Algorithm Design and Law: Potential and Challenges

(from a theoretical computer scientist's perspective)

Inbal Talgam-Cohen
Technion – Israel Institute of Technology

Beyond Differential Privacy Workshop
Simons Institute, May 2019

Acknowledgement

- This talk is based on ongoing discussions with Josef Drexl, Niva Elkin-Koren, Michal Feldman, Shafi Goldwasser, Noam Nisan, Ariel Porat and others [any errors my own]

Is There Such a Thing as Algorithm Design and Law?

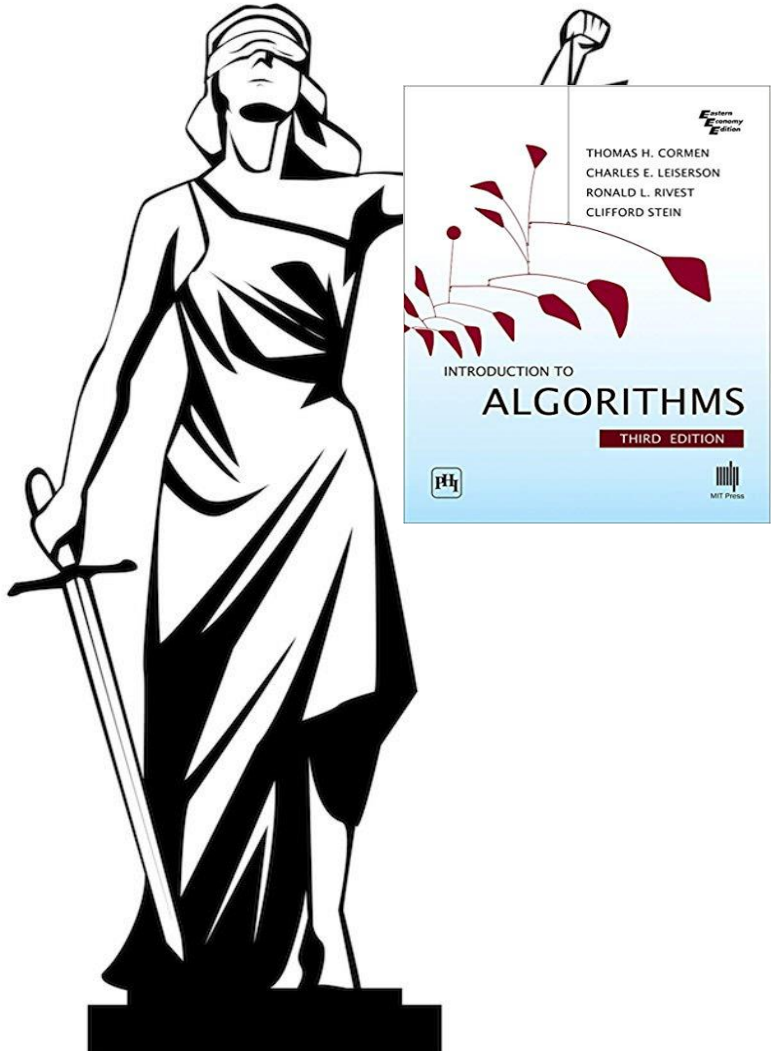

Is There Such a Thing as Algorithm Design and Law?

- Technological developments serve as catalysts for new fields of study

Disciplinary Examples

1. Technological developments:
   - Photography, recording, widespread newspaper circulation
- New theory needs:
   - A definition of the legal right to privacy [Warren-Brandeis 1890]

2. Technological development:
   - Cyberspace
- New theory needs:
   - Study of regulatory role of computer code (“code is law”) [Lessig’99]

Interdisciplinary Example: Algorithmic Game Theory

3. Technological development:
   - Online markets and the internet in general
- New theory needs:
   - Guiding theory for computer scientists engaging in market design and algorithm design in strategic environments [1999 seminal papers]

- Success story of cross-disciplinary collaboration
   - CS and econ researchers in daily interaction
   - Joint annual ACM conference on economics and computation (EC)

Is There Such a Thing as Algorithm Design and Law?

- Technological developments serve as catalysts for new fields of study

- Usage of algorithms in society is exploding
- Algorithms infiltrating – and increasingly governing – every aspect of our lives as individuals and society ("algorithms are law")

- Differential privacy and fairness can arguably be seen as (the first?) two research fields to arise of this

Differential Privacy and Fairness

4. Technological development:
   - Increasingly detailed electronic data about individuals
- New theory needs:
   - Mathematical definition of private data analysis [Dwork et al.'06]
   - A rich class of algorithms that satisfy this definition

5. Technological development:
   - Increasingly accurate algorithmic predictions
- New theory needs:
   - Mathematical definition(s) of fairness + algorithms satisfying fairness

Generalization: Algorithm Design and Law

- Technological development:
  - Far-reaching effects of algorithms on societal values (privacy, fairness and beyond)
- New theory needs:
  - Guiding theory for algorithm designers engaging in social engineering

- Is it time for a wider academic collaboration on algorithms and law (a la algorithmic game theory)?
  - Scope and content?
  - Potential gains and challenges?

This Talk

- Initial thoughts and directions to facilitate a discussion
- In spirit of the “unbaked ideas” session
- A biased/narrow viewpoint?
- Will try to go beyond privacy and fairness

Out of Scope for this Talk

- “Computational law” / “AI in service of law”
  - Automated dispute resolution
  - Automated legal reasoning
  - Legal text mining
  - Legal knowledge representation
  - E-government
- Existing communities and dedicated scientific events like ICAIL

Timeliness of Discussion

Mutual Interest

- Computer science stands to gain:
  - Access to decades of legal thought
  - New problems to apply our tools and way of thinking to
  - In line with workshop goal of "*surfacing problems that would benefit from the attention of the CS theory lens*"
- Law stands to gain:
  - Opportunities to expand and reshape legal doctrines
  - Inspiration from rigorous mathematical approaches
- Each community may need the other for its work to stay relevant

“Computer scientists cannot solve algorithmic fairness (and privacy in data analysis or any other issue of this sort) on their own.

On the other hand, these issues, in their current computation-driven large-scale incarnation, cannot be seriously addressed without major involvement of computer scientists.

Furthermore, what is needed is a true collaboration, rather than a division of work, where one community sub-contracts another for specific expertise."

Notable Signs of a Tightening Collaboration

- ACM Inaugural Symposium on “Computer Science and Law”
- Co-chaired by Pamela Samuelson, Daniel Weitzner
- October 2019 in New York
- Aims:
  - Bridge the divide between CS and law
  - Stimulate interest in the emerging field
  - Articulate a research agenda
  - Recommendations on how ACM and other institutions can support

Symposium Topics

- Security, privacy, encryption, and surveillance
- Cyber espionage, cyber war, and cyber diplomacy
- Cyber crime, cyber law enforcement, and digital forensics
- Freedom of expression online (or the lack thereof)
- Online market structure, platform monopolies, and antitrust law
- Online government services
- Digital intellectual property
- Legal informatics
- Automation of legal reasoning and legal services
- Fairness, accountability, transparency, and ethics (FATE) in machine learning and data mining
- Methodological compatibility and incompatibility between the discipline of computer science and the discipline of law

Notable Signs of a Tightening Collaboration

- Fall 2018 Course on “Law for Algorithms”
- Taught jointly at Berkeley, BU, Columbia and Harvard
- Open to law and CS students working jointly on assignments
- Instructed by Daniela Caruso, Ran Canetti, Stacey Dogan, Cynthia Dwork, Shafi Goldwasser, Martha Minow, Patricia Williams
- Topics include verifiability, proofs, identity, autonomy, consent, fairness, privacy, secure computation, governance, trust, voting and online platforms

Notable Signs of a Tightening Collaboration

- ACM Conference on Fairness, Accountability and Transparency (FAT*)
- From speaker instructions: “This is an interdisciplinary conference… Computer scientists – please note that this means we are encouraging you to give a different style of talk than you usually would at a conference, and defer the technical details to your paper.”

- [Compare to ACM Conference on Economics and Computation (EC)]

- Multiple additional events (including this workshop!)

Algorithm vs. Legal Doctrine

Algorithm

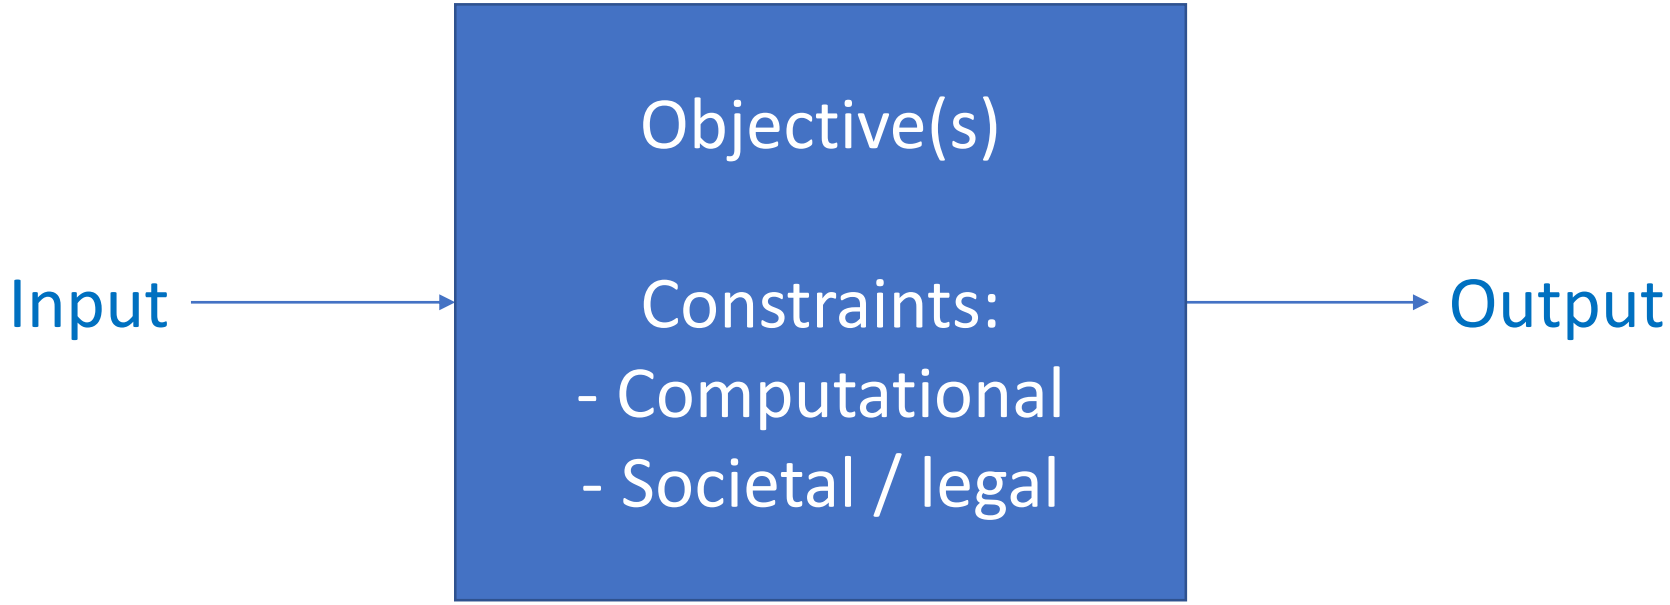

- Computational steps that map input to output, optimizing the objective subject to the constraints
- May use learning

Example: Autonomous Driving Policy

- Reinforcement learning

The Study of Algorithm Design

- Find good algorithms or show none exist
- What is a good algorithm?
- Example: While maintaining scalability and safety,
  - Algorithm 1 gets quickly from Berkeley to Stanford;
  - Algorithm 2 gets quickly from point A to B within San Francisco
- Which is better?
- Worst case approach: The algorithm is as good as its performance for its worst input
- [Compare to: “Hard cases make bad law”]

Legal Doctrine

- Legal view of same scenario:
  - Different drivers have competing rights to use the road
  - Competing societal values of useful transportation vs. safety
- Legal doctrines like negligence balance competing rights and values

- Negligence doctrine in a nutshell:
  - Driver must avoid negligence (not taking precautions that cost less than the expected damages) or pay damages where there is a duty of care

Open-ended

Algorithms Shaped by Legal Doctrines (and Vice Versa)

3 Case Studies

Encoding Legal Doctrines

- Should the safety constraint of the autonomous driving algorithm mathematically formulate the negligence doctrine?

- Same question relevant to almost any algorithmic task with implications to societal values

Algorithms Shaped by Legal Doctrines

Algorithmic Task	Societal Value	Math. Formulation	Legal Doctrine
Data analysis	Privacy	Differential	Right to privacy
Prediction	Fairness	Multiple notions	Anti-discrimination
Driving	Safety	Initial proposal	Negligence
Pricing	Free market	From economics	Antitrust
Content moderation	Freedom of speech	Largely open (?)	Copyright, hate speech, libel

Case Study 1: Formalizing Duty of Care

- "On a Formal Model of Safe and Scalable Self-driving Cars"
  - Shalev-Shwartz, Shammah and Shashua*
  - Working paper, 2017

*CS researchers; 1st and 3rd authors are founders of Mobileye

From [SSS'17]

- Introduce a model called RSS (Responsibility Sensitive Safety)
- RSS "formalizes an interpretation of duty of care from tort law" applicable to self-driving cars
- Designed to achieve 3 goals
  1. Compatibility with human interpretation of the law
  2. Useful (not over-defensive) driving
  3. Tractable verification
- If all agents follow RSS then guaranteed "utopia" (no accidents)

RSS Formalizes 5 Common Sense Rules

1. Do not hit someone from behind.
2. Do not cut-in recklessly.
3. Right-of-way is given, not taken.
4. Be careful of areas with limited visibility
5. If you can avoid an accident without causing another one, you must do it.

Example of RSS Model

- Simplest possible scenario:
  - Single-lane, straight road
  - Cars driving forward

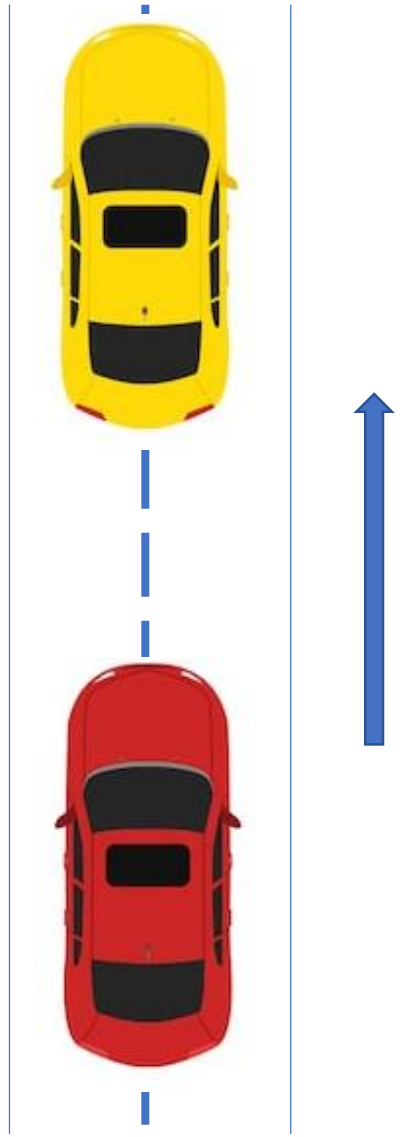

Example of RSS Model

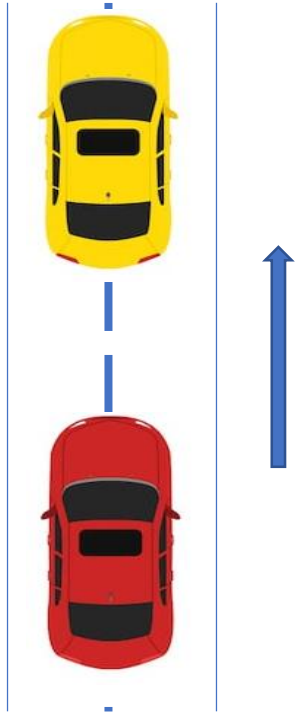

Define this as a "proper response" of Red

- Parameters:
  - ρ = reasonable response time
  - a = max. reasonable acceleration
  - x = max. reasonable break (dea eration)
  - n = min. reasonable break (de celeration)

- Definition: d is a safe distance for given velocities if assuming Yellow breaks at rate $\leq x$, and Red accelerates at rate $\leq a$ during response time ρ and breaks at rate $\geq n$ afterwards, then Red doesn't hit Yellow

Example of RSS Model

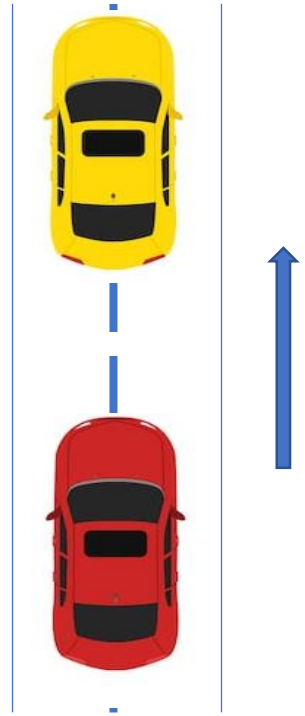

- Claim: By properly responding to a violation of safe distance, Red does not hit Yellow (assuming Yellow breaks at rate $\leq x$)
- This can be formally proved by induction on time intervals

- Scalability: Proper response can be verified with respect to every other car on the road individually ("star-shaped computations")
- [Compare to a duty on Yellow to slightly accelerate if safe distance is violated]

RSS Generalizes to Complex Scenarios

- Same principals of safe distance and proper response can be adapted to complex scenarios while remaining scalable:
  - Passing between multiple lanes
  - Driving in both directions
  - Compensating for improper behavior of others
  - Roundabouts, junctions, merging – right-of-way or traffic lights
  - Unstructured roads
  - Pedestrians
  - Occlusions

Discussion of Case Study 1

- Definition of proper response is supposed to formalize duty of care
  - Do you agree?
- The definition takes into account computational tractability
- It seems to work in experiments!

- Choice of parameters trades off usefulness and safety
  - May take into account the driver (human/robot), road conditions
  - [Will related to case study 3]

Case Study 2: Collusion by Pricing Algorithms

- "Artificial Intelligence, Algorithmic Pricing and Collusion"
  - Calvano, Calzolari, Denicolo and Pastorello*
  - Working paper, 2019

  *Economics researchers

- Collusion in economics is a reward-punishment scheme that leads to prices and profits above some competitive benchmark

Pricing Algorithms Learn to Collude

- Despite no communication!

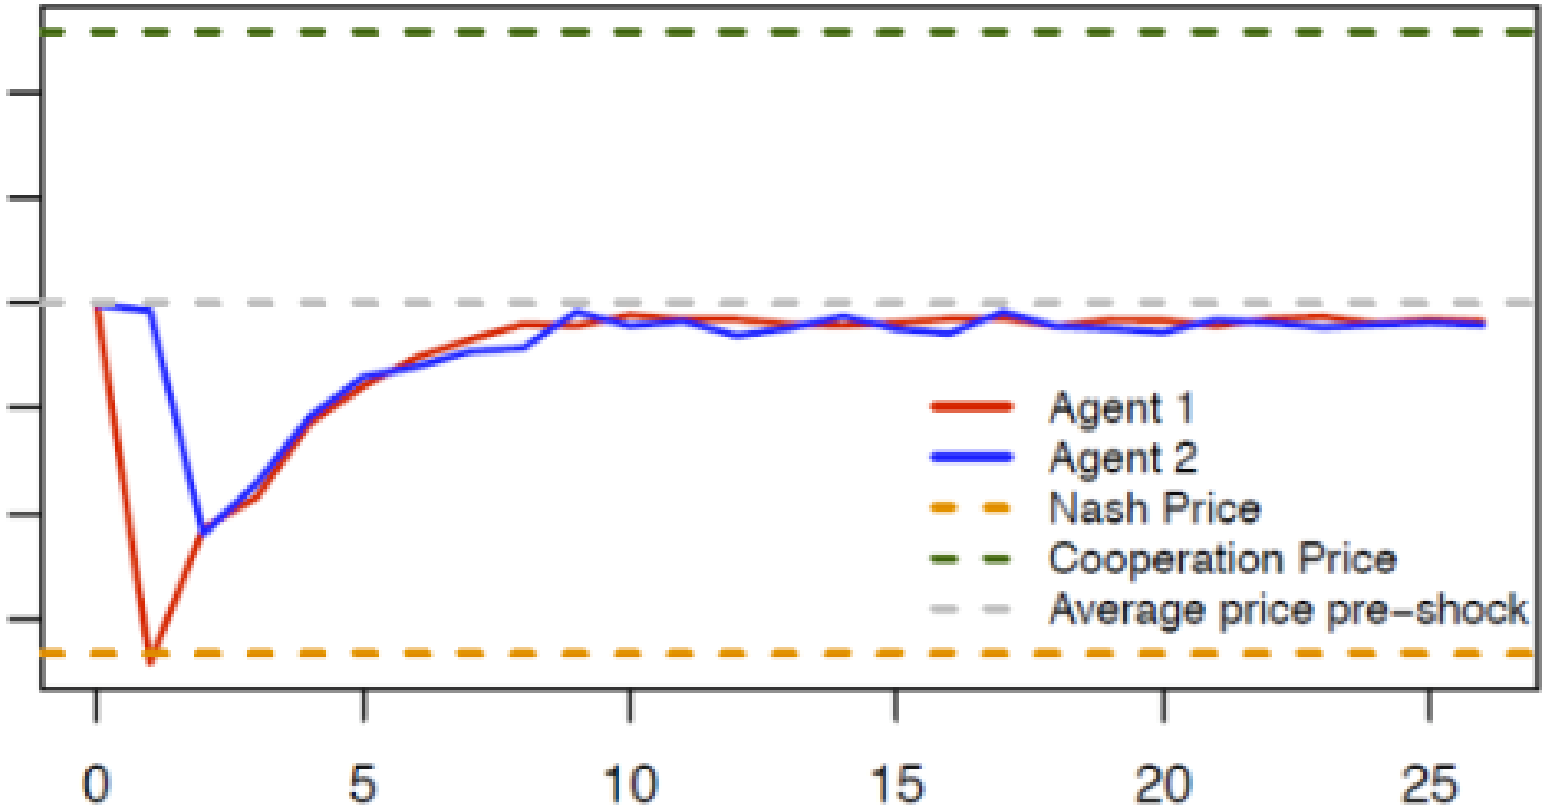

Algorithmic vs. Human Collusion

- Learning to collude persists with
  - Multiple algorithmic agents
  - Noisy information

- For humans, tacit collusion (absent communication) is hard to achieve
  - Seem to need explicit threats to punish deviation

Discussion of Case Study 2

- Pricing algorithms should be constrained not to collude
- But without hindering their flexibility to react to economic shocks
- Challenging problem!

- Current antitrust legislation circumvents it by forbidding explicit collusions (at least “wink and nod”)

- Simultaneous need to revise algorithms and law [Gal’18]

Case Study 3: Personalized Law

- Up till now we saw algorithms “getting closer” to legal doctrines
- What about laws getting closer to algorithms?
  - In particular, what if a law had input?

- “Personalizing Default Rules and Disclosure with Big Data”
  - Porat and Strahilevitz*
  - Michigan Law Review, 2014.

  *Law researchers

Example from Inheritance Law

- 55% of married fathers leave everything to their spouse
- 34% of married mothers leave everything to their spouse
- Most individuals leave no wills – default rule kicks in

- A gender-based default rule would better implement true preferences; lower expenses of drafting wills

- More ambitiously, the default rule could be fully personalized

Discussion of Case Study 3

- CS experience with personalization very relevant
  - E.g., personalized pricing, personalized ads
  - Privacy and fairness issues
  - Fragmentation, uncertainty

- The ultimately personalized law will probably be implemented by an algorithm

Discussion

Challenges to Algorithm Design and Law

1. Fundamentally different approaches?
   - Worst-case approach of CS too restrictive?
   - Open-ended approach of law too informal? [Elkin-Koren'16]
   - [Nissim et al.'17] attempts to bridge the gap in the context of privacy
2. Impossibility results, e.g.:
   - Incompatibility of natural fairness notions [Kleinberg et al.'17]
   - Security against singling out does not self-compose [Cohen-Nissim'19]
3. Mathematical modeling
   - Start with math-oriented areas of legal research?
4. Auditing/transparency/accountability

Additional Potential Areas of Mutual Interest

- Algorithmic content regulation
  - Algorithm's purpose is to implement legal doctrines like fair use
  - Industry-led
- Blockchain-based smart contracts – algorithms as legal texts
  - Issues of interpretation
  - Is contract law relevant?
- Proof notions
  - Are probabilistic proof notions applicable in law?
- Other (liquid democracy, net neutrality, …)

Summary

- Algorithms currently implementing balances among societal values that were traditionally addressed by legal doctrines

- Within CS – new avenues for algorithm design beyond differential privacy and fairness
  - 2 examples – negligence, antitrust

- Between CS and law – potential (e.g. personalized law) and challenges of joint research

Some Questions for Discussion

- What do you consider as “algorithm design and law”?
- What could you contribute to the other discipline? What would you want to get out of the conversation?
- Have you collaborated in the past with the other discipline? Gains and challenges?
- How do other disciplines (e.g. economics) fit in the picture?
- What’s next?
  - Proposal: Joint workshop dedicated to algorithm design and law

Thank you!

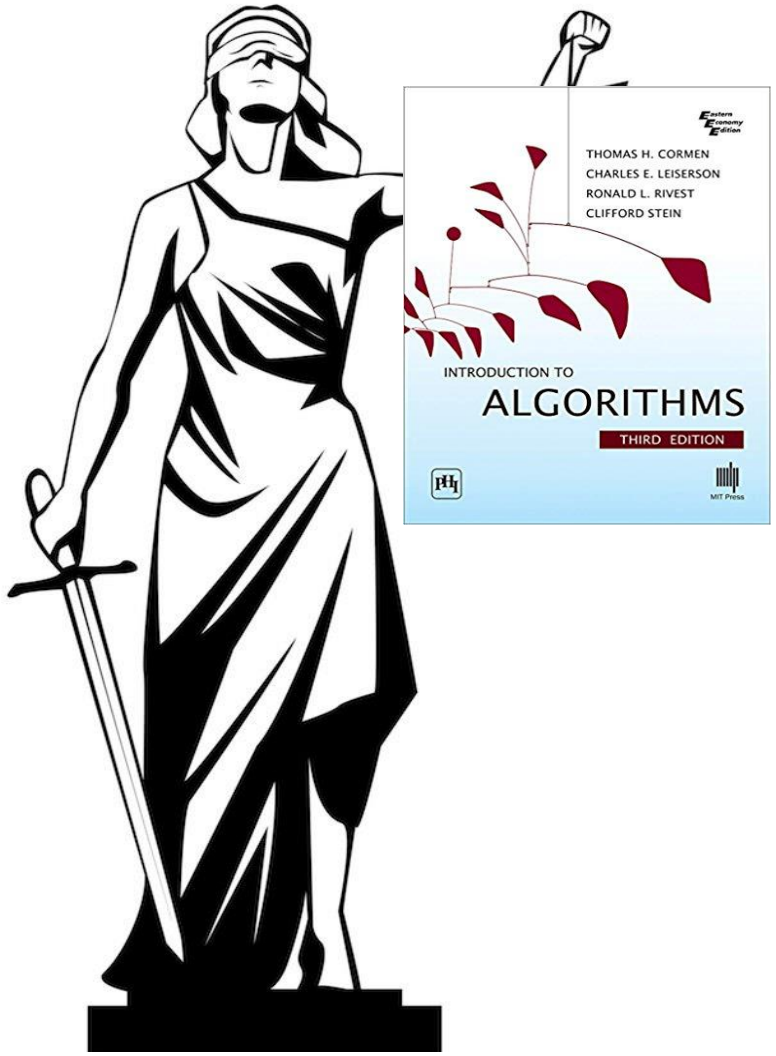
THOMAS H. CORMEN
CHARLES E. LEISERSON
RONALD L. RIVEST
CLIFFORD STEIN
INTRODUCTION TO
ALGORITHMS
THIRD EDITION