Valuing Individual Contributing Events (V-ICE) in Hockey

Ethan Douglas, Sean Clement, Nick Wan, and Ian Greengross

Hockey - more than just goals

Hockey is the ultimate team game. All of the parts sum to the whole, which for teams and players is often goals. However, goals are a rare event in hockey - only a few are scored every game. The majority of a skater's time on the ice is spent performing actions to eventually score (or prevent an opponent from scoring), yet for most traditional hockey stats such as points or assists, only the goal itself or the direct previous two passes are logged in the stat sheet. Zone entries, takeaways, and passes may be recorded, but they are not "evaluated" in the context of their relative contribution towards helping a team achieve its goals. We don't know how valuable each individual action was.

Thus, we set out to answer the question: "How can we evaluate the contributions of all players on the ice for all of their actions?"

Soccer is another sport where current mainstream methods for evaluation rely on rare goals. A common method in both soccer and hockey at advancing this question is through an "Expected Goals" model [1,2,3,4]. This gives shooters or teams credit for the "quality" of their shots, even if they did not go in. While an improvement, Expected Goals is still only an evaluation of shots, hardly the comprehensive contributions of a player.

Expected Threat was developed in soccer to quantify how much value there is in moving to a "threatening" portion of the pitch [5]. However, it is purely location-based - a player moving the puck or ball to a specific location will always have the same value, agnostic of how the player got there. VAEP measures how each action contributes to increasing the chance of an eventual goal (and decreasing the chance of an opposition goal), and allows for more context of *how* a play develops rather than just where it develops [6].

In hockey, Total Hockey Rating (THoR) was conceived to address this question, however THoR event values are fixed and thus less context can be incorporated into the framework [7]. Additionally, state-based Markov models have been applied, however these models rely on a series of future events rather than a time-based window - and even two events apart may be too far to have a real influence on a goal, if the player with the puck is inactive for a long window [8].

More complex models have been created to evaluate the "real time" value of a possession [9, 10]. However, these frameworks require tracking data, which while available to teams in the NHL, is not necessarily available to leagues at all levels of hockey.

Thus, we wanted a flexible framework that: values events by their contribution to goals, captures the context of how plays develop, and only relies on event-level data. We now introduce V-ICE - "Valuing Individual Contributing Events." The V-ICE framework evaluates how individual actions contribute to an eventual expected goal (or prevent an eventual opponent's expected goal) and outputs a single, interpretable value for every action a player takes on the puck in an event-level dataset. The key advantage to this approach is that all actions are weighted by their value to a team in goals, so different types of actions can be directly compared. Additionally, this methodology is directly transferable to high fidelity tracking datasets available to teams. V-ICE could contextualize on or off puck actions throughout the duration of a play even if no event takes place, as our general framework and target would remain unchanged.

V-ICE: How to compute it

Data

Our data for this project was an event-level dataset composed of Stathletes-tracked women's hockey data from the 2018 Olympics, NCAA, and the 2021 NWHL season. This included the location (X and Y coordinates) of "events" (faceoffs, passes, shots, etc.) and the relevant game contextual information (score, time, etc.) at the time of the event. For the following models we do not differentiate the three different leagues because while there are certainly differences in play quality and style between them, we found the sample size sufficiently small such that any differences between leagues was more than offset by the increased sample from combining them.

Shot Probability Model

The first model in the series is a shot probability model. Our shot probability model drew inspiration from both Expected Threat (xT) and VAEP [5,6]. In adapting the framework to hockey data, we opted to use a time-based window rather than event-based. This better matched our intuition and experience with how hockey goals are scored - multiple passes or events in quick succession can all contribute to scoring, yet even just two events, if significant time has passed, may have little or no relation to each other or to an eventual score. In order to determine the optimal window, we started with an empirical approach of 45 seconds, the average time between line changes [11]. We then tested the performance of the framework on this window, as well as 25 and 15 seconds, and found the model performed best on the 25 second window. It's possible that the optimal window may differ if this framework is tested on a larger dataset.

The features for the shot probability model were:

1) Period
2) Time remaining in a period
3) Amount of skaters on ice for the home and the away teams
4) Whether the player with possession is on the home team
5) The goal differential
6) X,Y coordinates of the event
7) The type of event
8) Whether or not the event location is behind the goal line.

Furthermore, the previous two events were included using all of the same features as well as the event type and the event details. Lastly, the target player for a given event also has their shot share and goal share, with respect to their team up through the previous game-period. This feature adds the advantage that players with high shot or goal share for their team are likely skilled, and getting the puck to higher skilled players will be value captured by our framework.

A gradient boosted decision tree model (LightGBM) was used with the max xG in the 25 second window as a target.

Expected Goals Model

Second, an Expected Goals (xG) model was built to quantify the "value" of a given shot. Our xG model was heavily derived from the public Moneypuck xG model (http://moneypuck.com/about.htm). We trained a gradient boosted model using the xgboost package for R using the following features:

1) X ,Y coordinates of the shot origin
2) Shot probability (output from the shot probability model)
3) The angle between the shot origin and the goal
4) Seconds remaining in the period
5) The coordinates of the previous event on the ice
6) The number of seconds since the last event
7) The amount of time between the last shot event
8) The last shot angle
9) The number of defenders
10) The powerplay numerical advantage
11) The angle difference between the current and last shot.

The model was selected using AUC, having a test AUC of .737. On the withheld test set the model had an overall accuracy of 90.8% and a balanced accuracy of 59.2%. We trained the model taking advantage of xgboost's scale_pos_weight option which allowed us to more heavily weight correctly assigned goals than correctly assigned non-goals. The ratio we assigned to this parameter was the square root of the number of negative cases divided by the number of positive cases.

Bringing it all together

The outputs from the xG model were then joined to the shots in the event-level dataset, and a new shot probability model was built with the same features as the one described above, however instead of a classification model predicting the probability of a shot in the next 25 seconds, our target was the max expected goal output over the next 25 seconds. An additional model was built with a target of total expected goals over the next 25 seconds. This model utilized the same gradient boosted decision tree (LightGBM) architecture as the shot probability model, and had a Root Mean Square Error 43% lower than a basic naive model predicting the average xG for each string of events. This framework allows us the flexibility to see how actions influence scoring chances, shot quality or shot quantity.

$$VICE_n = max\left(xG_{next\ 25s\ pos\ team} - xG_{next\ 25s\ opp\ team}\right)$$
$$VICE\ ADDED = V\ ICE_n - V\ ICE_{n-1}$$

The below chart is an example play, with the V-ICE outputs for each event.

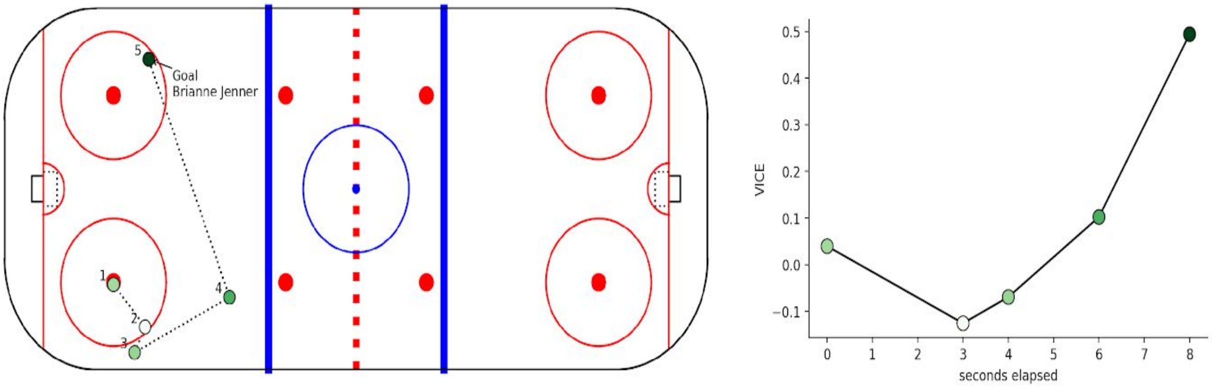

V-ICE: Event Values

		Predicted xG		
Eventr	**V-ICE**	**Total**	**Posession Team**	**Opponent Team**
Faceoff Win	0.04	0.54	0.63	0.09
Puck Recovery	−0.16	0.38	0.43	0.05
Play	0.05	0.43	0.41	−0.02
Play	0.17	0.60	0.62	0.02
Goal	0.40	1.00	1.00	0.00

Figure 1. Play diagram (left), event V-ICE values (right) and full model output (bottom) for an example play.

V-ICE: How to use it

This methodology can be applied by teams to both evaluate individual players, and to uncover strategic insights about a team's styles of play.

V-ICE: Top Skaters

		V-ICE	
	Skater	**Total**	**Per Game**
	Christina Putigna	7.3	1.0
	Mikyla Grant-Mentis	7.0	1.2
	Samantha Davis	6.7	1.0
	Melodie Daoust	6.3	0.9
	Sarah-Eve Coutu Godbout	4.8	0.8
	McKenna Brand	4.5	0.6
	Marie-Philip Poulin	4.4	0.5
	Mackenzie MacNeil	3.8	0.6
	Alyssa Wohlfeiler	3.4	0.9
	Brooke Boquist	3.4	0.8

Player Evaluation

An application of V-ICE is to determine which individuals contributed the most value to their team with their actions. Ideally we'd divide this number by the number of minutes the player spent on the ice, so that players with different playing times can be compared, however detailed participation data was absent from this dataset. Instead, we have divided the player's contributions by the number of games in which they recorded an event.The top ten players in the dataset regardless of league are listed in the table to the left.

While the data we have available does not cover a long enough window to perform stability analysis or other meta-checks on this player-level data, these results seem in line with expert consensus. We can see that Christina Putigna is the leading contributor in the NWHL to their team's expected goals. Close behind is Mikyla Grant-Mentis, the leading goal scorer of the 2020-2021 NWHL season. The highest ranking Olympic skater is Mélodie Daoust, who was awarded the Most Valuable Player award for the 2018 Olympics.

This player-level aggregation could be used to answer important broad questions as well, such as which line combinations contribute the most value to their team, or which positions are worthy of the highest contracts. However, we do not have sufficient data to answer these questions directly here.

We can also break this down further to aggregate a player's offensive and defensive contributions separately. For simplicity, we've considered defensive actions as all takeaways or puck recoverys occurring in the opposing team's offensive zone, and all other actions as offensive.

Figure 2. Player offensive and defensive contributions

Here we see that Christina Putigna, our highest ranking player by V-ICE, contributed as much defensively as she did offensively. This added defensive value is often not be captured in traditional stats like points, which is one possible reason why she leads our list

Even more granular, one could break down contributions by action type or region of the ice (e.g. defender with the most valuable contributions in front of the net). This allows coaches and talent evaluators to rank players by very specific skills their team may be in need of.

Tactics & Tendency Scouting

We know from our calculations that short direct and long indirect passes contribute the most to scoring chances. However, teams have different strengths and weaknesses, so in preparation for a game teams could use V-ICE to assess where their opponent loses value and try to exploit that matchup through player selection or modifying their on-ice defensive behavior. For instance - it would be useful to know how an upcoming opponent has most effectively created scoring chances. In the figure below we see that the Boston Pride's most successful short passes while in the offensive zone have been from the top corner.

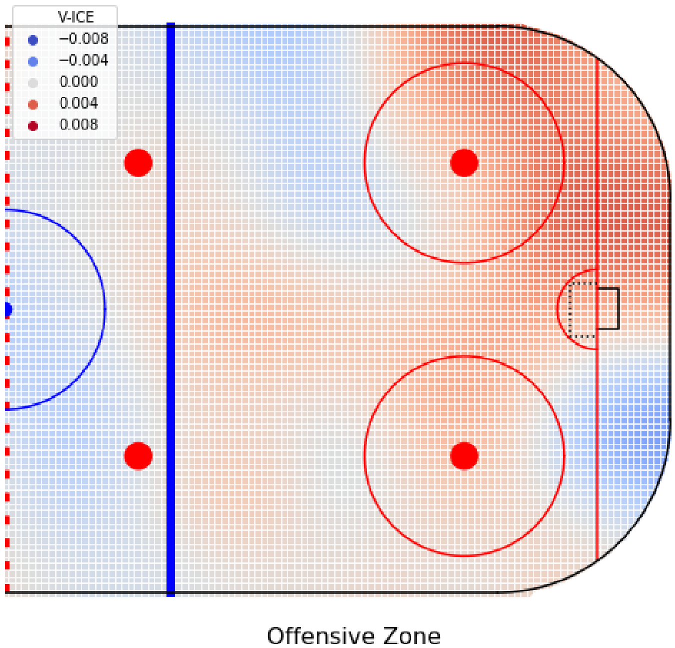

Figure 3. Boston Pride short direct passing V-ICE values

Conclusion

We've established a novel framework for Valuing Individual Contributing Events (V-ICE) so that every on-puck event in a hockey game can be analyzed for its contribution towards getting high quality scoring chances and thereby allowing us to see which players are best at driving those chances. We believe this system has advantages over current evaluations because, while goals ultimately determine the outcome of a game, consistently getting more quality chances will lead to more goals. We view V-ICE as a logical extension of VAEP using xG and shot probability instead of goals as it allows us to look not just as when we expect a goal to be scored, but why. This deeper look into the components of what makes a successful offensive attack in hockey could help teams not just with player evaluation, but in game strategy as well by evaluating their players' passing habits in a quantitative way in both terms of produced shot quality and quantity. Future work could explore the performance of this framework on a larger dataset as well as identify if separate models are needed for different levels of play.

References

1. Eggels, H. (2016). *Expected Goals in Soccer: Explaining Match Results using Predictive Analytics.* Eindhoven: Technische Universiteit.
2. Macdonald, B. (2012). An Expected Goals Model for Evaluating NHL Teams. *Sloan Sports Analytics Conference.* Boston: MIT Sloan.
3. Rathke, A. (2017). An examination of expected goals and shot efficiency in soccer. *Journal of Human Sport and Exercise*, S514-S529.
4. Tippett, J. (2019). *The Expected Goals Philosophy: A Game-Changing Way of Analysing.*
5. Singh, K. (2019, February 15). *Introducing Expected Threat (xT).* Retrieved from https://karun.in/blog/expected-threat.html
6. Decroos, T., Bransen, L., Van Hareen, J., & Davis, J. (2019). Actions Speak Louder than Goals:. *SIGKDD*.
7. Schuckers, M., & Curro, J. (2013). Total Hockey Rating (THoR): A comprehensive statistical rating of National Hockey League forwards and defensemen based upon all on-ice events. *Sloan Sports Conference.* Boston: MIT Sloan.
8. Routley, K., & Schulte, O. (2017). A Markov Game Model for Valuing Player Actions in Ice Hockey. *Data Mining and Knowledge Discovery*, 1735–1757.
9. Fernández, J., Bornn, L., & Cervone, D. (2019). Decomposing the Immeasurable Sport: A deep learning expected possession value framework for soccer. *Sloan Sports Analytics Conference.* Boston: MIT Sloan.
10. Cervone, D., D'Armour, A., Bornn, L., & Goldsberry, K. (2016). A Multiresolution Stochastic Process Model for Predicting Basketball Possession Outcomes. *Journal Of The American Statistical Association* , 585–599.
11. Jones, W. (n.d.). *Hockey Answered*. Retrieved March 4, 2021, from https://hockeyanswered.com/how-long-do-hockey-players-stay-on-the-ice-a-guide-to-shift-lengths/