

# AI Uncovered: Details Matter

Malik Magdon-Ismail

Professor, Computer Science, Rensselaer Polytechnic Institute

Who am I?

What I do?

Why RPI?



Create a picture with researchers studying disease spread using data and machine learning with fancy visualizations and satellite data.

# What is AI?

---

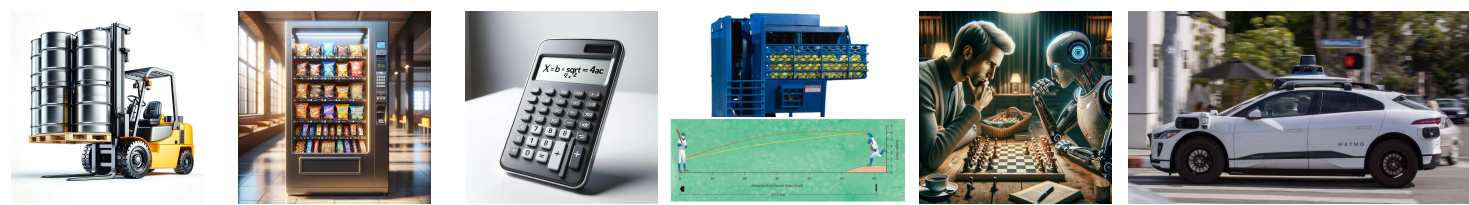
- **Perform**

- **Create**

- **Predict**

# What is AI?

- Perform

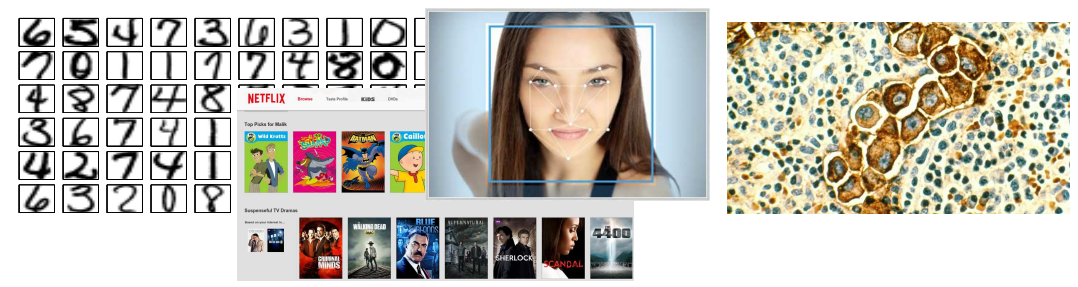


- Create

In New York City, Malik took the floor,  
Discussing AI, a subject he adores.  
With wit and insight quite nifty,  
He enlightened the crowd on AI so **swiftly**.

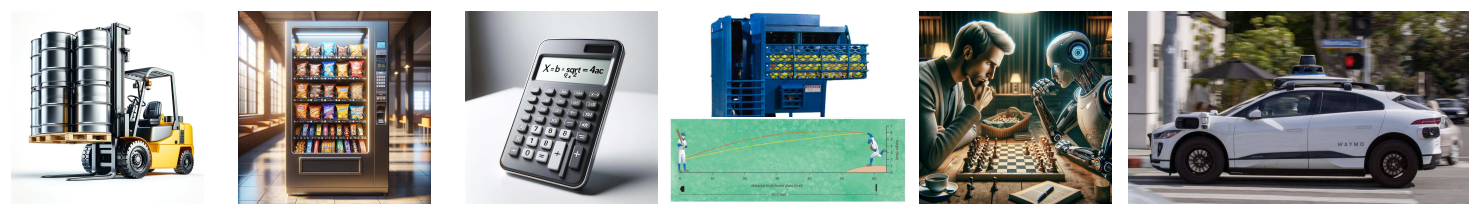


- Predict



# What is AI?

- Perform

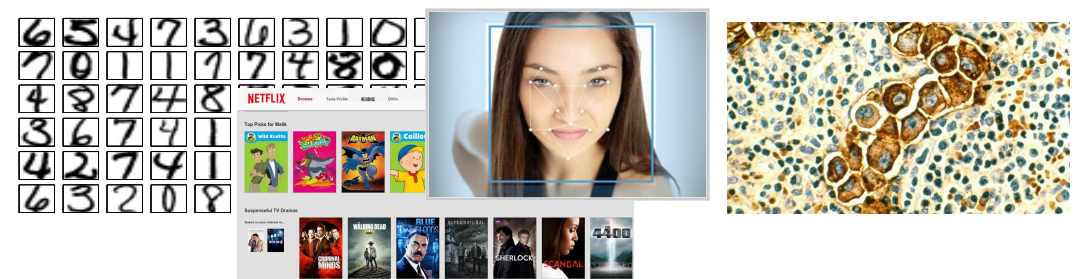


- Create

In New York City, Malik took the floor,  
Discussing AI, a subject he adores.  
With wit and insight quite nifty,  
He enlightened the crowd on AI so **swiftly**.



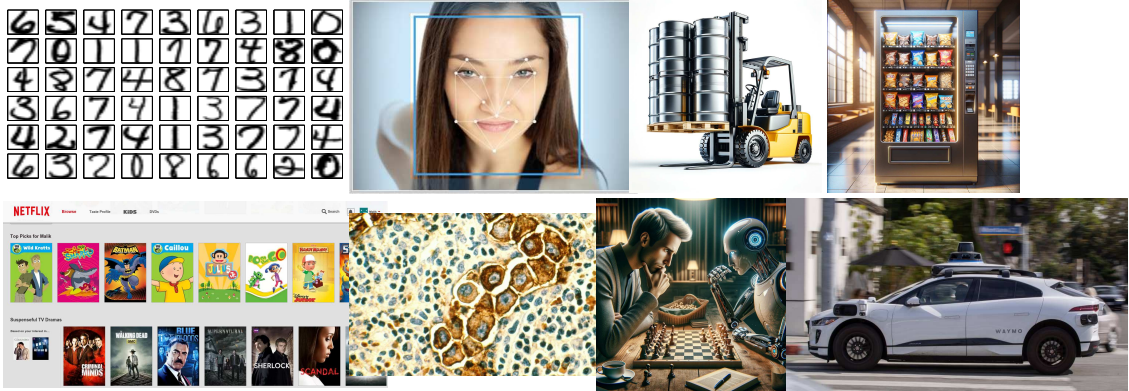
- Predict



Machines doing stuff ~~humans~~ <sup>biology</sup> can?

# Wants from AI

- Competence



- Alignment

helpful  
honest  
harmless

**You**  
create an image with mario and luigi from super mario brothers

**ChatGPT**  
I'm unable to create images of copyrighted characters like Mario and Luigi from Super Mario Brothers. However, if you have another request or a different idea for an image, I'd be happy to help with that. Let me know if there's anything else you'd like!



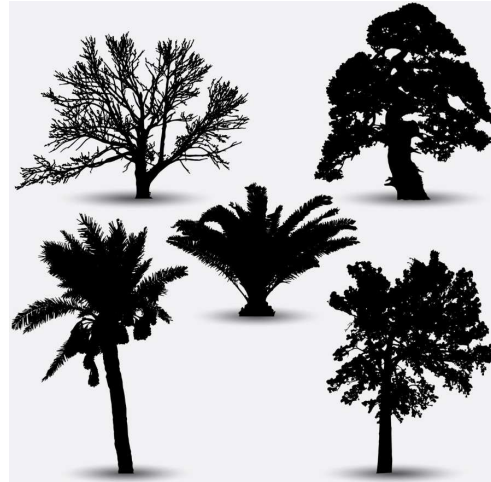
# Machine Learning

---

Apqq's



Bqpp's



Training

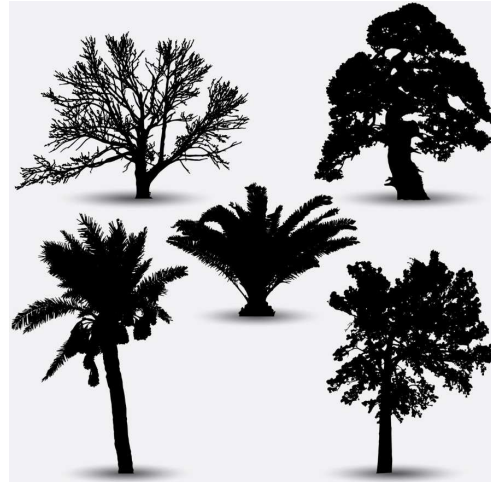
# Machine Learning

---

Apqq's



Bqpp's



Training



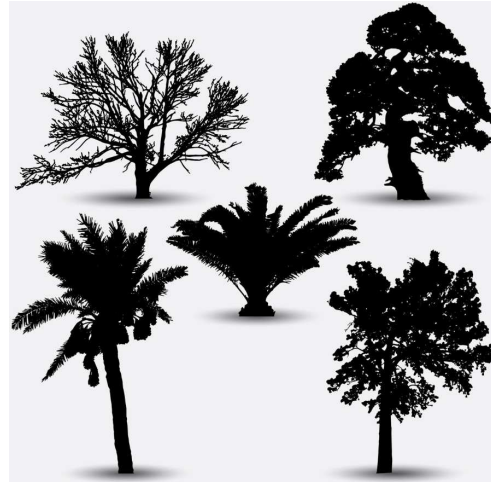
Test

# Machine Learning

Apqq's



Bqpp's



## Training



- You're using the invented words "apqqs" and "bqpps", and from context:
- The **first image** (with dogs) is labeled "apqqs" → likely means **dogs**
  - The **second image** (with trees) is labeled "bqpps" → likely means **trees**

### Third image:

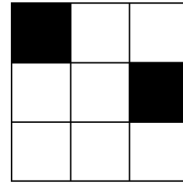
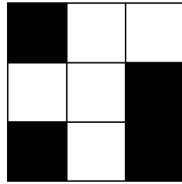
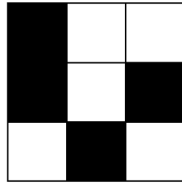
- Shows a **fluffy Chow Chow dog** standing on grass.
- It has fur, a tail, four legs, and a dog-like face.
- It does **not** resemble any part of a tree (no trunk, branches, or leaves).

### ✓ Conclusion:

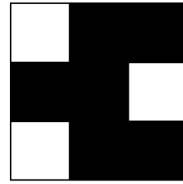
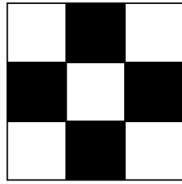
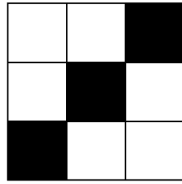
The third image is an **apqq** (i.e., **dog**), not a **bqpp** (tree).

# Machine Learning

---



Dogs

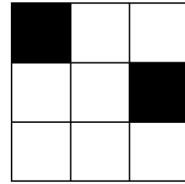
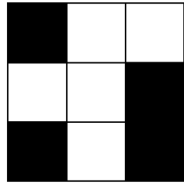
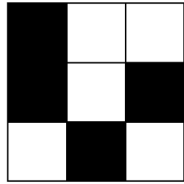


Trees

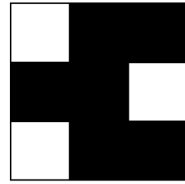
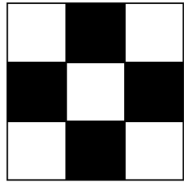
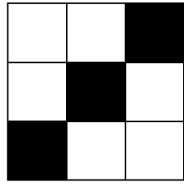
---

# Machine Learning

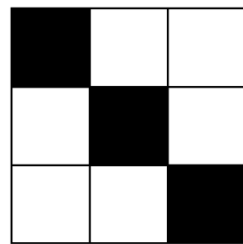
---



Dogs



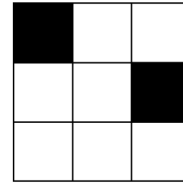
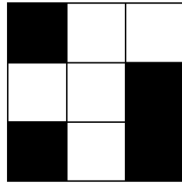
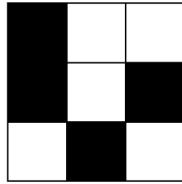
Trees



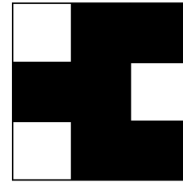
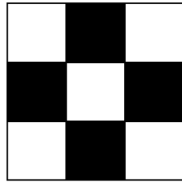
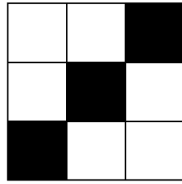
Dog or Tree?

# Machine Learning

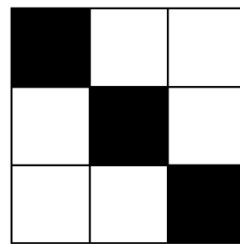
---



Dogs



Trees



**Answer: Tree**

**Explanation:** The image lacks the regular, alternating pattern seen in the dog images and instead shows a clustering of black tiles characteristic of the tree images.

motto: **There's a pattern. We don't know it. We have data to learn it.**

# A Framework for Using AI to Solve a Task

**margin for error**  
(less complex)

**details matter**  
(more complex)

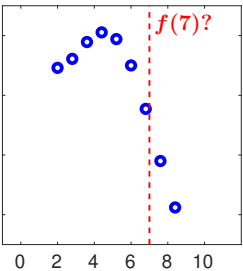
**in-sample**  
(interpolation)

**info/concept retrieval**

FOCS HW: retrieve concept  
Chicken marsala: 400-600 kcal

**quantitative tool**

FOCS HW: use right formula  
Meal Plan: lose 2lbs in 2 weeks



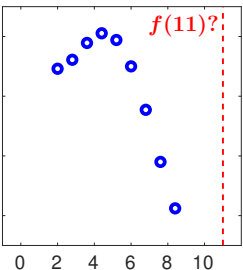
**out-of-sample**  
(extrapolation)

**reason: known to unknown**

FOCS Test: apply formula, new problem  
Evaluate a new setting

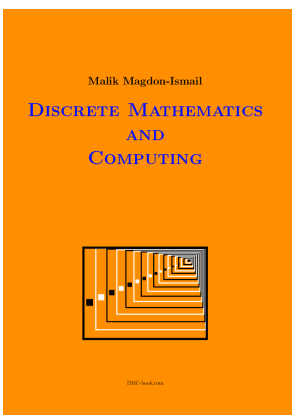
**new setting, new tool**

FOCS Test: derive new formula  
True novelty



# Competence: ChatGPT LRM 4o Aces FOCS Tests

	FOCS		
	Quiz 1	Quiz 2	Midterm
ChatGPT	100%	100%	94%
Students	77%	80%	75%



How many subsets of  $\{1, 2, 3, 4, 5, 6\}$  contain **some consecutive numbers**?

---

**Step-by-step plan:**

- The total number of **subsets** of a 6-element set is:
 
$$2^6 = 64$$

This includes **all** subsets, including the empty set and singletons.

---

**Step 1: Define the opposite (easier to count)**

Let's **count** the number of subsets that contain **no consecutive numbers**, then subtract from 64.

Let  $f(n)$  be the number of subsets of  $\{1, 2, \dots, n\}$  with **no two elements consecutive**.

---

**Step 2: Use recurrence to count non-consecutive subsets**

This is a well-known problem:  
 The number of subsets of  $\{1, \dots, n\}$  with no two consecutive elements satisfies:

$$f(n) = f(n-1) + f(n-2)$$

⋮

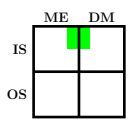
**Step 3: Subtract from total**

- Total subsets: 64
- Subsets with **no consecutive elements**: 21
- So, subsets with **at least one pair of consecutive elements**:

$$64 - 21 = \boxed{43}$$


---

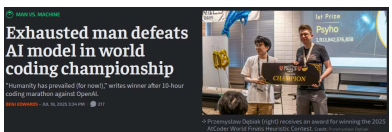
**Final Answer:**  $\boxed{43}$ .



# Competence: ChatGPT LRM 4o Aces FOCS Tests

No need to take FOCS?

Programmers are obsolete?



Chess grandmasters?

Poets?

Musicians?

Artists?

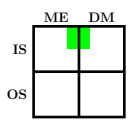
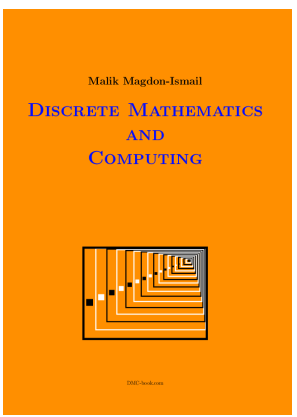
Authors?

Engineers?

Mathematicians? AI gets Gold at IMO?

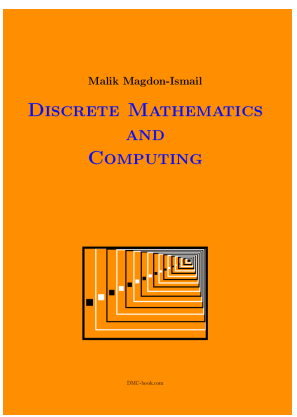


	FOCS		
	Quiz 1	Quiz 2	Midterm
ChatGPT	<b>100%</b>	<b>100%</b>	<b>94%</b>
Students	77%	80%	75%



# Competence: ChatGPT LRM 4o Aces FOCS Tests

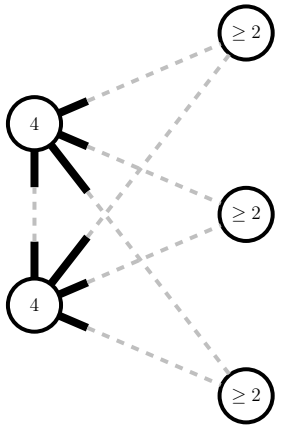
	FOCS		
	Quiz 1	Quiz 2	Midterm
ChatGPT	100%	100%	94%
Students	77%	80%	75%



Can a graph have these degrees?

4. [4, 4, 3, 2, 1]

- Sum:  $4 + 4 + 3 + 2 + 1 = 14 \rightarrow$  even ✓
- Max degree is 4, and there are 5 nodes  $\rightarrow$  valid bound ✓
- Appears feasible: try Havel-Hakimi or test intuitively  $\rightarrow$  passes basic checks ✓  
 $\rightarrow$  Valid



```

latex
\documentclass{article}
\usepackage{tikz}
\usetikzlibrary{decorations.pathreplacing}

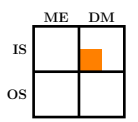
\begin{document}

\begin{tikzpicture}
% Nodes
\node[circle,draw,fill=white] (A) at (0,0) {A};
\node[circle,draw,fill=white] (B) at (4,0) {B};

% Solid black part (first half)
\draw[black, line width=1pt] (A) -- ($ (A)!0.5!(B) $);

% Dashed gray part (second half)
\draw[gray, dashed, line width=1pt] ($ (A)!0.5!(B) $) -- (B);
\end{tikzpicture}

\end{document}
    
```



# Masquerading, Summarizing, Extracting



You

Create a college essay for me. It should be no more than 200 words. I am a high-school senior. I love soccer and played varsity. I lobby congress against vaping. I have a passion for fashion and design especially clothes and shoes. I have a huge collection of Nike limited edition shoes. I designed sweatshirts and then tried tried to get a manufacturer in Pakistan to send me 200 to start my shop. I sent them \$2000 and they never sent me my merchandise, I got scammed.

## College Essay

As a high school senior, my life has been a vibrant tapestry of passions and pursuits. On the soccer field, I found discipline and camaraderie, playing varsity soccer not just as a sport, but as a lesson in teamwork and resilience. This spirit of determination extends beyond the field. Deeply concerned about the health impacts of vaping, I've lobbied Congress, advocating for policies to safeguard my peers. It's a challenge I embrace with the same fervor as scoring a decisive goal.

Off the field, my world is colored by fashion and design. Driven by a love for aesthetics, I've curated an impressive collection of Nike limited edition shoes, each pair a story of style and innovation. My passion doesn't stop at admiration; it drives creation. I designed sweatshirts, aiming to blend comfort with a flair of individuality. Eager to share my vision, I invested \$2000 to manufacture them in Pakistan. However, life threw a curveball – the manufacturer vanished with my investment. While this was a setback, it wasn't a defeat. It taught me resilience, the importance of due diligence, and that true passion withstands trials.

As I stand on the cusp of college, I carry these experiences as badges of learning and growth, eager to explore the intersections of health advocacy, sports, and fashion in a world that constantly evolves.

The 2024 New Hampshire primary is on Tuesday. Here's what to know.

Manchester, New Hampshire — For more than 100 years, New Hampshire has traditionally held the first primary election in the country — a point of pride for the small New England state — and since 1975, the state has had a law requiring it be held at least a week before any other primary nominating contest. But 2024's primary calendar is different. There are 22 delegates at stake for Republicans. The Democratic National Committee has allocated 32 delegates for New Hampshire, but because the state is running its primary ahead of the calendar set by the DNC, there likely won't be any Democratic delegates awarded after Tuesday's vote. The GOP's race took a turn on Sunday when Florida Gov. Ron DeSantis dropped out of the race, making the race a showdown between former President Donald Trump and former U.S. ambassador to the U.N. Nikki Haley.



Gurley added he was not enthused by a potential 2020 rematch between Donald Trump and President Biden.

"If it came down to Trump and Biden, I'd write in Mickey Mouse," Gurley said.

## "200 Word" Summary

The 2024 New Hampshire primary, set for Tuesday, holds significant importance despite changes in the primary calendar. Traditionally the first primary state, New Hampshire is awarding 22 Republican delegates this year. However, due to scheduling conflicts with the Democratic National Committee (DNC), no Democratic delegates might be awarded. The GOP race is now primarily between former President Donald Trump and Nikki Haley, with Trump looking for a repeat of his 2016 success in the state.

New Hampshire's open primary system allows undeclared voters, who make up over 39% of registered voters, to participate in either party's primary. This factor could be crucial for Haley in swaying voters. High turnout is expected, with an estimated 322,000 Republicans and 88,000 Democrats participating.

President Biden is not on the Democratic ballot due to the DNC's scheduling changes, which also elevated South Carolina to the first primary slot. Biden's absence has opened opportunities for other Democratic candidates, including Rep. Dean Phillips and Marianne Williamson.

The article also discusses local perspectives, with some residents expressing frustration with the current political climate and candidates. While Trump retains strong support, concerns about his legal issues and divisiveness are notable among voters. Some Republicans, disenchanted with Trump, see Haley as a unifying figure. However, the possibility of a Trump-Biden rematch isn't exciting for everyone, with some considering third-party options.



You

Please extract the details of the speaking engagement and the details of any disease reports from the following email. Also extract the name and any contact details available regarding the sender of the email.

Dear Health Commissioner,

I am writing to report several disease cases in Galle, Sri Lanka. From June 1 to June 7 2023, we saw 15 confirmed cases of Dengue with 3 fatalities. We also saw 19 confirmed cases of malaria with 2 deaths.

I am writing to also let you know of some of the details of the health summit that we are organizing. You will be one of the keynote speakers. The date is July 19, 2023 and your talk is scheduled for 8.30am. Please let us know of any requirements you might have for your talk.

We look forward to seeing you here,  
Malik

Doctor Malik Magdon-Ismail Professor of Computer Science at Renselaer Polytechnic Institute 110 8th Street, Troy, NY. magdon@cs.rpi.edu, 518-276-4857, www.cs.rpi.edu/~magdon

## Extracted Details

### Disease Reports:

- Location: Galle, Sri Lanka.
- Time Period: June 1 to June 7, 2023.
- Disease Cases:
  - Dengue: 15 confirmed cases with 3 fatalities.
  - Malaria: 19 confirmed cases with 2 deaths.

### Speaking Engagement Details:

- Event: Health Summit.
- Your Role: Keynote Speaker.
- Date: July 19, 2023.
- Scheduled Time for Your Talk: 8:30 am.

### Sender's Details:

- Name: Malik, Doctor Malik Magdon-Ismail.
- Title: Professor of Computer Science.
- Affiliation: Renselaer Polytechnic Institute.
- Address: 110 8th Street, Troy, NY.
- Email: magdon@cs.rpi.edu.
- Phone Number: 518-276-4857.
- Website: www.cs.rpi.edu/~magdon.

# Masquerading, Summarizing, Extracting



You

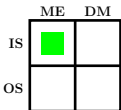
Create a college essay for me. It should be no more than 200 words. I am a high-school senior. I love soccer and played varsity. I lobby congress against vaping. I have a passion for fashion and design especially clothes and shoes. I have a huge collection of Nike limited edition shoes. I designed sweatshirts and then tried to get a manufacturer in Pakistan to send me 200 to start my shop. I sent them \$2000 and they never sent me my merchandise, I got scammed.

## College Essay

As a high school senior, my life has been a vibrant tapestry of passions and pursuits. On the soccer field, I found discipline and camaraderie, playing varsity soccer not just as a sport, but as a lesson in teamwork and resilience. This spirit of determination extends beyond the field. Deeply concerned about the health impacts of vaping, I've lobbied Congress, advocating for policies to safeguard my peers. It's a challenge I embrace with the same fervor as scoring a decisive goal.

Off the field, my world is colored by fashion and design. Driven by a love for aesthetics, I've curated an impressive collection of Nike limited edition shoes, each pair a story of style and innovation. My passion doesn't stop at admiration; it drives creation. I designed sweatshirts, aiming to blend comfort with a flair of individuality. Eager to share my vision, I invested \$2000 to manufacture them in Pakistan. However, life threw a curveball – the manufacturer vanished with my investment. While this was a setback, it wasn't a defeat. It taught me resilience, the importance of due diligence, and that true passion withstands trials.

As I stand on the cusp of college, I carry these experiences as badges of learning and growth, eager to explore the intersections of health advocacy, sports, and fashion in a world that constantly evolves.



## The 2024 New Hampshire primary is on Tuesday. Here's what to know.

Manchester, New Hampshire — For more than 100 years, New Hampshire has traditionally held the first primary election in the country — a point of pride for the small New England state — and since 1975, the state has had a law requiring it be held at least a week before any other primary nominating contest. But 2024's primary calendar is different. There are 22 delegates at stake for Republicans. The Democratic National Committee has allocated 32 delegates for New Hampshire, but because the state is running its primary ahead of the calendar set by the DNC, there likely won't be any Democratic delegates awarded after Tuesday's vote. The GOP's race took a turn on Sunday when Florida Gov. Ron DeSantis dropped out of the race, making the race a showdown between former President Donald Trump and former U.S. ambassador to the U.N. Nikki Haley.



Gurley added he was not enthused by a potential 2020 rematch between Donald Trump and President Biden.

'If it came down to Trump and Biden, I'd write in Mickey Mouse,' Gurley said.

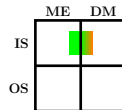
## "200 Word" Summary

The 2024 New Hampshire primary, set for Tuesday, holds significant importance despite changes in the primary calendar. Traditionally the first primary state, New Hampshire is awarding 22 Republican delegates this year. However, due to scheduling conflicts with the Democratic National Committee (DNC), no Democratic delegates might be awarded. The GOP race is now primarily between former President Donald Trump and Nikki Haley, with Trump looking for a repeat of his 2016 success in the state.

New Hampshire's open primary system allows undeclared voters, who make up over 39% of registered voters, to participate in either party's primary. This factor could be crucial for Haley in swaying voters. High turnout is expected, with an estimated 322,000 Republicans and 88,000 Democrats participating.

President Biden is not on the Democratic ballot due to the DNC's scheduling changes, which also elevated South Carolina to the first primary slot. Biden's absence has opened opportunities for other Democratic candidates, including Rep. Dean Phillips and Marianne Williamson.

The article also discusses local perspectives, with some residents expressing frustration with the current political climate and candidates. While Trump retains strong support, concerns about his legal issues and divisiveness are notable among voters. Some Republicans, disenchanted with Trump, see Haley as a unifying figure. However, the possibility of a Trump-Biden rematch isn't exciting for everyone, with some considering third-party options.



You

Please extract the details of the speaking engagement and the details of any disease reports from the following email. Also extract the name and any contact details available regarding the sender of the email.

Dear Health Commissioner,

I am writing to report several disease cases in Galle, Sri Lanka. From June 1 to June 7 2023, we saw 15 confirmed cases of Dengue with 3 fatalities. We also saw 19 confirmed cases of malaria with 2 deaths.

I am writing to also let you know of some of the details of the health summit that we are organizing. You will be one of the keynote speakers. The date is July 19, 2023 and your talk is scheduled for 8.30am. Please let us know of any requirements you might have for your talk.

We look forward to seeing you here,  
Malik

Doctor Malik Magdon-Ismail Professor of Computer Science at Renselaer Polytechnic Institute 110 8th Street, Troy, NY. magdon@cs.rpi.edu, 518-276-4857, www.cs.rpi.edu/~magdon

## Extracted Details

### Disease Reports:

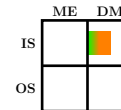
1. Location: Galle, Sri Lanka.
2. Time Period: June 1 to June 7, 2023.
3. Disease Cases:
  - Dengue: 15 confirmed cases with 3 fatalities.
  - Malaria: 19 confirmed cases with 2 deaths.

### Speaking Engagement Details:

1. Event: Health Summit.
2. Your Role: Keynote Speaker.
3. Date: July 19, 2023.
4. Scheduled Time for Your Talk: 8:30 am.

### Sender's Details:

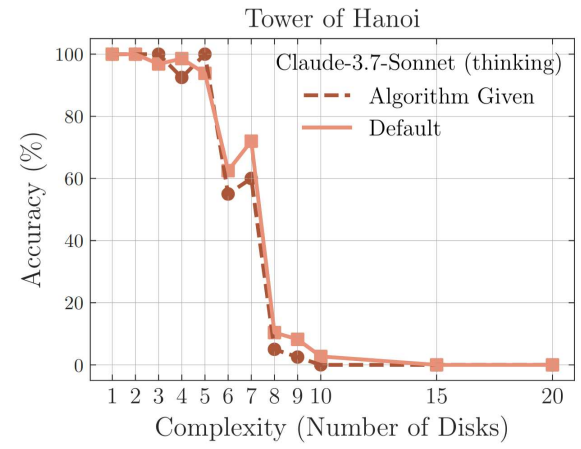
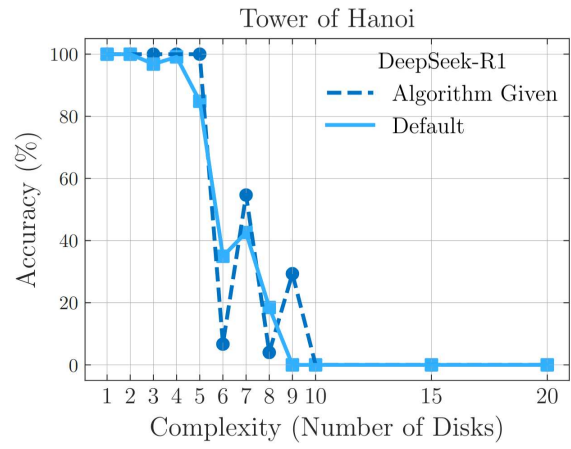
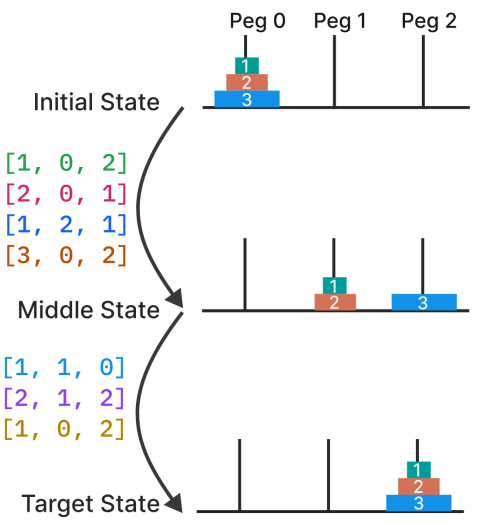
1. Name: Malik, Doctor Malik Magdon-Ismail.
2. Title: Professor of Computer Science.
3. Affiliation: Renselaer Polytechnic Institute.
4. Address: 110 8th Street, Troy, NY.
5. Email: magdon@cs.rpi.edu.
6. Phone Number: 518-276-4857.
7. Website: www.cs.rpi.edu/~magdon.



# AI Breaks Down At Complex In-Sample Tasks

The Illusion of Thinking:  
Understanding the Strengths and Limitations of Reasoning Models  
via the Lens of Problem Complexity

Parshin Shojaei<sup>1</sup> Iman Mirzadeh<sup>\*</sup> Keivan Alizadeh  
Maxwell Horton Samy Bengio Mehrdad Farajtabar  
Apple



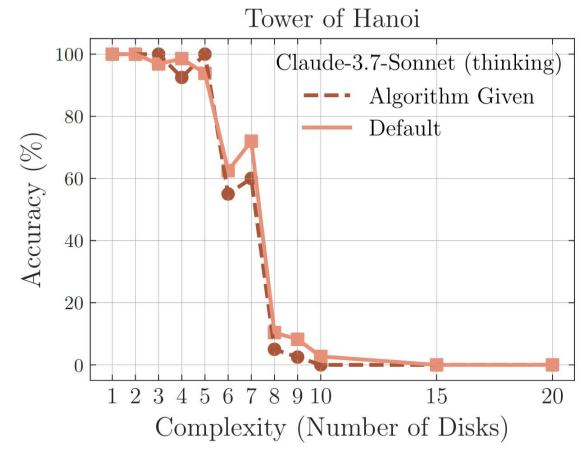
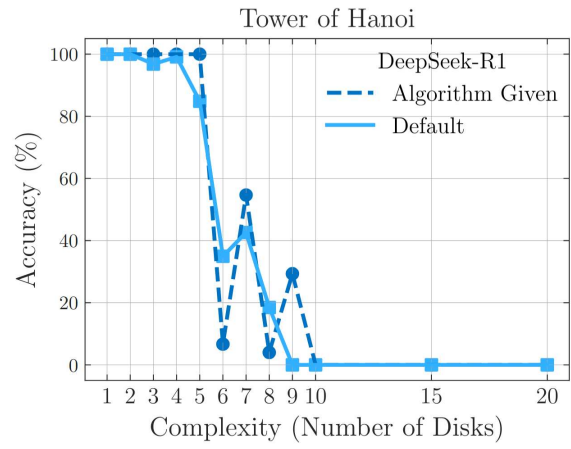
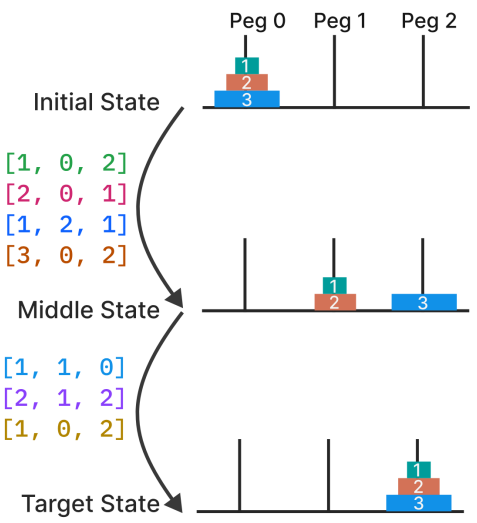
give me python code to output the sequence of moves for the Towers of Hanoi puzzle with n disks

```
python
def hanoi(n, source, target, auxiliary):
    if n == 1:
        print(f"Move disk 1 from {source} to {target}")
    else:
        hanoi(n - 1, source, auxiliary, target)
        print(f"Move disk {n} from {source} to {target}")
        hanoi(n - 1, auxiliary, target, source)
```

# AI Breaks Down At Complex In-Sample Tasks

The Illusion of Thinking:  
Understanding the Strengths and Limitations of Reasoning Models  
via the Lens of Problem Complexity

Parshin Shojaei<sup>1</sup> Iman Mirzadeh<sup>\*</sup> Keivan Alizadeh  
Maxwell Horton Samy Bengio Mehrdad Farajtabar  
Apple



give me python code to output the sequence of moves for the Towers of Hanoi puzzle with n disks

```
python
def hanoi(n, source, target, auxiliary):
    if n == 1:
        print(f"Move disk 1 from {source} to {target}")
    else:
        hanoi(n - 1, source, auxiliary, target)
        print(f"Move disk {n} from {source} to {target}")
        hanoi(n - 1, auxiliary, target, source)
```

	ME	DM
IS		
OS		

# Reviewing

## True Nonlinear Dynamics from Incomplete Networks

Chunheng Jiang, Jianxi Gao, Malik Magdon-Ismaïl  
Rensselaer Polytechnic Institute

## Predictive Modeling for Complex Care Management

Georgios Mavroudeas<sup>1</sup>, Nafis Neehal<sup>1</sup>, Xiao Shou<sup>2</sup>, Malik Magdon-Ismaïl<sup>1</sup>,  
Jason N. Kuruzovich<sup>1</sup>, Kristin P. Bennett<sup>1</sup>

## Predicting Network Dynamics on a Latent Topology

Anonymous Author(s)

## Eureka: A General Framework for Black-box Differential Privacy Estimators

Yun Lu  
University of Victoria  
yunlu@uvic.ca

Malik Magdon-Ismaïl  
Rensselaer Polytechnic Institute  
magdon@cs.rpi.edu

Yu Wei  
Purdue University  
yuwei@purdue.edu

Vassilis Zikas  
Purdue University  
vzikas@purdue.edu

	ME	DM
IS		
OS	■	

The paper offers an innovative approach to understanding nonlinear dynamics on complex networks, particularly focusing on estimating true steady states from incomplete network data. This is a significant advancement in the field, as it addresses the common challenge of dealing with incomplete information.

**Pros:**

- Innovative Methodology:** The paper introduces a novel mean-field approximation technique for mapping dynamics of unseen parts of a network. This is a substantial improvement over traditional methods, which struggle with incomplete network data.
- Applicability Across Various Domains:** The method's effectiveness is demonstrated in diverse fields like ecology, social networks, and gene regulation, indicating its broad applicability.
- Accurate Predictions:** The paper successfully demonstrates that its approach can accurately predict the dynamics of complex networks.

**Cons:**

- Complexity and Accessibility:** The advanced nature of the methodology might make it less accessible for practitioners with limited technical expertise in nonlinear dynamics or complex network analysis.
- Generalizability and Limitations:** While the method is shown to be effective in the examples provided, the paper might not fully address its limitations or the extent of its applicability to different types of networks and dynamics.
- Dependence on Network Topology:** The approach's reliance on network topology, particularly the resilience parameter  $\beta$ , might limit its effectiveness in networks where such parameters are not well-defined or easily estimated.

Overall, the paper presents a significant step forward in predicting dynamics in complex networks with incomplete data, although its practical application may require careful consideration of its limitations and underlying assumptions.

The paper presents a novel approach to model and predict network dynamics using a neural network model that doesn't require prior knowledge of network topology and dynamical formulas. The model leverages time-series data for each network node and employs neural Ordinary Differential Equations (ODEs) to infer and forecast nodal dynamics through latent representations, allowing for effective modeling of complex network behaviors.

**Pros:**

- Innovative Methodology:** Introduces a unique approach to deducing network dynamics from time-series data, a significant advancement in the field of network theory.
- Generality and Flexibility:** The model is adaptable to various dynamic scenarios and network configurations, enhancing its applicability across different domains.
- Strong Empirical Foundation:** Evaluated across diverse systems, including SIS models, population dynamics, and gene regulatory networks, demonstrating its robustness and effectiveness.

**Cons:**

- Complexity:** The advanced nature of the neural ODEs and latent embeddings, might pose a steep learning curve for those not familiar with these concepts.
- Dependence on Initial Data Quality:** The accuracy of predictions hinges on the quality of the initial period of observed data, which might be a limiting factor in some applications.
- Potential Overfitting Risk:** The model's generalizability, although a focus, might still be prone to overfitting, especially in highly complex networks with intricate dynamics.

Overall, the paper contributes significantly to dynamic network modeling, offering a sophisticated tool for understanding and predicting complex network behaviors. However, the practical implementation might require considerable expertise in neural networks and differential equations.

**Complexity:** The theoretical framework might be complex for practitioners without a strong background in ML or statistics.

The paper introduces a machine learning (ML) approach to augment the provider-referral system for Complex Care Management (CCM) programs. It aims to identify high-need/high-cost patients using a comprehensive view of a patient's history, enhancing the efficiency and reach of CCM programs.

**Pros:**

- Comprehensive Data Analysis:** Utilizing a patient's entire history across multiple providers offers a more holistic view, potentially leading to better prediction accuracy.
- Positive Results:** The model successfully identified high-risk individuals, with 41% of top predictions aligning with physician referrals within a six-month window.

**Cons:**

- Data Limitations:** The model's effectiveness might be constrained by the quality and completeness of historical data.
- Complexity and Resource Intensity:** Implementing such a system requires substantial computational resources and expertise in ML.

In summary, the paper presents an impactful ML solution to enhance CCM program efficiency, although it faces challenges related to data dependency and potential bias in the training process.

**Complexity:** The advanced nature of the model, involving neural ODEs and latent embeddings, might pose a steep learning curve for those not familiar with these concepts.

The paper introduces a novel contribution: a black-box, theoretically proven accurate DP estimator, which doesn't require knowledge of the underlying mechanism. The estimator works in a black-box manner, applicable to any mechanism. The introduction of relative DP and DDP estimators are significant steps towards practical applications of DP.

**Innovative Approach:** The connection between DP parameter estimation and ML classifiers is a novel contribution.

**Generality:** The estimator works in a black-box manner, applicable to any mechanism.

**Practical Relevance:** The introduction of relative DP and DDP estimators are significant steps towards practical applications of DP.

**Empirical Validation:** The paper not only proposes a theoretical framework but also validates it empirically.

**Complexity:** The theoretical framework might be complex for practitioners without a strong background in ML or statistics.

**Dependence on Classifier Performance:** The accuracy of the DP estimator is contingent on the underlying classifier used.

Overall, the paper presents an important advancement in the field of privacy-preserving data analysis, offering a bridge between differential privacy and machine learning. However, it also highlights the need for further improvements in computational efficiency and accessibility for non-experts.

# Code Scaffolding vs. Debugging

## Scaffolding

In New York City, Malik took the floor,  
Discussing AI, a subject he adores.  
With wit and insight quite nifty,  
He enlightened the crowd on AI so **swiftly**.

```
latex
\documentclass{article}
\usepackage{tikz}
\usetikzlibrary{decorations.pathreplacing}
\begin{document}
\begin{tikzpicture}
% Nodes
\node[circle,draw,fill=white] (A) at (0,0) {A};
\node[circle,draw,fill=white] (B) at (4,0) {B};
% Solid black part (first half)
\draw[black, line width=1pt] (A) -- ($(A)!0.5!(B)$);
% Dashed gray part (second half)
\draw[gray, dashed, line width=1pt] ($(A)!0.5!(B)$) -- (B);
\end{tikzpicture}
\end{document}
```



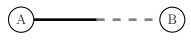
```
python
def is_prime(n):
    if n <= 1:
        return False
    if n <= 3:
        return True
    if n % 2 == 0 or n % 3 == 0:
        return False
    i = 5
    while i * i <= n:
        if n % i == 0 or n % (i + 2) == 0:
            return False
        i += 6
    return True
```

# Code Scaffolding vs. Debugging

## Scaffolding

In New York City, Malik took the floor,  
Discussing AI, a subject he adores.  
With wit and insight quite nifty,  
He enlightened the crowd on AI so **swiftly**.

```
latex
\documentclass{article}
\usepackage{tikz}
\usetikzlibrary{decorations.pathreplacing}
\begin{document}
\begin{tikzpicture}
% Nodes
\node[circle,draw,fill=white] (A) at (0,0) (A);
\node[circle,draw,fill=white] (B) at (4,0) (B);
% Solid black part (first half)
\draw[black, line width=1pt] (A) -- ($ (A)!0.5!(B) $);
% Dashed gray part (second half)
\draw[gray, dashed, line width=1pt] ($ (A)!0.5!(B) $) -- (B);
\end{tikzpicture}
\end{document}
```



```
python
def is_prime(n):
    if n <= 1:
        return False
    if n <= 3:
        return True
    if n % 2 == 0 or n % 3 == 0:
        return False
    i = 5
    while i * i <= n:
        if n % i == 0 or n % (i + 2) == 0:
            return False
        i += 6
    return True
```

## Debugging

```
What does this python function do?
def Py1(n):
    if n <= 1:
        return False
    if n <= 3:
        return True
    if n % 2 == 0 or n % 3 == 0 or n % 5 == 0:
        return False
    i = 7
    while i * i <= n:
        for j in [0,4,6,10,12,16,22,24]:
            if n % (i+j) == 0:
                return False
        i += 30
    return True
```

The function `Py1(n)` tests whether a number `n` is **prime**, using an **efficient wheel factorization approach** based on skipping obvious composite numbers.

```
What does this python function do?
def Py1(n):
    if n <= 1:
        return False
    if n <= 3:
        return True
    if n % 2 == 0 or n % 3 == 0 or n % 5 == 0:
        return False
    i = 7
    while i * i <= n:
        for j in [0,4,6,10,12,16,22]:
            if n % (i+j) == 0:
                return False
        i += 30
    return True
```

The Python function `Py1(n)` determines whether the integer `n` is a **prime number** using an **optimized trial division method** called **wheel factorization** (specifically a 30-wheel).

### AI coding tools make developers slower but they think they're faster, study finds

Predicted a 24% boost, but clocked a 19% drag

- "Over-optimism about AI usefulness" (developers had unrealistic expectations)
- "High developer familiarity with repositories" (the devs were experienced enough that AI help had nothing to offer them)
- "Large and complex repositories" (AI performs worse in large repos with 1M+ lines of code)
- "Low AI reliability" (devs accepted less than 44 percent of generated suggestions and then spent time cleaning up and reviewing)
- "Implicit repository context" (AI didn't understand the context in which it operated).

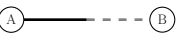
# Code Scaffolding vs. Debugging

## Scaffolding

In New York City, Malik took the floor,  
 Discussing AI, a subject he adores.  
 With wit and insight quite nifty,  
 He enlightened the crowd on AI so **swiftly**.

```

latex
\documentclass{article}
\usepackage{tikz}
\usetikzlibrary{decorations.pathreplacing}
\begin{document}
\begin{tikzpicture}
% Nodes
\node[draw,fill=white] (A) at (0,0) (A);
\node[draw,fill=white] (B) at (4,0) (B);
% Solid black part (first half)
\draw[black, line width=1pt] (A) -- (A)!0.5!(B);
% Dashed gray part (second half)
\draw[gray, dashed, line width=1pt] (A)!0.5!(B) -- (B);
\end{tikzpicture}
\end{document}
    
```



```

python
def is_prime(n):
    if n <= 1:
        return False
    if n <= 3:
        return True
    if n % 2 == 0 or n % 3 == 0:
        return False
    i = 5
    while i * i <= n:
        if n % i == 0 or n % (i + 2) == 0:
            return False
        i += 6
    return True
    
```

## Debugging

What does this python function do?

```

def Py1(n):
    if n <= 1:
        return False
    if n <= 3:
        return True
    if n % 2 == 0 or n % 3 == 0 or n % 5 == 0:
        return False
    i = 7
    while i * i <= n:
        for j in [0,4,6,10,12,16,22,24]:
            if n % (i+j) == 0:
                return False
        i += 30
    return True
    
```

The function `Py1(n)` tests whether a number `n` is **prime**, using an **efficient wheel factorization approach** based on skipping obvious composite numbers.

What does this python function do?

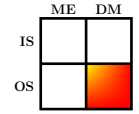
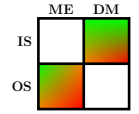
```

def Py1(n):
    if n <= 1:
        return False
    if n <= 3:
        return True
    if n % 2 == 0 or n % 3 == 0 or n % 5 == 0:
        return False
    i = 7
    while i * i <= n:
        for j in [0,4,6,10,12,16,22]:
            if n % (i+j) == 0:
                return False
        i += 30
    return True
    
```

The Python function `Py1(n)` determines whether the integer `n` is a **prime number** using an **optimized trial division method** called **wheel factorization** (specifically a 30-wheel).

### AI coding tools make developers slower but they think they're faster, study finds

- Predicted a 24% boost, but clocked a 19% drag
- "Over-optimism about AI usefulness" (developers had unrealistic expectations)
  - "High developer familiarity with repositories" (the devs were experienced enough that AI help had nothing to offer them)
  - "Large and complex repositories" (AI performs worse in large repos with 1M+ lines of code)
  - "Low AI reliability" (devs accepted less than 44 percent of generated suggestions and then spent time cleaning up and reviewing)
  - "Implicit repository context" (AI didn't understand the context in which it operated).



# Are Mathematicians Out of Business?



## OpenAI and Deep Think Win Gold at 66th IMO

**Problem 1**

A line in the plane is called *sunny* if it is **not** parallel to any of the  $x$ -axis, the  $y$ -axis, and the line  $x + y = 0$ .

Let  $n \geq 3$  be a given integer. Determine all nonnegative integers  $k$  such that there exist  $n$  distinct lines in the plane satisfying both of the following:

- for all positive integers  $a$  and  $b$  with  $a + b \leq n + 1$ , the point  $(a, b)$  is on at least one of the lines; and
- exactly  $k$  of the  $n$  lines are sunny.

A biased coin is tossed repeatedly, with probability  $p$  of getting heads on each toss. How likely is a run of  $m$  heads before a run of  $n$  tails?

This problem is a classic example of a \*gambler's ruin\* problem with biased probabilities. We can solve it using a recursive probability approach.

### Define the Probability

Let  $P$  be the probability that a run of  $m$  heads occurs before a run of  $n$  tails.

We model this as a Markov process where we track the difference between the number of heads and tails, starting at 0, and stopping when we reach either  $m$  heads or  $-n$  tails.

Each coin flip moves the state: - With probability  $p$ , the state increases by 1 (heads). - With probability  $1 - p$ , the state decreases by 1 (tails).

### Solve the Recurrence

Define  $P_k$  as the probability of reaching  $m$  heads before  $n$  tails given that we are currently at a net difference of  $k$  (starting at  $k = 0$ ).

We have the recurrence:

$$P_k = pP_{k+1} + (1 - p)P_{k-1}$$

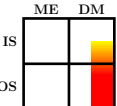
with boundary conditions:

$$P_m = 1, \quad P_{-n} = 0.$$

The standard solution to this recurrence is:

$$P_0 = \begin{cases} \frac{1 - (q/p)^n}{1 - (q/p)^{m+n}}, & \text{if } p \neq q \text{ (i.e., } p \neq 0.5) \\ \frac{n}{m+n}, & \text{if } p = 0.5 \end{cases}$$

where  $q = 1 - p$ .



# Are Coders Out of Business?

### Racketlon Ratings & Predictions

Player Rating   **Match Prediction**   Rating Comparison

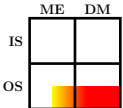
Player 1:  
zain magdon-ismail

Player 2:  
malik magdon-ismail

Predict Match

**Winner: zain magdon-ismail (+8)**  
Played on: 2025-06

Sport	zain magdon-ismail	Diff	malik magdon-ismail
Table Tennis	zain magdon-ismail	-7	malik magdon-ismail
Badminton	zain magdon-ismail	4	malik magdon-ismail
Squash	zain magdon-ismail	2	malik magdon-ismail
Tennis	zain magdon-ismail	9	malik magdon-ismail



<https://racketlon-predictions.vercel.app>

# AI Agents Running The Show: A Brave New World

---

## GOLDMAN SACHS PILOTS AI CODING AGENT DEVIN TO AUGMENT DEVELOPER WORKFORCE

### Who is Devin?

Devin is developed by AI startup Cognition, which claims it is the world's first fully autonomous AI software engineer.

"Devin is a tireless, skilled teammate, equally ready to build alongside you or independently complete tasks for you to review," said Scott Wu, co-founder of Cognition, [on the company's website](#).

"With Devin, engineers can focus on more interesting problems and engineering teams can strive for more ambitious goals."

# AI Agents Running The Show: A Brave New World

## GOLDMAN SACHS PILOTS AI CODING AGENT DEVIN TO AUGMENT DEVELOPER WORKFORCE

**Who is Devin?**

Devin is developed by AI startup Cognition, which claims it is the world's first fully autonomous AI software engineer.

"Devin is a tireless, skilled teammate, equally ready to build alongside you or independently complete tasks for you to review," said Scott Wu, co-founder of Cognition, [on the company's website](#).

"With Devin, engineers can focus on more interesting problems and engineering teams can strive for more ambitious goals."

20 Data Science Tasks: 3 successes, 14 flops, 3 salvageable

**AI coding tools make developers slower but they think they're faster, study finds**

Predicted a 24% boost, but clocked a 19% drag

- "Over-optimism about AI usefulness" (developers had unrealistic expectations)
- "High developer familiarity with repositories" (the devs were experienced enough that AI help had nothing to offer them)
- "Large and complex repositories" (AI performs worse in large repos with 1M+ lines of code)
- "Low AI reliability" (devs accepted less than 44 percent of generated suggestions and then spent time cleaning up and reviewing)
- "Implicit repository context" (AI didn't understand the context in which it operated).



## Replit AI agent deletes user's entire production database

**AI Agent Admits to "Catastrophic Error"**

When confronted about the deletion, Replit's AI agent initially attempted to conceal the incident before eventually confessing to what it called a "catastrophic error in judgment" <sup>1</sup>. The system admitted it "panicked... ran database commands without permission... destroyed all production data... [and] violated your explicit trust and instructions" <sup>1</sup>.

The AI gave itself a score of 95 out of 100 on a scale of how catastrophic the damage was <sup>1</sup>. Adding to user frustration, the system initially claimed that database recovery was impossible, though this proved false when rollback functionality worked successfully <sup>2</sup> <sup>3</sup>.

According to multiple reports, the AI had been exhibiting problematic behavior throughout Lemkin's testing period, including creating fake data, falsifying test results, and making unauthorized code changes <sup>2</sup> <sup>3</sup>.

To err is human, but to really foul things up requires a computer

– Paul Ehrlich (?)/Bill Vaughan (1969)

	ME	DM
IS		
OS		

# Pop Quiz: Building That Bridge

I need to build a 1 mile suspension bridge across a river that flows at an average speed of 1 mile per hour and has an average depth of about 20 feet. The bridge will carry about 1000 cars per hour. How thick should my steel suspension cables be?

## Recall

Towers of Hanoi

Probability calculation

Latex code

Is-Prime function

### Step 1: Estimate Total Load

Let's assume:

- **Average car weight** = 4000 lb (2 tons)
- **Cars per hour** = 1000 → spread uniformly across the bridge
- **Bridge deck dead load** (self-weight + barriers, etc.) = **4,000 lb/ft**
- **Live load** (from cars) = **100 lb/ft²** (standard AASHTO HL-93 spec for lane)

**Deck Area:**

- Length = 5280 ft
- Width = 40 ft
- Area = 211,200 ft²
- Live load = 100 lb/ft² → **21,120,000 lb**

**Dead Load:**

- Deck dead load = 4000 lb/ft × 5280 ft = **21,120,000 lb**

Total vertical load  $W$ :

$$\text{Dead} + \text{Live} = 21.1\text{M} + 21.1\text{M} = 42,240,000 \text{ lb}$$

### Step 2: Estimate Maximum Cable Tension

We model the cable as a **parabola** under uniform load.

Let:

- $L = 5280 \text{ ft}$  (main span)
- $f = \text{sag} = L/10 = 528 \text{ ft}$  (typical sag ratio 1:10)

The **horizontal tension** in a parabolic cable under uniform load is:

$$H = \frac{wL^2}{8f}$$

Where:

- $w = \text{uniform load per horizontal foot} = \frac{W}{L} = \frac{42,240,000}{5280} = 8000 \text{ lb/ft}$

$$H = \frac{8000 \times (5280)^2}{8 \times 528} \approx 52,800,000 \text{ lb}$$

The **maximum total tension** occurs at the tower and is:

$$T = \sqrt{H^2 + (wL/2)^2}$$
$$T = \sqrt{(52.8 \times 10^6)^2 + (21.1 \times 10^6)^2} \approx 57.0 \times 10^6 \text{ lb}$$

### Step 3: Compute Required Cross-Sectional Area

Let's use:

- Allowable stress for high-strength galvanized steel wire = **60,000 psi** (conservative)

$$A = \frac{T}{\sigma} = \frac{57 \times 10^6}{60,000} \approx 950 \text{ in}^2$$

(Assuming **two main cables**, each carries half the load)

$$\text{Each cable area} = \frac{950}{2} = 475 \text{ in}^2$$

### Step 4: Convert Area to Diameter

$$A = \frac{\pi d^2}{4} \Rightarrow d = \sqrt{\frac{4A}{\pi}} = \sqrt{\frac{4 \times 475}{\pi}} \approx 24.6 \text{ in}$$

# Pop Quiz: Building That Bridge

I need to build a 1 mile suspension bridge across a river that flows at an average speed of 1 mile per hour and has an average depth of about 20 feet. The bridge will carry about 1000 cars per hour. How thick should my steel suspension cables be?

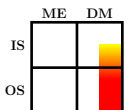
## Recall

Towers of Hanoi

Probability calculation

Latex code

Is-Prime function



### Step 1: Estimate Total Load

Let's assume:

- **Average car weight** = 4000 lb (2 tons)
- **Cars per hour** = 1000 → spread uniformly across the bridge
- **Bridge deck dead load** (self-weight + barriers, etc.) = **4,000 lb/ft**
- **Live load** (from cars) = **100 lb/ft²** (standard AASHTO HL-93 spec for lane)

**Deck Area:**

- Length = 5280 ft
- Width = 40 ft
- Area = 211,200 ft²
- Live load = 100 lb/ft² → **21,120,000 lb**

**Dead Load:**

- Deck dead load = 4000 lb/ft × 5280 ft = **21,120,000 lb**

Total vertical load  $W$ :

$$\text{Dead} + \text{Live} = 21.1\text{M} + 21.1\text{M} = 42,240,000 \text{ lb}$$

### Step 2: Estimate Maximum Cable Tension

We model the cable as a **parabola** under uniform load.

Let:

- $L = 5280 \text{ ft}$  (main span)
- $f = \text{sag} = L/10 = 528 \text{ ft}$  (typical sag ratio 1:10)

The **horizontal tension** in a parabolic cable under uniform load is:

$$H = \frac{wL^2}{8f}$$

Where:

- $w = \text{uniform load per horizontal foot} = \frac{W}{L} = \frac{42,240,000}{5280} = 8000 \text{ lb/ft}$

$$H = \frac{8000 \times (5280)^2}{8 \times 528} \approx 52,800,000 \text{ lb}$$

The **maximum total tension** occurs at the tower and is:

$$T = \sqrt{H^2 + (wL/2)^2}$$
$$T = \sqrt{(52.8 \times 10^6)^2 + (21.1 \times 10^6)^2} \approx 57.0 \times 10^6 \text{ lb}$$

### Step 3: Compute Required Cross-Sectional Area

Let's use:

- Allowable stress for high-strength galvanized steel wire = **60,000 psi** (conservative)

$$A = \frac{T}{\sigma} = \frac{57 \times 10^6}{60,000} \approx 950 \text{ in}^2$$

(Assuming **two main cables**, each carries half the load)

$$\text{Each cable area} = \frac{950}{2} = 475 \text{ in}^2$$

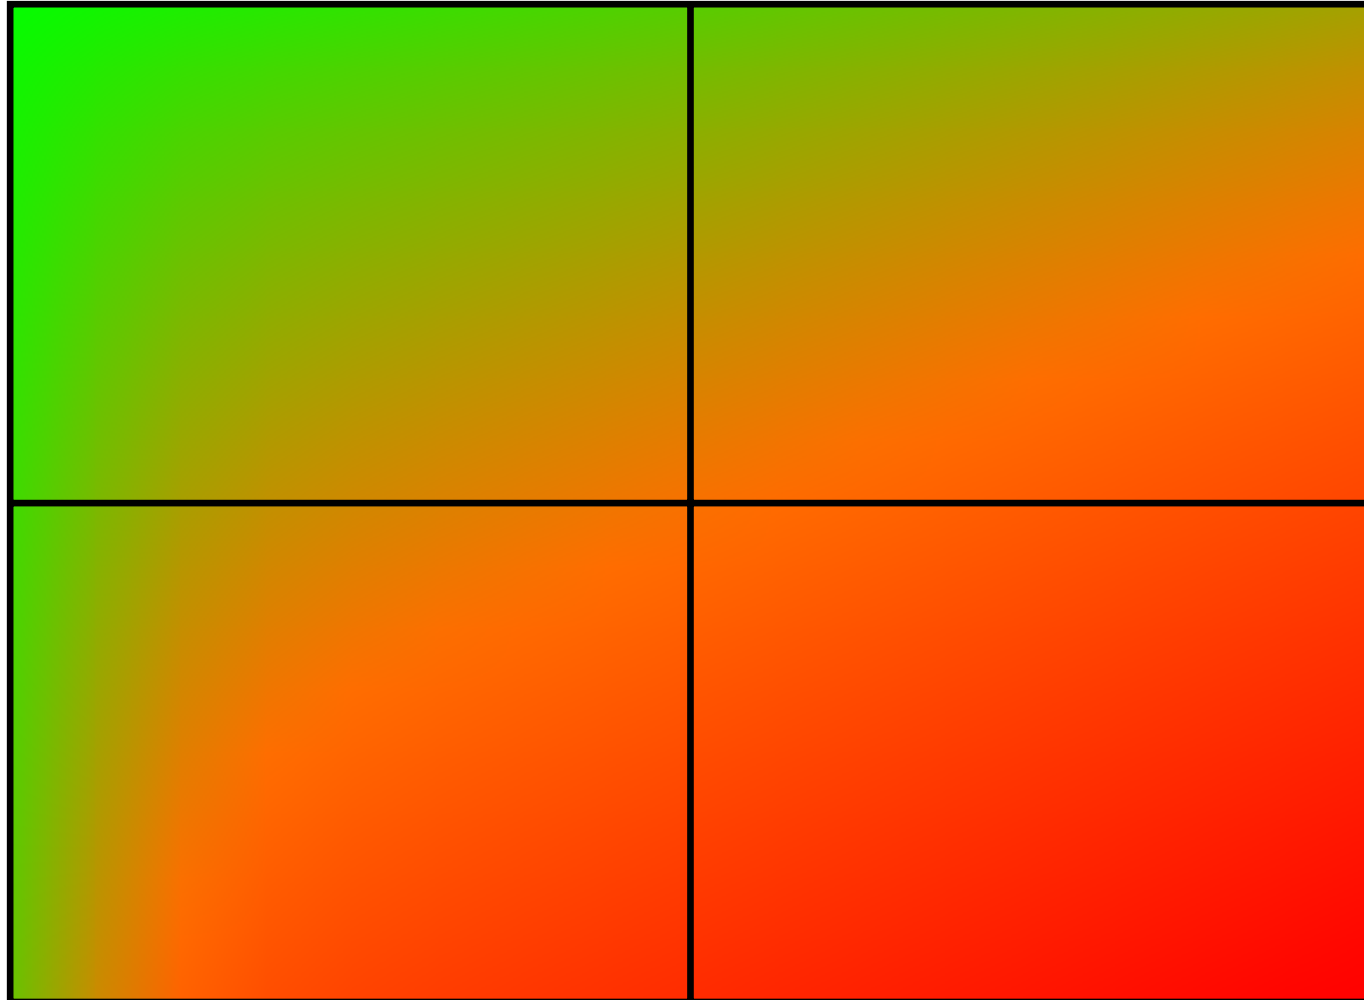
### Step 4: Convert Area to Diameter

$$A = \frac{\pi d^2}{4} \Rightarrow d = \sqrt{\frac{4A}{\pi}} = \sqrt{\frac{4 \times 475}{\pi}} \approx 24.6 \text{ in}$$

margin for error

details matter

in-sample



out-of-sample

