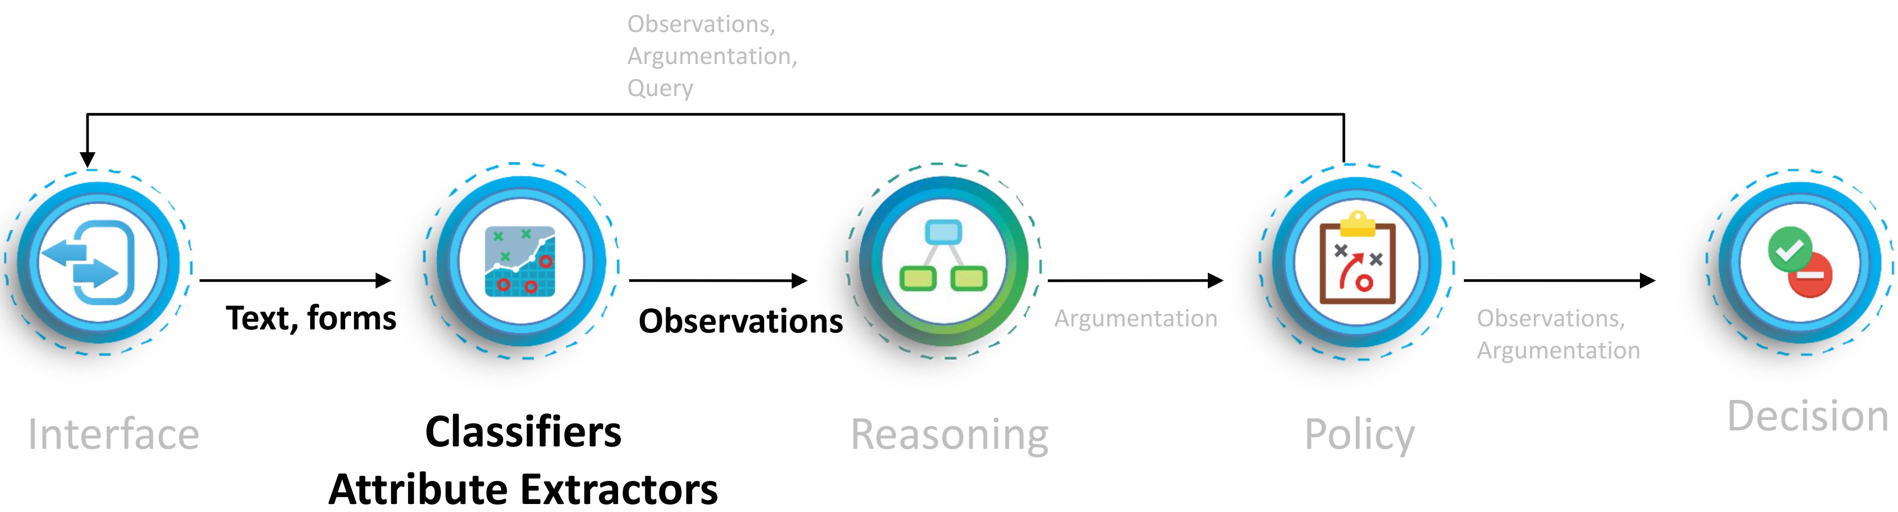

From Text to observations

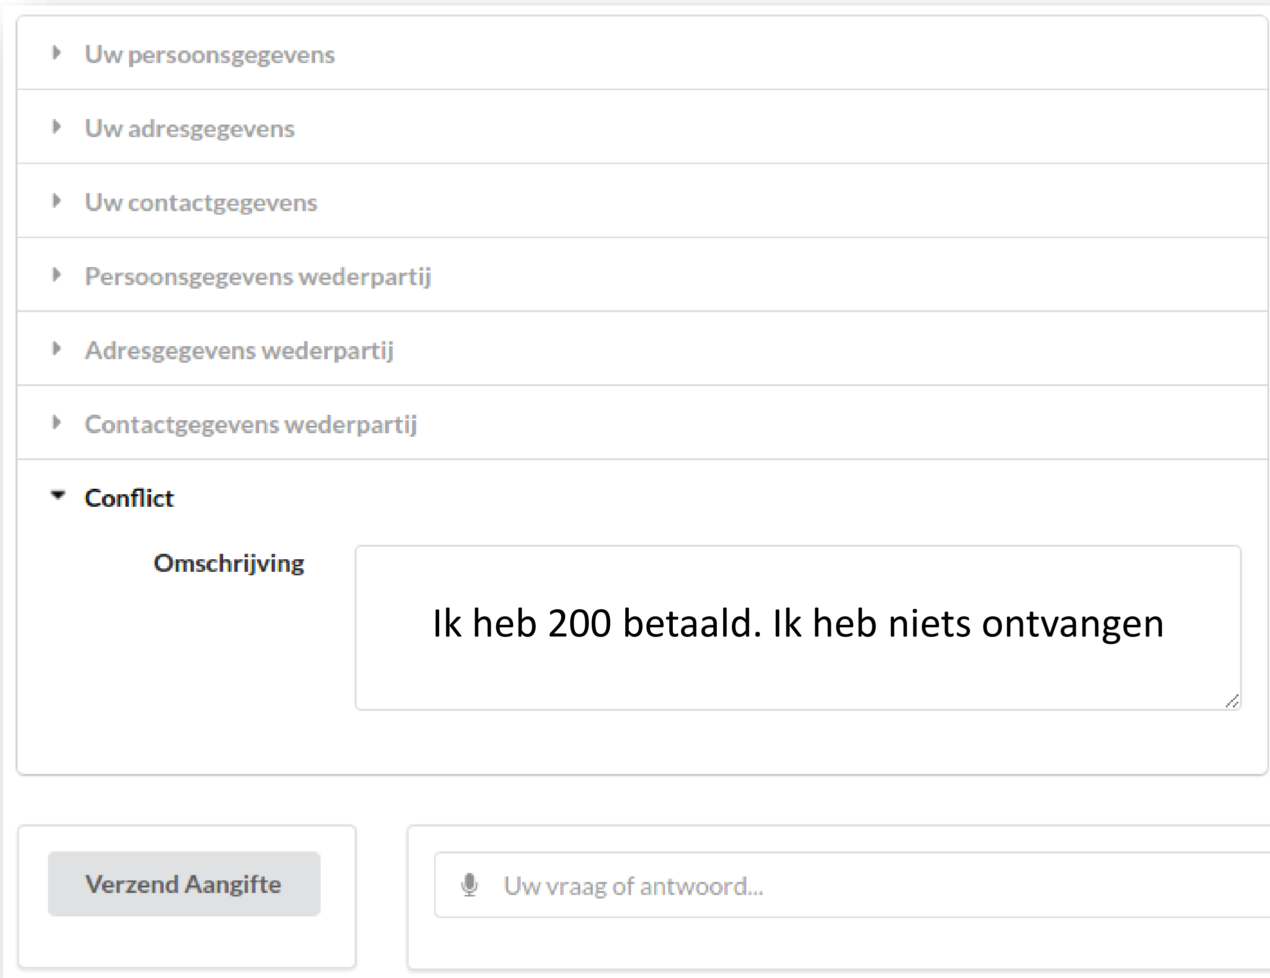

Interface

From Text to observations

I have paid 200.
I did not receive anything

"Pay" = yes AND "not" = no-> Paid
"Pay" = yes AND "not" = yes-> Not paid

<u>Observations in report</u>

Observation present?	**Yes**	**No**
Paid		
Not paid		
Received		
Not received		

Classifiers

From Text to observations

I have **paid** 200.
I did not receive anything

"Pay" = yes AND "not" = no-> Paid
"Pay" = yes AND "not" = yes-> Not paid

<u>Observations in report</u>

Observation present?	**Yes**	**No**
Paid	X	
Not paid		X
Received		
Not received		

Classifiers

From Text to observations

I have paid 200.
I did not receive anything

”Receive" = yes AND "not" = no-> Received
”Receive" = yes AND "not" = yes-> Not received

<u>Observations in report</u>

Observation present?	**Yes**	**No**
Paid	X	
Not paid		X
Received		
Not received		

Classifiers

From Text to observations

I have paid 200.
I did not receive anything

”Receive" = yes AND "not" = no-> Received
”Receive" = yes AND "not" = yes-> Not received

Observations in report

Observation present?	Yes	No
Paid	X	
Not paid		X
Received		X
Not received	X	

Classifiers

From Text to observations

- Classifications (rules) can be learnt
  - Supervised Learning: Give the AI enough examples so it learns to categorize phrases (can also be with "deep learning"!)
  - Tagging is done manually

From Text to observations

- Classifications (rules) can be learnt
  - Supervised Learning: Give the AI enough examples so it learns to categorize phrases (can also be with "deep learning"!)
  - Tagging is done manually

I paid 200	**Paid**
I have not paid	**Not paid**
I did not give them my money	**Not paid**
I transferred 100 euros	**Paid**
I gave him my money	**Paid**
I didn’t pay anything	**Not paid**
…	

From Text to observations

- After learning the AI can classify a new (unseen) sentence
  - AI has learned certain features of "Paid" and "Not paid" phrases

So I really didn't pay him anything
I have paid quite a lot of money
I didn't think about paying
I would pay him

From Text to observations

- After learning the AI can classify a new (unseen) sentence
  - AI has learned certain features of "Paid" and "Not paid" phrases

So I really **didn't pay him** anything	**Not paid**
I have **paid** quite a lot of money	**Paid**
I **didn't** think about **paying**	**Not paid**
I should **pay** him	**Paid**

  - Not always accurate!
  - Accuracy algorithm 80%-> 80% of the sentences is classified correctly as (Not) Paid
  - Confidence Classification 80%-> for a certain sentence, the algorithm is 80% sure that it is (Not) Paid

From Observations to arguments

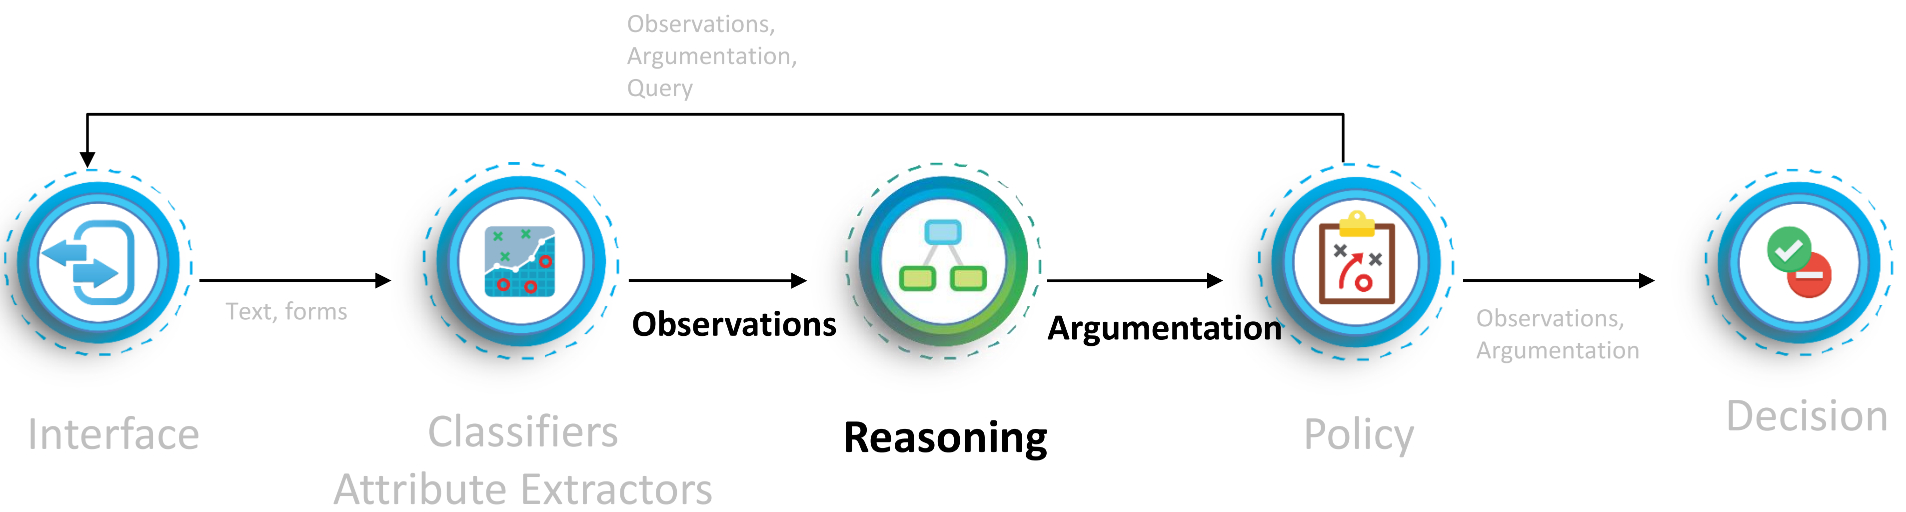

From Observations to arguments

- Arguments for/against possible fraud

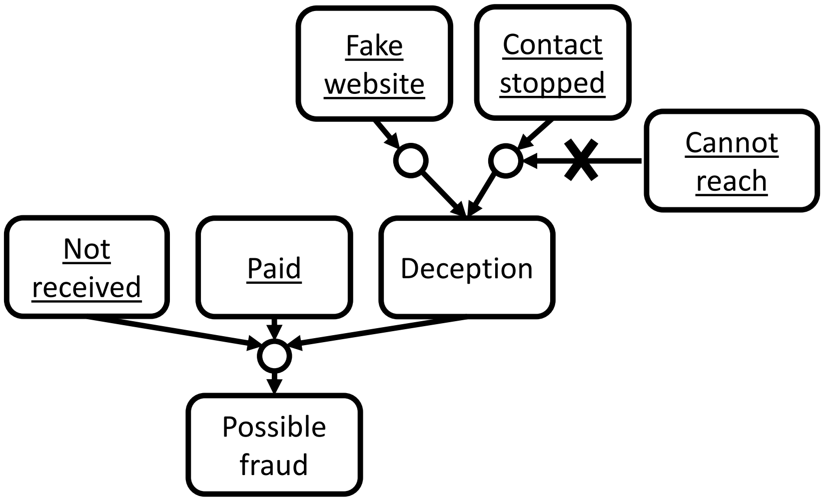

From Observations to arguments

- Arguments for/against possible fraud
  - If certain observations are present in the report...

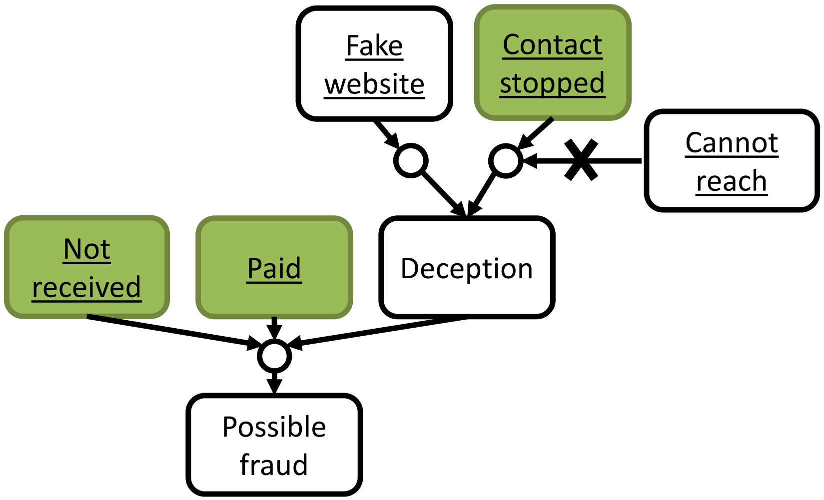

Reasoning

From Observations to arguments

- Arguments for/against possible fraud
  - …we can infer possible fraud

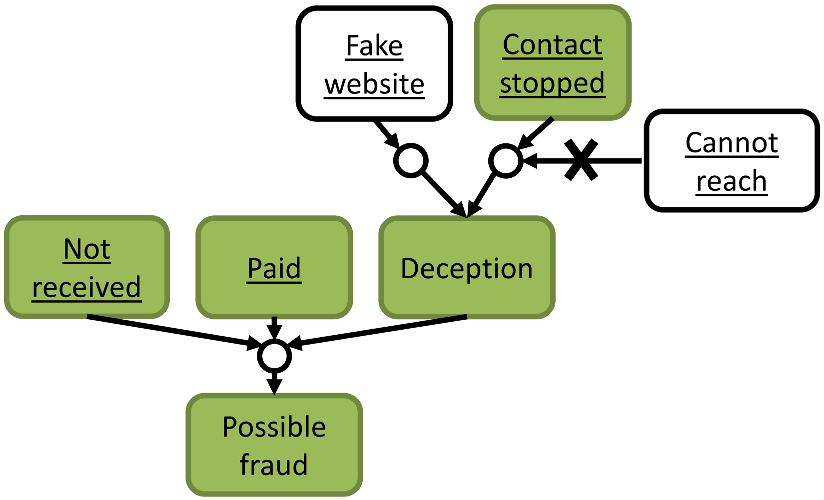

From Observations to arguments

- Arguments for/against possible fraud
  - Exceptions

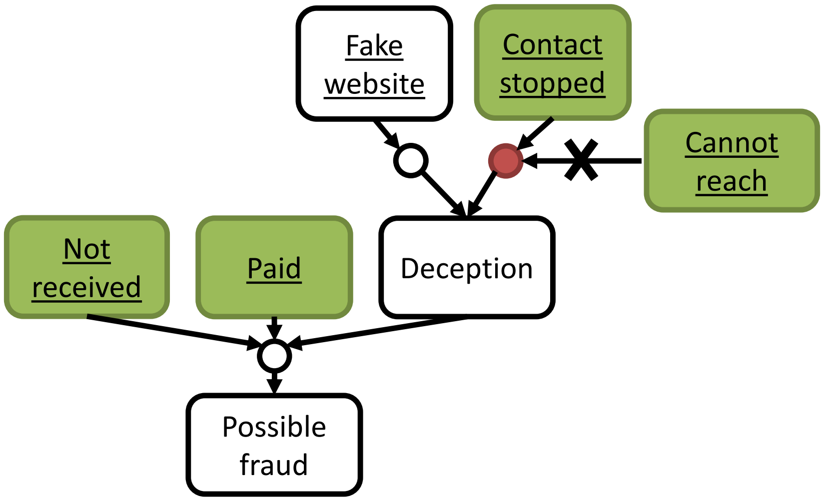

Reasoning

Van observaties naar argumenten

- Arguments are based on legislation, case law and expertise

- Explicit Knowledge has advantages
  - Transparency (for civilian, police, prosecution, judge)
  - Explicit Link Laws & Jurisprudence
  - Easier to adjust by police & Justice

From Observations to arguments

- Learning Arguments?
  - Label complete reports with fraud or non-fraud
  - Learning to classify new reports

Report 1; Name = Bart; Website = Alibaba; Conflict = "... I paid but didn't get anything... "

Possible fraud

Report 2; name=Floris; website=Alibaba; conflict="…Could get free iPhone have never received anything... "

Not Possible Fraud

Report 3; …
Report 4; …

- However...
  - Tagging is difficult (need experts)
  - Bad accuracy (65-70%)
  - Transparency disappears (more "black-box")

From arguments to Actions

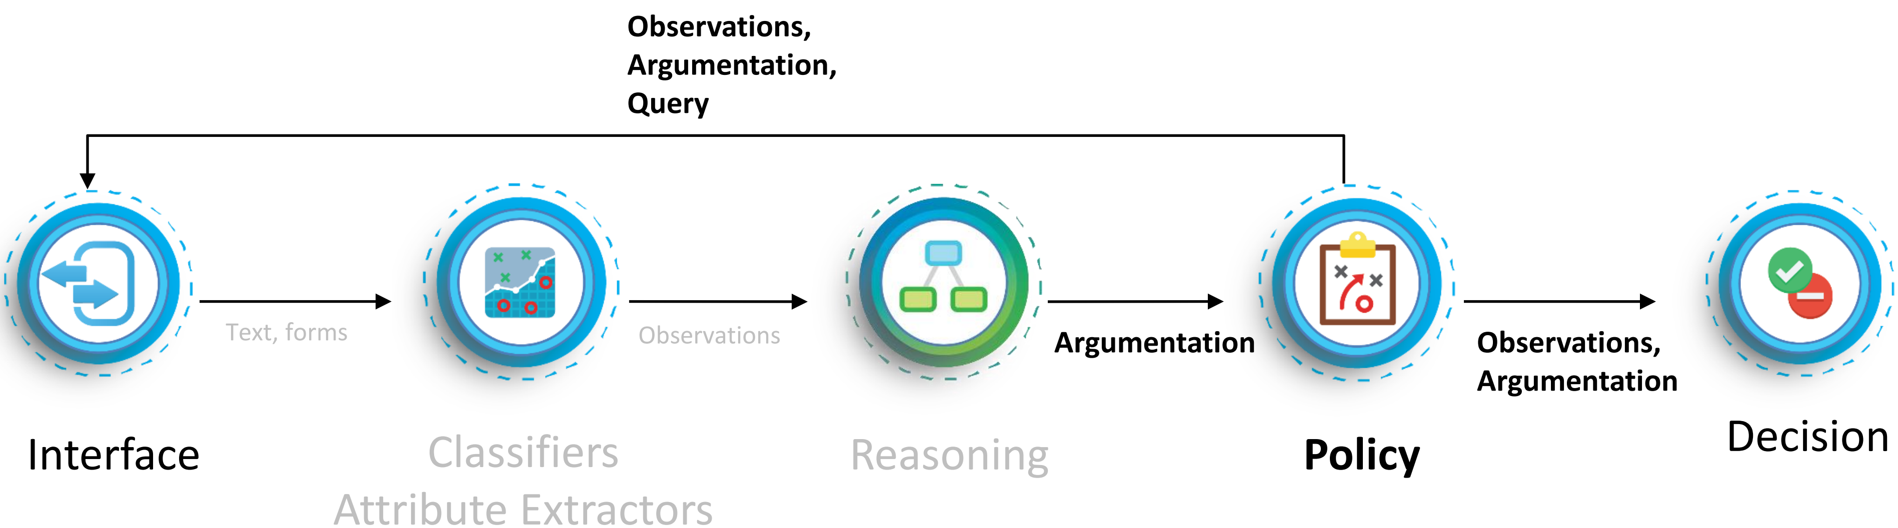

From arguments to actions

- Can you already conclude something? If not, what else should you ask for?

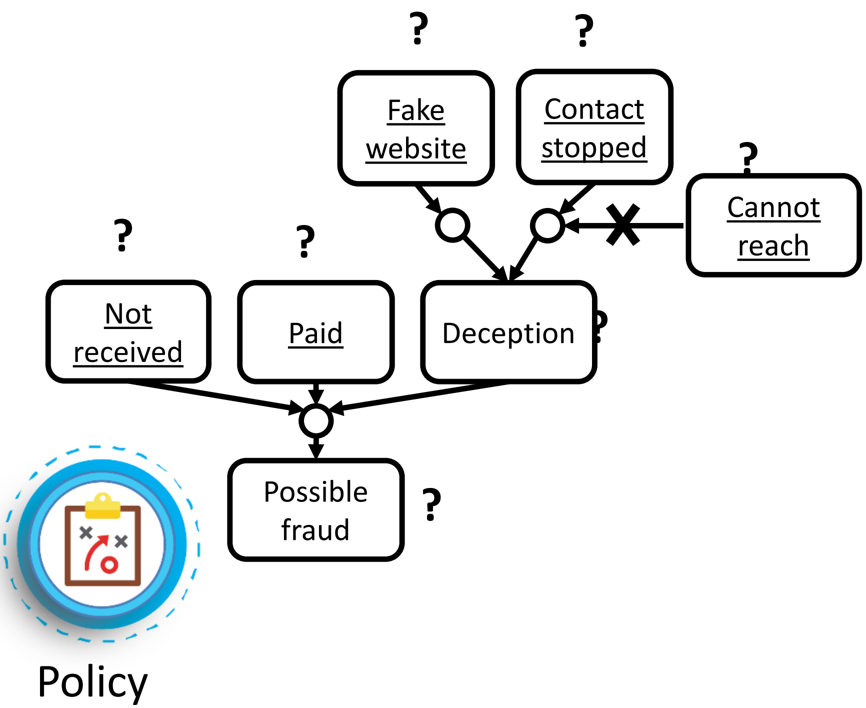

From arguments to actions

- Can you already conclude something? If not, what else should you ask for?
-

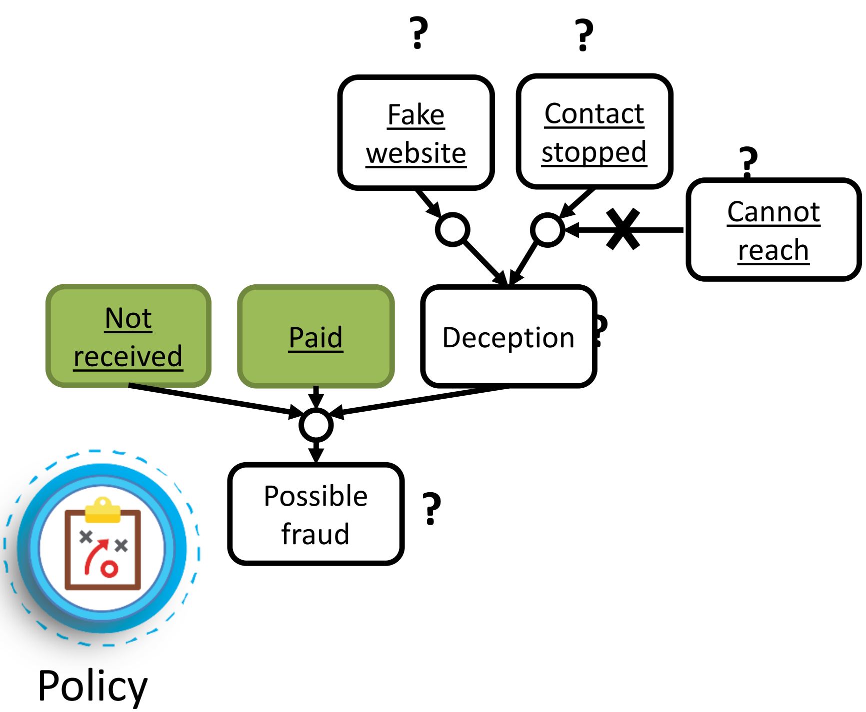

Observations in report

Observation present?	Yes	No
Paid	X	
Not paid		X
Received		X
Not received	X	

From arguments to actions

- Can you already conclude something? If not, what else should you ask for?
  - "Was there a fake website?"

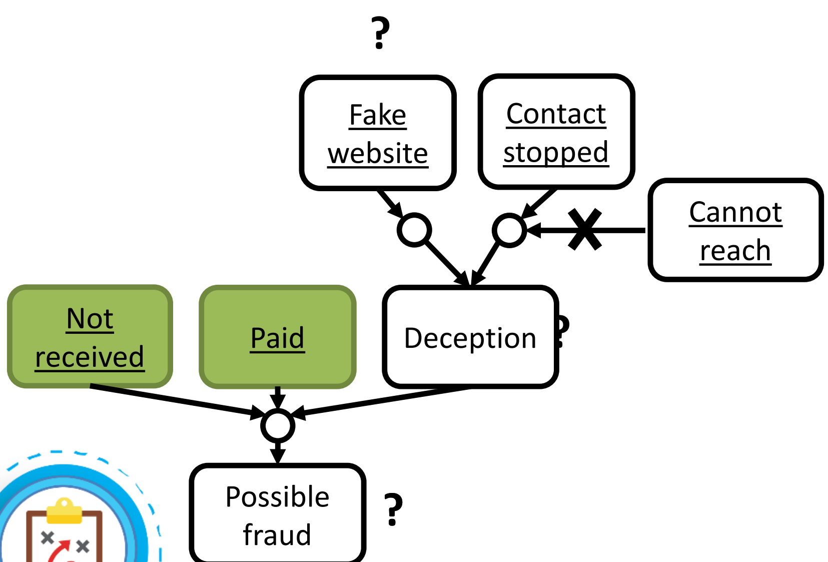

Policy

Observations in report

Observation present?	**Yes**	**No**
Paid	X	
Not paid		X
Received		X
Not received	X	

From arguments to actions

- Can you already conclude something? If not, what else should you ask for?
  - "Has the other party broken the contact?”
    - "Were you sufficiently available?"

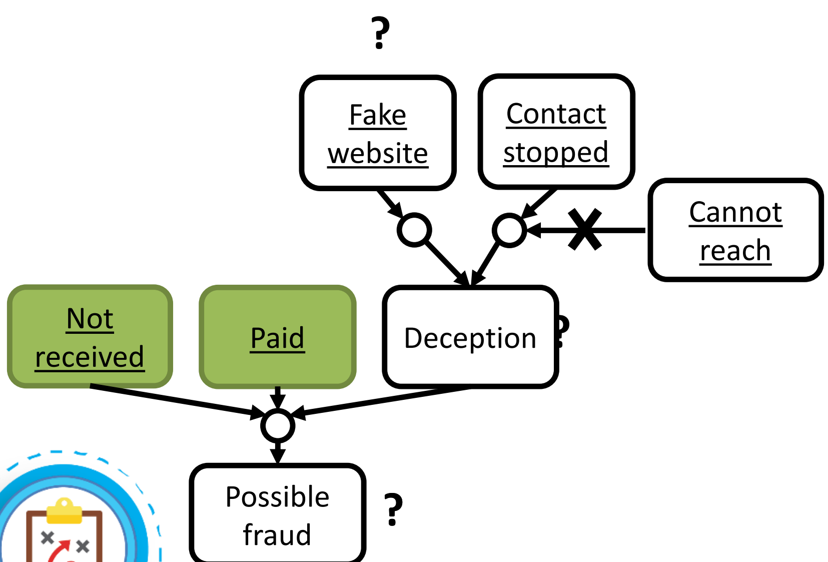

Policy

Observations in report

Observation present?	Yes	No
Paid	X	
Not paid		X
Received		X
Not received	X	

From arguments to actions

- Can you already conclude something? If yes, give a decision.
  - "You have paid and not received a product. The other party used a fake website. Thank you for your report, we will contact you a.s.a.p.. "

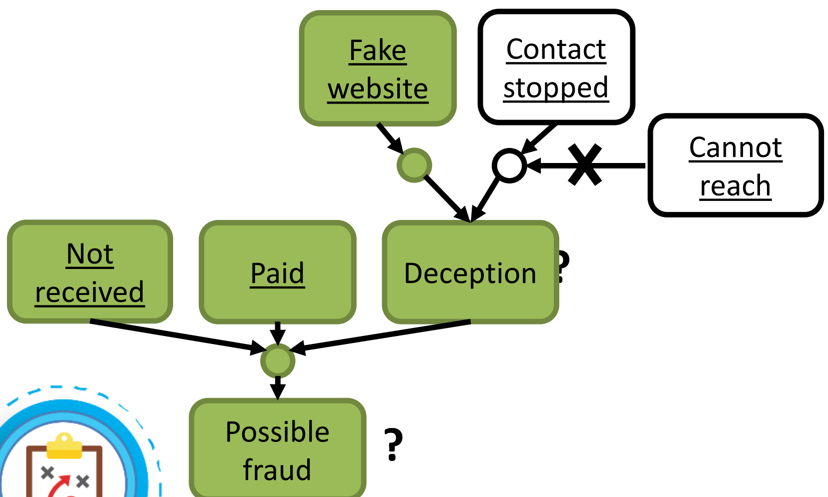

Policy

From arguments to actions

- Can you already conclude something? If yes, give a decision.
  - "You did not receive a product. The other party used a fake website. However, you have not paid, so it is not fraud. "

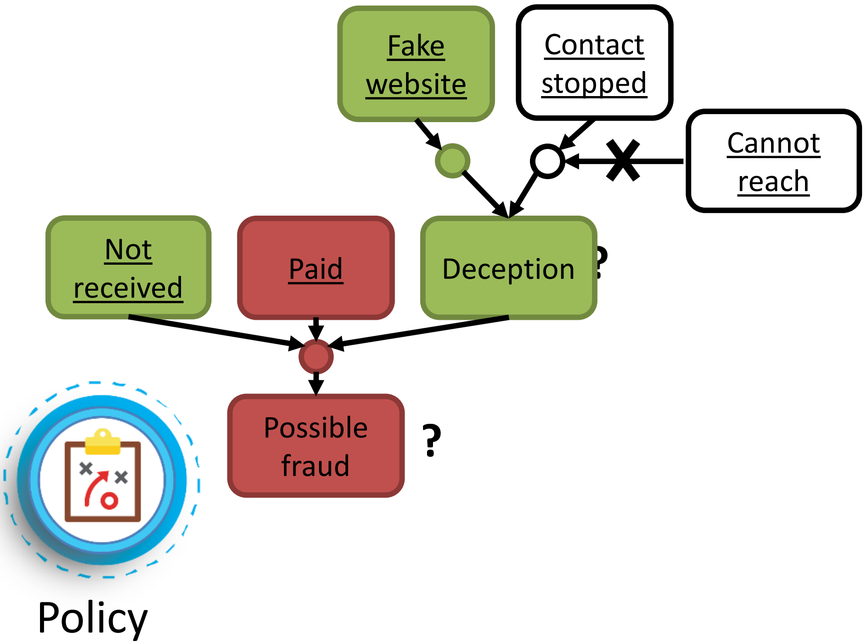

From arguments to actions

- Efficient search algorithm to determine the best question
  - If you know nothing, what should you ask first?

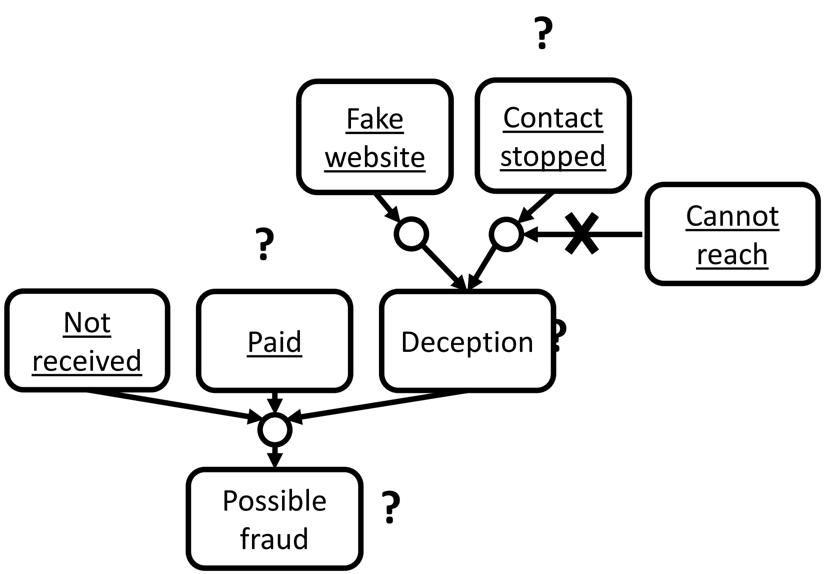

From arguments to actions

- Efficient search algorithm to determine the best question
  - If you know nothing you can better first ask "Paid?" instead of "Contact broken?" – Paid is always needed to infer the conclusion!

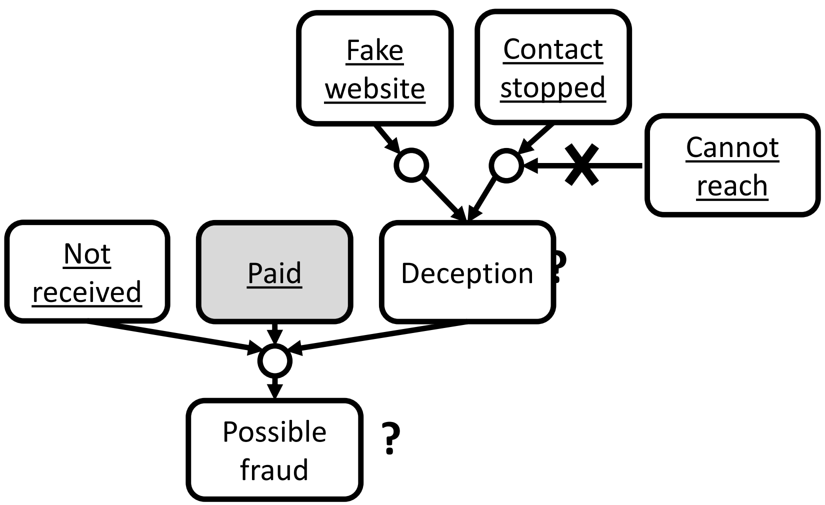

From arguments to actions

- Efficient search algorithm to determine the best question
  - But: you do not know in advance how citizens (users) will reply

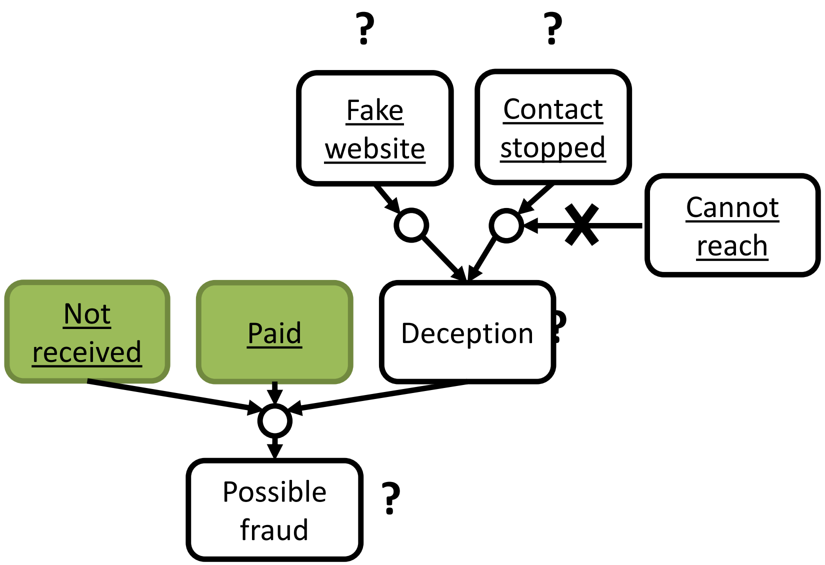

From arguments to actions

- Efficient search algorithm to determine the best question
  - Reinforcement Learning: Let the AI perform dialogues with real humans, "reward" if conclusion reached, "punish" if additional question is asked or dialogue is stopped

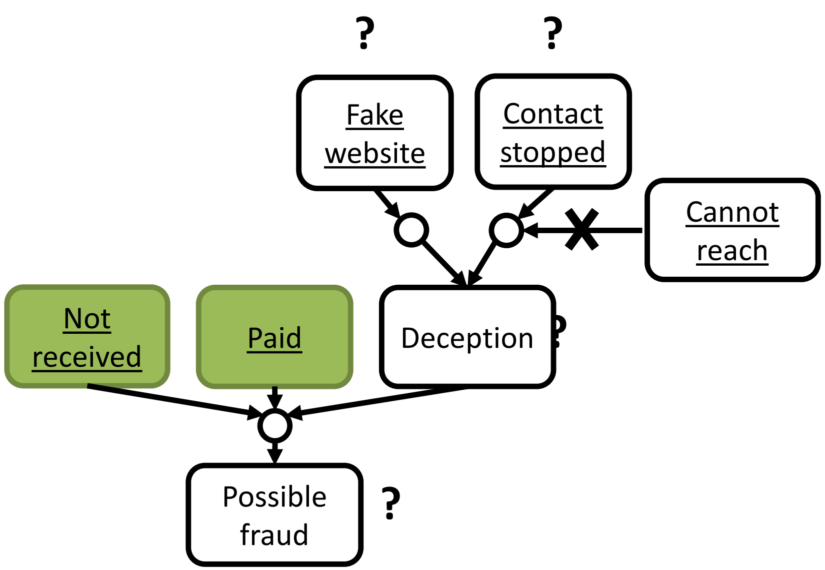

IA system architecture

- Requirements for the AI
  - Accurate: Minimize Mistakes
  - Transparency: Explanation of important decisions
  - Control: Can detect where errors are, keep improving
  - Efficient: Minimize unnecessary actions

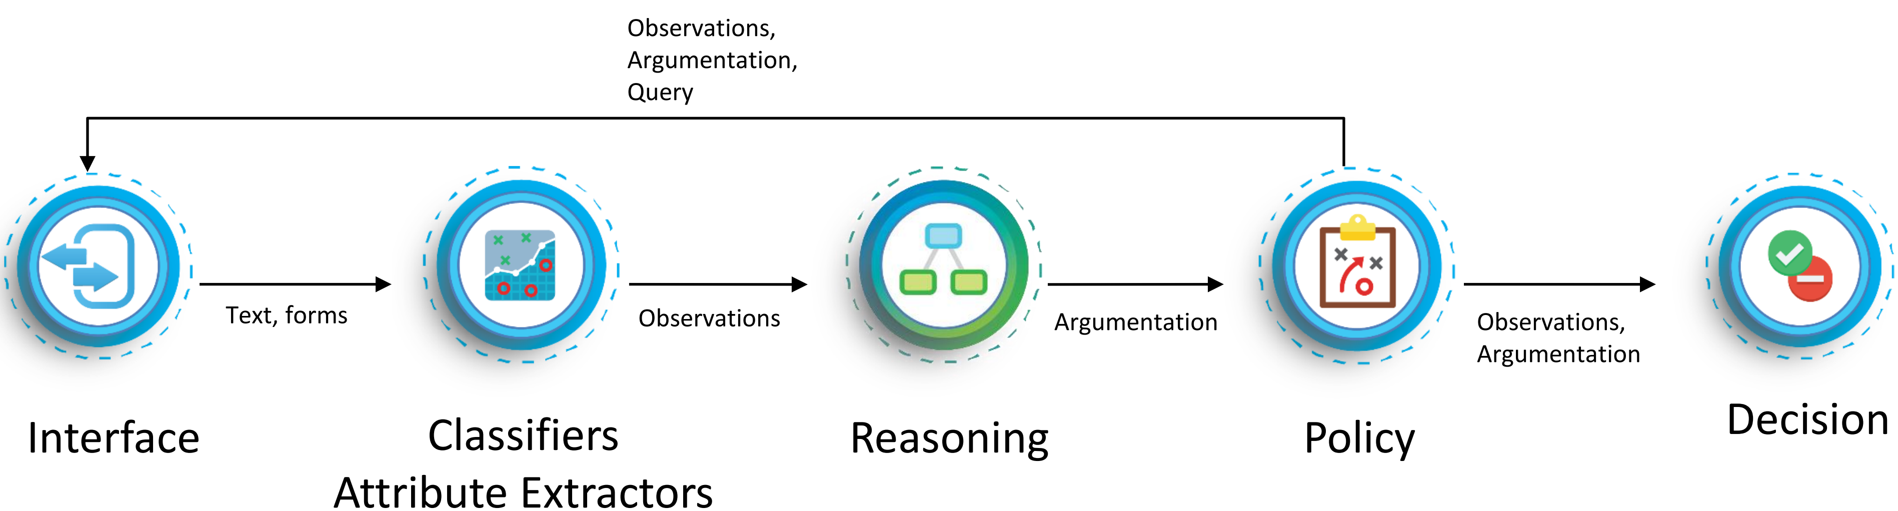

“Deep IA”?

- Supervised learning
  - Input: text of report, text of question or decision
  - A lot of data needed
  - Declaration text + question + decision
  - Black box
  - Unclear why a particular decision is taken

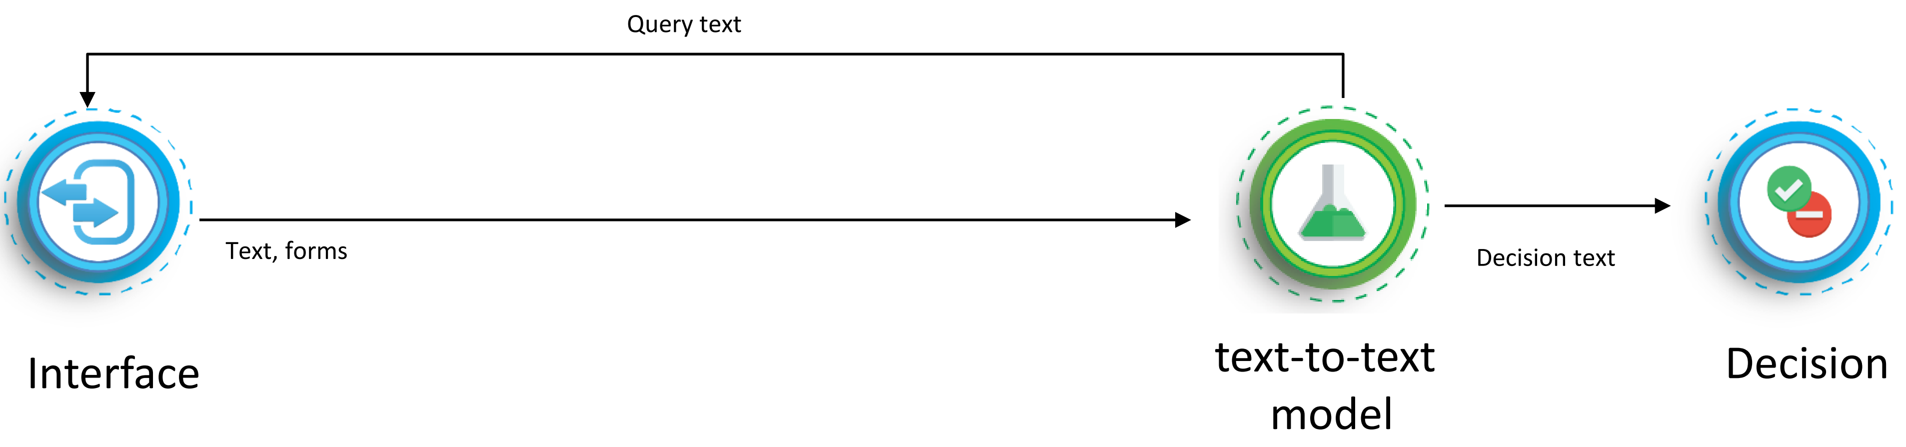

Police Lab AI

- Dialogues & chatbots
  - Citizen reports, Interpol reports & questions
- Explainable AI
  - Explains offender profiling to judges
- Crime scripting
  - Analyse and predict crime
- Networks and simulation
  - Simulate networks of terror cells and drug rings – what happens if you remove a person?
- Multimodal summaries
  - Summarize video, tekst, etc.
- Sensing
  - Information from cameras and sensors

Data science & AI for the legal field

- Smart search 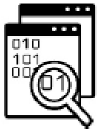
  - Information retrieval, decision support
  - Machine learning, symbolic knowledge
- (Predictive) legal analysis 
  - Jurimetrics, public administration, sociology
  - Statistics, machine learning
- Decision support 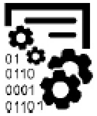
  - Decision support, expertsystemen, “robotrechter”
  - Statistiek, machine learning, symbolische kennis (bijv. regels)

Data science & AI for the legal field

- Smart search
  - Information retrieval, decision support
  - Machine learning, symbolic knowledge

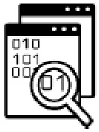

- (Predictive) legal analysis
  - Jurimetrics, public administration, sociology
  - Statistics, machine learning

- Decision support
  - Decision support, expertsystemen, “robotrechter”
  - Statistiek, machine learning, symbolische kennis (bijv. regels)

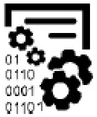

Simple search

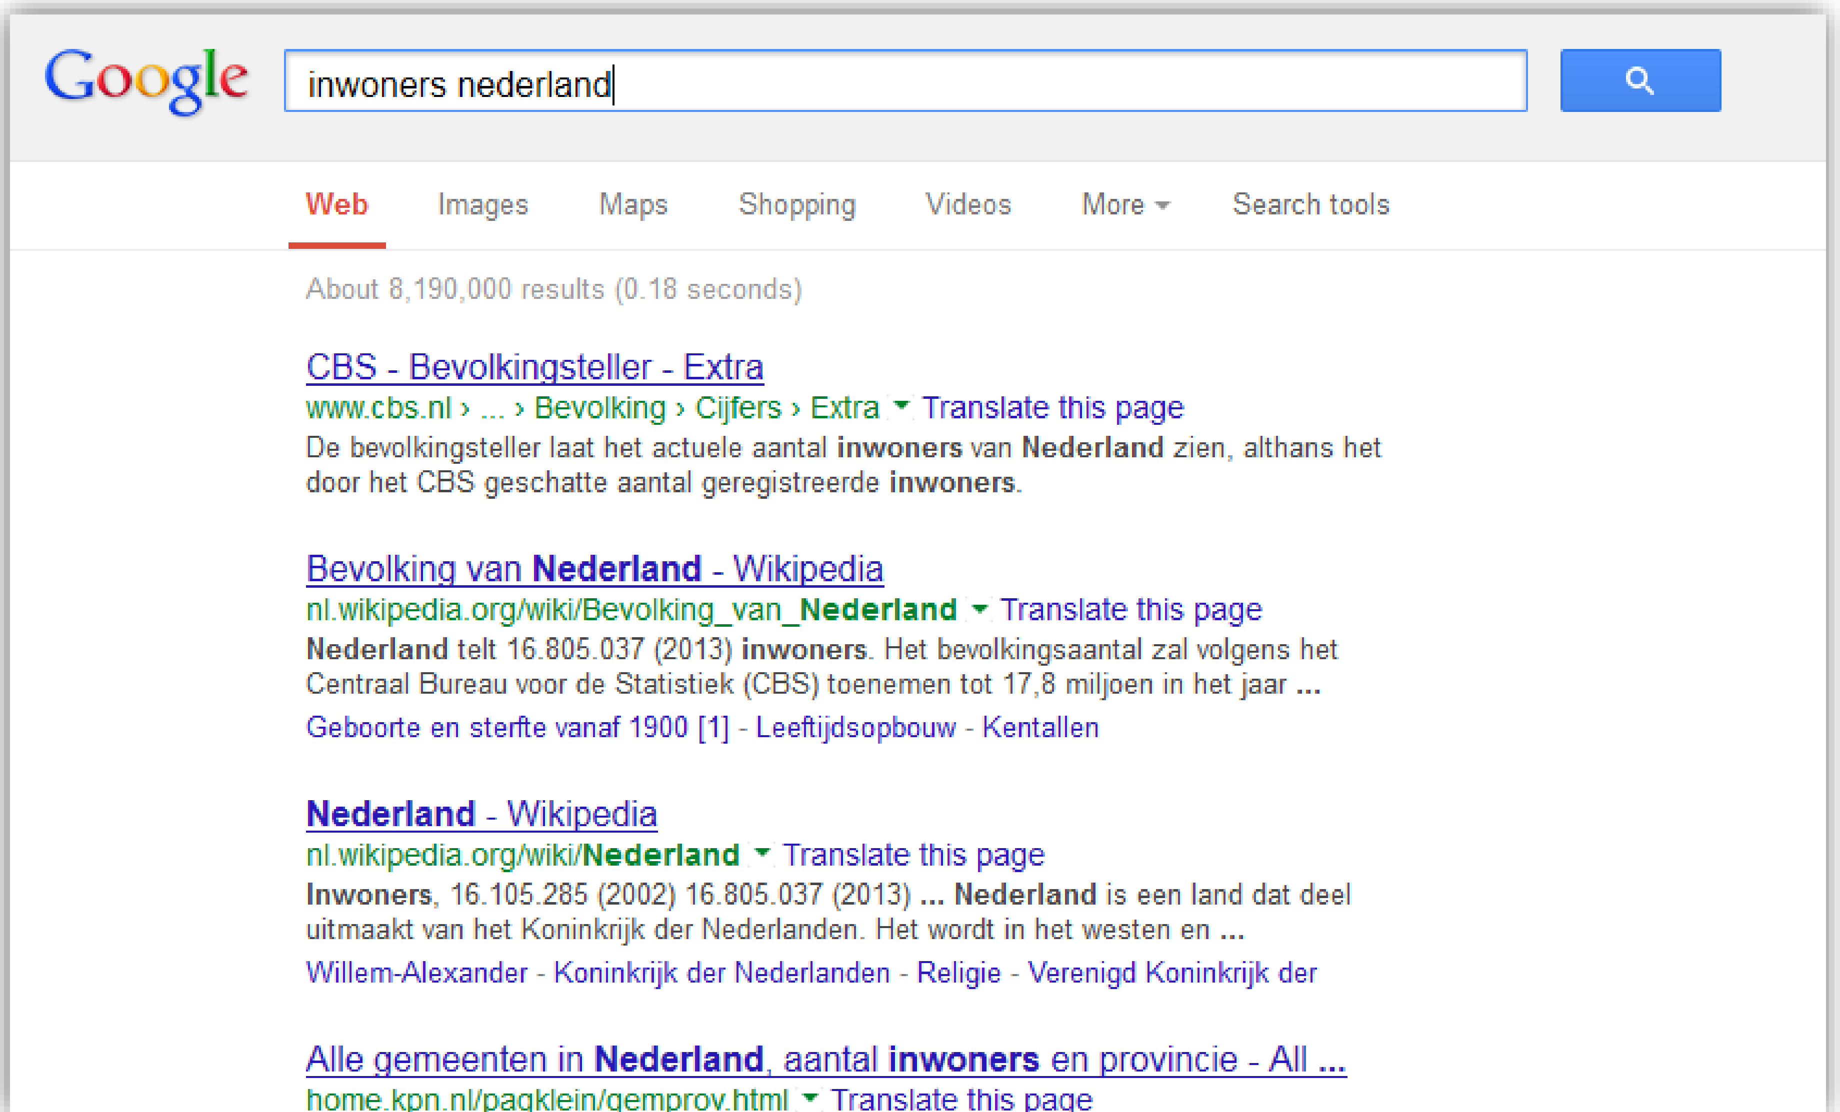

Smart (semantic) search

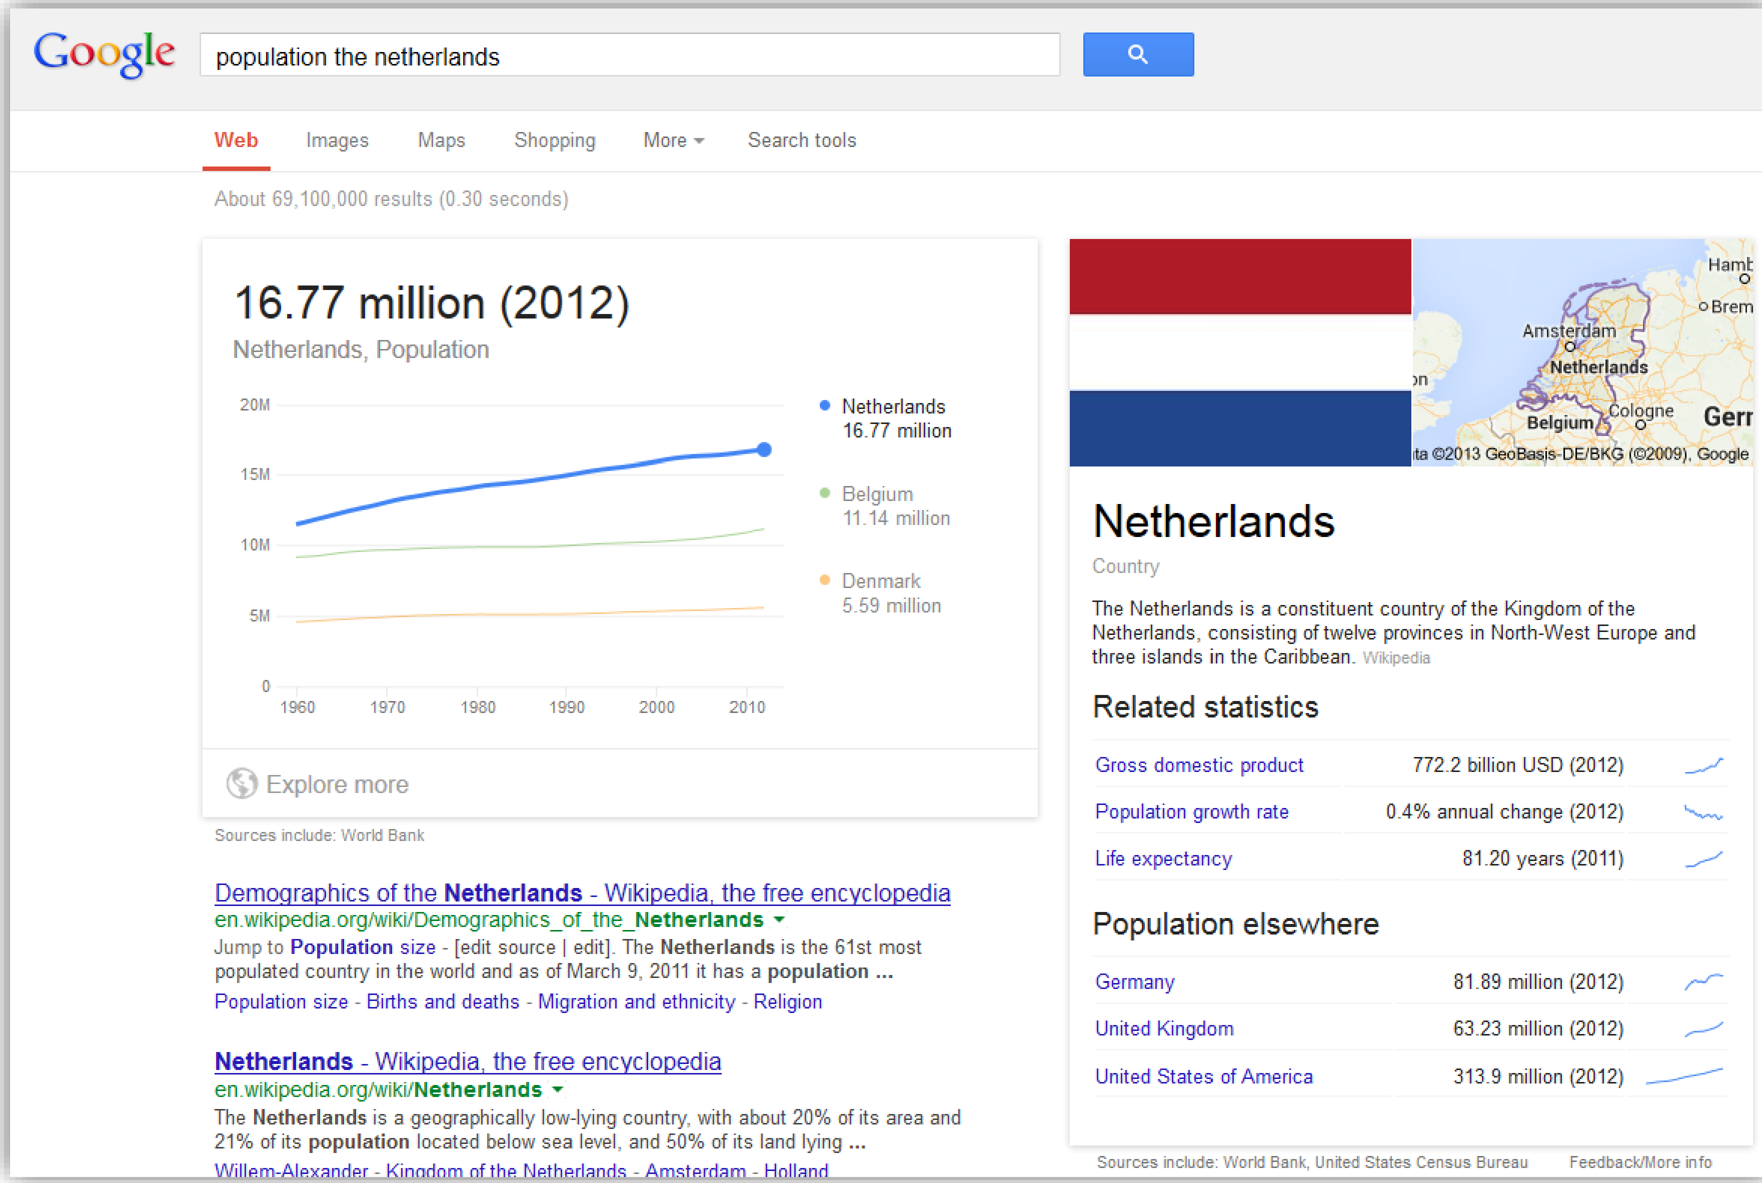

Smart search for the judiciary

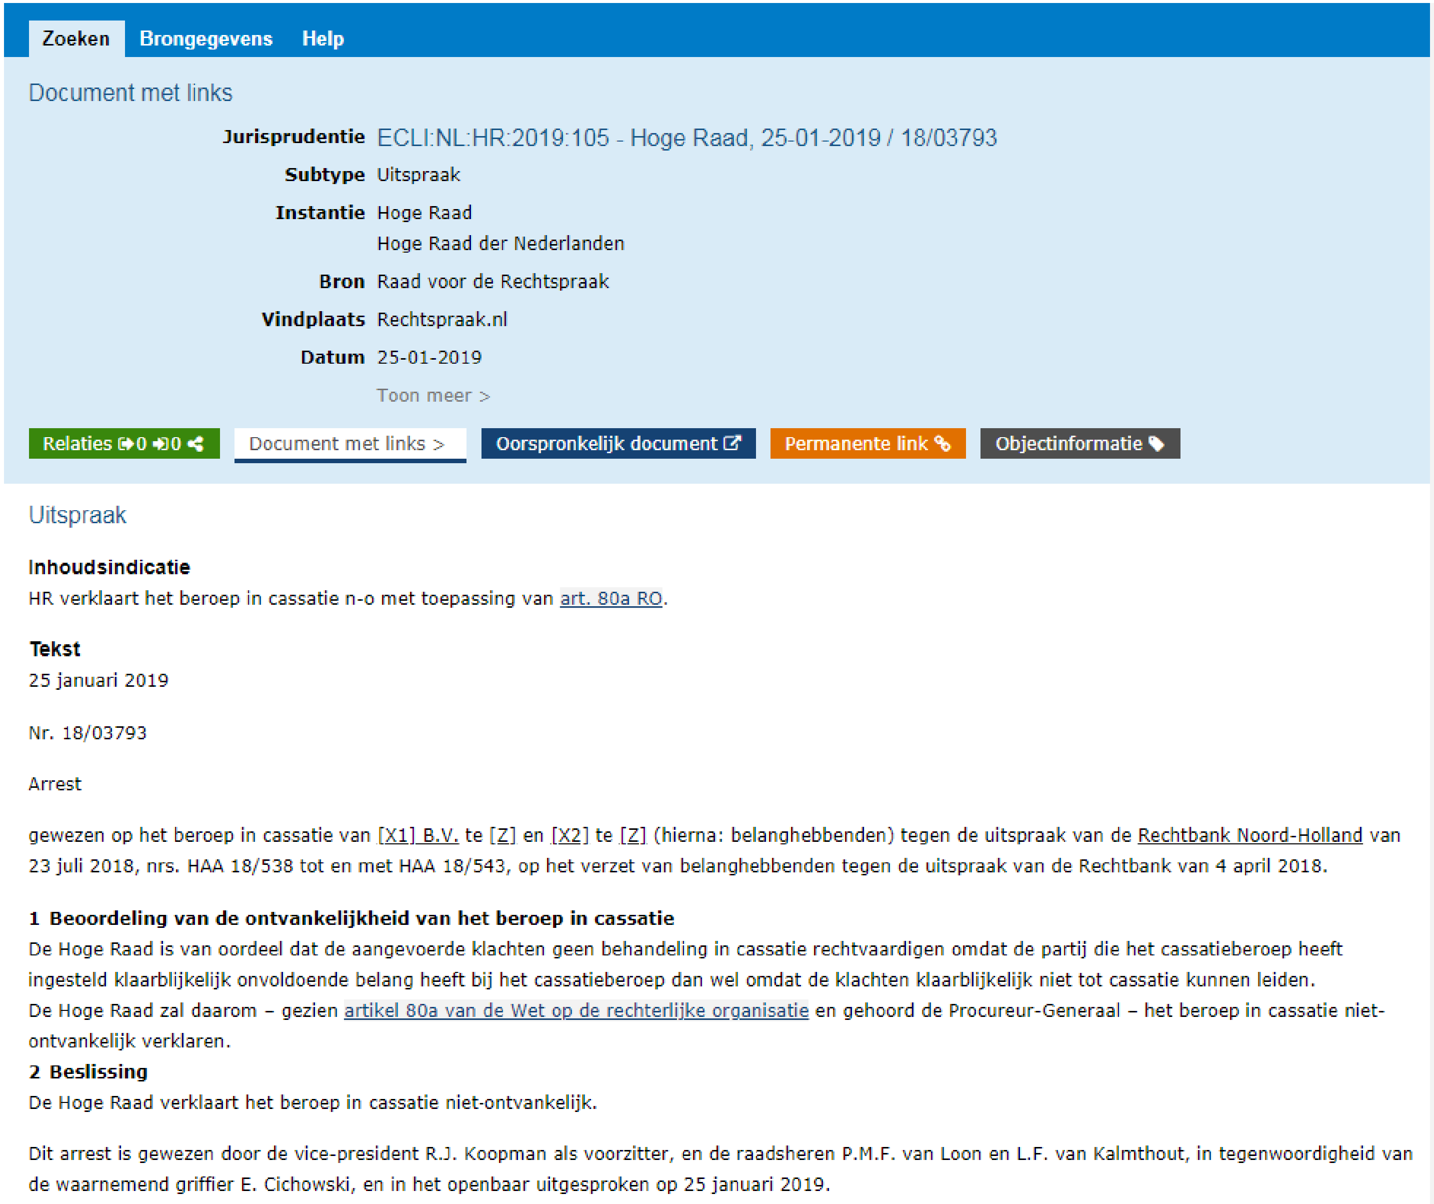

Smart search

- Needs structured data (Semantic Web)
- Knowledge acquisition bottleneck
  - What about Wikipedia? Huge knowledge engineering effort!
- Legal ontologies, linked data for the law

Data science & AI for the legal field

- Smart search
  - Information retrieval, decision support
  - Machine learning, symbolic knowledge

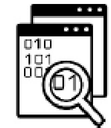

- (Predictive) legal analysis
  - Jurimetrics, public administration, sociology
  - Statistics, machine learning

- Decision support
  - Decision support, expertsystemen, “robotrechter”
  - Statistiek, machine learning, symbolische kennis (bijv. regels)

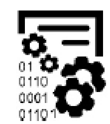

Legal analysis

- The costs of going to trial for judge X are as follows:

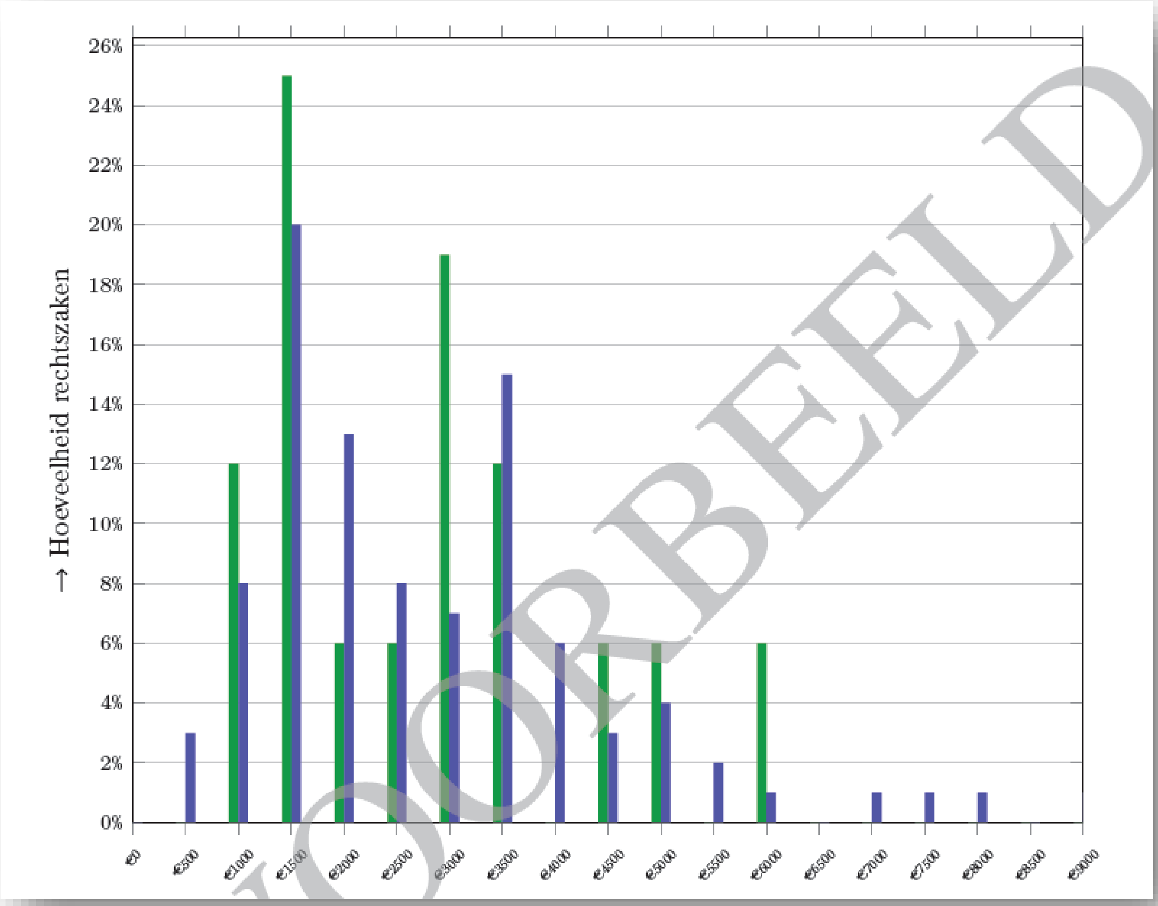

- Costs, probability of sentencing, etc.
- Allows for smart lawyering

Legal analysis

- Analysis of “metadata”
  - Number of cases, time taken, costs, …
- Analysis of case contents
  - Which arguments are given by the parties? Which laws are called on?
  - Argument & topic mining

Predictive legal analysis

RESEARCH ARTICLE

A general approach for predicting the behavior of the Supreme Court of the United States

Daniel Martin Katz[1,2]*, Michael J. Bommarito II[1,2], Josh Blackman[3]

1 Illinois Tech - Chicago-Kent College of Law, Chicago, IL, United States of America, **2** CodeX - The Stanford Center for Legal Informatics, Stanford, CA, United States of America, **3** South Texas College of Law Houston, Houston, TX, United States of America

* dkatz3@kentlaw.iit.edu

Abstract

Building on developments in machine learning and prior work in the science of judicial prediction, we construct a model designed to predict the behavior of the Supreme Court of the United States in a generalized, out-of-sample context. To do so, we develop a time-evolving random forest classifier that leverages unique feature engineering to predict more than 240,000 justice votes and 28,000 cases outcomes over nearly two centuries (1816-2015). Using only data available prior to decision, our model outperforms null (baseline) models at both the justice and case level under both parametric and non-parametric tests. Over nearly two centuries, we achieve 70.2% accuracy at the case outcome level and 71.9% at the justice vote level. More recently, over the past century, we outperform an *in-sample optimized* null model by nearly 5%. Our performance is consistent with, and improves on the general level of prediction demonstrated by prior work; however, our model is distinctive because it can be applied out-of-sample to the entire past and future of the Court, not a single term. Our results represent an important advance for the science of quantitative legal prediction and portend a range of other potential applications.

Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective

Nikolaos Aletras[1,2], Dimitrios Tsarapatsanis[3], Daniel Preoţiuc-Pietro[4,5] and Vasileios Lampos[2]

[1] Amazon.com, Cambridge, United Kingdom
[2] Department of Computer Science, University College London, University of London, London, United Kingdom
[3] School of Law, University of Sheffield, Sheffield, United Kingdom
[4] Positive Psychology Center, University of Pennsylvania, Philadelphia, United States
[5] Computer & Information Science, University of Pennsylvania, Philadelphia, United States

ABSTRACT

Recent advances in Natural Language Processing and Machine Learning provide us with the tools to build predictive models that can be used to unveil patterns driving judicial decisions. This can be useful, for both lawyers and judges, as an assisting tool to rapidly identify cases and extract patterns which lead to certain decisions. This paper presents the first systematic study on predicting the outcome of cases tried by the European Court of Human Rights based solely on textual content. We formulate a binary classification task where the input of our classifiers is the textual content extracted from a case and the target output is the actual judgment as to whether there has been a violation of an article of the convention of human rights. Textual information is represented using contiguous word sequences, i.e., N-grams, and topics. Our models can predict the court's decisions with a strong accuracy (79% on average). Our empirical analysis indicates that the formal facts of a case are the most important predictive factor. This is consistent with the theory of legal realism suggesting that judicial decision-making is significantly affected by the stimulus of the facts. We also observe that the topical content of a case is another important feature in this classification task and explore this relationship further by conducting a qualitative analysis.

Predictive legal analysis

RESEARCH ARTICLE

A general approach for predicting the behavior of the Supreme Court of the United States

Daniel Martin Katz[1,2]*, Michael J. Bommarito II[1,2], Josh Blackman[3]

1 Illinois Tech - Chicago-Kent College of Law, Chicago, IL, United States of America, **2** CodeX - The Stanford Center for Legal Informatics, Stanford, CA, United States of America, **3** South Texas College of Law Houston, Houston, TX, United States of America

* dkatz3@kentlaw.iit.edu

Abstract

Building on developments in machine learning and prior work in the science of judicial prediction, we construct a model designed to predict the behavior of the Supreme Court of the United States in a generalized, out-of-sample context. To do so, we develop a time-evolving random forest classifier that leverages unique feature engineering to predict more than 240,000 justice votes and 28,000 cases outcomes over nearly two centuries (1816-2015). Using only data available prior to decision, our model outperforms null (baseline) models at both the justice and case level under both parametric and non-parametric tests. Over nearly two centuries, we achieve 70.2% accuracy at the case outcome level and 71.9% at the justice vote level. More recently, over the past century, we outperform an *in-sample optimized* null model by nearly 5%. Our performance is consistent with, and improves on the general level of prediction demonstrated by prior work; however, our model is distinctive because it can be applied out-of-sample to the entire past and future of the Court, not a single term. Our results represent an important advance for the science of quantitative legal prediction and portend a range of other potential applications.

- Given features of the judges, predict whether they will rule for or against the party

- 70% accurate
  - Smart guess: 67%

Predictive legal analysis

- Given (text) parts of statements + pronunciation (label), classify unseen cases
  - 79% accurate
  - "Violation" predict is 84% accurate!

Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective

Nikolaos Aletras[1,2], Dimitrios Tsarapatsanis[3], Daniel Preoţiuc-Pietro[4,5] and Vasileios Lampos[2]

[1] Amazon.com, Cambridge, United Kingdom
[2] Department of Computer Science, University College London, University of London, London, United Kingdom
[3] School of Law, University of Sheffield, Sheffield, United Kingdom
[4] Positive Psychology Center, University of Pennsylvania, Philadelphia, United States
[5] Computer & Information Science, University of Pennsylvania, Philadelphia, United States

ABSTRACT

Recent advances in Natural Language Processing and Machine Learning provide us with the tools to build predictive models that can be used to unveil patterns driving judicial decisions. This can be useful, for both lawyers and judges, as an assisting tool to rapidly identify cases and extract patterns which lead to certain decisions. This paper presents the first systematic study on predicting the outcome of cases tried by the European Court of Human Rights based solely on textual content. We formulate a binary classification task where the input of our classifiers is the textual content extracted from a case and the target output is the actual judgment as to whether there has been a violation of an article of the convention of human rights. Textual information is represented using contiguous word sequences, i.e., N-grams, and topics. Our models can predict the court's decisions with a strong accuracy (79% on average). Our empirical analysis indicates that the formal facts of a case are the most important predictive factor. This is consistent with the theory of legal realism suggesting that judicial decision-making is significantly affected by the stimulus of the facts. We also observe that the topical content of a case is another important feature in this classification task and explore this relationship further by conducting a qualitative analysis.

Predictive legal analysis

- Given the text of the case (evidence + charge) predict youth or adult punishment
- 72% accurate
  - Smart guess: 70%
- More useful: what are the important factors for the decision?
  - Age of perpetrator, type of crime

Prediction of Adolescent Law Case Outcomes Using Unstructured and Structured Data

Nazli Ander
TU/e Student Number: 1236648

Accuracy of Classification Models

- In classification problems, the primary source for accuracy estimation is the ***confusion matrix***

		True Class	
		Positive	Negative
Predicted Class	Positive	True Positive Count (TP)	False Positive Count (FP)
	Negative	False Negative Count (FN)	True Negative Count (TN)

There are 100 positives and 100 negatives
Algorithm classifies 120 as positive, of which 90 are correct

TP = 90, FP = 30
FN = 10, TN = 70

Accuracy of Classification Models

- Recall: how many of the actual (true) positives were found by the algorithm?

		True Class	
		Positive	Negative
Predicted Class	Positive	True Positive Count (TP)	False Positive Count (FP)
	Negative	False Negative Count (FN)	True Negative Count (TN)

TPR/Recall

$$Recall = \frac{TP}{TP + FN}$$

There are 100 positives and 100 negatives
Algorithm classifies 120 as positive, of which 90 are correct

TP = 90, FP = 30
FN = 10, TN = 70

Recall = 90/100 = 90%

Accuracy of Classification Models

- Precision: of the actual (true) positives found, how many are correct?

		True Class		
		Positive	Negative	
Predicted Class	Positive	True Positive Count (TP)	False Positive Count (FP)	Precision
	Negative	False Negative Count (FN)	True Negative Count (TN)	

$$precision = \frac{TP}{TP + FP}$$

There are 100 positives and 100 negatives
Algorithm classifies 120 as positive, of which 90 are correct

TP = 90, FP = 30
FN = 10, TN = 70

Precision = 90/120 = 75%

Accuracy of Classification Models

- Recall vs precision

$$Recall = \frac{TP}{TP + FN} \qquad precision = \frac{TP}{TP + FP}$$

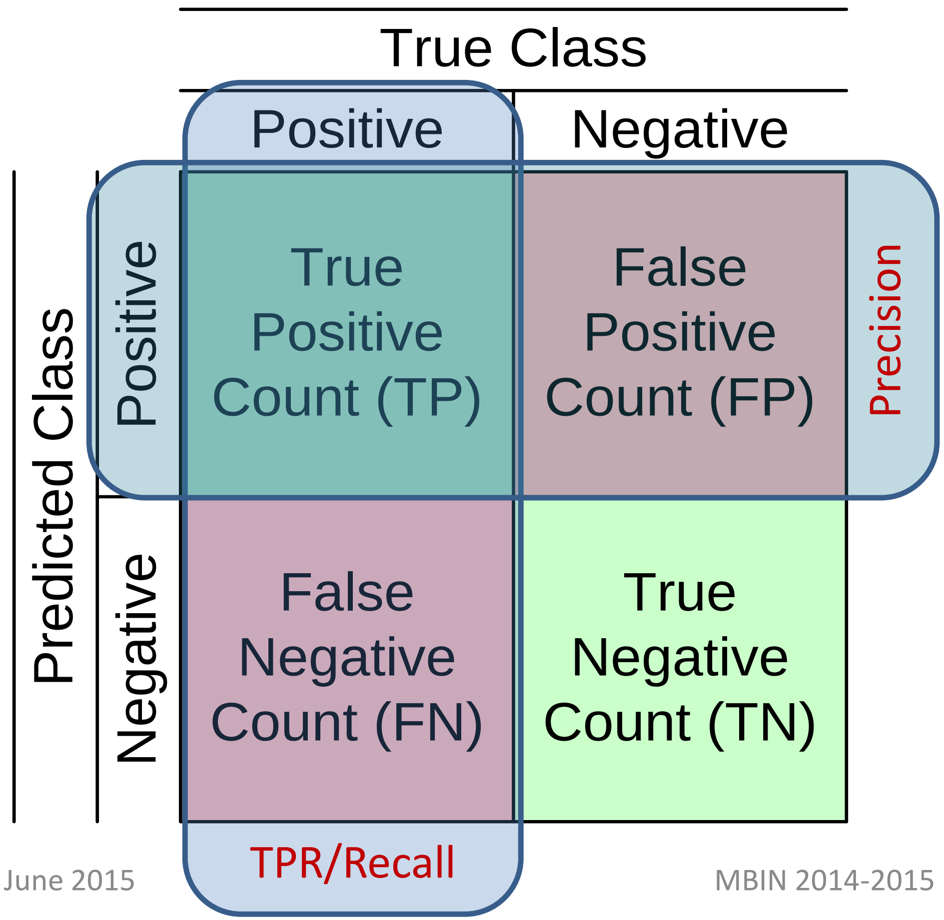

Which one is more important?

high precision: algorithm returned substantially more relevant results than irrelevant ones (but maybe not many)
high recall: algorithm returned most of the relevant results (but maybe also many irrelevant ones_

Accuracy of Classification Models

- Accuracy: how many predictions are actually (true) positives or negatives?

		True Class	
		Positive	Negative
Predicted Class	Positive	True Positive Count (TP)	False Positive Count (FP)
	Negative	False Negative Count (FN)	True Negative Count (TN)

Accuracy

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

There are 100 positives and 100 negatives
Algorithm classifies 120 as positive, of which 90 are correct

TP = 90, FP = 30
FN = 10, TN = 70

Accuracy = 160/200 = 80

Predictive legal analysis

- What does “prediction” really mean?

- 90% of criminal cases that end up in court result in “guilty” decision
  - Many innocents will not even be prosecuted
- Say we have 100 random cases, what is the accuracy if we predict “guilty”?
  - 90%

Predictive legal analysis

- What does “prediction” really mean?
- 90% of criminal cases that end up in court result in “guilty” decision
  - Many innocents will not even be prosecuted
- Say we have 100 random cases, what is the accuracy if we predict “guilty”?
  - 90%
  - Very high accuracy for “guilty”, but we will never find the ”innocent” cases!

Data science & AI for the legal field

- Smart search
  - Information retrieval, decision support
  - Machine learning, symbolic knowledge

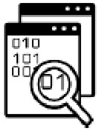

- (Predictive) legal analysis
  - Jurimetrics, public administration, sociology
  - Statistics, machine learning

- Decision support
  - Decision support, expert systems, “robojudge”
  - Statistics, machine learning, symbolic knowledge (e.g. rules)

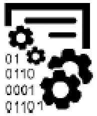

Traffic fine appeals

- Input: citizen appeal against a traffic fine
- Output:
  - Similar cases
  - Questions and advice for citizen
  - Draft decision

appeals

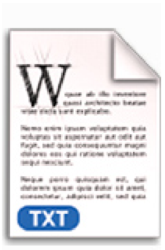

decision

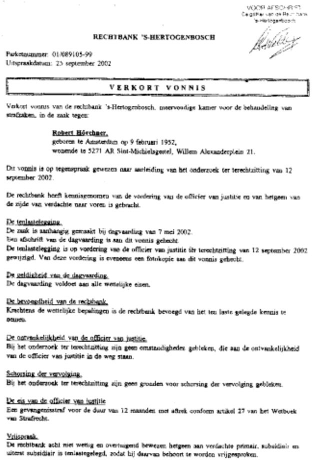

RECHTBANK 'S-HERTOGENBOSCH

Parketnummer: 01/089105-99
Uitspraakdatum: 23 september 2002

VERKORT VONNIS

Verkort vonnis van de rechtbank 's-Hertogenbosch, meervoudige kamer voor de behandeling van strafzaken, in de zaak tegen:

Robert Hövehner,
geboren te Amsterdam op 9 februari 1952,
wonende te 5271 AR Sint-Michielsgestel, Willem Alexanderplein 21.

Dit vonnis is op tegenspraak gewezen naar aanleiding van het onderzoek ter terechtzitting van 12 september 2002.

De rechtbank heeft kennisgenomen van de vordering van de officier van justitie en van hetgeen van de zijde van verdachte naar voren is gebracht.

De tenlastelegging.
De zaak is aanhangig gemaakt bij dagvaarding van 7 mei 2002.
Een afschrift van de dagvaarding is aan dit vonnis gehecht.
De tenlastelegging is op vordering van de officier van justitie ter terechtzitting van 12 september 2002 gewijzigd. Van deze vordering is eveneens een fotokopie aan dit vonnis gehecht.

De geldigheid van de dagvaarding.
De dagvaarding voldoet aan alle wettelijke eisen.

De bevoegdheid van de rechtbank.
Krachtens de wettelijke bepalingen is de rechtbank bevoegd van het ten laste gelegde kennis te nemen.

De ontvankelijkheid van de officier van justitie.
Bij het onderzoek ter terechtzitting zijn geen omstandigheden gebleken, die aan de ontvankelijkheid van de officier van justitie in de weg staan.

Schorsing der vervolging.
Bij het onderzoek ter terechtzitting zijn geen gronden voor schorsing der vervolging gebleken.

De eis van de officier van justitie
Een gevangenisstraf voor de duur van 12 maanden met aftrek conform artikel 27 van het Wetboek van Strafrecht.

Vrijspraak.
De rechtbank acht niet wettig en overtuigend bewezen hetgeen aan verdachte primair, subsidiair en uiterst subsidiair is tenlastegelegd, zodat hij daarvan behoort te worden vrijgesproken.

AI for law and police

- Current AI “boom” focuses on supervised, unsupervised and reinforcement learning.

- Supervised: distinguishing real weapons from toy weapons using example photos
- Unsupervised: Automatic clustering of Twitter/Weibo messages
- Reinforcement learning: Finding an optimal policy

AI for law and police

- Data-driven techniques are sensitive to the quality of data
- The quality of data is more important than the quantity
- Preparing data is more difficult than executing an algorithm on it
- You want to keep a practical application “fresh”: keep collecting and preparing data

AI for law and police

- Fear of AI
  - “black box”
  - Lawyers do not understand numbers & algorithms

Black box: the Chinese room

- Man in the room has a huge book, in which for every input Chinese sentence there is a Chinese output
- Man in the room does not understand Chinese

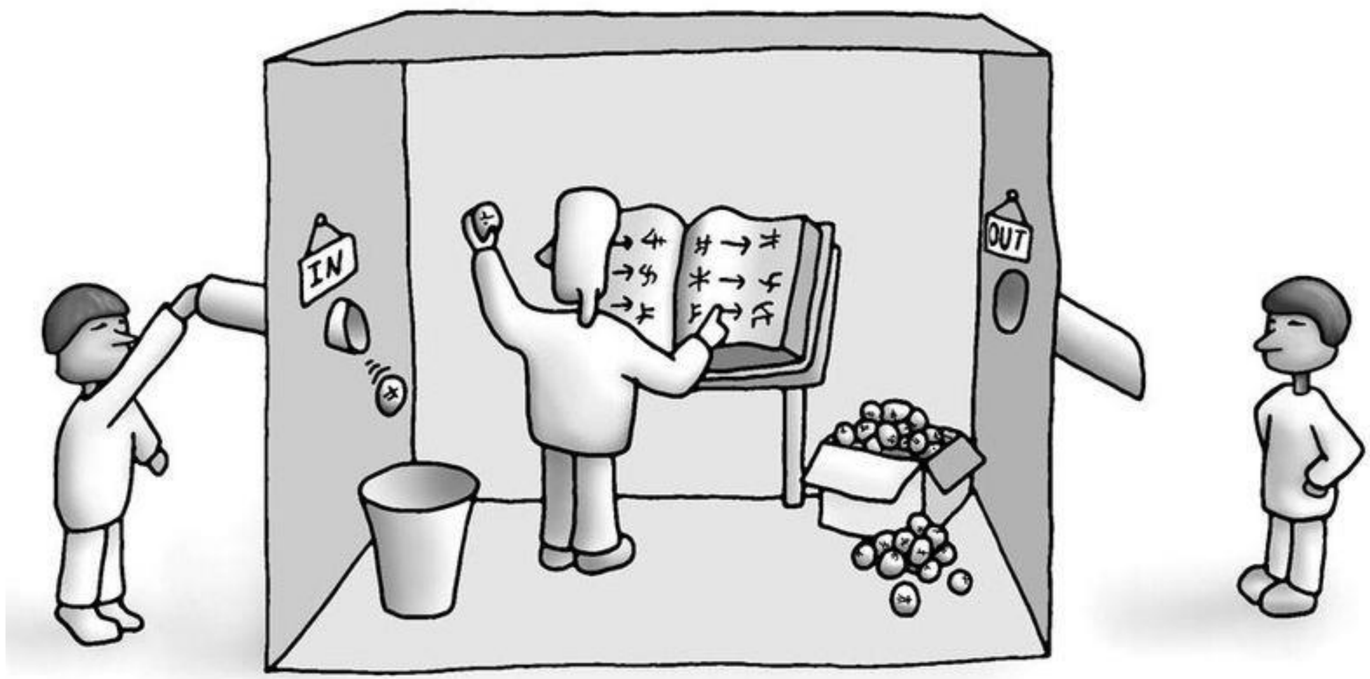

Black box: the Chinese Room

- The humanity of the person in the room adds nothing to the instruction book
- Protocol-based working is actually placing many Chinese rooms one after the other
- A.I. can replace the persons in the room
- What does this mean for the justice of the system?
  - **Many objections to A.I. also apply to modern bureaucracies.**

Numbers and algorithms

- Numbers and algorithms are very hard to understand
- But: do we know how other humans make their decision? What is the “accuracy” of human judges?
  - Human decision making works, but is also notoriously unreliable, particularly in hard/boundary cases!

AI for the legal field

- Legal field is lagging behind when it comes to AI
  - Conservative
  - Non-technical

- More work is needed
  - Data sets and resources
  - Young people who want to work on real problems
  - **Engineering** & philosophy