Estimating NHL Screen Assists from Public Replay Tracking

Rento Saijo

October 20, 2026

Abstract

We estimate offensive screen assists from public NHL goal-replay tracking and assess player attribution against independent Clear Sight Analytics (CSA) annotations. The snapshot covers 2023–24 through 2025–26 and completed 2026–27 games through October 19, 2026. A geometric baseline and ridge logistic regression are compared using separate seasons for development, model selection, and evaluation. On the reserved 2025–26 reference set, conservative attribution has 92.60% precision and 73.21% recall; a broader definition increases recall to 85.85%, with precision of 84.42%. Player rankings report both definitions, rates based on official appearances and ice time, and a replay for each credited goal. A separate mixed-effects analysis of reviewed CSA shots estimates standardized scoring rates of 6.15% with offensive screens and 2.07% with clear sight. The attribution results support a measure of screening contributions on goals, while the shot analysis places those contributions in the context of an observational association with scoring.

Keywords: ice hockey analytics, sports analytics, screen assists, tracking data, mixed-effects, logistic regression

1 Introduction

A teammate who screens the goaltender can contribute to a goal without receiving an official assist. The location and timing of that screen matter: a player near the crease may leave the puck visible, while a brief crossing of the sightline around release may interfere with the goalie's view. Recording this contribution therefore requires a definition tied to the scoring shot. We examine the movement of potential screeners over a short interval and evaluate whether the resulting credits identify independently annotated participants.

Screening and traffic have been measured from several perspectives. Cait (2019) identifies screen assists manually on goals, covering both stationary and moving screens. Using tracking data, Pitassi et al. (2025) examine traffic between the puck and goalposts and its relationship with shot outcomes, accounting for shot location; the outcomes considered include whether attempts reach the net. Iaboni et al. (2025) combine tracking with video-derived player poses to estimate how much of the net is visible and reachable from the puck's perspective. These approaches address quantities that can differ on the same play: occupancy of a shooting lane, the area of net available to a shot, and obstruction of the goaltender's view.

Our player estimates use public NHL goal replays to identify offensive teammates whose tracked positions resemble independently annotated screens. Clear Sight Analytics (CSA) supplies the private reference for development and evaluation, and official appearances and ice time provide exposure denominators for the rankings. An accompanying replay explorer connects every credited goal to the underlying play. Because these data condition on a goal being scored, we examine screening's association with scoring in a separate CSA shot population. The two analyses address complementary parts of the problem: attributing screens on goals and comparing scoring outcomes across observed shot contexts.

2 Data and study populations

2.1 Public goals and player exposure

The study uses a fixed snapshot of completed regular-season and playoff games in 2023–24, 2024–25, and 2025–26, together with 2026–27 games dated through October 19, 2026. Shootout events are excluded. Public data are acquired with `nhlscraper` version 0.7.0.9000 (Saijo et al. 2026), with the package revision and public database snapshot recorded in the repository. To relate player credits to exposure, we use official skater appearances and time on ice, assigning team membership separately for each game.

Each public replay provides planar puck and player coordinates, identities, and timestamps at an observed interval of 0.10 s. The sequences usually span about 12–14 s, but their endpoints do not identify shot release or goal time. We order observations by timestamp, keep the first observation when timestamps are duplicated, and follow player identifiers as they move between tracking slots. All coordinates are in feet on a rink centered at the origin.

The public population comprises 26,457 non-shootout goals. Stored replays cover 99.22% of the 2023–24 goals before tracking-quality exclusions; over all included seasons, 22,570 goals

have both a usable shot interval and candidate teammates. Table 1 separates these evaluated goals from empty-net goals and other unevaluable sequences. Goals whose tracking cannot support attribution retain an unevaluable status.

Table 1: Public goals after shootout exclusion, separated by replay eligibility. The 2026–27 sample covers completed games through October 19, 2026.

Season	Goals	Evaluated	Empty net	Unevaluable
2023–24	8,590	7,379	485	726
2024–25	8,435	7,098	568	769
2025–26	8,572	7,336	547	689
2026–27	860	757	44	59

2.2 Independent reference annotations

The reference comes from a read-only query of CSA's historical reporting database, with private records retained in memory. Eligible events are game-reviewed, quality-control complete, classified as shots, and linked to an identifiable goalie. A public goal matches a CSA event only when game identity, scorer, shooting team, home and away teams, goalie, and period agree. An imported event-identifier match permits a clock difference of at most one second. If that match is unavailable or invalid, the fallback requires a unique exact clock-and-scorer match and the same contextual agreement. We map participants through NHL identifiers and remove duplicate event–player–side memberships.

Event labels require explicit evidence in the CSA classifications. A positive event has an offensive-screen, mixed-screen, one-timer/screen, or screen/deflection classification and names at least one offensive teammate other than the scorer. A negative event has a clear-sight or defensive-only classification and records no offensive teammate. Screen-on-pass cases and screen-classified events without a named offensive teammate are excluded. Consequently, a missing participant row by itself cannot establish a negative event.

For events with recorded participants, the participant list defines the player-attribution reference. Other candidate teammates are treated as negatives, and named participants absent from usable tracking count as unrecovered positives. This construction assumes a complete participant list. The available documentation also leaves the operational threshold for brief screens and frame-level screen timing unspecified, so disagreements can reflect coding scope as well as detector error. The matched and evaluable populations in Table 2 reflect these eligibility rules; their screening prevalence can differ from prevalence across all NHL goals.

Table 2: CSA reference populations by season. Matched goals are counted before label and replay-quality restrictions.

Season	Matched goals	Benchmark goals	Screen-positive goals
2023–24	7,587	1,687	511
2024–25	7,245	1,730	528
2025–26	7,813	1,747	494

3 Estimating a screen assist

3.1 Release and arrival

Shot timing is reconstructed from the puck trajectory after rotating each sequence to place the defending goal at $x = 89$ ft. We identify crossings toward the goal plane and retain those whose interpolated lateral coordinate is within 4.50 ft of goal center, accommodating tracking displacement around the 6 ft goalmouth. Arrival is the first sample at or beyond the plane. It anchors a backward search of up to 2.00 s for release, with candidate intervals assessed using puck speed and agreement with the official shot location.

For an ordinary shot, a release candidate must move toward goal at at least 12 ft/s and lie within both 22 ft of the recorded shot location and 12 ft of the scorer's sensor. Puck–scorer separation must then increase by at least 2 ft over the next three samples, or over the remaining samples if arrival occurs sooner. Three samples ordinarily span 0.30 s. Among eligible candidates, we minimize location error divided by 10 ft plus scorer separation divided by 5 ft. The spatial tolerances allow for displacement between a stick release and a shoulder sensor, as well as imprecision in official shot locations; stick contact itself is unobserved.

Tips and deflections require a different reconstruction. We follow the incoming flight backward through consecutive fast-moving samples and check the credited scorer's position near the terminal deflection to approximate the original release. For a rebound goal, the scoring attempt and its recorded location define the sequence, so preceding traffic contributes only when it falls within that attempt's analysis window. Every sequence requires at least 0.50 s of preceding tracking. Critical intervals are rejected if they contain gaps longer than 0.21 s or lack puck or goalie positions; two candidate arrivals more than 0.50 s apart with matching costs differing by less than 0.25 leave the sequence ambiguous. Timing accuracy remains constrained by the 10 Hz sampling rate, planar coordinates, and absence of puck-touch annotations.

3.2 Finite sightline features

The geometry is defined by the tracked positions of the goalie, puck, and candidate player at time t, denoted g_t, p_t, and s_{it}, respectively. With $v_t = p_t - g_t$, we calculate

$$u_{it} = \frac{(s_{it} - g_t)^\mathsf{T} v_t}{v_t^\mathsf{T} v_t}, \qquad d_{it} = \left\| (s_{it} - g_t) - u_{it} v_t \right\|.$$

The projection condition $0 < u_{it} < 1$ places the candidate between the goalie and puck; we evaluate obstruction only when the segment also exceeds 2 ft. Features summarize the 0.50 s before release, the band from -0.20 to $+0.10$ s around release, and flight toward the goalie. Post-release windows end at arrival and exclude puck positions more than 1 ft beyond the goalie's longitudinal position. To calculate duration, we weight intervals centered on observed samples by their overlap with each window. Reacquisition time measures the interval since the most recent obstruction in the preceding 1.00 s, taking a capped value of 1.00 s if no obstruction is observed. Candidate players require at least three valid samples and 80% observed positions in the feature window.

The resulting 16 predictors capture minimum and release-time separation; obstruction duration under primary, narrower, and wider tolerances; longest continuous duration; and flight obstruction. They also describe alignment strength before, near, and after release, reacquisition time, goalie proximity, relative depth, crossing, shot distance, and flight time. Alignment strength for an effective radius r is

$$A_{it}(r) = \frac{\exp\{-d_{it}^2/(2r^2)\}}{1 + \|s_{it} - g_t\|/20},$$

and is zero outside the finite segment. The geometric baseline combines the peak alignment near release or during flight with pre-release obstruction duration D_i and reacquisition time R_i, both in seconds:

$$B_i = \max(A_{i,\text{near}}, A_{i,\text{flight}})\{0.25 + \min(1, D_i/0.30)\}\exp(-R_i/0.40).$$

Close alignment and sustained obstruction increase this score, as does a shorter interval for reacquisition before release.

3.3 Position and timing uncertainty

Interpreting this geometry requires accounting for what the sensors measure. The right-shoulder placement described by Radke et al. (2021) separates a tracked position from the player's body center and, for a goalie, from the eyes. Skating direction cannot recover orientation: players can move backward or sideways and can remain stationary. The replays also omit puck height, gaze, stance, and the visual obstruction caused by limbs or equipment. The finite sightline thus approximates a three-dimensional visual process with planar positions.

The primary effective tolerance is 1.50 ft. Alternatives of 1.00 and 2.00 ft assess how body extent and sensor displacement affect the estimates, and separate shifts of -0.10 and $+0.10$ s in estimated release examine timing sensitivity. These scenarios retain the selected model and thresholds. When a shifted feature window has insufficient tracking, the primary feature values are retained. The scenarios assess sensitivity to these changes without defining a calibrated distribution of anatomical error or covering every possible posture.

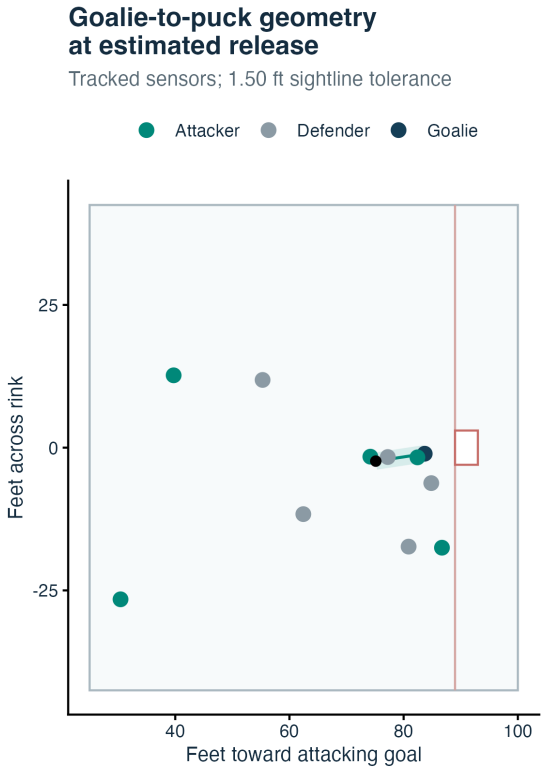


Figure 1: Estimated release in an actual replay. Sensor positions define the finite goalie-to-puck segment, and the band represents its effective geometric tolerance.

3.4 Model selection and attribution

Development, model selection, and final evaluation use separate seasons: 2023–24, 2024–25, and 2025–26, respectively. Keeping entire seasons separate also keeps games intact. The candidate models are the geometric baseline and ridge logistic regression fitted through `tidymodels` and `glmnet` to standardized public features, with an unpenalized intercept. Neither player identities nor CSA scoring probabilities enter the predictors. We compare five requested ridge penalties, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, and 1, along the fitted regularization path; requests below its minimum use the terminal coefficients. Player-level average precision determines selection, grouping tied scores and favoring the baseline when methods tie.

The selected detector supplies scores for two attribution rules. The conservative threshold maximizes

$$F_{0.5} = \frac{1.25\,PR}{0.25P + R},$$

where P and R are precision and recall. The broader threshold maximizes $F_1 = 2PR/(P + R)$ and is constrained to be no higher than the conservative threshold. Their selected values are 0.6155 and 0.2443, respectively. Coefficients and thresholds are frozen before final-season evaluation. Under either rule, each qualifying offensive teammate receives one credit per goal, with the scorer excluded. Several teammates can qualify on a goal, including official

assist recipients, and all conservative credits also appear in the broader total. At the goal level, at least one qualifying teammate produces a positive prediction.

Evaluation reports precision, recall, and average precision, together with calibration for the classifier. Threshold-specific 95% intervals account for clustering within games through a sandwich calculation for each ratio. For precision $\hat{P} = \sum_g T_g / \sum_g N_g$, for example, the estimated variance is

$$\widehat{\mathrm{Var}}(\hat{P}) = \frac{G}{G-1} \frac{\sum_g (T_g - \hat{P} N_g)^2}{(\sum_g N_g)^2}.$$

Here G denotes the number of games in the benchmark. Wald interval endpoints are truncated to $[0, 1]$. These intervals quantify benchmark sampling variation conditional on the chosen detector and reference labels, leaving annotation and model-selection uncertainty outside their scope.

Tracking-frame review helps connect the numerical features to possible failure mechanisms. We inspect 12 development-season player–goal cases across 11 goals, together with an additional unresolved trajectory. The cases cover strong alignment, participant disagreements, multiple potential screeners, tips, a rapid follow-up scoring attempt, and a crossing near an attribution threshold. Frames before, at, and after estimated release show how those situations enter the features. They cannot independently establish the goalie's view and supply no additional estimate of accuracy.

3.5 Exposure and rankings

Player totals and exposure share the same season, competition, team, and cutoff filters. For player i, the corresponding rates are

$$\text{Screen assists per game}_i = \frac{C_i}{G_i},$$
$$\text{Screen assists per 60 min}_i = \frac{3600 C_i}{T_i},$$

where C_i denotes estimated credits under the selected definition, G_i official appearances, and T_i ice time in seconds. Each traded player's records follow the team represented in that game, and eligibility includes players with zero detected credits. Counts have no exposure minimum. Rate rankings default to 25 games or 250 minutes for completed seasons and the combined all-seasons view, and to five games or 60 minutes for the ongoing season. Playoff-only rankings begin without a minimum. Readers can adjust these thresholds; calculating a rate always requires positive exposure.

Official exposure includes games whose goal replays are missing or unevaluable, so the rates combine detected credits from evaluable goals with all official appearances or ice time. They measure screen assists detected per game or unit of playing time under that coverage. Missing replays can depress the rates because unevaluable goals receive no imputed credits.

4 Screening and scoring

The goal-replay rankings describe contributions conditional on scoring. To examine the association between screening and scoring itself, we use a separate CSA shot population from 2024–25 and 2025–26, when release-type information is available. CSA describes its central coverage as shot sequences resulting in shots on goal (Clear Sight Analytics n.d.); this coverage should not be treated as a complete record of blocked and missed attempts. The primary comparison includes reviewed, quality-complete offensive-screen and clear-sight shots with a goalie present. It excludes screen-on-pass cases, posts, and deflection releases, as well as the separate mixed-screen, defensive-screen, one-timer/screen, and rebound categories. Each included season–strength–precise-location–release stratum must contain both screened and clear-sight observations, restricting the comparison to shared observed contexts.

We model scoring with the prespecified mixed-effects logistic model

$$\operatorname{logit}\{\Pr(Y_j = 1)\} = \alpha + \beta S_j + x_j^{\mathsf{T}}\gamma + b_{\text{shooter}(j)} + b_{\text{goalie}(j)} + b_{\text{game}(j)}.$$

The response Y_j indicates a goal, and S_j indicates an offensive screen. Adjustment through x_j includes season, regular-season/playoff status, strength, precise shot zone, release type, broad play context, first and second preceding-pass zones, and attacking and defending teams. Independent mean-zero normal distributions describe the shooter, goalie, and game intercepts. Missing pass and context values remain in the analysis as an explicit unknown category, and CSA scoring probabilities are excluded.

Estimation uses `tidymodels` with the `multilevelmod`/`lme4` engine and penalized iteratively reweighted least squares (`nAGQ = 0`). To put the two groups on a common covariate distribution, we predict screened and clear-sight outcomes over the screened shots in the shared strata. Standardization integrates the fitted random-intercept distributions using 41 deterministic normal-quantile nodes. We report the conditional contrast $\exp(\beta)$ as the adjusted odds ratio with a fixed-effect Wald interval, and use delta-method intervals for the standardized rates and their difference. Each interval calculation holds the variance components at their fitted estimates.

5 Results

5.1 Independent attribution performance

The ridge detector is selected with tuning-season average precision of 0.893, compared with 0.835 for the geometric baseline. Its reserved-season average precision is 0.911 at the player level and 0.931 at the goal level. The final benchmark comprises 1,747 goals and 7,260 player–goal units, of which 530 are annotated offensive teammate contributions. Conservative attribution recovers 388 of these contributions and assigns 31 credits to participants absent from the reference. The broader definition recovers 455 contributions, accompanied by 84 such disagreements.

Table 3: Reserved-season performance for 2025–26. Intervals are 95% and account for clustering within games.

Definition	Task	Precision (%)	Recall (%)
Conservative	Player	92.60 (90.04, 95.16)	73.21 (69.56, 76.85)
Conservative	Goal	93.70 (91.31, 96.09)	75.30 (71.51, 79.10)
Broader	Player	84.42 (81.41, 87.42)	85.85 (82.92, 88.78)
Broader	Goal	87.27 (84.35, 90.19)	87.45 (84.57, 90.33)

Both definitions identify screened goals with higher precision and recall than they identify individual participants (Table 3). Detecting a screened play therefore leaves a further attribution task. The held-out Brier score is 0.0171, but calibration varies across the score range: intermediate-score bins contain higher observed participant frequencies than their mean predictions, and the large concentration of nonparticipants in the lowest bin contributes to the low overall score. The classifier supplies scores for empirically selected attribution thresholds; calibration beyond the selected reference population remains unknown.

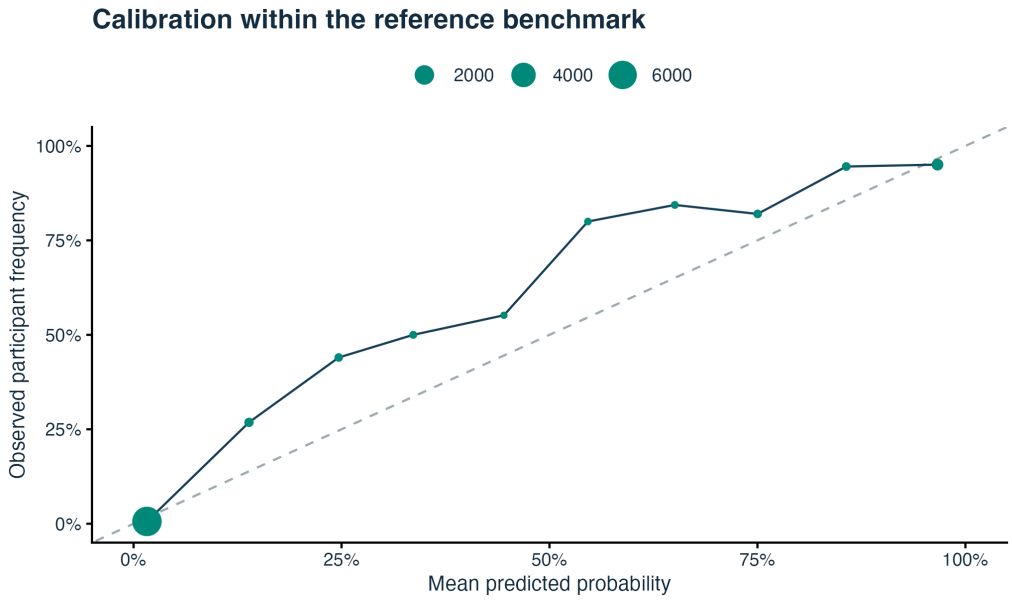


Figure 2: Reserved-set calibration, with point size reflecting bin count. Intermediate-score bins underpredict participant frequency; calibration beyond this selected benchmark remains unknown.

5.2 Player rankings and sensitivity

Over all included seasons and competitions, the detector assigns 2,211 conservative credits and 3,320 broader credits. For the 2025–26 regular season, Alexander Wennberg leads with 13 conservative credits, followed by Joel Eriksson Ek and Chris Kreider, each with 10. Their totals and exposure-based rates appear with the other leading players in Table 4. The

website opens on this most recent completed season, 2025–26, and links each credited goal to its replay.

Table 4: Leading conservative screen-assist totals in the 2025–26 regular season, with rates calculated from conservative credits.

Player	Conservative	Broader	Games played	Screen assists per game	Screen assists per 60 min
Alexander Wennberg	13	14	80	0.16	0.48
Joel Eriksson Ek	10	15	70	0.14	0.45
Chris Kreider	10	11	75	0.13	0.47
Anze Kopitar	9	9	67	0.13	0.42
Alex Tuch	7	9	79	0.09	0.28
Corey Perry	7	8	72	0.10	0.43
Jake Guentzel	7	8	81	0.09	0.26
Anders Lee	6	10	82	0.07	0.28
Charlie Coyle	6	9	82	0.07	0.24
Dmitri Voronkov	6	8	63	0.10	0.41

A player's zero total means that no contribution is detected under the selected definition and filters, even if that player screens on other shots. The development examples show how such distinctions can arise within a play. Alignment can change over a few samples as a player crosses the segment, and a named participant can remain outside the sensor-defined sightline while body position is unobserved. One trajectory remains unresolved because the tracked puck crosses the goal plane 6.38 ft from goal center, beyond the goalmouth tolerance; it retains an unevaluable status.

Changes in effective tolerance and estimated release timing alter the number of qualifying contributions (Table 5). Player totals thus depend on the measurement assumptions as well as the attribution threshold. The conservative and broader columns apply two operating definitions to the primary features; the sensitivity scenarios examine how additional positional and timing uncertainty affects credits with the detector held fixed.

Table 5: Sensitivity to geometry and release timing with the detector held fixed, combining the 2025–26 regular season and playoffs.

Scenario	Conservative	Broader
Primary: 1.50 ft	709	1,077
Tolerance: 1.00 ft	535	865
Tolerance: 2.00 ft	863	1,276
Release: 0.10 s earlier	800	1,209
Release: 0.10 s later	581	961

5.3 Adjusted scoring association

The shared-stratum population includes 3,697 screened shots and 54,690 clear-sight shots from 2,792 games, retaining 96.40% of otherwise eligible screened shots. Before adjustment, screened shots score at rates of 5.86% in 2024–25 and 6.45% in 2025–26. Clear-sight rates in the same seasons are 3.10% and 3.20%, respectively.

Adjustment and standardization to the screened-shot covariates give a screened scoring rate of 6.15 (5.43, 6.88) percent and a clear-sight rate of 2.07 (1.88, 2.25) percent, with 95% intervals in parentheses. Their standardized difference is 4.09 (3.34, 4.84) percentage points, and the conditional adjusted odds ratio is 3.47 (2.92, 4.13). These estimates describe the association between screening and scoring in the included shot population.

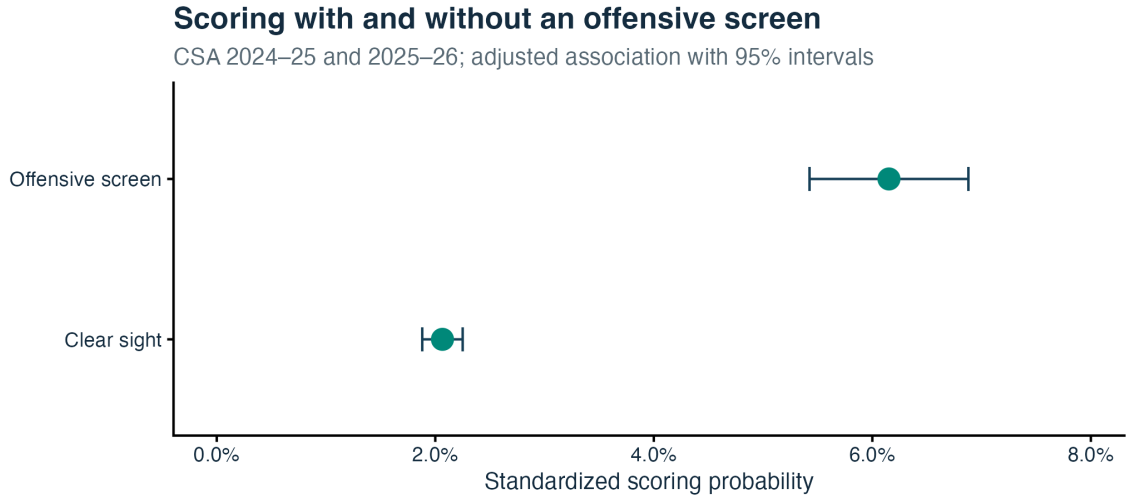


Figure 3: Standardized scoring rates and model-based 95% intervals. Both predictions use the covariate distribution of screened shots within common support.

The optimizer converges at a boundary fit, with the goalie random-intercept variance equal to zero. The shooter and game variances are approximately 0.0052 and 0.0049 on the log-odds scale. We retain that fit and calculate intervals conditional on the estimated variances. Unmeasured shot characteristics, detailed pass timing, player selection, and classification practices can also contribute to the observed association.

6 Discussion

The detector recovers many independently annotated screening participants from public replay tracking, giving the player rankings an empirical basis for attribution. The two thresholds expose a useful trade-off: conservative credits have higher precision, while broader credits recover more reference contributions and include more disagreements. Reading these totals alongside the credited replays connects a player's position in the rankings to the plays that produce it.

That position depends on opportunities as well as screening activity. Every credit conditions on a goal, leaving screens on saved, missed, and blocked shots outside the measure. Team shooting volume, offensive role, special-teams deployment, and teammate finishing all affect the opportunities to receive credit. Official appearances and ice time put the rates on consistent exposure denominators, but cannot remove those differences. The resulting counts and rates describe detected contributions on goals, with a player's wider screening activity left unmeasured.

The same care is needed when interpreting differences between players. Shoulder-mounted positions reduce a three-dimensional visual process to a goalie-to-puck segment, and brief crossings may occur between samples. Deflections, rebounds, and goalie contact complicate release inference; unevaluable trajectories may also be distributed unevenly across teams or types of play. The sensitivity results show that totals respond to effective tolerance and release timing, so an apparent difference between players can depend partly on assumptions about positions and events the tracking does not observe.

Reference agreement also has a defined scope. Explicit classifications and participant identities make evaluation possible, while the eligibility rules exclude ambiguous plays and alter screening prevalence. Errors in named participants and obstruction by unlisted teammates remain possible. Reported performance consequently describes agreement within the selected reference population. A broader collection of independently reviewed plays, combining tracking with video and clearer annotation rules, would help establish how well the method extends beyond that population.

The separate CSA shot analysis adds evidence that offensive screens are associated with scoring after adjustment for observed context. Its population and estimand differ from those of the goal-replay detector, so the adjusted odds ratio and standardized rate difference cannot be multiplied by a player's screen-assist total to obtain causal goals added. Estimating that quantity would require broader tracking, an account of shot selection, and a defensible counterfactual design. The present rankings support comparisons of attributed screens on goals, with credited replays and measured reference agreement available to assess those comparisons.

7 Data and code availability

The website brings together player rankings, replay exploration, methodology, and the complete paper. The repository supplies the R code and editable Quarto source, explicit detector coefficients, NHL-derived goal–player credits, player-game exposure, short replay segments, and aggregate findings. These public artifacts reproduce the published ranking arithmetic, figures, and report; provenance records the `nhlscraper` revision and public database snapshot. Re-estimating the CSA-dependent models requires authorized access because CSA records, event-level labels, and fitted private model objects are not distributed. All published results refer to the fixed snapshot through October 19, 2026.

References

Cait, J. (2019), '2019–20 screen assists: Measuring impactful net-front presence', Medium. https://medium.com/@justin.cait/2019-20-screen-assists-measuring-impactful-net-front-presence-updated-monthly-d5984e1d5990

Clear Sight Analytics (n.d.), 'What we do'. Accessed October 20, 2026. https://www.csahockey.com/what-we-do

Iaboni, E., Negulescu, S., Pitassi, M., Nazemi, A., Bright, J., Chomko, V., Clausi, D. A., Dickinson, S. & Brecht, T. (2025), Position paper: New views of shots – towards measures of net visibility and reachability, *in* 'Linköping Hockey Analytics Conference, Research Track'. https://cs.uwaterloo.ca/~brecht/papers/linhac-2025-netviz.pdf

Pitassi, M., Iaboni, E., Lodhi, F. & Brecht, T. (2025), Examining the impact of traffic on shot attempts in ice hockey, *in* 'International Sports Analytics Conference and Exhibition'. https://cs.uwaterloo.ca/~brecht/papers/isace-2025-traffic.pdf

Radke, D., Radke, D., Brecht, T. & Pawelczyk, A. (2021), Passing and pressure metrics in ice hockey, *in* 'Artificial Intelligence for Sports Analytics (AISA) Workshop at IJCAI '21'. https://cs.uwaterloo.ca/~brecht/papers/hockey-aisa-2021.pdf

Saijo, R., Skytte, L. & Pallotta, J. (2026), *nhlscraper: Scraper for National Hockey League Data*. R package version 0.7.0.9000. https://rentosaijo.github.io/nhlscraper/