

Imitation Games

Avi Wigderson

Institute for Advanced Study

Plan

The original Imitation Game

Intelligence

The main idea, generalizations, speculations

Modern Imitation Games

Cryptography


Randomness

Combinatorics and number theory

Privacy

Definitions and major consequences

Conclusions



Alan Turing 1912-1954

Computer science

Computing revolution

Enigma machine

Biological modeling

Artificial Intelligence

Intelligence

[Turing'50] *Computing machinery and intelligence*


How to study the question "can machines think?"

1) Define machine & think. (philosophical/ontological)

2) Test it!

(operational/behavioral)

[Turing '50] Turing's test: the original imitation game




H




H

“Definition”: Robot is intelligent if the referee makes a similar guess in both worlds

[Turing '50] Turing's test: the original imitation game




H




R

“Definition”: Robot is intelligent if the referee makes a similar guess in both worlds

Turing '50] Turing's test: the original imitation game




92%




93%

“Definition”: Robot is intelligent if the referee makes a similar guess in both worlds

Turing '50] Turing's test: the original imitation game




75%




74%

“Definition”: Robot is intelligent if every human referee makes a similar guess in both worlds

Turing '50] Turing's test: the original imitation game



P₀



P₁

“Definition”: Robot is intelligent if for every human referee, $P_0 \approx P_1$

Cognition, emotions, ...

Can a computer ~~xxx~~ behave as if it feels

- Empathy, fear, pain, ... ?
- Conscious, self-aware, ... ?

How about me? How about my dog?

What is the difference between carbon-based and silicon-based machines ?

Objective, ontological definitions may not be testable, falsifiable, have universal meaning, ...

Subjective, behavioral definitions may be precise, testable, operational, useful, ...,

revolutionary

Rigorous Imitation Games

- Theory of Computation
- Discrete Mathematics

Paradigm: two things are the same if they cannot be told apart by any reasonable test

Formal definitions of central notions - primary

Theorems, proofs, constructions, theories,...


Power of the **Imitation Game** paradigm - enable math science technology policy!

Cryptography

Encryption: the first 2000 years

Enc(ATTACK) = XQQXZH
Enc(RETREAT) = OBQOBXQ

Caesar's cipher




Enigma machine



Is **Enc** secure?

What is "secure"?

Secure encryption



Definition: Enc is secure if for every m_0 & m_1 and for every efficient algorithm $P_0 \approx P_1$

Theory: Shooting for the moon (and getting there!) Defining secret. Shows Enc must be probabilistic!

Modern cryptography

last 4 decades

Public-key Encryption

On-line shopping

Contract signing

Secret exchange

Zero-knowledge proofs

Internet elections

Poker on telephone

Blockchains & digital currency

.....

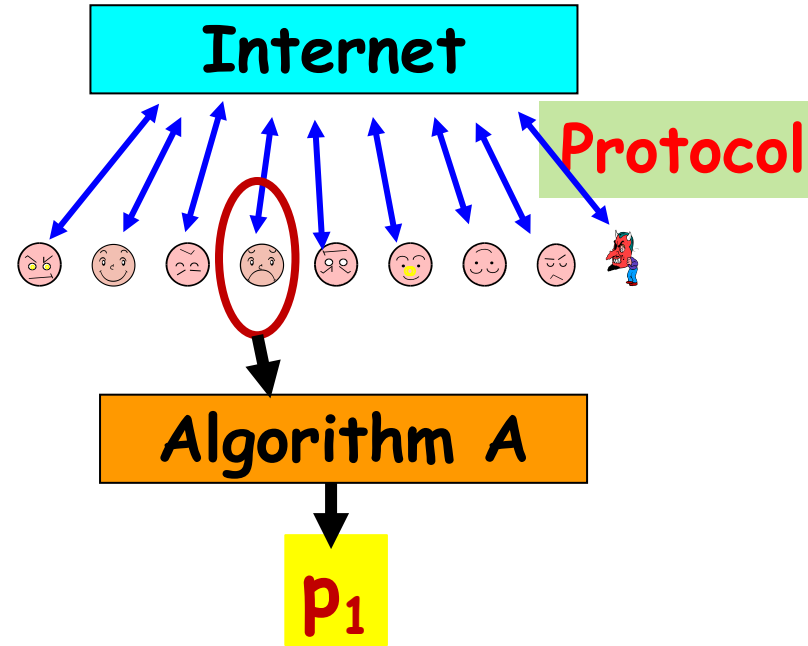
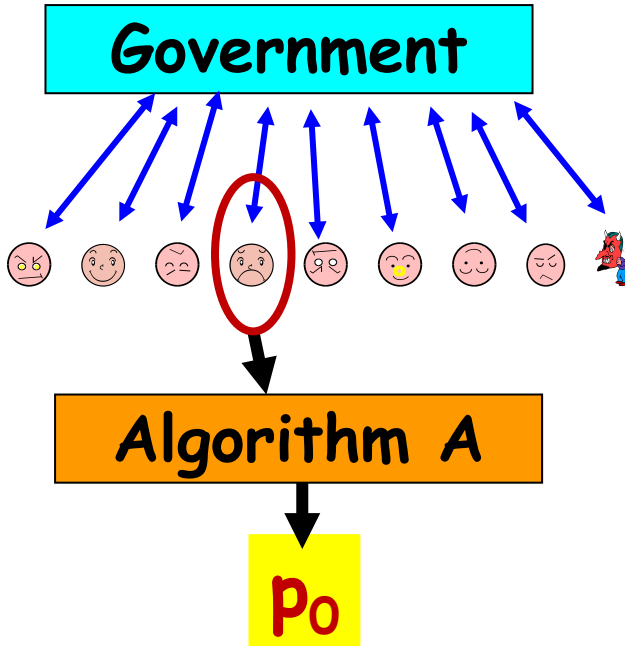
Everything (w/out physical implements!)

Is a given
protocol secure?

What is "secure"?

[Yao'86, Goldreich
-Micali-W '87]

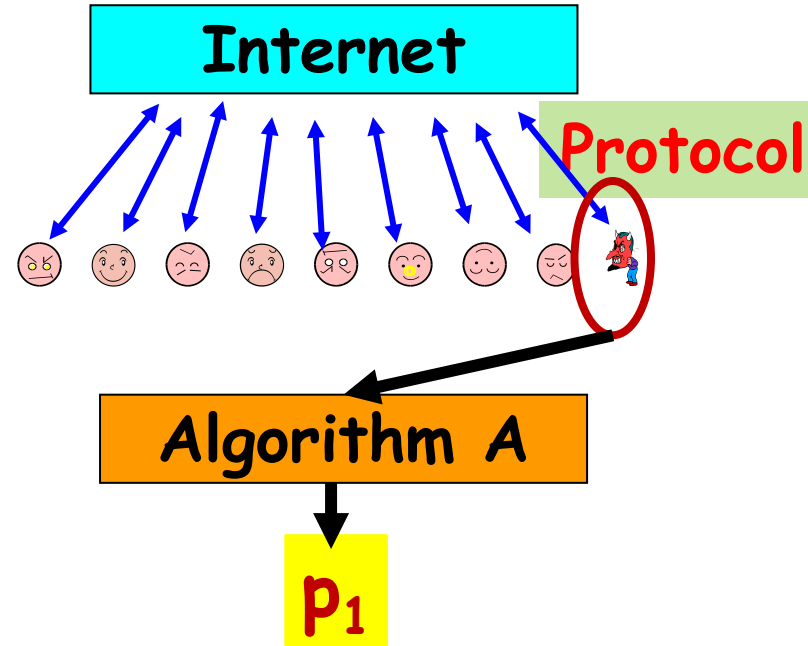
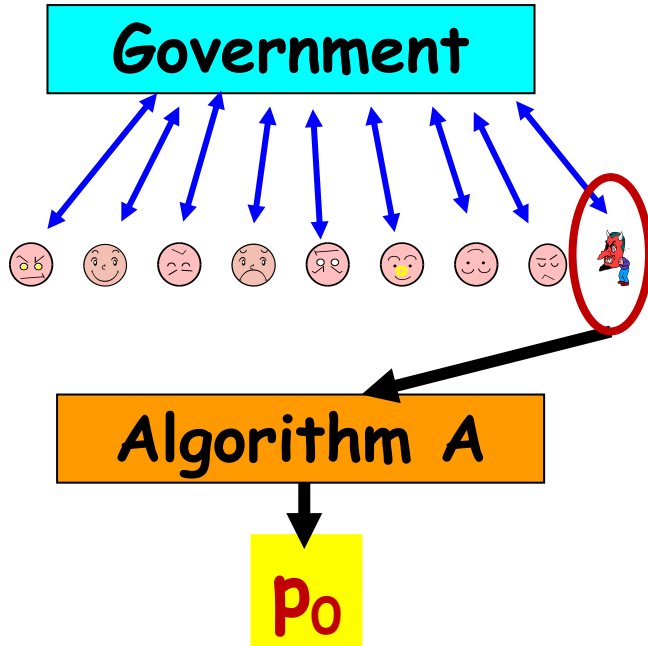
(e.g.) Secure elections



Definition: Protocol is secure if for every efficient algorithm, and every $P_0 \not\approx P_1$ player

[Yao'86, Goldreich
-Micali-W '87]

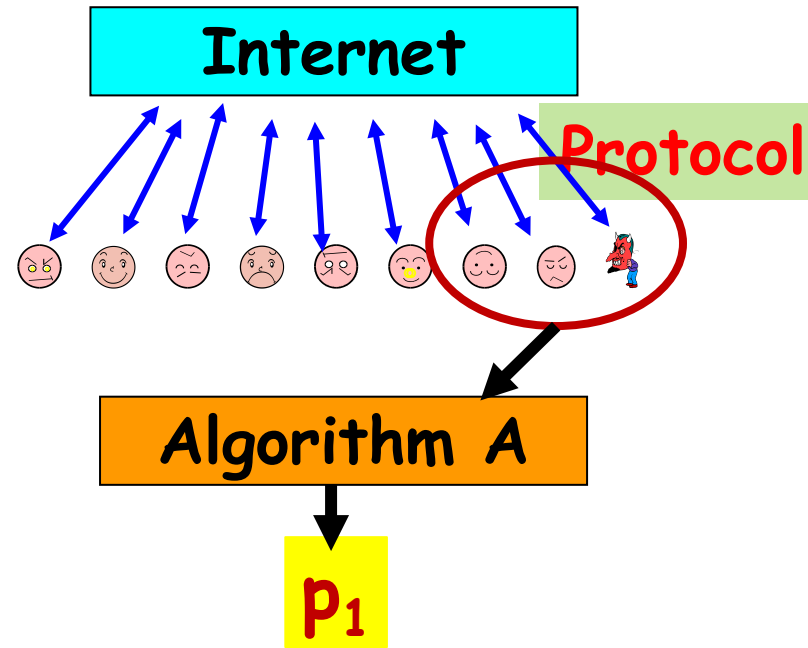
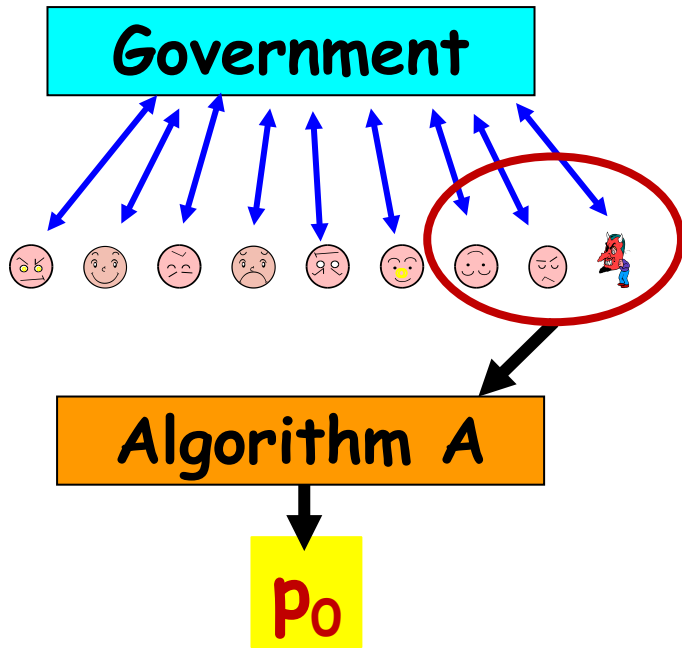
Secure elections



Definition: Protocol is secure if for every efficient algorithm, and every $P_0 \approx P_1$ player

[Yao'86, Goldreich
-Micali-W '87]

Secure elections



Definition: Protocol is secure if for every efficient algorithm, and every subset P_0 & P_1 ers,
Theory: Crypto! Zero-Knowledge, Secure protocol design for any problem, under simple assumptions
Practice: Internet security, shopping,

Randomness

The amazing utility of randomness

Nature seems to supply us perfect randomness



Unbiased, independent bits

which we use for numerous applications

- Sampling
- Scientific experiments & simulations
- Probabilistic algorithms
- Cryptography
- Game theory
- Gambling
- ...

Seemingly much faster than deterministic ones for many problems

Where are the random bits from?

this power real?

Is

Universe has:


- Perfect randomness



Applications

Everything!

- Weak random sources



Biased, dependent bits

??



- No randomness

??

[Blum-Micali '81
Yao'82]

Computational pseudo-randomness



P_0



P_1

Definition:  is pseudo-random

if for every efficient algorithm, $P_0 \approx P_1$

Theory: Composition, Amplification,

Hardness vs. Randomness paradigm


Practice: Pseudo-random generators, Crypto

Universe has:

- Perfect randomness



- Weak random sources



Biased, dependent bits



Major theories

Unexpected benefits

- No randomness

Assuming "P≠NP": there are hard problems

Applications

Everything!

all prob algorithms

Extractor theory

[B,SV,NZ,T,...,GUV,DW,...]

purifying randomness

all prob algorithms

Hardness vs. Randomness

[BM,Y,...NW,IW,...]

Every fast prob alg has a deterministic counterpart




Structural pseudo-randomness (discrete math, number theory,...)

[Ramsey '30] *"Every large enough system must have some structure"*



Periodic sequences



Which subsets contains long periodic sequences?

[Szemerédi '75] All "dense" subsets do!


[Green-Tao '04] The primes numbers do!

[Easy] Random subsets do!

Which subsets "look" random?

[Szemerédi '75] Every "dense" network does!!

Networks, Clusters, Bonds

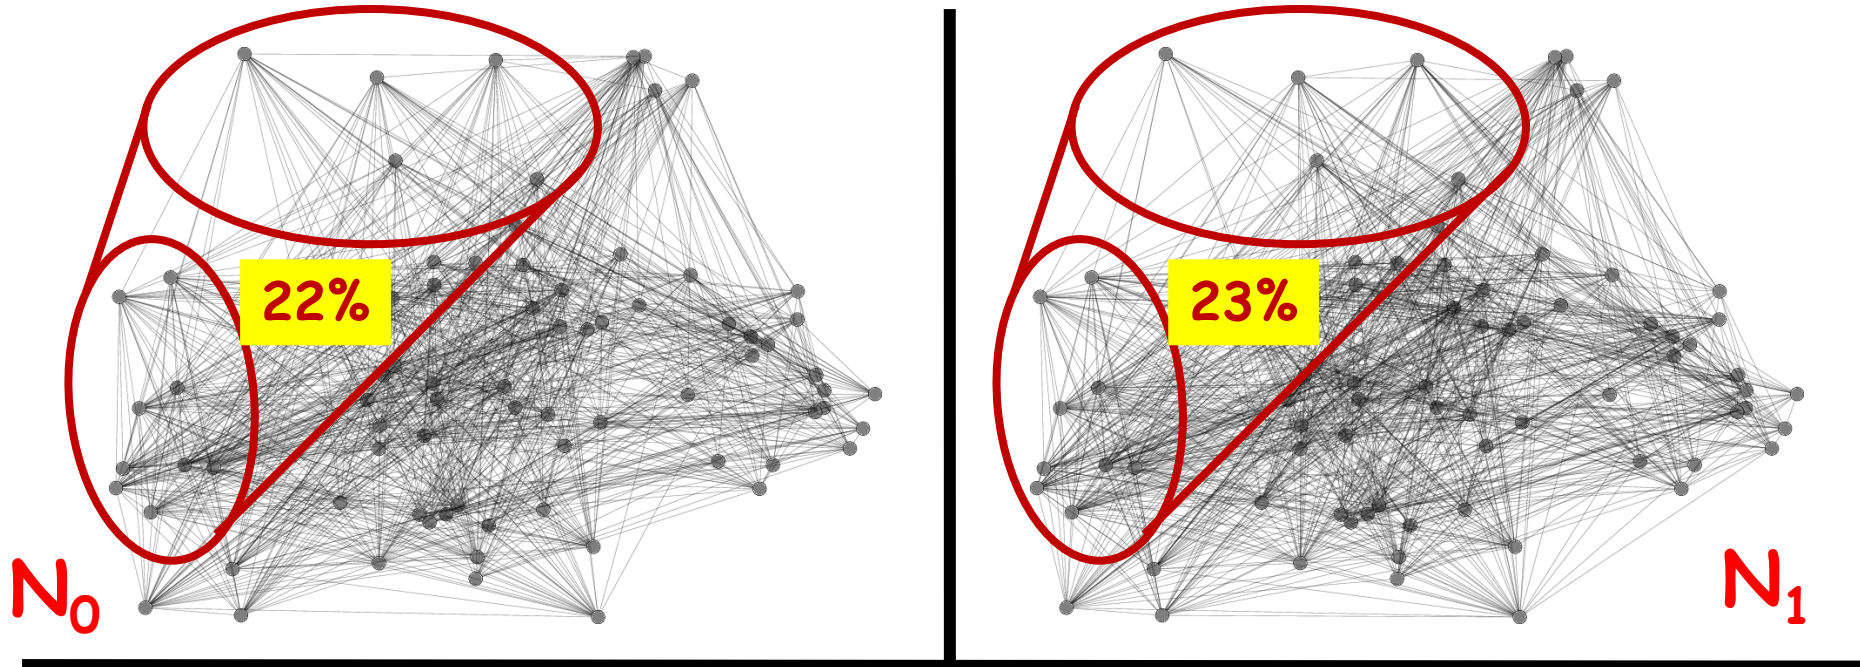


p = bond: fraction of actual connections between clusters, out of all possible

Roads, Internet, Facebook, ...

[Szemerédi'75]

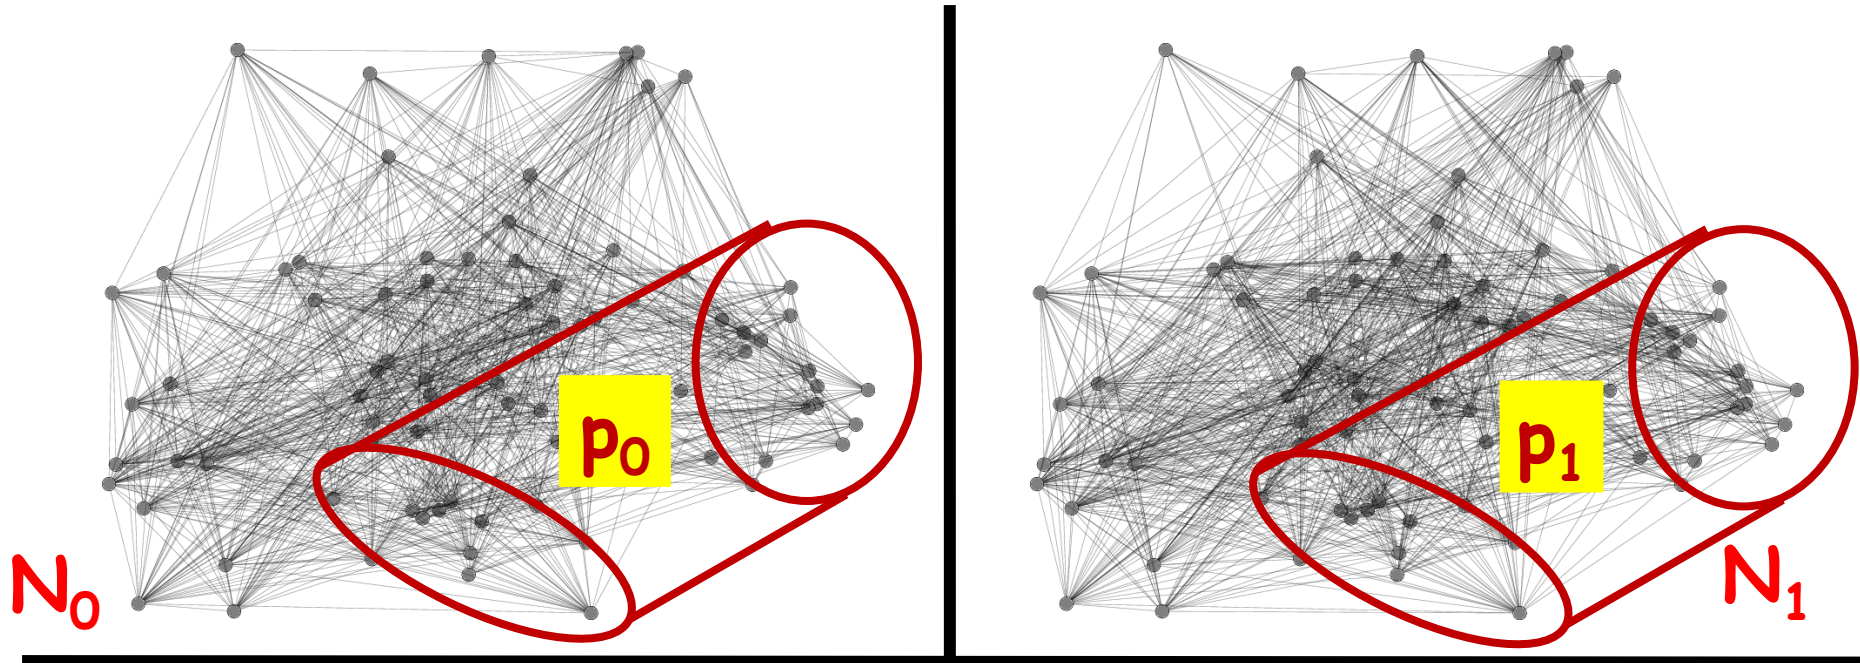
Bond-similarity of networks



Definition: N_0 & N_1 are **bond-similar** if **every** two communities are equally bonded $p_0 \approx p_1$

[Szemerédi'75]


Bond-similarity of networks



Definition: N_0 & N_1 are **bond-similar** if **every** two communities are equally bonded $p_0 \approx p_1$

[Szemerédi'75]

Regularity



N_0


Szemerédi's Regularity Lemma:
Every network is bond-similar to a **random** network

Proof: *Every* network has a **tiny** model

[Szemerédi'75]

Regularity

	●	●	●
●	0.8	0.1	0.4
●	0.1	0.3	0.3
●	0.4	0.3	0.2



N_0


Szemerédi's Regularity Lemma:
Every network is bond-similar to a **random** network

Proof: *Every* network has a **tiny** model

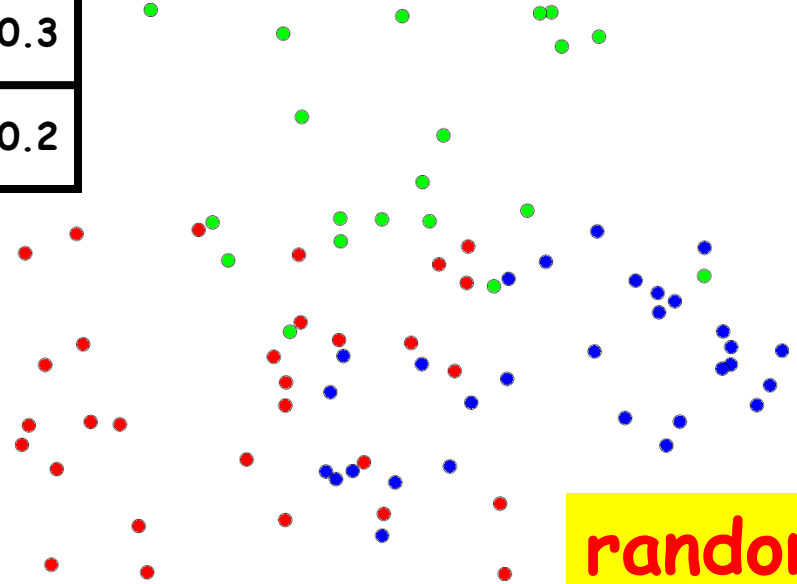
[Szemerédi'75]

Regularity

	●	●	●
●	0.8	0.1	0.4
●	0.1	0.3	0.3
●	0.4	0.3	0.2



N_0



random



Szemerédi's Regularity Lemma:
Every network is bond-similar to a **random** network

Proof: *Every* network has a **tiny** model

[Szemerédi'75]

Regularity

	●	●	●
●	0.8	0.1	0.4
●	0.1	0.3	0.3
●	0.4	0.3	0.2



Szemerédi's Regularity Lemma:
Every network is bond-similar to a **random** network

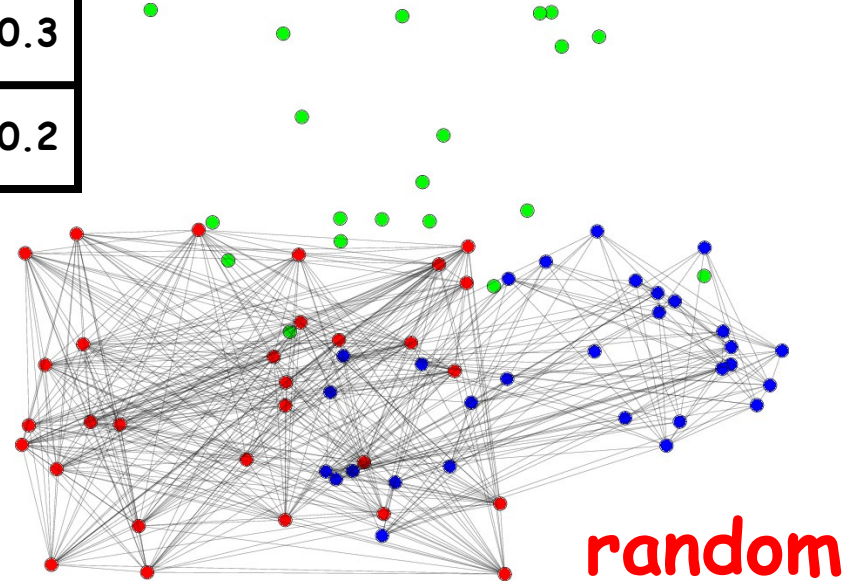

Proof: *Every* network has a **tiny** model

[Szemerédi'75]

Regularity

● ● ●

●	0.8	0.1	0.4
●	0.1	0.3	0.3
●	0.4	0.3	0.2



Szemerédi's Regularity Lemma:
Every network is bond-similar to a **random** network



Proof: *Every* network has a **tiny** model

[Szemerédi'75]

Regularity

● ● ●

●	0.8	0.1	0.4
●	0.1	0.3	0.3
●	0.4	0.3	0.2



Szemerédi's Regularity Lemma:
Every network is bond-similar to a **random** network



Proof: *Every* network has a **tiny** model

[Szemerédi'75]

Regularity

● ● ●

●	0.8	0.1	0.4
●	0.1	0.3	0.3
●	0.4	0.3	0.2



Szemerédi's Regularity Lemma:
Every network is bond-similar to a **random** network



Proof: *Every* network has a **tiny** model

[Szemerédi'75]

Regularity

● ● ●

●	0.8	0.1	0.4
●	0.1	0.3	0.3
●	0.4	0.3	0.2



Szemerédi's Regularity Lemma:
Every network is bond-similar to a **random** network



Proof: *Every* network has a **tiny** model

[Szemerédi'75]

Regularity

● ● ●

●	0.8	0.1	0.4
●	0.1	0.3	0.3
●	0.4	0.3	0.2



Szemerédi's Regularity Lemma:

Every network is bond-similar to a **random** network

Proof: **Every** network has a **tiny** model

Regularity Galore

Dense graphs, sequences “look” random

True for other objects, any set of tests!

[Szemerédi, Thomasson, Chung-Graham-Wilson,
Green-Tao-Ziegler, Gowers] Transference
principle

[Impagliazzo, Reingold-Trevisan-Tulisani-
Vadhan] Dense model theorem

Structure vs. Randomness dichotomy

- Discrete Math
- Number Theory
- Complexity Theory
-
- PDEs
- Ergodic Theory
- Analysis

Privacy

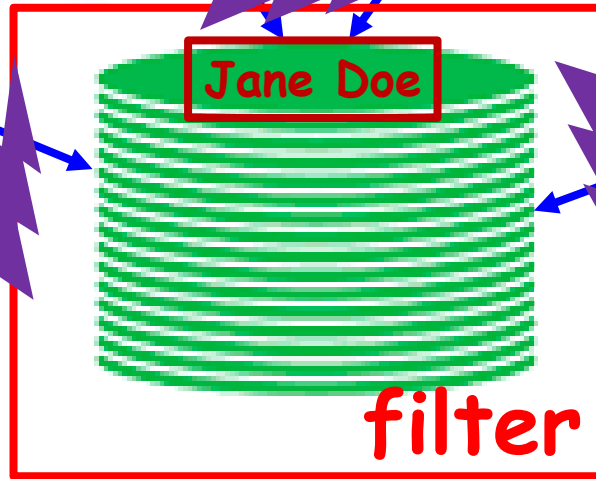

Privacy

versus

Utility

Why should I
- Participate?
- Tell the truth?

Welfare gaps
Urban planning
Drug/treatment



Data
analysts




Database D: Many
individual records, e.g.
census/medical data

- Restrict
queries
- Filter
answers




Differential Privacy




Definition: A **filter** is **differentially private** if for **every** adjacent D_0 & D_1 and **every** legal query Q , the outputs P_0 and P_1 are indistinguishable to an adversary.

Theory: Composition, many queries, higher accuracy, resistance to future knowledge

Adaptive data analysis



sample



Data expensive!

filter



Science/ML

Adaptively:

- Tune parameters
- Form hypotheses
- Generate queries

- Data reuse
- Overfitting
- No generalization

Crisis: "Why most published research findings are wrong"
Ioannidis'05 PLoSMed

Theorem: Using a differentially private filter can prevent overfitting, ensure statistical validity and generalization, despite data reuse

Summary

Things are identical unless you can tell them apart (eg diamonds, Picassos, news... real/fake)

Behavioral, "subjective" definitions (like Imitation Games) are extremely powerful. May be useful in philosophical and operational understanding other fundamental notions.


ToC is a modeling science, essentially tuned to the underlying process and its resources. Definitions of central notions are primary -

Book ad

MATHEMATICS + COMPUTATION

A THEORY REVOLUTIONIZING
TECHNOLOGY AND SCIENCE

Avi Wigderson



- Published by Princeton University Press
- Free (forever) on my website
- Comments welcome!