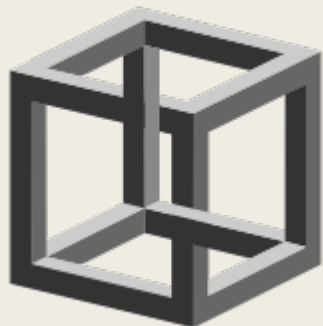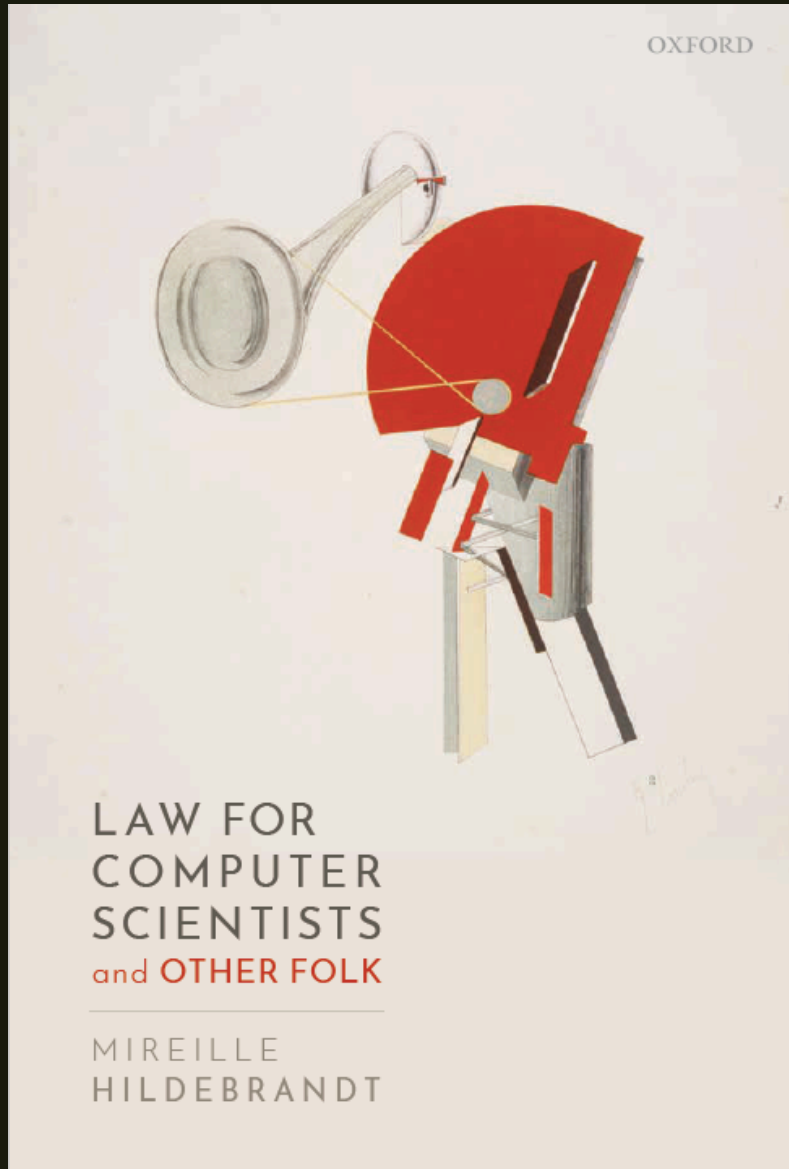# DISTANT READING OF LEGAL TEXT AND ACCESS TO LAW

Mireille Hildebrandt

Professor of 'Interfacing Law and Technology',
Faculty of Law & Criminology, Vrije Universiteit Brussel

Professor of 'Smart Environments, Data Protection and the Rule of Law,
Science Faculty, Radboud University Nijmegen

COHUBICOL

Law for CS & Other Folk

Oxford University Press forthcoming 2020

lawforcomputerscientists.pubpub.org

Law for Computer Scientists

Mireille Hildebrandt

LAW FOR CS & OTHER FOLK

AVAILABLE AT
HTTPS://LAWFORCOMPUTERSCI
ENTISTS.PUBPUB.ORG

Fundamental principles of the Rule of Law

- Access to justice
- Transparency
- Contestability
- Legal certainty

Fundamental principles of the Rule of Law

Does computational law enhance or reduce **access to justice?**

Does computational law enhance or reduce **human agency?**

How to ensure that

➢ *those who create our new onlife world*

➢ *are under the Rule of Law?*

INNOVATION OF LEGAL METHOD

'It would be nice if all of the data which sociologists require could be enumerated because then we could run them through IBM machines and draw charts as the economists do. However, not everything that can be counted counts, and not everything that counts can be counted'.

William Cameron, Informal Sociology, 1963, p. 13

25/11/19

"I think you'll find that mine is bigger…"
timoelliott.com

BIG LAW
&
BIG DATA

COHUBICOL

What's Next?

1. Human and machine reading of law
2. Pandora's box
3. NLP of legal text
4. The force of law and the force of technology
5. Interaction between Law and CS

What's Next?

1. **Human and machine reading of law**
2. NLP of legal text
3. Pandora's box
4. The force of law and the force of technology
5. Interaction between law and CS

COHUBICOL

The difference that makes a difference (Bateson)

- **Wetten.nl**
  Information Retrieval
  (wonderful resource)
  - Information as content or communication
  - Information as novelty,
    compared to knowledge background
  - Connecting legislation with
    its history and relevant case law

- **Sources of law:**
  the authentic legal 'text'
  that defines positive law
  - Information as 'informare', shaping societal
    architecture
  - Information with performative effect
    (it does what it describes)
  - Legal conditions and legal effect:
    the choice architecture of human society

Anecdotics from the European Legal Space

- Art. 33 Loi Reforme de la Justice: prohibition to use judges names for analytics

- Art. 52(e)(2) draft *Medienstaatsvertrag*: search must be 'discrimination-free'

- CRvB, 15 mei 2019, ECLI:NL:CRVB:2019:1737 (CBBS algorithmic decision-system)

LOI n° 2019-222 du 23 mars 2019 de programmation 2018-2022 et de réforme pour la justice (1)

Section 3 : Concilier la publicité des décisions de justice et le droit au respect de la vie privée

Article 33

Les données d'identité des magistrats et des membres du greffe

■ ne peuvent faire l'objet d'une réutilisation ayant pour objet ou pour effet

■ d'évaluer, d'analyser, de comparer ou de prédire

■ leurs pratiques professionnelles réelles ou supposées.

La violation de cette interdiction est punie (....).

§ 52 Medienplattformen und Benutzeroberflächen

(e) Auffindbarkeit in Benutzeroberflächen

(2) Gleichartige Angebote oder Inhalte dürfen bei der Auffindbarkeit, insbesondere der Sortierung, Anordnung oder Abbildung auf Benutzeroberflächen,

■ nicht ohne sachlich gerechtfertigten Grund unterschiedlich behandelt werden;

■ ihre Auffindbarkeit darf nicht unbillig behindert werden.

Zulässige Kriterien für eine Sortierung oder Anordnung sind insbesondere Alphabet, Genres oder Nutzungsreichweite.

Eine Sortierung oder Anordnung soll in mindestens zwei verschiedenen Varianten angeboten werden.

Alle Angebote müssen mittels einer Suchfunktion diskriminierungsfrei auffindbar sein.

Einzelheiten regeln die Landesmedienanstalten durch Satzungen und Richtlinien.

CRvB, 15 mei 2019, ECLI:NL:CRVB:2019:1737

CBBS
algorithmic decision system

- Het **Claimbeoordelings- en Borgingssysteem (CBBS)** wordt door verzekeringsartsen en arbeidsdeskundigen van **Uitvoering Werknemersverzekeringen (UWV)** gebruikt bij de WAO/WAZ/Wajong claimbeoordelingen.

- CBBS heeft een tweeledige functie.
  - Het is primair een instrument voor het **uitvoeren van de claimbeoordeling**.
  - Daarnaast levert het systeem **feedback** over deze beoordeling.

- CBBS vervangt het Functie Informatie Systeem (FIS). Sinds 1 januari 2002 worden alle WAO/WAZ/Wajong-claimbeoordelingen uitgevoerd met behulp van CBBS.

What is 'computational law'?

- **Data-Driven 'Law' (inductive)**

  Use of predictive analytics on legal text (case law, statutes, regulation)

  - *Argumentation mining*
  - *Prediction of judgement*
  - *Based on NLP (text mining) or random forests (mining of judges votes) (both supervised ML but otherwise very different assumptions)*

- **Code-Driven 'Law' (deductive)**

  Self-executing algorithmic decision-making

  - *Smart regulation (blockchain)*
  - *'Traditional' decision-support (decision-trees)*
  - *Based on IFTTT logic, painstakingly interpreted and translated*

COHUBICOL

What is text-driven law?

- Is modern positive law technologically mediated?
- Yes: technologies of the word = text
- **Modern positive law = text-driven law**

What is text-driven law?

- **Text-driven normativity followed orality:**
    - *Distantiation in time and space: author-reader-text-meaning*
    - *Evokes the need for interpretation (death of the author emancipates the text)*

What is text-driven law?

- **The cybernetics of text-driven normativity (control at a distance)**
    - Uniformity of the text across time and space (jurisdiction extended)
    - Natural language is generative because it is ambiguous (feature not bug)
    - Need for interpretation implies argumentation and contestation
    - Legal certainty: combination of foreseeability and contestability
    - Text-driven normativity generates closure as well as openings
    - Rule of Law as an affordance of text-driven normativity
    - We cannot take for granted that code- or data-driven law has similar affordances

Text-driven law = based on *close reading*

- Individual access, reconfiguring own understanding
- Creating a theatre of debate in the back of one's mind
- The 'monologue intérieur'
- Narrative and argumentative structure of human mind
- Abstract thought depends on script (external memory)
- Sapir Whorf thesis: language uses shapes the mind (both grammar & vocabulary)

COHUBICOL

NLP:
TEXT MINING *AS* DISTANT READING

BECENTRAL AI and the LAW Hildebrandt

NLP: text mining as distant reading

FRANCO MORETTI *DISTANT READING*

How does a literary historian end up thinking in terms of **z-scores, principal component analysis, and clustering coefficients**? The essays in *Distant Reading* led to a new and often contested paradigm of literary analysis. In presenting them here Franco Moretti reconstructs his intellectual trajectory, the theoretical influences over his work, and explores the polemics that have often developed around his positions.

From the evolutionary model of "Modern European Literature," through the geo-cultural insights of "Conjectures of World Literature" and "Planet Hollywood," to the quantitative findings of "Style, inc." and the abstract patterns of "Network Theory, Plot Analysis," the book follows two decades of conceptual development, organizing them around the metaphor of **"distant reading,"** that has come to define—well beyond the wildest expectations of its author—a growing field of unorthodox literary studies.
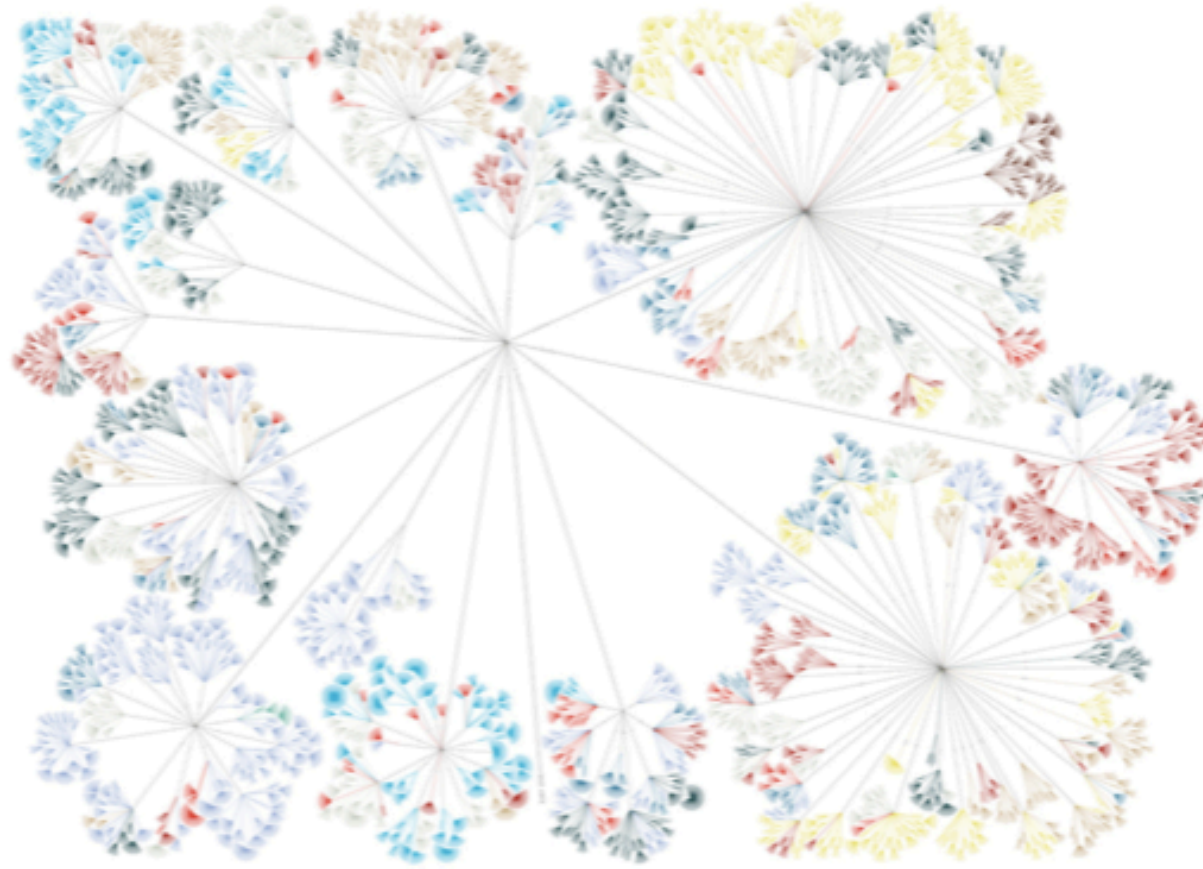
COHUBICOL

Figure 2: Distant reading example shows the structure of and the themes in Jack Kerouac's "On the Road" (Figure reproduced with permission from Posavec [Pos07]).

What's Next?

1. Human and machine reading of law
2. Pandora's box
3. NLP of legal text
4. The force of law and the force of technology
5. Interaction between law and CS

COHUBICOL

Pandora's Box
Difficult to get toothpaste back into the tube

■ **Once legal tech is employed, it may transform how we understand 'law'**

■ Citron: technological due process
  – *Interpretation, translation and execution are conflated*
  – *Enacting, applying, adjudicating law collapses into one big deal*
  – *Checks and balances get lost, redress becomes more difficult*
  – *Those who design the code are legislator, executive and court all at once*

COHUBICOL

Pandora's Box
Difficult to get toothpaste back into the tube

■ **Once legal tech is employed, it may transform how we understand 'law'**

■ The accessibility of the binding legal texts within the European Legal space could:
- *enable forum shopping that may generate a Delaware effect*

Pandora's Box
Difficult to get toothpaste back into the tube

- **Once legal tech is employed, it may transform how we understand 'law'**

- Increasing use of automated decision systems within public administration will pressure legislatures
  - *to articulate statutory law in a way that is amenable to 'codification'*

COHUBICOL

Pandora's Box
Difficult to get toothpaste back into the tube

■ **Once legal tech is employed, it may transform how we understand 'law'**

■ The urge to provide 'easy access to clear and consistent law' in combination with 'eTranslation technologies' may result in

– *monolingualism to the extent that training data focus on English translation*

– *consistent misinterpretation due to the bugs inherent in eTranslation*

Pandora's Box
Black boxing access to law?

- **Once legal tech is employed, it may transform how we understand 'law'**

- Technical standardisation will open Pandora's box because the law will serve as training data for predictive analytics:
    - *Both for case law of the European courts and for national courts*
    - *This will further increase the ability to engage in forum shopping*
    - *It will also increase the use of legal tech by e.g. Big Law*

COHUBICOL

Pandora's Box
Black boxing access to law?

■ **Once legal tech is employed, it may transform how we understand 'law'**

■ Developing and/or purchasing legal analytics is a costly affair
– *if Big Law gains an advantage this will endanger the foundations of both law and the Rule of Law*
– *argumentation (based on close reading) will in part be replaced by correlation (based on distant reading)*

COHUBICOL

Pandora's Box
Black boxes access to law?

1. *intentional secrecy*

   – *trade secrets, IP rights, public security*

2. *current education invests in writing and reading natural language, not in code or ML*

   – *monopoly of the new clerks, the end of democracy, unless ...*

2. *kmismatch between math-optimization in high-dimensional ML and human semantics*

   – *when it comes to law and justice we cannot settle for 'computer says no'*

   – Cp. *https://journals.sagepub.com/doi/abs/10.1177/2053951715622512*

COHUBICOL

Pandora's Box
Black boxing access to law?

- **Once legal tech is employed, it may transform how we understand 'law'**

- Economic incentives will prioritize proprietary analytics, which will co-opt open source initiatives (e.g. Aletras et al)
  - *This will generate black boxes that in point of fact reduce the accessibility of the sources of law*
  - *While also halting and disrupting the development of law, as these systems can only be trained on historical data*

COHUBICOL

What's Next?

1. Human and machine reading of law
2. Pandora's box
3. NLP of legal text
4. The force of law and the force of technology
5. Interaction between law and CS

COHUBICOL

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.

COHUBICOL

■ Assumption: text extracted from published judgments are a proxy for applications lodged with the Court

- *why? published judgments = low hanging fruit*
- *problem: as authors state, facts may be articulated by court to fit the conclusion*

COHUBICOL

PUBLICATION BIAS
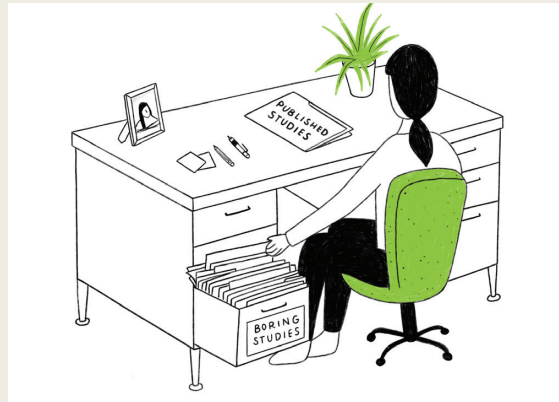
*How interesting a research finding is affects how likely it is to be published, distorting our impression of reality.*

· · · · · · · · · · · · ·

For every study that shows statistically significant results, there may have been many similar tests that were inconclusive. However, significant results are more interesting to read about and are therefore more likely to get published. Not knowing how many 'boring' studies were filed away impacts our ability to judge the validity of the results we read about. When a company claims a certain activity had a major positive impact on growth, other companies may have tried the same thing without success, so they don't talk about it.

G GECKOBOARD

■ Cases held inadmissible or struck out beforehand are not reported, which entails that a text-based predictive analysis of these cases is not possible.

  – *why? admissible cases = low hanging fruit*

  – *problem: these cases would probably make a difference which now remains invisible*

SURVIVORSHIP BIAS
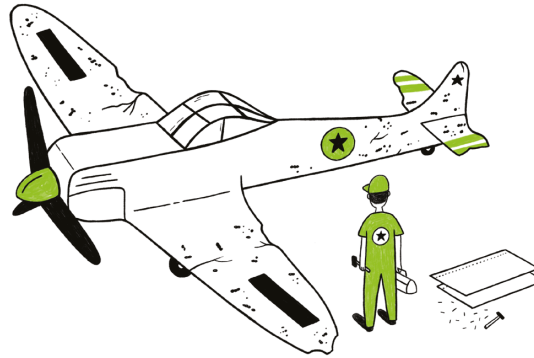
Drawing conclusions from an incomplete set of data, because that data has 'survived' some selection criteria.

When analyzing data, it's important to ask yourself what data you don't have. Sometimes, the full picture is obscured because the data you've got has survived a selection of some sort. For example, in WWII, a team was asked where the best place was to fit armour to a plane. The planes that came back from battle had bullet holes everywhere except the engine and cockpit. The team decided it was best to fit armour where there were no bullet holes, because planes shot in those places had not returned.

G GECKOBOARD

■ Data on cases related to art. 3, 6, 8 ECHR

– *why? provided the most data to be scraped, and sufficient cases for each*

– *problem: impact of framing of the case remains invisible (think e.g. art. 5, 7, 9, 10, 14)*

■ Dataset = publicly available

- – *for each article: all cases [apart from non-English judgments]*
- – *equal amount of violation/non-violation cases*
- – *text extraction by using regular expressions, excluding operative provisions*

COHUBICOL

- Prediction is defined as a binary classification task: yes/no violation:
  - *using each set of textual features to train SVM classifiers*
  - *linear kernal to check the weight learned for each feature*
  - *violation cases labelled +1, non-violation labelled -1*
  - *features with positive weights indicative of violation, with negative indicative of non-violation*

  - *trained and tested by 10-fold cross validation, a held-out of 10% for testing*
  - *performance computed as **mean accuracy** after 10-fold cross-validation*

Table 2 **Accuracy of the different feature types across articles.** Accuracy of predicting violation/non-violation of cases across articles on 10-fold cross-validation using an SVM with linear kernel. Parentheses contain the standard deviation from the mean. Accuracy of random guess is .50. **Bold** font denotes best accuracy in a particular Article or on Average across Articles.

Feature Type		Article 3	Article 6	Article 8	Average
N-grams	Full	.70 (.10)	.82 (.11)	.72 (.05)	.75
	Procedure	.67 (.09)	.81 (.13)	.71 (.06)	.73
	Circumstances	.68 (.07)	.82 (.14)	.77 (.08)	.76
	Relevant law	.68 (.13)	.78 (.08)	.72 (.11)	.73
	Facts	.70 (.09)	.80 (.14)	.68 (.10)	.73
	Law	.56 (.09)	.68 (.15)	.62 (.05)	.62
Topics		**.78 (.09)**	.81 (.12)	.76 (.09)	.78
Topics and circumstances		.75 (.10)	**.84 (0.11)**	**.78 (0.06)**	.79

COHUBICOL

- **Circumstances and topics are best predictors, combined works best**
  - *law has lowest performance*
    - discussion: facts more important than law
    - legal formalism and realism: evidence that legal realism is realistic

COHUBICOL

This is **nonsense** for 2 reasons:

1. as indicated by the authors the facts, formulated by the court, may be tuned to the outcome

2. in many cases there is no law section due to an inadmissibility judgment

3. To seriously make sense, one would need the facts of 'cases' that did not reach the court...

What's Next?

1. Human and machine reading of law
2. Pandora's box
3. NLP of legal text
4. The force of law and the force of technology
5. Interaction between law and CS

The force of law and the force of technology

■ *The European Forum of Official Gazettes:*
  – *Text with legal effect*
  – *Legal effect = the force of law*
  – *The force of law = performative speech act in text*

■ *Text-mining, predictive analytics, 'codification of law':*
  – *Operations with legal effect?*
  – *Does code generate legal effect?*
  – *Or does it thrive on the force of technology?*

COHUBICOL

Whiteboxing predictive legal tech?

- used as a means to provide feedback to lawyers, clients, prosecutors, courts

- could involve a sensitivity analysis, modulating facts, legal precepts, claims

- as a domain for experimentation, developing new insights, argumentation patterns, testing alternative approaches

- could detect missing information (facts, legal arguments), helping to improve the outcome of cases

- can be used to improve the acuity of human judgment, if not used to replace it

- if used to replace, it should not be confused with law; *then is becomes administration – the difference is crucial, critical and pertinent*

- cp. http://www.vikparuchuri.com/blog/on-the-automated-scoring-of-essays/

COHUBICOL

'Codification' under the Rule of Law?

■ *Automated decision-making is not law, but public administration*

■ *It cannot be 'legal by design', but may contribute to legal protection by design*

■ *Automated decision-making in public administration must be brought under the Rule of Law (connection with art. 22 GDPR, legal remedies in administrative law):*

- – *Democratic legitimation (representation, deliberation, participation)*
- – *Resistability (otherwise not law but administration)*
- – *Contestable in a court of law (under the Rule of Law)*

COHUBICOL

COUNTING AS A HUMAN BEING
IN THE ERA OF COMPUTATIONAL LAW

COHUBICOL

SAY CUBICLE • THINK WITTGENSTEIN'S CUBE

NEWS	ON THE PROJECT	RESEARCH BLOG	COMPUTATIONAL LAW	LEGAL PROTECTION	PRESS	RESEARCH OUTCOME

NOW HIRING @Radboud:
2 postdoctoral researchers in CS
for foundational research into 'legal tech'

This is your chance to dig into the **fundamental assumptions underlying computer science,** teasing out the **implications** they may have for **real life applications**, notably those of 'legal tech'. The combination of research into the **theory of computer science** and the **opportunity to make a difference** in the legal domain provides a unique opening for those willing to address the societal

What's Next?

1. Human and machine reading of law
2. Pandora's box
3. NLP of legal text
4. The force of law and the force of technology
5. Interaction between Law and CS

DATA FALLACIES TO AVOID
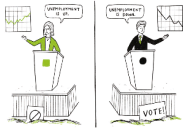
CHERRY PICKING
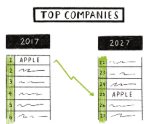Selecting results that fit your claim and excluding those that don't.

DATA DREDGING
Repeatedly testing new hypotheses against the same set of data, failing to acknowledge that most correlations will be the result of chance.

SURVIVORSHIP BIAS
Drawing conclusions from an incomplete set of data, because that data has 'survived' some selection criteria.

COBRA EFFECT
Setting an incentive that accidentally produces the opposite result to the one intended. Also known as a Perverse Incentive.
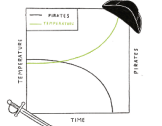
FALSE CAUSALITY
Falsely assuming when two events appear related that one must have caused the other.

GERRYMANDERING
Manipulating the geographical boundaries used to group data in order to change the result.

SAMPLING BIAS
Drawing conclusions from a set of data that isn't representative of the population you're trying to understand.

GAMBLER'S FALLACY
Mistakenly believing that because something has happened more frequently than usual, it's now less likely to happen in future (and vice versa).

HAWTHORNE EFFECT
The act of monitoring someone can affect their behaviour, leading to spurious findings. Also known as the Observer Effect.

REGRESSION TOWARDS THE MEAN
When something happens that's unusually good or bad, it will revert back towards the average over time.

SIMPSON'S PARADOX
When a trend appears in different subsets of data but disappears or reverses when the groups are combined.

MCNAMARA FALLACY
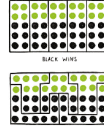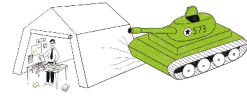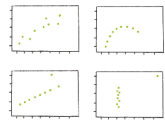Relying solely on metrics in complex situations and losing sight of the bigger picture.

OVERFITTING
Creating a model that's overly tailored to the data you have and not representative of the general trend.

PUBLICATION BIAS
Interesting research findings are more likely to be published, distorting our impression of reality.
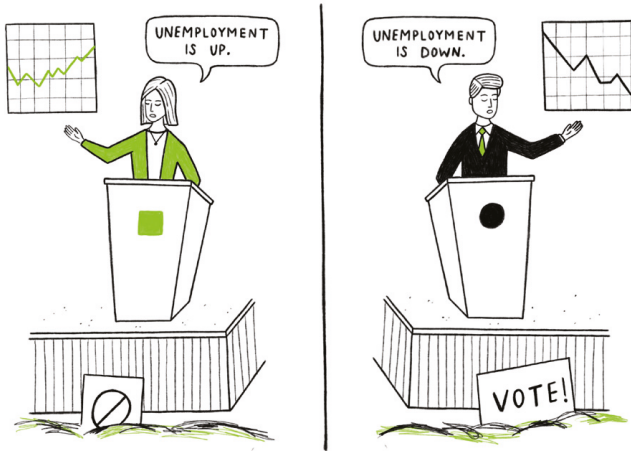
DANGER OF SUMMARY METRICS
Only looking at summary metrics and missing big differences in the raw data.

Precaution

CHERRY PICKING

The practice of selecting results that fit your claim, and excluding those that don't. The worst and most harmful example of being dishonest with data.

• • • • • • • • • • •

When making a case, data adds weight – whether a study, experiment or something you've read. However, people often only highlight data that backs their case, rather than the entire body of results. It's prevalent in public debate and politics where two sides can both present data that backs their position. Cherry Picking can be deliberate or accidental. Commonly, when you're receiving data second hand, there's an opportunity for someone choosing what data to share to distort the truth to whatever opinion they're peddling. When on the receiving end of data, it's important to ask yourself: 'What am I not being told?'.

G GECKOBOARD

Feature space matters

DANGER OF SUMMARY METRICS
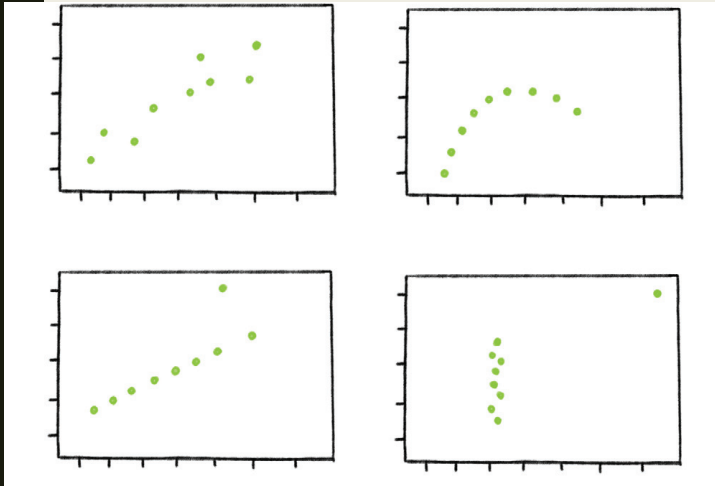
*It can be misleading to only look at the summary metrics of data sets.*

· · · · · · · · · · · · ·

To demonstrate the effect, statistician Francis Anscombe put together four example data sets in the 1970s. Known as Anscombe's Quartet, each data set has the same mean, variance and correlation. However, when graphed, it's clear that each of the data sets are totally different. The point that Anscombe wanted to make is that the shape of the data is as important as the summary metrics and cannot be ignored in analysis.
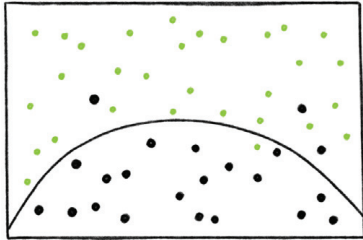
FLEXBOARD

HUBICOL

Can lawyers Understand What is not Visualised?

OVERFITTING
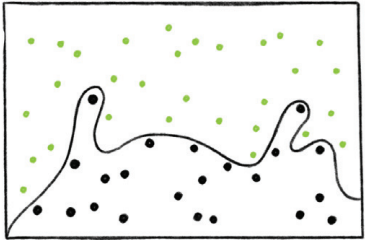
OVERFITTING   JUST RIGHT

A more complex explanation will often describe your data better than a simple one. However, a simpler explanation is usually more representative of the underlying relationship.

· · · · · · · · · · · · · ·

When looking at data, you'll want to understand what the underlying relationships are. To do this, you create a model that describes them mathematically. The problem is that a more complex model will fit your initial data better than a simple one. However, they tend to be very brittle: They work well for the data you already have, but try too hard to explain random variations. Therefore, as soon as you add more data, they break down. Simpler models are usually more robust and better at predicting future trends.
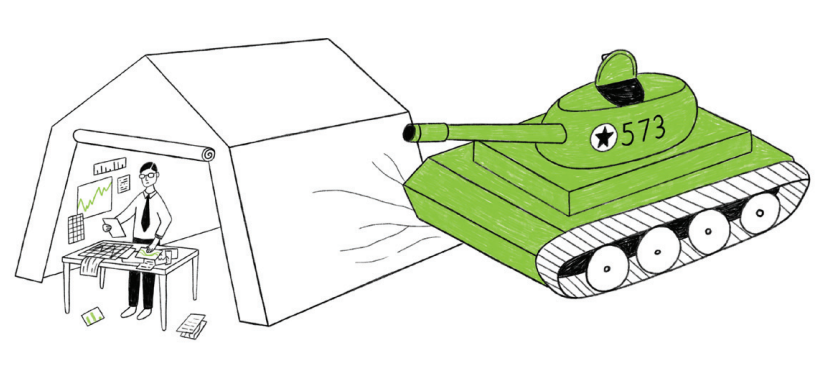
THE SKETCHBOARD

HUBICOL

Can lawyers check where to draw the line?

MCNAMARA FALLACY

Relying solely on metrics in complex situations can cause you to lose sight of the bigger picture.

· · · · · · · · · · · · ·

Named after Robert McNamara, the U.S. Secretary of Defense (1961-1968), who believed truth could only be found in data and statistical rigor. The fallacy refers to his approach of taking enemy body count as the measure of success in the Vietnam War. Obsessing over it meant that other relevant insights like U.S. body count, territorial gains, and the shifting mood of the general public were ignored. When analyzing complex phenomena, we're often forced to use a metric as proxy for success. However, dogmatically optimizing for this number and ignoring all other information is risky.

THE CATCH BOARD
HUBICOL

Can lawyers detect what data has not been taken into account?

REGRESSION TOWARD THE MEAN

TOP COMPANIES

2017 → 2027
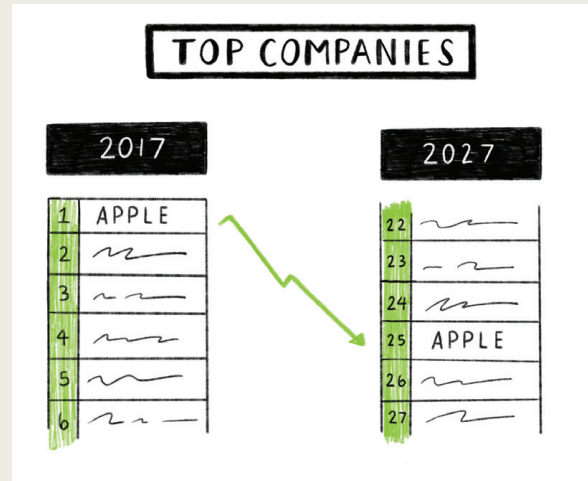
When something happens that's unusually good or bad, over time it will revert back towards the average.

· · · · · · · · · · · · ·

Anywhere that random chance plays a part in the outcome, you're likely to see regression toward the mean. For example, success in business is often a combination of both skill and luck. This means that the best performing companies today are likely to be much closer to average in 10 years time, not through incompetence but because today they're likely benefitting from a string of good luck – like rolling a double-six repeatedly.

GO THE BOARD

HUBICOL

Do lawyers understand the statistics? (that determines the output)

HAWTHORNE EFFECT
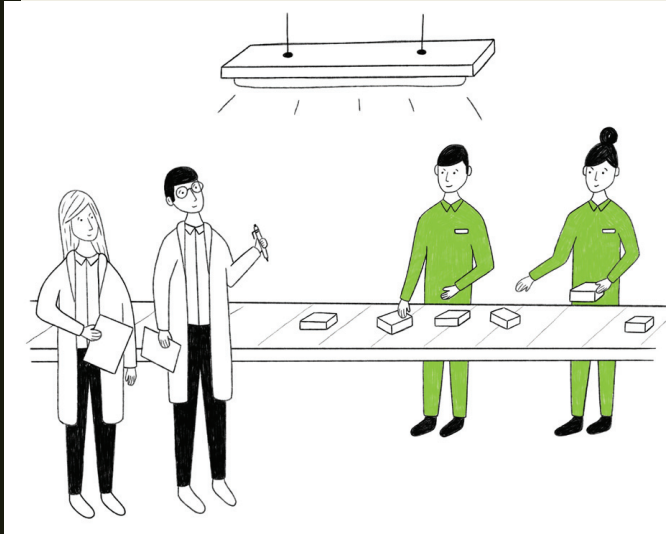
When the act of monitoring someone can affect that person's behavior. Also known as the Observer Effect.

· · · · · · · · · · · · ·

In the 1920s at Hawthorne Works, an Illinois factory, a social sciences experiment hypothesised that workers would become more productive following various changes to their environment such as working hours, lighting levels and break times. However, it turned out that what actually motivated the workers' productivity was someone taking an interest in them. When using human research subjects, it's important to analyze the resulting data with consideration for the Hawthorne Effect.

GHUBICOL

Will prediction of judgement be gamed by Big Tech or Big Law?

GAMBLER'S FALLACY
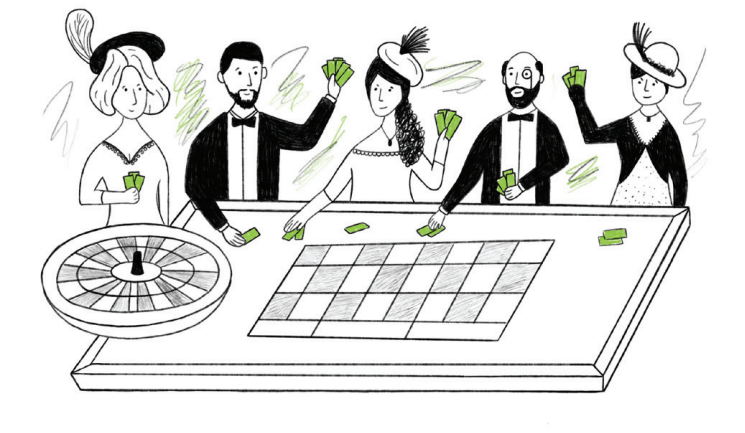
The mistaken belief that because something has happened more frequently than usual, it's now less likely to happen in future and vice versa.

• • • • • • • • • • • •

This is also known as the Monte Carlo Fallacy because of an infamous example that occurred at a roulette table there in 1913. The ball fell in black 26 times in a row and gamblers lost millions betting against black, assuming the streak had to end. However, the chance of black is always the same as red regardless of what's happened in the past, because the underlying probability is unchanged. A roulette table has no memory. When tempted by this fallacy, remind yourself that there's no rectifying force in the universe acting to 'balance things out'!

SEED BOARD

HUBICOL

25/11/19

What about the inductive fallacy?

DATA DREDGING

Slice your data in enough different ways and you'll observe some correlations purely as a result of chance. Data dredging is the failure to acknowledge that the correlation was in fact the result of chance.

· · · · · · · · · · ·

Tests for statistical significance only work if you've defined your hypothesis upfront. Historically, this has been a problem with clinical trials where researchers have 'data-dredged' their results and switched what they were testing for. It explains why so many results published in scientific journals have subsequently been proven to be wrong. To avoid this, it's now becoming standard practice to register clinical trials, stating in advance what your primary endpoint measure is.

G GECKOBOARD

Do lawyers recognize P-hacking when they see it?

We shall see that most current theory of machine learning rests on the *crucial assumption* that the distribution of training examples is identical to the distribution of test examples. Despite our need to make this assumption in order to obtain theoretical results, it is important to keep in mind that this assumption must often be violated in practice.

Tom Mitchell *Machine Learning*

We shall see that most current theory of machine learning rests on the crucial assumption that *the distribution of training examples is identical to the distribution of test examples.* Despite our need to make this assumption in order to obtain theoretical results, it is important to keep in mind that this assumption must often be violated in practice.

Tom Mitchell *Machine Learning*

Another – even more fundamental – assumption of machine learning

is that of an underlying mathematical reality

that maps human intercourse

COHUBICOL