FROM MACHINE LEARNING TO OPTIMIZATION IN SPORTS ANALYTICS

Oliver Schulte

School of Computing Science

Simon Fraser University

Vancouver, Canada

SPORTS ANALYTICS

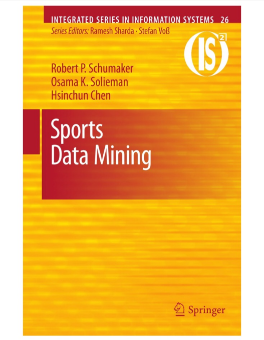

Sport

- Tracking
- Evaluation
- Prediction

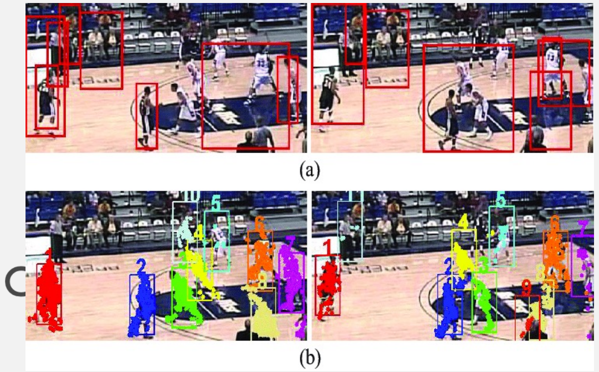

OVERVIEW

Success Probabilities		Action Values	Player Ranking	Learning Success Probabilities	Optimization
Defining Success Probabilities	A fundamental problem in sports analytics	Success Probabilities ↓ Action Values	Action Values ↓ Player Ranking	Classifier, Reinforcement Learning	Finding optimal players, tactics, wagers

SUCCESS PROBABILITIES

WHAT IS SUCCESS?

- An outcome (binary event) that a team wants to bring about.
- Can be defined according to the interest of the analyst/coach/player.

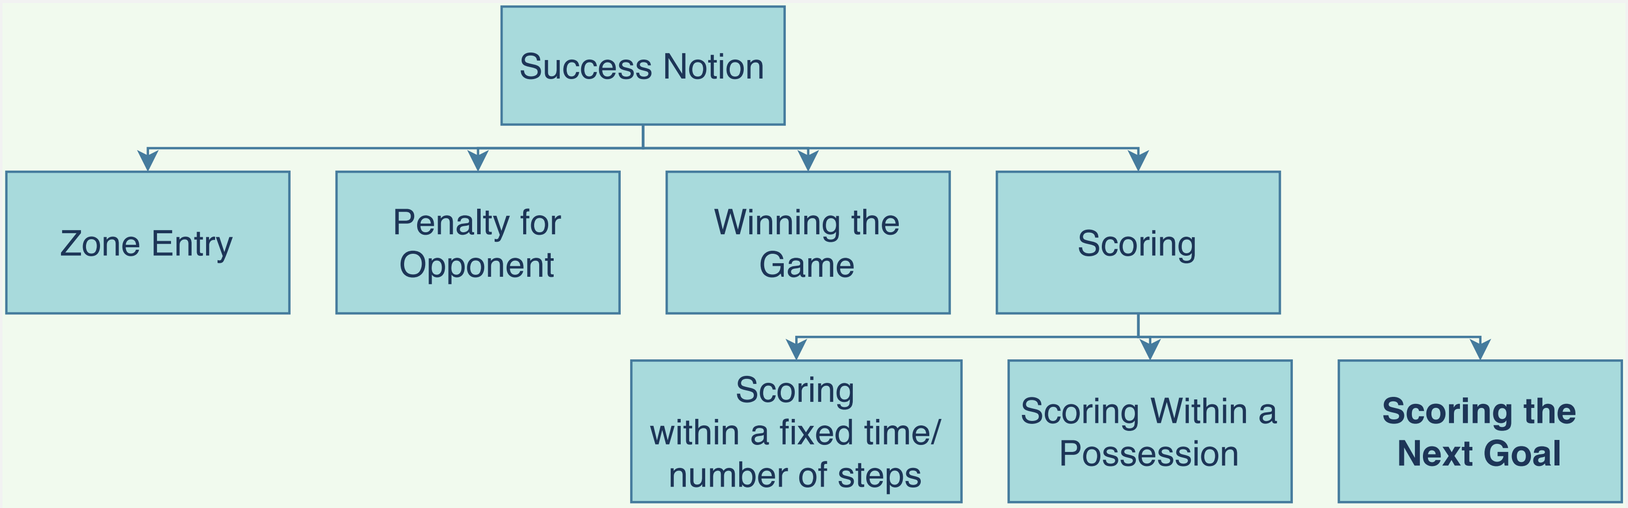

SUCCESS PROBABILITY TICKER

- Assigns to each time *t* in the match an estimated probability of *future* success
- Example: Win probability in NHL (Pettigrew 2015)

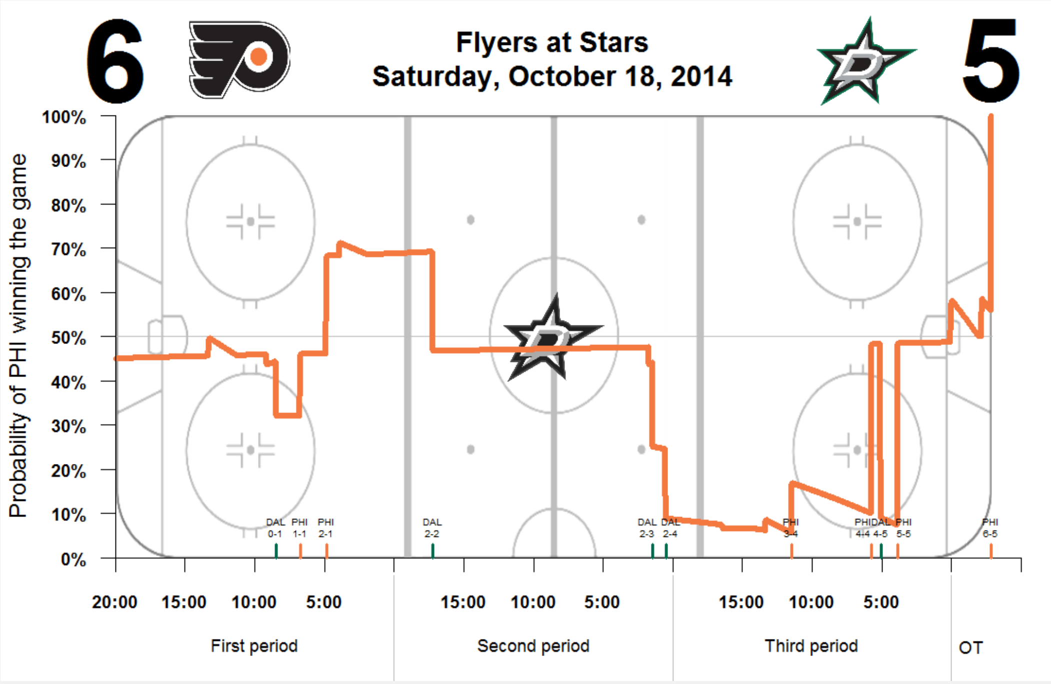

S. Pettigrew, "Assessing the offensive productivity of NHL players using in-game win probabilities," 2015. 9th Annual MIT Sloan Sports Analytics Conference

SCORE IN POSSESSION

http://www.lukebornn.com/sloan_epv_curve.mp4

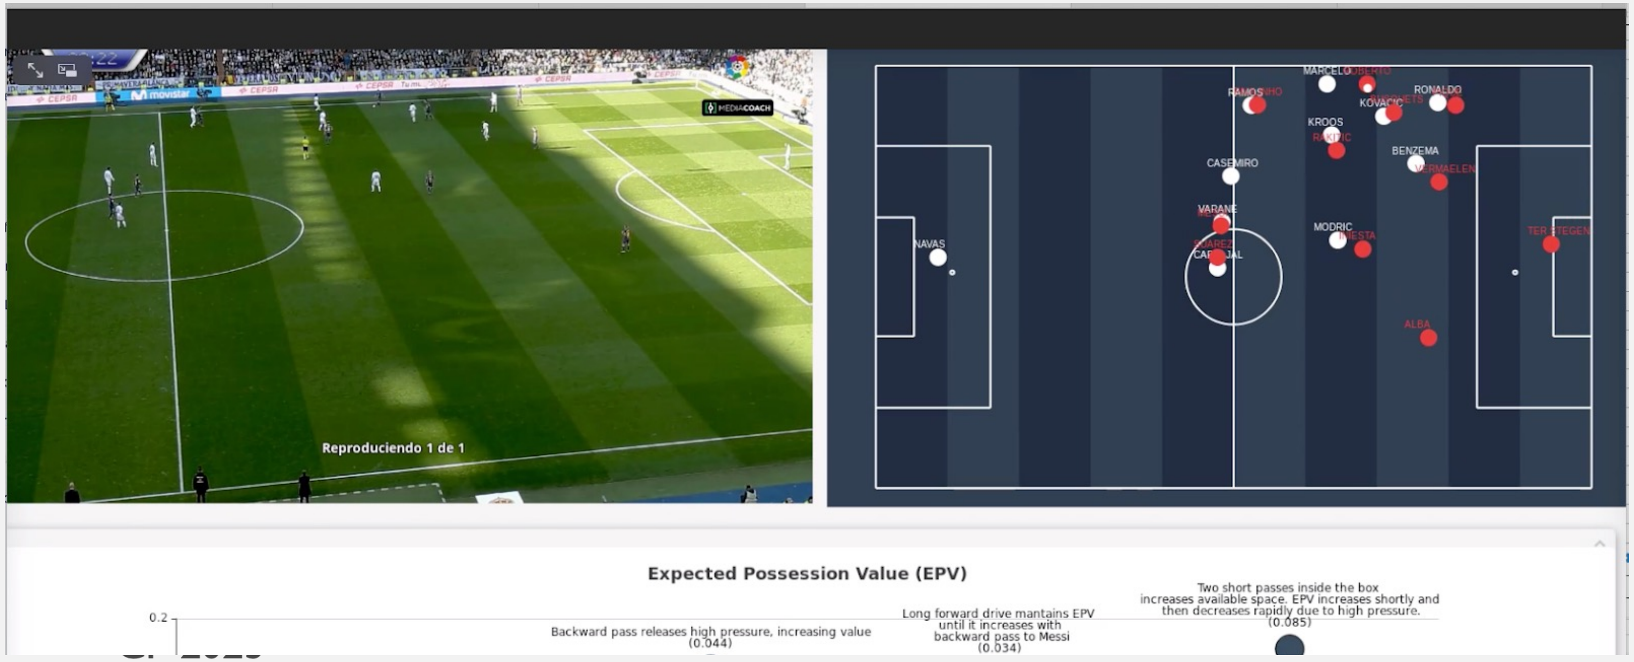

NEXT GOAL PROBABILITIES

- Y-axis: the chance of scoring the next goal

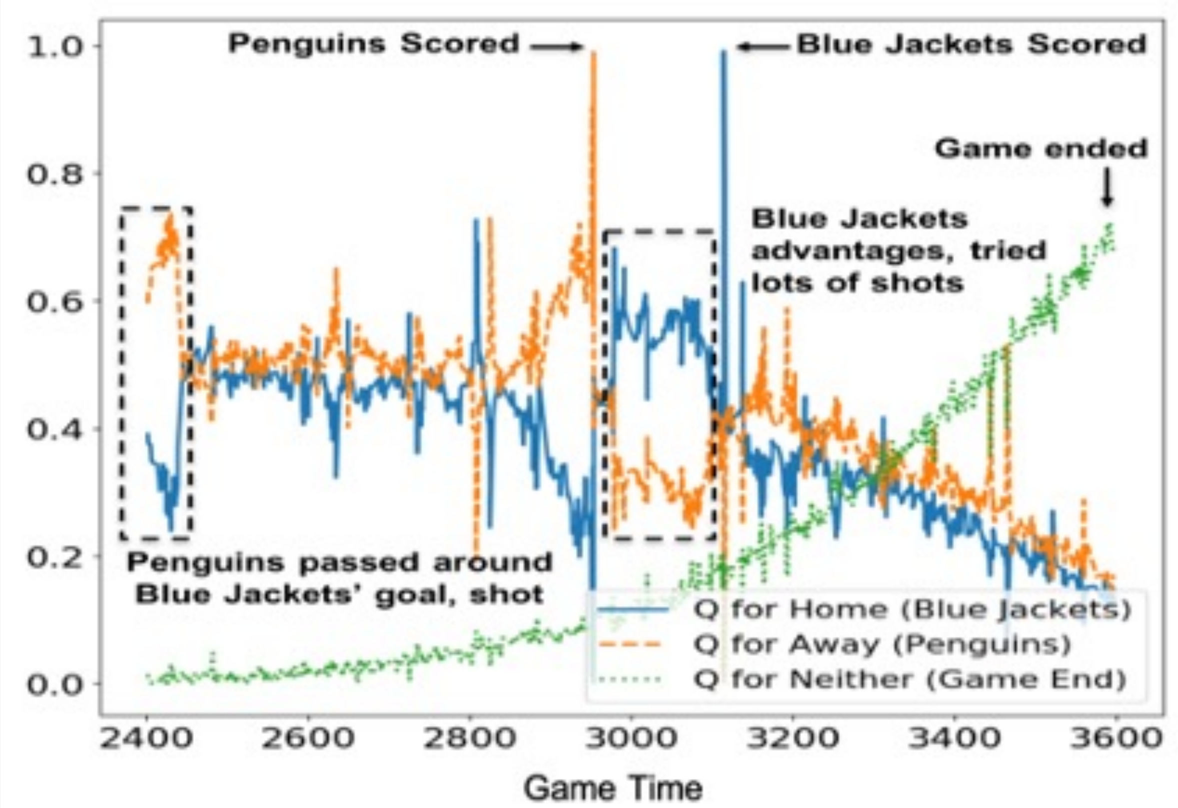

G. Liu and O. Schulte, "Deep Reinforcement Learning in Ice Hockey for Context-Aware Player Evaluation," in *IJCAI-18*, 2018-07, pp. 3442–3448.

FROM SUCCESS PROBABILITIES TO ACTION VALUES

ACTION VALUES

- Success probabilities can be used to **evaluate players and actions**
- impact(action) =
  [success probability after action – success probability before action]
- "We assert that most questions that coaches, players, and fans have about basketball, particularly those that involve the offense, can be phrased and answered in terms of EPV [i.e. expected future success]." Cervone, Bornn et al. 2014.

D. Cervone, A. D'Amour, L. Bornn, and K. Goldsberry, "POINTWISE: Predicting points and valuing decisions in real time with NBA optical tracking data," in *8th Annual MIT Sloan Sports Analytics Conference, February*, 2014, vol. 28.

THE IMPACT OF AN ACTION

- The impact of an action a_t performed at time $t>0$ is the difference in successive success probabilities:

  $\text{impact}(a_t) = p_t - p_{t-1}$

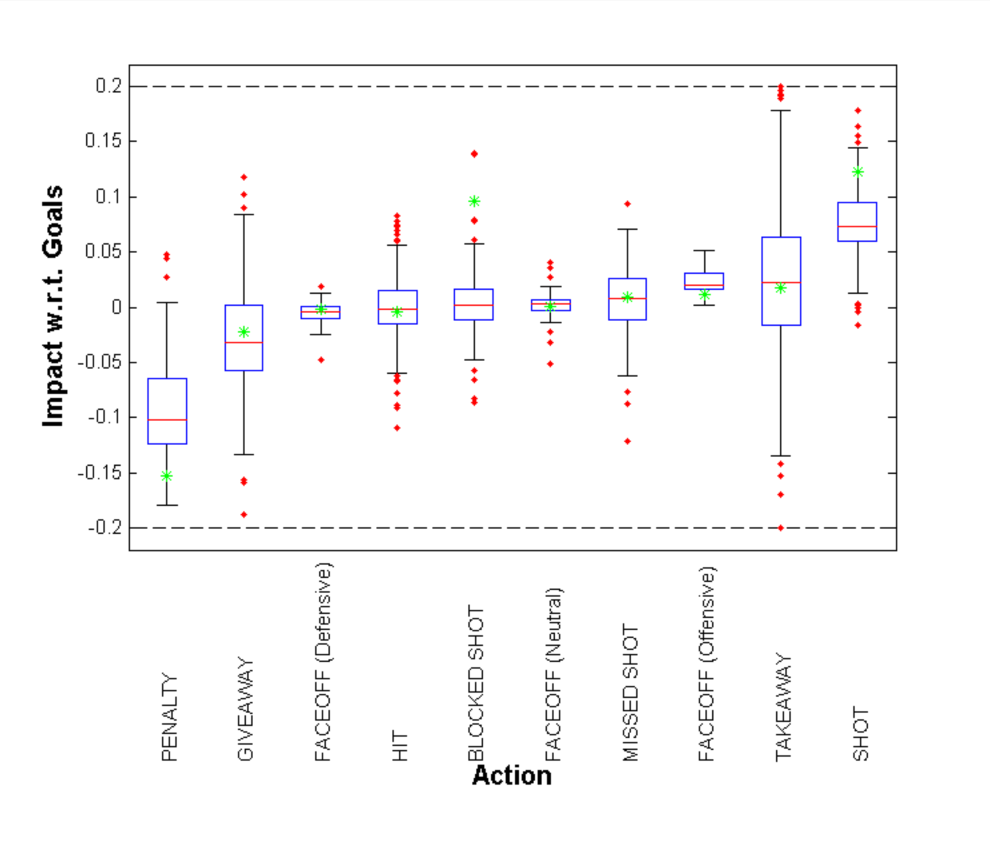

* = THOR baseline

FROM ACTION IMPACT TO PLAYER RANKINGS

RESULTS 2013-2014 SEASON NHL

Name	Goal Impact	Points	+/-	Salary
Jason Spezza	29.64	66	-26	$5,000,000
Jonathan Toews	28.75	67	25	$6,500,000
Joe Pavelski	27.20	79	23	$4,000,000
Marian Hossa	26.12	57	26	$7,900,000
Patrick Sharp	24.43	77	12	$6,500,000
Sidney Crosby	24.23	104	18	$12,000,000
Claude Giroux	23.89	86	7	$5,000,000
Tyler Seguin	23.89	84	16	$4,500,000

Player performance =
total impact of all actions performed
Jason Spezza: high goal impact, low +/-.

- plays very well on poor team (Ottawa Senators 2013).
- Requested transfer for 2014-2015 season.

PLAYER RANKING

- 2015-16 NHL season
- Johnny Gaudreau and Mark Scheifele drew salaries below what their GIM rank would suggest.
- Later they received a $5M+ contract for the 2016-17 season.

Name	GIM	Assists	Goals	Points	Team	Salary
Taylor Hall	96.40	39	26	65	EDM	$6,000,000
Joe Pavelski	94.56	40	38	78	SJS	$6,000,000
Johnny Gaudreau	94.51	48	30	78	CGY	$925,000
Anze Kopitar	94.10	49	25	74	LAK	$7,700,000
Erik Karlsson	92.41	66	16	82	OTT	$7,000,000
Patrice Bergeron	92.06	36	32	68	BOS	$8,750,000
Mark Scheifele	90.67	32	29	61	WPG	$832,500
Sidney Crosby	90.21	49	36	85	PIT	$12,000,000
Claude Giroux	89.64	45	22	67	PHI	$9,000,000
Dustin Byfuglien	89.46	34	19	53	WPG	$6,000,000
Jamie Benn	88.38	48	41	89	DAL	$5,750,000
Patrick Kane	87.81	60	46	106	CHI	$13,800,000
Mark Stone	86.42	38	23	61	OTT	$2,250,000
Blake Wheeler	85.83	52	26	78	WPG	$5,800,000
Tyler Toffoli	83.25	27	31	58	DAL	$2,600,000
Charlie Coyle	81.50	21	21	42	MIN	$1,900,000
Tyson Barrie	81.46	36	13	49	COL	$3,200,000
Jonathan Toews	80.92	30	28	58	CHI	$13,800,000
Sean Monahan	80.92	36	27	63	CGY	$925,000
Vladimir Tarasenko	80.68	34	40	74	STL	$8,000,000

CORRELATIONS WITH STANDARD STATS

- GIM: our ranking (goal impact metric)
- Takeaway: high correlation with standard stats
  - e.g. 0.93 with points

methods	Assist	Goal	GWG	OTG	SHG	PPG	S
+/-	0.236	0.204	0.217	0.16	0.095	0.099	0.118
GAR	0.527	0.633	0.552	0.324	0.191	0.583	0.549
WAR	0.516	0.652	0.551	0.332	0.192	0.564	0.532
EG	0.783	0.834	0.704	0.448	0.249	0.684	0.891
SI	0.869	0.745	0.631	0.411	0.27	0.591	0.898
GIM-T1	0.873	0.752	0.682	0.428	0.291	0.607	0.877
GIM	**0.875**	**0.878**	**0.751**	**0.465**	**0.345**	**0.71**	**0.912**

methods	Point	SHP	PPP	FOW	P/GP	TOI	PIM
+/-	0.237	0.159	0.089	-0.045	0.238	0.141	0.049
GAR	0.622	0.226	0.532	0.16	0.616	0.323	0.089
WAR	0.612	0.235	0.531	0.153	0.605	0.331	0.078
EG	0.854	0.287	0.729	0.28	0.702	0.722	0.354
SI	0.869	0.37	0.707	0.185	0.655	**0.955**	**0.492**
GIM-T1	0.902	0.384	0.736	0.288	0.738	0.777	0.347
GIM	**0.93**	**0.399**	**0.774**	**0.295**	**0.749**	0.835	0.405

LEARNING SUCCESS PROBABILITY MODELS

Try this at home

MACHINE LEARNING APPROACH

Deployment

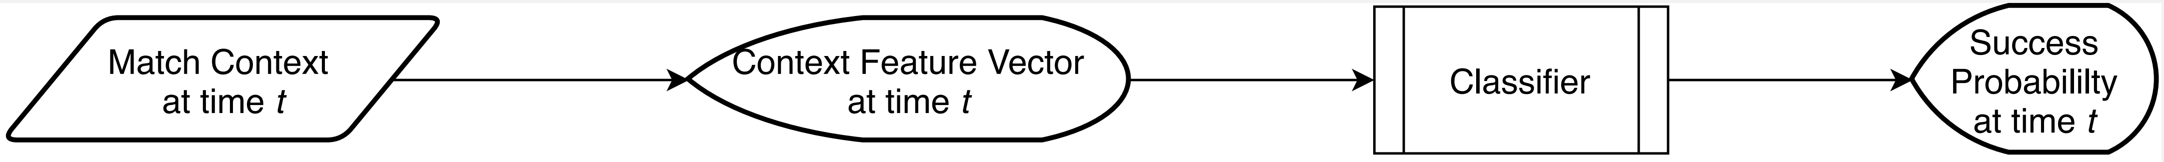

Training

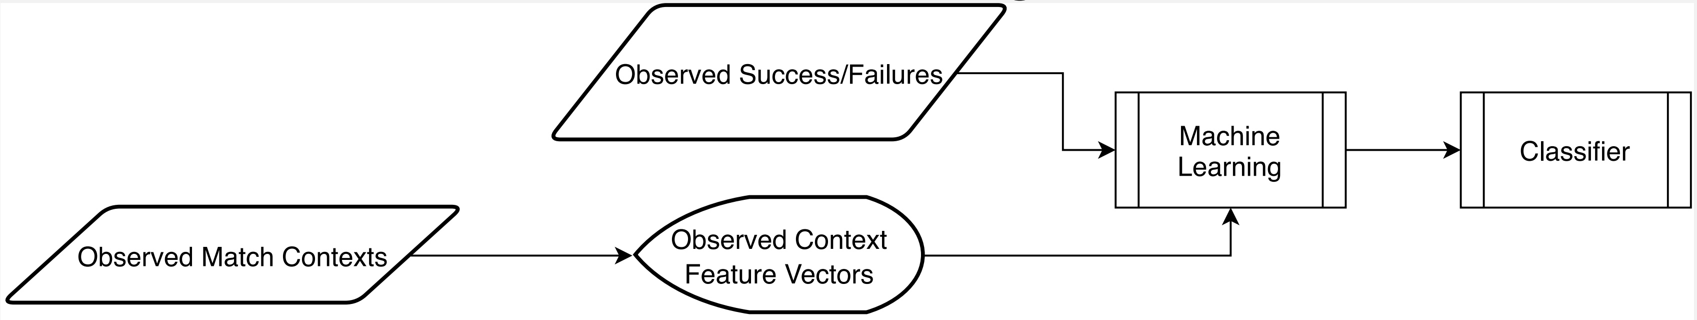

Oliver Schulte. (2022). Valuing Actions and Ranking Hockey Players with Machine Learning (Extended Abstract). Linköping Hockey Analytics Conference, Linköping, Sweden (2-9). https://ecp.ep.liu.se/index.php/linhac/issue/view/67

EVENT DATA

- Illustrate approaches with event data
  - [available from nhl.com]
  - also [pre-crawled]
- Less work on tracking data
- [Sportlogiq.com] is Montreal-based data provider

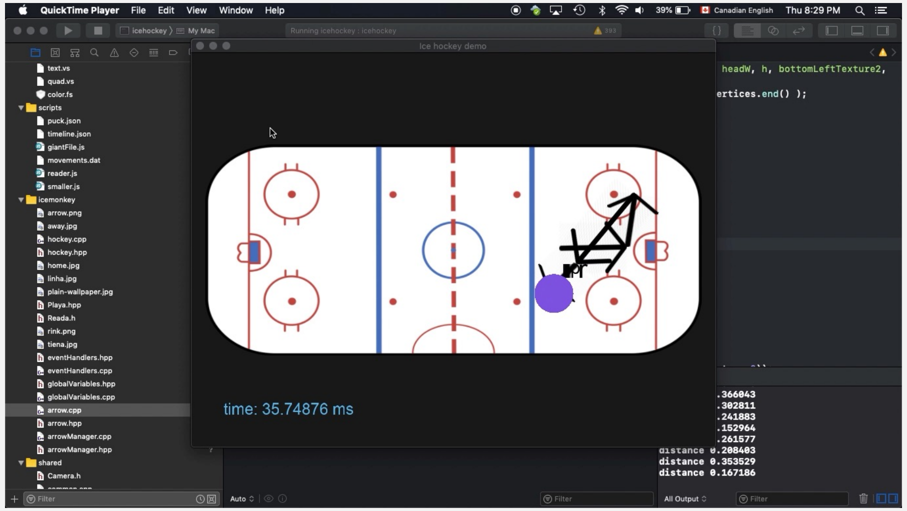

DATASET STATISTICS 2015-16

Number of Teams	**30**
Number of Players	2,233
Number of Games	1,140
Number of Events	3.3M

SPECIFYING LOOK-AHEAD

- Suppose we start with event data
- For each time *t*, add a success column Y_t depending on whether the team succeeded after time *t*.
  - E.g., did they score the next goal?
- Could also annotate whether they score the next goal in *k* steps (Descrooset al. 2019) or fixed time interval (Shuckers and Curro 2013 THoR)

T. Decroos, L. Bransen, J.V. Haaren, and J. Davis, "Actions Speak Louder than Goals: Valuing Player Actions in Soccer," *(KDD-19)*, 2019, pp. 1851–1861.
M. Schuckers and J. Curro, "Total Hockey Rating (THoR): A comprehensive statistical rating of National Hockey League forwards and defensemen based upon all on-ice events," 2013. 7th Annual MIT Sloan Sports Analytics Conference

EXAMPLE

game	player	Period	team	x	y	Manpower	Action	Y=next goal
849	402	1	15	-9.5	1.5	Even	Recovery	0
849	402	1	15	-24.5	-17	Even	Carry	0
849	417	1	16	-75.5	-21.5	Even	Check	1
849	402	1	15	-79	-19.5	Even	Pass	0
849	413	1	16	-92	-32.5	Even	Turnover	1
849	413	1	16	-92	-32.5	Even	Pass	1
849	389	1	16	-98	0	Even	**Goal**	**1**

CONTEXT

- Context can be represented as a feature vector
  - E.g. score differential, manpower differential
- What do to about the previous match history?
- Simple Approach:
  - Fix a sliding window size *k* (common values are 3,4,10).
  - Use previous *k* events context for current event
  - Statsbomb approach

EXAMPLE: SLIDING WINDOW

k=2

Manpower	Action	Y=next goal
Even	Turnover	1
Even	Pass	1
Even	**Goal**	1

MP(-2)	Action(-2)	MP(-1)	Action(-1)	Manpower(0)	Action(0)	Y=next goal
*	*	*	*	Even	Turnover	1
*	*	Even	Turnover	Even	Pass	1
Even	Turnover	Even	Pass	Even	**Goal**	1

TRAINING A CLASSIFIER

- Given a list of pairs (context vector, success target), can simply run any classifier (R/Weka/Scikit-Learn)
  - E.g. logistic regression, gradient boosted decision trees
- Excerpt from logistic model tree: success probability of home team
- Alternative Approaches to Classification:
  - Recurrent Neural Networks (handles sequences)
  - Reinforcement Learning (value function learning)
    - My own work, RealAnalytics

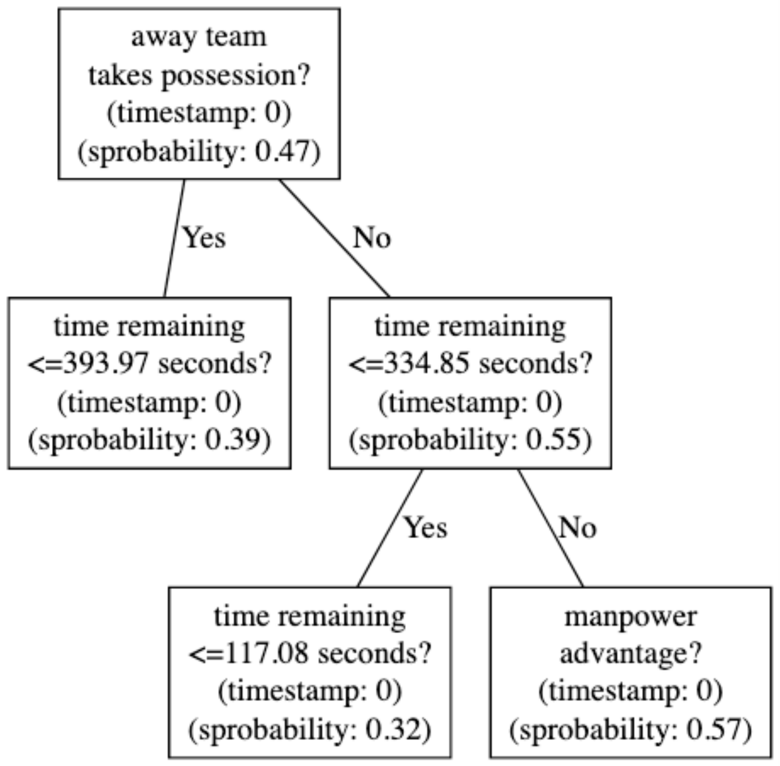

LEARNING SUCCESS PROBABILITIES: OVERVIEW

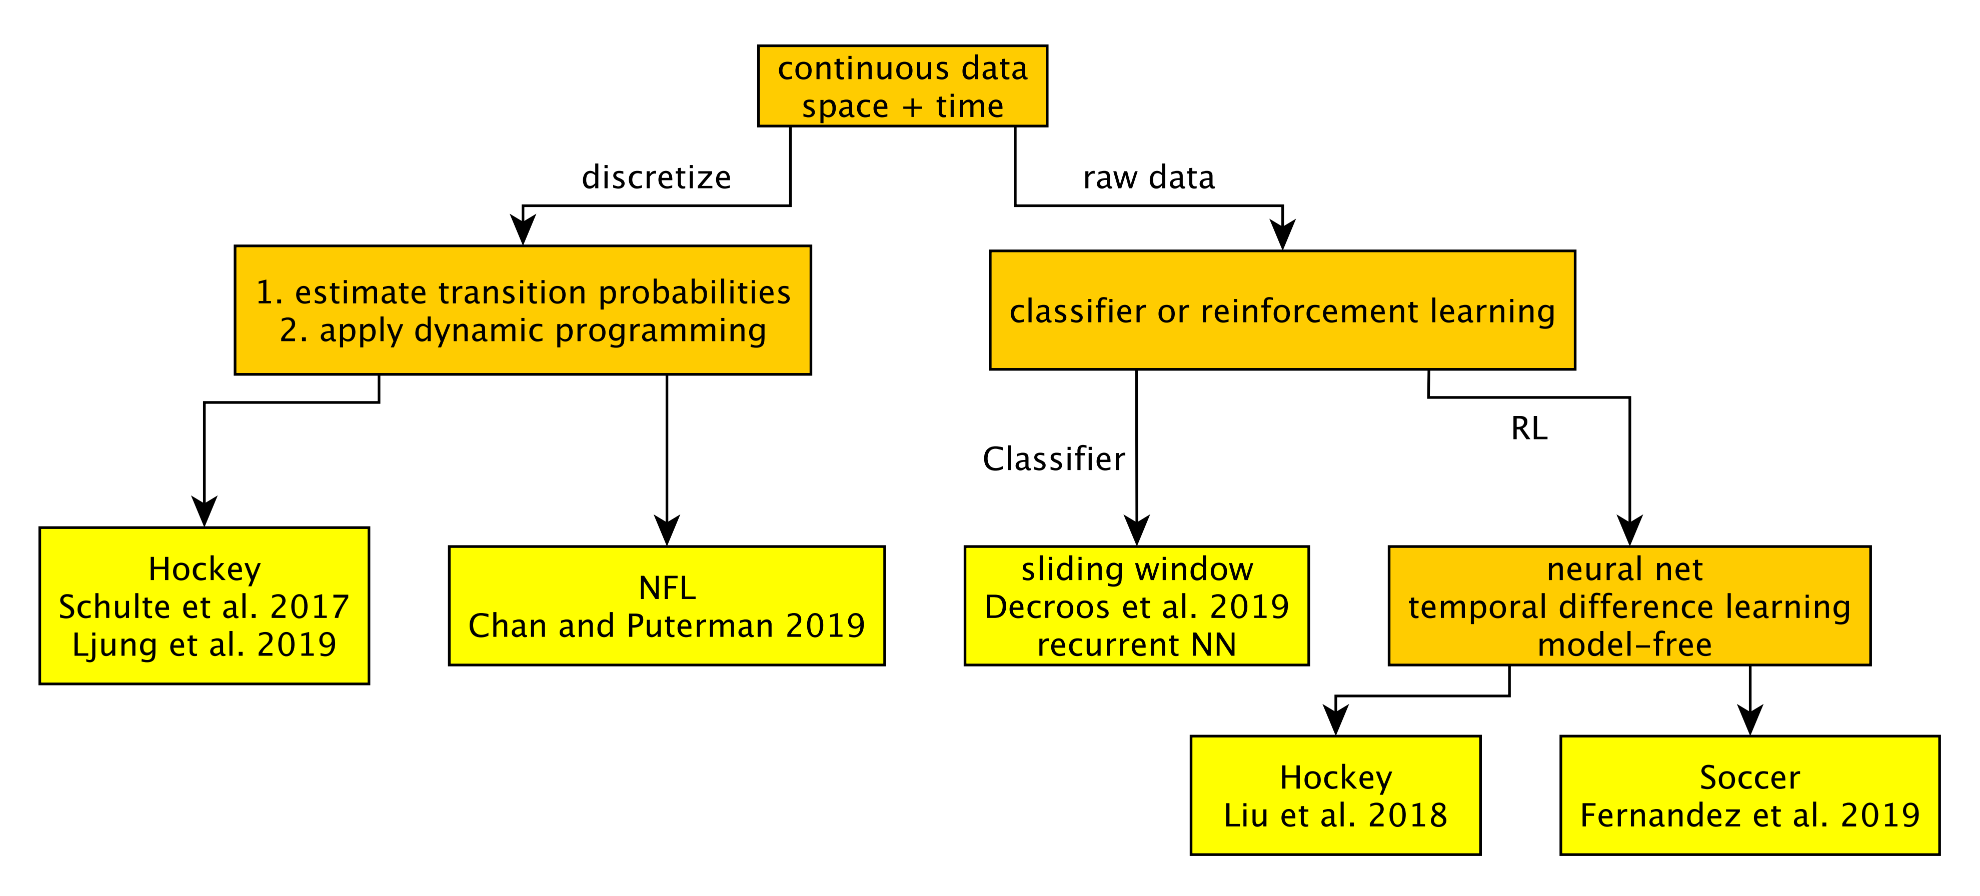

OPTIMIZATION PROBLEMS

OPTIMIZATION MAKES SPORTS ANALYTICS ACTIONABLE

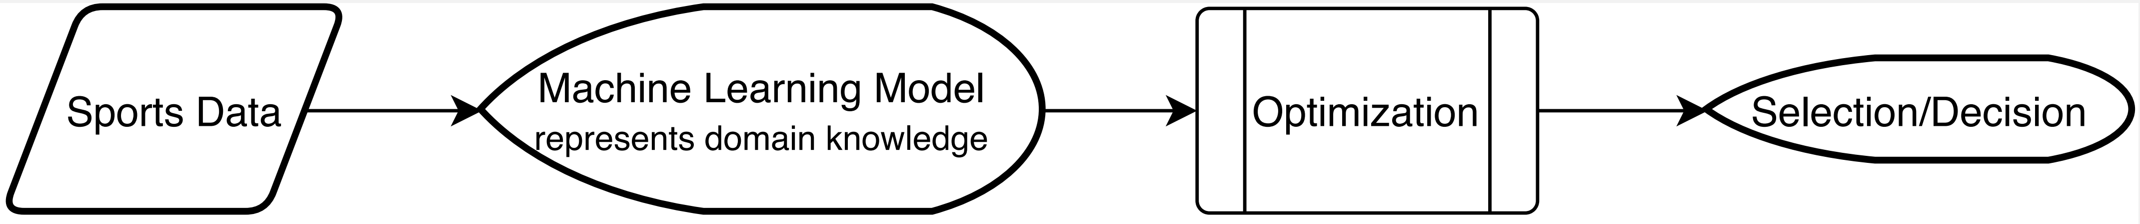

FANTASY LEAGUE/HOCKEY POOL

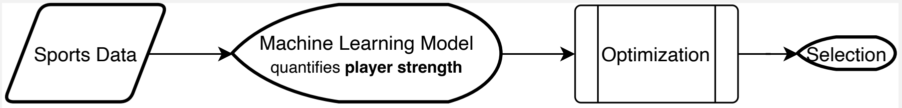

- See Tauhid Zaman's talk.
- You "draft" a team of *m* players (the *lineup*).
- Your players earn points through the season (goals + assists in real games)
- At the end of the season the "manager" with the most points wins the pool.
- Variations:
  - Players take turns drafting
  - add constraints on selections, e.g. at least 1 goalie, at most 4 forwards.
  - Lemmer (2013) applies integer programming with constraints
- Estimated 35M people play fantasy sports in North America (Becker and Sun 2016)

Becker, Adrian, and Xu Andy Sun. "An analytical approach for fantasy football draft and lineup management." *Journal of Quantitative Analysis in Sports* 12.1 (2016): 17-30.
Lemmer, Hermanus Hofmeyr. "Team selection after a short cricket series." European Journal of Sport Science 13.2 (2013): 200-206

REFERENCES

- https://www.officepools.com/
- Summers, Amy E., Tim B. Swartz, and Richard A. Lockhart. "Optimal drafting in hockey pools." *Statistical Thinking in Sports* (2007): 275-288. Defines the problem and some basic techniques.
- https://www.sloansportsconference.com/event/fantasy-sports-analytics MIT Sloan Sports Analytics has talks and panels on fantasy play

MONEYBALL

- The drafting problem but for real teams and players
- Commercial Software: https://octothorpesoftware.com
- Optimal Lineups *for a specific opponent* seems to be a new problem
  - e.g. optimal lineup in Cricket for England against India vs. against Australia
  - Perera et al. study optimal lineups against average opponents

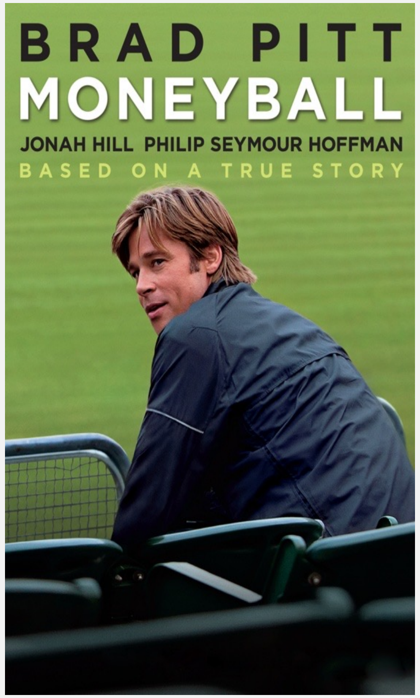

Perera, Harsha, Jack Davis, and Tim B. Swartz. "Optimal lineups in Twenty20 cricket." *Journal of Statistical Computation and Simulation* 86.14 (2016): 2888-2900.

TACTICS: FINDING OPTIMAL ACTIONS

- Basketball: go for 3 points (risky) or 2 points (safer)
- Basketball: doubling (e.g. put 2 defenders on LeBron James) Wang et al. 2018
- Hockey: when to pull the goalie? Beaudoin and Swartz 2010
- Often involves *counterfactual* questions: What if we tried a tactic that has never been tried before?
  - Related to *off-line reinforcement learning*

Wang et al. (2018) "The Advantage of Doubling: A Deep Reinforcement Learning Approach to Studying the Double Team in the NBA," 2018 MIT Sloan Sports Analytics Beaudoin, David, and Tim B. Swartz. "Strategies for pulling the goalie in hockey." *The American Statistician* 64.3 (2010): 197-204.

FINDING OPTIMAL BETS

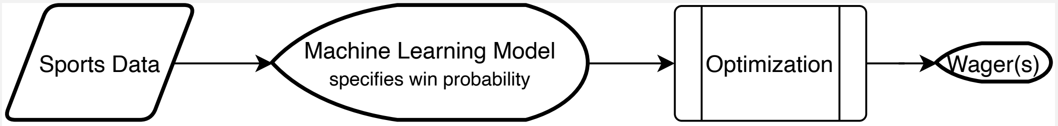

- Single bet scenario. Given:
  - Bankroll (e.g. $100)
  - Win probability p (from ML model)
  - Odds b from bookmaker (e.g. b= 2 if 2-1 return on winning bet)
- Output: fraction f of bankroll to be wagered
- *Kelly Criterion*: The maximum expected payoff is achieved by
  $f^* = p - (1-p)/b$

- Example
- $b = 2, p=1/2$ (even strength)
- F* = ½-(1/2/2)= ½-1/4 = ¼

➢ With $100 bankroll, you should bet $25

EXTENSIONS

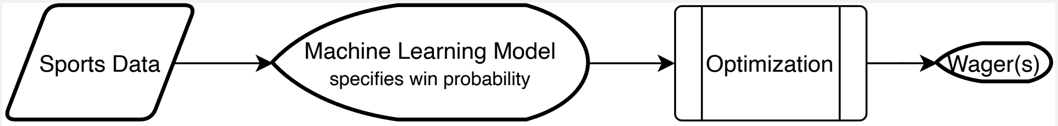

- Add constraints, e.g.
  - Maximum on bets
  - No risk of ruin (losing entire bankroll)
- Place bets with different bookmakers
- Bets on different matches and at different times (betting lines move) Insley et al. 2004
  - Related to portfolio management Lien et al. 2023

Insley, Robin, Lucia Mok, and Tim Swartz. "Issues related to sports gambling." *Australian & New Zealand Journal of Statistics* 46.2 (2004): 219-232.
"Contrastive Learning and Reward Smoothing for Deep Portfolio Management" Yun-Hsuan Lien, Yuan-Kui Li, Yu-Shuen Wang. IJCAI 2023

CONCLUSION

- Learning to <u>predict success probabilities</u> is a fundamental task for sports analytics
- Powerful approach to <u>action values</u> and <u>player ranking</u>
- Different machine learning models can be used
  - Classification, recurrent neural networks, *reinforcement learning*
- Machine learning provides the domain knowledge
- Optimization makes the domain knowledge actionable
  - Optimizing players, tactics, wagers

THANK YOU!

Kurt Routley

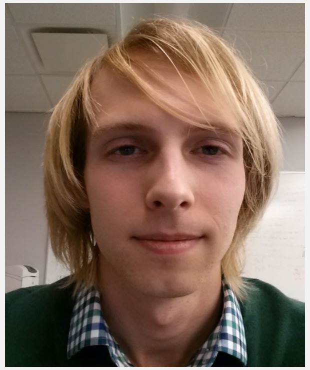

Zeyu Zhao

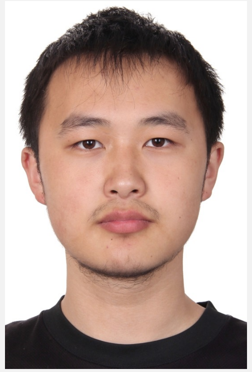

Guiliang Liu

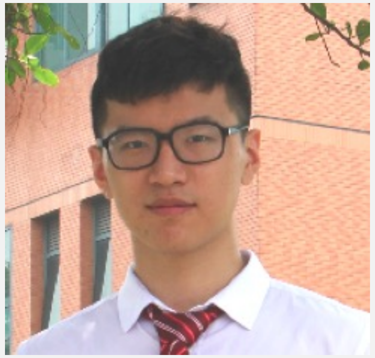

Pascal Poupart

BACKUPS

REINFORCEMENT LEARNING

See My [Aggregate Intellect Talk]

STATE TRANSITION PROBABILITIES

- Step 1: Estimate the probabilities of getting one from match state to the other
- Basketball Demo
- In our discrete NHL models, we estimated state transition probabilities for 1.3M states
- Step 2: Estimate the chances of reaching a success state using *dynamic programming*

MULTI-STEP TRANSITION PROBABILITIES

To compute $P_{l+1}(s_1|s_0)$: the probability of reaching state s_1 from s_0

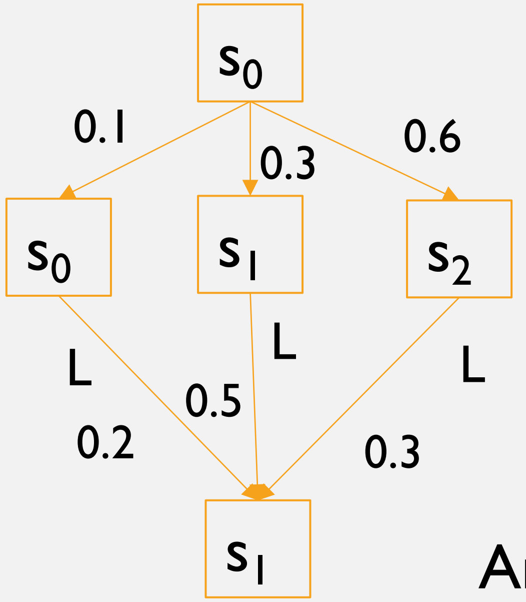

1-step transition probabilities $P(s^*|s_0)$

l-step transition probabilities $P(s_1|s^*)$

Answer: $0.1 \times 0.2 + 0.3 \times 0.5 + 0.6 \times 0.3$

DYNAMIC PROGRAMMING

- Input: State Transition Probabilities
- Output: Probability of Future Success for every match state
- For lookahead L = 1,…
  - Compute probability of success in L+1 steps using 1-step state transitions and L step success probabilities from previous lookahead
  - Terminate at convergence or at fixed bound
- For the NHL, our computation converged at L = 13
- Xthreat Visualization

MONTE CARLO VS. TEMPORAL DIFFERENCE

- Target = final outcome
- "Monte Carlo Learning"
- E.g. if a possession ends in a goal, then outcome = $target_t$ = 1
- Standard with sports analysts

+ Leverages supervised methods (e.g. classifiers)

- Ignores temporal dependencies and dynamics

- $Target_t = v_{t+1}$
- "Temporal Difference Learning"
- *Connects predictions at different times and for different actions*
- Standard method in RL

TOY EXAMPLE

Current Match State	Action	Current Score (Goal)	Estimated Next Goal Chance	TD-target	TD-error	MC-target	MC-error
[1-0, even strength, DZ]	Carry	0	55%	62%	$(55\%-62\%)^2$	1	$(55\%-100\%)^2$
[1-0, powerplay, OZ, …]	Pass	0	62%	75%	$(62\%-75\%)^2$	1	$(62\%-100\%)^2$
[1-0, powerplay, OZ, …]	Shot	0	75%	1	$(75\%-100\%)^2$	1	$(75\%-100\%)^2$
[2-0, ES]	Face-off	1					