Detecting Abnormal User Behavior Through Pattern-mining Input Device Analytics

Ignacio X. Domínguez, Alok Goel, David L. Roberts, and Robert St. Amant

North Carolina State University

Abnormal

- Abnormal:[1]
  - Different from what is normal or average
  - **Unusual, especially in a way that causes problems**

- Practical examples of abnormal behavior detection:
  - Bots
  - Not proper attention to the task
  - Intrusion
  - Knowledge

[1] "abnormal." Merriam-Webster.com. 2015. http://www.merriam-webster.com (6 Apr 2015).

exile toy
the cake is a lie...
8699A070
horse
3V5GPQ
round
mark
234417
Wall St. Gets Urge to Merge
FREE
gimpy
spade
welcome
porter
spitfires
SMTL
solution
CAXA
DV2P
$1,223,017,697
2ACZA
163J
5B8581
April
sound
GOLD
vatman
rSUlmFyNb7sDW82Gje
bixe
there
glorious
exile
Punish cats
DESCRIBE what LOVE feels like IN 200 words.
January
PAWN
kill
1sharp long
WASHINGTON minate its small airport
WELL
5Z28AF
0 4
marcus
593744
die
bad
summoned
canvas
passkey
policial
war 22833
punchbroad
exile
get million
157
code
47.56
stabbings
4:30
1897
nite
cuteness
tame
game
moan
retard
Crucify
3:07:02
aeriest
s756
sage
ball
542642
HATE
7e5h4s
lime
sandwich
168441
june
vacation bf7wv3
punish now
flad
dead boy
sick fridge
occupy
base
808720
row now now
психиатрия понимает
trying
pancakes
99999
the cake is a lie...
Гладить
VmARg
wire
SXGPT
we unite thing
547692
erased
contribute
sage
KOHUCESA
V793TR
10 + 9
unyu!
ygrmxas
refresh
llyzatio
269806
madness
sparta!
known
UNDREAM
hfnhv
yum death
подчинись Джобсу
Propheten,
keith
628149
captcha
HRAI
slow
bazooka
प्राजापत्यं; we
shook
borsch
PAB
uob82Wp
iw62K
shi
kill
dshack
XKG TYH

Human Interactive Proofs (HIPs)

- Completely Automated Public Turing test to tell Computers and Humans Apart (CAPTCHA)
- Disruptive
- Adds cognitive burden to the user
- Single-point check
- Not applicable to every domain

Human Observational Proofs (HOPs)

- Observe behavior to make sure it looks like something a human would produce
- Problems solved:
  - Unobtrusive
  - Constant
- Can we do better?

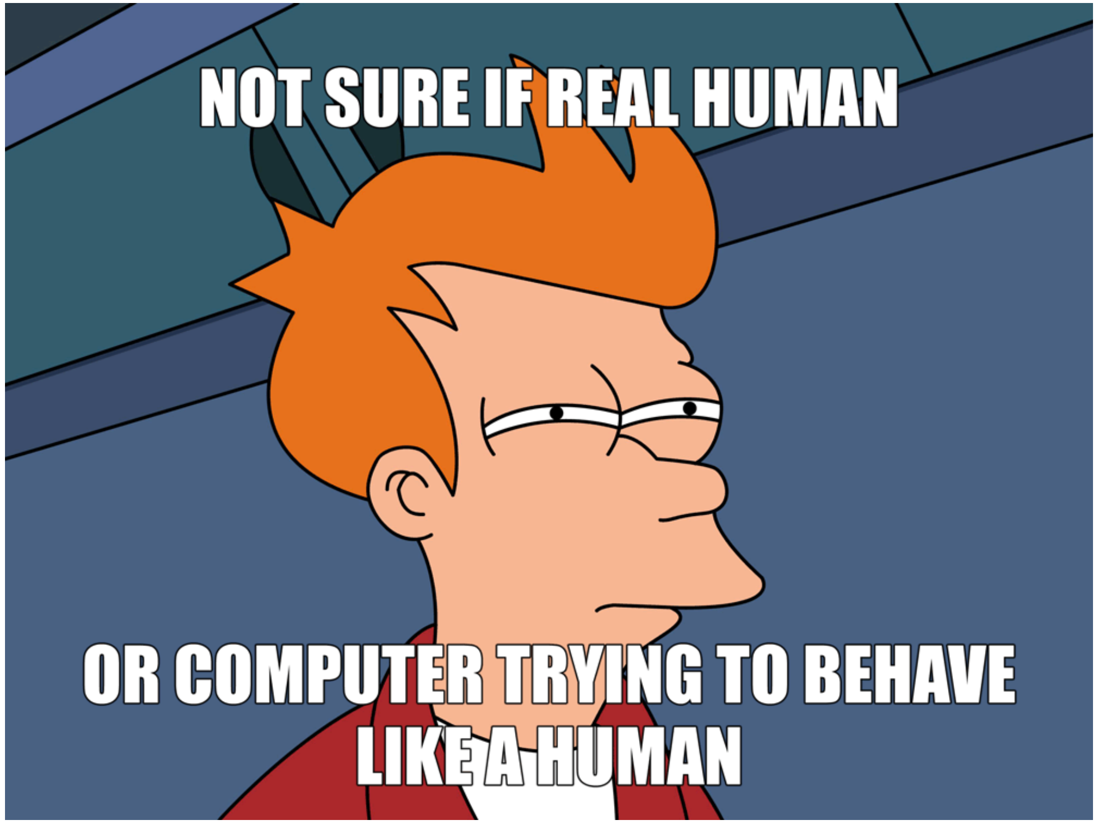

Idea: Human Subtlety Proofs

- Expand on the idea of HOPs
- Use cognitive models of interaction to classify behavior
- Improvements:
  - More precise
  - More expressive (can identify cognitive state)

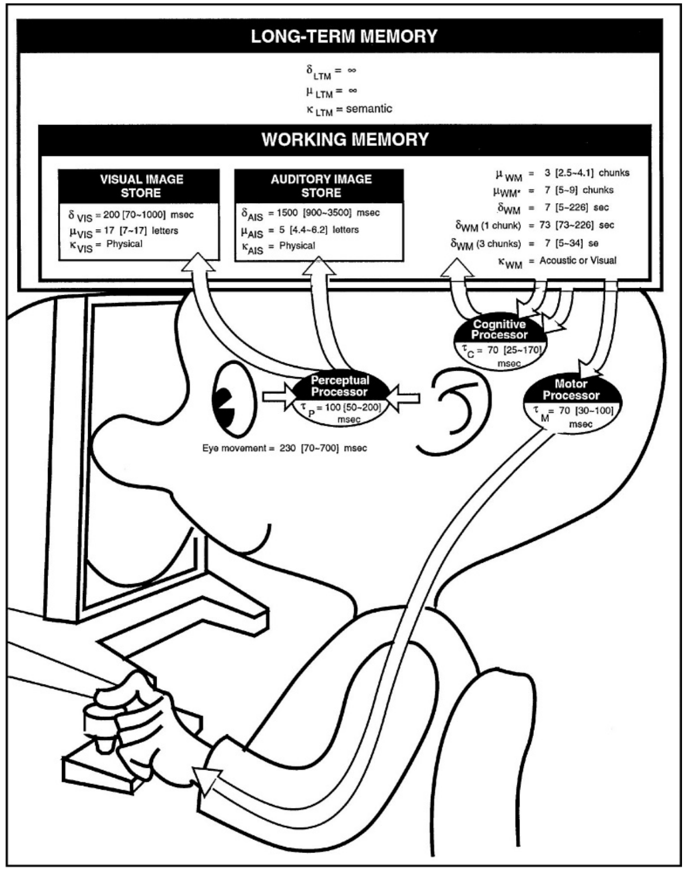
LONG-TERM MEMORY
$\delta_{LTM} = \infty$
$\mu_{LTM} = \infty$
κ_{LTM} = semantic
WORKING MEMORY
VISUAL IMAGE STORE
δ_{VIS} = 200 [70~1000] msec
μ_{VIS} = 17 [7~17] letters
κ_{VIS} = Physical
AUDITORY IMAGE STORE
δ_{AIS} = 1500 [900~3500] msec
μ_{AIS} = 5 [4.4~6.2] letters
κ_{AIS} = Physical
μ_{WM} = 3 [2.5~4.1] chunks
μ_{WM^*} = 7 [5~9] chunks
δ_{WM} = 7 [5~226] sec
δ_{WM} (1 chunk) = 73 [73~226] sec
δ_{WM} (3 chunks) = 7 [5~34] se
κ_{WM} = Acoustic or Visual
Cognitive Processor
τ_C = 70 [25~170] msec
Perceptual Processor
τ_P = 100 [50~200] msec
Motor Processor
τ_M = 70 [30~100] msec
Eye movement = 230 [70~700] msec

You Are Here

Hypotheses

Hypotheses

1. **Different cognitive processes** will translate into **differences in how people use input devices**

2. Those differences **cannot be hidden** by people, even if they try

Evaluation

Using games

- Simulating real-world complexities

  e.g., Ben Schneiderman's Direct manipulation[1]

- Tightly control variables

- Fun!

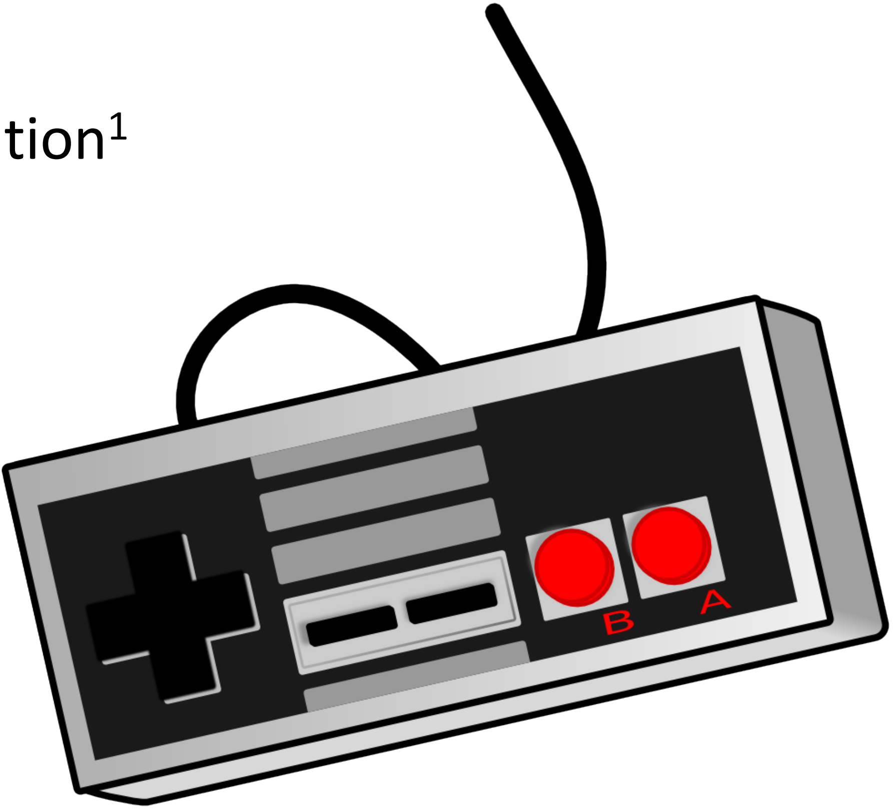

[1] Ben Shneiderman and Pattie Maes. 1997. Direct manipulation vs. interface agents. *interactions* 4, 6 (November 1997), 42-61. http://doi.acm.org/10.1145/267505.267514

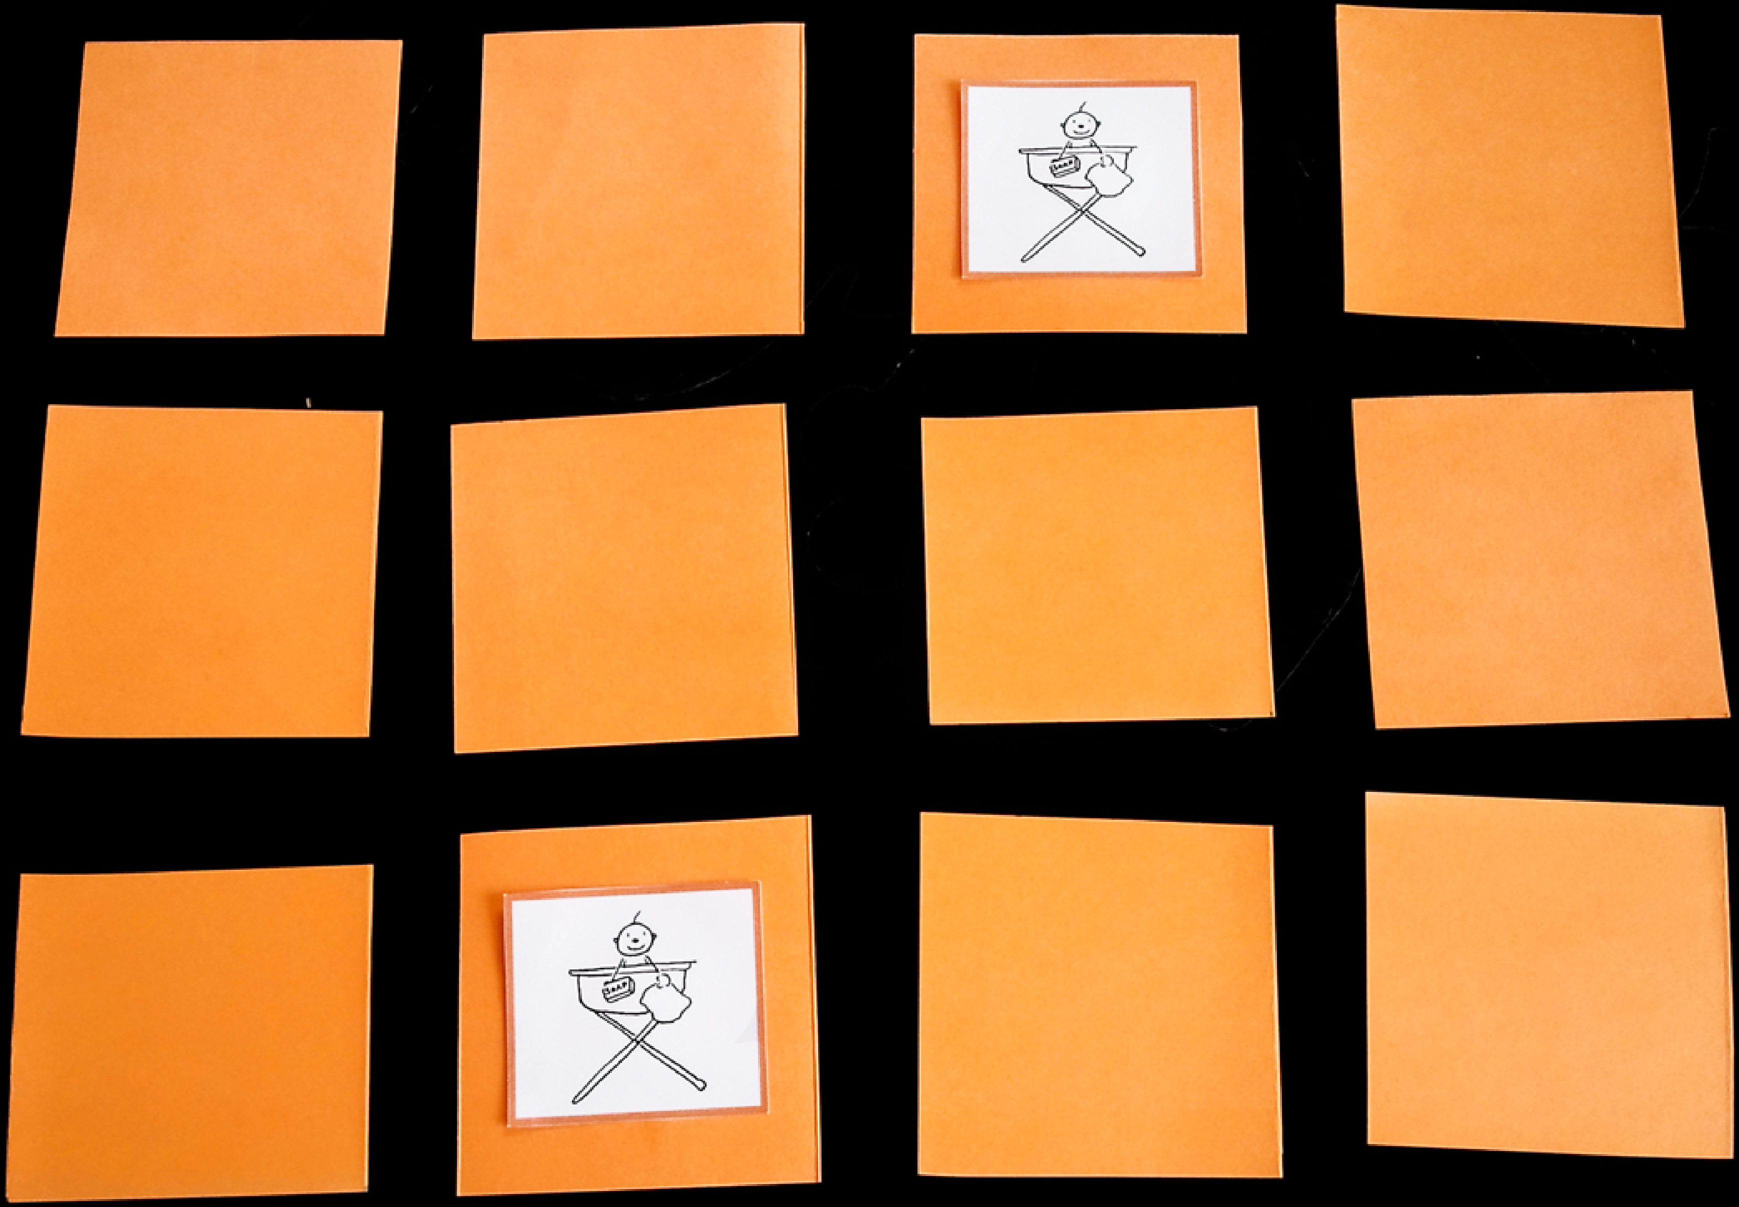

The Concentration Game

- Web-based (Flash)
- 16 (4x4) 100-pixel tiles
- Letters instead of pictures
  - Helvetica Neue LT Std 65 Medium
- Random positions

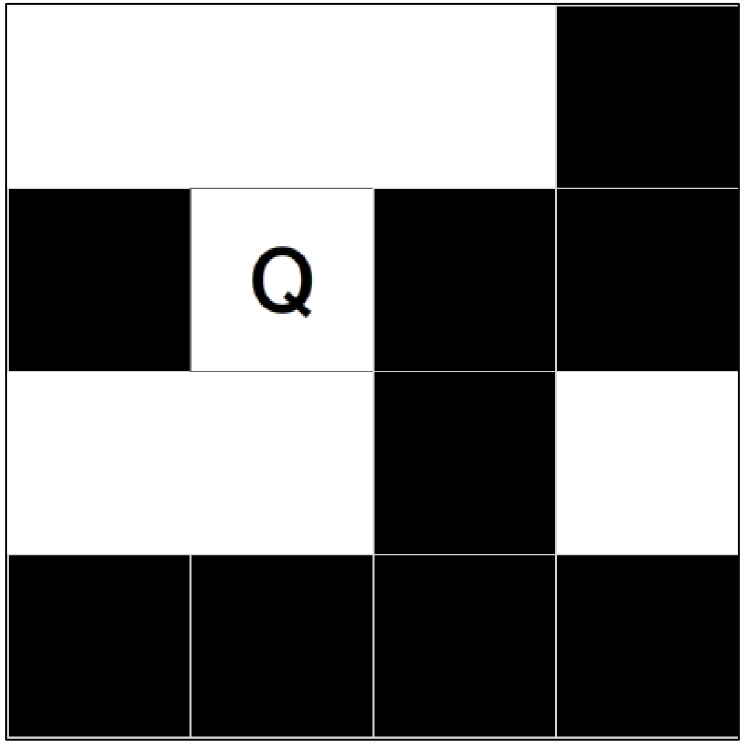

The Concentration Game With a Twist

Normal mode

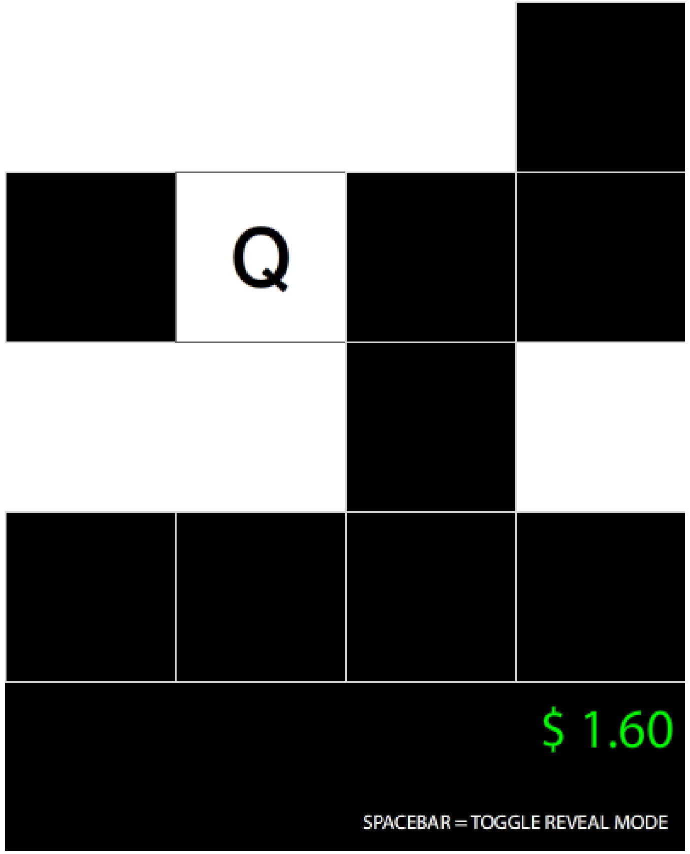

SPACEBAR = TOGGLE REVEAL MODE

Reveal mode

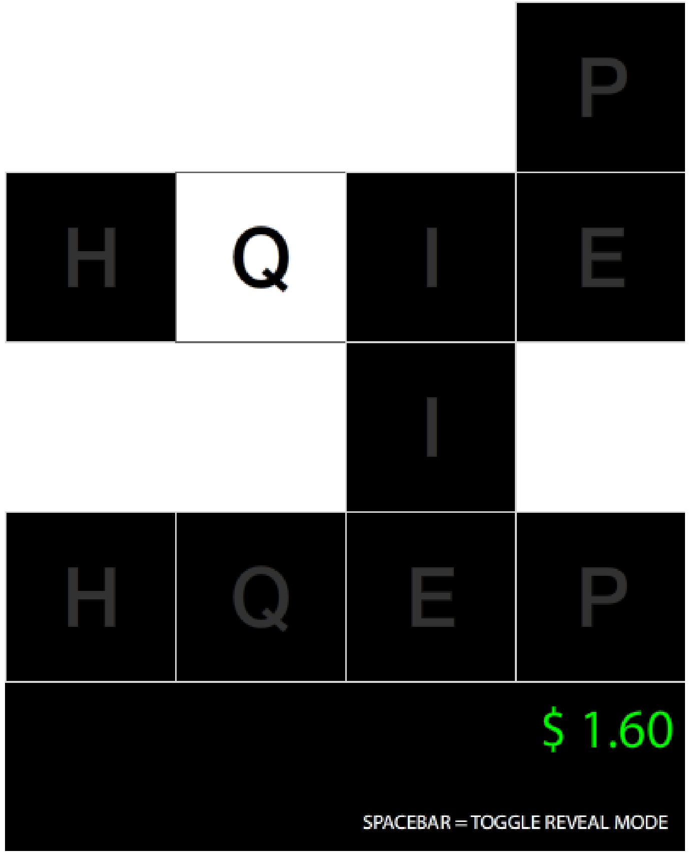

SPACEBAR = TOGGLE REVEAL MODE

Reveal Mode

- Does not interfere with mouse patterns (uses space bar)
- The same mechanics are required to accomplish the same goal
- Relies on visual search rather than on memory recall
  - Therefore, the cognitive process required to solve the task is different

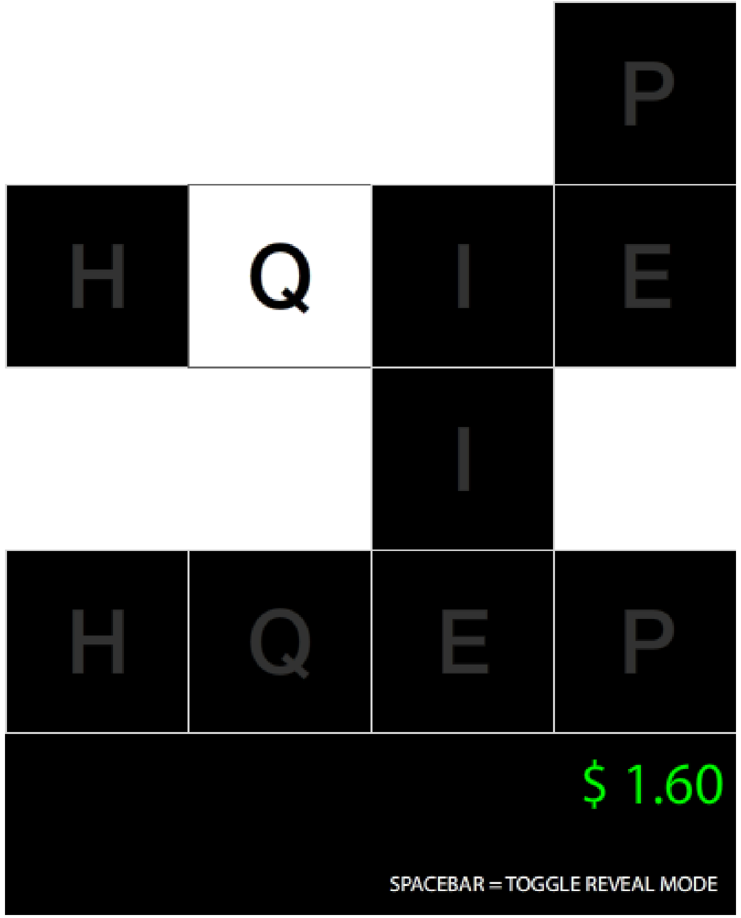

Experimental Conditions

1. Reveal mode disabled

2. Reveal mode discouraged

   Detection module enabled

3. Reveal mode encouraged

   Detection module disabled

4. Reveal mode enabled

   No mention of reveal mode or detection module in instructions

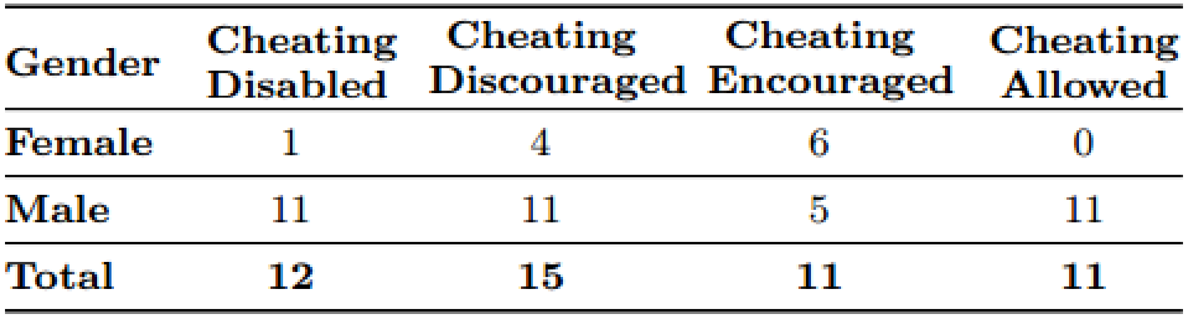

Gender	Cheating Disabled	Cheating Discouraged	Cheating Encouraged	Cheating Allowed
Female	1	4	6	0
Male	11	11	5	11
Total	**12**	**15**	**11**	**11**

Analysis and Results

Different Types of Round

- No reveal

  Reveal mode was never active during the round

- Full reveal

  Reveal mode was always active during the round

- Partial reveal

  Reveal mode was toggled at least once during the round

- Mixed reveal

  Full reveal + partial reveal

Different Types of Round

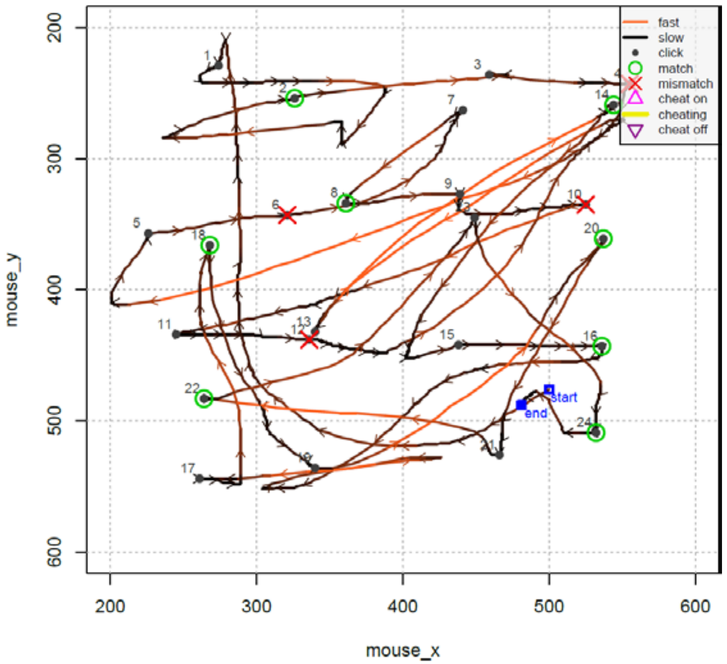

No reveal

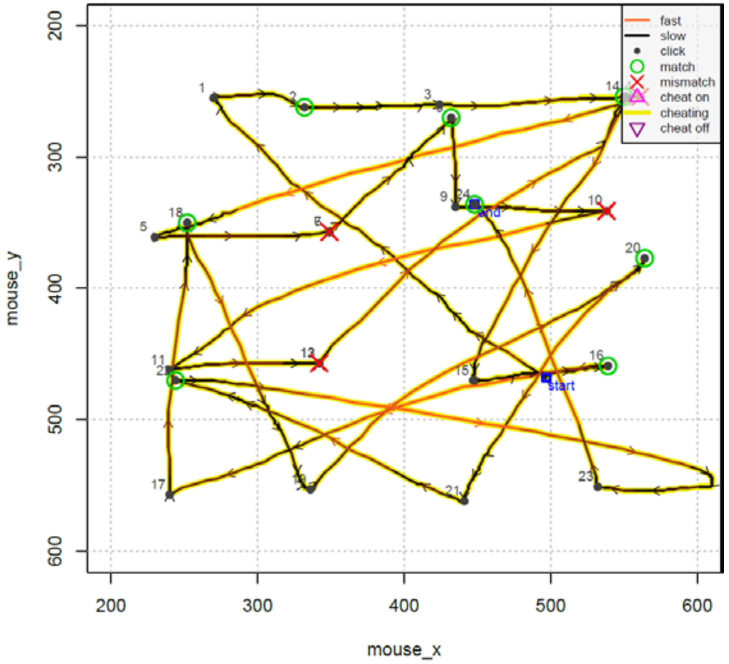

Full reveal

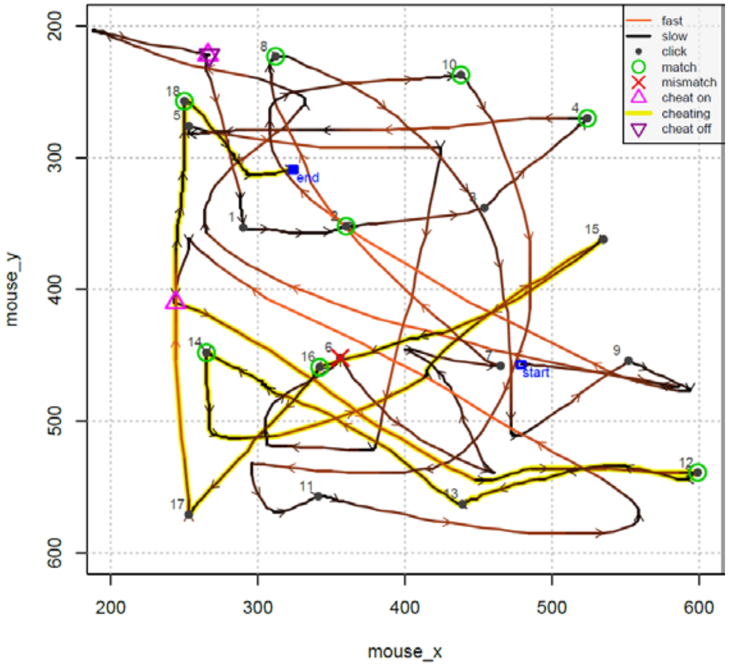

Partial reveal

Three Separate Analyses

- Analysis 1

  No reveal vs. Mixed reveal

- Analysis 2

  No reveal vs. Full reveal

- Analysis 3

  No reveal vs. Full reveal vs. Partial reveal

Method

- Random forest classifier
- 1000 estimators
- 10-fold cross-validation

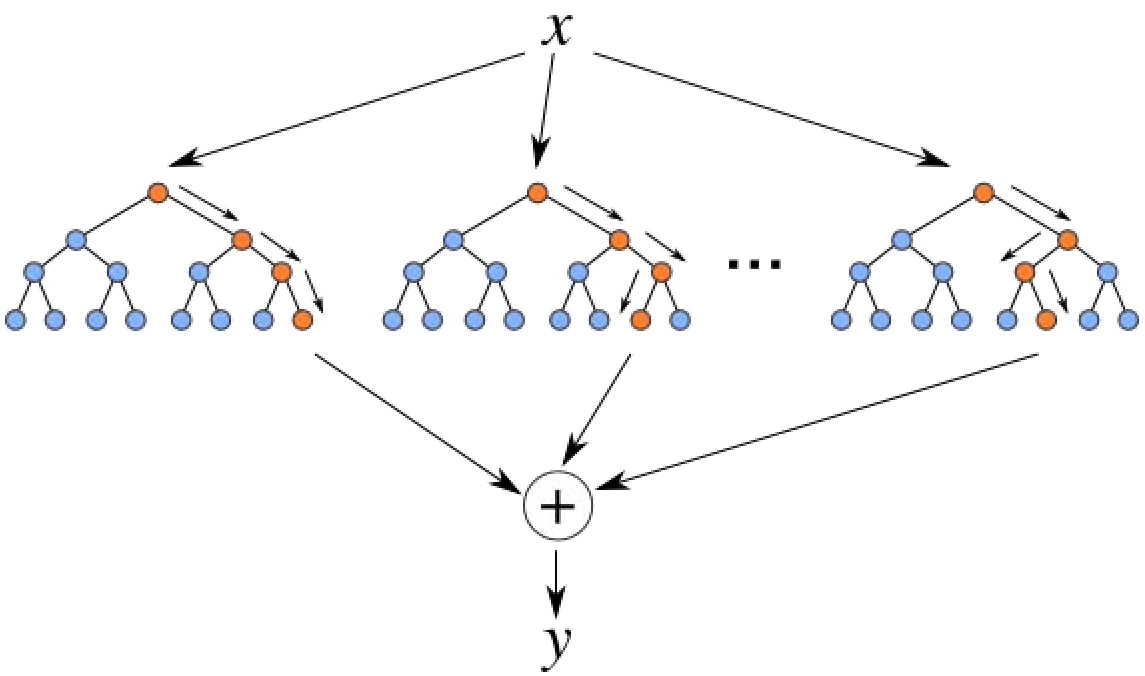

Attributes

- Time between clicks (ms)
- Time between a click and a succeeding mouse movement (ms)
- Count of change in direction of mouse motion
- Screen region hover count
- Task completion time (ms)
- Total number of clicks

Analysis 1

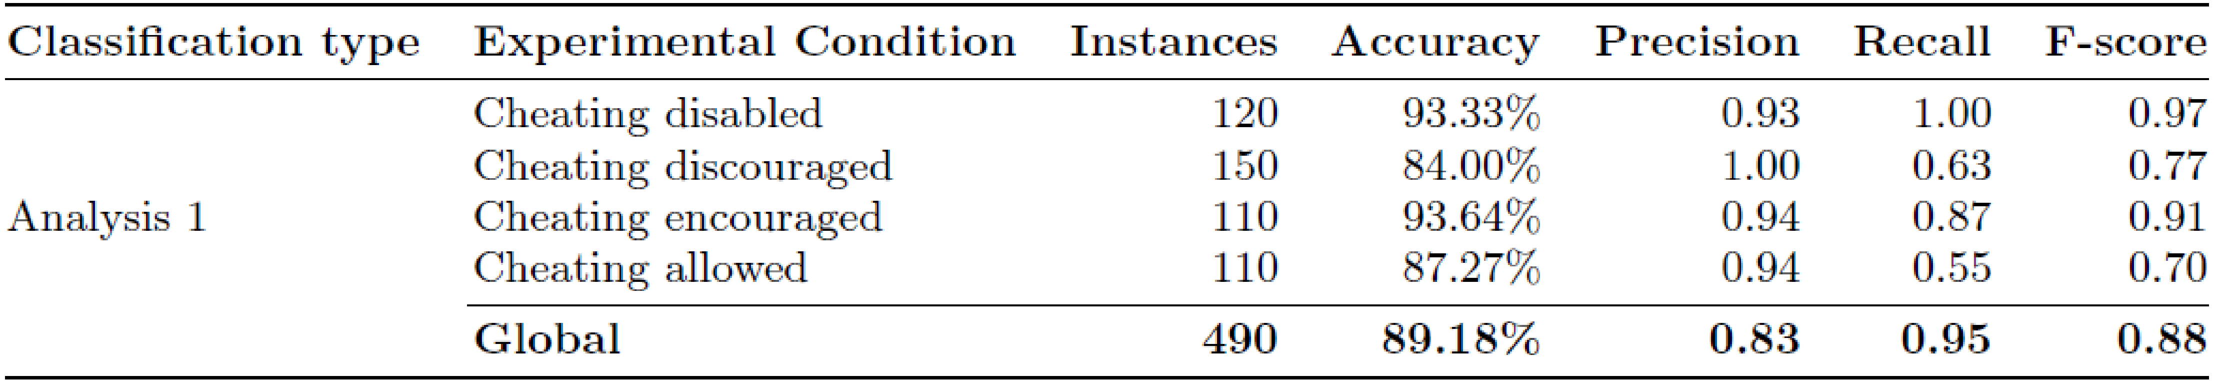

Classification type	Experimental Condition	Instances	Accuracy	Precision	Recall	F-score
Analysis 1	Cheating disabled	120	93.33%	0.93	1.00	0.97
	Cheating discouraged	150	84.00%	1.00	0.63	0.77
	Cheating encouraged	110	93.64%	0.94	0.87	0.91
	Cheating allowed	110	87.27%	0.94	0.55	0.70
	Global	**490**	**89.18%**	**0.83**	**0.95**	**0.88**

- Classes
  - No reveal (43.67%)
  - Mixed reveal (56.33%)

- Can detect different input device usage patterns (H1)
- Even if people try to hide their behavior, can still detect these patterns (H2)

Analysis 2

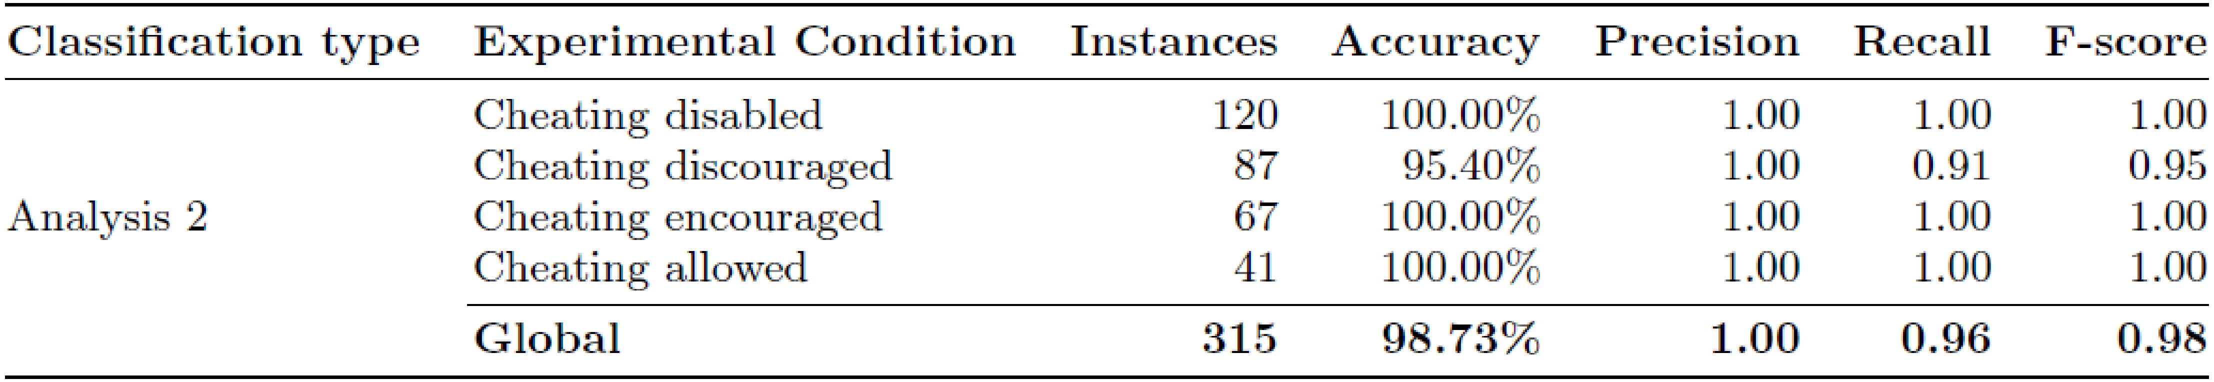

Classification type	Experimental Condition	Instances	Accuracy	Precision	Recall	F-score
Analysis 2	Cheating disabled	120	100.00%	1.00	1.00	1.00
	Cheating discouraged	87	95.40%	1.00	0.91	0.95
	Cheating encouraged	67	100.00%	1.00	1.00	1.00
	Cheating allowed	41	100.00%	1.00	1.00	1.00
	Global	**315**	**98.73%**	**1.00**	**0.96**	**0.98**

- Classes
  - No reveal (67.94%)
  - Full reveal (32.06%)

- More accurate than Analysis 1
  - 98.73% vs. 89.18%
  - Mixed reveal is more fuzzy
- A few false negatives - missed a few
- Can detect different input device usage patterns (H1)
- Even if people try to hide their behavior, can still detect these patterns (H2)

Analysis 3

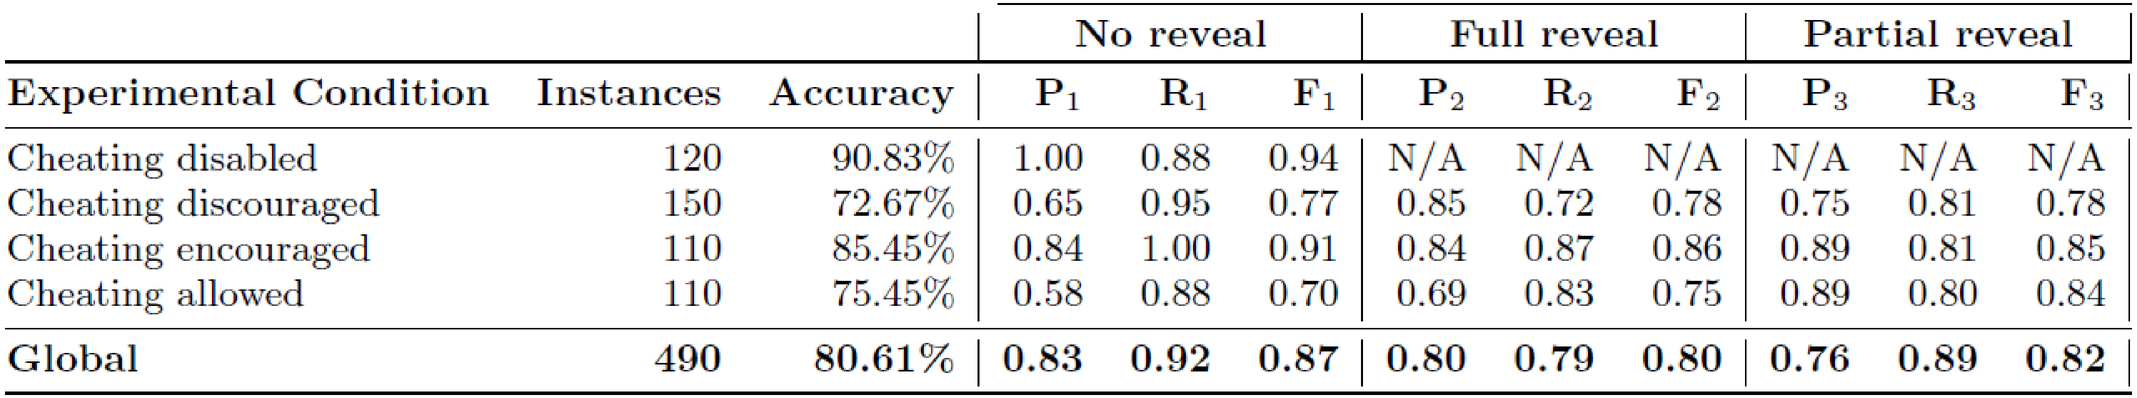

			No reveal			Full reveal			Partial reveal		
Experimental Condition	Instances	Accuracy	P_1	R_1	F_1	P_2	R_2	F_2	P_3	R_3	F_3
Cheating disabled	120	90.83%	1.00	0.88	0.94	N/A	N/A	N/A	N/A	N/A	N/A
Cheating discouraged	150	72.67%	0.65	0.95	0.77	0.85	0.72	0.78	0.75	0.81	0.78
Cheating encouraged	110	85.45%	0.84	1.00	0.91	0.84	0.87	0.86	0.89	0.81	0.85
Cheating allowed	110	75.45%	0.58	0.88	0.70	0.69	0.83	0.75	0.89	0.80	0.84
Global	**490**	**80.61%**	**0.83**	**0.92**	**0.87**	**0.80**	**0.79**	**0.80**	**0.76**	**0.89**	**0.82**

- Classes
  - No reveal (43.67%)
  - Full reveal (20.61%)
  - Partial reveal (35.71%)

- Can detect different input device usage patterns (H1)
- Even if people try to hide their behavior, can still detect these patterns (H2)

Limitations

- Not validated on other domains
- Only considers entire rounds
- Different tasks may produce interaction patterns that are difficult to differentiate
- Does not consider task-specific metrics

Human Subtlety Proofs: Reprise

- Expand on the idea of HOPs

- Use cognitive models of interaction to classify behavior

- Improvements:
  - More precise
  - More expressive (can identify cognitive state)

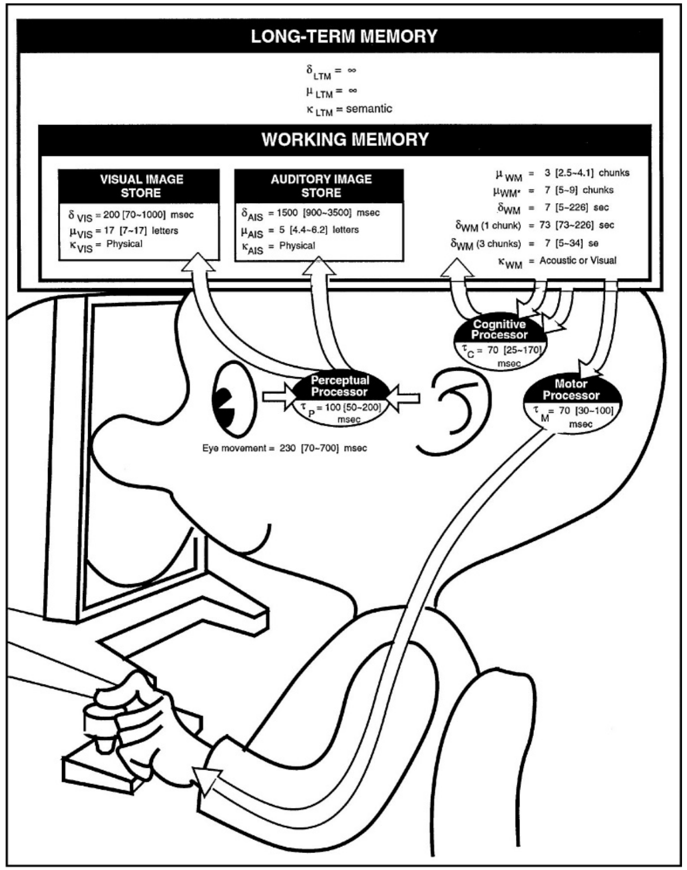

LONG-TERM MEMORY
$\delta_{LTM} = \infty$
$\mu_{LTM} = \infty$
κ_{LTM} = semantic
WORKING MEMORY
VISUAL IMAGE STORE
δ_{VIS} = 200 [70~1000] msec
μ_{VIS} = 17 [7~17] letters
κ_{VIS} = Physical
AUDITORY IMAGE STORE
δ_{AIS} = 1500 [900~3500] msec
μ_{AIS} = 5 [4.4~6.2] letters
κ_{AIS} = Physical
μ_{WM} = 3 [2.5~4.1] chunks
μ_{WM^*} = 7 [5~9] chunks
δ_{WM} = 7 [5~226] sec
δ_{WM} (1 chunk) = 73 [73~226] sec
δ_{WM} (3 chunks) = 7 [5~34] se
κ_{WM} = Acoustic or Visual
Cognitive Processor
τ_C = 70 [25~170] msec
Perceptual Processor
τ_P = 100 [50~200] msec
Motor Processor
τ_M = 70 [30~100] msec
Eye movement = 230 [70~700] msec

Conclusions and Future Work

Conclusions

- By introducing reveal mode, mouse interaction patterns changed

- We were able to detect these different mouse interaction patterns
  - This confirms Hypothesis 1

- When discouraging reveal mode, people who used it tried to conceal their behavior
  - We can still detect the use of reveal mode with high accuracy.
    - This confirms Hypothesis 2

Future Work

- See if accuracy is improved by including task-specific metrics
- Online detection
- Explore other domains
    - Same physical manifestations of cognitive processes?
    - More traditional tasks
- Explore other types of input devices
    - Typing game
    - Combinations of input devices

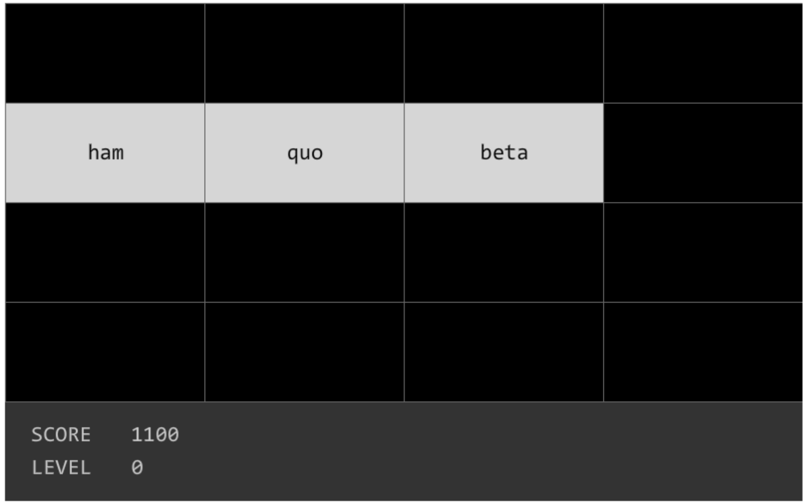

ham quo beta

SCORE 1100
LEVEL 0

Detecting Abnormal User Behavior Through Pattern-mining Input Device Analytics

Ignacio X. Domínguez, Alok Goel, David L. Roberts, and Robert St. Amant

{ignacioxd, agoel2}@ncsu.edu, {robertsd, stamant}@csc.ncsu.edu

Q & A

http://ciigar.csc.ncsu.edu/

Appendices

Descriptive Statistics

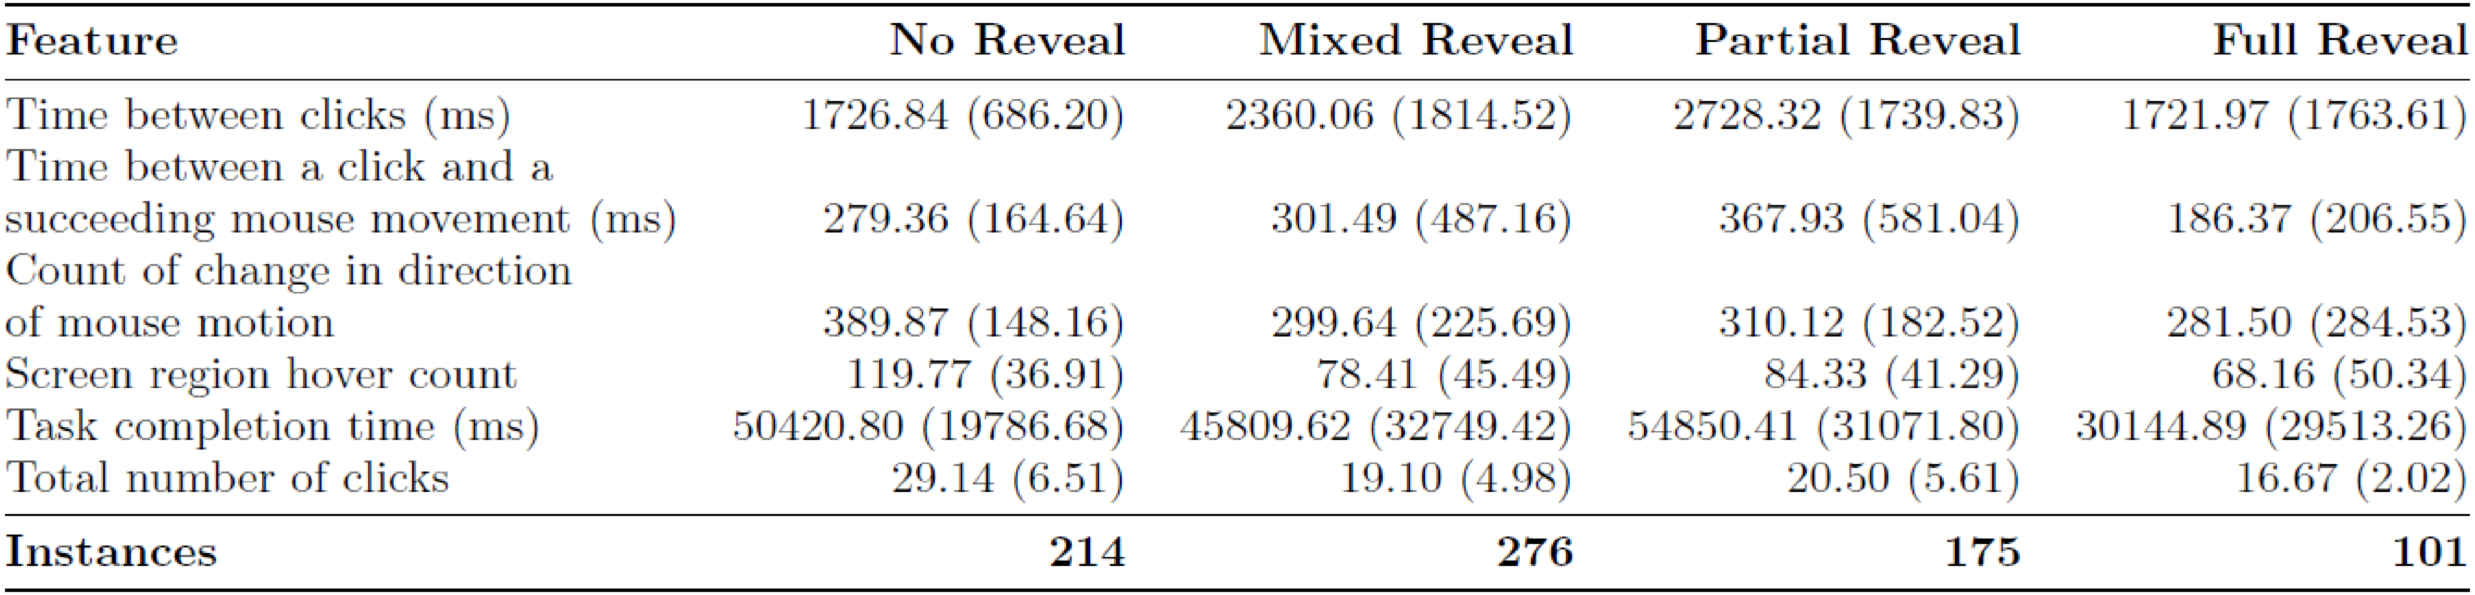

Feature	No Reveal	Mixed Reveal	Partial Reveal	Full Reveal
Time between clicks (ms)	1726.84 (686.20)	2360.06 (1814.52)	2728.32 (1739.83)	1721.97 (1763.61)
Time between a click and a succeeding mouse movement (ms)	279.36 (164.64)	301.49 (487.16)	367.93 (581.04)	186.37 (206.55)
Count of change in direction of mouse motion	389.87 (148.16)	299.64 (225.69)	310.12 (182.52)	281.50 (284.53)
Screen region hover count	119.77 (36.91)	78.41 (45.49)	84.33 (41.29)	68.16 (50.34)
Task completion time (ms)	50420.80 (19786.68)	45809.62 (32749.42)	54850.41 (31071.80)	30144.89 (29513.26)
Total number of clicks	29.14 (6.51)	19.10 (4.98)	20.50 (5.61)	16.67 (2.02)
Instances	**214**	**276**	**175**	**101**

Averages (and Std)

Data Distribution Across Classes

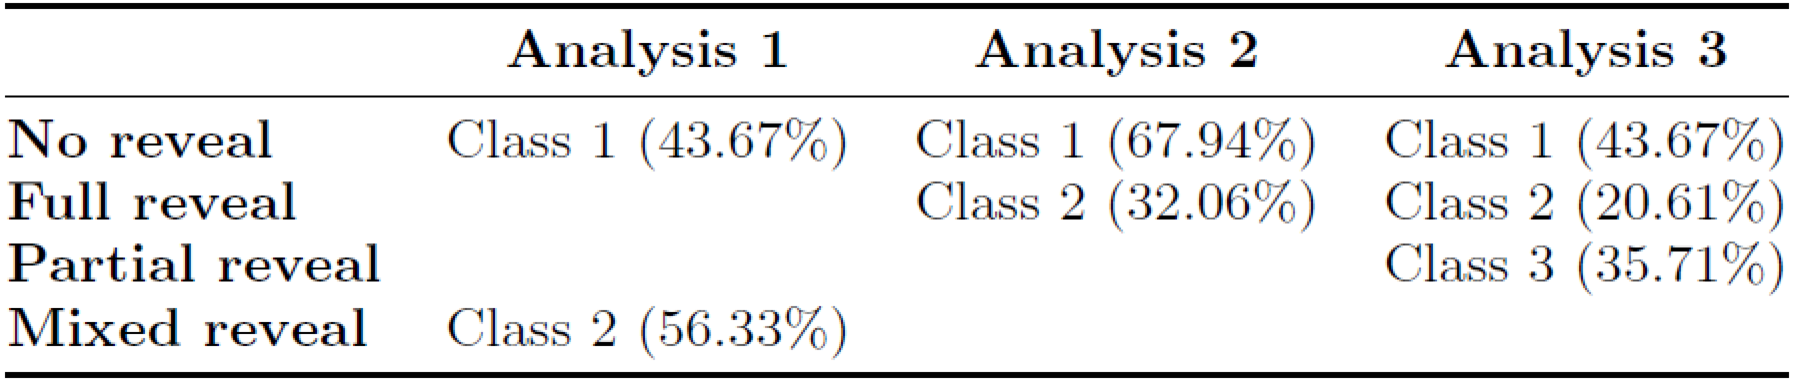

	Analysis 1	**Analysis 2**	**Analysis 3**
No reveal	Class 1 (43.67%)	Class 1 (67.94%)	Class 1 (43.67%)
Full reveal		Class 2 (32.06%)	Class 2 (20.61%)
Partial reveal			Class 3 (35.71%)
Mixed reveal	Class 2 (56.33%)		