Generalized Bradley-Terry Models for Score Estimation from Paired Comparisons

Julien Fageot[1,2], **Sadegh Farhadkhani**[2], **Lê-Nguyên Hoang**[1,3], **and Oscar Villemaud**[1,2]
[1] Tournesol Association
[2] EPFL
[3] Calicarpa
julien.fageot@gmail.com, sadegh.farhadkhani@epfl.ch, len@tournesol.app, oscar.villemaud@epfl.ch

Abstract

Many applications, e.g. in content recommendation, sports, or recruitment, leverage the comparisons of alternatives to score those alternatives. The classical Bradley-Terry model and its variants have been widely used to do so. The historical model considers binary comparisons (victory/defeat) between alternatives, while more recent developments allow finer comparisons to be taken into account. In this article, we introduce a probabilistic model encompassing a broad variety of paired comparisons that can take discrete or continuous values. We do so by considering a well-behaved subset of the exponential family, which we call the family of *generalized Bradley-Terry (GBT) models*, as it includes the classical Bradley-Terry model and many of its variants. Remarkably, we prove that all GBT models are guaranteed to yield a strictly convex negative log-likelihood. Moreover, assuming a Gaussian prior on alternatives' scores, we prove that the maximum a posteriori (MAP) of GBT models, whose existence, uniqueness and fast computation are thus guaranteed, varies monotonically with respect to comparisons (the more A beats B, the better the score of A) and is Lipschitz-resilient with respect to each new comparison (a single new comparison can only have a bounded effect on all the estimated scores). These desirable properties make GBT models appealing for practical use. We illustrate some features of GBT models on simulations.

Introduction

In many settings, alternatives are rather compared than individually scored. Typically, in chess, football, tennis, judo, or cycling, individuals and teams compete against one another. Similarly, students and job candidates are arguably easier to compare, rather than to assess directly. In fact, comparative judgments are implicitly performed all the time in online applications, as users often have to select content, applications, or products to consume or purchase, within a set of proposed alternatives. However, ranking all (or a subset of top) alternatives is often demanded. Many sport competitions identify a current number-one player or team, job candidates are ordered for hiring procedures and recommendation AIs must select a handful of content to recommend. Such rankings are often produced based on scoring inferred from comparisons. Scores also allow to reflect the fact that an alternative vastly outperforms a particularly bad alternative, while it is only slightly better than a third alternative.

Transforming comparisons into scores is not straightforward, especially when the comparisons are noisy. Typically, better sport teams are sometimes unfortunate, and end up losing against less competitive teams, which is sometimes called the "beauty of sport". In other applications, the comparative judgments may vary because they are made by different individuals, or simply because it is hard for humans to remain consistent in their sequential judgments.

Contributions

In this paper, we introduce and analyze a natural and well-behaved family of probabilistic models that convert observed comparisons into individual scoring. Essentially, our family of models, which we call the generalized Bradley-Terry (GBT) family, is obtained by considering a subset of the exponential family of particular interest. Interestingly, a practitioner then merely needs to define how they expect two equally good alternatives to be compared, to effortlessly construct a unique model of our family. We show that the GBT family generalizes the well-known and widely studied Bradley-Terry model (Bradley and Terry 1952) (which is limited to binary comparisons in its historical version), as well as other more recent models (Guo et al. 2012; Kristof et al. 2019). In fact, we highlight the generality of our GBT models by briefly analyzing multiple noteworthy instances, including the K-nary-GBT, the Gaussian-GBT, the Uniform-GBT, and the Poisson-GBT.

Remarkably, as our key contribution, we prove that, given a Gaussian Bayesian prior on alternatives' scores, all GBT models are guaranteed to yield several desirable properties. Namely, first, all yield a *strongly convex negative log-posterior*, which means that the maximum-a-posteriori (MAP) can be quickly computed by any optimizer for strongly convex losses. Second, we prove that the MAP scores vary *monotonically* with the comparisons. This guarantees that alternatives will always be incentivized to "win comparisons" by as large of a margin as possible. Finally, we prove that the MAP scores are *Lipschitz-resilient* to each new comparison, thereby guaranteeing that any outlier comparison will only have a limited impact on the estimated scores. This is especially important in the context of online applications, where misclicks are extremely common.

These properties of the GBT models make them appealing to practitioners. In fact, they have been deployed on the online collaborative Tournesol platform (Hoang et al. 2021), which aims to construct a secure, ethical and collaboratively governed content recommendation algorithm. To do so, the Tournesol platform asks its contributors to provide comparative judgments of which of a pair of videos should be more often recommended by the Tournesol recommender. They reportedly use a GBT model, among other tools, to transform such comparisons into video scores. An openly available implementation of GBT, under AGPL license, was made available in the python package called `solidago`.

Related Works

An important motivation behind this work is to lay the theoretical foundations of the method of estimating personal scores from paired comparisons made by Tournesol users (Hoang et al. 2021). We therefore take inspiration from similar methods applied in contexts where they have proven their worth. One classical approach consists of defining a probabilistic model that conceives the comparisons as random variations from intrinsic scores (modeled as latent variables) of these comparisons (David 1963).

The Bradley-Terry model (Bradley and Terry 1952), which follows the pioneering work of Zermelo (Zermelo 1929), proposes to consider this task from binary comparisons (victory or defeat). These ideas are the root of Elo rating system used in chess (Élö 1978) and other competitive contexts[1]. A similar approach was followed independently by Thurstone (Thurstone 1927) in psychophysics. The Bradley-Terry and Thurstonian models have been generalized in many ways, for instance in order to include ties (Davidson 1970; Rao and Kupper 1967). Refinements of the Elo rating such as Glicko (Glickman 1999) or TrueSkill (Herbrich, Minka, and Graepel 2006) have been introduced and practically used, together with online or Bayesian techniques to compute score estimators (Hunter 2004; Cattelan 2012).

Recent approaches considered different comparison models, including unbounded Poisson (Maher 1982) and Skellam (i.e. symmetric Poisson) models (Karlis and Ntzoufras 2009), or continuous-domain Gaussian models (Guo et al. 2012; Maystre, Kristof, and Grossglauser 2019; Kristof et al. 2019). To the best of our knowledge, there is no theory covering all of these models which are based on different modeling of the comparisons. This article proposes precisely to introduce a unified framework including them all.

Outline

The rest of the paper is organized as follows. First, we define the setting, and introduce the GBT models. We especially stress the importance of the *cumulant-generating functions*, wherein lie so many of the well-behaved properties of GBT models. We then introduce MAP estimators based on GBT models, given a Gaussian prior on scores, and highlight their basic computational and statistical properties. Next, we define the *monotonicity* of score estimators and prove that any GBT MAP estimator has this desirable property. In the following section, we (re)define *Lipchitz-resilience* to user's modifications, and show that GBT MAP estimators with bounded scores are Lipschitz-resilient. We then exemplify GBT models, from the historical binary Bradley-Terry model to continuous-domain ones, provide illustrative simulations, and finally conclude in the last section.

[1]For instance, the Elo rating of tennis players are computed here: https://tennisabstract.com/reports/wta_elo_ratings.html.

Generalized Bradley-Terry Models

In this section, we introduce the setting and the GBT models. We then redefine cumulant-generating functions and highlight a remarkable result about these functions, which will prove, as a corollary, that the MAP of GBT models can be efficiently computed. Finally, we draw the connection with the historical Bradley-Terry model.

The Setting

Consider a set $\mathcal{A} = \{a\}_{a\in\mathcal{A}}$ of alternatives with cardinal $A = |\mathcal{A}|$. We assume that comparisons r_{ab} between alternatives a and b have been made, for a (potentially small) subset of pairs of alternatives. A positive comparison $r_{ab} > 0$ means that b is judged better than a, and increases when the judgment is more pronounced. We assume that $r_{ab} = -r_{ba}$ (i.e. a beats b if and only if b is beaten by a).

We denote by $\mathcal{C}$ the set of pairs $(ab) \in \mathcal{A}^2$ that have been compared, and by $C = |\mathcal{C}| \geq 1$ its cardinal. For $a \in \mathcal{A}$, let $\mathcal{A}_a = \{b \in \mathcal{A},\ (ab) \in \mathcal{C}\} \subset \mathcal{A}$ be the set of alternatives that have been compared with a and A_a its cardinal. Then, $C \leq A(A-1)$, where the equality corresponds to the case where all the alternatives have been compared. If so, we will have that $\mathcal{A}_a = \mathcal{A}\backslash\{a\}$ and $A_a = A-1$ for any $a \in \mathcal{A}$.

The comparisons are then characterized by an antisymmetric *comparison matrix*[2]

$$\mathrm{R} = (r_{ab})_{(ab)\in\mathcal{C}}. \tag{1}$$

Our goal is to attribute a score θ_a to each alternative $a \in \mathcal{A}$ from the paired comparison data R, i.e. to construct a function $\Theta^* : \mathrm{R} \mapsto \Theta^*(\mathrm{R}) \in \mathbb{R}^A$ that yields desirable computational, monotonicity and resilience properties.

GBT Models as an Exponential Subfamily

To infer scores from comparisons, as is often done in probabilistic models, we first define a distribution of comparisons given scores. Our proposal is to focus on a particular subset of the widely studied *exponential family* (Barndorff-Nielsen 2014). Recall that this family is of the form

$$p(r|\theta) = f(r)g(\theta)\exp(h(\theta)k(r)) \tag{2}$$

for some functions f, g, h, k, where r is a random variable whose law is parameterized by θ. Essentially, we define GBT models as the subfamily of such models where $h(\theta) = \theta$ and $k(r) = r$. In fact, since $g(\theta)$ is merely a normalization factor, this amounts to defining $p(r|\theta) \propto f(r)\mathrm{e}^{r\theta}$. Since the

[2]The comparison matrix is not a matrix in the classical sense, since its entries r_{ab} are only defined for $(ab) \in \mathcal{C}$ and not $(ab) \in \mathcal{A}^2$ in general.

distribution $p(r|\theta)$ is fully determined by f, assuming it is normalized so that $\int_{\mathbb{R}} f(r)\mathrm{d}r = 1 = \int_{\mathbb{R}} \mathrm{d}F(r)$, we call f the *root law* of the GBT model, and F is the associated cdf. More precisely, we define GBT models as follows.

Definition 1. *For any probability law f over $\mathbb{R}$ with finite exponential moments (i.e. $\int_{\mathbb{R}} e^{r\theta} dF(r) < \infty$ for any $\theta \in \mathbb{R}$), we define the* f-GBT model *as follows. For any hidden scores $\theta \in \mathbb{R}^A$, the comparisons R_{ab} given θ are assumed to be independent and each only depends on the score difference $\theta_{ab} \triangleq \theta_a - \theta_b$, with $p(r_{ab}|\Theta) \propto f(r_{ab})e^{r_{ab}\theta_{ab}}$.*

Equivalently, this corresponds to defining, for any $(ab) \in \mathcal{C}$ and any $\Theta \in \mathbb{R}^A$,

$$p(r_{ab}|\Theta) = \frac{\mathrm{e}^{\theta_{ab}r_{ab}} f(r_{ab})}{\int_{\mathbb{R}} \mathrm{e}^{\theta_{ab}r} \mathrm{d}F(r)}, \tag{3}$$

which is well-defined when the *root law* f has finite exponential moments. Note that, assuming $\theta_a = \theta_b$, the (normalized) *root law* f is then exactly the probability distribution of a comparison r_{ab}. In other words, f describes the distribution of comparisons between alternatives of similar quality. This makes it natural to expect f to be symmetric with respect to zero, which implies that its support $\textsc{Supp}(f)$ is as well. Using the independence of the comparisons conditionally to the scores, we easily deduce the following result.

Proposition 1. *Under the f-GBT model, the comparisons r_{ab} are independent conditionally to $\Theta \in \mathbb{R}^A$ and $r_{ab}|\Theta = r_{ab}|\theta_{ab}$. Moreover, we have, for any (R, Θ)*

$$p(\mathrm{R}|\Theta) = \prod_{(ab)\in\mathcal{C}} p(r_{ab}|\theta_{ab}) = \prod_{(ab)\in\mathcal{C}} \frac{f(r_{ab})\,\mathrm{e}^{r_{ab}\theta_{ab}}}{\int_{\mathbb{R}} \mathrm{e}^{\theta_{ab}r} \mathrm{d}F(r)} \tag{4}$$

Cumulant-Generating Functions

Let us now introduce the cumulant-generating function (Kenney and Keeping 1951) derived from the *root law* f, which will be central to our analysis of f-GBT models.

Definition 2. *Let f be a probability law over $\mathbb{R}$ with finite exponential moment. Its* cumulant generating function *is defined for any $\theta \in \mathbb{R}$ by*

$$\Phi_f(\theta) = \log\left(\int_{\mathbb{R}} \mathrm{e}^{\theta r} \mathrm{d}F(r)\right), \tag{5}$$

where we recall that F is the cdf of f.

The cumulant generating function of f, also called the log-partition function (Wainwright, Jaakkola, and Willsky 2005), will play an important role for GBT models. By extension, we say that Φ_f is the cumulant generating function of the f-GBT model.

The Taylor series expansion of Φ_f provides the cumulants of f. The function Φ is known for many classical probability laws[3] and has been extensively studied, in particular in large deviations theory (Dembo and Zeitouni 2009). We recall some of its main properties in Theorem 1, whose proof is given in the Appendix for the sake of completeness. We recall that $\textsc{Supp}(f)$ is the support of f and we denote by $r_{max} = \sup \textsc{Supp}(f) \in (0, \infty]$.

[3]https://en.wikipedia.org/wiki/Moment-generating_function\#Examples

Theorem 1 ((Dembo and Zeitouni 2009))**.** *For any root law f with finite exponential moments, the cumulant-generating function Φ_f is non-negative, strictly convex, even, and infinitely smooth over $\mathbb{R}$. Its derivative Φ'_f is a strictly increasing odd bijection from $\mathbb{R}$ to its image, which is $(-r_{\max}, r_{\max})$ as soon as $r_{\max} < \infty$.*

As we will see in the next section, the cumulant-generating function Φ_f yields numerous key statistics of the maximum-a-posteriori score estimator.

Discrete and Continuous GBT Models

For concreteness, we distinguish two types of GBT models, depending on the discreteness of the comparisons' domain. As we will see, the discreteness of f defines the discreteness of f-GBT.

Discrete GBT models. Let the root law f be of the form

$$f(r) = \sum_{k\in\mathcal{K}} \mathbb{P}[r_k|0]\delta_{r_k}(r), \tag{6}$$

where δ_x is the Dirac distribution on x and $\mathcal{K}$ is a countable (possibly finite) subset of $\mathbb{R}$ with no accumulation point. The discrete root law has an associated probability mass function $\mathbb{P}[r|0]$. The random variable $\mathrm{R}|\Theta$ is then also discrete with identical support $\textsc{Supp}(f)$. According to (4), its probability mass function $P(\mathrm{R}|\Theta)$ is given by

$$P(\mathrm{R}|\Theta) = \prod_{(ab)\in\mathcal{C}} \frac{\mathbb{P}[r_{ab}|0]\mathrm{e}^{r_{ab}\theta_{ab}}}{\sum_{r\in\textsc{Supp}(f)} \mathbb{P}[r|0]\mathrm{e}^{r\theta_{ab}}}. \tag{7}$$

Continuous GBT models. In the continuous setting, the comparison matrix $\mathrm{R}|\Theta$ also admits a probability density function given by

$$p(\mathrm{R}|\Theta) = \prod_{(ab)\in\mathcal{C}} \frac{f(r_{ab})\mathrm{e}^{r_{ab}\theta_{ab}}}{\int_{\textsc{Supp}(f)} \mathrm{e}^{r\theta_{ab}} f(r)\mathrm{d}r}. \tag{8}$$

Historical Bradley-Terry as Bernoulli-GBT

Recall that the classical (Bradley and Terry 1952) model is characterized by the fact that comparisons follows the following distribution:[4]

$$P(r_{ab} = 1|\theta_a - \theta_b) = \frac{\mathrm{e}^{\theta_a}}{\mathrm{e}^{\theta_a} + \mathrm{e}^{\theta_b}} \tag{9}$$

and $P(r_{ab} = 1|\theta_a, \theta_b) + P(r_{ab} = -1|\theta_a, \theta_b) = 1$.

Here, we highlight the fact that this is an instance of the discrete GBT models. Namely, consider the Bernouilli *root law* $f = \frac{\delta_1 + \delta_{-1}}{2}$. Its cumulant-generating function is

$$\Phi_{\text{Bernoulli}}(\theta) = \log\left(\frac{\mathrm{e}^{\theta} + \mathrm{e}^{-\theta}}{2}\right) = \log(\cosh\theta). \tag{10}$$

Using (3) and $\theta_{ab} = \theta_a - \theta_b$, we deduce

$$P(r_{ab} = 1|\Theta) = \frac{\mathrm{e}^{\theta_{ab}}}{2\cosh(\theta_{ab})} = \frac{\mathrm{e}^{\theta_a}}{\mathrm{e}^{\theta_a} + \mathrm{e}^{\theta_b}}. \tag{11}$$

We recover (9), hence the fact that (Bradley and Terry 1952) is the Bernoulli-GBT model.

[4]This is one possible parametrization, using an exponential score function. An alternative is to consider that $P(a > b) = \frac{p_a}{p_a + p_b}$ with $p_a > 0$ the score value of a.

MAP Scores Based on GBT Models

In this section, we take a Bayesian approach to define a maximum-a-posteriori estimator based on GBT models, and prove some of the basic resulting properties.

Bayesian Model for Score Estimation

Consider a normal prior $\mathcal{N}(0, \sigma^2 I)$ on Θ. In the continuous setting, the posterior density function of Θ conditionally to the comparisons R is, using Bayes law,

$$p(\Theta|\mathrm{R}) = \frac{p(\mathrm{R}|\Theta)p(\Theta)}{p(\mathrm{R})}. \tag{12}$$

In the discrete setting, probability density functions are replaced by probability mass functions.

Remark. Using (8) in (12), we can interpret the GBT model with a Gaussian prior (or a prior from any law of the exponential family) as an exponential family for the random variable Θ, where R is the natural parameter.

Maximum A Posteriori Estimator

We define the *negative log-posterior*, either for discrete or continuous comparisons, as

$$\mathcal{L}(\Theta|\mathrm{R}) = -\log p(\Theta|\mathrm{R}). \tag{13}$$

Interestingly, this yields a loss function that directly depends on the cumulant-generating function, as

$$\mathcal{L}(\Theta|\mathrm{R}) = \frac{1}{2\sigma^2}\sum_{a\in\mathcal{A}} \theta_a^2 + \sum_{(ab)\in\mathcal{C}} \left(\Phi_f(\theta_{ab}) - r_{ab}\theta_{ab}\right). \tag{14}$$

In particular, the nice properties of Φ_f and the fact that the Gaussian prior is turned into a quadratic regularization imply that the negative log-posterior $\mathcal{L}$ is well-behaved.

Proposition 2. *For any comparison matrix* R*, the negative log-posterior* $\mathcal{L}(\cdot|\mathrm{R})$ *is* $(1/\sigma^2)$*-strongly convex, and thus admits a unique minimizer* $\Theta^*_{f,\sigma^2}(\mathrm{R}) \in \mathbb{R}^A$.

Since it is the mode of the posterior, $\Theta^*_{f,\sigma^2}(\mathrm{R})$ is commonly known as the *maximum a posteriori (MAP) estimator*. Crucially, for any f-GBT model with a normal prior, Proposition 2 guarantees that the MAP is well-defined and fastly computable by any strongly convex optimizer. In fact, it is used in Tournesol to estimate individual scores from user comparisons (Anonymous 2023).

First Properties of Score Estimators

In the following, we show that the maximum likelihood score estimator $\Theta^*_{f,\sigma^2}(\mathrm{R})$ has zero mean in Proposition 3, we provide its first two moments in Proposition 4, and we give a bound on its supremum norm in Proposition 5. The proofs are in the appendix.

Proposition 3. *For any comparison matrix* R*, the MAP estimator* $\theta^* = \Theta^*_{f,\sigma^2}(\mathrm{R})$ *verifies* $\sum_{a\in\mathcal{A}} \theta^*_a = 0$.

We denote by $\mathbb{E}[g(\mathrm{R})|\Theta]$ and $\mathbb{V}[g(\mathrm{R})|\Theta]$ the mean and the variance of $g(\mathrm{R})$ conditionally to the scores. Also, denote by $\mathbb{C}\mathrm{ov}[g(\mathrm{R}), h(\mathrm{R})|\Theta]$ the covariance of $g(\mathrm{R})|\Theta$ and $h(\mathrm{R})|\Theta$.

Proposition 4. *Let* $(ab), (cd) \in \mathcal{C}$ *with* $(cd) \notin \{(ab), (ba)\}$*. Then,*

$$\mathbb{E}[r_{ab}|\Theta] = \Phi_f'(\theta_{ab}) \quad \textit{and} \tag{15}$$
$$\mathbb{V}[r_{ab}|\Theta] = \Phi_f''(\theta_{ab}), \quad \mathbb{C}\mathrm{ov}[r_{ab}, r_{cd}|\Theta] = 0.$$

Proposition 5. *If* $r_{\max} < \infty$*, then the MAP estimator* $\Theta^*_{f,\sigma^2}(\mathrm{R})$ *verifies, for any* $a \in \mathcal{A}$,

$$|\theta^*_a(\mathrm{R})| \leq 2A_a r_{\max}\sigma^2, \tag{16}$$

and therefore

$$\|\Theta^*_{f,\sigma^2}(\mathrm{R})\|_\infty \leq 2r_{\max}\sigma^2 \sup_{a\in\mathcal{A}} A_a \leq 2r_{\max}\sigma^2(A-1). \tag{17}$$

Monotonicity of MAP Score Estimators

In this section, we prove a desirable property of MAP estimators for all GBT models with a Gaussian prior. Namely, we show that the more an alternative wins comparisons, the better it is scored.

Partial Orders over Paired Comparisons

Let us first formalize a partial order between comparisons associated to a given alternative, which captures the idea that a is better compared when all the other comparisons between alternatives different from a are fixed. Note that this partial order is only defined for comparison matrices R and R' sharing the same set of compared pairs $\mathcal{C}$.

Definition 3. *For any alternative* $a \in \mathcal{A}$*, we say that two comparison matrices* R *and* R' *satisfy*

$$\mathrm{R} \leq_a \mathrm{R}' \tag{18}$$

if (i) $r_{ab} \leq r'_{ab}$ *for all* b *such that* $(ab) \in \mathcal{C}$ *and (ii)* $r_{cd} = r'_{cd}$ *for* $(cd) \in \mathcal{C}$ *with* $a \notin \{c, d\}$*. The relation*

$$\mathrm{R} <_a \mathrm{R}' \tag{19}$$

means that (i) is strict for at least one b *such that* $(ab) \in \mathcal{C}$.

We formalize the notion that a score estimator is consistent with the partial orders of Definition 3.

Definition 4. *We say that an estimator* $\widehat{\Theta}(\mathrm{R})$ *is* increasing *with respect to the* R *if, for any* $a \in \mathcal{A}$,

$$\mathrm{R} \leq_a \mathrm{R}' \Longrightarrow \widehat{\theta}_a(\mathrm{R}) \leq \widehat{\theta}_a(\mathrm{R}'), \tag{20}$$

and strictly increasing *with respect to* R *if, for any* $a \in \mathcal{A}$,

$$\mathrm{R} <_a \mathrm{R}' \Longrightarrow \widehat{\theta}_a(\mathrm{R}) < \widehat{\theta}_a(\mathrm{R}'). \tag{21}$$

Elementary Monotonicity Criteria

We provide criteria for the monotonicity of score estimators for continuous and discrete comparisons, in Proposition 6 and Proposition 7 that are proved in the Appendix.

Proposition 6. *We suppose that* R *is continuous-domain and that the estimator* $\widehat{\Theta}(\mathrm{R})$ *is differentiable with respect to* r_{ab} *for any* $(ab) \in \mathcal{C}$*. Then,* $\widehat{\Theta}(\mathrm{R})$ *is increasing with respect to* R *if and only if, for any* $(ab) \in \mathcal{C}$ *and any* R,

$$\partial_{r_{ab}} \widehat{\theta}_a(\mathrm{R}) \geq 0. \tag{22}$$

It is moreover strictly increasing if and only if the inequalities are strict.

For discrete comparisons, we define finite-difference operators over scores as follows. For any R and $(ab) \in \mathcal{C}$, we define R'_{ab} as the comparison with identical scores, except at position (ab) where the score is increased from r_k to r_{k+1} (or remains unchanged if r_K is reached) and reduces symmetrically the comparison at position (ba). Then, we define the operator $\Delta_{(ab)}$ over functions $F : \textsc{Supp}(f)^C \to \mathbb{R}$,

$$\Delta_{(ab)}F(\mathrm{R}) = \frac{F(\mathrm{R}^{\mathrm{up}}_{(ab)}) - F(\mathrm{R})}{r_{k+1} - r_k} \tag{23}$$

if $r_k < r_K$ and $\Delta_{(ab)}F(\mathrm{R}) = 0$ if $r_k = r_K$.

Proposition 7. *We suppose that* R *is discrete-domain. Then,* $\widehat{\Theta}(\mathrm{R})$ *is increasing with respect to* R *if and only if, for any* $(ab) \in \mathcal{C}$ *and any* R,

$$\Delta_{(ab)}\widehat{\theta}_a(\mathrm{R}) \geq 0. \tag{24}$$

It is moreover strictly increasing if and only if the inequalities are strict.

Monotonicity of GBT Estimators

We show that the monotonicity of GBT estimators is automatically satisfied. The proof is in the Appendix.

Theorem 2. *For any f-GBT model with a Gaussian prior, the MAP estimator* $\Theta^*(\mathrm{R})$ *is strictly increasing with* R *in the sense of Definition 4.*

The proof relies on the monotonicity criteria for continuous (Proposition 6) and discrete (Proposition 7) GBT models. In order to show that $\partial_{r_{ab}}\widehat{\theta}_a(\mathrm{R}) \geq 0$, we analyze the gradient relation $\nabla\mathcal{L}(\Theta^*_{f,\sigma^2}(\mathrm{R})|\Theta) = 0$. By applying $\partial_{\theta_{ab}}$ to it, we obtain a linear system on $\partial_{\theta_{ab}}\Theta^*_{f,\sigma^2}(\mathrm{R})$. The study of this linear system relies on the properties of diagonally-dominant matrices and leads to the desired result.

Note that our proof yields a slightly more general result, as the monotonicity actually holds for all coordinate-independent priors (i.e. θ_a is a priori independent from θ_b for $a \neq b$) which yield a strongly convex negative log-prior. In fact, it can be extended to convex negative log-prior, or even to the maximum likelihood estimator (i.e. no prior), if we consider inferred scores with values in $[-\infty, +\infty]$.

Impact of New Comparisons

Assume that R and R' are two comparison matrices over $\mathcal{C}$ and $\mathcal{C}'$ respectively, such that

$$\mathcal{C}' = \mathcal{C} \cup \{(ab), (ba)\} \tag{25}$$

where $(ab) \notin \mathcal{C}$. We assume that $r_{cd} = r'_{cd}$ for any $(cd) \in \mathcal{C}$, hence $\mathcal{C}$ and $\mathcal{C}'$ only differs by the new comparison. We evaluate the impact of this comparison r'_{ab} over the score vector $\Theta^*_{f,\sigma^2}(\mathrm{R}')$.

Proposition 8. *For* R *and* R' *as defined above, we have the equivalences*

$$\Theta^*_{f,\sigma^2}(\mathrm{R}) = \Theta^*_{f,\sigma^2}(\mathrm{R}') \iff r'_{ab} = \Phi'_f(\theta^*_{ab}(\mathrm{R})), \tag{26}$$

$$\Theta^*_{f,\sigma^2}(\mathrm{R}) < \Theta^*_{f,\sigma^2}(\mathrm{R}') \iff r'_{ab} > \Phi'_f(\theta^*_{ab}(\mathrm{R})), \tag{27}$$

$$\Theta^*_{f,\sigma^2}(\mathrm{R}) > \Theta^*_{f,\sigma^2}(\mathrm{R}') \iff r'_{ab} < \Phi'_f(\theta^*_{ab}(\mathrm{R})). \tag{28}$$

The proof is provided in the Appendix. Now, Φ'_f can be interpreted as a conversion function between score differences θ_{ab} and comparisons r_{ab}. This is highlighted by the relation (26) and the fact that $\mathbb{E}[r_{ab}|\Theta] = \Phi'_f(\theta_{ab})$ (Proposition 4). When $r_{\max} = 1$, we observe that Φ'_f is a sigmoid function, i.e. an increasing bijection from $\mathbb{R}$ to $(-1, 1)$.

Lipschitz-Resilience of Score MAP Estimators

Among the motivation for the generalized Bradley-Terry model, we aim at a scoring method from expressed pair comparisons which controlled impact of user's decision. In the more global Tournesol pipeline (Hoang et al. 2021), the individual scoring method is used at a first step in a global scoring method for global scoring from individual comparisons (Allouah et al. 2022). We formalize the notion of Lipschitz-resilience to the user's updates and provide criteria to determine if a given GBT model is resilient. We show that the Lipschitz-resilience is guaranteed as soon as the comparison domain $\textsc{Supp}(f)$ is bounded.

Lipschitz-Resilience to User Modifications

The Lipschitz-resilience of an estimator captures its ability to be limitedly modified by changing or adding new paired comparisons for the user.

Let $\mathrm{R} \in \textsc{Supp}(f)^C$ and $\mathrm{R}' \in \textsc{Supp}(f)^{C'}$ be two comparison matrices over some possibly distinct $\mathcal{C}$ and $\mathcal{C}'$ of respective size C and C'. We define the *symmetric difference* of $\mathcal{C}$ and $\mathcal{C}'$ as $\mathcal{C}\Delta\mathcal{C}' = (\mathcal{C}\backslash\mathcal{C}') \cup (\mathcal{C}\backslash\mathcal{C}')$ and denote its cardinal by

$$\Delta_{\mathrm{domain}}(\mathrm{R}, \mathrm{R}') = |\mathcal{C}\Delta\mathcal{C}'|. \tag{29}$$

The set $\mathcal{C}\Delta\mathcal{C}'$ is made of pairs (ab) that are in one of the two set $\mathcal{C}$ and $\mathcal{C}'$ and not on the other. The number $\Delta_{\mathrm{domain}}(\mathrm{R}, \mathrm{R}')$ therefore quantifies the number of comparisons needed to be added or removed to transform $\mathcal{C}$ into $\mathcal{C}'$.

We define the matrices $\tilde{\mathrm{R}} = (r_{ab})_{(ab)\in\mathcal{C}\cap\mathcal{C}'}$ and $\tilde{\mathrm{R}}' = (r'_{ab})_{(ab)\in\mathcal{C}\cap\mathcal{C}'}$, both in $\textsc{Supp}(f)^{|\mathcal{C}\cap\mathcal{C}'|}$, which coincide with R and R' on $\mathcal{C}\cap\mathcal{C}'$. We recall that the $L0$ *"norm"* of a vector is the number of its non-zero entries. We then define

$$\Delta_{\mathrm{entries}}(\mathrm{R}, \mathrm{R}') = \|\tilde{\mathrm{R}} - \tilde{\mathrm{R}}'\|_0, \tag{30}$$

which measures the number of entries $(ab) \in \mathcal{C} \cap \mathcal{C}'$ on which R and R' differ. We also set

$$\Delta(\mathrm{R}, \mathrm{R}') = \Delta_{\mathrm{domain}}(\mathrm{R}, \mathrm{R}') + \Delta_{\mathrm{entries}}(\mathrm{R}, \mathrm{R}'), \tag{31}$$

which counts the number of elementary modifications (removing, adding, or changing a comparison) from R to R'.

Definition 5. *An estimator* $\widehat{\Theta}(\mathrm{R})$ *is said to be* L-Lipschitz-resilient *for some* $L > 0$ *for the Euclidean norm if, for any comparison matrices* R, R',

$$\|\widehat{\Theta}(\mathrm{R}) - \widehat{\Theta}(\mathrm{R}')\|_2 \leq L\Delta(\mathrm{R}, \mathrm{R}') \tag{32}$$

If $\widehat{\Theta}$ is L-Lipschitz-resilient, then L bounds the possible impact on the score $\widehat{\Theta}(\mathrm{R})$ for single modifications of the comparison matrix (single update over one comparison or addition of a new comparison). This impact is measured in

GBT model	Parameter	$F/f(r\|0)$	$\textsc{Supp}(f)$	$\Phi(\theta)$
Binary	-	$\frac{1}{2}$	$\{\pm 1\}$	$\log(\cosh\theta)$
K-nary	$K \geq 2$	$\frac{1}{K}$	$\{-1, \dots, +1\}$	$\log\left(\frac{\sinh\left(\frac{K\theta}{K-1}\right)}{K\sinh\left(\frac{\theta}{K-1}\right)}\right)$
Poisson	$\lambda > 0$	$\frac{e^{-\lambda}\lambda^{\|k\|}}{\|k\|!}(\delta_{k=0} + \frac{\delta_{k\neq 0}}{2})$	$\mathbb{Z}$	$\lambda\cosh(\theta)$
Gaussian	$\sigma_0^2 > 0$	$g_{\sigma_0^2}(r) = \frac{1}{\sqrt{2\pi}\sigma_0}e^{-r^2/2\sigma_0^2}$	$\mathbb{R}$	$\frac{\sigma_0^2\theta^2}{2}$
Beta	$\beta > 0$	$\frac{\Gamma(2\beta)}{4^\beta\Gamma(\beta)^2}(1-r^2)^{\beta-1}$	$(-1, 1)$	$\log\left(1 + \sum_{k\geq 1}\left(\prod_{n=0}^{k}\frac{\beta+n}{2\beta+n}\right)\frac{\theta^{2k}}{(2k)!}\right)$
Uniform	$\beta = 1$	$\frac{1}{2}$	$(-1, 1)$	$\log\left(\frac{\sinh\theta}{\theta}\right)$
	$\beta = 2$	$\frac{3}{4}(1-r^2)$	$(-1, 1)$	$\log\left(\frac{3(\theta\cosh\theta - \sinh\theta))}{\theta^3}\right)$

Table 1: Examples of GBT models, with their cumulant-generating functions.

terms of ℓ_2-norm. When the comparison sets $\mathcal{C} = \mathcal{C}'$ coincide[5], the bound (32) is simply

$$\|\widehat{\Theta}(\mathrm{R}) - \widehat{\Theta}(\mathrm{R}')\|_2 \leq L\|\mathrm{R} - \mathrm{R}'\|_0. \tag{33}$$

Lipschitz-Resilience Guarantee for GBT Models

Theorem 3. *For any root law f and given a Gaussian prior $\mathcal{N}(0, \sigma^2 I)$, the MAP estimator for the f-GBT model is $(4\sqrt{2}r_{\max}\sigma^2)$-Lipschitz-resilient, i.e.*

$$\sup_{\mathrm{R}\neq\mathrm{R}'} \frac{\|\Theta^*_{f,\sigma^2}(\mathrm{R}) - \Theta^*_{f,\sigma^2}(\mathrm{R}')\|_2}{\Delta(\mathrm{R}, \mathrm{R}')} \leq 4\sqrt{2}r_{\max}\sigma^2. \tag{34}$$

In the GBT model, contrary to the historical BT model, we regularize the scores using a Gaussian prior on Θ. We see in Theorem 3 that the Lipschitz-resilience constant explodes when $\sigma^2 \to \infty$ (i.e. with no regularization). The regularization leads to controllable user's modifications, where the prior variance σ^2 plays a crucial role. Moreover, the Gaussian and Poisson BT models (see examples below), for which the comparison domain $\textsc{Supp}(f) = \mathbb{R}$ or $\mathbb{Z}$ is unbounded, are not Lipschitz-resilient (we provide a proof in the Appendix). The boundedness of the comparisons is a key ingredient to the Lipschitz-resilience.

Examples of GBT Models

The examples detailed in this section are listed in Table 1, together with their corresponding comparison domain and their cumulant-generating function. Each of these models can be used to provide score estimators based on paired comparisons that are all strictly increasing with respect to R according to Theorem 2. They are moreover all resilient for bounded comparisons, while the Gaussian and Poisson-GBT model are not (see Proposition 9 below and the appendix).

The Gaussian-GBT Model

The Gaussian GBT model is characterized by a Gaussian root law $f = g_{\sigma_0^2}$. This model has already been studied (Guo et al. 2012) and applied (Kristof et al. 2019). We summarize its main properties in Proposition 9. The proof and the closed form expression of $\Theta^*(\mathrm{R})$ are given in the appendix.

[5]This ensures that $\mathrm{R} - \mathrm{R}'$ is well-defined.

Proposition 9. *The Gaussian-GBT model with variance σ_0^2 is such that*

$$\mathrm{R}|\Theta \sim \mathcal{N}\left((\sigma_0\theta_{ab})_{(ab)\in\mathcal{C}}, \sigma_0^2\mathrm{Id}\right). \tag{35}$$

The MAP estimator $\Theta^*(\mathrm{R})$ is linear, strictly increasing with respect to R*, and is* not *Lipschitz-resilient.*

The Uniform-GBT Model

The Uniform-GBT model corresponds to choosing the uniform probability law $f = 1/2$ on $[-1, 1]$ as the root law of the model. We recall that the prior variance is denoted by σ^2. The Uniform-GBT model is used for the current version of the Tournesol pipeline (Anonymous 2023).

Proposition 10. *The cumulant generating function of the Uniform-GBT model is $\Phi(\theta) = \log(\sinh(\theta)/\theta)$. The MAP estimator $\Theta^*(\mathrm{R})$ is strictly increasing with respect to* R *and $4\sqrt{2}\sigma^2$-Lipschitz-resilient.*

Remark. The derivative $\Phi'(\theta) = \coth(\theta) - \frac{1}{\theta}$ of Φ is known as the Langevin function (Cohen 1991).

Empirical Simulations

We propose three experiments that illustrate interesting properties of GBT models. The generative data model is itself a GBT model, and does not simulate a realistic situation corresponding to real data. These simulations allow us to measure three aspects of the model: (i) the impact of the sparsity of the comparison graph, (i) the impact of the discretization level, and (iii) the impact of the regularization parameter (prior variance) on the quality of the reconstruction. The results are expressed in terms of the normalized mean-square error against the true scores $\Theta^\dagger$, given by

$$\textsc{NormError} = \mathbb{E}\left[\frac{\|\Theta(\mathrm{R}) - \Theta^\dagger\|_2^2}{\|\Theta^\dagger\|_2^2}\right] \tag{36}$$

We use Monte-Carlo simulations to obtain normalized mean-square errors for various GBT MAP estimators with Gaussian prior. Plots in Figure 1 depict the mean and standard deviation of the Monte-Carlo simulations. We consider $A = 500$ alternatives. The ground-truth scores $\Theta^\dagger \in \mathbb{R}^{500}$

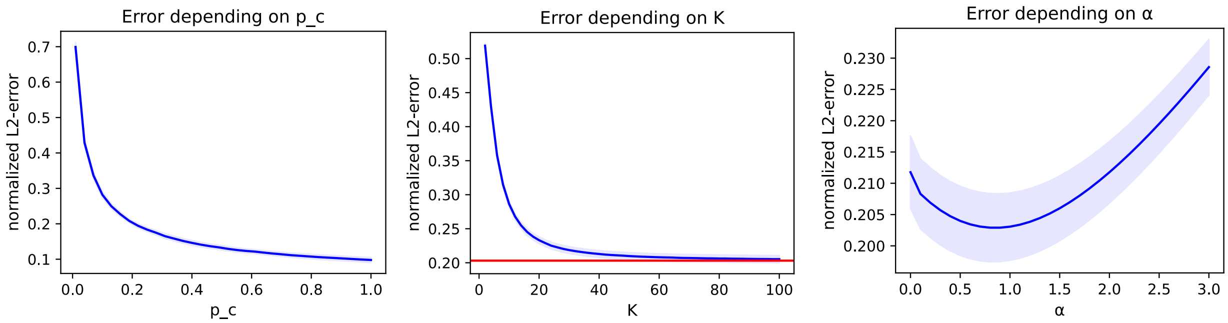

Figure 1: Left: Normalized mean-square error with respect to the sparsity parameter p_c for Erdös-Rényi comparison graphs. Middle: Normalized mean-square errors NormError_K using K-nary-GBT MAP estimators on the data generated via the Uniform-GBT model. Blue: $K \mapsto \text{NormError}_K$; Red: NormError for the Uniform-GBT model. Right: Normalized mean-square error with respect to the regularization scale $\frac{1}{\sigma^2}$.

are generated as i.i.d. Gaussian random variables with variance $\sigma^{\dagger 2} = 1$. All experiments are run with ten seeds from 1 to 10. The code and details of the experiments are available at `https://github.com/sadeghfarhadkhani/GBT`.

(i) Impact of the graph sparsity. We generate the comparisons $\mathrm{R}|\Theta^\dagger$ using the Uniform-GBT model over $[-1, 1]$. The comparison set $\mathcal{C}$ is generated as an Erdös-Rényi random graph $\mathcal{G}(A = 500, p_c)$ where the nodes are different alternatives and each edge corresponds to a comparison which is randomly activated independently from other edges with probability $p_c \in [0, 1]$. We estimate the normalized mean-square error of MAP estimators based on the Uniform-GBT model with variance $\sigma^2 = 1$ for different values of p_c. The results are depicted in Figure 1 (left). With no surprise, the sparsity of the graph strongly impacts the reconstruction performance.

(ii) K-nary-GBT models against Uniform-GBT model. The data are generated using the Uniform-GBT model on a comparison graph following the Erdös-Rényi random graph $\mathcal{G}(A = 500, p_c = 0.2)$. For any integer $K \geq 2$, we estimate $\Theta_K(\mathrm{R})$ as the MAP estimator of the K-nary-GBT model with variance $\sigma^2 = 1$. We also compute $\Theta(\mathrm{R})$ as the MAP estimator of the Uniform-GBT model with the same variance. We show the evolution of the normalized mean-square error with respect to K in Figure 1 (middle). The error decays with respect to K, and converges to the limit value corresponding to the Uniform-GBT reconstruction model. These results advocate for discretized comparison models when discretization is at play. The value $K = 21$ was chosen on the Tournesol platform (Anonymous 2023), as a compromise between greater finesse on the comparisons and a restricted choice of possible comparisons for users.

(iii) Impact of the prior variance over the mean-square error. Recall from (14) that the regularization factor is scaled by $\frac{1}{\sigma^2}$ where σ^2 is the variance of the prior for the score estimation model. Again, we use the Uniform-GBT model on a comparison graph following the Erdös-Rényi random graph $\mathcal{G}(A = 500, p_c = 0.2)$, to generate the data. We then estimate the normalized mean-square error of MAP estimators based on the Uniform-GBT model for different values of σ. The results are given in Figure 1 (right). We observe that the error for small $\frac{1}{\sigma^2}$ (i.e., little regularization) is slightly better than for $\sigma = \infty$ ($\frac{1}{\sigma^2} = 0$, i.e., no regularization), which is in favor of using a Gaussian Bayesian prior for the scores. This is a second advantage of the Bayesian approach in addition to Lipschitz-resilience. A natural conjecture that our plot raises, which we leave open, is whether the optimal prior variance for score inference matches the variance of the score distribution, i.e. $\sigma = \sigma^\dagger$.

Conclusion and Future Perspective

In this paper, we generalized the historical Bradley-Terry model, by defining a family of so-called GBT (generalized Bradley-Terry) probabilistic models, each of which transforms comparisons into scores. This family is parameterized by a *root law*, which models how two equally good alternatives are expected to be compared. Remarkably, for any such prior, we proved that the derived GBT model is guaranteed to feature numerous desirable properties, such as strict convexity, monotonicity and Lipschitz-resilience. Because of these compelling features, GBT models seem desirable to deploy in many practical applications, as they have been on Tournesol (Hoang et al. 2021) to turn contributors' comparisons of video recommendability into scores.

Having said this, GBT models raise further intriguing research questions. For one thing, there may be other appealing models that also have the convexity, monotonicity and Lipschitz-resilience properties. We leave open the problem of identifying the set of all models with such properties. We also leave open the analysis of the statistical error of the GBT scores. A related question is that of estimating the uncertainty on the GBT scores, given comparisons. We stress that these analyses are challenging as they depend on the graph $\mathcal{C}$ of comparisons. In fact, another related question is that of optimizing the graph to minimize the statistical error, which is known as the *active learning* problem. Finally, a general analysis of how the GBT scores depend on the root law is also left open.

References

Allouah, Y.; Guerraoui, R.; Hoang, L.; and Villemaud, O. 2022. Robust Sparse Voting. *CoRR*, abs/2202.08656.

Anonymous. 2023. Permissionless Collaborative Algorithmic Governance with Security Guarantees. *Under submission*.

Barndorff-Nielsen, O. 2014. *Information and exponential families: in statistical theory*. John Wiley & Sons.

Bradley, R. A.; and Terry, M. E. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4): 324–345.

Cattelan, M. 2012. Models for paired comparison data: A review with emphasis on dependent data. *Statistical Science*, 27(3): 412–433.

Cohen, A. 1991. A Padé approximant to the inverse Langevin function. *Rheologica acta*, 30: 270–273.

David, H. A. 1963. *The method of paired comparisons*, volume 12. London.

Davidson, R. R. 1970. On extending the Bradley-Terry model to accommodate ties in paired comparison experiments. *Journal of the American Statistical Association*, 65(329): 317–328.

Dembo, A.; and Zeitouni, O. 2009. *Large deviations techniques and applications*, volume 38. Springer Science & Business Media.

Durrett, R. 2019. *Probability: theory and examples*, volume 49. Cambridge university press.

Glickman, M. E. 1999. Parameter estimation in large dynamic paired comparison experiments. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 48(3): 377–394.

Guo, S.; Sanner, S.; Graepel, T.; and Buntine, W. 2012. Score-based Bayesian skill learning. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2012, Bristol, UK, September 24-28, 2012. Proceedings, Part I 23*, 106–121. Springer.

Herbrich, R.; Minka, T.; and Graepel, T. 2006. TrueSkill™: a Bayesian skill rating system. *Advances in neural information processing systems*, 19.

Hoang, L.; Faucon, L.; Jungo, A.; Volodin, S.; Papuc, D.; Liossatos, O.; Crulis, B.; Tighanimine, M.; Constantin, I.; Kucherenko, A.; Maurer, A.; Grimberg, F.; Nitu, V.; Vossen, C.; Rouault, S.; and El-Mhamdi, E. 2021. Tournesol: A quest for a large, secure and trustworthy database of reliable human judgments. *CoRR*, abs/2107.07334.

Horn, R. A.; and Johnson, C. R. 2012. *Matrix analysis*. Cambridge university press.

Hunter, D. R. 2004. MM algorithms for generalized Bradley-Terry models. *The annals of statistics*, 32(1): 384–406.

Karlis, D.; and Ntzoufras, I. 2009. Bayesian modelling of football outcomes: using the Skellam's distribution for the goal difference. *IMA Journal of Management Mathematics*, 20(2): 133–145.

Kenney, J.; and Keeping, E. 1951. Cumulants and the cumulant-generating function. *Mathematics of Statistics, Princeton, NJ*.

Kristof, V.; Quelquejay-Leclère, V.; Zbinden, R.; Maystre, L.; Grossglauser, M.; and Thiran, P. 2019. A User Study of Perceived Carbon Footprint. *arXiv preprint arXiv:1911.11658*.

Maher, M. J. 1982. Modelling association football scores. *Statistica Neerlandica*, 36(3): 109–118.

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*, 1236–1246.

Peña, J. M. 1995. M-matrices whose inverses are totally positive. *Linear algebra and its applications*, 221: 189–193.

Rao, P.; and Kupper, L. L. 1967. Ties in paired-comparison experiments: A generalization of the Bradley-Terry model. *Journal of the American Statistical Association*, 62(317): 194–204.

Thurstone, L. L. 1927. A law of comparative judgment. *Psychological review*, 34(4): 273.

Wainwright, M. J.; Jaakkola, T. S.; and Willsky, A. S. 2005. A new class of upper bounds on the log partition function. *IEEE Transactions on Information Theory*, 51(7): 2313–2335.

Zermelo, E. 1929. Die Berechnung der Turnier-Ergebnisse als ein Maximumproblem der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 30: 436–460.

Élö, A. E. 1978. *The Rating of Chessplayers, Past and Present*. New York: Arco Pub.