DotaCounters Model

Statistical analysis of predictive efficiency

How to improve your winrate using data

DotaCounters Team

February 27, 2026

Abstract

This document presents the statistical analysis of the **DotaCounters** predictive model, a system that calculates in real-time the relative strength of hero picks in Dota 2. Our model processes raw match data from aggregators (OpenDota, Stratz) through proprietary algorithms to derive numerical counter scores and synergy scores. Data is filtered to include only matches at or above Archon rank to ensure quality. We tested the model on a dataset of **150,054 ranked matches** to quantitatively measure its statistical significance and demonstrate how players can use it to improve their winrate.

1 Introduction

Dota 2 is a complex team game where the hero pick phase has a significant impact on the final outcome. **We tested the DotaCounters model** on a dataset of **150,054 classified matches** to quantitatively measure its statistical significance.

1.1 The problem

During the pick phase, players must make rapid decisions about which hero to select. The key questions are:

1. Does this hero counter the enemy heroes?
2. Does this hero have synergy with my allies?
3. How **quantitatively** better is this choice compared to an alternative?
4. Can I play this hero **decently**? (spoiler: no, you can't)

1.2 Our solution

DotaCounters combines derived counter scores and synergy scores to calculate a strength for each possible pick.

For individual hero recommendations, the score combines:

1. **Matchup component**: counter scores against already-picked enemy heroes
2. **Synergy component**: synergy scores with already-selected allied heroes, transformed via *power law*

The formula for complete 5vs5 match analysis will be introduced in the **Match forecast model** section.

1.3 General applicability

While this document focuses on Dota 2, the underlying model structure is isomorphic to any team-based competitive scenario where individual entities (players, heroes, athletes) exhibit pairwise advantages and synergies. The same mathematical framework can be adapted to other multiplayer games, team sports, or any domain where historical performance data can inform optimal team composition and counter-strategy selection.

1.4 Notation

- Let H be our candidate hero
- Let $\mathcal{E} = \{E_1, \ldots, E_k\}$ be the set of k already-picked enemy heroes ($0 \leq k \leq 5$)
- Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ be the set of m already-selected allied heroes ($0 \leq m \leq 4$)

- Let $(H \times E)$ be the counter score of hero H against enemy hero E
- Let MS_H be the matchup sum of hero H
- Let $(H \cdot A)$ be the synergy score of hero H with ally hero A
- Let SS_H be the synergy sum of hero H

2 Definitions and $\mathcal{S}(H)$

2.1 Definitions

Counter score $(H \times E)$: Represents the statistical tendency of hero H to perform well against hero E based on historical match outcomes. The exact computation is proprietary. It satisfies the following properties:

- *Anticommutation relation*: $(H \times E) + (E \times H) = 0$ (antisymmetry)
- *Null self-score*: $(H \times H) = 0$ (no counter relationship with itself)

Matchup sum MS_H: The matchup sum of hero H is the sum of counter scores against all picked enemies:

$$MS_H = \sum_{i=1}^{k} (H \times E_i) \tag{1}$$

Synergy score $(H \cdot A)$: Represents the statistical tendency of heroes H and A to perform well when played on the same team, based on historical match outcomes. The exact computation is proprietary. It satisfies the following properties:

- *Commutation relation*: $(H \cdot A) - (A \cdot H) = 0$ (symmetry)
- *Null self-score*: $(H \cdot H) = 0$ (no synergy relationship with itself)

Synergy sum SS_H: The total synergy contribution of hero H to the team, accounting for both the benefit H receives from allies and the benefit allies receive from H. To properly filter non-linear noise, the power law transformation Φ must be applied to each pair individually before summation:

$$SS_H = 2 \cdot \sum_{j=1}^{m} \Phi(H \cdot A_j) \tag{2}$$

This definition ensures consistency with the 5vs5 match score formula, where each synergy pair is counted twice.

2.2 $\mathcal{S}(H)$

The final score for a single hero recommendation is:

$$\boxed{\mathcal{S}(H) = MS_H + \alpha \cdot SS_H}$$

where:

- $\Phi(x) = \text{sgn}(x) \cdot |x|^{\beta}$ is the *power law* transformation with $\beta = 2.5$ (signed exponentiation, optimized for high synergy discrimination)
- $\alpha = 0.1$ is the synergy scaling factor (optimized to balance amplified synergy contributions)

The optimal values for α and β were determined through extensive grid search optimization on the full dataset. A comprehensive analysis of this parameter optimization is provided in **Appendix A**.

2.3 Why the power law?

The power law transformation $\Phi(x) = \text{sgn}(x) \cdot |x|^{\beta}$ with $\beta = 2.5$ is crucial because:

- **Low values**: when $|(H \cdot A)| < 1$, the power law further reduces the contribution, filtering out statistical *noise*
- **High values**: when $|(H \cdot A)| \gg 1$, the power law exponentially amplifies the signal, capturing *decisive* synergies

Heroes are susceptible to synergies only when these assume very high values in magnitude. For intermediate values, synergies mainly add noise that worsens predictions. The power law thus acts as a **non-linear filter** that preserves strong signals and attenuates background noise.

3 Match forecast $\mathcal{S}_{\text{match}}$

In the following analysis, we consider the Radiant team as the allied team and the Dire team as the enemy team. All scores are computed from the Radiant perspective: positive scores indicate Radiant advantage, negative scores indicate Dire advantage.

3.1 Defining M, S_a and S_e

For a complete match with full teams (used in our 5vs5 benchmark analysis), let:

- Radiant $= \{R_1, R_2, R_3, R_4, R_5\}$ be the Radiant team
- Dire $= \{D_1, D_2, D_3, D_4, D_5\}$ be the Dire team

The **Matchup score sum** M for the match is the sum of individual matchup sums:

$$M = \sum_{i=1}^{5} MS_{R_i} = \sum_{i=1}^{5}\sum_{j=1}^{5}(R_i \times D_j) \qquad (3)$$

The **Ally team synergy** S_a (Radiant synergies):

$$S_a = \frac{1}{2}\sum_{i=1}^{5} SS_{R_i} = \sum_{i=1}^{5}\sum_{j=1}^{5}\Phi(R_i \cdot R_j) \qquad (4)$$

The **Enemy team synergy** S_e (Dire synergies):

$$S_e = \frac{1}{2}\sum_{i=1}^{5} SS_{D_i} = \sum_{i=1}^{5}\sum_{j=1}^{5}\Phi(D_i \cdot D_j) \qquad (5)$$

Note: The double counting is intentional: each hero in a pair receives synergy benefit from the other, consistent with the single hero formula (2) where SS_H includes a factor of 2.
The terms with $i = j$ are null due to the null self-score property.

3.2 $\mathcal{S}_{\text{match}}$

The final **match score** used in our benchmark is:

$$\boxed{\mathcal{S}_{\text{match}} = M + \alpha \cdot [S_a - S_e] \pm C}$$

where:

- $C = 17$ is the map constant (Radiant bias due to map asymmetry, side-dependent and time-variable due to game balance updates). Use $+C$ for Radiant perspective, $-C$ for Dire perspective
- $\alpha = 0.1$ is the synergy scaling factor (optimized value)
- $\beta = 2.5$ for the power law transformation (optimized value)

M is **antisymmetric** with respect to team exchange: swapping Radiant and Dire flips the sign. This follows from $(H \times E) = -(E \times H)$).

S_a and S_e are individually invariant under team exchange (synergies depend on internal team composition, not the opponent). However, their difference $[S_a - S_e]$ is antisymmetric, intuitively representing which team synergizes better.

The constant C is inherently side-dependent: $+17$ for Radiant, -17 for Dire. Its derivation will become clearer when analyzing the match score distribution in the next section.

Since $\mathcal{S}_{\text{match}}$ is a sum of antisymmetric terms, it is itself antisymmetric under team exchange.

This is expected: if team R has an advantage $\mathcal{S}_{\text{match}}$ over team D, then team D has a disadvantage $\mathcal{S}_{\text{match}}$ (or advantage $-\mathcal{S}_{\text{match}}$) over team R.

4 Statistical analysis and match score distribution

We tested the model on a database of **150,054 ranked matches**, comparing predictions with actual results.

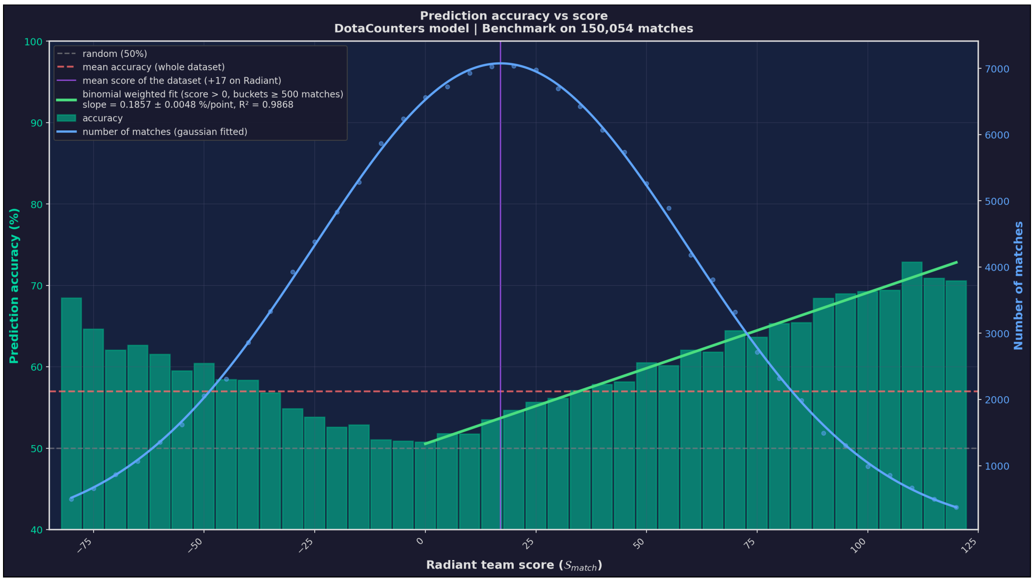

Figure 1: Relationship between Radiant team score and prediction accuracy. The green line shows the weighted linear fit. The blue curve is the Gaussian fit of the match distribution.

Figure 1 reveals the quantitative relationship between predicted score and actual accuracy. The **linear relationship** is the key feature of this analysis:

- **Slope**: 0.1857 ± 0.0048 % per point – each additional point of score increases accuracy by 0.186%
- **R²**: 0.9868 – the linear model explains 98.68% of variance, confirming the model's predictive power
- **Statistical significance**: the standard error (0.0048) confirms the slope estimate is highly precise

This linear relationship allows us to make quantitative predictions: for any given match score $|S_{\text{match}}|$, the expected accuracy is approximately:

$$\text{Accuracy}(|S_{\text{match}}|) \approx 50.0\% + 0.186\% \times |S_{\text{match}}|$$

Each score bin corresponds to a win probability that becomes **increasingly accurate** as we move away from zero. Matches near score = 0 (balanced matchups) have accuracy close to 50% (essentially random), while matches with high absolute scores ($|S_{\text{match}}| > 100$) reach 70%+ accuracy.

The blue curve shows the Gaussian distribution of matches across score ranges. The majority of matches cluster near the center (score $\approx +17$, reflecting the current map asymmetry favoring Radiant), precisely where prediction accuracy is lowest. This distribution explains why overall model accuracy appears modest: most matches fall in the challenging "balanced" regime where the model has limited discriminative power.

The +17 bias corresponds to a baseline Radiant winrate of approximately 53.2% according to our linear model ($50.0\% + 0.186\% \times 17 = 53.2\%$), which closely matches the observed Radiant winrate of 53.62% in the database (80,453 wins out of 150,054 matches). This confirms that the constant $C = 17$ effectively captures the intrinsic Radiant advantage, ensuring that the model's baseline prediction (when all other factors are neutral) yields an accuracy $\geq 50\%$.

Benchmark results

Metric	Accuracy
Global accuracy	57.00%
Accuracy for $\lvert S_{\text{match}}\rvert > 25$	60.2%
Accuracy for $\lvert S_{\text{match}}\rvert > 50$	63.8%
Accuracy for $\lvert S_{\text{match}}\rvert > 75$	67.5%
Accuracy for $\lvert S_{\text{match}}\rvert > 100$	71.3%
Accuracy for $\lvert S_{\text{match}}\rvert > 125$	74.0%
Accuracy for $\lvert S_{\text{match}}\rvert > 150$	79.3%

Conclusions

- Our model is **statistically efficient** ($R^2 = 0.9865$)
- Every 10 points in $\mathcal{S}_{\text{match}}$ improve accuracy by **1.86%**

5 Practical application: improving winrate

Our website includes a **desktop client** that analyzes ally and enemy picks in real-time during draft phases, showing the best available heroes sorted by our $\mathcal{S}(H)$ formula. It also provides a match forecast when the draft phase is over, using our $\mathcal{S}_{\text{match}}$ formula.

5.1 Calculating winrate improvement

Here's how the model translates into **quantitative advantage**:

Practical example: Last pick

Imagine you are **last pick** and need to choose between two heroes, A and B:

$\mathcal{S}(A) = -30$ (instinctive pick)
$\mathcal{S}(B) = +30$ (recommended by our model)

In case you are wondering, similar or worse scenarios are more than frequent. Players often rely on instinct, meta trends, or personal comfort picks without considering the specific counter dynamics of the current match.

The score difference is:

$$\Delta s = \mathcal{S}(B) - \mathcal{S}(A) = +60 \tag{6}$$

The winrate improvement is calculated as:

$$\Delta \text{WR} = \Delta s \times 0.186\% = 60 \times 0.186\% = +\mathbf{11.16}\% \tag{7}$$

This means that choosing hero B over hero A in this specific match context increases your win probability by over 11 percentage points.

The considerations made for the last pick become even more relevant in **Captain's Mode**, where a single user (the captain) sequentially selects allied heroes for the entire team. As information accumulates during the draft phase, the model provides increasingly pertinent recommendations: starting from raw winrate for the first pick (where minimal information about enemy strategy is available) to the optimal scenario of last pick, which guarantees the winrate delta calculated above. Each successive pick benefits from the accumulated knowledge of both teams' strategies, allowing the model to refine its suggestions and maximize the cumulative score advantage.

Of course, choosing heroes you are comfortable with remains important: a well-played hero is always preferable to a statistically strong one you cannot execute effectively. The statistical information provided by our model serves as **decision support** during the pick phase, not as a replacement for player skill and preference. Users may not always select the top recommended hero, but rather the highest-ranked among those they feel confident playing. This is precisely why the desktop client includes **whitelist and blacklist** functionality: to filter recommendations based on personal hero pools, ensuring that suggestions are both statistically sound and practically viable for each individual player.

Appendix

6 Appendix A: Parameter optimization (α and β)

The model has two key hyperparameters that require empirical optimization:

- **Alpha (α)**: Synergy scaling factor, determining the weight of synergy scores relative to matchup scores
- **Beta (β)**: Power law exponent, controlling how synergy values are transformed

6.1 Optimization methodology

We performed a comprehensive grid search optimization over the full dataset of 150,054 matches, testing 651 combinations:

- α: 0.0 to 2.0 (step 0.1) $\rightarrow$ 21 values
- β: 0.0 to 3.0 (step 0.1) $\rightarrow$ 31 values
- Total combinations: $21 \times 31 = 651$ benchmark runs
- Execution time: 16-22 minutes

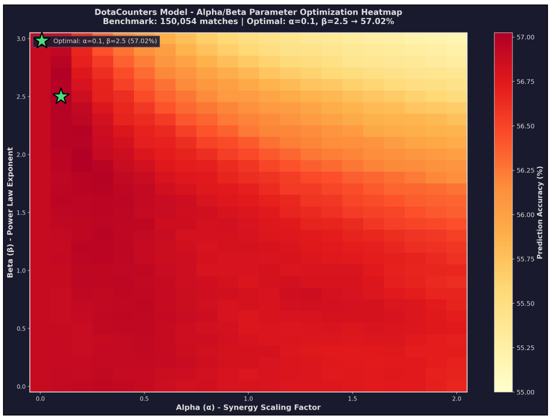

Figure 6.1: Heatmap showing prediction accuracy for all combinations of α and β. The green star marks the optimal configuration. Warmer colors (red/orange) indicate higher accuracy, cooler colors (yellow) indicate lower accuracy.

6.2 Optimization results

Figure 6.1 shows the complete optimization landscape:

- **Optimal configuration**: $\alpha = 0.1$, $\beta = 2.5 \rightarrow$ **57.02% accuracy**
- **Baseline (no synergies)**: $\alpha = 0.0 \rightarrow 56.90\%$ accuracy
- **Improvement**: +0.12% accuracy gain

6.3 Interpretation

The optimized $\beta = 2.5$ (higher than a simple square) aggressively amplifies high synergy values while suppressing low values, providing effective noise filtering. The lower $\alpha = 0.1$ compensates for this stronger transformation, preventing synergy scores from dominating the matchup component.

Critical observation: The naive configuration $\alpha = 1.0$, $\beta = 1.0$ (no power law, no scaling) yields only 56.81% accuracy, *worse than the baseline* (56.90% with no synergies). This demonstrates that raw synergy scores are comparable to noise and must be filtered through the power law transformation. Without proper filtering ($\beta < 2$), synergies actually *degrade* prediction quality rather than improve it.

The +0.12% improvement may appear modest at first glance, but it must be considered in context: this gain is measured on the **overall accuracy**, which is inherently pulled down by the fact that the majority of matches fall within the "balanced" regime at the center of the Gaussian distribution, where predictions naturally lose accuracy. Improving the baseline accuracy across this challenging distribution is therefore more significant than the raw percentage suggests.