Course Introduction

About me

- My name is Radoslav Ivanov
  - Call me Rado
- Undergrad degree in CS and ECON from Colgate in 2011
- Got my PhD in CS from UPenn in 2017
- My research is on safe and secure autonomous systems
  - Verification of neural networks
  - Attack-resilient sensor fusion
  - Context-aware detection and estimation
- Started at RPI in Jan. 2022

Impressive Progress in Autonomy

Rensselaer

Control

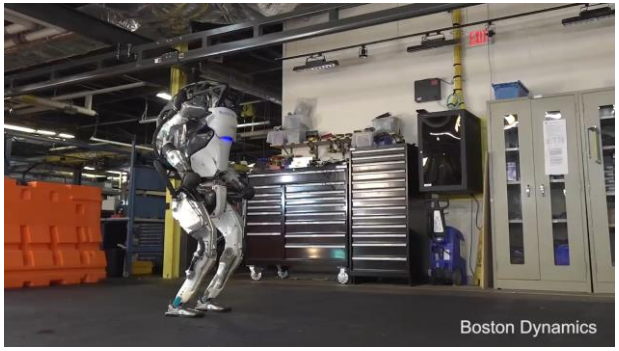

Boston Dynamics

Perception

YOLO v. 3

Learning

Kormushev, Calinon, Caldwell, IROS'10

JPL-Caltech, DARPA Robotics Challenge

Zhu, Zhou, Daniilidis, ICCV'15

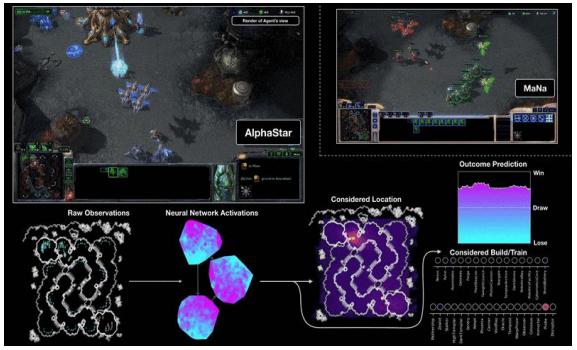

DeepMind

But we're not there yet…

Rensselaer

Tesla Driver Was on Autopilot Eating a Bagel When He Smashed into a Fire Truck

A National Transportation Safety Board investigation found the driver had hands off the wheel and ignored warnings in the 2018 crash.

By Clifford Atiyeh SEP 4, 2019

- The driver of a 2014 Tesla Model S that ran into the back of a fire engine in California in 2018 was using Autopilot at the time, according to a National Transportation Safety Board (NTSB) report this week.
- The agency's investigators reported that the driver was having breakfast while he let Autopilot take over the driving; his hands were not on the steering wheel, and he did not brake prior to the crash.

Uber self-driving car involved in fatal crash couldn't detect jaywalkers

The system had several serious software flaws, the NTSB said.

Steve Dent, @stevetdent 11.06.19 in Transportation 32 Comments 2136 Shares

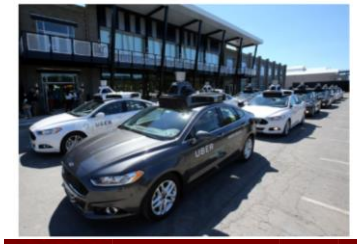

Sponsored Links

QuickBooks® - Get Pair QuickBooks® Online

Biotech Sector Action A Today

ym: Australian open 2020 - T | Дискусия на живо: Аустралиън | Live Scores | Australian Open

320936/waymo-self-driving-car-crash-arizona

GOOGLE | TECH | TRANSPORTATION

Waymo self-driving minivan involved in crash in Arizona

Minor injuries reported

By Andrew J. Hawkins | @andyjayhawk | May 4, 2018, 4:52pm EDT

SHARE

Boeing 737 Max Lion Air crash caused by series of failures

25 October 2019 Share

Lion Air plane crash

The Washington Post | PostTV | Politics | Opinions | Local | Sports | Natio

National Security

Home > Collections > Surveillance

Iran says it downed U.S. stealth drone; I acknowledges aircraft downing

By Greg Jaffe and Thomas Erdbrink, December 04, 2011

A secret U.S. surveillance drone that went missing last week in western Afgha have crashed in Iran, in what may be the first case of such an aircraft ending u an adversary.

Iran's news agencies asserted that the nation's defense forces brought down t Iranian reports said was an RQ-170 stealth aircraft. It is designed to penetrate defenses that could see and possibly shoot down less-sophisticated Predator a

A stealthy RQ-170 drone played a critical role in surveilling the compound in Osama bin Laden was hiding in the months before the raid in which he was ki SEALs in May.

U.S. officials acknowledged Sunday that a drone had been lost near the Irania declined to say what kind of aircraft was missing.

The Iranian government has not released any pictures of the recovered aircraf was downed by defense forces after it flew across the border and into the cour unnamed Iranian defense official said in one report that a cyberattack caused

U.S. officials cast doubt on th Iranian assertions. "We have no indication that it was brought down by hostil senior Pentagon official, speaking on the condition of anonymity to discuss se activity.

BBC NEWS | News | Sport | Weather | Capital | Culture | Autos

NEWS TECHNOLOGY

Home | US & Canada | Latin America | UK | Africa | Asia | Europe | Mid-East | Business | Health | Sci/Environme

25 July 2013 Last updated at 19:04 ET

Share

Car hackers use laptop to control standard car

By Zoe Kleinman

Technology reporter, BBC News

Brian Donohue Follow @TheBrianDonohue July 30, 2013 , 3:26 pm

A 213-foot luxury yacht veered off course while cruising in the Mediterranean Sea this

Neural Network (NN) Vulnerabilities

- Neural networks increasingly used in safety-critical systems

Perception (autonomous cars)

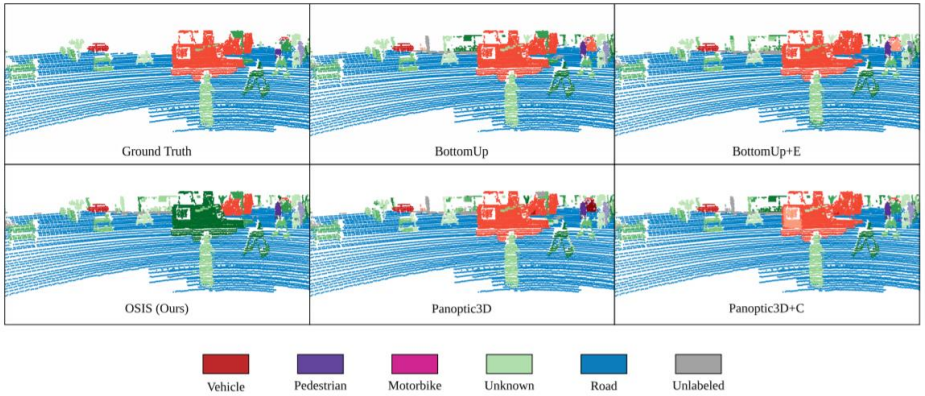

Wong et al., CoRL'19

Control (air traffic avoidance)

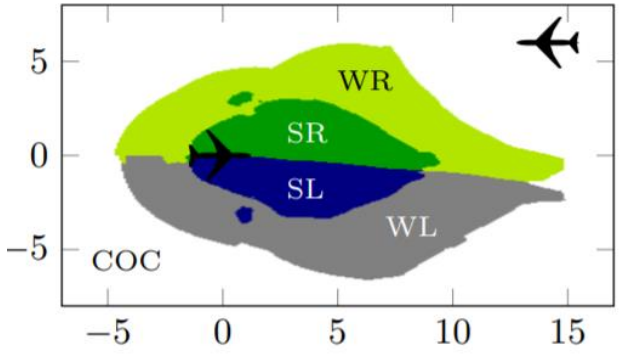

Katz et al., CAV '17

- Safety concerns discovered in both domains

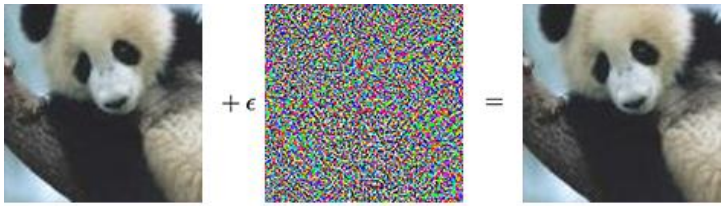

Goofellow et al., ICRL'15

Table 2: Verifying properties of the ACAS Xu networks.

	Networks	Result	Time	Stack	Splits
ϕ_1	41	UNSAT	394517	47	1522384
	4	TIMEOUT			
ϕ_2	1	UNSAT	463	55	88388
	35	SAT	82419	44	284515
ϕ_3	42	UNSAT	28156	22	52080
ϕ_4	42	UNSAT	12475	21	23940

Cyber-Physical Systems (CPS)

Rensselaer

Tight coupling between **communication**, **computation** and interaction with the **physical world**

Aircraft

Autonomous Cars

Medical CPS

Military

Smart Grids

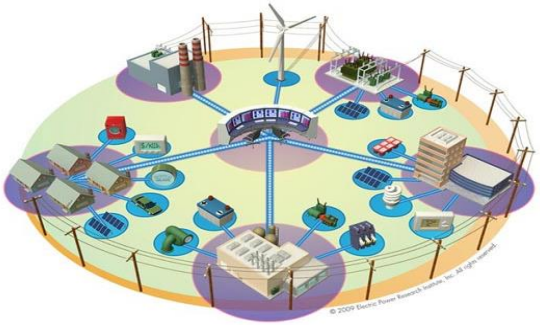

Robotics

A standard CPS design

F1/10 Autonomous Racing Competition, ES Week 2016

Problem: How do we know car won't crash?

- How do we build safe algorithms?
- How do we analyze algorithms?
- What about "black-box" components such as neural networks?
- How do we convince other people car is safe (assurance argument)?

Why is safe autonomy so hard?

Topics

Rensselaer

- Discrete objects.
- Reasoning about discrete objects
- Counting discrete objects
- Randomness: probability
- What can we compute?
- What can we compute efficiently?

our language will be mathematics . . .
. . . it will be everywhere

Course Mechanics

- Meeting time: MTh 10am-noon
  - Each lecture will last around 75-90 minutes
  - We will meet in Sage 3303
- Recitation sessions: W 10am-4pm (depending on your section)
  - Sessions 10am-3pm will meet in Aemos Eaton 215
  - Session 3-4pm will meet in Lowe 3039
  - Led by TAs
- Instructor office hours: M 1-2pm, W 4-5pm, F 9-10am
  - Lally 309
  - Office hours will be in person unless noted otherwise
  - **Office hour on Wed, Jan. 10: 3-4pm**

Course Mechanics, cont'd

Rensselaer

- We will have 3 quizzes
  - Each quiz will be an hour long, multiple choice only
  - Dates are (all at 8am during the Wednesday slot):
    - Feb. 7
    - Mar. 27
    - Apr. 17
- Midterm exam (at 8am during the Wednesday slot)
  - two-hour exam on Feb. 28
- Final exam
  - Scheduled by the Registrar
- Locations TBD

Course Mechanics, cont'd

Rensselaer

- TAs: Michael Cleversley (clevem), Sharmishtha Dutta (duttas), Shuhang Tan (tans5), Yuchen Zhang (zhangy94)
  - Office hours: TBD
  - Will be monitoring Piazza and grading
- Mentors: Wilde Chu (chuw7), Mohamed Lashuel (lashum), Michael Lyga (lygam), David Wang (wangd14), Ilan Beyen (beyeni), Emma Huntington (huntie), Matthew Hurtado (hurtam), Jun Kim (kimj43), Zain Magdon-Ismail (magdoz2), Fatih Orhan (orhanf), Justin Ottesen (ottesj), Shimu Pan (pans), Eric Scheer (scheee2)
  - Office hours: by appointment
  - Will be helping with marking

Course Mechanics, cont'd

Rensselaer

- Course website: http://cs.rpi.edu/~ivanor/focs/S24/focs.html
- All course information will be posted on the website
  - lecture slides
  - homework assignments (already posted)
  - practice exams
- Homework assignments and submissions will be through Submitty (submitty.cs.rpi.edu)

Course Mechanics: Piazza

- We will be using Piazza for questions and discussions
  - Signup link: https://piazza.com/rpi/spring2024/csci2200/home
- Please request access if you are not enrolled already
- Please use Piazza for questions and discussion
  - I won't monitor Submitty/Webex for questions

Book

- Discrete Mathematics and Computing, M. Magdon-Ismail
  - Very well written book
  - We will try to cover 1 chapter per lecture
- The book is required
  - Won't be able to cover all examples/exercises from the book
  - You will need to read each chapter and work out the exercises on your own
  - Homework problems will be assigned from the book also
  - Problems also posted on course website
    - If you have an older version of the book, the problem numbers may be different!
    - Prof. Magdon-Ismail will exchange a new copy of the book for an old one if you got the older version

Course Mechanics, cont'd

Rensselaer

- This is a theory course
  - Despite its current trajectory, CS is still based in math
  - Need to understand algorithms
- Homeworks will be problem sets from the book
  - Submit through Submitty
    - Please make sure you have access now
    - Course page was just set up this morning so it may be online yet
  - There will be 10 homeworks total

Grading

- Homework: 10% (1% each)
- 3 Quizzes: 30% (10% each)
- Midterm: 30%
- Final: 30%
- In-class pop quizzes (roughly one per week): 2% bonus
  - Please bring your laptops to class so you can submit these
  - Will get full 2% if you do well on 60% of pop quizzes
  - Will get 1% if you do well on 40% of pop quizzes
  - If you do the pop-quizzes, you'll make up points for 2 homeworks!

Homework Grading

- We will use the following homework grading scheme:
  - Correct solution (MUST show your work): 100%
  - You made significant progress (MUST show your work): 80%
  - You understood the problem and made a reasonable attempt but something went wrong: 50%
  - Not in the right ball park or serious error: 0%

Prerequisites

Rensselaer

- CS II (data structures)
- Calc I
- Calc II is strongly recommended!

Background

Rensselaer

- Programming, numbers, geometry, algebra, calculus, . . .
- What is the minimum element in the set {8, 9, 3, 10, 19}?

  3
- Does this set of *positive* numbers have a minimum element:

  {25, 97, 107, 100, 18, 33, 99, 27, 2014, 2200, 23, …}
  - Any (non-empty) set containing only positive integers has a minimum element.
  - Well-ordering principle
- Can you build an algorithm to find it?
  - Suppose after first 1M numbers the min is 10
  - When do you declare you have found the minimum?
  - Any algorithm you come up with may not always terminate

Two-Contact Covid on a Grid

- Need to model the spread process
- A square gets infected if two or more neighbors are infected
- Given initial gray infections, who ultimately gets infected?
- Minimum infections to infect everyone?
- Given few vaccines, who to immunize?
- **Answers involve discrete math.**

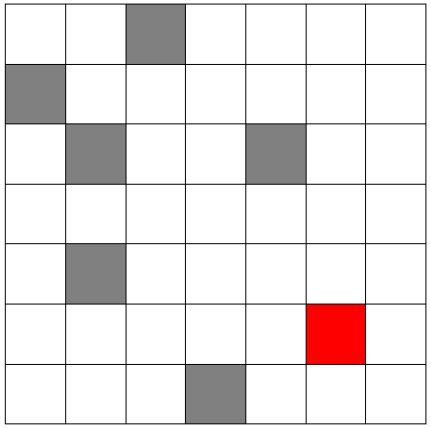

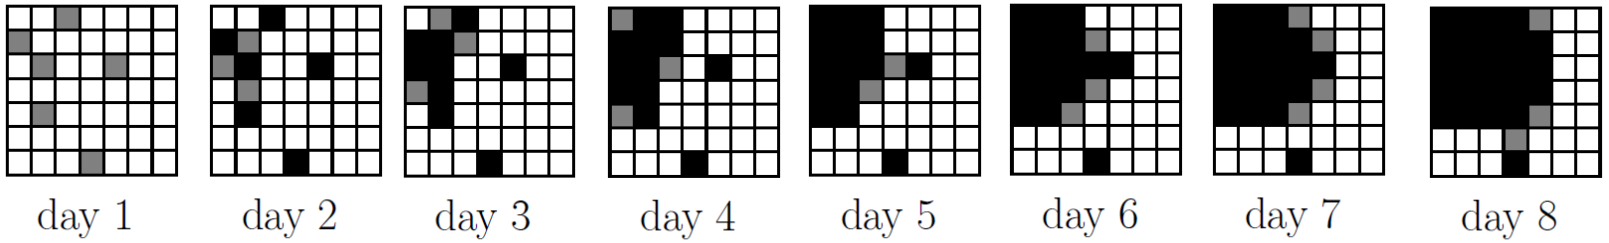

Two-Contact Covid on a Grid

- Given initial gray infections, who ultimately gets infected?
  - Last two columns don't get infected
- Minimum infections to infect everyone?
  - Turns out 6 isn't enough
- Given few vaccines, who to immunize?
  - Look at first few susceptible people

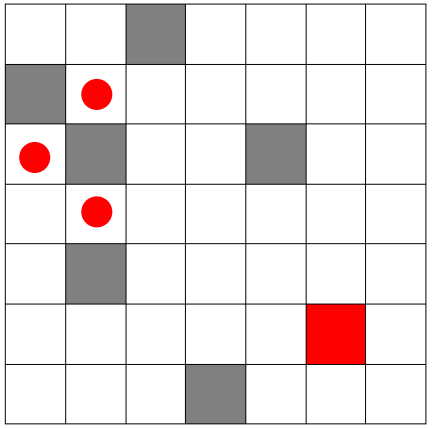

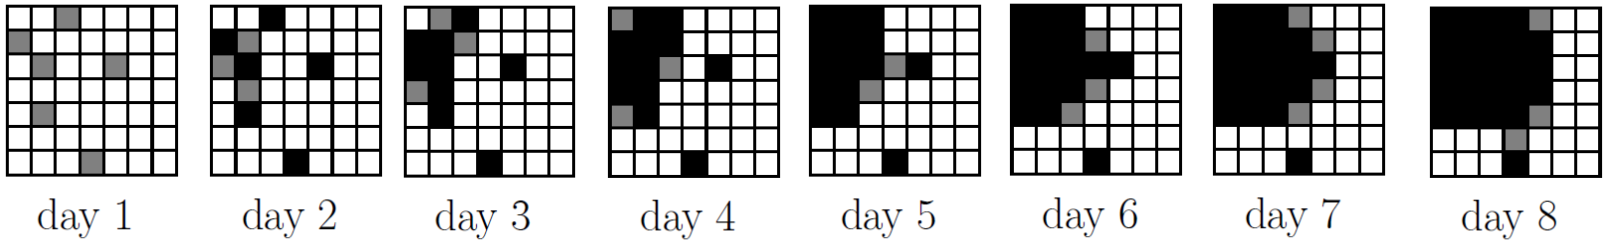

Scheduling Speed Dates

- In each round 4 people “group”-speed-date around a table. (4 rounds in all)

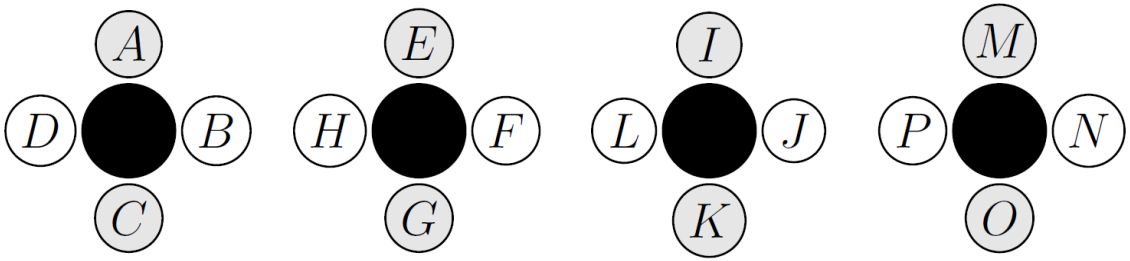

- How to organize the rounds so that people meet as many people as possible?
- Do you care about average or minimum number of meetups per person?
- Can everyone meet at least 10 people?
- What happens if you assign tables randomly?
- **Answers involve discrete math.**

Friendship Networks and Ads

- People are circles and links are friendships.

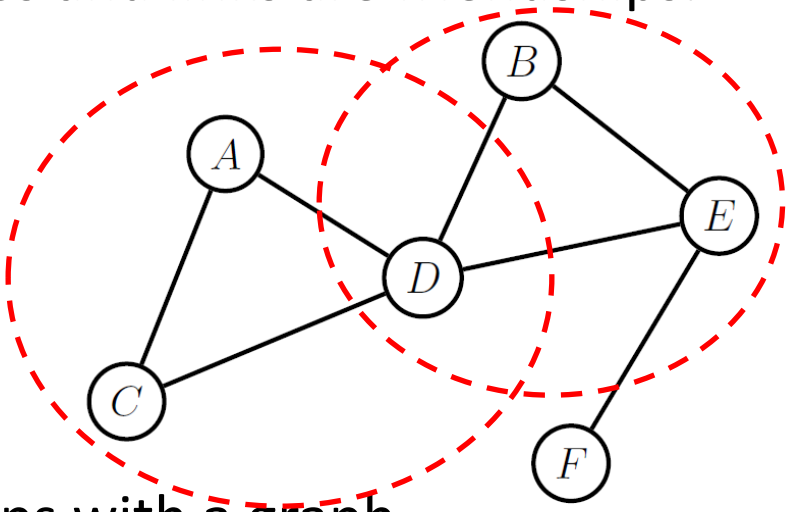

Cliques of size 3

- Model friendships with a graph
- Who would you advertise to? You wish to maximize adoption of your new technology.
- Can you find all cliques of certain size?
- **Answers involve discrete math.**

Modeling Computers

Rensselaer

- Desktop, smartphone, fitbits, smart watches
- We have deep questions:
- What can we compute?
- What *can't* we compute?
  - Think about finding the smallest number in an infinite set
  - Turns out it cannot be computed! Ever!
    - In other words, for any algorithm, I can give you a set for which your algorithm won't ever terminate
- How fast?
- **Answers involve discrete math.**

Post's Correspondence Problem (PCP)

Rensselaer

- **PCP**: Consider 3 dominos

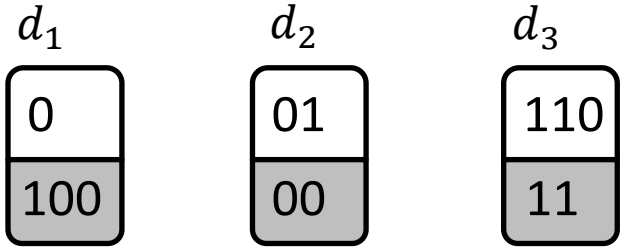

- Can I arrange dominos (using multiple copies of each) so that top and bottom strings match

$$d_3 d_2 d_3 d_1 =$$

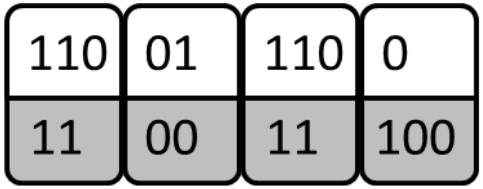

- INPUT: Dominos $\{d_1, d_2, \ldots, d_n\}$.
- **Challenge ($1000 reward)**: write a program that takes a txt file of dominos and **ALWAYS TERMINATES** with the correct answer
  - Correct sequence of dominoes or says it can't be done!
  - Prove your program is correct!

Matching Subset Sums

Rensselaer

- Consider the set of numbers:

$$\{3,7,9,5,4,10,13\}$$

- What are all the subsets that sum up to 22?

$$\{10,7,5\}, \{9,13\}, \{3,9,10\}, \{3,5,4,10\}, \ldots$$

  - Quite a few
  - Enumerating all subsets is hard (i.e., slow)

- Are there distinct subsets that sum up to the same thing?

$$\{10,7,5\}, \{9,13\}$$

- If there are none, can you **prove** it?

Why do we need proofs?

Rensselaer

- Computers used in safety-critical systems
  - Self-driving cars, medical systems
- Need to prove algorithms and their implementations are correct!
  - Some things we can prove in the worst case
    - E.g., find all people who get infected
  - Sometimes all we can give are probabilistic guarantees
    - World is probabilistic anyway
- In this class, we'll look at proof techniques for both cases

Computing is Mathematics

Rensselaer

- "Too few people recognize that the high technology so celebrated today is essentially a mathematical technology."
- "A programmer must *demonstrate* that their program has the required properties. If this comes as an afterthought, it is all but certain that they won't be able to meet this obligation. Only if this obligation influences the design is there hope to meet it…
- "The required techniques of effective reasoning are pretty formal, but as long as programming is done by people who don't master them, the software crisis will remain with us and will be considered an incurable disease. And you know what incurable diseases do: they invite the quacks and charlatans in, who in this case take the form of Software Engineering Gurus."

  – Edsger Dijkstra

- Rado (Dijkstra): Turns out coding is very difficult (especially in the case of large systems), but I mostly agree

Polya's Mouse

- "A mouse tries to escape from an old fashioned cage. After many futile attempts bouncing back-and-forth, thumping his body against the cage bars, he finally finds one place where the bars are *slightly* wider apart. The mouse, bruised and battered escapes through this small opening, and to his elation, finds freedom." – Polya

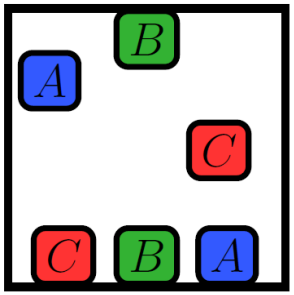

- Connect tiles of the same letter with wires. Wires cannot cross, enter tiles, or leave the box. How can it be done? If it can't be done, why not?
- Don't be quick to dismiss either conclusion. Try this and that. Fiddle around until you understand the problem and the difficulty. Patience.
- To solve such problems, "*You need brains and good luck. But, you must also sit tight and wait til l you get a bright idea.*" – Polya.

Getting Good at Discrete Math

Rensselaer

- The *professional*'s workflow in addressing a discrete math problem:

1. Model the problem your are trying to solve using a discrete mathematical object.
2. Tinker with easy cases to build an understanding of the model.
3. Based on the tinkering, formulate a conjecture about your problem/model.
4. Prove the conjecture and make it a theorem. You now *know* something new.

Acknowledgements

- Slides are based on lectures developed by Prof. Malik Magdon-Ismail (RPI)