Quantifying the Drivers of Serve Effectiveness in Men's Tennis

Harry Wieand[1], Yijia Chen[2], Jason Zhou[3], and Nicole Lin[4]

[1]Boston University Academy
[2]Nanjing Foreign Language School
[3]Wilbraham Monson Academy
[4]South Hills Academy

October 31, 2025

Abstract

Tennis is one of the few sports in which the server initiates every point, making the serve the only stroke entirely under a player's control. A powerful, well-placed serve can create immediate advantages, but identifying which characteristics most effectively translate serves into winning points remains unclear. Using point-by-point data from the U.S. Open men's singles (2018—2019; 2021—2024), we evaluate how first-serve speed, accuracy, and spatial variation (serve-location entropy) shape point outcomes. Regression results show that speed has the strongest relationship with serve efficiency, while accuracy in the service box and variation in serve location have smaller but still statistically significant associations. In out-of-sample tests with an 80/20 train–test split, our model achieves a root mean squared error (RMSE) of 0.066 (vs. a SD of 0.084 for observed outcomes) and a correlation of $r = 0.39$ between predicted and actual outcomes. Overall, serve speed has the strongest relationship with serve efficiency, followed by accuracy and unpredictability.

1 Introduction

The serve is the only stroke in tennis that a player executes entirely under their own control, initiating every point and setting the tone for each rally. An effective serve combines three essential elements: velocity (which compresses the returner's reaction window and ability to control their return), accuracy (ensuring the ball lands within the service box), and spatial variation (unpredictability of serve direction). Together, these components determine whether the server gains an immediate advantage or merely initiates a neutral exchange.

Aggregate performance statistics consistently underscore the centrality of the serve in match outcomes. Yet while this overall link between "good serving" and winning is well established, the **relative** contribution of its underlying components – speed, accuracy, and placement variability – remains less clear. In particular, it is uncertain whether unpredictability in serve location offers meaningful benefits once speed and accuracy are accounted for, and how these elements interact.

This study addresses that gap using detailed point-by-point data from U.S. Open men's singles (2018–2019; 2021–2024). We introduce an outcome metric, *serve efficiency*, defined as the proportion of first-serve points won within **the first three shots of the rally** (i.e., serve, return, and the server's next shot). We also formalize *serve-location entropy* as a measure of unpredictability in serve placement. Leveraging regression with standardized predictors, we estimate the associations of serve speed, first-serve-in percentage, and location entropy with serve efficiency. We then assess out-of-sample predictive performance via an 80/20 train–test split. Together, our results quantify the tactical balance between power, control, and variation in professional men's tennis.

2 Literature Review

2.1 Serve Location

Prior research on serve placement has examined spatial patterns and the biomechanics of concealment. Using Hawk-Eye tracking, Kolbinger and Lames (2013) documented distinct lateral and vertical serve distributions in elite men's singles and showed that first serves vary more spatially than second serves. Additionally, Reid et al. (2011) found that professionals adjust toss direction and racket kinematics when targeting T, body, or wide serves, enabling placement concealment without sacrificing accuracy. More recently, Fitzpatrick et al. (2024) combined large-scale Hawk-Eye data across tournaments and reported that serves struck toward extreme locations and at higher speeds produce greater pressure on the return. However, most studies classify serve placement into coarse zones (T, body, wide) or rely on player-level averages, providing limited traction on **how unpredictable** a server's distribution is from point to point. The notion of **serve-location entropy** is therefore conceptually appealing but empirically underexplored and rarely evaluated alongside other serve attributes.

2.2 Serve Speed

A complementary line of work emphasizes velocity as a primary driver of serve effectiveness. Using hundreds of professional matches, Mecheri et al. (2016) showed that faster serves are strongly associated with higher point-winning rates, particularly on quicker surfaces. Brown (2021) similarly reported positive correlations between serve speed and points won on first and second serves, with some sex-specific differences in magnitude. These findings position speed as a central determinant of serving effectiveness.

2.3 Open Questions

Despite advances on placement patterns and speed, few studies have jointly modeled speed, placement unpredictability, and accuracy to isolate each component's incremental contribution to point outcomes. Moreover, the practical question of whether unpredictability (entropy) adds value *conditional* on speed and first-serve-in percentage remains largely unanswered. Our study contributes by (i) operationalizing serve-location entropy at the point level, (ii) introducing a short-rally outcome metric (serve efficiency) that targets the immediate advantage servers intend to create, and (iii) evaluating all three components together on out-of-sample data.

3 Data and Exploratory Analysis

3.1 Data Details

We analyzed point-by-point data from the 2018–2019 and 2021–2024 U.S. Open men's singles tournaments, obtained from Jeff Sackmann's public tennis database. For each point, the dataset includes serve speed, serve direction, rally length, and other match context. To isolate serving ability, analyses were restricted to first-serve points only. This choice avoids confounding from second-serve tactics and from extended baseline exchanges.

Our primary response variable is *serve efficiency*, defined as the proportion of first-serve points won within three total shots – specifically, the serve, return, and the server's next stroke. Intuitively, this measures how often a player's serve produces an immediate or near-immediate advantage, rather than merely initiating a neutral rally. Formally,

$$\text{Serve Efficiency} = \frac{\text{Number of first-serve points won within 3 total shots}}{\text{Total number of first-serve points played}}. \quad (1)$$

We considered three explanatory variables:

1. First-Serve Percentage: the proportion of first serves that land in.
2. Serve Speed (mph): the average speed (in mph) of first serves.

3. Serve-Location Entropy: a measure of unpredictability in serve placement. We partition the service target into C non-overlapping locations and compute entropy as

$$\text{Serve Entropy} = -\sum_{i=1}^{C} p_i \log_2 p_i, \tag{2}$$

where p_i is the fraction of first serves landing in location i (terms with $p_i = 0$ contribute 0). Higher values of location entropy indicate more varied, less predictable placement.

We compute each variable as aggregates at the player level, using all first-serve points for each player across years. We now proceed with exploratory analysis.

3.2 Exploratory Data Analysis

We begin by examining the distributions of serve efficiency, first-serve percentage, serve speed, and serve-location entropy in Figures 1 and 2. Serve efficiency is centered near 0.35 and heavily concentrated around the mean, indicating that many players generate similar value from their first serves. Furthermore, average serve speed exhibits moderate left skew, first-serve percentage is approximately symmetric about 0.6, and serve-location entropy demonstrates notable left skew.

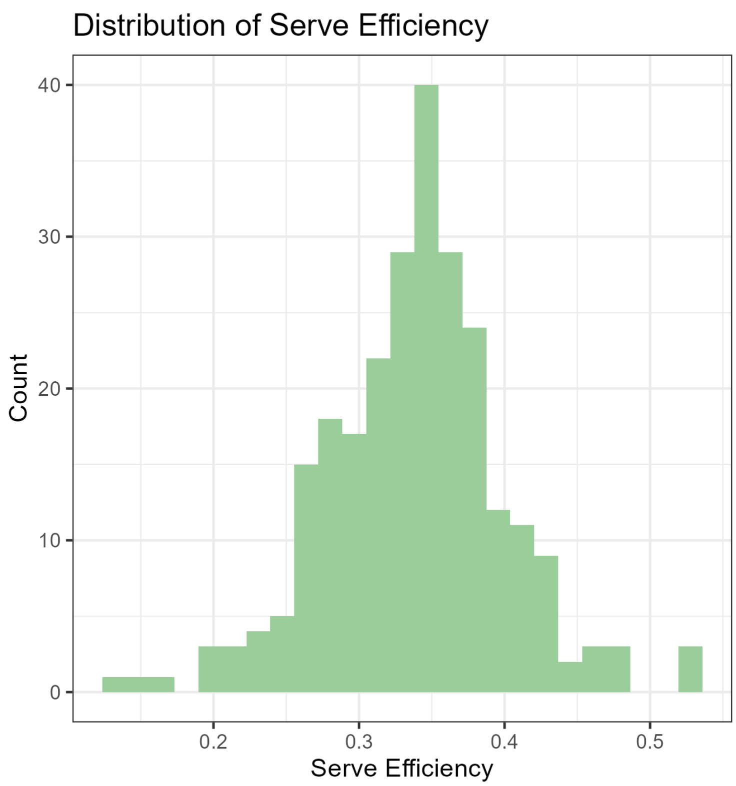

Figure 1: Distribution of serve efficiency across all players (U.S. Open men's singles, 2018–2019 and 2021–2024).

To further investigate entropy, we plot the distribution of serve locations for the five highest- and lowest-entropy servers in Figure 4. High-entropy players distribute serves evenly from the T to the wide line, forcing opponents to defend the full width of the box. In contrast, low-entropy players concentrate their serves in one or two preferred zones.

Finally, we examine correlations between the explanatory variables in Figure 3. These correlations are modest: speed is slightly negatively related to first-serve percentage, and entropy is weakly positively related

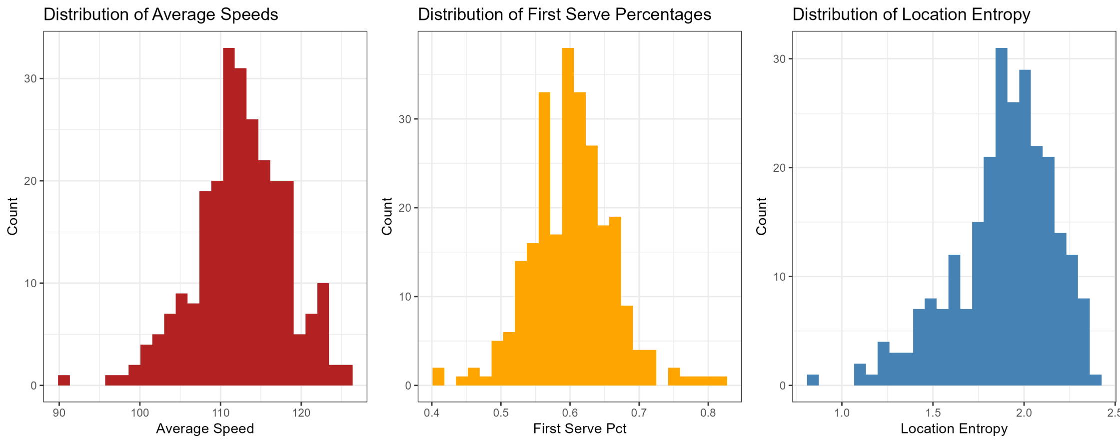

Figure 2: Distributions of the three explanatory variables.

to speed. Importantly, these patterns suggest no severe multicollinearity.

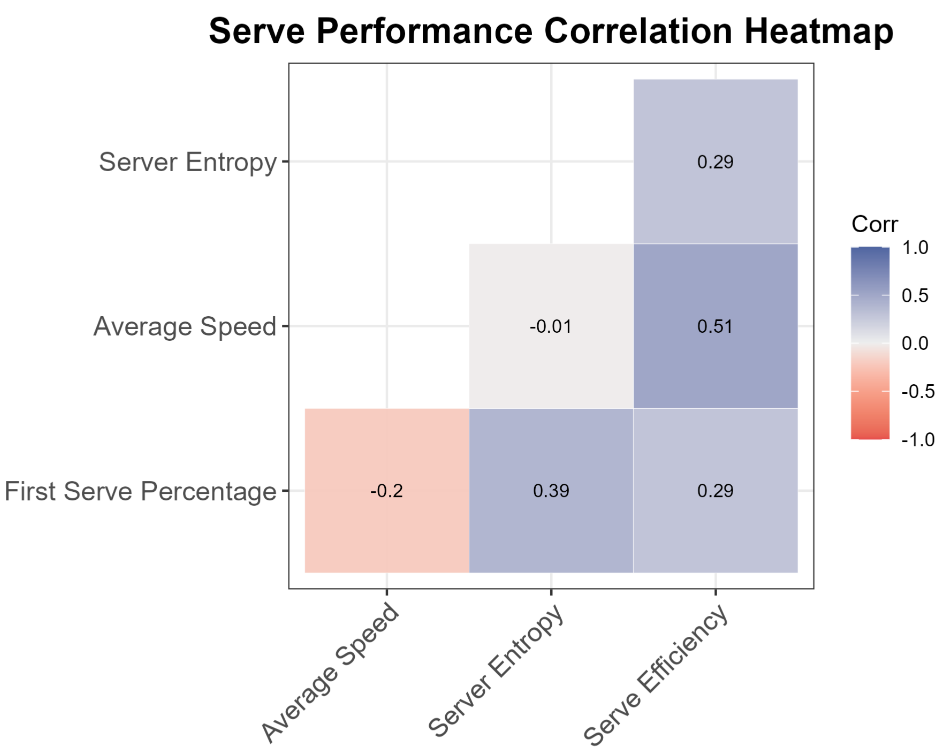

Figure 3: Correlation matrix between the three explanatory variables.

4 Methodology

4.1 Modeling Strategy

We leverage regression to estimate associations between serve characteristics and serve efficiency. First, we explore the marginal associations between each variable and the outcome with locally weighted smoothing (LOESS). Next, we fit a multivariate linear regression with all three standardized predictors (first-serve

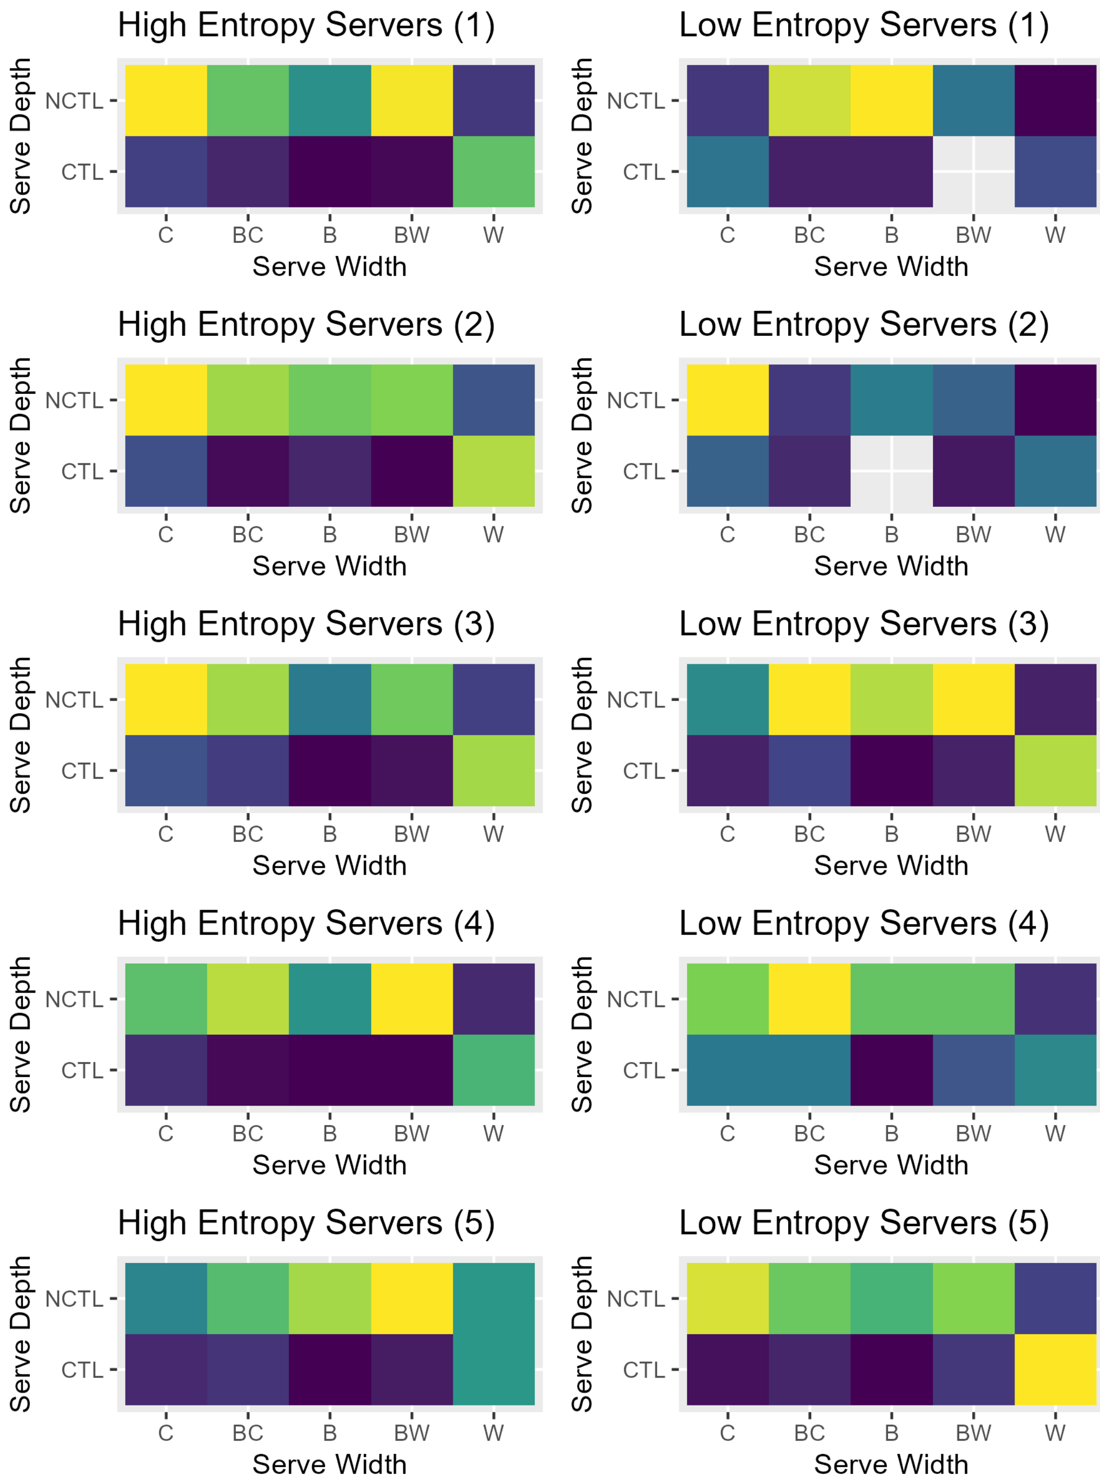

Figure 4: Distribution of serve locations for servers with top 5 highest entropies (left) and 5 lowest entropies (right). Lighter colors (approaching yellow) indicate greater concentration of serves in that location, while darker colors (approaching purple) indicate lower concentration.

percentage, serve speed, and location entropy) simultaneously to assess each factor's incremental contribution conditional on the others. Predictors were standardized to mean 0 and standard deviation 1 prior to modeling; the outcome (serve efficiency) remained on its original scale to preserve interpretability.

4.2 Validation and Evaluation

To assess predictive validity, we randomly partitioned the point-level data into a training set (80%) and a test set (20%), recomputing player-level aggregates separately within each split. We trained an OLS model on the training set, then used these estimated coefficients to generalize predictions of serve efficiency on the held-out test data. We summarize model performance by root mean squared error (RMSE) and the Pearson correlation (r) between predicted and observed serve efficiency in the test set.

4.3 Model Assumptions

We examined standard OLS assumptions for the multivariate linear model. Figure 5a demonstrates that there is no pronounced relationship between fitted values and residuals, consistent with approximate homoskedasticity and linearity. Additionally, Figure 5b indicates that residuals are generally close to normal (with some heavy-tailed behavior), which together with our collinearity check in Figure 3 indicates that our model is a good fit for the data.

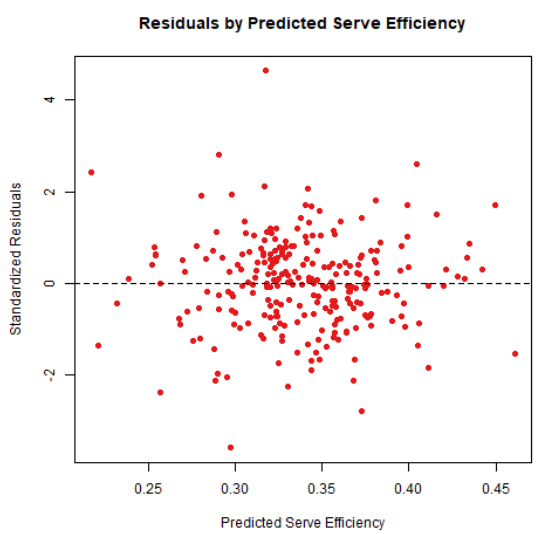

(a) Residual plot for multivariate regression model.

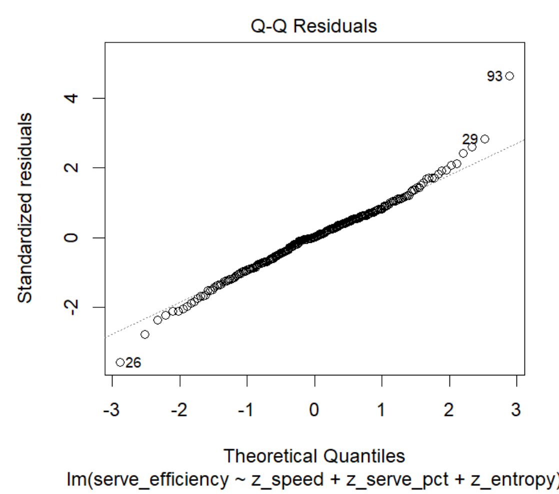

(b) QQ plot for multivariate regression model.

Figure 5: Diagnostic plots for the multivariate regression model: (a) residual plot and (b) QQ plot.

5 Results

5.1 Univariate Associations

LOESS smoothing of serve efficiency against each standardized predictor is pictured in 6, revealing the following marginal relationships:

1. **Serve Speed**: Serve speed is strongly and positively associated with serve efficiency; faster first serves consistently correspond to higher short-rally success.
2. **First-Serve Percentage**: First-serve-in percentage has a positive but smaller association with serve efficiency, with a plateau (and potentially a slight downturn) in the upper tail.

3. **Serve-Location Entropy**: Serve-location entropy shows a smaller positive association with serve efficiency, with a plateau in the upper tail.

Together, these findings indicate that serving power has the strongest positive association with serve efficiency, while first serve percentage and entropy demonstrate weaker positive correlations.

Figure 6: Serve efficiency plotted against each of the three explanatory variables with a LOESS curve and 95% confidence intervals overlaid.

5.2 Multivariate Analysis

Our multivariate OLS model linking all three predictors explains a substantial share of variation in serve efficiency between players: predicted vs. observed serve efficiencies are strongly correlated (r = 0.692), meaning that our model explains about $r^2 \approx 48\%$ of the variance in serve efficiency. When looking at the scatterplot of predicted vs. actual serve efficiency in Figure 7, the linear pattern reinforces that our model predicts serve efficiency well.

Coefficient estimates in Table 1 indicate that serve speed shows the largest association on serve efficiency, with an estimated coefficient of 0.063. First-serve percentage and location entropy have smaller, but still positive, estimated coefficients ($0.023 and 0.013$, respectively). The p-values for all three predictors are close to 0, with speed and first-serve-in percentage having the most significant relationships with serve efficiency ($p < 0.0001$) followed by serve-location entropy ($p < 0.01$). This highlights the predictors' statistical significance as key explanatory factors, especially for speed and first-serve percentage.

Table 1: Estimated coefficients from multivariate regression model.

Term	Estimated Coef.	Std. Error	t-value	p-value
(Intercept)	0.286989301	0.005196998	55.22213033	< 0.0001
Average Speed (std)	0.063477077	0.005297181	11.98318053	< 0.0001
First-Serve % (std)	0.022829983	0.003505690	6.51226474	< 0.0001
Serve Entropy (std)	0.013004929	0.004118859	3.15741066	< 0.01

5.3 Out-of-Sample Testing

On the held-out test set, our model achieved a root mean squared error (RMSE) of 0.066 relative to an observed standard deviation of 0.084 in serve efficiency. This indicates that our model achieved a substantial reduction in error when predicting serve efficiency out-of-sample. Furthermore, the correlation between predicted and observed efficiencies (r) was 0.39. Comparing the predicted vs. observed serve efficiencies to the $y = x$ reference line in Figure 7 shows that our model tends to over-shrink extreme predictions: under-predicting very efficient servers and over-predicting very inefficient ones.

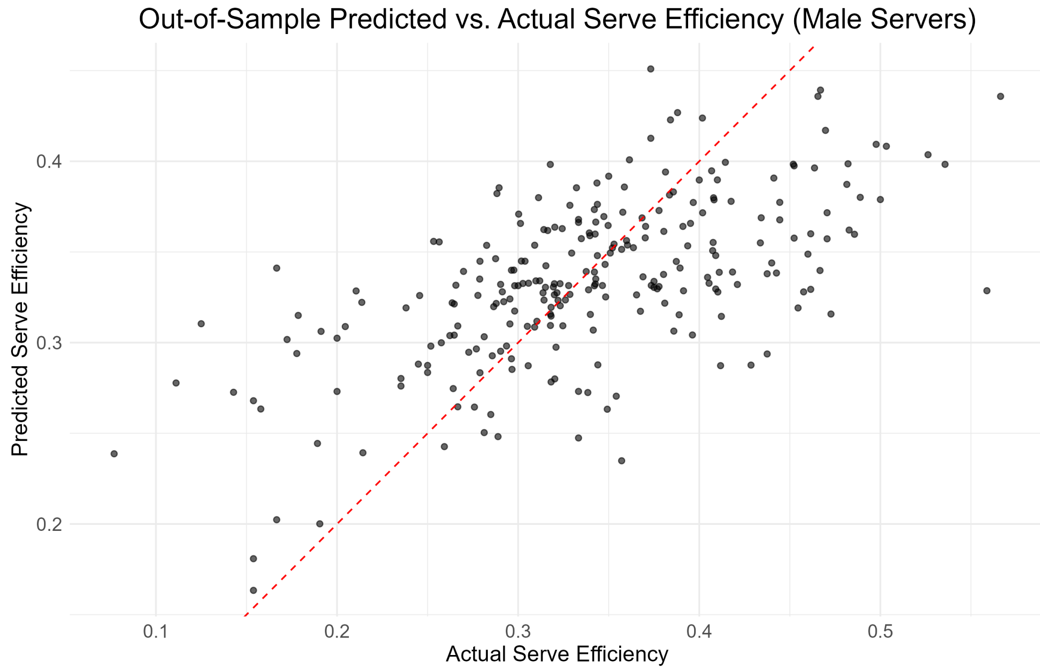

Figure 7: Predicted vs. actual serve efficiency. The red dotted line represents the $y = x$ reference line.

6 Discussion

6.1 Conclusions

This study provides a quantitative framework for understanding how different serve characteristics predict point-level success in elite men's tennis. Across six years of U.S. Open data, we find that serve speed is the strongest predictor of serve efficiency (the share of first-serve points won within three shots). First-serve-in percentage is also meaningfully associated with serve success, but to a smaller extent than speed. Placement unpredictability (serve-location entropy) is also a statistically significant predictor of serve efficiency, but less so than either of the other two predictors. Out-of-sample performance is moderate (RMSE = 0.066 vs a SD of 0.084; r = 0.39), and diagnostic plots suggest that the model captures general patterns but falls short when predicting outlying players.

6.2 Limitations and Future Directions

Our methodology has a few notable limitations. First of all, our analysis is exclusive to the U.S. Open, which is played exclusively on hard courts. This may emphasize the value of speed relative to slower surfaces like clay. Additionally, our findings apply exclusively to first serves, as we left second serves unmodeled. Finally, we do not yet incorporate opponent characteristics, score pressure, or environmental factors that could reasonably influence server tendency and performance. These limitations offer opportunities for future work extending to other tournaments and surfaces, incorporating second-serve behavior, as well as the inclusion of additional contextual and player-level covariates. Methodologically, exploring non-linear terms or interactions between variables may provide additional insights into the predictors of serve efficiency.

7 Reproducibility

The full code and results are in our Github repository for this project: https://github.com/wharton-moneyball/serve-performance.

8 Acknowledgments

We would like to express our gratitude to Moneyball Academy, run by Professor Abraham Wyner through the Wharton Sports Analytics and Business Initiative (WSABI), for providing the opportunity and resources that made this research possible. We would also like to thank Aiwen Li, Tianshu Feng, and Jonathan Pipping for their generous mentorship, guidance, and support throughout this project.

References

Brown, E. G. (2021). A faster serve has more impact on success for female elite tennis players than males. *International Journal of Performance Analysis in Sport*, 21(4):600–610.

Fitzpatrick, A., Stone, J. A., Choppin, S., and Kelley, J. (2024). Analysing hawk-eye ball-tracking data to explore successful serving and returning strategies at wimbledon. *International Journal of Performance Analysis in Sport*, 24(3):251–268.

Kolbinger, O. and Lames, M. (2013). Ball trajectories in tennis - lateral and vertical placement of right handed men's singles serves. *International Journal of Performance Analysis in Sport*, 13(3):750–758.

Mecheri, S., Rioult, F., Mantel, B., Kauffmann, F., and Benguigui, N. (2016). The serve impact in tennis: First large-scale study of big hawk-eye data. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 9(5):310–325.

Reid, M., Whiteside, D., and Elliott, B. (2011). Serving to different locations: set-up, toss, and racket kinematics of the professional tennis serve. *Sports Biomechanics*, 10(4):407–414.