From Outcome Prediction to Advantage Construction: A Temporal and Spatial Attribution Framework for *League of Legends*

Henrique de Macedo Airoso da Silva[1] **, Yuri Kaszubowski Lopes**[1]

[1]PPGCAP – Universidade do Estado de Santa Catarina (UDESC),
Joinville – SC – Brazil

henrique.silva21@edu.udesc.br, yuri.lopes@udesc.br

Abstract. Introduction: *Prior League of Legends (LoL) studies often emphasize winner prediction from aggregate indicators. This paper instead treats a match as a temporal process in which advantage is built, stabilized, or contested.* ***Objective:*** *The work asks whether a LoL match can be represented as a causal minute-level sequence and analyzed through spatial attribution, temporal state estimation, and short-horizon forecasting.* ***Methodology:*** *Match and timeline data are reconstructed into minute-level states using only information observable up to each minute. The representation combines composition, observed state, lane pressure, objectives, itemization, temporal variation, and gravity-informed descriptors. Gold is treated as an emergent macroindicator rather than as the explanatory center.* ***Results:*** *The state-estimation model reached 0.821 ROC AUC at 15 minutes and 0.872 at 20 minutes. Forecasting improved from 0.536 weighted F1 at a 1-minute horizon to 0.678 at 5 minutes. Spatial artifacts showed outcome-aligned regional control in intermediate windows, while lock-in remained exploratory rather than a validated detector.*

Keywords: League of Legends, temporal modeling, game analytics, competitive advantage, short-horizon forecasting.

1. Introduction

League of Legends (LoL) is a Multiplayer Online Battle Arena (MOBA) game. In this genre, two teams compete on a shared map and try to destroy the opposing base. Each player controls one character, called a champion, with a role in combat, farming, vision, movement, and objective control. A match is not decided by a single action. It develops through repeated local decisions, such as where players move, which areas they control, when they fight, and which objectives they secure.

For this reason, predicting only the final winner gives an incomplete view of the match. A model may correctly predict that one team wins, but still fail to explain how the advantage was created. In practical analysis, the important question is often not only "who wins?" but also "when did the match become favorable?", "where was the advantage located?", and "which short-term events could still change the state?"

Outcome prediction in competitive games is a relevant topic at the intersection of artificial intelligence, game analytics, and player modeling. LoL is a recurrent testbed in MOBA-oriented artificial intelligence research [Costa et al. 2024]. Its public match-history and timeline data support empirical reconstruction of in-game states [Riot Games, Inc. 2024], while its role-based team structure and sequential decision process make it suitable for modeling player contribution, pressure transfer, and temporal advantage formation [Bahrololloomi et al. 2022]. In this paper, a match is treated as a sequence of competitive states rather than as a final label alone.

The proposed framework reads the same match sequence at three levels. First, a spatial attribution layer describes where map control is concentrated and which champions contribute to local influence. Second, a temporal state-estimation model estimates the final win probability from the observable history at each minute. Third, a short-horizon forecasting model predicts plausible local events in the next few minutes. Together, these components describe how advantage becomes visible over time.

Gold is handled carefully. In LoL, gold is an in-game resource used to buy items and strengthen champions. It is often predictive [Hitar-García et al. 2022], but it is also a consequence of many mechanisms, such as farming, takedowns, towers, neutral objectives, and map pressure. Snowball and comeback-oriented studies support the interpretation that economic advantage interacts with broader match dynamics [Jung e Kim 2022], [Lee e Kim 2025]. The present framework therefore treats gold as an aggregate indicator, not as the explanatory center.

2. Related Work

Prior work on LoL includes pre-game prediction, in-game state prediction, temporal modeling, and explainable analytics [Costa et al. 2024]. Pre-game studies estimate outcomes from champion selection or historical context [Costa et al. 2021], [Cong et al. 2024]. In-game predictors often rely on compact indicators such as gold, kills, objectives, and player-level statistics [Do et al. 2021], [Chowdhury et al. 2025]. Temporal approaches model match evolution rather than static snapshots [Silva et al. 2018]. Explainable and KPI-oriented studies investigate coaching indicators, player behavior, vision, map movement, and contribution analysis [Hojaji et al. 2025]. Related MOBA work also shows the tension between predictive performance and interpretability in complex team games [Yang et al. 2020].

A related formulation uses Percent Elapsed Time (PET), where the model predicts after a fixed fraction of the match has passed [Junior e Campelo 2023]. This is useful when the question is how predictable the outcome becomes after a given percentage of the match. The present work asks a different question: what is the competitive state at absolute minute k, using only information available up to that minute? Absolute time is important because LoL objectives, item timings, and map pressure are tied to concrete minutes, not only to match duration percentages.

The gap addressed here is not the absence of winner prediction. The gap is that many formulations do not explicitly separate three analytical needs: locating spatial advantage, estimating the current global state, and forecasting near-future local events. This paper builds a framework around that separation.

3. Objective

The objective is to formulate LoL match analysis as a temporal and structurally interpretable problem. Instead of treating a match only as a final win or loss, the work models it as a sequence of minute-level states. At each minute, the framework asks what is currently observable, where the advantage is concentrated, how likely the final outcome is, and which events may happen next.

The central research question is:

How can a LoL match be represented as a causal temporal sequence in order to explain how competitive advantage is built, stabilized, or contested over time?

The working hypothesis is that a three-level representation gives a richer account of match dynamics than a single outcome classifier centered on aggregate indicators: spatial attribution identifies where map pressure is concentrated and which champions contribute to local control; temporal state estimation estimates final win probability from the observable history at each minute; and short-horizon forecasting predicts local developments under the same causal boundary. The paper is organized around four research questions: how to represent minute-level causal states, which signals describe advantage beyond aggregate gold, how state estimation and forecasting complement each other, and how spatial attribution makes map pressure interpretable through champion-level contribution and regional control.

4. Methodology

4.1. Data reconstruction and causal contract

The dataset is built from Challenger-ranked matches through the LoL Developer's (Riot Games) Application Programming Interface (API) [Riot Games, Inc. 2024], storing match details and timeline data. Challenger matches are used because high-ranked games are expected to contain more coherent decision-making and more legible pressure exchanges [Bahrololloomi et al. 2022]. Preprocessing reconstructs each match as an ordered sequence of minute-level states. Observations are grouped by match, ordered chronologically, and transformed into temporal matrices. Padding and masking are applied only at the batch level.

The central rule is causal: at minute k, the models can only use information observed up to minute k. No future-minute information can be used as an input. This rule is important because the goal is not to explain the match after it is over, but to estimate what could have been known at each moment.

The feature space is organized into interpretable groups: composition features describe champion identities, roles, and team structure; observed-state features describe gold, experience, combat, farming, vision, position, objectives, structures, and itemization; temporal features encode minute-to-minute changes; and spatial descriptors summarize pressure, regional dominance, and champion contribution. The groups are interpreted as observable mechanisms rather than isolated causal explanations. Pressure denotes local competitive influence over a map region, derived from signals indicating whether one team can contest or control it more effectively.

4.2. Causal task formulation

The notation below formalizes the two supervised tasks under the same causal sequence constraint. The state-estimation task follows conditional probability estimation in supervised learning, while the forecasting task follows sequence-to-label and multi-label temporal modeling principles [Pedregosa et al. 2011], [Hochreiter e Schmidhuber 1997], [Vaswani et al. 2017], [Vardakis et al. 2026].

Let $\mathcal{M}$ denote a set of matches. Each match $m \in \mathcal{M}$ is observed as a sequence with total length T_m. At each minute k, with $1 \leq k \leq T_m$, the match state is represented by a feature vector $\mathbf{x}_k^{(m)} \in \mathbb{R}^d$. Let $y^{(m)} \in \{0, 1\}$ indicate whether the blue side wins match m.

The state-estimation task estimates the final blue-side win probability $\hat{p}_k^{(m)}$ at minute k, conditioned only on the observable history:

$$f_\theta^{(\text{state})} : \mathbf{x}_{1:k}^{(m)} \mapsto \hat{p}_k^{(m)} = \mathbb{P}\left(y^{(m)} = 1 \mid \mathbf{x}_{1:k}^{(m)}\right). \tag{1}$$

Here, $\mathbf{x}_{1:k}^{(m)} = \left(\mathbf{x}_1^{(m)}, \ldots, \mathbf{x}_k^{(m)}\right)$ is the history available from the beginning of the match up to minute k. The value $\hat{p}_k^{(m)} \in [0, 1]$ is the estimated probability that the blue side eventually wins. Thus, $1 - \hat{p}_k^{(m)}$ is the estimated probability that the red side wins. The index k denotes the current minute and the length of the observed prefix.

The forecasting task predicts local developments over short horizons:

$$f_\phi^{(\text{forecast})} : \mathbf{x}_{1:k}^{(m)} \mapsto \hat{\mathbf{z}}_{k+h|k}^{(m)}, \qquad h \in \{1, 3, 5\}. \tag{2}$$

The output $\hat{\mathbf{z}}_{k+h|k}^{(m)} \in [0, 1]^C$ is a vector of predicted probabilities. Each component estimates the probability that one event type occurs within the future interval $(k, k + h]$, such as combat, tower damage, a neutral objective, lane-pressure reversal, or spatial-pressure shift.

For state estimation, each prefix of match m defines one supervised example. This construction is analogous to treating each observed prefix as a temporally valid supervised instance, while preserving match-level grouping to avoid leakage across splits [Birant e Birant 2022], [Silva et al. 2018]:

$$e_{B,k}^{(m)} = \left(X_{B,1:k}^{(m)}, y^{(m)}\right), \qquad X_{B,1:k}^{(m)} \in \mathbb{R}^{k \times d_B}, \qquad y^{(m)} \in \{0, 1\}. \tag{3}$$

The matrix $X_{B,1:k}^{(m)}$ contains k observed time steps, each represented by d_B features. The target $y^{(m)}$ is constant across all prefixes of the same match because the model asks, at each minute, how likely the final outcome is from the information observed so far.

The state-estimation dataset is:

$$\mathcal{D}_B = \left\{\left(X_{B,1:k}^{(m)}, y^{(m)}\right) : m \in \mathcal{M},\ 1 \leq k \leq T_m\right\}. \tag{4}$$

For forecasting, each valid minute and horizon pair defines one multi-label example, following the broader formulation of event prediction from in-game temporal data [Vardakis et al. 2026]:

$$e_{C,k,h}^{(m)} = \left(X_{C,1:k}^{(m)}, \mathbf{z}_{k,h}^{(m)}\right), \qquad X_{C,1:k}^{(m)} \in \mathbb{R}^{k \times d_C}. \tag{5}$$

The observed target vector is:

$$\mathbf{z}_{k,h}^{(m)} = \left(z_{k,h,1}^{(m)}, \ldots, z_{k,h,C}^{(m)}\right) \in \{0, 1\}^C. \tag{6}$$

Each component $z^{(m)}_{k,h,c}$ indicates whether event type c occurs within the interval $(k, k+h]$. Therefore, $\mathbf{z}^{(m)}_{k,h}$ is binary and observed, while $\hat{\mathbf{z}}^{(m)}_{k+h|k}$ is probabilistic and predicted.

The forecasting dataset is:

$$\mathcal{D}_C = \left\{ \left(X^{(m)}_{C,1:k}, \mathbf{z}^{(m)}_{k,h} \right) : m \in \mathcal{M},\ h \in \{1,3,5\},\ 1 \leq k \leq T_m - h \right\}. \tag{7}$$

The condition $1 \leq k \leq T_m - h$ ensures that the future interval exists inside the observed match.

4.3. Spatial attribution layer

The spatial attribution layer is gravity-informed, but it is not a physics model: gravity is an analogy for spatial influence and structural concentration. Advantage construction is defined as the accumulation, displacement, or stabilization of observable signals that make one team better able to control regions, contest objectives, or restrict the opponent's options. A champion with strong local presence is treated as a source of influence; nearby regions accumulate pressure from allied and enemy champions. The resulting field indicates which side controls each region and which champion contributes most to that control. This formulation is conceptually related to potential-field methods, multi-agent influence fields in games, and spatio-temporal control analysis in team sports [Khatib 1986], [Hagelbäck e Johansson 2009], [Gudmundsson e Horton 2017], [Spearman 2018].

Let i denote a champion at time k. Structural mass $m_i(k)$ is a weighted transformation of observable signals. It is defined here as an operational representation of player-state intensity rather than as a physical quantity:

$$m_i(k) = \sum_{r=1}^{R} \alpha_r\, \psi_r(x_{i,r}(k))\,. \tag{8}$$

The index r identifies one signal, such as level, farm, damage, vision, or itemization. Thus, $x_{i,r}(k)$ is the value of the r-th signal for champion i, $\psi_r(\cdot)$ is its transformation, and $\alpha_r \geq 0$ is its weight.

Mass is propagated over spatial and relational interactions, following the intuition that competitive influence depends not only on an isolated unit but also on nearby allies, enemies, and relational context [Battaglia et al. 2018], [Duan et al. 2022]:

$$\mathrm{infl}_i(k) = \sum_{j \in \mathcal{N}(i)} \omega_{ij}(k)\, m_j(k), \tag{9}$$

where $\mathcal{N}(i)$ denotes neighboring champions and $\omega_{ij}(k)$ is a context-dependent interaction weight. This influence is separated into allied support, adversarial pressure, and net contextual influence.

A champion's attractor potential combines local mass, propagated influence, local pressure, and objective context. This equation defines a task-specific potential score, inspired by potential-field reasoning but grounded in observable match variables [Khatib 1986], [Hagelbäck e Johansson 2009]:

$$\phi_i(k) = \beta_1 m_i(k) + \beta_2 \mathrm{infl}_i(k) + \beta_3 p_i(k) + \beta_4 o_i(k). \tag{10}$$

Here, $p_i(k)$ denotes team-aligned local pressure terms and $o_i(k)$ denotes objective context. A high value of $\phi_i(k)$ means that the champion is an important local source of control.

At the regional level, let $M_R^{blue}(k)$ and $M_R^{red}(k)$ be the aggregated mass assigned to region R. The signed regional density follows the logic of spatial dominance and regional control measures used in spatio-temporal team analysis [Gudmundsson e Horton 2017], [Spearman 2018]:

$$D_R^{\Delta}(k) = \left(M_R^{blue}(k) - M_R^{red}(k)\right)\left(1 + \gamma_2 C_R(k) + \gamma_3 O_R(k)\right). \tag{11}$$

Positive values indicate blue-side dominance and negative values indicate red-side dominance.

The local field aggregates champion-level potential over map coordinates through spatial kernels, which makes local control attributable to individual units [Khatib 1986], [Hagelbäck e Johansson 2009], [Gudmundsson e Horton 2017]:

$$F_{\Delta}(x, y, k) = \sum_{i \in blue} \phi_i(k) K_i(x, y) - \sum_{j \in red} \phi_j(k) K_j(x, y), \tag{12}$$

where $K_i(x, y)$ is a spatial kernel centered at champion i's position. Each map cell is assigned to the champion with the largest local contribution. This makes pressure attributable: the framework identifies not only which side controls a region, but also which champion contributes most to that local control.

Finally, the lock-in score summarizes structural concentration. It is treated as a descriptive stability diagnostic rather than as a validated irreversible-state detector, consistent with the need to evaluate comeback and snowball dynamics empirically [Jung e Kim 2022], [Lee e Kim 2025]:

$$L(k) = \lambda_1 A(k) + \lambda_2 D(k) + \lambda_3 S(k) + \lambda_4 \Delta c(k) + \lambda_5 B(k). \tag{13}$$

Here, $A(k)$ is global regional asymmetry, $D(k)$ is aggregate impact density, $S(k)$ is regional stability, $\Delta c(k)$ is separation between team centers of mass, and $B(k)$ summarizes dominant-attractor gaps. In this paper, $L(k)$ is exploratory. It is not treated as a validated detector of irreversible match states.

4.4. Temporal models and evaluation protocol

The state-estimation model uses causal temporal encoders selected among multiscale Temporal Convolutional Networks with attention, grouped transformers, and hybrid TCN-transformer backbones, following sequence-modeling principles from recurrent and attention-based temporal architectures [Hochreiter e Schmidhuber 1997], [Vaswani et al. 2017]. In the evaluated run, the selected backbone was an MSTCN-attention variant. Splits are performed by match, scalers are fitted only on the training set, and padding is handled through masks. The time-weighted loss is an empirical risk objective with temporal weights and valid-step masking:

$$\mathcal{L}_{\text{state}}(\theta) = \frac{\sum_m \sum_{k=1}^{T_m} \lambda_k\, w_k^{(m)}\, \ell\left(\hat{p}_k^{(m)}, y^{(m)}\right)}{\sum_m \sum_{k=1}^{T_m} \lambda_k\, w_k^{(m)}}. \tag{14}$$

Here, λ_k is the temporal weight, $w_k^{(m)}$ is the valid-step mask, and $\ell(\cdot)$ is the binary loss. Temperature scaling is applied after training to improve calibration [Guo et al. 2017]. The evaluated run used temperature scaling with a temperature of 0.751.

The short-horizon forecasting model uses the same causal sequence discipline, but replaces the binary state-estimation head with a multi-label head, following event-prediction formulations for MOBA and in-game temporal data [Vardakis et al. 2026]. In the evaluated run, the forecast task contained 45 binary targets across three horizons. Targets are defined through short-horizon columns, degenerate targets are pruned, and training uses masked multi-label Binary Cross-Entropy with Logits. The final training run used balanced target weighting and executed 14 epochs.

The evaluation does not claim state-of-the-art superiority over external single-state classifiers. This is a limitation of the present study. Instead, the experiment evaluates whether causal reconstruction, forecasting, and spatial attribution answer the research question more directly than a winner-only classifier. For state estimation, evaluation uses ROC AUC, Average Precision, Brier Score, ECE, risk sets, and trajectory stability. For forecasting, metrics are reported by horizon and event type. For spatial attribution, evaluation is descriptive and diagnostic. A preliminary gravity-GNN diagnostic is also reported, but only as supporting evidence; it is not treated as the main validation of the spatial layer.

Metric interpretation is task-dependent. ROC AUC measures ranking quality; Average Precision is useful under imbalance; Brier Score and ECE evaluate probabilistic calibration; and Weighted F1/recall summarize thresholded multi-label behavior, but must be read with prevalence because frequent and rare events produce different regimes.

Metrics at minute k are computed only over the temporal risk set:

$$\mathcal{R}_k = \{m \in \mathcal{M} : T_m \geq k\}, \tag{15}$$

where $\mathcal{R}_k$ contains matches still active at minute k. This prevents late-window evaluation from assigning late-game observations to matches that ended earlier.

Temporal volatility is:

$$\text{Volatility}^{(m)} = \frac{1}{T_m - 1} \sum_{k=2}^{T_m} \left| \hat{p}_k^{(m)} - \hat{p}_{k-1}^{(m)} \right|. \tag{16}$$

The analysis also considers boundary flips, balanced-zone time, and stabilization minute.

5. Results

5.1. State-estimation results

The state-estimation evaluation uses 450 test matches. Table 1 reports performance at different minutes. The number of valid matches decreases after 20 minutes because some games end earlier.

Table 1. State-estimation performance and temporal risk sets.

Window	$\lvert\mathcal{R}_k\rvert$	ROC AUC	AP	Brier	ECE	Accuracy
10 min	450	0.710	0.693	0.218	0.035	0.651
15 min	450	0.821	0.821	0.173	0.044	0.744
20 min	426	0.872	0.872	0.149	0.074	0.784
25 min	372	0.886	0.878	0.136	0.042	0.812
Final	450	0.996	0.997	0.031	0.096	0.976

In practical terms, the model is uncertain in the earliest windows because both teams still have many possible paths. By 10 minutes, it begins to separate favorable and unfavorable states. At 15 minutes, it reaches 0.821 ROC AUC and 0.821 AP, meaning that eventual-winning prefixes are ranked above eventual-losing prefixes with substantially better-than-random discrimination. At 20 minutes, discrimination increases to 0.872 ROC AUC and 0.872 AP over 426 still-active matches. The final window is almost deterministic and should not be treated as the main evidence of early usefulness.

Trajectory diagnostics reinforce this interpretation. The model's main value is not one isolated prediction, but the probability path that shows how the match moves from uncertainty toward a more stable state. In one inspected real match trajectory, the model reached a temporal stability score of 0.960, with mean minute-to-minute probability change of 0.032 and one decision-boundary flip. Across the evaluated timeline summary, the mean ROC AUC was 0.801, mean Average Precision was 0.790, mean Brier Score was 0.174, and mean ECE was 0.095.

5.2. Short-horizon forecasting results

The forecast model answers a different question from the state model. It does not ask who will win the match. It asks which local events are likely to happen soon. Table 2 shows that the 3-minute and 5-minute horizons are easier than the 1-minute horizon. This does not mean that longer horizons are always better; it means that events have more time to materialize, increasing signal and prevalence. Thus, forecasting performance must be interpreted jointly with the horizon and the event family.

Table 2. Short-horizon forecasting performance.

Horizon	Targets	Weighted F1	Mean AP	Mean ROC AUC	Recall
1 min	15	0.536	0.521	0.847	0.687
3 min	15	0.648	0.660	0.852	0.746
5 min	15	0.678	0.687	0.828	0.769

In Table 3, *Targets* denotes how many binary event labels are included in each family. Each label is a yes/no question about whether a specific type of event occurs within a future window. It gives a clearer interpretation for readers who do not know LoL-specific target names. Combat and structure-related events were easier because they are more frequent and more strongly connected to the current match state. Objective-related events also performed well. Lane-pressure reversals were difficult because they are rare and often depend on short tactical choices. Their mean prevalence was only 0.033, which explains the weak mean AP of 0.067 and mean F1 of 0.120. Future work should test imbalance-aware losses, such as focal loss [Lin et al. 2017], and grouped heads by event family.

Table 3. Forecasting performance by event family.

Family	Targets	Prevalence	Mean AP	Mean F1
Combat	3	0.913	0.946	0.960
Objectives	24	0.198	0.763	0.730
Structure	3	0.464	0.817	0.823
Spatial shift	6	0.476	0.635	0.668
Lane pressure	9	0.033	0.067	0.120

5.3. Spatial attribution results

The spatial attribution layer is evaluated as an exploratory structural layer rather than as an independent predictive model. Here, map pressure denotes local competitive influence over a region. The artifacts make this influence attributable through field concentration, regional ownership, champion-level contribution, and local basin size. The main variables include `field_share`, `dominant_local_cells`, `field_peak`, `signed_regional_impact_density`, `enemy_pressure_received`, and `center_distance_norm`.

The evaluated analytical artifacts used the current build manifest and a smoke rendering profile, ensuring that the visual evaluation consumed the artifacts generated by the same run. For the inspected match, the highest mean and maximum regional density occurred in a jungle region. The dominant regions changed by segment: a river-adjacent region was dominant early, while a jungle region became dominant in mid and late segments. This supports the interpretation that spatial pressure is not static; it moves across the map as objectives, rotations, and champion positions change.

The attribution tables covered 39 matches, 1080 minute-level states, and 10800 champion-minute rows. Internal consistency checks showed that `field_share` summed to one per minute, local cell ownership covered the evaluated grid, and signed regional density matched regional mass difference. The side aligned with the eventual winner owned the dominant field in approximately 54% of states before 10 minutes, 65% between 10 and 14 minutes, 61% between 15 and 19 minutes, and 58% after 20 minutes. Signed mass alignment followed a similar pattern, reaching approximately 66% between 10 and 14 minutes.

These values are not standalone prediction results. They indicate that spatial organization becomes more aligned with the final outcome after early-game noise decreases. In less game-specific terms, the spatial layer shows whether the eventual winner tends to control more influential regions as the match develops. It also identifies which champion acts as a local source of control, without requiring the reader to know the champion's name or game lore. For this reason, champion-specific examples are interpreted by function, such as mobile engage champions, frontline champions, or durable zone-control champions, rather than by name.

The lock-in score also remained conservative. In the inspected analytical evaluation, the raw peak occurred at minute 0, but the eligible competitive peak occurred at minute 21 with score 0.330 after applying the eligibility window from minute 10 onward. No lock-in candidate satisfied the stronger decision rule. This supports the methodological decision to treat `lock_in_score` as a continuous structural concentration diagnostic, not as a binary detector of irreversible match states.

A preliminary supervised gravity-GNN diagnostic gives additional evidence for caution. This diagnostic follows the general motivation of graph-based relational learning, but is used here only as supporting evidence rather than as the main validation layer [Bisberg e Ferrara 2022], [Battaglia et al. 2018]. On the test split, the GNN reached event micro-F1 of 0.744 and objective micro-F1 of 0.556, but region accuracy was only 0.200. The win head also remained weak, with ROC AUC of 0.487 and Average Precision of 0.660. These results support the paper's interpretation: the spatial layer is useful for attribution and diagnostics, but it is not yet a validated standalone prediction model.

5.4. Interpretation of the hypothesis

The results support the working hypothesis. The state-estimation model provides a global probability trajectory: it shows how the match moves from uncertainty toward a more stable competitive state. The forecasting model provides local short-term expectations: it shows which kinds of events are likely to happen soon. The spatial attribution layer provides map-level explanation: it shows where control is concentrated and which champions contribute to it.

Together, the three levels describe advantage construction rather than only outcome classification. The framework does not merely say who is likely to win. It also describes when the advantage becomes visible, where it is located, and which local events may reinforce or weaken it.

6. Conclusion

This paper presented a temporal and structurally interpretable formulation of *League of Legends* match analysis. Instead of centering representation on aggregate macrovariables alone, the framework treats advantage as an emergent configuration shaped by farming, combat pressure, objective control, itemization, tempo, and spatial influence. This supports state estimation, short-horizon forecasting, and spatial attribution under a shared causal reconstruction discipline.

The empirical results support this framing in three ways. First, the state-estimation model produced meaningful probability trajectories across absolute-time windows, with clear improvement from 10 to 20 minutes. Second, the forecasting model performed better on broad combat, objective, structure, and spatial-shift events than on rare lane-pressure reversals. Third, the spatial attribution layer made map pressure interpretable through champion-level contribution, signed regional dominance, and local basins of influence. Its lock-oriented outputs remained exploratory and should not yet be treated as a validated decision rule.

Taken together, these results support modeling a LoL match as a temporal process in which structural advantage is built, stabilized, and sometimes redistributed. The main limitations are the absence of direct comparison with competitive external baselines, the exploratory status of the spatial layer, and the difficulty of rare-event forecasting. Future work should add stronger baselines and ablations, improve rare-event forecasting through imbalance-aware objectives, and validate lock-in through comeback-oriented evaluation.

References

Bahrololloomi, F., Smit, M., Ghotbi, E., van der Vegt, W., e van Deursen, A. (2022). Machine learning analysis of player performance for predicting victory in league of legends. In *Proceedings of the 14th International Conference on Computer Supported Education (CSEDU 2022)*, volume 1, pages 63–74.

Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gulcehre, C., Song, F., Ballard, A., Gilmer, J., Dahl, G., Vaswani, A., Allen, K. R., Nash, C., Langston, W., Dyer, C., Heess, N., Wierstra, D., Kohli, P., Botvinick, M., Vinyals, O., Li, Y., e Pascanu, R. (2018). Relational inductive biases, deep learning, and graph networks. *arXiv preprint*.

Birant, K. e Birant, D. (2022). Multi-objective multi-instance learning: A new approach to machine learning for esports. *Entropy*, 24(1):123.

Bisberg, A. J. e Ferrara, E. (2022). Gcn-wp–semi-supervised graph convolutional networks for win prediction in esports. In *2022 IEEE Conference on Games (CoG)*, pages 449–456. IEEE.

Chowdhury, S., Ahsan, M., e Barraclough, P. (2025). Applications of linear and ensemble-based machine learning for predicting winning teams in league of legends. *Applied Sciences*, 15(10).

Cong, S., Wu, H., Xiong, S., Wu, Y., e Zuo, L. (2024). Predicting the outcome of moba game: a case study on pregame prediction for league of legends. In *2024 8th Asian Conference on Artificial Intelligence Technology (ACAIT)*, pages 1366–1372.

Costa, L. M., Drachen, A., Souza, F. C. M., e Xexéo, G. (2024). Artificial intelligence in moba games: A multivocal literature mapping. *IEEE Transactions on Games*, 16(1):5–18.

Costa, L. M., Mantovani, R. G., Souza, F. C. M., e Xex'eo, G. (2021). Feature analysis to league of legends victory prediction on the picks and bans phase. In *2021 IEEE Conference on Games (CoG)*, pages 1–5.

Do, T. D., Wang, S. I., Yu, D. S., McMillian, M. G., e McMahan, R. P. (2021). Using machine learning to predict game outcomes based on player-champion experience in league of legends. *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, pages 28–34.

Duan, L., Li, S., Zhang, W., e Wang, W. (2022). Moba game item recommendation via relation-aware graph attention network. In *2022 IEEE Conference on Games (CoG)*, pages 338–344.

Gudmundsson, J. e Horton, M. (2017). Spatio-temporal analysis of team sports – a survey. *ACM Computing Surveys*, 50(2).

Guo, C., Pleiss, G., Sun, Y., e Weinberger, K. Q. (2017). On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330.

Hagelbäck, J. e Johansson, S. (2009). Using multi-agent potential fields in real-time strategy games. In *2009 IEEE Symposium on Computational Intelligence and Games (CIG)*, pages 191–198.

Hitar-García, J. A., Valls-Vargas, J., Cuesta-Vargas, F. J., e Alba, E. (2022). Machine learning methods for predicting league of legends game outcome. In *2022 IEEE Conference on Games (CoG)*, pages 406–413.

Hochreiter, S. e Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.

Hojaji, F., McIlroy, R. E., Dupuy, A., Pedroni, G., Toth, A. J., e Campbell, M. J. (2025). Deep learning techniques for identifying kpis in league of legends: Win prediction, map navigation, and vision control. *Computers in Human Behavior Reports*, 19.

Jung, C. e Kim, H. (2022). Win prediction from the snowball effect perspectives. In *2022 IEEE Games, Entertainment, Media Conference (GEM)*.

Junior, J. B. S. e Campelo, C. E. C. (2023). League of legends: Real-time result prediction.

Khatib, O. (1986). Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1):90–98.

Lee, J. e Kim, N. (2025). Development of machine learning-based indicators for predicting comeback victories using the bounty mechanism in moba games. *Electronics*, 14(7).

Lin, T.-Y., Goyal, P., Girshick, R., He, K., e Dollár, P. (2017). Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., e Duchesnay, E. (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830.

Riot Games, Inc. (2024). Riot games api. `https://developer.riotgames.com/`. Acessado em: 02 de julho de 2024.

Silva, A. L. C., Pappa, G. L., e Chaimowicz, L. (2018). Continuous outcome prediction of league of legends competitive matches using recurrent neural networks. In *Proceedings of SBGames 2018*. Computing Track – Short Papers.

Spearman, W. (2018). Beyond expected goals. In *Proceedings of the 12th MIT Sloan Sports Analytics Conference*, Boston.

Vardakis, M., Margetis, G., Chatzakis, I., Apostolakis, K. C., e Stephanidis, C. (2026). Prediction of moba game events based on in-game data. *Entertainment Computing*, 57:101091.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., e Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.

Yang, Z., Pan, Z., Wang, Y., Cai, D., Liu, X., Shi, S., e Huang, S.-L. (2020). Interpretable real-time win prediction for honor of kings, a popular mobile moba esport. *arXiv preprint*.