

PREDICTING HUMAN DECISION-MAKING: FROM PREDICTION TO RECOMMENDATIONS

*DECISION MAKING AND RECOMMENDER SYSTEMS,
BOZEN-BOLZANO, ITALY 2018*




Ariel Rosenfeld, PhD
Dept. of Management
Bar-Ilan University
Israel
arielros1@gmail.com
www.arielrosenfeld.com

BASED ON THE BOOK

- Free just for you 😊

<http://tinyurl.com/predicting-human-DM>

- Bottom of the page...



AGENDA

- Introduction to human decision making
- Expert-Based Prediction
- Data-Based Prediction
- From Predictions to Recommendations
- Conclusions



INTRODUCTION TO HUMAN DECISION MAKING

What are exactly are we talking about?

HUMAN DECISION-MAKING

- Complex Cognitive process.
 - A lot of theoretical questions
 - A lot of practical implications
-
- Should we try and Understand or Predict?



UNDERSTANDING DM

- “...much have been learned about how and why we make decisions... yet, we are still discovering the tip of the iceberg...”


Rorie AE, Newsome WT. A general mechanism for decision-making in the human brain? *TRENDS in Cognitive Sciences*. 9(2):41-43 (2009).

Gold JI, Shadlen MN. The neural basis of decision making. *Annual Reviews in Neuroscience*. (30):535-574 (2007).

Feng S, Holmes P, Rorie A, Newsome WT. Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLOS Computational Biology*. 5(2): (2009).

Kiani R, Shadlen MN. Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*. 324(5928):759-764 (2009).





J. P. Gallivan, D. A. McLean, K. F. Valyear, C. E. Pettypiece, J. C. Culham. **Decoding Action Intentions from Preparatory Brain Activity in Human Parieto-Frontal Networks.** *Journal of Neuroscience*, 2011;

PREDICTING DM

- Easier to validate.
- Observable - No intrusive investigation.
- Great theoretical and practical benefit
 - To the Agents community:
 - Enhancing Human Interaction with Software\Robots...
 - Training People
 - Replacing People
 - Supporting People
 - Learning from People
 - ...



WHAT DRIVES HUMAN DM?

- Among others...
 - **Past experience** (Juliussan, Karlsson, & Gärling, 2005, Sagi, & Friedland, 2007)
 - **MANY Cognitive biases** (e.g., Marsh, & Hanlon, 2007; Nestler. & von Collani, 2008; Stanovich & West, 2008; West et al., 2008, Epley, & Gilovich, 2006).
 - **Individual differences such as Age, cognitive abilities, gender...** (de Bruin, Parker, & Fischhoff, 2007; Finucane, Mertz, Slovic, & Schmidt, 2005, Reed, Mikels, & Simon, 2008)
 - **Decision Complexity** (Goldstein & Gigerenzer, 2002; Hilbig & Pohl, 2008)
 - **Social aspects** (Acevedo and Krueger , 2004).


ARE PEOPLE PREDICTABLE?

- **Expert-Driven**

- Decision-Theory
- Game-Theory
- Etc.

- **Data-Driven**

- Statistics
- Machine learning
- Etc.





EXPERT-DRIVEN PARADIGM

Using normative rules, expert knowledge, behavioral sciences,
etc...

SINGLE DECISION MAKER (DECISION THEORY)

- Decision Theory =
Probability theory + Utility Theory
(deals with chance) (deals with outcomes)
- Fundamental idea
 - The MEU (Maximum expected utility) principle
 - Weigh the utility of each outcome by the probability that it occurs




MULTIPLE DECISION MAKERS (GAME THEORY)

Game Theory by
Maschler et al. 2013.

- The mathematical theory of **interaction between self-interested agents (“players”)**.
- Each player must consider how each other player will act in order to make its optimal choice: hence **strategic** considerations
- If a system has a single designer/owner, then game theoretic analysis is probably inappropriate



John von
Neumann



John
Nash




Heinrich
Freiherr
von
Stackelberg

EXAMPLE

- Trump and Clinton meet in a presidential debate
- They must each choose between **debating issues** or **making insults**.
- What should Clinton do. . . ?
- How well she is perceived to do will depend (in part) on the choice Trump makes...
- What are the possible outcomes here? How do the candidates rank them?

NASH EQUILIBRIUM



- A *strategy profile* is a list (s_1, s_2, \dots, s_n) of the strategies each player is using
- If each strategy is a best response given the other strategies in the profile, the profile is a *Nash equilibrium*
- Why is this important?
 - If we assume players are rational, they will play Nash strategies
 - Even less-than-rational play will often *converge* to Nash in repeated settings.
- Movie!

Pure Nash Equilibrium



	Issues	Insults
Issues	0,0	-2,+2
Insults	+2,-2	+1, -1


Mixed Nash Equilibrium



	Issues	Insults
Issues	-1,+1	+1,-1
Insults	+1,-1	-1, +1

EXERCISE

- <https://www.youtube.com/watch?v=x5Q6-wMx-K8>
- Movie - Rock Paper Scissors Lizard Spock.



	Rock	Paper	Scissors	Lizard	Spock
Rock	0, 0	-1, 1	1, -1	1, -1	-1, 1
Paper	1, -1	0, 0	-1, 1	-1, 1	1, -1
Scissors	-1, 1	1, -1	0, 0	1, -1	-1, 1
Lizard	-1, 1	1, -1	-1, 1	0, 0	1, -1
Spock	1, -1	-1, 1	1, -1	-1, 1	0, 0

DOES GT REALLY WORK?

- In SOME cases where*:
 1. **Social norms don't play** (e.g., when **incentives** are sufficiently large, then they can override norms).
 2. The game is **sufficiently simple**.
 3. **Sufficient experience** (e.g., opportunity for trial-and-error learning).

* Ken Binmore, Does Game Theory Work?, MIT Press, 2007.

ARE PEOPLE RATIONAL?



**Daniel
Kahneman**



**Israel
Aumann**



**Vernon
Smith**

LET'S TRY IT OURSELVES!



WHAT HAPPENED HERE?



- You go to see the new StarWars movie. Naturally, you want to buy popcorn (5\$).

What if -


1. You go to the popcorn stand and once you open your wallet you notice that a 5\$ bill that you put before is missing. You must have dropped it on the way to the movies (no way to recover it).
2. A few moments after purchasing the popcorn you bump into a friend and your popcorn is spilled all over on the floor.

- Would you buy popcorn?
 - Given that you have another 5\$ at your disposal.




TWO DECISION MAKERS


- The Ultimatum game



● Proposer ● Responder




● Proposer
● Responder





WHAT ARE WE MISSING?

- Irrationalities attributed to
 - sensitivity to context
 - lack of knowledge of own preferences
 - the effects of complexity
 - the interplay between emotion and cognition
 - the problem of self control
 - Etc.



IRRATIONALLY PREDICTABLE?

- To a certain extent – YES!
- Behavioral sciences provide empirical observations and explanatory theories.
- Most models specify general criteria that are context sensitive but usually **do not provide specific parameters or mathematical definitions.**



IRRATIONALLY PREDICTABLE?

- To a certain extent – NO!
- So why not extend the behavioral models?
- Extending classic normative models such as the Prospect Theory requires additional non-trivial assumptions and/or parameters.
- Let's give it a try...

IRRATIONALITY

- Kahneman and Tversky (1979) identified the most important deviations from the assumption that people maximize expected return.
- Prospect theory (extended to cumulative prospect theory), 1992
- Reflection effect.
- Overweighting of rare events.
- Loss aversion.

BOUNDED RATIONALITY

- Quantal Response Equilibrium [McKelvey & Palfrey 1995]
- Level- k [Costa-Gomes et al. 2001]
- Cognitive Hierarchy [Camerer et al. 2004]
- Noisy introspection [Goeree & Holt 2004]
- Quantal Lk, Quantal CH [Stahl & Wilson 1994; Camerer et al.]



QUANTAL RESPONSE

- Highly useful.
- A lot of experimental evidence support it.

Perfect Response: $EU^{adversary}(j) = CaptureProb \times Penalty + (1 - CaptureProb) \times Reward$


Quantal Response(QR) [McFadden 73]: Stochastic Choice, Better Choice More likely

$$\text{Adversary's probability of choosing target } j = \frac{e^{\lambda \cdot (EU^{adversary}(x, j))}}{\sum_{j'=1}^T e^{\lambda \cdot (EU^{adversary}(x, j'))}}$$




- **Best response:** Maximum utility action is always played
- **Quantal** (“softmax”) **response:** High-utility actions played often, low-utility actions played rarely

DIFFERENT VALUES OF LAMBDA



HEURISTICS




- Rules of thumb.
- Simple heuristics can explain some of the anomalies that motivated the development of normative approaches in the first place (Brandstätter, Gigerenzer, and Hertwig, 2006).
 - E.g., the *get-something effect* (where people prefer lotteries where values greater than zero are guaranteed) can be easily explained by a simple heuristic (e.g., Venkatraman, Payne, & Huettel, 2014).

HEURISTICS

- Heuristics serve as a framework in which satisfactory decisions are made quickly and with ease (Shah & Oppenheimer, 2008).
- In various cases, produce the optimal results (Nokes & Hacker, 2007).

ALBERT IT DOES MOVE

- Despite the popularization of
Machine Learning approaches...
- Decision Theory
- Game theory
- Psychological models
- Heuristics



Galileo Galilei

NORMATIVE PREDICTION

- Modeling tool
 - E.g., Security Games
- Recommendation tool
 - E.g., Suboptimal advice provision

SECURITY GAMES

- A major success story!

LAX AIRPORT CASE



Milind Tambe

Eight Inbound Roads, Eight Terminals: Limited Staff, Canines

Where and when to set up checkpoints?

Where and when to do canine patrols?



Pita, James, et al. "Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport.", 2008.

ARMOR: DEPLOYED AT LAX 2007

- “Assistant for Randomized Monitoring Over Routes”
 - Problem 1: Schedule vehicle checkpoints
 - Problem 2: Schedule canine patrols
- Randomized schedule: (i) target weights; (ii) surveillance

ARMOR-Checkpoints



ARMOR-K9



STACKELBERG SECURITY GAMES (SSGS): DEFENDER VS ADVERSARY DEFENDER'S OPTIMAL RANDOMIZED STRATEGY



Adversary

	Terminal #1	Terminal #2
Terminal #1	5, -3	-1, 1
Terminal #2	-5, 5	2, -1

LAX BASED GAME

Stackelberg security games

➔ *Defender (rational)*

- Commit to a strategy first

➔ *Adversary (bounded rational)*

- Observe defender's strategy
- Attack one of targets

Gates	Gate 1	Gate 2	Gate 3	Gate 4	Gate 5	Gate 6	Gate 7	Gate 8
Your Rewards	10	1	9	2	3	10	2	4
Your Penalties	-1	-5	-9	-2	-5	-8	-5	-3
Probability of No Guard	0,95	0,70	0,50	0,65	0,70	0,35	0,65	0,50
Probability of Guard	0,05	0,30	0,50	0,35	0,30	0,65	0,35	0,50
Guards' Rewards	5	6	2	8	4	2	1	4
Guards' Penalties	-2	-3	-2	-2	-3	-3	-3	-2

Game Interface



- ☐ Must Be Scheduled
- ☐ Must Not Be Scheduled
- ☒ At Least One Scheduled
- ☐ Unrestrict

Apply

Manually adjust the generated schedule

Add

or

Remove

<

November, 2007

>

Sun	Mon	Tue	Wed	Thu	Fri	Sat
28	29	30	31	1	2	3
4	5	6	7	8	9	10
11	12	13				
	19	20	21	22	23	24
25	26	27	28	29	30	1
2	3	4	5	6	7	8

Today: 11/13/2007

Set First Day of Schedule

Days to Schedule

7

Patrols to Schedule

0

Simultaneous Patrols

1

Randomness: Uncalculated

Tuesday

Wednesday

Thursday

Friday

Saturday

Sunday

Monday


At-Least

Checkpoint #:	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	
8:00-10:00 AM																																			none
10:00-12:00 AM																																			1
12:00-2:00 PM																																			2
2:00-4:00 PM																																			
4:00-6:00 PM																																			
6:00-8:00 PM																																			
8:00-10:00 PM																																			

Remove At-Least

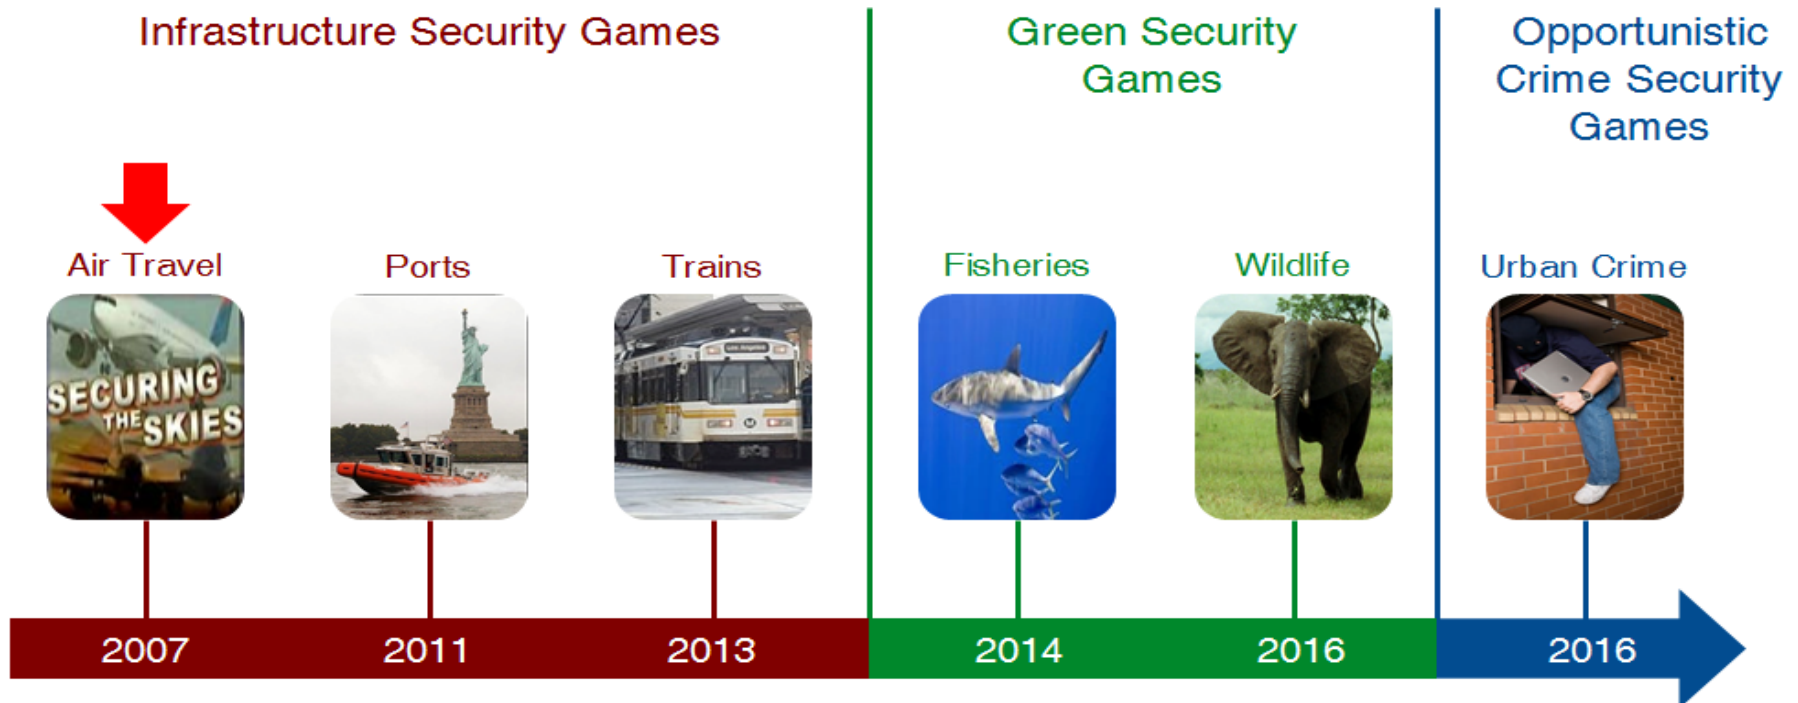
ARE ATTACKERS REALLY RATIONAL?

- Not completely.
- AMT workers – Not even expected. ..
- BUT – even security experts don't play according to the Stakelberg equilibrium.



SECURITY GAMES STATUS

- **Game Theoretic underpinnings prevail.**
- **The focus shifted from assuming “classic” rationality to bounded rationality behavior models (e.g., SUQR)**



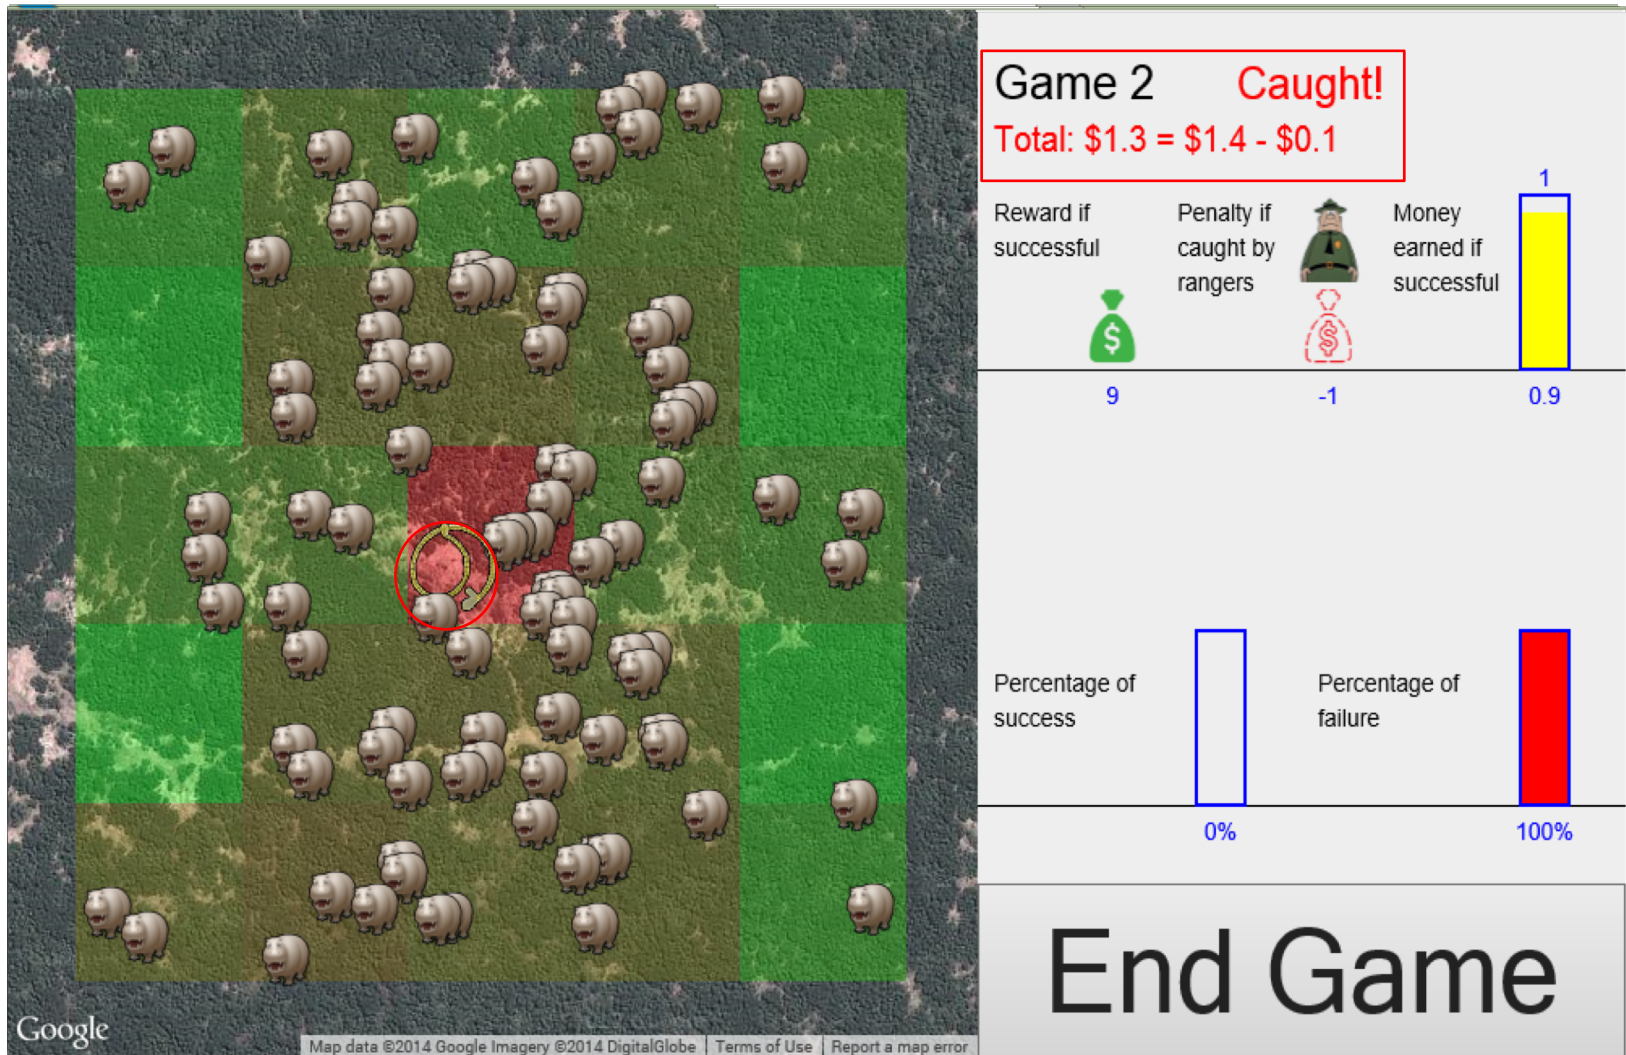
ONE STEP FURTHER

Subjective Utility Quantal Response(SUQR) [Nguyen 13]:


$$SEU^{adversary}(j) = w_1 \times \text{Capture Prob} + w_2 \times \text{Reward} + w_3 \times \text{Penalty}$$

$$\begin{array}{l} \text{Adversary's} \\ \text{probability of} \\ \text{choosing target } j \end{array} = \frac{e^{SEU^{adversary}(x,j)}}{\sum_{j'=1}^M e^{SEU^{adversary}(x,j')}}$$


BOUNDED RATIONALITY OF POACHERS



REPEATED GAMES [KAR, 2016]



Repeated games on AMT:
35 weeks, 40 human subjects
10,000 emails!



WHY DOES GAME THEORY PERFORM BETTER HERE?

Weaknesses of Previous Methods



Human Schedulers

Predictable patterns, e.g., LAX, US Coast Guard

Scheduling efforts and cognitive burden



Simple random (e.g., dice roll):

Repeatedly fails in deployments, e.g., officers to sparsely crowded terminals

Trillions of patrolling strategies, selecting important ones?

Incorporating learned adversary models, planning?

BIASING RECOMMENDATIONS

- A lot of examples from Economical settings (not only...)
- There are several well-known (and well-studied) cognitive biases of human buyers (e.g., anchoring effect, Bandwagon Effect, etc.)
- Capitalizing on these human biases combined with economical search theory for selecting when and which information to disclose and what price to set

ONLINE SHOPPING

- Display prices sequentially, adding some delay between any two prices.
- Anchoring
- Ordering
- etc.

Booking.com



Hajaj et al. Enhancing comparison shopping agents through ordering and gradual information disclosure, JAAMAS 2017.

Sarne et al. Improving Comparison Shopping Agents' Competence through Selective Price Disclosure, Electronic Commerce Research and Applications 2015.



DATA-DRIVEN PARADIGM

"The best predictor of future behavior is ... past behavior"

RECALL THAT...

- Experimental studies of choice behavior document distinct, and sometimes contradictory, deviations from maximization.
- Could those deviations be predicted using machine learning?

SETTING

Supervised

- Labeled decision making settings.
 - In setting x , person y choose z
 - ...
- Construct a model $F(.)$
 - Such that F approximates the real decision making outcomes
 - Generalizes to new decision making settings


Unsupervised

- Unlabeled decision making settings:
 - In setting x , person y
 - ...
- Identifying underlying structure:
 - Clusters
 - Association rules
 - ...

Reinforcement Learning

BUT... IT'S (SUPPOSED TO BE) EASY!

- Deep networks should answer all our prayers!
- Deep learning has demonstrated the possibility of stunning predictive performance via learning features and digesting large amounts of data.
- Could we automatically search for decision-making models?




ALPHAGO

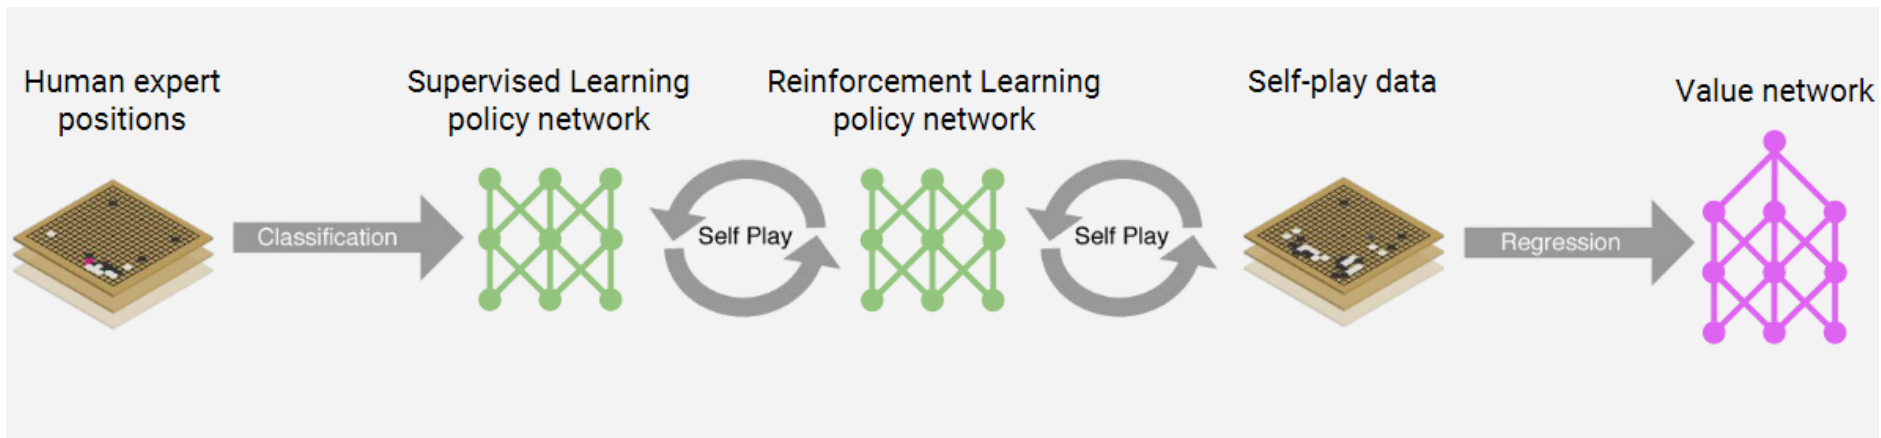
1997: IBM's Deep Blue beats
chess champion Garry Kasparov



2016: Google AlphaGo beats
Go Champion Lee Se-dol

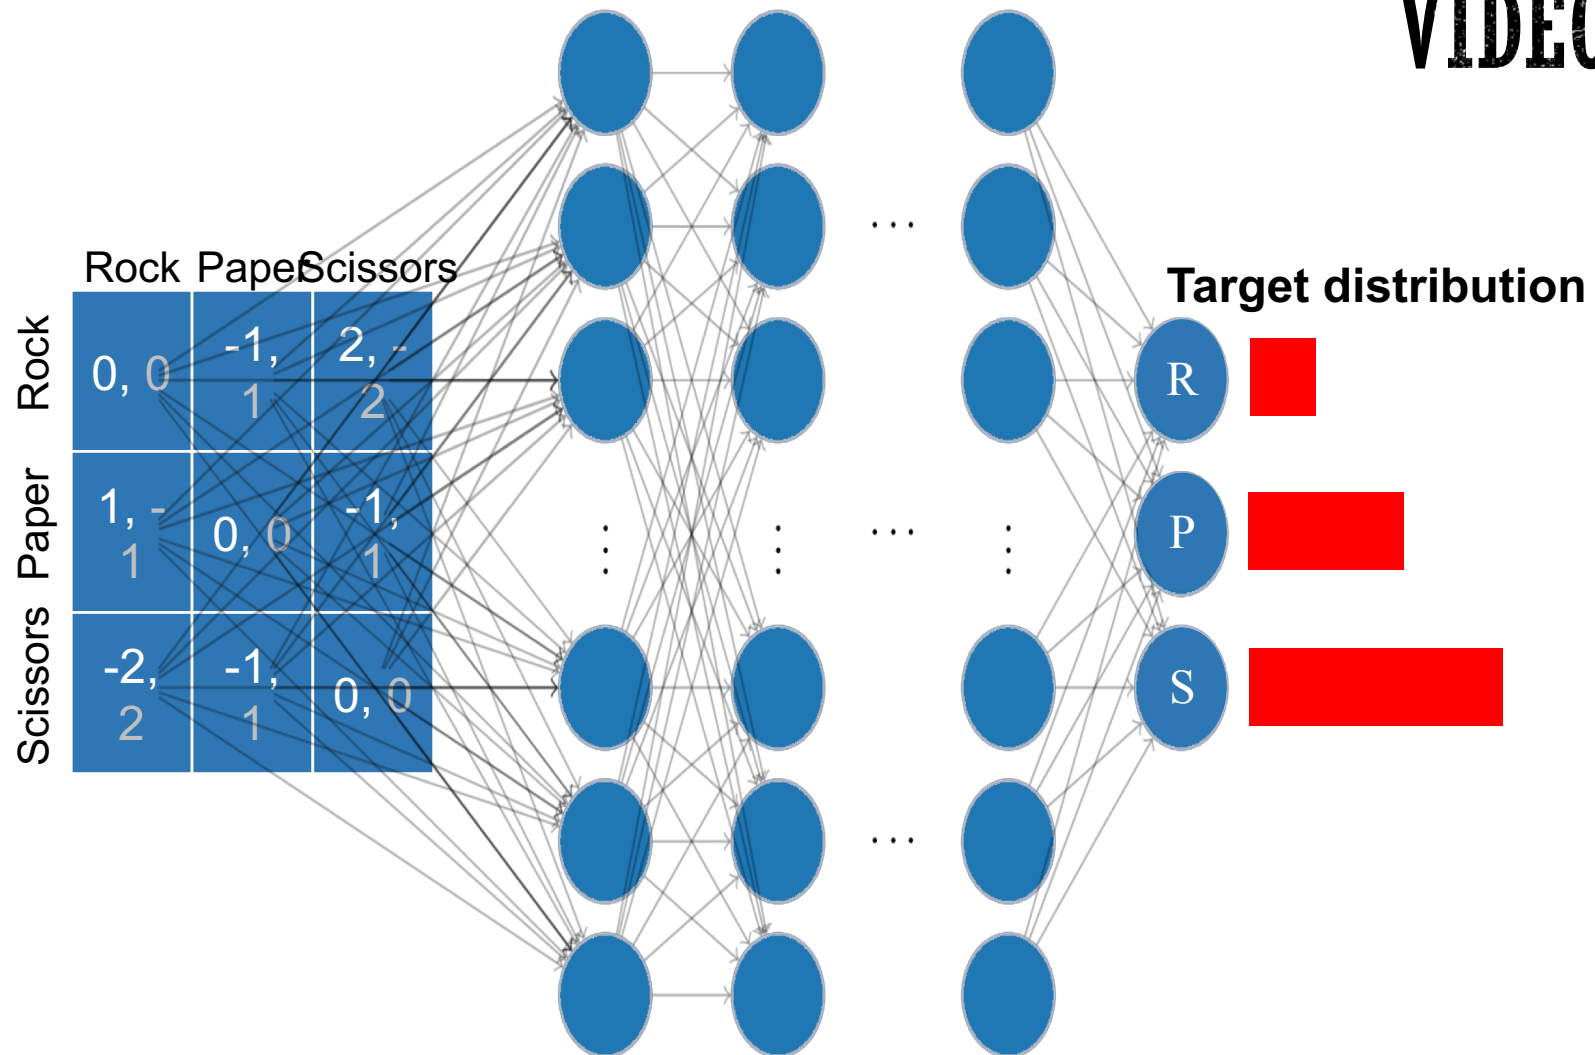


Silver et al. (Google Deepmind), Mastering the game of Go with deep neural networks and tree search 2016



Silver et al. (Google Deepmind), Mastering the game of Go with deep neural networks and tree search 2016

VIDEO!



**However,
all that glitters is not gold. At least not yet.**

PROB. 1 - DEEP LEARNING IS A DARK ART

- Not a “one-size-fits-all” solution
- Not an “off-the-shelf” solution
- A LOT of tricks.

...and then you
pour some data
inside...



PROB. 2 - WHERE IS BIG DATA COMING FROM?

■ Observational data:

- posts to social media sites
- digital pictures and videos
- GPS trails
- Transaction records
- cell phones
- traffic.
- Etc.



- **HARD TO OBTAIN FOR MANY DECISION SETTINGS!**

BACK IN 2015.... TIME TO GRADUATE?



My Past Advisor



I have 6 conference papers
and 2 journal papers.

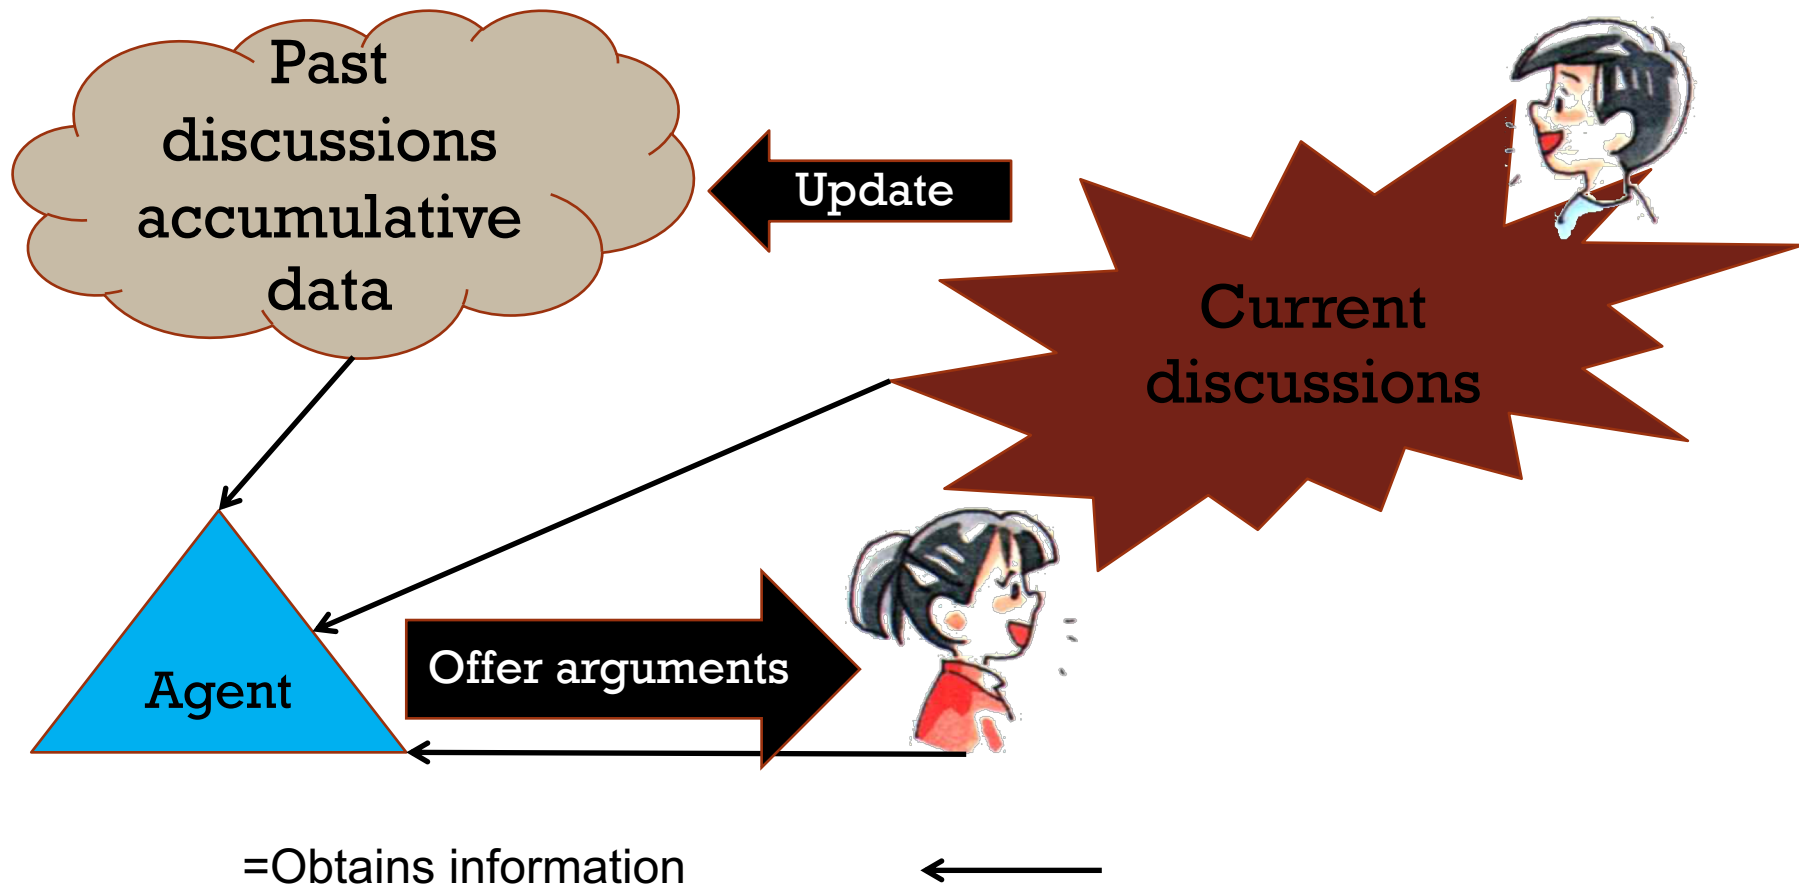
They are great papers, but you
can do more.

I guess you are right...



Me

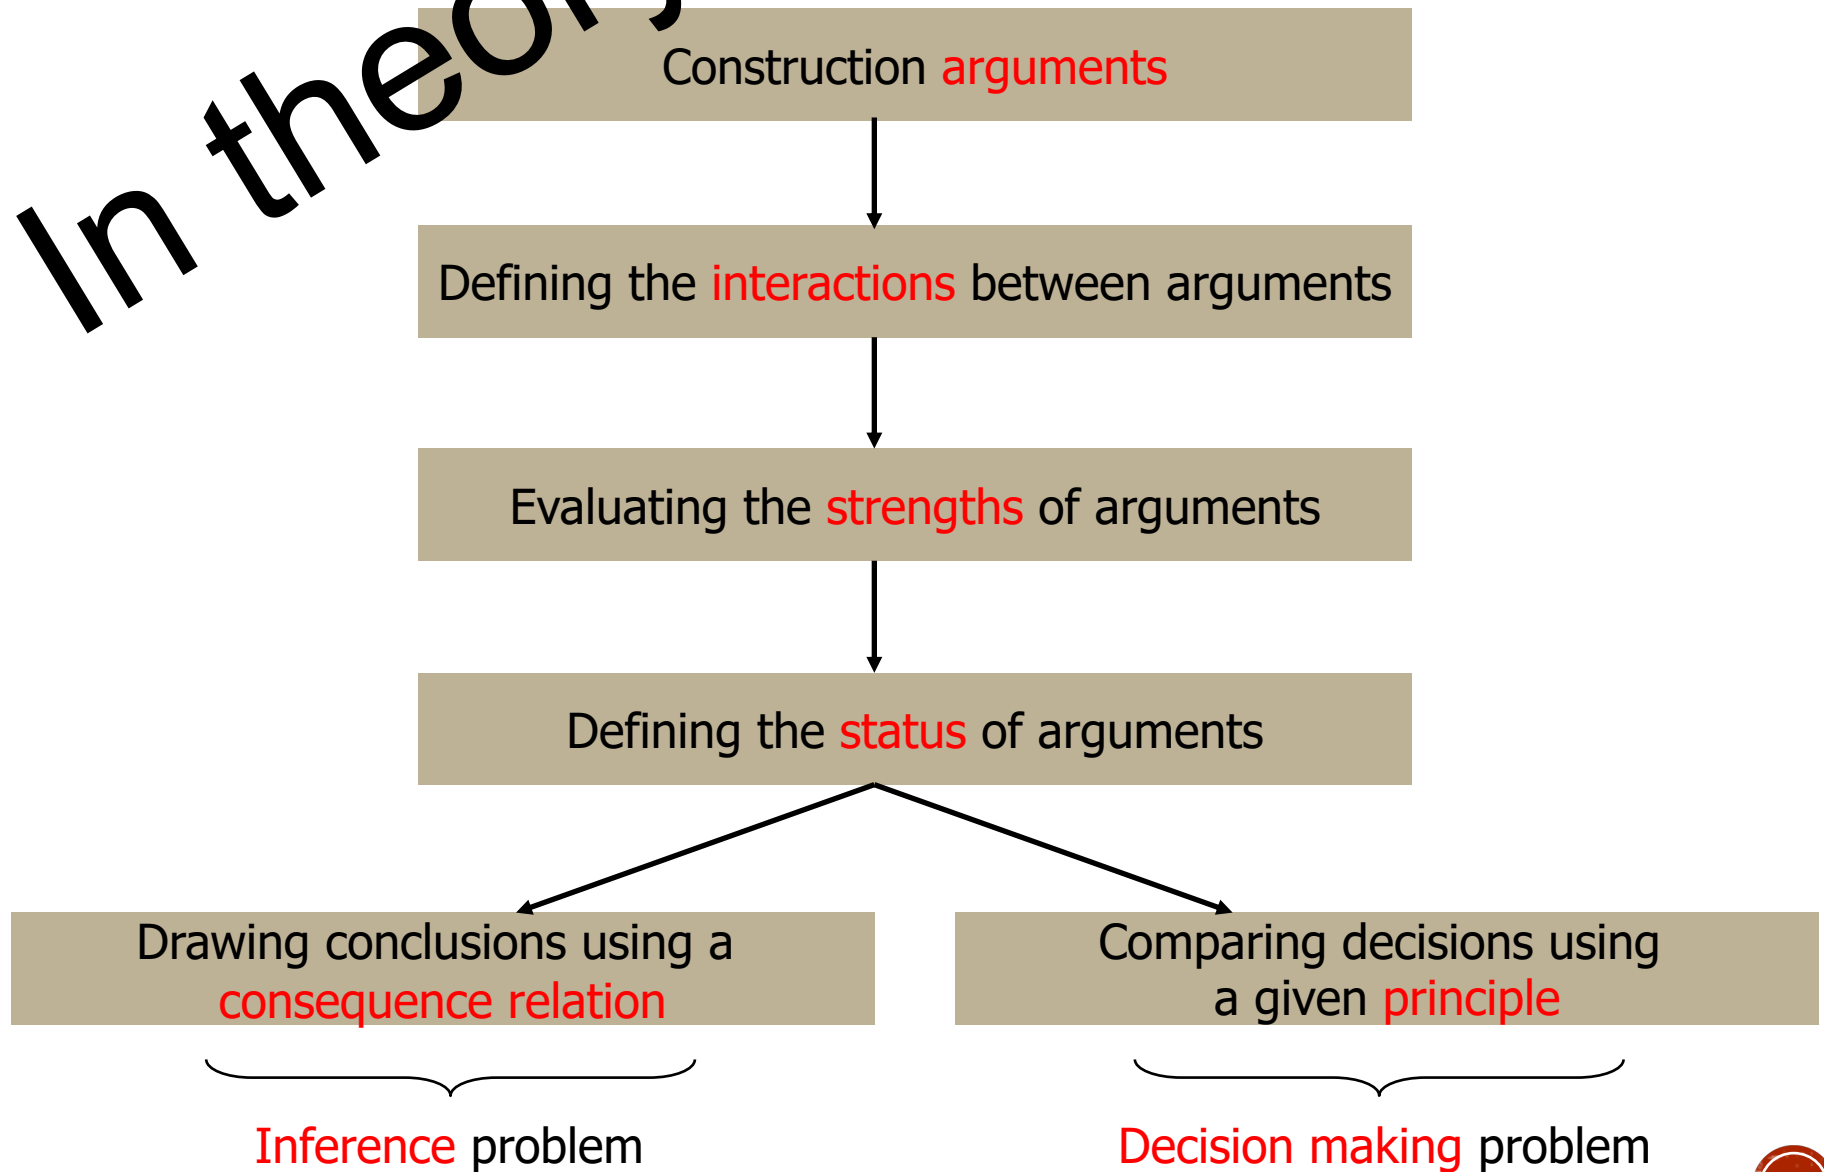
Agent Supports Discussions



PREDICTING HUMAN ARGUMENTATIVE BEHAVIOR


- Three (Major) options:
 - Argumentation theory.
 - Heuristics.
 - Machine Learning.

In theory...




ARGUMENTATION THEORY?


- Extensions?
- Validity values?
- Justification value?
- People do not reason logically.
- There is temporal nature of argumentation which is not captured.




Dung




Wyner



Reed



Cayrol



Parsons



Giacomin


ARGUMENTATION THEORY?

DATA COLLECTION OF 6 FICTIONAL CASES

- 64 participants from Amazon Turk;
 - Age average: 38.5
 - 21 females; 17 males
 - 3 with Phd



ARGUMENTATION THEORY?



Arvapally et al, 2012


ARGUMENTATION THEORY? TRANSCRIPTION OF REAL DISCUSSIONS

- Penn TreeBank Project (1995) conversation database:
 - CAPITAL PUNISHMENT (33)
 - TRIAL BY JURY (31)



Less than 30% were justified.


Predicting Arguments



Given 5 choices of a subject, calculate the average of each feature, and predict the 6th one.

CULTURE BASED?

- 78 Computer Science students.
 - CS-77% > AMT 72%
- Exactly the same features as AMT.
- Can learn from one and predict to the other.




ALTRUISM PREDICTION

1,067 Mechanical Turk workers X 20 decisions = 21,340 decision total

	Selfish option	Cooperative Option
Participant Payoff	x	x-c
Recipient Payoff	y	y+c*f

Ziv Epstein, Alex Peysakhovich and Dave Rand. **The Good, the Bad, and the Unflinchingly Selfish:** Cooperative Decision-Making can be Predicted with high Accuracy when using only Three Behavioral Types, 2016

MODEL BUILDING



- Capturing individual cooperation history leads to much higher predictive power
- Model using 3 cooperative types yields approximately same predictive power as fully heterogeneous model.

WHAT ARE THE THREE GROUPS?



most cooperative
they cooperated in 6
or more of their 15
training decisions



**intermediately
cooperative**
they cooperated at in
at least 1 to 6 of their
15 training
observations



least cooperative
they cooperated in
none of their 15
training observations

Attempted to predict cooperative type from these demographics.

- Maximum AUC achieved was 0.54. Not very good!
- Suggests cooperativeness is orthogonal to other demographics and is a natural kind (i.e. the cooperative phenotype)


HYBRID APPROACH

- Incorporating psychological models and data science in service of predicting human behavior.
- Aspiration Level
- Anchoring Bias
- Availability Heuristic
- Etc.

Noti, Gali, et al. "Behavior-Based Machine-Learning: A Hybrid Approach for Predicting Human Decision Making." *arXiv preprint arXiv:1611.10228* (2016).

Rosenfeld, Avi, et al. "Combining psychological models with machine learning to better predict people's decisions." *Synthese* 189.1 (2012): 81-93.

COMBINING BEHVAIORAL SCIENCE IN ML




- Prediction: Find $f(\cdot)$ such that $f(x) \approx y$
 - Machine learning people are the experts on this!
- What is x ?
 - Behavioral scientists are many times the experts on this!

Example for trials 6-15:

Please choose 'A' or 'B':

A:
3 with certainty

B:
4 with probability 0.8
0 with probability 0.2



You selected 'A' and your payoff is 3

Had you selected 'B', your payoff would have been 4


Erev et al. From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience, Psychological Review, 2017

WHAT IS X?

- “Objective” features
 - 11 parameters Defining the choice problem
- “Naïve features”
 - Difference in EVs, Difference in SDs...
- “Psychological features”:
 - Pessimism:
$$\text{diffMins} = \text{Min}_B - \text{Min}_A$$

= minimal outcome of B – minimal outcome of A
 - Minimization of regret:
$$pB_better = P[F_B^{-1}(x) > F_A^{-1}(x)] - P[F_A^{-1}(x) > F_B^{-1}(x)]$$

= *P[B providing better outcome than A] – P[A providing better outcome than B]*






FROM PREDICTION TO RECOMMENDATION

“A prediction model is only as good as it's agent's performance.”

PREDICTION MODELS IN PRACTICE

**“All models are wrong.
Some are useful.”
-George Box**



How “useful” are these models? How would you even measure that?

PREDICTION MODELS IN PRACTICE


“A prediction model is only as good as
it’s agent’s performance.”
- Ariel Rosenfeld



WHAT MAKES A GOOD AGENT?


- Usually,

Maximize Your
Expected Utility




ARGUMENTATION RECOMMENDATION POLICIES

- How can we assist a human while arguing?
 - No pro-active approach.
- Should we suggest the predicted arguments?
- How to maintain a good hit-rate while offering novel arguments?




AGENTS



- PRD: Prediction [17 chats]
- REL: Relevant [17 chats]
- WRL: Weakly related [17 chats]
- PRH: Prediction (2) + Relevant (1) [17 chats]
- TRY: Theory [17 chats]
- REP: PRH with repeated arguments [17 chats]
- TLA: TL agent [17 chats]
- RND: Random [17 chats]

NORMALIZED ACCEPTANCE RATE



USER SATISFACTION


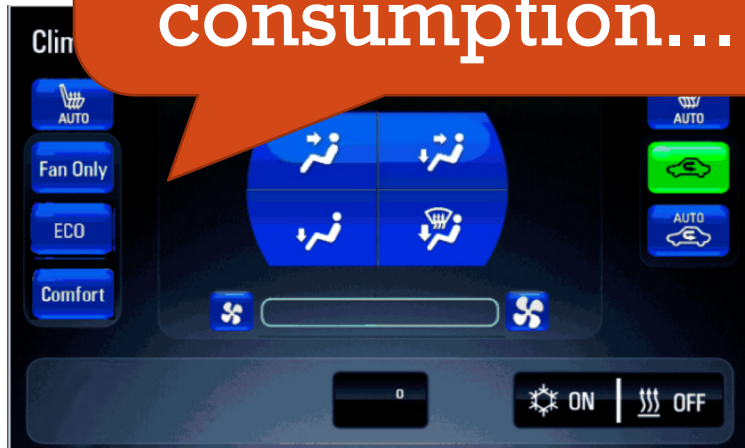
Agent	Neutral	Positive	V. Positive
Prediction	5	10	2
Relevant	5	12	0
Weakly related	16	1	0
Prediction + Relevant	0	12	5
Theory	13	3	1
Random	14	3	0
Repeated	12	5	0
TLA	13	3	1

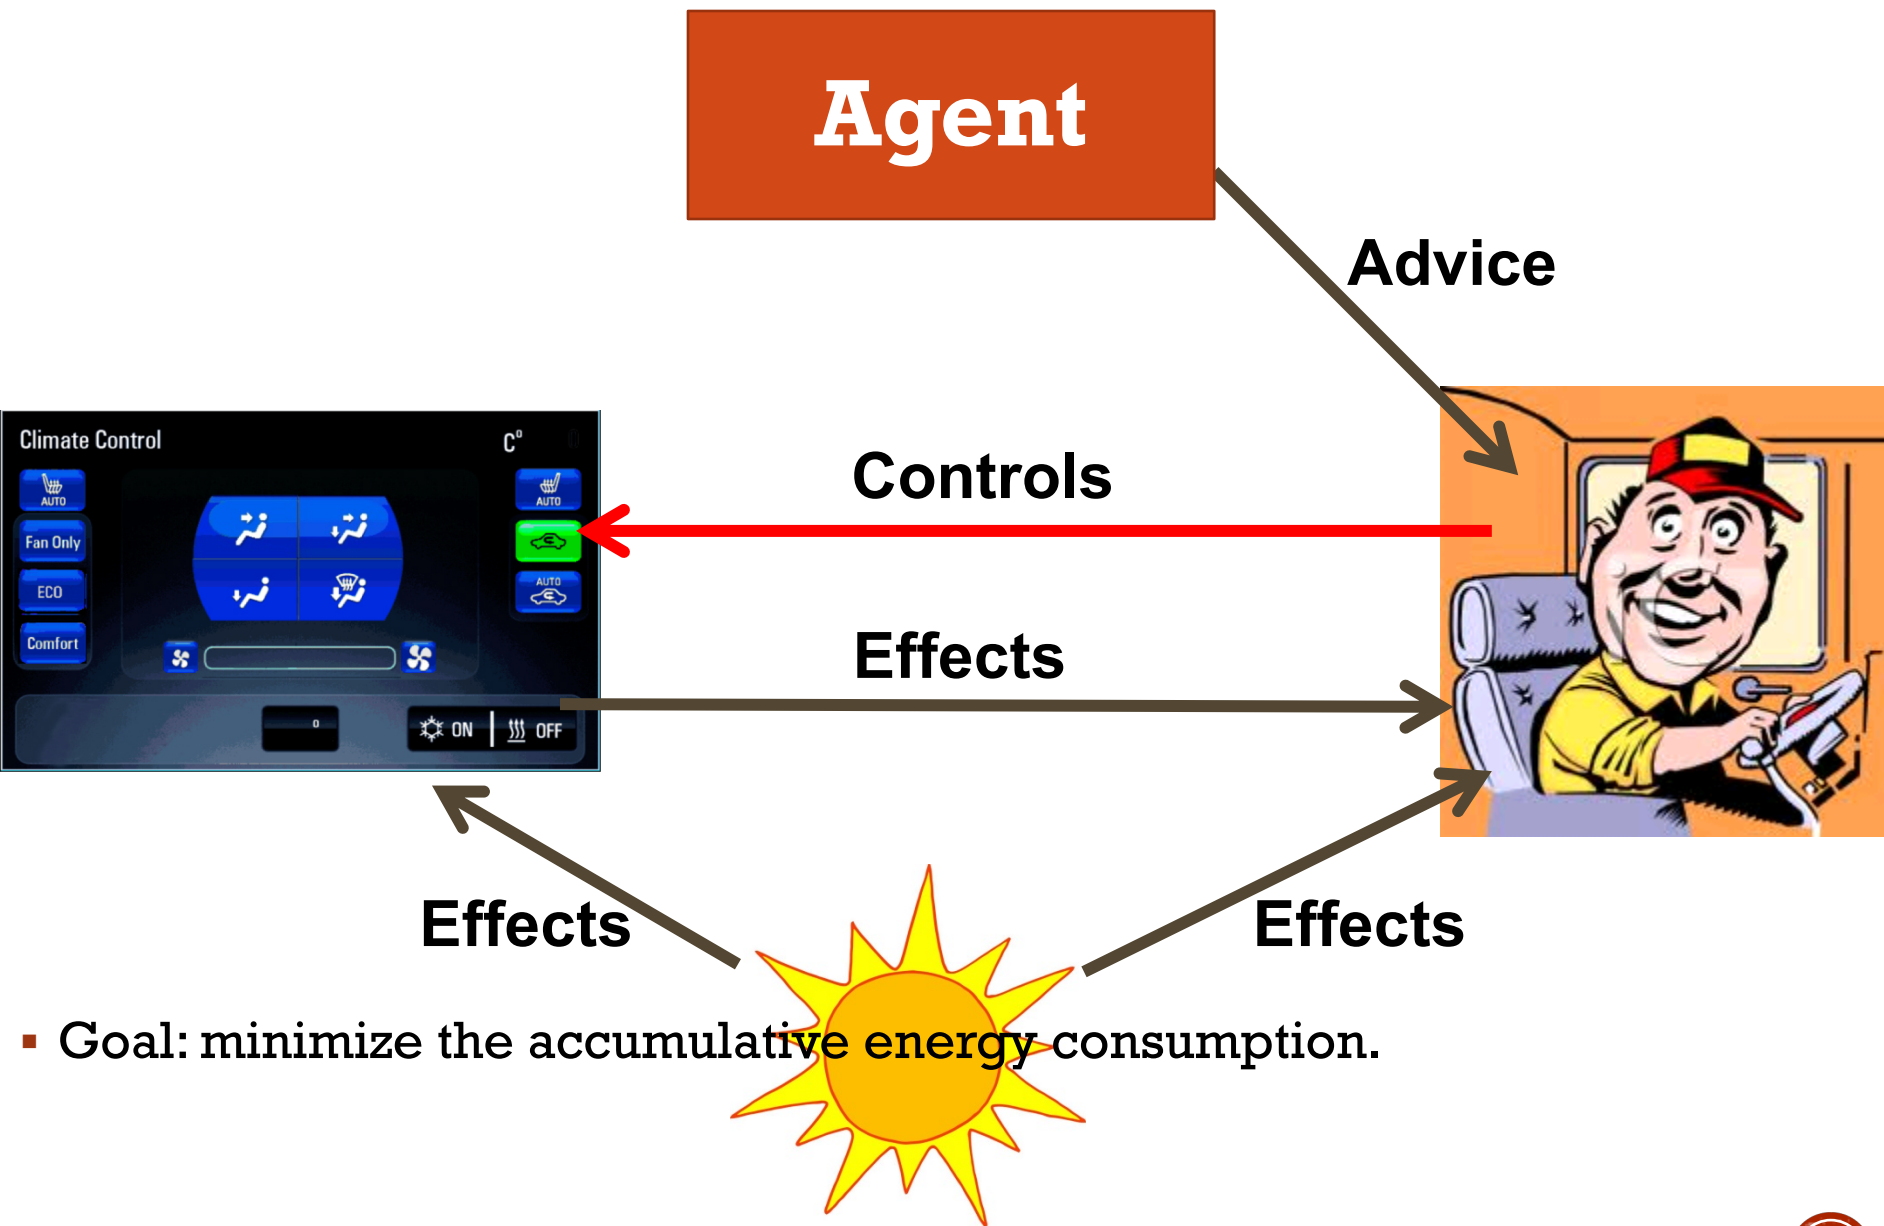
PARTIALLY CONFLICTING INTERESTS

- Driver's and system's goals are **partially** conflicting.

Let's minimize
energy
consumption...


I'm Hot!





EVALUATION

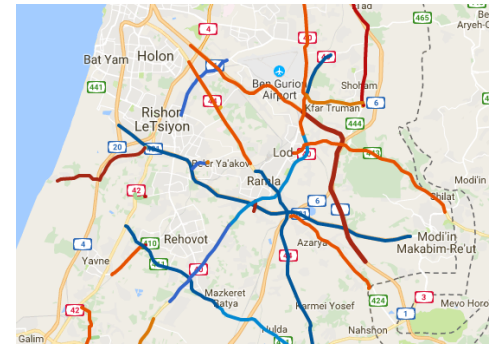
- 45 drivers - 15 per condition, 3 rounds.



The lower the better.

MITIGATING TRAFFIC ACCIDENTS IN ISRAEL

- Predicting locations and times of higher\lower risk.
- Drivers **react** to police presence in both time and space – marginal effects
- Prediction calls for new optimizing allocation though master-slave optimization.



Red = higher risk, Blue = lower risk



Rosenfeld, Maksimov and Kraus, When Security Games Hit Traffic: Optimal Traffic Enforcement under One Sided Uncertainty ,IJCAI 2017

Rosenfeld, Maksimov and Kraus, Optimal Crusier-Drone Traffic Enforcement ,IJCAI 2018

RECIPROCAL RECOMMENDER SYSTEMS

A Reciprocal Recommender System (RRS) recommends people to people.

Potential applications:


- Job recruiting.
- Online-dating.



Klenierman, Rosenfeld, Ricci and Kraus. Optimally Balancing Receiver and Recommended Users' Importance in Reciprocal Recommender Systems, RecSys 2018



THE RECOMMENDATION PROCESS




RRS should increase *successful interactions* by considering the interest of both sides.

How should the system balance the interests of both sides?





PREVIOUS RRS RECOMMENDATION METHODS

- **Interaction-based Collaborative Filtering:** preferences elicited from the users' interactions.
- **All previous methods:** constant and equal importance to both sides.



USER'S VARIANCE IN INTERACTIONS

Users vary in selectivity and popularity.




RECIPROCAL WEIGHTED SCORE (RWS)

Our method balances two scores:

- CF : The service user's interest.
Estimated by interaction-based collaborative filtering.
- PR : The likelihood for positive reply.
Estimated by an Adaptive Boosting prediction model.

$$RWS_{x,y} = \alpha_x (CF_{x,y}) + (1 - \alpha_x) PR_{y,x}$$




WEIGHT OPTIMIZATION

$$RWS_{x,y} = \alpha_x (CF_{x,y}) + (1 - \alpha_x) PR_{y,x}$$

Optimization Problem:

Observed from x 's interaction history.

Given all of user's viewed by x , find optimal weight which will rank x 's successful interactions highest.



$$\begin{aligned} & \underset{\alpha_x}{\text{minimize}} && \sum_{v \in V_x} \mathbb{1}_{v \in \text{SuccInter}_x} \text{Rank}_v (RWS_{x,*}(\alpha_x)) \\ & \text{subject to} && 0 \leq \alpha_x \leq 1 \end{aligned}$$

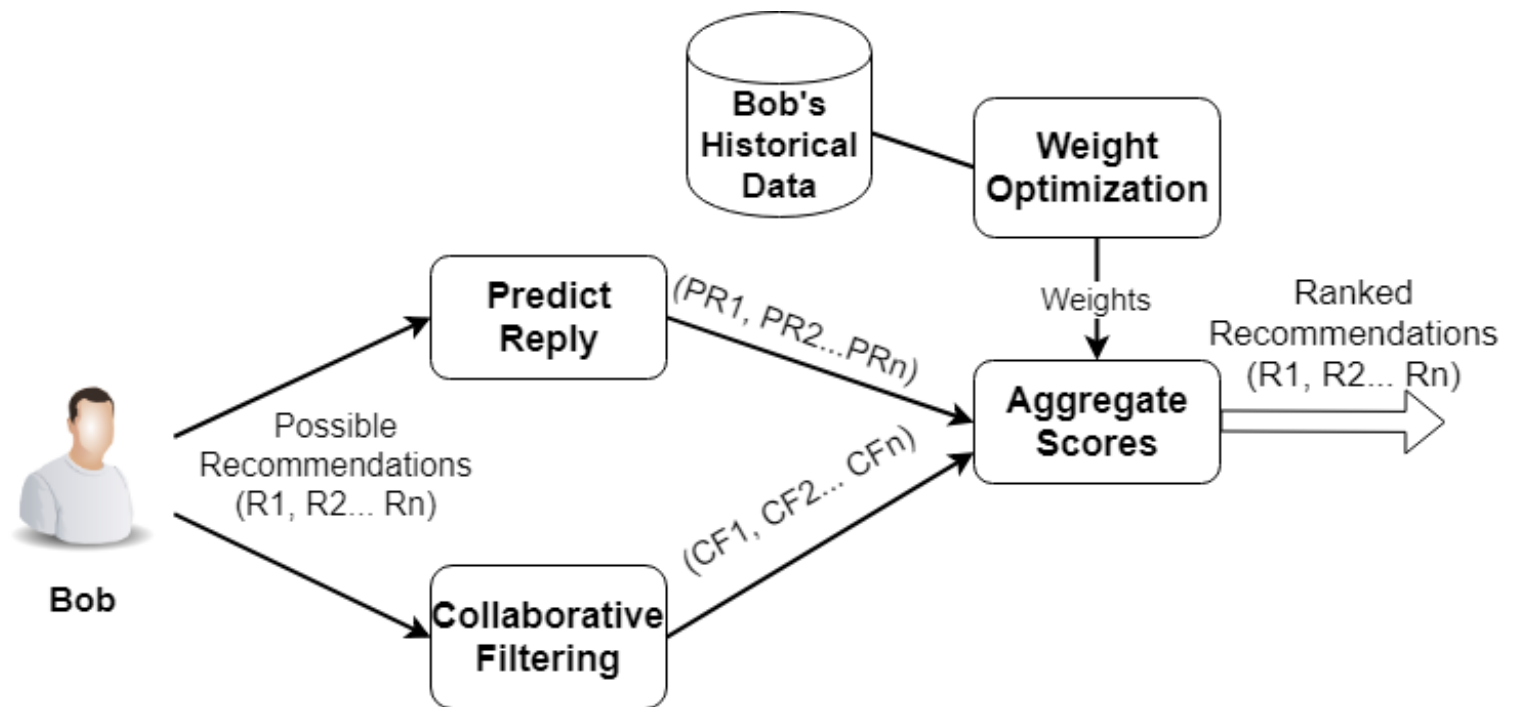


REPLY PREDICTION MODEL

- 35,000 messages classified to either:
 - Positive reply.
 - Negative reply or no-reply at all.
- 54 features:
 - Sender's and Receiver's features: public profile, activity and popularity.
- Adaptive Boosting Classifier.
 - AUC: 0.833

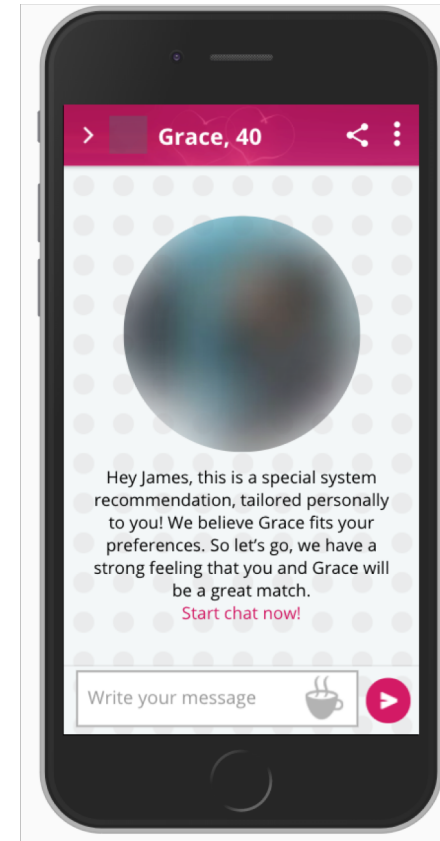
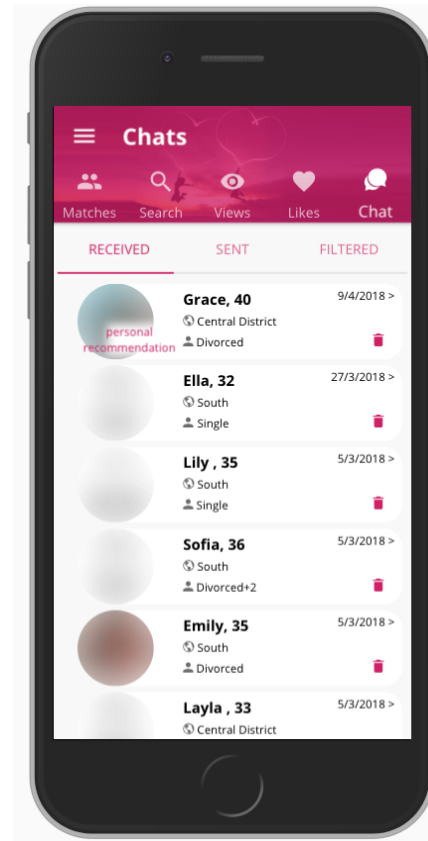


RECOMMENDATION SYSTEM DIAGRAM




EXPERIMENTAL SETUP

- Evaluation in Doovdevan, an operational online dating mobile-app.
- Two conditions:
 - RWS.
 - Baseline.
- Each participant received three recommendations.





RESULTS

Successful Interactions



OPTIMIZED WEIGHT GENDER DIFFERENCES



SUMMARY

- In RRSs, the interests of both sides of the recommendation should be considered.
- Our method (RWS) finds an optimal balance of both side's interests, tailored individually for each user by his history.
- We evaluated RWS in an online dating application and found it is effective in increasing successful interactions.





CONCLUSIONS

What should I do?

STEP 1

- “Take a social science crash course”
- Form a firm grasp on what makes YOUR people tick, it may not always be what you’d expect.
- The Important Questions:
 - Which assumptions were made\validated by pervious works in the realm?
 - Did you check behavioral sciences for answers?

STEP 2

- “In DATA we trust”.
- Is there available data?
- Collect contextual data (it might take more effort than you’d think!)
- The Important Questions:
 - Which is the appropriate data collection method?
 - Is the data reliable and extensive?
 - Did you examine the ecological validity of YOUR assumptions?

STEP 3

- “Put on your lab coat”.
- Carefully choose a prediction method.
- The Important Questions:
 - Do we have enough quality data? Should we collect more?
 - Which machine learning procedures are appropriate for the data?
 - Did you account for statistical factors such as minority cases, time series effects, outliers, etc.?
 - Did you make sure you avoid overfitting?

STEP 4


- “Don’t stop ‘till you get enough”
- Enhance your prediction model
 - Consider Hybrid Approach
 - Consider collecting more data
 - Adding heuristics
 - Etc.

STEP 5

- “Put the agent where the prediction is”
- Test your agent. Mediocre prediction model may suffice.
 - Optimization.
 - Reasoning

TAKE HOME MESSAGE

- Predicting Human Decision Making is **Hard**.
 - Don't expect to reach 90% acc. (even with cheating)
 - Normative Models can help.
 - Machine Learning is very useful
- Human Prediction is important for turning recommendations into impact.



THANKS

www.arielrosenfeld.com

<http://tinyurl.com/predicting-human-DM>

- Ariel Rosenfeld:
arielros1@gmail.com
- Looking for collaborations
 - Human-Agent Interaction
 - Explainable AI
 - AI for social good
 - Etc...
- Email me!





That's all Folks!