Nonparametric Estimation in the Dynamic Bradley-Terry Model

Heejong Bong* **Wanshan Li*** **Shamindra Shrotriya*** **Alessandro Rinaldo**
Department of Statistics and Data Science
Carnegie Mellon University

Abstract

We propose a time-varying generalization of the Bradley-Terry model that allows for nonparametric modeling of dynamic global rankings of distinct teams. We develop a novel estimator that relies on kernel smoothing to preprocess the pairwise comparisons over time and is applicable in sparse settings where the Bradley-Terry may not be fit. We obtain necessary and sufficient conditions for the existence and uniqueness of our estimator. We also derive time-varying oracle bounds for both the estimation error and the excess risk in the model-agnostic setting where the Bradley-Terry model is not necessarily the true data generating process. We thoroughly test the practical effectiveness of our model using both simulated and real world data and suggest an efficient data-driven approach for bandwidth tuning.

1 Introduction and Prior Work

1.1 Pairwise Comparison Data and the Bradley-Terry Model

Pairwise comparison data are very common in daily life, especially in cases where the goal is to rank several objects. Rather than directly ranking all objects simultaneously, it is usually much easier and more efficient to first obtain results of pairwise comparisons and then use them to derive a *global* ranking across all individuals in a principled manner. Since the required global rankings are not directly observable, developing a statistical framework for the estimating rankings is a challenging unsupervised learning problem. One such statistical model for deriving global rankings using pairwise comparisons was presented in the classic paper (Bradley and Terry, 1952), and thereafter commonly referred to as the Bradley-Terry model in the literature. A similar model was also studied by Zermelo (Zermelo, 1929). The Bradley-Terry model is one of the most popular models to analyze paired comparison data due to its interpretability and computational efficiency in parameter estimation. The Bradley-Terry model along with its variants has been studied and applied in various ranking applications across many domains. This includes the ranking of sports teams (Masarotto and Varin, 2012; Cattelan et al., 2013; Fahrmeir and Tutz, 1994), scientific journals (Stigler, 1994; Varin et al., 2016), and the quality of several brands (Agresti, 2013; Radlinski and Joachims, 2007), to name a few.

In order to introduce the Bradley-Terry model, suppose that we have N distinct teams, each with a positive score s_i, $i \in [N] := \{1, \ldots, N\}$, quantifying its propensity to be picked or win over other items. The model postulates that the comparisons between different pairs are independent and the results of the comparisons between a given pair, say team i and team j, are independent and identically distributed Bernoulli random variables, with *winning probability* defined as

$$\mathbb{P}(i \text{ beats } j) = \frac{s_i}{s_i + s_j}, \; \forall\, i, j \in [N] \tag{1}$$

A common way to parametrize the model is to set, for each i, $s_i = \exp(\beta_i)$, where $(\beta_1, \ldots, \beta_N)$ are real parameters such that $\sum_{i \in [N]} \beta_i = 0$ (this latter condition is to ensure identifiability). In this case, equation (1) is usually expressed as $\mathsf{logit}(\mathbb{P}(i \text{ beats } j)) = \beta_i - \beta_j$, where, for $x \in (0, 1)$, $\mathsf{logit}(x) := \log \frac{x}{1-x}$.

1.2 The time-varying (dynamic) Bradley-Terry Model

In many applications it is very common to observe paired comparison data spanning over multiple (discrete) time periods. A natural question of interest is then to understand how the *global* rankings *vary* over time i.e. *dynamically*. For example, in sports ana-

*Equal contribution.

Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108.

lytics the performance of teams often changes across match rounds and thus explicitly incorporating the time-varying dependence into the model is crucial. In particular the paper (Fahrmeir and Tutz, 1994) considers a state-space generalization of the Bradley-Terry model to analyze sports tournaments data. In a similar manner, bayesian frameworks for the dynamic Bradley-Terry model are studied further in (Glickman, 1993; Glickman and Stern, 1998; Lopez et al., 2018). Such dynamic ranking analysis is becoming increasingly important because of the rapid growth of openly available time-dependent paired comparison data.

Our main focus in this paper is to tackle the problem of generalizing the Bradley-Terry model to the time-varying setting with statistical guarantees. Our approach estimates the *changes* in the Bradley-Terry model parameters over time nonparametrically. This enables the derivation of time-varying global dynamic rankings with minimal parametric assumptions. Unlike previous time-varying Bradley-Terry estimation approaches, we seek to establish guarantees in the model-agnostic setting where the Bradley-Terry model is not the true data generating process. This is in contrast to more assumption-heavy parametric frequentist dynamic Bradley-Terry models including (Cattelan et al., 2013). Our method is also computationally efficient compared to some state-space methods including but not limited to (Fahrmeir and Tutz, 1994; Glickman and Stern, 1998; Maystre et al., 2019).

2 Time-varying Bradley-Terry Model

2.1 Model Setup

In our time-varying generalization of the original Bradley-Terry model we assume that N distinct teams play against each other at possibly different times, over a given time interval which, without loss of generality, is taken to be $[0, 1]$. The result of a game between team i and team j at time t is determined by the timestamped winning probability $p_{ij}(t) := \mathbb{P}(i \text{ defeats } j \text{ at time } t)$ which we assume arises from a distinct Bradley-Terry model, one for each time point. In detail, for each i, j, and t

$$\mathsf{logit}(p_{ij}(t)) = \beta_i(t) - \beta_j(t),\ \forall\, i, j \in [N], t \in [0, 1] \quad (2)$$

where $\boldsymbol{\beta}(t) = \mathrm{vec}(\beta_1(t), \beta_2(t), \ldots, \beta_N(t)) \in \mathbb{R}^N$ is an unknown vector such that $\sum_i \beta_i(t) = 0$.

We observe the outcomes of M timestamped pairwise matches among the N teams $\{(i_m, j_m, t_m) : m \in [M]\}$. Here (i_m, j_m, t_m) indicates that team i_m and team j_m played at time t_m, where $t_1 \leq t_2 \leq \ldots \leq t_M$. The result of the m-th match can be expressed in a $N \times N$ data matrix $X^{(m)}$ as follows:

$$X_{ij}^{(m)} = \begin{cases} \mathbf{1}(i \text{ defeats } j \text{ at } t_m) & \\ \sim \mathrm{Bernoulli}(p_{ij}(t_m)) & \text{for } i = i_m \text{ and } j = j_m \\ 1 - X_{ji}^{(m)} & \text{for } i = j_m \text{ and } j = i_m \\ 0 & \text{elsewhere} \end{cases} \quad (3)$$

Our goal is to estimate the underlying parameters $\boldsymbol{\beta}(t)$ where $t \in [0, 1]$, and then derive the corresponding global ranking of teams. In order to make the estimation problem tractable we assume that the parameters $\boldsymbol{\beta}(t)$ vary smoothly as a function of time $t \in [0, 1]$. It is worth noting that the naive strategy of estimating the model parameters separately at each observed discrete time point on the original data will suffer from two major drawbacks: (i) it will in general not guarantee smoothly varying estimates and, perhaps more importantly, (ii) computing the maximum likelihood estimator (MLE) of the parameters in the static Bradley-Terry model may be infeasible due to sparsity in the data (e.g., at each time point we may observe only one match), as we discuss below in 4. To overcome these issues we propose a nonparametric methodology which involves kernel-smoothing the observed data over time.

2.2 Estimation

Our approach in estimating time-varying global rankings is described in the following three step procedure:

1. *Data pre-processing:* Kernel smooth the pairwise comparison data across all time periods. This is used to obtain the smoothed pairwise data at each time t:

$$\tilde{X}(t) = \sum_{m=1}^{M} W_h(t_m, t) X^{(m)}, \quad (4)$$

   where W_h is an appropriate kernel function with bandwidth h, which controls the extent of data smoothing. The higher the value of h is, the smoother $\tilde{X}_{ij}(t)$ is over time.

2. *Model fitting:* Fit the regular Bradley-Terry model on the smoothed data $\tilde{X}_{ij}(t)$. The model estimates the performance of each team at time t using the estimated score vector

$$\widehat{\boldsymbol{\beta}}(t) = \arg\min_{\boldsymbol{\beta}: \sum_i \beta_i(t) = 0} \widehat{\mathcal{R}}(\boldsymbol{\beta}; t) \quad (5)$$

   minimizing the negative log-likelihood risk

$$\widehat{\mathcal{R}}(\boldsymbol{\beta}; t) = \sum_{i,j: i \neq j} \frac{\tilde{X}_{ij}(t) \log(1 + \exp(\beta_j(t) - \beta_i(t)))}{\sum_{i,j: i \neq j} \tilde{X}_{ij}(t)} \quad (6)$$

3. *Derive global rankings*: Rank each team at time t by its score from $\hat{\boldsymbol{\beta}}(t)$.

We observe that if $t = t_1 = t_2 = \cdots = t_M$ then step 1 reduces to the original (static) Bradley-Terry model. In this case,

$$\begin{aligned} \tilde{X}_{ij}(t) &= W_h(t,t) \sum_{m=1}^{M} \mathbf{1}(i_m = i, j_m = j) \\ &= W_h(t,t) X_{ij} \end{aligned} \tag{7}$$

where $X_{ij} = \#\{i \text{ defeated } j\}$. Thus, fitting the model on $\tilde{X}(t)$ in step 2 is equivalent to the original method on data X. In this sense our proposed ranking approach represents a time-varying generalization of the original Bradley-Terry model.

This *data pre-processing* is a critical step in our method and is similar to the approach adopted in (Zhou et al., 2010) where it was used in the context of estimating smoothed time varying undirected graphs. This approach has two main advantages. First, applying kernel smoothing on the input pairwise comparison data enables borrowing of information across timepoints. In sparse settings, this reduces the data requirement at each time point to meet necessary and sufficient conditions required for the Bradley-Terry model to have a unique solution as detailed in Section 4. Second, kernel smoothing is computationally efficient in a single dimension and is readily available in open source scientific libraries.

3 Our Contributions

Our main contributions in this paper are summarized as follows:

1. We obtain necessary and sufficient conditions for the existence and uniqueness for our time-varying estimator $\widehat{\boldsymbol{\beta}}(t)$ for each $t \in [0,1]$ simulataneously. See Section 4.

2. We extend the results of Simons and Yao (1999) and obtain statistical guarantees for our proposed method in the form of convergence results of the estimated model parameters *uniformly* over all times. We express such guarantees in the form of oracle inequalities in the model-agnostic setting where the Bradley-Terry model is not necessarily the true data generating process. See Section 5.

3. We apply our estimator with an data-driven tuned (by LOOCV) hyperparameter successfully to simulations and to real life applications including a comparison to 5 seasons of NFL ELO ratings. See Section 6 and Section 7.

4 Existence and uniqueness of solution

The existence and uniqueness of solutions for model (5) is not guaranteed in general. This is an innate property of the original Bradley-Terry model (Bradley and Terry, 1952). As pointed out in (Ford, 1957) existence of the MLE for the Bradley-Terry model parameters demands a sufficient amount of pairwise comparison data so that there is enough information of relative performance between any pair of two teams for parameter estimation purposes. For example, if there is a team which has never been compared to the others, there is no information which the model can exploit to assign a score for the team. As such its derived rank could be arbitrary. In addition if there are several teams which have never outperformed the others then the Bradley-Terry model would assign negative infinity for the performance of these teams. It would not be possible to compare amongst them for global ranking purposes. In all such cases, the model parameters are not estimable.

Ford (1957) derived the necessary and sufficient condition for the existence and uniqueness of the MLE in the original Bradley-Terry model. Below we show how this condition can also be adapted to guarantee the existence and uniqueness of the solution in our time-varying Bradley-Terry model. The condition can be stated for each time t in terms of the corresponding kernel-smoothed data $\tilde{X}(t)$ as follows:

Condition 4.1. *In every possible partition of the teams into two nonempty subsets, some team i in the first set and some team j in the second set satisfy $\tilde{X}_{ij}(t) > 0$. Or equivalently, for each ordered pair (i,j), there exists a sequence of indices $i_0 = i, i_1, \ldots, i_n = j$ such that $\tilde{X}_{i_{k-1}i_k}(t) > 0$ for $k = 1, \ldots, n$.*

Remark 1. If we regard $[|\tilde{X}(t)|_{ij}]$ as the adjacency matrix of a (weighted) directed graph, then Condition 4.1 is equivalent to the *strong connectivity* of the graph.

Under condition 4.1 we obtain the following existence and uniqueness theorem on the solution set of the time-varying Bradley-Terry model.

Theorem 4.1. *If the smoothed data $\tilde{X}(t)$ satisfies Condition 4.1, then the solution of* (5) *uniquely exists at time t.*

Hence, in the proposed time-varying Bradley-Terry model we do not require the strong conditions of (Ford, 1957) to be met at each time point, but simply require the *aggregated conditions* in Theorem 4.1 to hold. This is a significant weakening of the data requirement. For example, even if one team did not play any game in a match round – a situation that would preclude the MLE in the standard Bradley-Terry model – it is still possible to assign a rank to this team in their missing round, as long a game is recorded in another round (with at least one win and one loss). In this sense, the kernel-smoothing of the data in our time-varying Bradley-Terry model reduces the required sample complexity for a unique solution.

5 Statistical Properties of the Time-varying Bradley-Terry Model

5.1 Preliminaries

Existing results (Simons and Yao, 1999; Negahban et al., 2017) demonstrate the consistency of the estimated *static* Bradley-Terry scores provided that the data were generated from the Bradley-Terry model. However this assumption may be too restrictive for data generation processes in real world applications. In the rest of this section, we will consider *model-agnostic* time-varying settings where the Bradley-Terry model is not necessarily the true pairwise data generating model. In order to investigate the statistical properties of the proposed estimator, we impose the following relatively mild assumptions.

Assumption 5.1. *Each pair of teams (i,j) play $T^{(i,j)}$ times at time points $\{t_k^{(i,j)}, k = 1, 2, \ldots, T^{(i,j)}\}$, where each $T^{(i,j)} > 0$ satisfy the following conditions, for fixed constants $T > 0$ and $0 < D_m \leq 1 \leq D_M$:*

1. $T^{(i,j)} > T$;
2. *for every interval $(a, b) \subset [0, 1]$,*

$$\begin{aligned} \lfloor D_m(b-a)T^{(i,j)} \rfloor \leq& |\{k : t_k^{(i,j)} \in (a,b)\}| \\ \leq& \lceil D_M(b-a)T^{(i,j)} \rceil. \end{aligned} \tag{8}$$

We remark that the second condition further implies that

$$\begin{aligned} &t_1^{(i,j)} \leq \frac{1}{D_m T^{(i,j)}},\ t_{T^{(i,j)}}^{(i,j)} \geq 1 - \frac{1}{D_m T^{(i,j)}}, \\ &\frac{1}{D_M T^{(i,j)}} \leq t_{k+1}^{(i,j)} - t_k^{(i,j)} \leq \frac{1}{D_m T^{(i,j)}} \end{aligned} \tag{9}$$

for $k = 1, 2, \ldots, T^{(i,j)} - 1$.

Assumption 5.1 allows for different team pairs to play against each other a different number of times and for the game times to be spaced irregularly, though in a controlled manner. To enable statistical analyses on time-varying quantities, we further require that the winning probabilities satisfy a minimal degree of smoothness and that their rate of decay is controlled.

Assumption 5.2. *For any i, j, the function $t \in [0,1] \mapsto p_{ij}(t)$ is Lipschitz with universal constant L_p and uniformly bounded below $p_{\min} > 0$ which is dependent to N and T.*

The quantity $p_{\min}$ need not be bounded away from 0 as function of T and N. However, in order to guarantee estimability of the model parameters in time-varying settings, we will need to control the rate at which it is allowed to vanish. See Theorem 5.1 below.

Finally, we assume that the kernel used to smooth over time satisfy the following regularity conditions, which are quite standard in the nonparametric literature.

Assumption 5.3. *The kernel function $W : (-\infty, \infty) \to (0, \infty)$ is a symmetric function such that*

$$\begin{gathered} \int_{-\infty}^{\infty} W(x)dx = 1,\ \int_{-\infty}^{\infty} |x|W(x)dx < \infty \\ \mathcal{V}(W) < \infty,\ \mathcal{V}(|\cdot|W) < \infty \end{gathered} \tag{10}$$

where $\mathcal{V}(f(x))$ is the total variation of a function f. For each $s, t \in [0,1]$ we further write

$$W_h(s,t) = \frac{1}{h} W\left(\frac{s-t}{h}\right). \tag{11}$$

It is easy to see that these conditions imply

$$\|W\|_\infty = \sup_x W(x) < \infty. \tag{12}$$

Thus, without loss of generality, we assume that $\|W\|_\infty \leq 1$; the general case can be handled by modifying the constants accordingly. The use of kernels satisfying the above assumptions is standard in nonparametric problems involving Hölder continuous functions of order 1, such as the winning probabilities function of Assumption 5.2.

5.2 Existence and uniqueness of solutions

Simons and Yao (1999) showed that the necessary and sufficient condition for the existence and uniqueness of the MLE in the original Bradley-Terry model is satisfied asymptotically almost surely under minimal assumptions. Below, we show that this type of result can be extended to our more general time-varying settings.

Theorem 5.1.

$$\begin{aligned} &\mathbb{P}(\textit{Condition 4.1 is satisfied at every } t) \\ &\geq 1 - 4N \exp\left(-\frac{NT}{2} p_{min}\right). \end{aligned} \tag{13}$$

Remark 2. As we remarked above, $p_{\min}$ needs not be bounded away from zero, but is allowed to vanish slowly enough in relation to N and T so that condition (5.1) is fulfilled as long as $\frac{1}{p_{\min}} = o\left(\frac{NT}{2\log N}\right)$.

5.3 Oracle Properties

In our general agnostic time-varying setting the Bradley-Terry model is not assumed to be the true data generating process. It follows that, for each t, there is no true parameter to which to compare the estimator defined in (5). Instead, we may compare it to the *projection* parameter $\boldsymbol{\beta}^*(t) \in \mathbb{R}^N$, which is the best approximation to the winning probabilities at time t using the dynamic Bradley Terry model; see (2). In detail, the oracle parameter is defined as

$$\boldsymbol{\beta}^*(t) = \arg\min_{\boldsymbol{\beta}:\sum_i \beta_i(t)=0} \mathcal{R}(\boldsymbol{\beta}; t) \tag{14}$$

where

$$\mathcal{R}(\boldsymbol{\beta}; t) = \frac{1}{\binom{N}{2}} \sum_{i,j:i\neq j} p_{ij}(t) \log(1 + \exp(\beta_j(t) - \beta_i(t))) \tag{15}$$

We note that, when the winning probabilities obey a Bradley-Terry model, the projection parameter corresponds to the true model parameters.

Next, for each fixed time $t \in (0, 1)$ and $h > 0$, we introduce two quantities, namely $M(t)$ and $\delta_h(t)$, that can be thought of as conditions numbers of sort, affecting directly both the estimation and prediction accuracy of the proposed estimator. In detail, we set

$$M(t) = \max_{i,j:i\neq j} \exp(\beta_i^*(t) - \beta_j^*(t))$$

and

$$\delta_h(t) = \max_i \sum_{j:j\neq i} \left| \frac{\tilde{T}_{ij}(t)}{\tilde{T}_i(t)} - \frac{1}{N-1} \right|.$$

where $\tilde{T}_{ij}(t) = \tilde{X}_{ij}(t) + \tilde{X}_{ji}(t)$ and $\tilde{T}_i(t) = \sum_{j:j\neq i} \tilde{T}_{ij}(t)$. The ratio $M(t)$ quantifies the maximal discrepancy in winning scores among all possible pairs at time t, and, as shown in Simons and Yao (1999), determines the consistency rate of the MLE in the traditional Bradely-Terry model (see also the comments following Remark 3 below). The quantity $\delta_h(t)$ is instead a measure of regularity in how evenly the teams play against each other. In particular $\delta_h(t) = 0$, for all t and h when there is a constant number of matches among each pair of teams, for each time. Since we allow for the possibility of different number of matches between teams and across time, it becomes necessary to quantity such degree of design irregularity.

In order to verify the quality of the proposed estimator $\widehat{\boldsymbol{\beta}}(t)$, we will consider high-probability oracle bound on estimation error $\|\widehat{\boldsymbol{\beta}}(t) - \boldsymbol{\beta}^*(t)\|_\infty$. In the following theorems, we present both a point-wise and a uniform in t version of this bound in the asymptotic regime of $T, N \to \infty$ and under only minimal assumptions on the ground-truth winning probabilities $p_{ij}(t)$'s.

Theorem 5.2. *Let $\gamma = \gamma(T, N, p_{\min})$ be the probability that Condition 4.1 fails and suppose that the kernel bandwidth is chosen as*

$$h = \max\left\{ \frac{1}{T^{1+\eta}}, \left(\frac{36(1-p_{\min})\log N}{C_s^2 D_m (N-1)T} \right)^{\frac{1}{3}} \right\},$$

for any $\eta > 0$ and some universal constant C_s depending only to D_m, D_M, and W. Then, for each fixed time $t \in (0, 1)$ and sufficiently large N and T,

$$\|\hat{\boldsymbol{\beta}}(t) - \boldsymbol{\beta}^*(t)\|_\infty \le 48 M(t) \left(\delta_h(t) + C_s h\right) \tag{16}$$

with probability at least $1 - \frac{2}{N} - \gamma$ as long as the right hand side is smaller than $\frac{1}{3}$.

Next, we strengthen our previous result, which is valid for each fixed time t, to a uniform guarantee over the entire time course.

Theorem 5.3. *Let $\gamma = \gamma(T, N, p_{\min})$ be the probability that Condition 4.1 fails and suppose that the kernel bandwidth is chosen as*

$$h = \max\left\{ \frac{1}{T^{1+\eta}}, \left(\frac{36(1-p_{\min})\log(NT^{3+3\eta})}{C_s^2 D_m (N-1)T} \right)^{\frac{1}{3}} \right\}$$

and that W is L_W-Lipschitz. Then, for sufficiently large N and T,

$$\sup_{t\in[0,1]} \|\hat{\boldsymbol{\beta}}(t) - \boldsymbol{\beta}(t)^*\|_\infty \le 48 \sup_{t\in[0,1]} M(t) \left(\delta_h(t) + C_s h\right) \tag{17}$$

with probability at least $1 - \frac{2h^3}{N} - \gamma$ as long as the right hand side is smaller than $\frac{1}{3}$.

Remark 3. The rate of point-wise convergence for the estimation error implied by the previous result is

$$M(t)\left(\delta_h(t) + \max\left\{\frac{1}{T^{1+\eta}}, \left(\sqrt{\frac{\log N}{NT}}\right)^{\frac{2}{3}}\right\}\right), \tag{18}$$

while the rate for uniform convergence is

$$M(t)\left(\delta_h(t) + \max\left\{\frac{1}{T^{\frac{1}{2}+\eta}}, \left(\sqrt{\frac{\log(NT^{1+\eta})}{NT}}\right)^{\frac{2}{3}}\right\}\right). \tag{19}$$

Importantly, as we can see in the previous results, the proposed time varying estimator $\hat{\boldsymbol{\beta}}(t)$ is consistent only provided that the design regularity parameter $\delta_h(t)$ goes to zero. Of course, if all teams play each other a constant number of times, then $\delta_h(t) = 0$ automatically. In general, however, the impact of the design on the estimation accuracy needs to be assessed on a case-by-case basis.

The rate (18) should be compared with the convergence rate to the true parameters under the static Bradley Terry model, which (Simons and Yao, 1999) show to be $O_p(\max_{i,j:i\neq j} \exp(\beta_i^* - \beta_j^*)\sqrt{\log N/NT})$. Thus, not surprisingly, in the more challenging dynamic settings with smoothly varying winning probabilities the estimation accuracy decreases. The exponent of $\frac{2}{3}$ in the rate (18) matches the familiar rate for estimating Hölder continuous function of order 1.

From (18) and (19) we observe that the desired oracle property on estimated parameters requires rate constraints on $M(t)$. These constraints appear to be strong assumptions without a direct connection to $p_{ij}(t)$'s in our model-agnostic setting. Instead, we circumvent this issue by introducing a more interpretable condition number K (or $p_{\min}$), dependent on N, T, given by

$$K = \exp\left(\frac{1}{p_{\min}}\right)$$

and proving that for each fixed time $t \in (0,1)$ our desired oracle property follows from a bound on $M(t)$.

Theorem 5.4. *Under the conditions in Theorem 5.2*

$$\|\hat{\boldsymbol{\beta}}(t) - \boldsymbol{\beta}^*(t)\|_\infty \leq 72K\left(\delta_h(t) + C_s h\right) \tag{20}$$

and under the conditions Theorem 5.3

$$\sup_{t\in[0,1]} \|\hat{\boldsymbol{\beta}}(t) - \boldsymbol{\beta}(t)^*\|_\infty \leq 72K \sup_{t\in[0,1]} \left(\delta_h(t) + C_s h\right) \tag{21}$$

with probability at least $1 - \frac{2}{N} - \gamma$.

Remark 4. We note that assuming $p_{\min}$ to be bounded away from 0 ensures $\sup_{t\in[0,1]} \|\hat{\boldsymbol{\beta}}(t) - \boldsymbol{\beta}^*(t)\|_\infty \to 0$, with high probability. This assumption means that no team is uniformly dominated by or dominates others (since this implies $1 - p_{\min}(t)$ is bounded away from 1). This is a reasonable assumption in real-world data such as sports match histories where teams are screened to be competitive with each other. Therefore it is reasonable to only consider matches between teams that do not have vastly different skills, which is reflected in winning probabilities that are bounded away from $\{0, 1\}$.

In summary, our proposed method achieves high-probability oracle bounds on the estimation error in our general model-agnostic time-varying setting. We provide the proofs for the stated theorems in Appendix Section 9.1.

6 Experiments

We compare our method with some other methods on synthetic data*. We consider both cases where the pairwise comparison data are generated from the Bradley-Terry model and from a different, nonparametric model.

*Code available at `https://github.com/shamindras/bttv-aistats2020`

6.1 Bradley-Terry Model as the True Model

First we conduct simulation experiments in which the Bradley-Terry model is the true model. Given the number of teams N and the number of time points M, the synthetic data generation process is as follows:

1. For $i \in [N]$, simulate $\boldsymbol{\beta}_i \in \mathbb{R}^M$ as described below;
2. For $1 \leq i < j \leq N$ and $t \in [M]$, set $n_{ij}(t)$ and simulate $X(t)$ by $x_{ij}(t) \sim \text{Binom}\Big(n_{ij}(t), 1/\{(1 + \exp[\beta_j(t) - \beta_i(t)]\}\Big)$ and $x_{ji}(t) = n_{ij}(t) - x_{ij}(t)$.

For each $i \in [N]$, we generate $\boldsymbol{\beta}_i \in \mathbb{R}^M$ from a Gaussian process $\text{GP}(\mu_i(t), \sigma_i(t,s))$ as follows:

1. Set the values of the mean process $\mu_i(t)$ for all $t \in [M]$ and get mean vector $\boldsymbol{\mu}_i = (\mu_i(1), \ldots, \mu_i(M))$;
2. Set the values of the variance process $\sigma_i(t,s)$ at $(s,t) \in [M]^2$, to derive $\Sigma_i \in \mathbb{R}^{M\times M}$;
3. Generate a sample $\boldsymbol{\beta}_i$ from Normal$(\boldsymbol{\mu}_i, \Sigma_i)$.

In our experiment, we generate the parameter $\boldsymbol{\beta}$ via a Gaussian process for $N = 50$, $M = 50$ and $n_{ij}(t) = 1$ for all t. See appendix for full details. We compare the true $\boldsymbol{\beta}$, the win rate, and $\hat{\boldsymbol{\beta}}$ by different methods in Fig. 1.

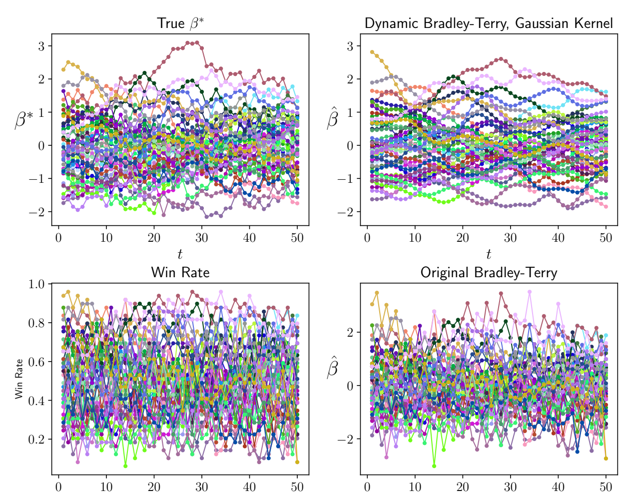

Figure 1: Comparison of $\boldsymbol{\beta}$ and different estimators. First row: true $\boldsymbol{\beta}$ (left), $\hat{\boldsymbol{\beta}}$ with our dynamic BT (right); second row: win rate (left), original BT (right).

We use LOOCV to select the kernel parameter h. The CV curve can be found in Section 6 of Appendix. In this relatively sparse simulated data we see that our dynamic Bradley-Terry estimator $\hat{\boldsymbol{\beta}}$ recovers the comprehensive global ranking of each team in line with the

Estimator	Rank Diff	LOO Prob	LOO nll
Win Rate	3.75	0.44	-
Original BT	3.75	0.37	0.56
Dynamic BT	2.29	0.37	0.55

Table 1: Comparison of Different estimators with results based on 20 repeats. Rank Diff. is the average absolute difference between the estimated and true rankings of teams. LOO Prob means average leave-one-out prediction error of win/loss. LOO nll means average leave-one-out negative log-likelihood.

original Bradley-Terry model. We also observe that due to the kernel-smoothing that our estimator has relatively more stable paths over time.

Table 1 compares the three estimators across key metrics. The results are averaged over 20 repeats. As expected, in this sparse data setting, our dynamic Bradley-Terry method performs better than the original Bradley-Terry model.

6.2 Model-Agnostic Setting

In our second experiment we adopt a model-agnostic setting where we assume that Bradley-Terry model is not necessarily the true data generating model, as described in Section 5.1. With the same notation, for $1 \leq i < j \leq N$ and $t \in [M]$ we first simulate $p_{ij}(t)$, and then set $n_{ij}(t)$ and simulate $X(t)$ by $x_{ij}(t) \sim \text{Binom}\Big(n_{ij}(t), p_{ij}(t)\Big)$ and $x_{ji}(t) = n_{ij}(t) - x_{ij}(t)$. To generate a *smoothly changing* $p_{ij}(t)$, we again use Gaussian process. Specifically, first we generate $p_{ij}(t)$ for $1 \leq i < j \leq N$ and $t \in [M]$ from a Gaussian process. Then we scale those $p_{ij}(t)$'s uniformly to make the values fall within a range $[p_l, p_u]$. In our experiment we set $M = 50$, $N = 50$, $[p_l, p_u] = [0.05, 0.95]$ and $n_{ij}(t) = 1$ for all t. The projection parameter $\boldsymbol{\beta}^*$, the win rate, and $\widehat{\boldsymbol{\beta}}$ by different methods are compared in Fig. 2. Again the kernel parameter h in our model is selected by LOOCV.

Estimator	Rank Diff	LOO Prob	LOO nll
Win Rate	10.68	0.49	-
Original BT	10.70	0.49	0.71
Dynamic BT	5.48	0.49	0.68

Table 2: Comparison of different estimators when the underlying model is not the Bradley-Terry model.

By comparing curves in Fig. 2, we note that our estimator $\hat{\boldsymbol{\beta}}$ recovers the global rankings better than the win rate and the original Bradley-Terry model, and produces relatively more stable paths over time. The same conclusion is confirmed by Table 2, which compares the three estimators in some metrics with 20 repetitions.

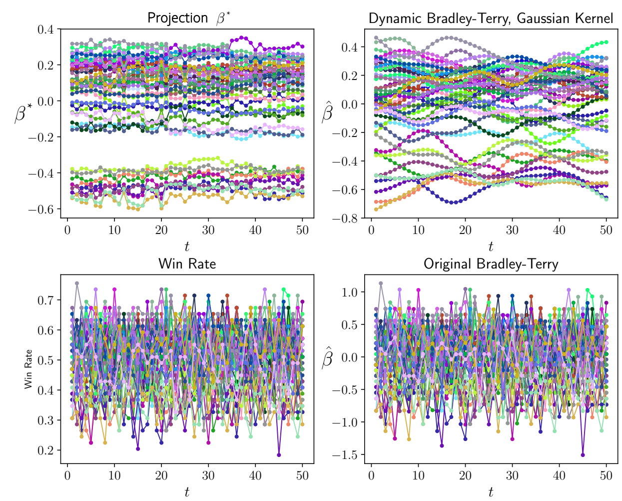

Figure 2: Comparison of $\boldsymbol{\beta}^*$ and different estimators when the underlying model is not the Bradley-Terry model. First row: projection $\boldsymbol{\beta}^*$ (left), our dynamic BT (right); second row: win rate (left), original BT (right).

Remark 5. In this sparse data setting where $n_{ij}(t)$ is fairly small, the original Bradley-Terry model performs worse than our model for two reasons: 1. Condition 4.1 can fail to hold at some time points, whence the MLE does not exist; 2. even when the MLE exists, it can fluctuate significantly over time because of the relatively small sample size at each time point. As we show in Section 9.3.4 in the appendix, when $M = 50$, $N = 50$ and $n_{ij}(t) = 1$, the MLE exists with fairly high frequency. Still, our model performs much better than the original Bradley-Terry model.

Remark 6. Since our method is aimed at obtaining accurate estimates of *smoothly* changing beta/rankings with strong prediction power, it may fail to capture some changes in rankings, especially when these changes are relatively small (as in the present case). However the winrate and original Bradley-Terry methods perform much worse, as they appear to miss some true ranking changes while returning many false change points.

Additional details about experiments, including running time efficiency in simulated settings, can be found in the Appendix Section 9.3.

7 Application - NFL Data

In order to test our model in practical settings we consider ranking National Football League (NFL) teams over multiple seasons. Specifically we source 5 seasons of openly available NFL data from 2011-2015 (inclusive) using the `nflWAR` package (Yurko et al., 2018). Each NFL season is comprised of $N = 32$ teams play-

rank	2011		2012		2013		2014		2015	
	ELO	BT	ELO	BT	ELO	BT	ELO	BT	ELO	BT
1	GB	GB	NE	HOU	SEA	SEA	SEA	DEN	SEA	CAR
2	NE	SF	DEN	ATL	SF	DEN	NE	ARI	CAR	DEN
3	NO	NO	GB	SF	NE	NO	DEN	NE	ARI	NE
4	PIT	NE	SF	CHI	DEN	KC	GB	SEA	KC	CIN
5	BAL	DET	ATL	GB	CAR	SF	DAL	DAL	DEN	ARI
6	SF	BAL	SEA	NE	CIN	NE	PIT	GB	NE	GB
7	ATL	PIT	NYG	DEN	NO	IND	BAL	PHI	PIT	MIN
8	PHI	HOU	CIN	SEA	ARI	CAR	IND	SD	CIN	KC
9	SD	CHI	BAL	BAL	IND	ARI	ARI	DET	GB	PIT
10	HOU	ATL	HOU	IND	SD	CIN	CIN	KC	MIN	SEA
Av. Diff.	4.2		5.0		3.5		4.3		3.4	

Table 3: BT within season vs. ELO NFL top 10 rankings. Blue: perfect match, yellow: top 10 match. Our dynamic BT model is fitted on 16 rounds of each season, and the ranking of a season is based on the ranking at the last round.

ing $M = 16$ games each over the season. This means that at each point in time t the pairwise comparison matrix based on scores across all 32 teams is sparsely populated with only 16 entries. We fit our time-varying Bradley Terry estimator over all 16 rounds in the season using a standard Gaussian Kernel and tune h using the LOOCV approach described in section 9.2. In order to gauge whether the rankings produced by our model are reasonable we compare our season-ending rankings (fit over all games played in that season) with the relevant openly available NFL ELO ratings (Paine, 2015). The top 10 season-ending rankings from each method across NFL seasons 2011-2015 are summarized in Table 3.

Based on Table 3 we observe that we roughly match between 6 to 10 of the top 10 ELO teams consistently over all 5 seasons. There is often misalignment with specific ranking values across both ranking methods. We note that the unlike our method, the NFL ELO rankings use pairwise match data and also additional features including an adjustment for margin of victory. This demonstrates an advantage of our model in only requiring the minimal time-varying pairwise match data and smoothness assumptions to deliver comparable results to this more feature rich ELO ranking method. Furthermore, since our model aggregates data across time it can provide a reasonable minimalist ranking benchmark in modern sparse time-varying data settings with limited “expert knowledge” e.g. e-sports.

8 Conclusion

We propose a time-varying generalization of the Bradley-Terry model that captures temporal dependencies in a nonparametric fashion. This enables the modeling of dynamic global rankings of distinct teams in settings in which the parameteres of the ordinary Bradley Terry model would not be estimable.

From a theoretical perspective we adapt the results of (Ford, 1957) to obtain the necessary and sufficient condition for the existence and uniqueness of our Bradley-Terry estimator in the time-varying setting. We extend the previous analysis of (Simons and Yao, 1999) to derive oracle inequalities on for our proposed method for both the estimation error and the excess risk under smoothness conditions on the winning probabilities. The resulting rates of consistency are of nonparametric type.

From an implementation perspective we provide a general strategy for tuning the kernel bandwidth hyperparameter using an efficient data-driven approach specific to our unsupervised time-varying setting. Finally, we test the practical effectiveness of our estimator under both simulated and real world settings. In the latter case we separately rank 5 consecutive seasons of open National Football League (NFL) team data (Yurko et al., 2018) from 2011-2015. Our NFL ranking results compare favourably to the well-accepted NFL ELO model rankings (Paine, 2015). We thus view our nonparametric time-varying Bradley-Terry estimator as a useful *benchmarking tool* for other feature-rich time-varying ranking models since it simply relies on the minimalist time-varying score information for modeling.

References

Agresti, A. (2013). *Categorical data analysis*. Wiley Series in Probability and Statistics. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, third edition.

Bradley, R. A. and Terry, M. E. (1952). Rank analysis of incomplete block designs. I. The method of paired comparisons. *Biometrika*, 39:324–345.

Cattelan, M., Varin, C., and Firth, D. (2013). Dynamic Bradley-Terry modelling of sports tournaments. *J. R. Stat. Soc. Ser. C. Appl. Stat.*, 62(1):135–150.

Fahrmeir, L. and Tutz, G. (1994). Dynamic stochastic models for time-dependent ordered paired comparison systems. *Journal of the American Statistical Association*, 89(428):1438–1449.

Ford, Jr., L. R. (1957). Solution of a ranking problem from binary comparisons. *Amer. Math. Monthly*, 64(8, part II):28–33.

Glickman, M. E. (1993). *Paired comparison models with time-varying parameters*. ProQuest LLC, Ann Arbor, MI. Thesis (Ph.D.)–Harvard University.

Glickman, M. E. and Stern, H. S. (1998). A state-space model for national football league scores. *Journal of the American Statistical Association*, 93(441):25–35.

Lopez, M. J., Matthews, G. J., and Baumer, B. S. (2018). How often does the best team win? A unified approach to understanding randomness in North American sport. *Ann. Appl. Stat.*, 12(4):2483–2516.

Masarotto, G. and Varin, C. (2012). The ranking lasso and its application to sport tournaments. *Ann. Appl. Stat.*, 6(4):1949–1970.

Maystre, L., Kristof, V., and Grossglauser, M. (2019). Pairwise comparisons with flexible time-dynamics. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1236–1246.

Negahban, S., Oh, S., and Shah, D. (2017). Rank centrality: ranking from pairwise comparisons. *Oper. Res.*, 65(1):266–287.

Paine, N. (2015). NFL Elo Ratings Are Back! `https://fivethirtyeight.com/features/nfl-elo-ratings-are-back/`.

Radlinski, F. and Joachims, T. (2007). Active exploration for learning rankings from clickthrough data. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '07, pages 570–579, New York, NY, USA. ACM.

Raghavan, P. (1988). Probabilistic construction of deterministic algorithms: approximating packing integer programs. *Journal of Computer and System Sciences*, 37(2):130–143.

Simons, G. and Yao, Y.-C. (1999). Asymptotics when the number of parameters tends to infinity in the Bradley-Terry model for paired comparisons. *Ann. Statist.*, 27(3):1041–1060.

Stigler, S. M. (1994). Citation patterns in the journals of statistics and probability. *Statist. Sci.*, 9:94–108.

Varin, C., Cattelan, M., and Firth, D. (2016). Statistical modelling of citation exchange between statistics journals. *J. Roy. Statist. Soc. Ser. A*, 179(1):1–63.

von Luxburg, U. (2007). A tutorial on spectral clustering. *Stat. Comput.*, 17(4):395–416.

Yurko, R., Ventura, S., and Horowitz, M. (2018). nflwar: A reproducible method for offensive player evaluation in football. *arXiv preprint arXiv:1802.00998*.

Zermelo, E. (1929). Die Berechnung der Turnier-Ergebnisse als ein Maximumproblem der Wahrscheinlichkeitsrechnung. *Math. Z.*, 29(1):436–460.

Zhou, S., Lafferty, J., and Wasserman, L. (2010). Time varying undirected graphs. *Mach. Learn.*, 80(2-3):295–319.