Contrib

Learning individual performance scores from team outcomes in Counter-Strike

XuHaoran 2026-09-18

Abstract. In professional Counter-Strike, players share a team result but are rated individually. CONTRIB learns a shared scoring rule from team outcomes, combining a statistical base with a bounded adjustment. Scores add to team totals, and score differences can be traced to their components. Across Counter-Strike 2 (CS2) and Counter-Strike: Global Offensive (CS:GO), scores agree closely with stored HLTV ratings, including among teammates. The adjustment changes individual comparisons and reduces that agreement. Match examples illustrate differences among winning teammates and strong performances in defeat. Player and match scores are available in My CS Events.

1 Overview

Two teammates win the same map. One has more kills; the other, more damage. Who played better? Contrib learns a shared scoring rule from team results, bringing these statistics together while letting readers examine what drives the gap between two players' scores.

2 Scoring method

Contrib evaluates completed maps rather than predicting future matches. The same scoring rule applies to every player (Figure 1).

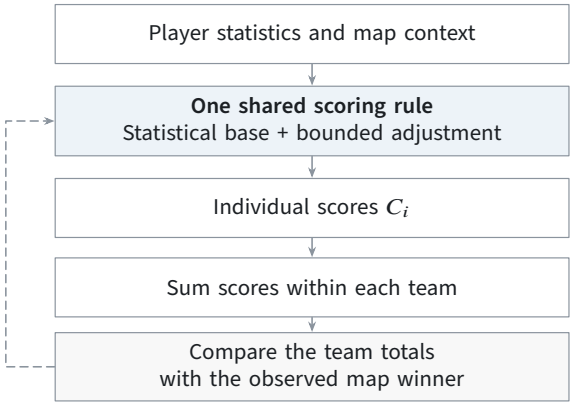


Figure 1. Scoring and learning. Solid arrows follow score construction; the dashed arrow marks feedback from the observed team result.

Score construction. The statistical base summarizes recorded performance, while a bounded adjustment captures patterns across statistics and map context. On the reported score scale, let S_i denote the statistical base and A_i the adjustment for player i. Then

$$C_i = S_i + A_i. \tag{1}$$

Learning from team outcomes. Team totals are compared with observed winners to learn the shared rule. Individual reference ratings informed earlier development but are not training targets for the reported model.

Reading the scores. Compare scores within CS2 or within CS:GO: higher scores indicate stronger performances on Contrib's scale. Zero is a reference point, not an absence of contribution. A negative score does not mean a player caused the team to lose. Scores and score gaps are not percentages or win probabilities.

Related approaches

HLTV's Rating 3.0 accounts for equipment and uses Round Swing to value changes in round win probability [1].

Existing systems connect players to team success in different ways. Round Win Shares (RWS) allocates damage and objective credit within won rounds, with zero credit for the losing side [2]. Perfect World pairs the statistical Rating Pro with Win Effect, which uses learned win probabilities to value events [3]. Leetify values changes in round win probability and shares credit for kills across damage, flash, and trade contributions [4, 5]. FACEIT likewise describes a learned model that values actions through changes in win probability [6].

In research, Xenopoulos et al. value damage events through changes in learned round win probabilities [7]. Xu and Moka estimate player-specific coefficients from lineups and results, linking differences between team sums to win probability [8].

Contrib uses a shared rule; lineup models fit player-specific coefficients (Figure 2).

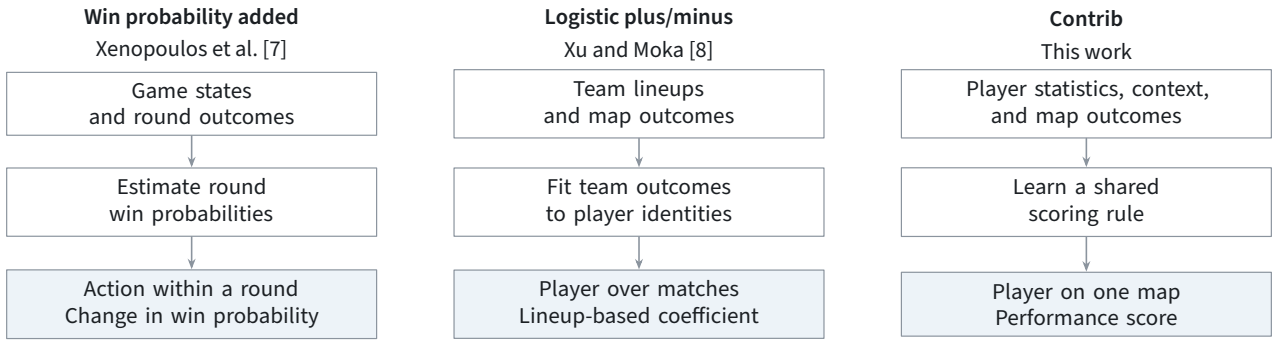


Figure 2. Method comparison. Rows show the inputs, learning task, and unit being scored.

3 Evidence and interpretation

Evaluation scope. Results are reported separately for CS:GO and CS2. The evaluation includes matches used during earlier development, and stored HLTV ratings lack a full version audit. It is therefore neither an untouched blind test nor a controlled Rating 3.0 comparison.

Reference agreement. Contrib broadly follows the reference rankings (Table 1). Within-team agreement is especially relevant for individual evaluation: all five players share the same result, so separating winners from losers cannot explain their ordering. Agreement is lower within teams than across all player–maps in both eras. Broad alignment with the reference therefore still leaves room for different judgments about which teammate performed better.

Effect of the bounded adjustment. The statistical base agrees more closely with the reference than the full score in both eras. The adjustment thus moves some comparisons away from the reference; this does not by itself establish greater accuracy.

Game era	CS2		CS:GO	
Measure	Base	Full	Base	Full
Overall ρ	0.929	0.923	0.968	0.955
Within-team ρ	0.836	0.830	0.903	0.883

Table 1. Agreement with individual reference ratings. Base retains only the statistical component of the same model; Full uses the complete score. Spearman ρ compares scores with stored HLTV ratings: overall across player–maps, or within teams of five and then averaged. Values are descriptive point estimates.

Reading the scores in matches

Three matches used during development illustrate map scores, teammate differences, and series averages.

Astana 2026. Individual rankings shift across three team wins (Figure 3).

Player	Dust2 16 : 12	Mirage 13 : 7	Ancient 13 : 10
donk	+0.45	+0.68	+1.12
tN1R	+0.21	+1.03	−0.14
sh1ro	−0.41	+0.67	+0.42
magixx	−0.17	+0.85	−0.33
zont1x	−0.04	−0.64	−0.03



Figure 3. Spirit's map scores, Astana 2026 final. Rows show players; columns show maps and team round scores. The shared color scale uses unrounded values; labels show two decimals.

Across the series, sh1ro ranks above magixx despite lower average damage per round (ADR). His advantage in kills, kill–death differential, and multikills outweighs the damage gap, while the adjustment narrows his lead (Figure 4).

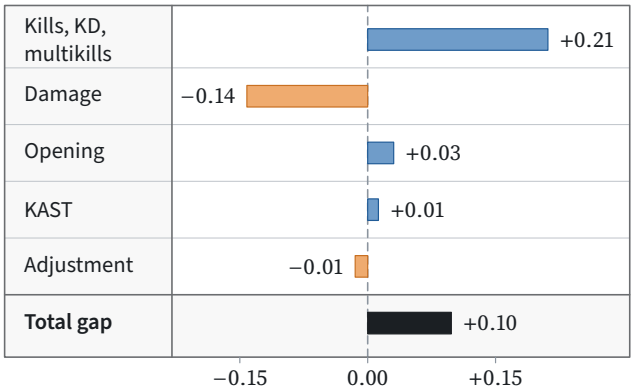


Figure 4. Astana series score decomposition. sh1ro (+0.16) minus magixx (+0.07); rightward terms favor sh1ro. All series scores are round-weighted means. KD and Opening denote kill–death and opening-kill differentials. KAST is the percentage of rounds with a kill, assist, survival, or traded death. Gray marks the total. Bars use unrounded values; rounded labels need not sum exactly.

IEM Rio 2026. KSCERATO leads all ten players (+0.96) despite FURIA's 0–2 loss. On Mirage, FURIA's team total (+1.03) exceeds Falcons' (+0.69) in an 11–13 loss. Strong aggregate performances can coexist with defeat; the sequence of rounds still matters.

PGL Major Stockholm 2021. b1t scores +0.23 on Ancient and −0.15 on Nuke. Weighting by each map's round count gives

$$\overline{C}_{\text{b1t}} = \frac{27C_{\text{Ancient}} + 41C_{\text{Nuke}}}{68} \approx 0.00.$$

The mean uses unrounded scores and hides two different map performances.

4 Discussion

The cases invite a broader discussion about what a performance score should reward. A shared rule treats identical recorded inputs alike, but players may produce those numbers under different tactical demands. This makes the opportunities behind a performance worth examining alongside its statistical output.

Roles and resources. Positions and entry duties shape those opportunities. For an AWPer, weapon access and investment could help put output in context, as Rating 3.0's economy adjustments suggest [1]. An in-game leader (IGL) raises a different issue: information and decisions may help teammates without appearing in personal statistics [9]. These are reasons to look beyond role labels, toward observations of resource use, decisions, and communication.

Timing and format. Leetify's treatment of saves and closely spaced kills illustrates why event order and later rounds matter [5, 10]. MR12 (12-round regulation halves) increases the share of pistol rounds [11], making pistol, conversion, and full-buy rounds natural settings for closer comparison.

Higher-order derived statistics. An open question is whether these features could make individual performance evaluations more stable. Capturing relationships among statistics might reduce sensitivity to fluctuations in any single measure. It would be worth examining whether this leads to more consistent assessments across maps and opponents while still reflecting real changes in performance.

Match review. These questions are best taken back to specific matches. Contrib shows the adjustment's effect on a score while leaving its internal reasoning unexplained; improving one statistic also need not raise the total. Independently reviewing players' decisions alongside the score breakdown would help assess whether a changed ranking is justified. Team outcomes alone cannot uniquely assign individual credit or establish causality. The discussion here concerns top-level professional play; how well these relationships carry over to casual games remains unvalidated.

5 Conclusion

Contrib links team outcomes to individual scores whose differences can be inspected. The results show broad agreement with existing ratings, while the cases illustrate how the scores distinguish teammates and describe strong performances in defeat.

Explore and discuss. My CS Events is an iOS app for exploring player and match scores in top-tier Counter-Strike. Feedback and suggestions are welcome and may be incorporated into future versions of Contrib. Feel free to share the report and continue the discussion, with a link back to the site.

© 2026 XuHaoran. All rights reserved.

References

[1] NER0cs. Introducing Rating 3.0. HLTV.org, August 20, 2025.

[2] FACEIT. FACEIT CS2 Advanced Stats. March 31, 2025; updated May 29, 2025.

[3] Perfect World. A potentially transformative innovation: a new CS:GO rating system arrives on the Perfect World platform. April 19, 2021. In Chinese; title translated.

[4] A. Ekman. Leetify Rating—a quantum leap for measuring CS:GO performance. Leetify, September 16, 2021.

[5] D. Hsieh. How does Leetify Rating actually work? Leetify, November 25, 2024.

[6] FACEIT. FACEIT Season 8—FACEIT Rating FAQ. April 2, 2026; updated May 26, 2026.

[7] P. Xenopoulos, H. Doraiswamy, and C. Silva. Valuing Player Actions in Counter-Strike: Global Offensive. arXiv:2011.01324, 2020.

[8] H. Xu and S. Moka. Rating Players of Counter-Strike: Global Offensive Based on Plus/Minus value. arXiv:2409.05052, 2024.

[9] NER0cs. Why are modern IGLs so aggressive? HLTV.org, June 2, 2022.

[10] J. Anderson. Leetify Rating Upgrade. Leetify, March 18, 2022.

[11] NER0cs. m0NESY's aggression, pistol rounds in MR12, and other stats from Sydney so far. HLTV.org, October 19, 2023.