A Novel Approach to Quantifying Zone Exits

By Arun Ramji and Ethan Batiuk

1 Introduction

In ice hockey, successfully exiting the defensive zone with puck possession is a critical factor in determining game outcomes. Failed zone exits and turnovers leave defenses vulnerable to counterattacks without defensive structure, increasing the likelihood of opposing high quality chances against. Data analysis from the Waterloo Warriors hockey team shows that exit rates are one of their most correlated statistics with wins, underscoring the importance of zone exits. However, limited research has explored the sequential dynamics of plays leading to successful exits versus turnovers or how the opposing team's offensive zone actions affect breakout success. Additionally, prior studies typically only credit the player completing the exit, overlooking contributions from other defensive zone actions, such as puck battles or passes preceding the final exit pass. This project addresses these gaps by developing a predictive model to identify play sequences that maximize successful zone exits, enabling us to quantify every action in the breakout sequence and provide actionable insights for team strategies and player evaluations.

Prior work, such as Corey Sznajder's hand-tracked microstats, has established the value of controlled zone exits and transition play. Subsequent studies have mapped optimal exit locations and developed expected goals (xG) models for player decisions. However, most rely on regression-based approaches or static metrics that overlook temporal dependencies in play sequences. Our approach employs a sequence-aware Long Short-Term Memory (LSTM) model, trained on event data from Linköping University, to predict zone exit success. By comparing our model's performance to traditional hockey exit statistics and existing predictive models, we quantify the added value of sequential modeling. This report details our methodology, findings, and conclusions, highlighting key play sequences, player contributions, and their correlation to generating shots, providing a novel framework for how to evaluate a breakout sequence.

2 Background

Ice hockey is a fast-paced sport where puck possession and territorial control significantly influence scoring opportunities. A zone exit occurs when a team moves the puck out of its defensive zone, ideally with control (e.g., via a pass or carry). Controlled exits are more likely to lead to shots and goals compared to uncontrolled exits, such as dump-outs, which involve clearing the puck without retaining possession. The opposing team's offensive zone actions, such as forechecking or positioning, can disrupt exits and impact breakout success. Traditional metrics, like exit success rates, focus on the final exit event, often ignoring preceding actions that contribute to the outcome. Advanced analytics, including microstats and xG models, have quantified the impact of exits but rely on static or hand-tracked data, limiting their ability to capture sequential patterns.

Our project uses event data from hockey games, including player actions (e.g., event name, type, x, y coordinates, success of the play), to model the dynamics of zone exits. By correlating exits to shot creation (shots generated post-exit), based on work by Thibaud Chatel, we evaluate their broader impact on the game. Unlike prior regression-based approaches, our sequence-aware model accounts for temporal dependencies, providing a more comprehensive picture of how breakouts occur. This framework is relevant to coaches, analysts, and sports scientists seeking to optimize team performance.

3 Algorithm

To model the sequential nature of zone exits in ice hockey, we selected a Long Short-Term Memory (LSTM) network, which captures temporal dynamics in play sequences. Zone exit attempts consist of events—puck retrieval, passes, dekes, and exit attempts—where success depends on their order and interaction. Unlike tree-based models that miss chronological dependencies, LSTMs inherently account for event order, handling variable-length sequences (2–30 events) effectively. The LSTM's gating mechanism prioritizes recent events while considering earlier actions to assign value to subtle contributions.
We chose an LSTM over a transformer-based approach due to its efficiency for our dataset. With sequence lengths of 2–30 events, transformers' strength in modeling long-range dependencies is unnecessary, and their quadratic complexity ($O(L^2)$ attention weights)

adds computational overhead. LSTMs process sequences in a single pass and using pack_padded_sequence avoids padding inefficiencies. However, transformers could be considered for larger datasets and tasks requiring multiple outputs (e.g., predicting shot quantity or quality post-exit).

For data capture, we structured sequences to reflect play flow, including the two events prior to the defensive team's first puck touch to contextualize exits. Sequences start when the exiting team gains possession in the defensive zone, with faceoffs treated as resets and final minutes excluded due to score effects. For exit attempts that did not follow a tagged zone entry, we included all touches in the defensive zone by the exiting team that led up to the attempt. We used embeddings for event and type features to capture underlying similarities between them, alongside five numeric features: spatial coordinates, score state, event outcome, and player position. A player rating metric was initially developed but removed to avoid evaluation bias, it increased performance of model by 2%. Our padding strategy with pack_padded_sequence ensures the LSTM ignores padded tokens, maintaining efficiency across varying sequence lengths.

4 Findings

4.1 Model Performance

The dataset was split into training, development and testing sets, with the test set comprising the following outcomes for zone exit attempts:

- **Successful Controlled Exits**: 17,502
- **Failed Exit Attempts**: 16,779
- **Successful Uncontrolled Exits**: 7,104

The performance of the Long Short-Term Memory (LSTM) model was evaluated using standard metrics, as shown in Table 1.

Metric	Value
Cross-Entropy Loss	0.4567
Accuracy	0.7818
ROC-AUC	0.9328
Brier Score	0.2846

Table 1: Evaluation metrics for the LSTM-based zone exit prediction model, assessing its ability to predict successful and failed zone exits.

The model performs robustly for the first publicly available model of its kind, achieving an accuracy of 0.7818 and a ROC-AUC of 0.9328, indicating strong discriminative power. Figure 1 illustrates the model's classification performance through a confusion matrix and reliability curves. The confusion matrix shows that the model correctly predicts 1,781 controlled exits, 2,210 uncontrolled exits, and 817 failed exits, with some misclassifications (e.g., 423 controlled exits predicted as uncontrolled). The reliability curves indicate that the model is well-calibrated, particularly for the "controlled" class, with predicted probabilities closely aligning with empirical accuracy.

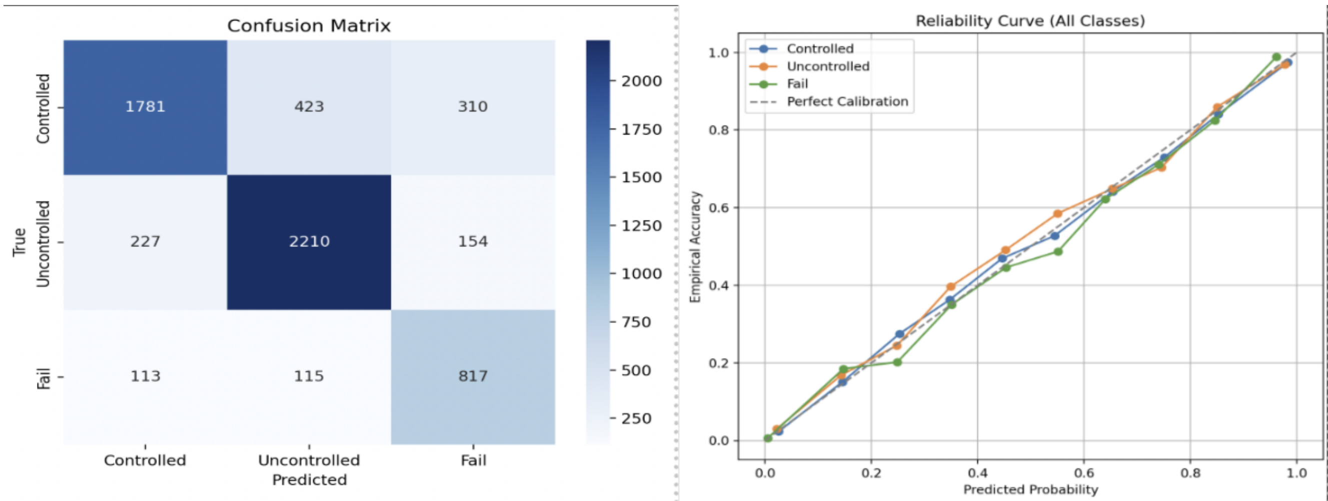

Figure 1: Model performance evaluation. (Left) A confusion matrix comparing true labels (rows: Controlled, Uncontrolled, Fail) against predicted labels (columns), with values indicating the number of instances (e.g., 1,781 true controlled exits predicted correctly). (Right) Reliability curves for all classes (Controlled: blue, Uncontrolled: orange, Fail: green), plotting predicted probability (x-axis) against empirical accuracy (y-axis), with a dashed line representing perfect calibration.

4.2 Classification Metrics by Class

To further assess the model's performance across classes, we computed precision, recall, and F1-score for each outcome (controlled, uncontrolled, fail). Figure 2 shows that the model achieves balanced performance, with the "uncontrolled" class having the highest scores (precision: ~0.82, recall: ~0.84, F1-score: ~0.83), followed by "controlled" and "fail." The slightly lower scores for "fail" (precision: ~0.78, F1-score: ~0.76) suggest that failed exits are more challenging to predict.

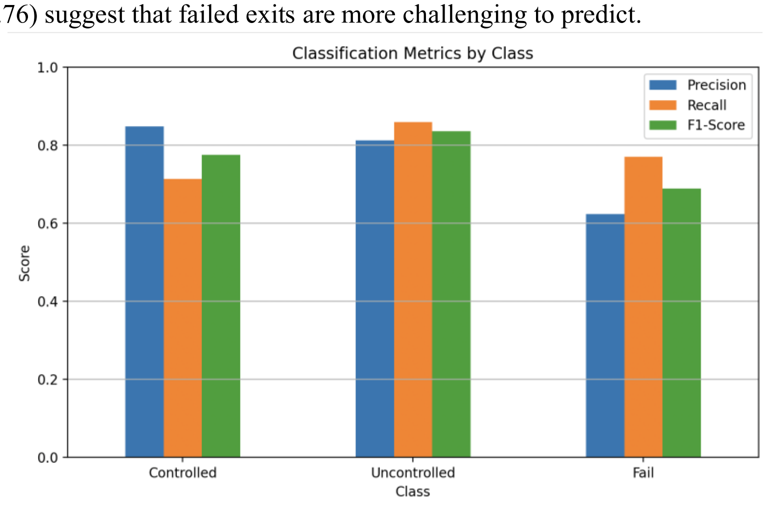

Figure 2: Classification metrics by class. A bar chart displaying precision (blue), recall (orange), and F1-score (green) for each class (Controlled, Uncontrolled, Fail), with scores ranging from 0.0 to 1.0 on the y-axis. The "Uncontrolled" class shows the highest metrics, while "Fail" has the lowest, indicating prediction challenges.

4.3 Player Rating Metric

A key application of the model is the development of a player rating metric based on zone exit contributions. Drawing on prior work by Thibaud Chatel and Mathieu Brosseau, who used a hockey decision tree model to assign expected shot values to zone exits, we assign shot generation values to exit outcomes. Unlike their approach, which differentiated exits by type (e.g., pass or carry), we simplify the framework by assigning average shot values to three outcomes:

- **Controlled Exit**: 0.38 shots
- **Uncontrolled Exit**: 0.16 shots
- **Failed Exit**: 0.07 shots

Using these values, we calculate a player's expected value added (EV) by weighing each puck touch according to the model's predicted probabilities of a controlled, uncontrolled, or failed exit. This metric quantifies a player's overall impact on successful zone exits, offering a novel approach to player evaluation.

This stat allows for a more comprehensive understanding of who is generating value by helping exit the puck.
Here are the top 5 and bottom 5 players, adjusted for touches in terms of expected shot generation, methodology referenced above. For a comprehensive list of player-level findings, refer to playerrating.py.

Player ID (Lowest)	Model EV	Player ID (Highest)	Model EV
71713.0	-0.103	270437.0	0.022
316087.0	-0.102	548693.0	0.018
509158.0	-0.090	608705.0	0.016
464287.0	-0.087	561317.0	0.014
461057.0	-0.075	348173.0	0.012

Table 2: 5 lowest and 5 highest players for expected shots created

We believe the stat has promise as a metric that captures players value in exiting the puck more than traditionally counting zone exits. As seen below if you apply our shot expected model to assigning exit credit the traditional way there is a moderate positive correlation, showing our stat has promise and that it goes a layer deeper.

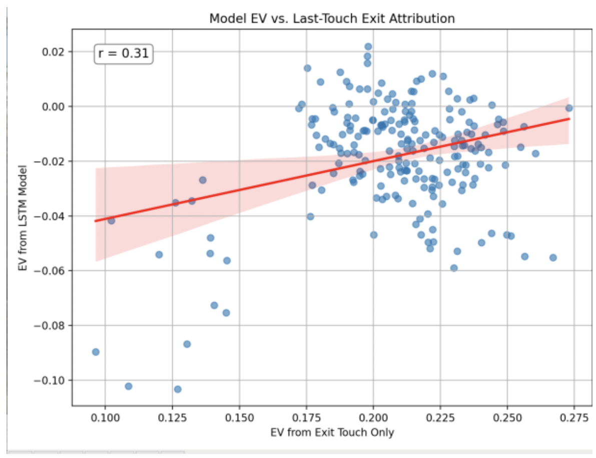

Figure 3: Model EV shift vs last touch EV shift

4.4 Event-Level Insights

Event-level analysis is another tool that can be used from this model, and it can be found in analysis.py. Our top takeaways are that when the team in the defensive zone has the puck, the only event that increases their expected shot value from an exit perspective is a pass. From an offensive team perspective, the worst thing you can do is dump the puck in and not be the first team to get possession, giving an average of 0.15 plus expected shot value to the opposing team. Additional takeaways and insights can be found by running analysis.py and reviewing the generated statistics.

4.5 Value Distribution Within Sequences

Building on the event-level insights, we further examined when expected value (EV) gains occur during a zone exit sequence. By analyzing the distribution of EV changes across sequence positions, as shown in the table below, we demonstrate the importance of our model by highlighting how preceding events contribute meaningfully to determining the outcome of a zone exit.

Sequence Position	% of Total EV Change
0–10%	+0.48%
10–20%	+0.79%
20–30%	+4.47%
30–40%	+11.46%
40–50%	+4.40%
50–60%	+29.81%
60–70%	+13.05%
70–80%	+11.50%
80–90%	+6.60%
90–100%	+17.43%

Table 3: Distribution of Expected Value (EV) Changes Across Sequence Positions

5 Summary of Key Ideas

This study introduces a novel approach to analyzing zone exits in ice hockey. Traditional metrics often focus solely on the final exit action, neglecting the sequential dynamics of preceding events, such as puck battles or passes, and the influence of opposing teams' offensive zone strategies. To address these gaps, the study employs a Long Short-Term Memory (LSTM) model, leveraging event data from Linköping University to predict zone exit success with high accuracy (0.7818) and ROC-AUC (0.9328), outperforming static regression-based models by capturing temporal dependencies. Additionally, a new player rating metric quantifies individual contributions to zone exits based on expected shot generation, offering a more nuanced evaluation of player impact compared to conventional exit counts. Event-level analysis further reveals that passes significantly enhance expected shot value for the defending team, while failed dump-ins by the offensive team concede substantial shot value to opponents, underscoring the strategic importance of precise puck management during breakouts.

6 Future Directions

6.1 Enhanced Data Capture and Analysis

To improve the useability of this model for the offensive team, our data capture could be expanded to more events then the last two plays the offensive team makes, allowing for deeper insights into how the actions the offensive team takes with the puck effect the defensive team's chances of exiting.

6.2 Better Tracking Data

Incorporating full player tracking data, capturing the location of every player on the ice, would significantly enhance the quality of the model and the depth of insights generated. This would allow for analysis of off-puck movement and could uncover valuable insights into areas such as which forecheck schemes most effectively reduce the odds of a successful breakout.

6.3 Refining the Player Rating System

The expected shot value framework could be significantly improved by incorporating additional metrics that better capture the impact of zone exit outcomes. Our current approach, based on prior work, does not consider the opposing team's likelihood of generating a shot following any of the three possible outcomes. Integrating this context—either through direct shot probability or an expected goals (xG) model—would meaningfully enhance the stat's accuracy and relevance. Furthermore, we believe this metric warrants deeper investigation. Exploring its correlation with player-level Corsi or traditional zone exit statistics could provide valuable insight into its validity and predictive power.

GIT: https://github.com/arunramji4/Linac25