Busting the Paper Ballot: Voting Meets Adversarial ML

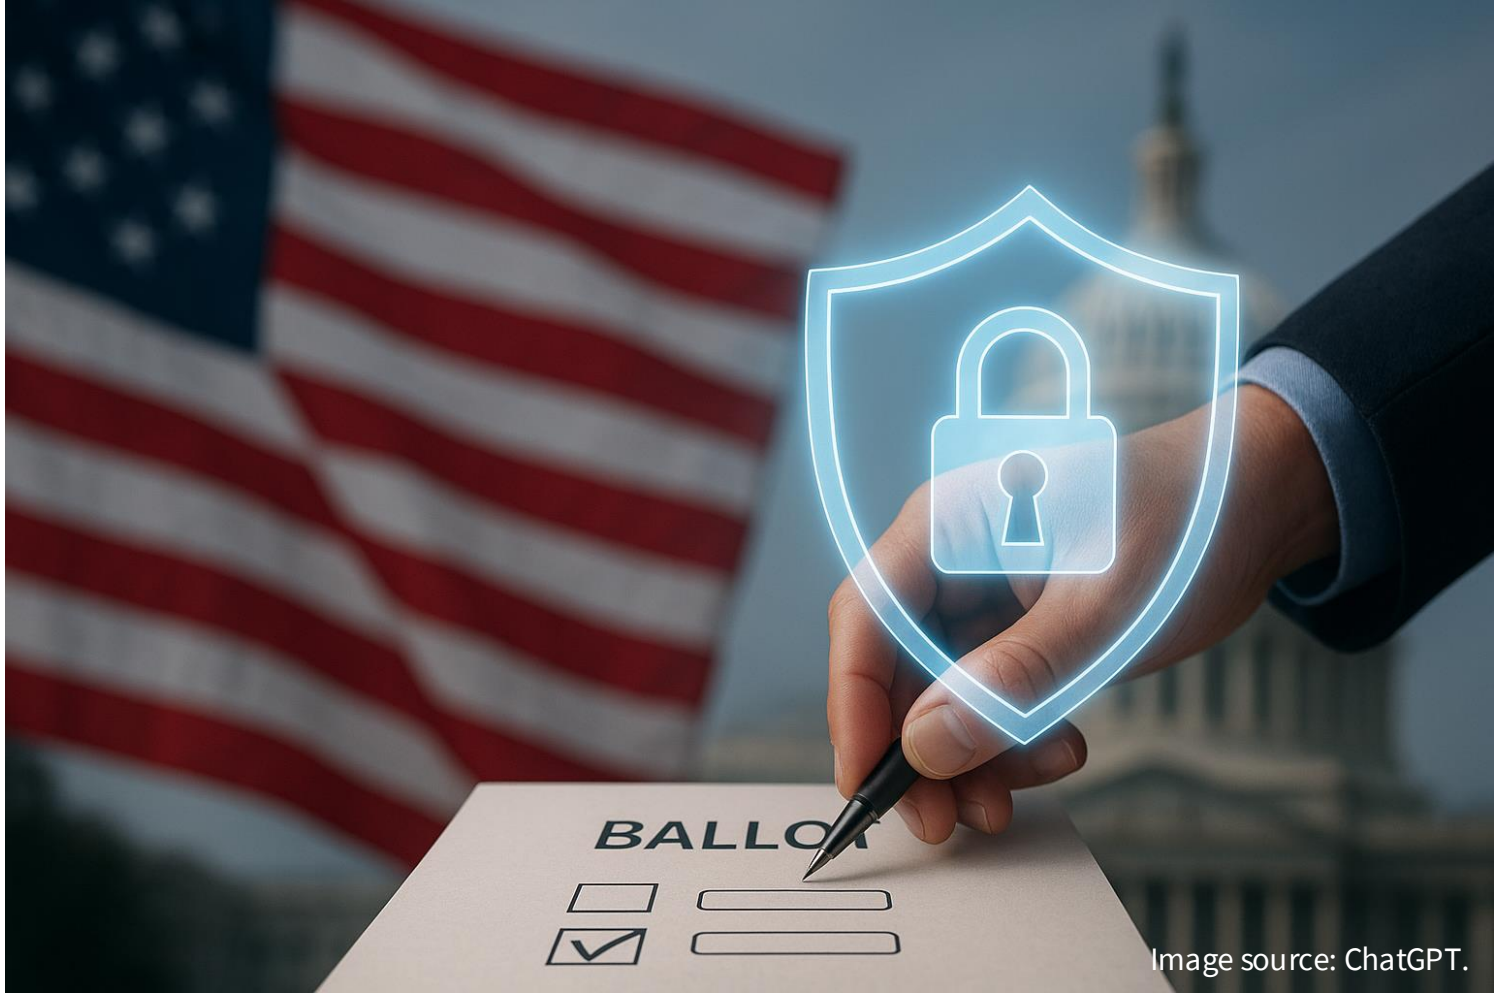

Image source: ChatGPT.

Aayushi Verma (University of Connecticut)

Joint work with: Kaleel Mahmood, Caleb Manicke, Ethan Rathbun, Nicholas Stamatakis, Sohaib Ahmad, Benjamin Fuller, Laurent Michel

Just a *few* **tampered** ballots can **swing** a **close** election.

This is a cautionary tale about using machine learning in critical systems.

Disclaimer: The images of ballots in this presentation are for illustrative purposes only. They are not intended as endorsements of any candidate, political party, or policy. All tested models are entirely created by our team for research purposes and are not intended to represent any vendor product(s).

This is our threat model and adversary.

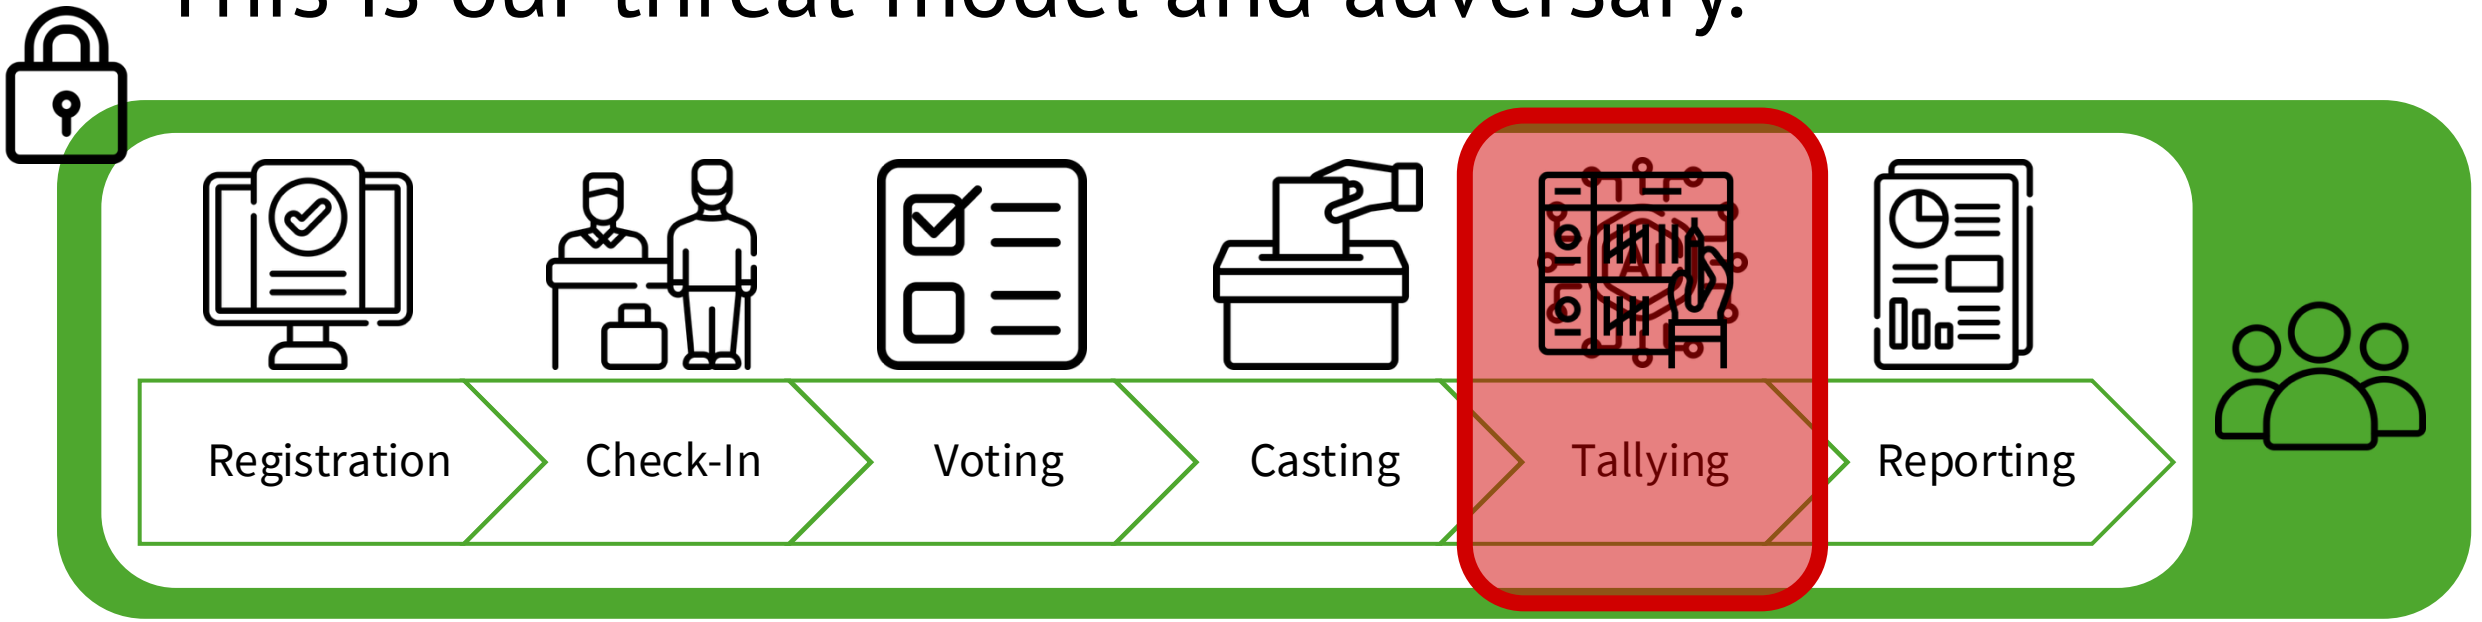

Logistics & operations

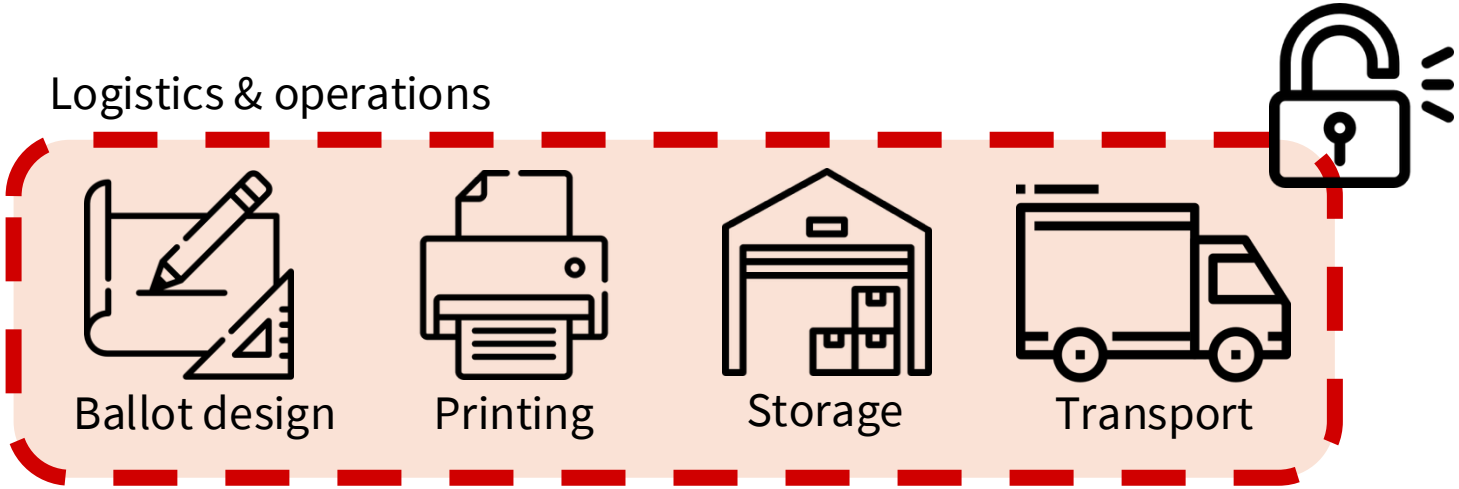

- ✓ Full **knowledge** of ML system (but no access)
- ✓ **Access** to physical pipeline.

Image sources: Flaticon.

Visual representation of attack.

- **Goal:** real Loser wins by at least 0.5%.
- **Attack:** create enough tampered ballots, mix with real ballots, watch chaos unfold.

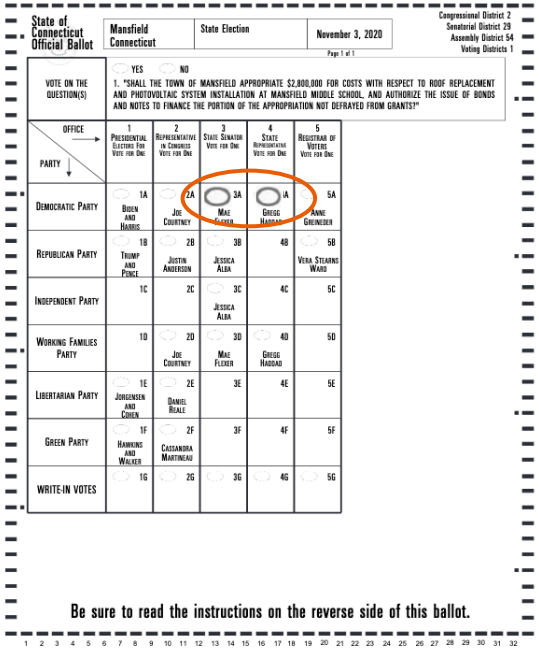

State of Connecticut Official Ballot | Mansfield Connecticut | State Election | November 3, 2020

Congressional District 2
Senatorial District 29
Assembly District 54
Voting Districts 1

Page 1 of 1

VOTE ON THE QUESTION(S)

YES NO

1. "SHALL THE TOWN OF MANSFIELD APPROPRIATE $2,800,000 FOR COSTS WITH RESPECT TO ROOF REPLACEMENT AND PHOTOVOLTAIC SYSTEM INSTALLATION AT MANSFIELD MIDDLE SCHOOL, AND AUTHORIZE THE ISSUE OF BONDS AND NOTES TO FINANCE THE PORTION OF THE APPROPRIATION NOT DEFRAYED FROM GRANTS?"

OFFICE / PARTY	1 Presidential Electors For Vote for One	2 Representative in Congress Vote for One	3 State Senator Vote for One	4 State Representative Vote for One	5 Registrar of Voters Vote for One
Democratic Party	1A Biden and Harris	2A Joe Courtney	3A Mae Flexer	4A Gregg Haddad	5A Anne Greinecker
Republican Party	1B Trump and Pence	2B Justin Anderson	3B Jessica Alba	4B	5B Vera Stearns Ward
Independent Party	1C	2C	3C Jessica Alba	4C	5C
Working Families Party	1D	2D Joe Courtney	3D Mae Flexer	4D Gregg Haddad	5D
Libertarian Party	1E Jorgensen and Cohen	2E Daniel Reale	3E	4E	5E
Green Party	1F Hawkins and Walker	2F Cassandra Martineau	3F	4F	5F
WRITE-IN VOTES	1G	2G	3G	4G	5G

Be sure to read the instructions on the reverse side of this ballot.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32

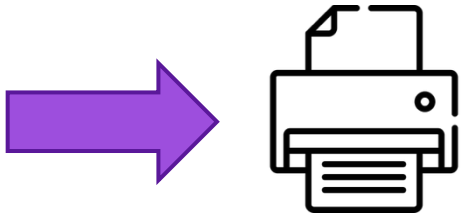

Printing

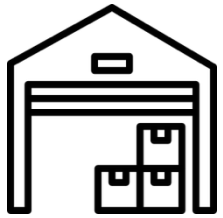

Storage

Transport

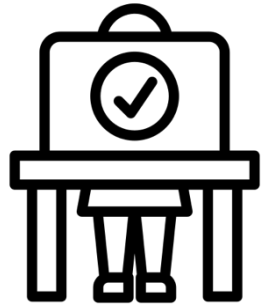

Real Loser → more votes

Real Winner → less votes

Election flipped!

Attack happens *before* election day.

Votes are classified by individual bubbles.

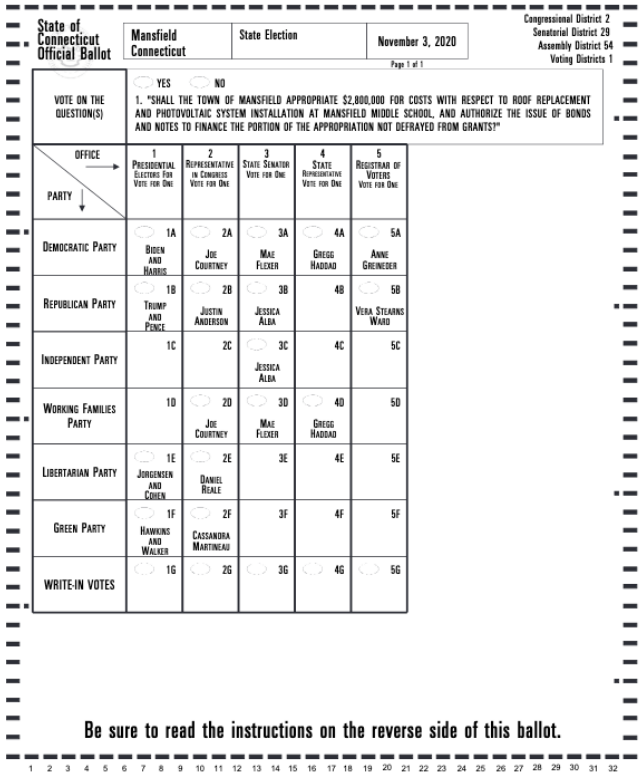
State of Connecticut Official Ballot

Mansfield Connecticut | State Election | November 3, 2020

Congressional District 2
Senatorial District 29
Assembly District 54
Voting Districts 1

Page 1 of 1

VOTE ON THE QUESTION(S)

YES NO

1. "SHALL THE TOWN OF MANSFIELD APPROPRIATE $2,800,000 FOR COSTS WITH RESPECT TO ROOF REPLACEMENT AND PHOTOVOLTAIC SYSTEM INSTALLATION AT MANSFIELD MIDDLE SCHOOL, AND AUTHORIZE THE ISSUE OF BONDS AND NOTES TO FINANCE THE PORTION OF THE APPROPRIATION NOT DEFRAYED FROM GRANTS?"

PARTY \ OFFICE	1 Presidential Electors For Vote for One	2 Representative in Congress Vote for One	3 State Senator Vote for One	4 State Representative Vote for One	5 Registrar of Voters Vote for One
Democratic Party	1A Biden and Harris	2A Joe Courtney	3A Mae Flexer	4A Gregg Haddad	5A Anne Greineder
Republican Party	1B Trump and Pence	2B Justin Anderson	3B Jessica Alba	4B	5B Vera Stearns Ward
Independent Party	1C	2C	3C Jessica Alba	4C	5C
Working Families Party	1D	2D Joe Courtney	3D Mae Flexer	4D Gregg Haddad	5D
Libertarian Party	1E Jorgensen and Cohen	2E Daniel Reale	3E	4E	5E
Green Party	1F Hawkins and Walker	2F Cassandora Martineau	3F	4F	5F
WRITE-IN VOTES	1G	2G	3G	4G	5G

Be sure to read the instructions on the reverse side of this ballot.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32

Real ballot from Nov. 2020 CT state election (town of Mansfield). Image source: ct.gov.

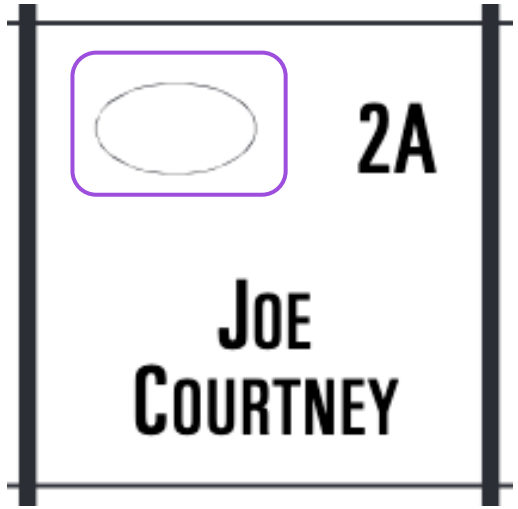

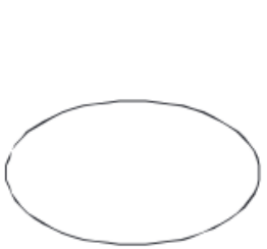

Empty bubble = non-vote

Marked bubble = vote

Binary image classification problem!

Existing works: [Barretto et al., 2021. "Improving the Accuracy of Ballot Scanners Using Supervised Learning", Bernhard et al., 2020. "Can Voters Detect Malicious Manipulation of Ballot Marking Devices?", Bernhard et al., 2019. "UnclearBallot: Automated Ballot Image Manipulation".]

Attack Phase One: adversarial images to fool classifier.

Goal: human-*imperceptibility*.

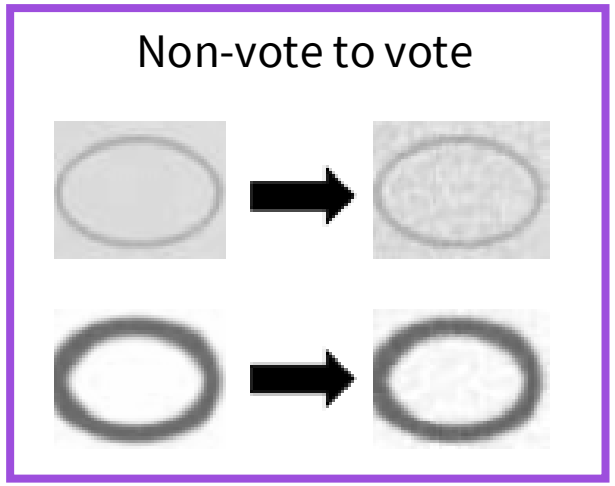

Attack specs:

- APGD (Auto-Projected Gradient Descent)
- L_∞ norm
- DLR loss
- ε = 8/255

Full attack experiments: (for non-vote to vote, and vote to non-vote)

- APGD, MIM, PGD, FGSM
- ε = {0, 4, 8, 16, 32, 64, 128, 255} / 255
- CE and DLR loss

[Croce et al., 2020. "Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks".]

Attack Phase Two: print tampered ballots (simulated).

Adversarial bubbles

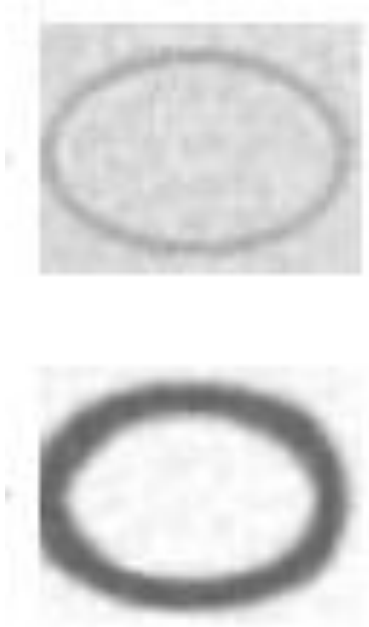

Image sources: Tecisoft, Amazon.

Not all models can resist this physical attack.

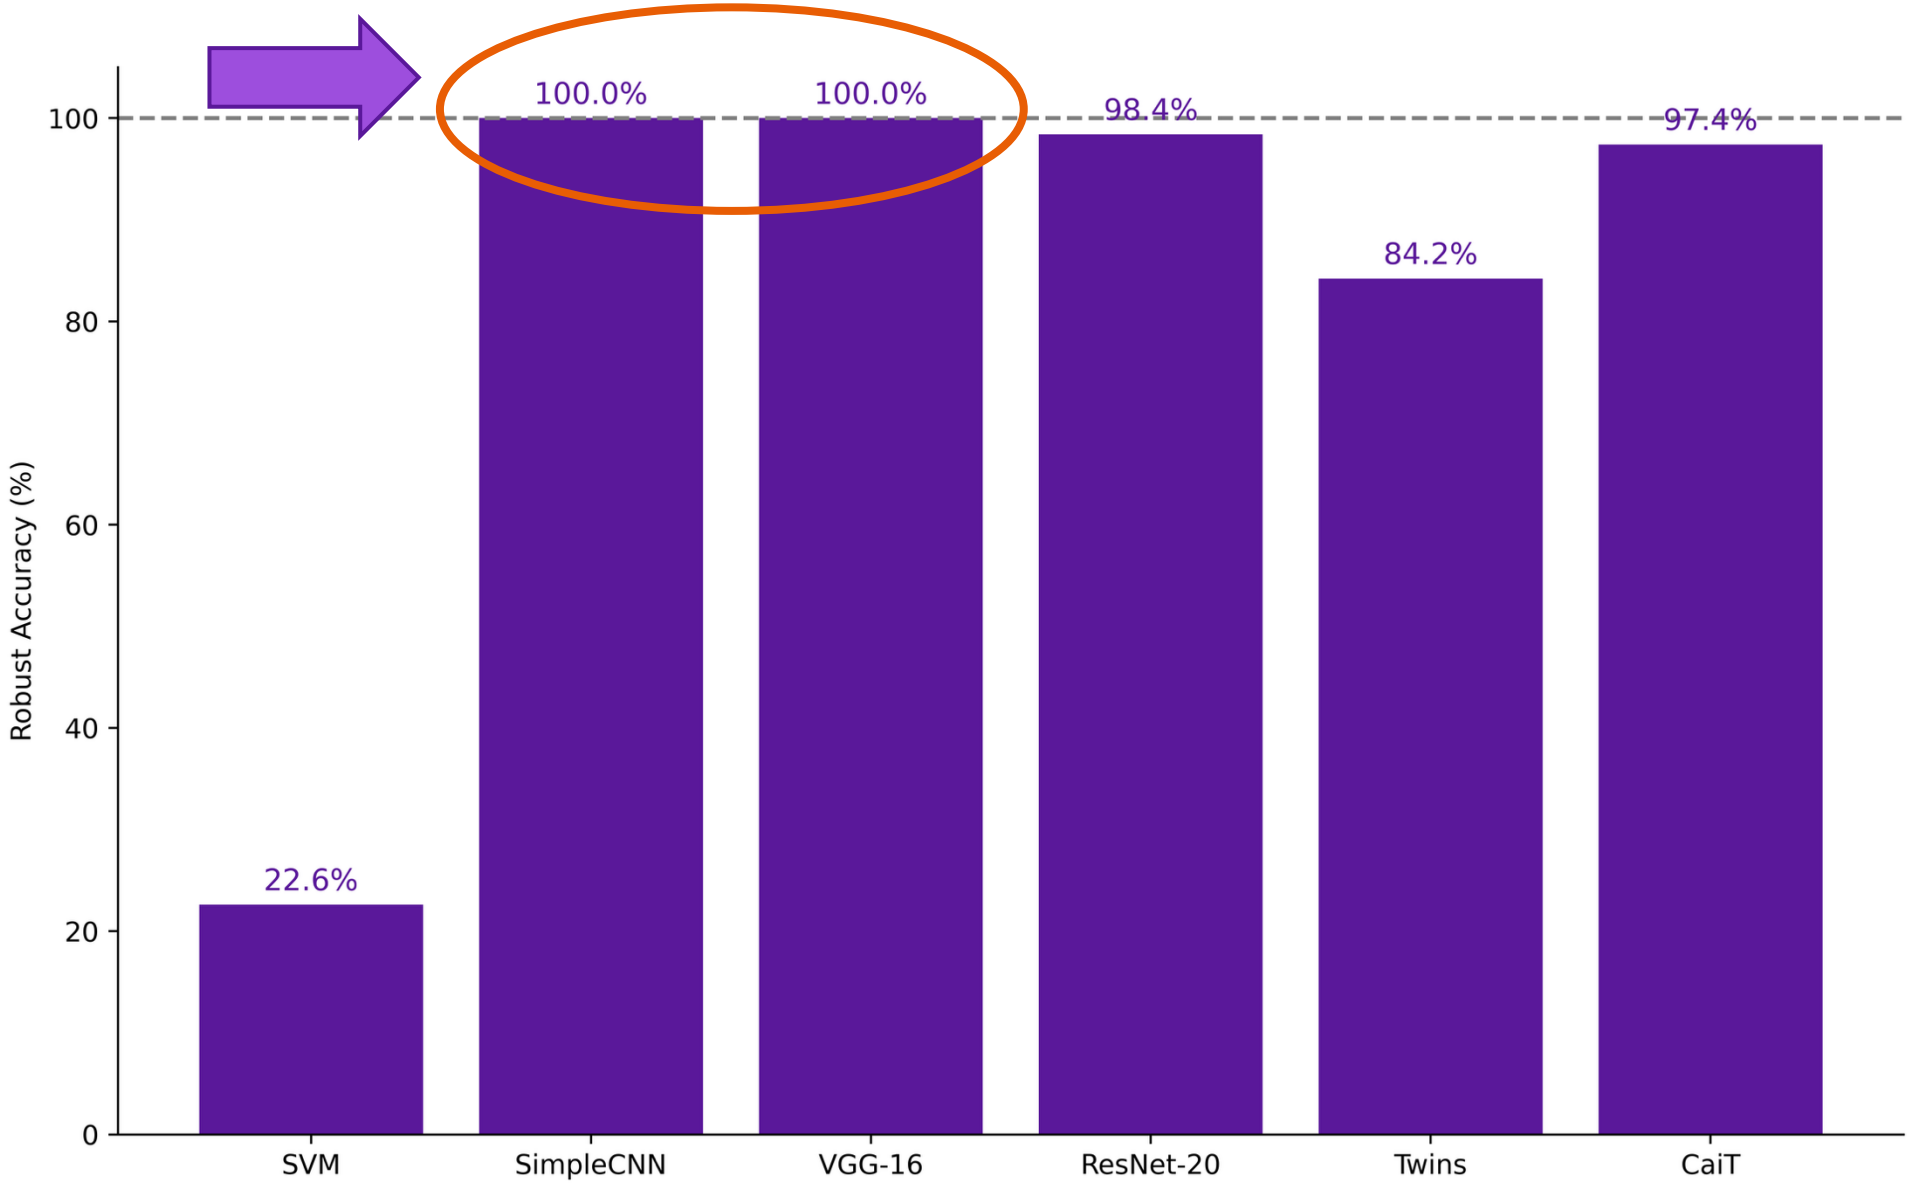

This is just a baseline – attacks only get better!

Adversary needs *just* 1 successful example.

Adversarial bubble

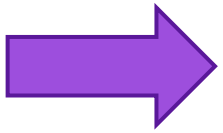

State of Connecticut Official Ballot | Mansfield Connecticut | State Election | November 3, 2020

Congressional District 2
Senatorial District 29
Assembly District 54
Voting Districts 1

Page 1 of 1

VOTE ON THE QUESTION(S)

YES NO

1. "SHALL THE TOWN OF MANSFIELD APPROPRIATE $2,800,000 FOR COSTS WITH RESPECT TO ROOF REPLACEMENT AND PHOTOVOLTAIC SYSTEM INSTALLATION AT MANSFIELD MIDDLE SCHOOL, AND AUTHORIZE THE ISSUE OF BONDS AND NOTES TO FINANCE THE PORTION OF THE APPROPRIATION NOT DEFRAYED FROM GRANTS?"

OFFICE / PARTY	1 PRESIDENTIAL ELECTORS FOR VOTE FOR ONE	2 REPRESENTATIVE IN CONGRESS VOTE FOR ONE	3 STATE SENATOR VOTE FOR ONE	4 STATE REPRESENTATIVE VOTE FOR ONE	5 REGISTRAR OF VOTERS VOTE FOR ONE
DEMOCRATIC PARTY	1A BIDEN AND HARRIS	2A JOE COURTNEY	3A MAE FLEXER	4A GREGG HADDAD	5A ANNE GREINEDER
REPUBLICAN PARTY	1B TRUMP AND PENCE	2B JUSTIN ANDERSON	3B JESSICA ALBA	4B	5B VERA STEARNS WARD
INDEPENDENT PARTY	1C	2C	3C JESSICA ALBA	4C	5C
WORKING FAMILIES PARTY	1D	2D JOE COURTNEY	3D MAE FLEXER	4D GREGG HADDAD	5D
LIBERTARIAN PARTY	1E JORGENSEN AND COHEN	2E DANIEL REALE	3E	4E	5E
GREEN PARTY	1F HAWKINS AND WALKER	2F CASSANDRA MARTINEAU	3F	4F	5F
WRITE-IN VOTES	1G	2G	3G	4G	5G

Be sure to read the instructions on the reverse side of this ballot.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32

Adversarial ballots can change the margin by 2.5%.

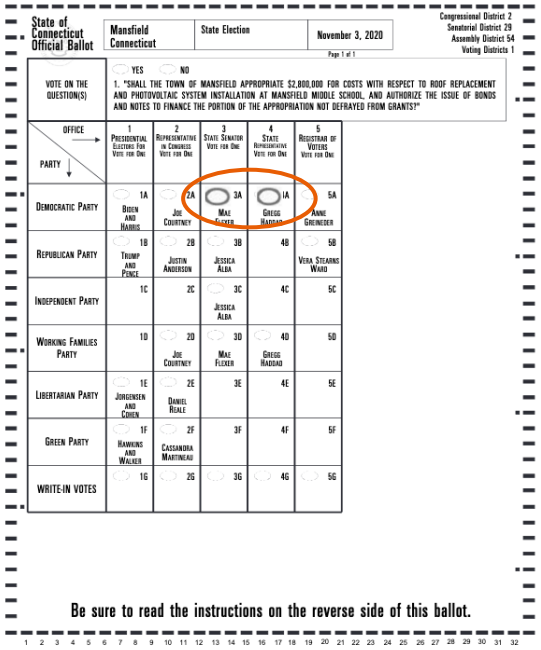

State of Connecticut Official Ballot

Mansfield Connecticut | State Election | November 3, 2020

Congressional District 2
Senatorial District 29
Assembly District 54
Voting Districts 1

Page 1 of 1

VOTE ON THE QUESTION(S)

YES NO

1. "SHALL THE TOWN OF MANSFIELD APPROPRIATE $2,800,000 FOR COSTS WITH RESPECT TO ROOF REPLACEMENT AND PHOTOVOLTAIC SYSTEM INSTALLATION AT MANSFIELD MIDDLE SCHOOL, AND AUTHORIZE THE ISSUE OF BONDS AND NOTES TO FINANCE THE PORTION OF THE APPROPRIATION NOT DEFRAYED FROM GRANTS?"

PARTY \ OFFICE	1 Presidential Electors For Vote for One	2 Representative in Congress Vote for One	3 State Senator Vote for One	4 State Representative Vote for One	5 Registrar of Voters Vote for One
Democratic Party	1A Biden and Harris	2A Joe Courtney	3A Mae Flexer	4A Gregg Haddad	5A Anne Greinecker
Republican Party	1B Trump and Pence	2B Justin Anderson	3B Jessica Alba	4B	5B Vera Stearns Ward
Independent Party	1C	2C	3C Jessica Alba	4C	5C
Working Families Party	1D	2D Joe Courtney	3D Mae Flexer	4D Gregg Haddad	5D
Libertarian Party	1E Jorgensen and Cohen	2E Daniel Reale	3E	4E	5E
Green Party	1F Hawkins and Walker	2F Cassandra Martinead	3F	4F	5F
Write-in Votes	1G	2G	3G	4G	5G

Be sure to read the instructions on the reverse side of this ballot.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32

Race left empty → 😈 ✅

Vote for preferred candidate → 😈 ✅

Voted for opponent → overvote*. 😈 ✅

*typically. Other possibilities: new ballot requested, ballot nullified.

Election day scenario:

- Local, 2-party race
- Tight margin of 2%
- 12% of ballots left blank
- Mail-in ballots only

	Real Results	**Adversarial Interference**
Real Loser	0.395 **(-2%)**	0.407 **(+0.5%)**
Real Winner	0.415	0.402

Final Takeaways

Just a *few* **tampered** ballots can **swing** a **close** election.

ML in voting → risky!

For trusted election security:

- NEVER replace humans in the loop!
- View physical world pipeline as part of defense.
- More transparency, and shared models/data!

Request to community:
How can we better defend these models for trustworthiness?

Existing works: [Barretto et al., 2021. "Improving the Accuracy of Ballot Scanners Using Supervised Learning", Bernhard et al., 2020. "Can Voters Detect Malicious Manipulation of Ballot Marking Devices?", Bernhard et al., 2019. "UnclearBallot: Automated Ballot Image Manipulation".]

Work-In-Progress

- Created new dataset, augmented with synthetic marks to model realistic voter behavior
- Investigating physical-world transformations and deformations
- Investigating defense mechanisms

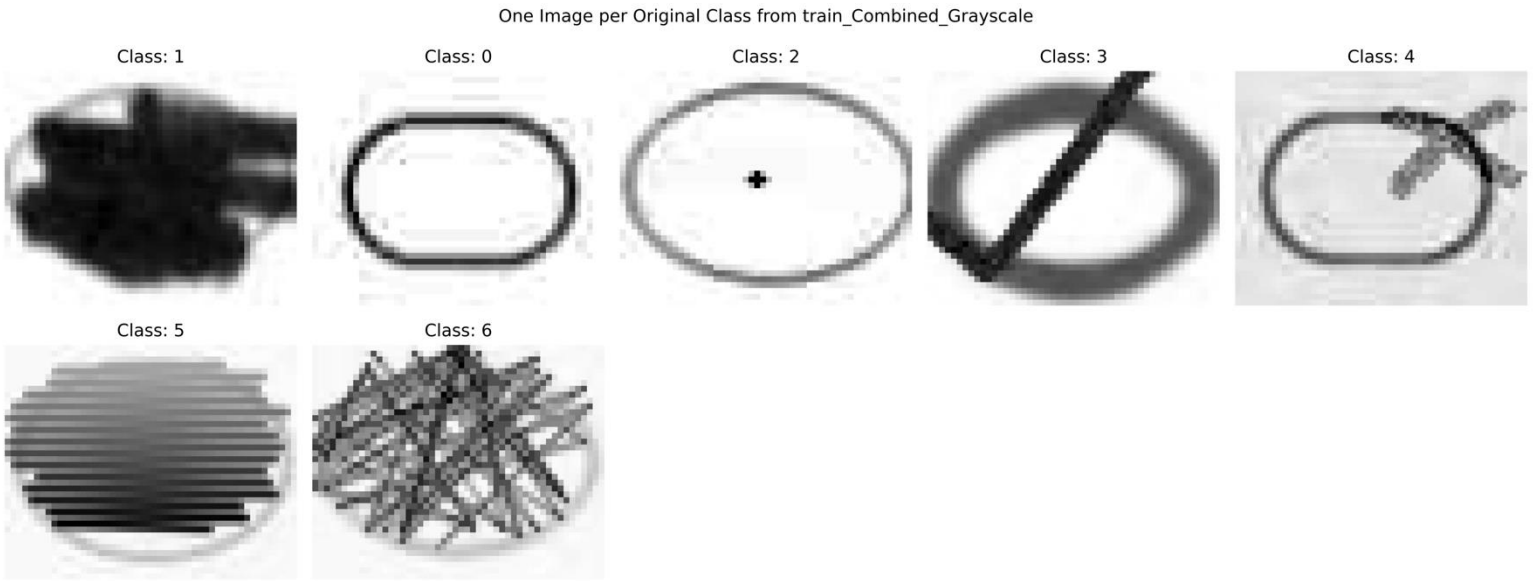

Thank you! Check out our work here:

Aayushi Verma

PhD student @UConn

aayushi.verma@uconn.edu

awesomecosmos

Our paper!

https://dl.acm.org/doi/10.1145/3719027.3744882

Backup Slides for Q&A

Work-In-Progress

- Question: how to reverse-engineer “ideal” adversarial bubble which survives physical world printing process
- Follow-up question: how to build defenses?