

Elo-MMR: A Rating System for Massive Multiplayer Competitions

Aram Eftekar
Vancouver, BC, Canada
aramebtech@gmail.com

Paul Liu
Stanford University
Stanford, CA, USA
paul.liu@stanford.edu

ABSTRACT

Skill estimation mechanisms, colloquially known as rating systems, play an important role in competitive sports and games. They provide a measure of player skill, which incentivizes competitive performances and enables balanced match-ups. In this paper, we present a novel Bayesian rating system for contests with many participants. It is widely applicable to competition formats with discrete ranked matches, such as online programming competitions, obstacle courses races, and video games. The system's simplicity allows us to prove theoretical bounds on its robustness and run-time. In addition, we show that it is *incentive-compatible*: a player who seeks to maximize their rating will never want to underperform. Experimentally, the rating system surpasses existing systems in prediction accuracy, and computes faster than existing systems by up to an order of magnitude.

CCS CONCEPTS

• Information systems → Learning to rank; • Computing methodologies → Learning in probabilistic graphical models.

KEYWORDS

rating system, skill estimation, mechanism design, competition, bayesian inference, robust, incentive-compatible, elo, glicko, trueskill

ACM Reference Format:

Aram Eftekar and Paul Liu. 2021. Elo-MMR: A Rating System for Massive Multiplayer Competitions. In *Proceedings of the Web Conference 2021 (WWW '21)*, April 19–23, 2021, Ljubljana, Slovenia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3442381.3450091>

1 INTRODUCTION

Competitions, in the form of sports, games, and examinations, have been with us since antiquity. Many competitions grade performances along a numerical scale, such as a score on a test or a completion time in a race. In the case of a college admissions exam or a track race, scores are standardized so that a given score on two different occasions carries the same meaning. However, in events that feature novelty, subjectivity, or close interaction, standardization is difficult. The Spartan Races, completed by millions of runners, feature a variety of obstacles placed on hiking trails around the world [10]. Rock climbing, a sport to be added to the 2020 Olympics, likewise has routes set specifically for each competition. DanceSport, gymnastics, and figure skating competitions have a panel of

judges who rank contestants against one another; these subjective scores are known to be noisy [31]. In all these cases, scores can only be used to compare and rank participants at the same event. Players, spectators, and contest organizers who are interested in comparing players' skill levels across different competitions will need to aggregate the entire history of such rankings. A strong player, then, is one who consistently wins against weaker players. To quantify skill, we need a **rating system**.

Good rating systems are difficult to create, as they must balance several mutually constraining objectives. First and foremost, rating systems must be accurate, in that ratings provide useful predictors of contest outcomes. Second, the ratings must be efficient to compute: within video game applications, rating systems are predominantly used for matchmaking in massively multiplayer online games (such as Halo, CounterStrike, League of Legends, etc.) [24, 28, 35]. These games have hundreds of millions of players playing tens of millions of games per day, necessitating certain latency and memory requirements for the rating system [11]. Third, rating systems must be **incentive-compatible**: a player's rating should never increase had they scored worse, and never decrease had they scored better. This is to prevent players from regretting a win, or from throwing matches to game the system. Rating systems that can be gamed often create disastrous consequences to player-base, potentially leading to the loss of players [3]. Finally, the ratings provided by the system must be human-interpretable: ratings are typically represented to players as a single number encapsulating their overall skill, and many players want to understand and predict how their performances affect their rating [20].

Classically, rating systems were designed for two-player games. The famous Elo system [17], as well as its Bayesian successors Glicko and Glicko-2, have been widely applied to games such as Chess and Go [20–22]. Both Glicko versions model each player's skill as a real random variable that evolves with time according to Brownian motion. Inference is done by entering these variables into the Bradley-Terry model [13], which predicts probabilities of game outcomes. Glicko-2 refines the Glicko system by adding a rating volatility parameter. Unfortunately, Glicko-2 is known to be flawed in practice, potentially incentivizing players to lose in what's known as "volatility farming". In some cases, these attacks can inflate a user's rating *several hundred points* above its natural value, producing ratings that are essentially impossible to beat via honest play. This was most notably exploited in the popular game of Pokemon Go [3]. See Section 5.1 for a discussion of this issue, as well as an application of this attack to the Topcoder rating system.

The family of Elo-like methods just described only utilize the binary outcome of a match. In settings where a scoring system provides a more fine-grained measure of match performance, Kovalchik [26] has shown variants of Elo that are able to take advantage of score information. For competitions consisting of several set

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '21, April 19–23, 2021, Ljubljana, Slovenia

© 2021 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-8312-7/21/04.

<https://doi.org/10.1145/3442381.3450091>

tasks, such as academic olympiads, Forišek [18] developed a model in which each task gives a different “response” to the player: the total response then predicts match outcomes. However, such systems are often highly application-dependent and hard to calibrate.

Though Elo-like systems are widely used in two-player settings, one needn’t look far to find competitions that involve much more than two players. In response to the popularity of team-based games such as CounterStrike and Halo, many recent works focus on competitions that are between two teams [14, 23, 25, 27]. Another popular setting is many-player contests such as academic olympiads: notably, programming contest platforms such as Codeforces, Topcoder, and Kaggle [6, 8, 9]. As with the aforementioned Spartan races, a typical event attracts thousands of contestants. Programming contest platforms have seen exponential growth over the past decade, collectively boasting millions of users [5]. As an example, Codeforces gained over 200K new users in 2019 alone [2].

In “free-for-all” settings, where N players are ranked individually, the Bayesian Approximation Ranking (BAR) algorithm [33] models the competition as a series of $\binom{N}{2}$ independent two-player contests. In reality, of course, the pairwise match outcomes are far from independent. Thus, TrueSkill [24] and its variants [16, 28, 30] model a player’s performance during each contest as a single random variable. The overall rankings are assumed to reveal the total order among these hidden performance variables, with various methods used to model ties and teams. For a textbook treatment of these methods, see [34]. These rating systems are efficient in practice, successfully rating userbases that number well into the millions (the Halo series, for example, has over 60 million sales since 2001 [4]).

The main disadvantage of TrueSkill is its complexity: originally developed by Microsoft for the popular Halo video game, TrueSkill performs approximate belief propagation, which consists of message passing on a factor graph, iterated until convergence. Aside from being less human-interpretable, this complexity means that, to our knowledge, there are no proofs of key properties such as runtime and incentive-compatibility. Even when these properties are discussed [28], no rigorous justification is provided. In addition, we are not aware of any work that extends TrueSkill to non-Gaussian performance models, which might be desirable to limit the influence of outlier performances (see Section 5.2).

It might be for these reasons that popular platforms such as Codeforces and Topcoder opted for their own custom rating systems. These systems are not published in academia and do not come with Bayesian justifications. However, they retain the formulaic simplicity of Elo and Glicko, extending them to settings with much more than two players. The Codeforces system includes ad hoc heuristics to distinguish top players, while curbing rampant inflation. Topcoder’s formulas are more principled from a statistical perspective; however, it has a volatility parameter similar to Glicko-2, and hence suffers from similar exploits [18]. Despite their flaws, these systems have been in place for over a decade, and have more recently gained adoption by additional platforms such as CodeChef and LeetCode [1, 7].

Our contributions. In this paper, we describe the Elo-MMR rating system, obtained by a principled approximation of a Bayesian model similar to Glicko and TrueSkill. It is fast, embarrassingly parallel, and makes accurate predictions. Most interesting of all, its simplicity

allows us to rigorously analyze its properties: the “MMR” in the name stands for “Massive”, “Monotonic”, and “Robust”. “Massive” means that it supports any number of players with a runtime that scales linearly; “monotonic” is a synonym for incentive-compatible, ensuring that a rating-maximizing player always wants to perform well; “robust” means that rating changes are bounded, with the bound being smaller for more consistent players than for volatile players. Robustness turns out to be a natural byproduct of accurately modeling performances with heavy-tailed distributions, such as the logistic. TrueSkill is believed to satisfy the first two properties, albeit without proof, but fails robustness. Codeforces only satisfies incentive-compatibility, and Topcoder only satisfies robustness.

Experimentally, we show that Elo-MMR achieves state-of-the-art performance in terms of both prediction accuracy and runtime on industry datasets. In particular, we process the entire Codeforces database of over 300K rated users and 1000 contests in well under a minute, beating the existing Codeforces system by an order of magnitude while improving upon its accuracy. Furthermore, we show that the well-known Topcoder system is severely vulnerable to volatility farming, whereas Elo-MMR is immune to such attacks. A difficulty we faced was the scarcity of efficient open-source rating system implementations. In an effort to aid researchers and practitioners alike, we provide open-source implementations of all rating systems, dataset mining, and additional processing used in our experiments at <https://github.com/EbTech/Elo-MMR>.

Organization. In Section 2, we formalize the details of our Bayesian model. We then show how to estimate player skill under this model in Section 3, and develop some intuitions of the resulting formulas. As a further refinement, Section 4 models skill evolutions from players training or atrophying between competitions. This modeling is quite tricky as we choose to retain players’ momentum while preserving incentive-compatibility. While our modeling and derivations occupy multiple sections, the system itself is succinctly presented in Algorithms 1 to 3. In Section 5, we perform a volatility farming attack on the Topcoder system and prove that, in contrast, Elo-MMR satisfies several salient properties, the most critical of which is incentive-compatibility. Finally, in Section 6, we present experimental evaluations, showing improvements over industry standards in both accuracy and speed. An extended version of this paper, with additional proofs and experiments, can be found at <https://arxiv.org/abs/2101.00400>.

2 A BAYESIAN MODEL FOR MASSIVE COMPETITIONS

We now describe the setting formally, denoting random variables by capital letters. A series of competitive **rounds**, indexed by $t = 1, 2, 3, \dots$, take place sequentially in time. Each round has a set of participating **players** \mathcal{P}_t , which may in general overlap between rounds. A player’s **skill** is likely to change with time, so we represent the skill of player i at time t by a real random variable $S_{i,t}$.

In round t , each player $i \in \mathcal{P}_t$ competes at some **performance** level $P_{i,t}$, typically close to their current skill $S_{i,t}$. The deviations $\{P_{i,t} - S_{i,t}\}_{i \in \mathcal{P}_t}$ are assumed to be i.i.d. and independent of $\{S_{i,t}\}_{i \in \mathcal{P}_t}$.

Performances are not observed directly; instead, a ranking gives the relative order among all performances $\{P_{i,t}\}_{i \in \mathcal{P}_t}$. In particular, ties are modelled to occur when performances are exactly equal,

a zero-probability event when their distributions are continuous.¹ This ranking constitutes the observational **evidence** E_t for our Bayesian updates. The rating system seeks to estimate the skill $S_{i,t}$ of every player at the present time t , given the historical round rankings $E_{\leq t} := \{E_1, \dots, E_t\}$.

We overload the notation \Pr for both probabilities and probability densities: the latter interpretation applies to zero-probability events, such as in $\Pr(S_{i,t} = s)$. We also use colons as shorthand for collections of variables differing only in a subscript: for instance, $P_{:,t} := \{P_{i,t}\}_{i \in \mathcal{P}_t}$. The joint distribution described by our Bayesian model factorizes as follows:

$$\Pr(S_{:,t}; P_{:,t}; E_t) \quad (1)$$

$$= \prod_i \Pr(S_{i,0}) \prod_{i,t} \Pr(S_{i,t} | S_{i,t-1}) \prod_{i,t} \Pr(P_{i,t} | S_{i,t}) \prod_t \Pr(E_t | P_{:,t}),$$

where $\Pr(S_{i,0})$ is the initial skill prior,

$\Pr(S_{i,t} | S_{i,t-1})$ is the skill evolution model (Section 4),

$\Pr(P_{i,t} | S_{i,t})$ is the performance model, and

$\Pr(E_t | P_{:,t})$ is the evidence model.

For the first three factors, we will specify log-concave distributions (see Definition 3.1). The evidence model, on the other hand, is a deterministic indicator. It equals one when E_t is consistent with the relative ordering among $\{P_{i,t}\}_{i \in \mathcal{P}_t}$, and zero otherwise.

Finally, our model assumes that the number of participants $|\mathcal{P}_t|$ is large. The main idea behind our algorithm is that, in sufficiently massive competitions, from the evidence E_t we can infer very precise estimates for $\{P_{i,t}\}_{i \in \mathcal{P}_t}$. Hence, we can treat these performances as if they were observed directly.

That is, suppose we have the skill prior at round t :

$$\pi_{i,t}(s) := \Pr(S_{i,t} = s | P_{i,<t}). \quad (2)$$

Now, we observe E_t . By Equation (1), it is conditionally independent of $S_{i,t}$, given $P_{i,\leq t}$. By the law of total probability,

$$\Pr(S_{i,t} = s | P_{i,<t}, E_t)$$

$$= \int \Pr(S_{i,t} = s | P_{i,<t}, P_{i,t} = p) \Pr(P_{i,t} = p | P_{i,<t}, E_t) dp$$

$$\rightarrow \Pr(S_{i,t} = s | P_{i,\leq t}) \quad \text{almost surely as } |\mathcal{P}_t| \rightarrow \infty.$$

The integral is intractable in general, since the performance posterior $\Pr(P_{i,t} = p | P_{i,<t}, E_t)$ depends not only on player i , but also on our belief regarding the skills of all $j \in \mathcal{P}_t$. However, in the limit of infinite participants, Doob's consistency theorem [19] implies that it concentrates at the true value $P_{i,t}$. Since our posteriors are continuous, the convergence holds for all s simultaneously.

Indeed, we don't even need the full evidence E_t . Let $E_{i,t}^L = \{j \in \mathcal{P} : P_{j,t} > P_{i,t}\}$ be the set of players against whom i lost, and $E_{i,t}^W = \{j \in \mathcal{P} : P_{j,t} < P_{i,t}\}$ be the set of players against whom i won. That is, we only see who wins, draws, and loses against i . $P_{i,t}$ remains identifiable using only $(E_{i,t}^L, E_{i,t}^W)$, which will be more convenient for our purposes.

Passing to the limit in which $P_{i,\leq t}$ is identified serves to justify several common simplifications made by total-order rating systems. Firstly, since $P_{i,\leq t}$ is a sufficient statistic for predicting $S_{i,t}$, it may

be said that $(E_{i,\leq t}^L, E_{i,\leq t}^W)$ are “almost sufficient” for $S_{i,t}$: any additional information, such as from domain-specific scoring systems, becomes redundant for the purposes of skill estimation. Secondly, conditioned on $P_{i,\leq t}$, the posterior skills $S_{:,t}$ are independent of one another. As a result, there are no inter-player correlations to model, and a player's posterior is unaffected by rounds in which they are not a participant. Finally, if we've truly identified $P_{i,t}$, then rounds later than t should not prompt revisions in our estimate for $P_{i,t}$. This obviates the need for expensive whole-history update procedures [15, 16], for the purposes of present skill estimation.²

Finally, a word on the rate of convergence. Suppose we want our estimate to be within ϵ of $P_{i,t}$, with probability at least $1 - \delta$. By asymptotic normality of the posterior [19], it suffices to have $O(\frac{1}{\epsilon^2} \log \frac{1}{\delta})$ participants.

When the initial prior, performance model, and evolution model are all Gaussian, treating $P_{i,t}$ as certain is the *only* simplifying approximation we will make; that is, in the limit $|\mathcal{P}_t| \rightarrow \infty$, our method performs *exact* inference on Equation (1). In the following sections, we focus some attention on generalizing the performance model to non-Gaussian log-concave families, parametrized by location and scale. We will use the logistic distribution as a running example and see that it induces robustness; however, our framework is agnostic to the specific distributions used.

The prior **rating** $\mu_{i,t}^\pi$ and posterior rating $\mu_{i,t}$ of player i at round t should be statistics that summarize the player's prior and posterior skill distribution, respectively. We'll use the mode: thus, $\mu_{i,t}$ is the maximum a posteriori (MAP) estimate, obtained by setting s to maximize the posterior $\Pr(S_{i,t} = s | P_{i,\leq t})$. By Bayes' rule,

$$\mu_{i,t}^\pi := \arg \max_s \pi_{i,t}(s),$$

$$\mu_{i,t} := \arg \max_s \pi_{i,t}(s) \Pr(P_{i,t} | S_{i,t} = s). \quad (3)$$

This objective suggests a two-phase algorithm to update each player $i \in \mathcal{P}_t$ in response to the results of round t . In phase one, we estimate $P_{i,t}$ from $(E_{i,t}^L, E_{i,t}^W)$. By Doob's consistency theorem, our estimate is extremely precise when $|\mathcal{P}_t|$ is large, so we assume it to be exact. In phase two, we update our posterior for $S_{i,t}$ and the rating $\mu_{i,t}$ according to Equation (3).

3 SKILL ESTIMATION IN TWO PHASES

3.1 Performance estimation

In this section, we describe the first phase of Elo-MMR. For notational convenience, we assume all probability expressions to be conditioned on the **prior context** $P_{i,<t}$, and omit the subscript t .

Our prior belief on each player's skill S_i implies a prior distribution on P_i . Let's denote its probability density function (pdf) by

$$f_i(p) := \Pr(P_i = p) = \int \pi_i(s) \Pr(P_i = p | S_i = s) ds, \quad (4)$$

where $\pi_i(s)$ was defined in Equation (2). Let

$$F_i(p) := \Pr(P_i \leq p) = \int_{-\infty}^p f_i(x) dx,$$

be the corresponding cumulative distribution function (cdf). For the purpose of analysis, we'll also define the following “loss”, “draw”,

¹The relevant limiting procedure is to treat performances within ϵ -width buckets as ties, and letting $\epsilon \rightarrow 0$. This technicality appears in the proof of Theorem 3.2.

²As opposed to *historical* skill estimation, which is concerned with $P(S_{i,t} | P_{i,\leq t'})$ for $t' > t$. Whole-history methods can take advantage of future information.

and “victory” functions:

$$\begin{aligned} l_i(p) &:= \frac{d}{dp} \ln(1 - F_i(p)) = \frac{-f_i(p)}{1 - F_i(p)}, \\ d_i(p) &:= \frac{d}{dp} \ln f_i(p) = \frac{f_i'(p)}{f_i(p)}, \\ v_i(p) &:= \frac{d}{dp} \ln F_i(p) = \frac{f_i(p)}{F_i(p)}. \end{aligned}$$

Evidently, $l_i(p) < 0 < v_i(p)$. Now we define what it means for the deviation $P_i - S_i$ to be log-concave.

DEFINITION 3.1. *An absolutely continuous random variable on a convex domain is **log-concave** if its probability density function f is positive on its domain and satisfies*

$$f(\theta x + (1 - \theta)y) > f(x)^\theta f(y)^{1-\theta}, \quad \forall \theta \in (0, 1), x \neq y.$$

Log-concave distributions appear widely, and include the Gaussian and logistic distributions used in Glicko, TrueSkill, and many others. We’ll see inductively that our prior π_i is log-concave at every round. Since log-concave densities are closed under convolution [12], the independent sum $P_i = S_i + (P_i - S_i)$ is also log-concave. Log-concavity is made very convenient by the following lemma, proved in the extended version of this paper:

LEMMA 3.1. *If f_i is continuously differentiable and log-concave, then the functions l_i, d_i, v_i are continuous, strictly decreasing, and*

$$l_i(p) < d_i(p) < v_i(p) \text{ for all } p.$$

For the remainder of this section, we fix the analysis with respect to some player i . As argued in Section 2, P_i concentrates very narrowly in the posterior. Hence, we can estimate P_i by its MAP, choosing p so as to maximize:

$$\Pr(P_i = p \mid E_i^L, E_i^W) \propto f_i(p) \Pr(E_i^L, E_i^W \mid P_i = p).$$

Define $j > i, j < i, j \sim i$ as shorthand for $j \in E_i^L, j \in E_i^W, j \in \mathcal{P} \setminus (E_i^L \cup E_i^W)$ (that is, $P_j > P_i, P_j < P_i, P_j = P_i$), respectively. The following theorem yields our MAP estimate:

THEOREM 3.2. *Suppose that for all j , f_j is continuously differentiable and log-concave. Then the unique maximizer of $\Pr(P_i = p \mid E_i^L, E_i^W)$ is given by the unique zero of*

$$Q_i(p) := \sum_{j>i} l_j(p) + \sum_{j\sim i} d_j(p) + \sum_{j<i} v_j(p).$$

The proof appears in the extended version of this paper. Intuitively, we’re saying that the performance is the balance point between appropriately weighted wins, draws, and losses. Let’s look at two specializations of our general model, to serve as running examples in this paper.

Gaussian performance model. If both S_j and $P_j - S_j$ are assumed to be Gaussian with known means and variances, then their independent sum P_j will also be a known Gaussian. It is analytic and log-concave, so Theorem 3.2 applies.

We substitute the well-known Gaussian pdf and cdf for f_j and F_j , respectively. A simple binary search, or faster numerical techniques such as the Illinois algorithm or Newton’s method, can be employed to solve for the unique zero of Q_i .




Figure 1: L_2 versus L_R for typical values (left). Gaussian versus logistic probability density functions (right).

Logistic performance model. Now we assume the performance deviation $P_j - S_j$ has a logistic distribution with mean 0 and variance β^2 . In general, the rating system administrator is free to set β differently for each contest. Since shorter contests tend to be more variable, one reasonable choice might be to make $1/\beta^2$ proportional to the contest duration.

Given the mean and variance of the skill prior, the independent sum $P_j = S_j + (P_j - S_j)$ would have the same mean, and a variance that’s increased by β^2 . Unfortunately, we’ll see that the logistic performance model implies a form of skill prior from which it’s tough to extract a mean and variance. Even if we could, the sum does not yield a simple distribution.

For experienced players, we expect S_j to contribute much less variance than $P_j - S_j$; thus, in our heuristic approximation, we take P_j to have the same form of distribution as the latter. That is, we take P_j to be logistic, centered at the prior rating $\mu_j^\pi = \arg \max \pi_j$, with variance $\delta_j^2 = \sigma_j^2 + \beta^2$, where σ_j will be given by Equation (8). This distribution is analytic and log-concave, so the same methods based on Theorem 3.2 apply. Define the scale parameter $\bar{\delta}_j := \frac{\sqrt{3}}{\pi} \delta_j$. A logistic distribution with variance δ_j^2 has cdf and pdf:

$$\begin{aligned} F_j(x) &= \frac{1}{1 + e^{-(x - \mu_j^\pi)/\delta_j}} = \frac{1}{2} \left(1 + \tanh \frac{x - \mu_j^\pi}{2\delta_j} \right), \\ f_j(x) &= \frac{e^{(x - \mu_j^\pi)/\delta_j}}{\delta_j \left(1 + e^{(x - \mu_j^\pi)/\delta_j} \right)^2} = \frac{1}{4\bar{\delta}_j} \operatorname{sech}^2 \frac{x - \mu_j^\pi}{2\bar{\delta}_j}. \end{aligned}$$

The logistic distribution satisfies two very convenient relations:

$$\begin{aligned} F_j'(x) &= f_j(x) = F_j(x)(1 - F_j(x))/\delta_j, \\ f_j'(x) &= f_j(x)(1 - 2F_j(x))/\bar{\delta}_j, \end{aligned}$$

from which it follows that

$$d_j(p) = \frac{1 - 2F_j(p)}{\bar{\delta}} = \frac{-F_j(p)}{\bar{\delta}} + \frac{1 - F_j(p)}{\bar{\delta}} = l_j(p) + v_j(p).$$

In other words, a tie counts as the sum of a win and a loss. This can be compared to the approach (used in Elo, Glicko, BAR, Topcoder, and Codeforces) of treating each tie as half a win plus half a loss.³

³Elo-MMR, too, can be modified to split ties into half win plus half loss. It’s easy to check that Lemma 3.1 still holds if $d_j(p)$ is replaced by $w_l l_j(p) + w_o v_j(p)$ for some $w_l, w_o \in [0, 1]$ with $|w_l - w_o| < 1$. In particular, we can set $w_l = w_o = 0.5$. The results in Section 5 won’t be altered by this change.

Finally, putting everything together:

$$Q_i(p) = \sum_{j \geq i} l_j(p) + \sum_{j \leq i} v_j(p) = \sum_{j \geq i} \frac{-F_j(p)}{\delta_j} + \sum_{j \leq i} \frac{1 - F_j(p)}{\delta_j}.$$

Our estimate for P_i is the zero of this expression. The terms on the right correspond to probabilities of winning and losing against each player j , weighted by $1/\delta_j$. Accordingly, we can interpret $\sum_{j \in \mathcal{P}} (1 - F_j(p))/\delta_j$ as a weighted expected rank of a player whose performance is p . Similar to the performance computations in Codeforces and Topcoder, P_i can thus be viewed as the performance level at which one's expected rank would equal i 's actual rank.

3.2 Belief update

Having estimated $P_{i,t}$ in the first phase, the second phase is rather simple. Ignoring normalizing constants, Equation (3) tells us that the pdf of the skill posterior can be obtained as the pointwise product of the pdfs of the skill prior and the performance model. When both factors are differentiable and log-concave, then so is their product. Its maximum is the new rating $\mu_{i,t}$; let's see how to compute it for the same two specializations of our model.

Gaussian performance model. When the skill prior and performance model are Gaussian with known means and variances, multiplying their pdfs yields another known Gaussian. Hence, the posterior is compactly represented by its mean $\mu_{i,t}$, which coincides with the MAP and rating; and its variance $\sigma_{i,t}^2$, which is our **uncertainty** regarding the player's skill.

Logistic performance model. When the performance model is non-Gaussian, the multiplication does not simplify so easily. By Equation (3), each round contributes an additional factor to the belief distribution. In general, we allow it to consist of a collection of simple log-concave factors, one for each round in which player i has participated. Denote the participation history by

$$\mathcal{H}_{i,t} := \{k \in \{1, \dots, t\} : i \in \mathcal{P}_k\}.$$

Since each player can be considered in isolation, we'll omit the subscript i . Specializing to the logistic setting, each $k \in \mathcal{H}_t$ contributes a logistic factor to the posterior, with mean p_k and variance β_k^2 . We still use a Gaussian initial prior, with mean and variance denoted by p_0 and β_0^2 , respectively. Postponing the discussion of skill evolution to Section 4, for the moment we assume that $S_k = S_0$ for all k . The posterior pdf, up to normalization, is then

$$\begin{aligned} \pi_0(s) & \prod_{k \in \mathcal{H}_t} \Pr(P_k = p_k \mid S_k = s) \\ & \propto \exp\left(-\frac{(s-p_0)^2}{2\beta_0^2}\right) \prod_{k \in \mathcal{H}_t} \operatorname{sech}^2\left(\frac{\pi}{\sqrt{12}} \frac{s-p_k}{\beta_k}\right). \end{aligned} \quad (5)$$

Maximizing the posterior density amounts to minimizing its negative logarithm. Up to a constant offset, this is given by

$$L(s) := L_2\left(\frac{s-p_0}{\beta_0}\right) + \sum_{k \in \mathcal{H}_t} L_R\left(\frac{s-p_k}{\beta_k}\right),$$

where $L_2(x) := \frac{1}{2}x^2$ and $L_R(x) := 2 \ln\left(\cosh \frac{\pi x}{\sqrt{12}}\right)$.

$$\text{Thus, } L'(s) = \frac{s-p_0}{\beta_0^2} + \sum_{k \in \mathcal{H}_t} \frac{\pi}{\beta_k \sqrt{3}} \tanh\left(\frac{(s-p_k)\pi}{\beta_k \sqrt{12}}\right). \quad (6)$$

L' is continuous and strictly increasing in s , so its zero is unique: it is the MAP μ_t . Similar to what we did in the first phase, we can solve for μ_t with binary search or other root-solving methods.

We pause to make an important observation. From Equation (6), the rating carries a rather intuitive interpretation: Gaussian factors in L become L_2 penalty terms, whereas logistic factors take on a more interesting form as L_R terms. From Figure 1, we see that the L_R term behaves quadratically near the origin, but linearly at the extremities, effectively interpolating between L_2 and L_1 over a scale of magnitude β_k .

It is well-known that minimizing a sum of L_2 terms pushes the argument towards a weighted mean, while minimizing a sum of L_1 terms pushes the argument towards a weighted median. With L_R terms, the net effect is that μ_t acts like a robust average of the historical performances p_k . Specifically, one can check that

$$\begin{aligned} \mu_t &= \frac{\sum_k w_k p_k}{\sum_k w_k}, \text{ where } w_0 := \frac{1}{\beta_0^2} \text{ and} \\ w_k &:= \frac{\pi}{(\mu_t - p_k)\beta_k \sqrt{3}} \tanh\left(\frac{(\mu_t - p_k)\pi}{\beta_k \sqrt{12}}\right) \text{ for } k \in \mathcal{H}_t. \end{aligned} \quad (7)$$

w_k is close to $1/\beta_k^2$ for typical performances, but can be up to $\pi^2/6$ times more as $|\mu_t - p_k| \rightarrow 0$, or vanish as $|\mu_t - p_k| \rightarrow \infty$. This feature is due to the thicker tails of the logistic distribution, as compared to the Gaussian, resulting in an algorithm that resists drastic rating changes in the presence of a few unusually good or bad performances. We'll formally state this *robustness* property in Theorem 5.7.

Estimating skill uncertainty. While there is no easy way to compute the variance of a posterior in the form of Equation (5), it will be useful to have some estimate σ_t^2 of uncertainty. There is a simple formula in the case where all factors are Gaussian. Since moment-matched logistic and normal distributions are relatively close (c.f. Figure 1), we apply the same formula:

$$\frac{1}{\sigma_t^2} := \sum_{k \in \{0\} \cup \mathcal{H}_t} \frac{1}{\beta_k^2}. \quad (8)$$

4 SKILL EVOLUTION OVER TIME

Factors such as training and resting will change a player's skill over time. If we model skill as a static variable, our system will eventually grow so confident in its estimate that it will refuse to admit substantial changes. To remedy this, we introduce a skill evolution model, so that in general $S_t \neq S_{t'}$ for $t \neq t'$. Now rather than simply being equal to the previous round's posterior, the skill prior at round t is given by

$$\pi_t(s) = \int \Pr(S_t = s \mid S_{t-1} = x) \Pr(S_{t-1} = x \mid P_{<t}) dx. \quad (9)$$

The factors in the integrand are the skill evolution model and the previous round's posterior, respectively. Following other Bayesian rating systems (e.g., Glicko, Glicko-2, and TrueSkill [21, 22, 24]), we model the skill diffusions $S_t - S_{t-1}$ as independent zero-mean Gaussians. That is, $\Pr(S_t \mid S_{t-1} = x)$ is a Gaussian with mean x and some variance γ_t^2 . The Glicko system sets γ_t^2 proportionally to

the time elapsed since the last update, corresponding to a continuous Brownian motion. Codeforces and Topcoder simply set γ_t to a constant when a player participates, and zero otherwise, corresponding to changes that are in proportion to how often the player competes. Now we are ready to complete the two specializations of our rating system.

Gaussian performance model. If both the prior and performance distributions at round $t - 1$ are Gaussian, then the posterior is also Gaussian. Adding an independent Gaussian diffusion to our posterior on S_{t-1} yields a Gaussian prior on S_t . By induction, the skill belief distribution forever remains Gaussian. This Gaussian specialization of the Elo-MMR framework lacks the R for robustness (see Theorem 5.7), so we call it Elo-MM χ .

Logistic performance model. After a player's first contest round, the posterior in Equation (5) becomes non-Gaussian, rendering the integral in Equation (9) intractable.

A very simple approach would be to replace the full posterior in Equation (5) by a Gaussian approximation with mean μ_t (equal to the posterior MAP) and variance σ_t^2 (given by Equation (8)). As in the previous case, applying diffusions in the Gaussian setting is a simple matter of adding means and variances.

With this approximation, no memory is kept of the individual performances P_t . Priors are simply Gaussian, while posterior densities are the product of two factors: the Gaussian prior, and a logistic factor corresponding to the latest performance. To ensure robustness (see Section 5.2), μ_t is computed as the argmax of this posterior *before* replacement by its Gaussian approximation. We call the rating system that takes this approach Elo-MMR(∞).

As the name implies, it turns out to be a special case of Elo-MMR(ρ). In the general setting with $\rho \in [0, \infty)$, we keep the full posterior from Equation (5). Since we cannot tractably compute the effect of a Gaussian diffusion, we seek a heuristic derivation of the next round's prior, retaining a form similar to Equation (5) while satisfying many of the same properties as the intended diffusion.

4.1 Desirable properties of a “pseudodiffusion”

We begin by listing some properties that our skill evolution algorithm, henceforth called a “pseudodiffusion”, should satisfy. The first two properties are natural:

- *Incentive-compatibility.* First and foremost, the pseudodiffusion must not break the incentive-compatibility of our rating system. That is, a rating-maximizing player should never be motivated to lose on purpose. (Theorem 5.5).
- *Rating preservation.* The pseudodiffusion must not alter the arg max of the belief density. That is, the rating of a player should not change: $\mu_t^\pi = \mu_{t-1}$.

In addition, we borrow four properties of Gaussian diffusions:

- *Correct magnitude.* Pseudodiffusion with parameter γ^2 must increase the skill uncertainty, as measured by Equation (8), by γ^2 .
- *Composability.* Two pseudodiffusions applied in sequence, first with parameter γ_1^2 and then with γ_2^2 , must have the same effect as a single pseudodiffusion with parameter $\gamma_1^2 + \gamma_2^2$.
- *Zero diffusion.* In the limit as $\gamma \rightarrow 0$, the effect of pseudodiffusion must vanish, i.e., not alter the belief distribution.

- *Zero uncertainty.* In the limit as $\sigma_{t-1} \rightarrow 0$ (i.e., when the previous rating μ_{t-1} is a perfect estimate of S_{t-1}), our belief on S_t must become Gaussian with mean μ_{t-1} and variance γ^2 . Finer-grained information regarding the prior history $P_{\leq t}$ must be erased.

In particular, Elo-MMR(∞) fails the *zero diffusion* property because it simplifies the belief distribution, even when $\gamma = 0$. In the proof of Theorem 4.1, we'll see that Elo-MMR(0) fails the *zero uncertainty* property. Thus, it is in fact necessary to have ρ strictly positive and finite. In Section 5.2, we'll come to interpret ρ as a kind of inverse momentum.

4.2 A heuristic pseudodiffusion algorithm

Each factor in the posterior (see Equation (5)) has a parameter β_k . Define a factor's **weight** to be $w_k := 1/\beta_k^2$, which by Equation (8) contributes to the **total weight** $\sum_k w_k = 1/\sigma_t^2$. Here, unlike in Equation (7), w_k does not depend on $|\mu_t - p_k|$.

The approximation step of Elo-MMR(∞) replaces all the logistic factors by a single Gaussian whose variance is chosen to ensure that the total weight is preserved. In addition, its mean is chosen to preserve the player's rating, given by the unique zero of Equation (6). Finally, the diffusion step of Elo-MMR(∞) increases the Gaussian's variance, and hence the player's skill uncertainty, by γ_t^2 ; this corresponds to a decay in the weight.

To generalize the idea, we interleave the two steps in a continuous manner. The approximation step becomes a **transfer step**: rather than replace the logistic factors outright, we take away the same fraction from each of their weights, and *place the sum of removed weights onto a new Gaussian factor*. In order for this operation to preserve ratings, the new factor must be centered at μ_{t-1} . Since Gaussian pdfs compose, the prior Gaussian factor can be combined with the new one. The diffusion step becomes a **decay step**, reducing each factor's weight by the same fraction, chosen such that the overall uncertainty is increased by γ_t^2 .

To make the idea precise, we generalize the posterior from Equation (5) with fractional **multiplicities** ω_k ; the k 'th factor is raised to the power ω_k . As a result, Equations (6) and (8) become:

$$L'(s) = \frac{\omega_0(s - p_0)}{\beta_0^2} + \sum_{k \in \mathcal{H}_t} \frac{\omega_k \pi}{\beta_k \sqrt{3}} \tanh \frac{(s - p_k) \pi}{\beta_k \sqrt{12}},$$

$$\frac{1}{\sigma_t^2} := \sum_{k \in \{0\} \cup \mathcal{H}_t} w_k, \quad \text{where } w_k := \frac{\omega_k}{\beta_k^2}. \quad (10)$$

For $\rho \in [0, \infty]$, the Elo-MMR(ρ) algorithm continuously and simultaneously performs transfer and decay, with transfer proceeding at ρ times the rate of decay. Holding β_k fixed, changes to ω_k can be described in terms of changes to w_k :

$$\dot{w}_0 = -r(t)w_0 + \rho r(t) \sum_{k \in \mathcal{H}_t} w_k,$$

$$\dot{w}_k = -(1 + \rho)r(t)w_k \quad \text{for } k \in \mathcal{H}_t,$$

where the arbitrary decay rate $r(t)$ can be eliminated by a change of variable $d\tau = r(t)dt$. After some time $\Delta\tau$, the total weight will

Algorithm 1 Elo-MMR(ρ, β, γ)

for all rounds t **do**
 for all players $i \in \mathcal{P}_t$ in parallel **do**
 if i has never competed before **then**
 $\mu_i, \sigma_i \leftarrow \mu_{\text{newcomer}}, \sigma_{\text{newcomer}}$
 $p_i, w_i \leftarrow [\mu_i], [1/\sigma_i^2]$
 diffuse(i, γ, ρ)
 $\mu_i^\pi, \delta_i \leftarrow \mu_i, \sqrt{\sigma_i^2 + \beta^2}$
 for all $i \in \mathcal{P}_t$ in parallel **do**
 update(i, E_t, β)

Algorithm 2 diffuse(i, γ, ρ)

$\kappa \leftarrow (1 + \gamma^2/\sigma_i^2)^{-1}$
 $w_G, w_L \leftarrow \kappa^\rho w_{i,0}, (1 - \kappa^\rho) \sum_{k \geq 0} w_{i,k}$
 $p_{i,0} \leftarrow (w_G p_{i,0} + w_L \mu_i) / (w_G + w_L)$
 $w_{i,0} \leftarrow \kappa (w_G + w_L)$
for all $k > 0$ **do**
 $w_{i,k} \leftarrow \kappa^{1+\rho} w_{i,k}$
 $\sigma_i \leftarrow \sigma_i / \sqrt{\kappa}$

Algorithm 3 update(i, E, β)

$p \leftarrow \text{zero}_x \left(\sum_{j \leq i} \frac{1}{\delta_j} \left(\tanh \frac{x - \mu_j^\pi}{2\delta_j} - 1 \right) + \sum_{j \geq i} \frac{1}{\delta_j} \left(\tanh \frac{x - \mu_j^\pi}{2\delta_j} + 1 \right) \right)$
 $p_i.\text{push}(p)$
 $w_i.\text{push}(1/\beta^2)$
 $\mu_i \leftarrow \text{zero}_x \left(w_{i,0} (x - p_{i,0}) + \sum_{k > 0} \frac{w_{i,k} \beta^2}{\beta} \tanh \frac{x - p_{i,k}}{2\beta} \right)$

have decayed by a factor $\kappa := e^{-\Delta\tau}$, resulting in the new weights:

$$w_0^{\text{new}} = \kappa w_0 + (\kappa - \kappa^{1+\rho}) \sum_{k \in \mathcal{H}_t} w_k,$$

$$w_k^{\text{new}} = \kappa^{1+\rho} w_k \quad \text{for } k \in \mathcal{H}_t.$$

In order for the uncertainty to increase from σ_{t-1}^2 to $\sigma_{t-1}^2 + \gamma_t^2$, we must solve $\kappa/\sigma_{t-1}^2 = 1/(\sigma_{t-1}^2 + \gamma_t^2)$ for the decay factor:

$$\kappa_t = \left(1 + \frac{\gamma_t^2}{\sigma_{t-1}^2} \right)^{-1}.$$

Algorithm 1 details the full Elo-MMR(ρ) rating system. The main loop runs whenever a round of competition takes place. First, new players are initialized with a Gaussian prior. Then, changes in player skill are modeled by Algorithm 2; note how the Gaussian prior is updated into a weighted combination with the newly created factor. Given the round rankings E_t , the first phase of Algorithm 3 solves an equation to estimate P_t . Finally, the second phase solves another equation for the rating μ_t .

The hyperparameters ρ, β, γ are domain-dependent, and can be set by standard hyperparameter search techniques. For convenience, we assume β and γ are fixed and use the shorthand $\bar{\beta}_k := \frac{\sqrt{3}}{\pi} \beta_k$. Whereas our exposition used global round indices, here a subscript k corresponds to the k 'th round in player i 's participation history.

THEOREM 4.1. *Algorithm 2 with $\rho \in (0, \infty)$ meets all of the properties listed in Section 4.1.*

PROOF. We go through each of the six properties in order.

- *Incentive-compatibility.* This property will be stated in Theorem 5.5. To ensure that its proof carries through, the relevant facts to note here are that the pseudodiffusion algorithm ignores the performances p_k , and centers the transferred Gaussian weight at the rating μ_{t-1} , which is trivially monotonic in μ_{t-1} .
- *Rating preservation.* Recall that the rating is the unique zero of L' in Equation (10). To see that this zero is preserved, note that the decay and transfer operations multiply L' by constants (κ_t and κ_t^ρ , respectively), before adding the new Gaussian term, whose contribution to L' is zero at its center.
- *Correct magnitude.* Follows from our derivation for κ_t .
- *Composability.* Follows from *correct magnitude* and the fact that every pseudodiffusion follows the same differential equations.
- *Zero diffusion.* As $\gamma \rightarrow 0$, $\kappa_t \rightarrow 1$. Provided that $\rho < \infty$, we also have $\kappa_t^\rho \rightarrow 1$. Hence, for all $k \in \{0\} \cup \mathcal{H}_t$, $w_k^{\text{new}} \rightarrow w_k$.
- *Zero uncertainty.* As $\sigma_{t-1} \rightarrow 0$, $\kappa_t \rightarrow 0$. The total weight decays from $1/\sigma_{t-1}^2$ to γ^2 . Provided that $\rho > 0$, we also have $\kappa_t^\rho \rightarrow 0$, so these weights transfer in their entirety, leaving behind a Gaussian with mean μ_{t-1} , variance γ^2 , and no additional history. \square

5 THEORETICAL PROPERTIES

In this section, we see how the simplicity of the Elo-MMR formulas enables us to rigorously prove that the rating system is incentive-compatible, robust, and computationally efficient.

5.1 Incentive-compatibility

To demonstrate the need for incentive-compatibility, let's look at the consequences of violating this property in the Topcoder and Glicko-2 rating systems. These systems track a "volatility" for each player, which estimates the variance of their performances. A player whose recent performance history is more consistent would be assigned a lower volatility score, than one with wild swings in performance. The volatility acts as a multiplier on rating changes; thus, players with an extremely low or high performance will have their subsequent rating changes amplified.

While it may seem like a good idea to boost changes for players whose ratings are poor predictors of their performance, this feature has an exploit. By intentionally performing at a weaker level, a player can amplify future increases to an extent that more than compensates for the immediate hit to their rating. A player may even "farm" volatility by alternating between very strong and very weak performances. After acquiring a sufficiently high volatility score, the strategic player exerts their honest maximum performance over a series of contests. The amplification eventually results in a rating that exceeds what would have been obtained via honest play. This type of exploit was discovered in Glicko-2 as applied to the Pokemon Go video game [3]. Table 5.3 of [18] presents a milder violation in Topcoder competitions.

To get a realistic estimate of the severity of this exploit, we performed a simple experiment on the first five years of the Codeforces contest dataset (see Section 6.1). In Figure 2, we plot the rating evolution of the world's #1 ranked competitive programmer, Gennady Korotkevich, better known as *tourist*. In the *control* setting, we plot his ratings according to the Topcoder and Elo-MMR(1) systems. We contrast these against an *adversarial* setting, in which we have *tourist* employ the following strategy: for his first 45 contests, *tourist* plays normally (exactly as in the unaltered data). For his




Figure 2: Volatility farming attack on the Topcoder system.

next 45 contests, tourist purposely falls to last place whenever his Topcoder rating is above 2975. Finally, tourist returns to playing normally for an additional 15 contests.

This strategy mirrors the Glicko-2 exploit documented in [3], and does not require unrealistic assumptions (e.g., we don't demand tourist to exercise very precise control over his performances). Compared to a consistently honest tourist, the volatility farming tourist ended up 523 rating points ahead by the end of the experiment, with almost 1000 rating points gained in the last 15 contests alone. Transferring the same sequence of performances to the Elo-MMR(1) system, we see that it not only is immune to such volatility-farming attacks, but it also penalizes the dishonest strategy with a rating loss that decays exponentially once honest play resumes.

Recall that a key purpose of modeling volatility in Topcoder and Glicko-2 was to boost rating changes for inconsistent players. Remarkably, Elo-MMR achieves the same effect: we'll see in Section 5.2 that, for $\rho \in [0, \infty)$, Elo-MMR(ρ) also boosts changes to inconsistent players. And yet, we'll now prove that no strategic incentive for purposely losing exists in *any* version of Elo-MMR.

To this end, we need a few lemmas. Recall that, for the purposes of the algorithm, the performance p_i is defined to be the unique zero of the function $Q_i(p) := \sum_{j>i} l_j(p) + \sum_{j\sim i} d_j(p) + \sum_{j<i} v_j(p)$, whose terms l_i, d_i, v_i are contributed by opponents against whom i lost, drew, or won, respectively. Wins (losses) are always positive (negative) contributions to a player's performance score:

LEMMA 5.1. *Adding a win term to $Q_i(\cdot)$, or replacing a tie term by a win term, always increases its zero. Conversely, adding a loss term, or replacing a tie term by a loss term, always decreases it.*

PROOF. By Lemma 3.1, $Q_i(p)$ is decreasing in p . Thus, adding a positive term will increase its zero whereas adding a negative term will decrease it. The desired conclusion follows by noting that, for all j and p , $v_j(p)$ and $v_j(p) - d_j(p)$ are positive, whereas $l_j(p)$ and $l_j(p) - d_j(p)$ are negative. \square

While not needed for our main result, a similar argument shows that performance scores are monotonic across the round standings:

THEOREM 5.2. *If $i > j$ (that is, player i beats j) in a given round, then player i and j 's performance estimates satisfy $p_i > p_j$.*

PROOF. If $i > j$ with i, j adjacent in the rankings, then

$$Q_i(p) - Q_j(p) = \sum_{k\sim i} (d_k(p) - l_k(p)) + \sum_{k\sim j} (v_k(p) - d_k(p)) > 0.$$

for all p . Since Q_i and Q_j are decreasing functions, it follows that $p_i > p_j$. By induction, this result extends to the case where i, j are not adjacent in the rankings. \square

What matters for incentives is that performance scores be *counterfactually* monotonic; meaning, if we were to alter the round standings, a strategic player will always prefer to place higher:

LEMMA 5.3. *In any given round, holding fixed the relative ranking of all players other than i (and holding fixed all preceding rounds), the performance p_i is a monotonic function of player i 's prior rating and of player i 's rank in this round.*

PROOF. Monotonicity in the prior rating follows directly from monotonicity of the self-tie term d_i in Q_i . Since an upward shift in the rankings can only convert losses to ties to wins, monotonicity in contest rank follows from Lemma 5.1. \square

Having established the relationship between round rankings and performance scores, the next step is to prove that, even with hindsight, players will always prefer their performance scores to be as high as possible:

LEMMA 5.4. *Holding fixed the set of contest rounds in which a player has participated, their current rating is monotonic in each of their past performance scores.*

PROOF. The player's rating is given by the zero of L' in Equation (10). The pseudodiffusions of Section 4 modify each of the β_k in a manner that does not depend on any of the p_k , so they are fixed for our purposes. Hence, L' is monotonically increasing in s and decreasing in each of the p_k . Therefore, its zero is monotonically increasing in each of the p_k .

This is almost what we wanted to prove, except that p_0 is not a performance. Nonetheless, it is a function of the performances: specifically, a weighted average of historical ratings which, using this same lemma as an inductive hypothesis, are themselves monotonic in past performances. By induction, the proof is complete. \square

Finally, we conclude that a rating-maximizing player is always motivated to improve their round rankings, or raw scores:

THEOREM 5.5 (INCENTIVE-COMPATIBILITY). *Holding fixed the set of contest rounds in which each player has participated, and the historical ratings and relative rankings of all players other than i , player i 's current rating is monotonic in each of their past rankings.*

PROOF. Choose any contest round in player i 's history, and consider improving player i 's rank in that round while holding everything else fixed. It suffices to show that player i 's current rating would necessarily increase as a result.

In the altered round, by Lemma 5.3, p_i is increased; and by Lemma 5.4, player i 's post-round rating is increased. By Lemma 5.3 again, this increases player i 's performance score in the following round. Proceeding inductively, we find that performance scores and ratings from this point onward are all increased. \square

In the special cases of Elo-MM χ or Elo-MMR(∞), the rating system is “memoryless”: the only data retained for each player are the current rating $\mu_{i,t}$ and uncertainty $\sigma_{i,t}$; detailed performance history is not saved. In this setting, we present a natural monotonicity theorem. A similar theorem was previously stated for the Codeforces system, albeit in an informal context without proof [8].

THEOREM 5.6 (MEMORYLESS MONOTONICITY). *In either the Elo-MM χ or Elo-MMR(∞) system, suppose i and j are two participants of round t . Suppose that the ratings and corresponding uncertainties satisfy $\mu_{i,t-1} \geq \mu_{j,t-1}$, $\sigma_{i,t-1} = \sigma_{j,t-1}$. Then, $\sigma_{i,t} = \sigma_{j,t}$. Furthermore:*

If $i > j$ in round t , then $\mu_{i,t} > \mu_{j,t}$.

If $j > i$ in round t , then $\mu_{j,t} - \mu_{j,t-1} > \mu_{i,t} - \mu_{i,t-1}$.

PROOF. The new contest round will add a rating perturbation with variance γ_t^2 , followed by a new performance with variance β_t^2 . As a result,

$$\sigma_{i,t} = \left(\frac{1}{\sigma_{i,t-1}^2 + \gamma_t^2} + \frac{1}{\beta_t^2} \right)^{-\frac{1}{2}} = \left(\frac{1}{\sigma_{j,t-1}^2 + \gamma_t^2} + \frac{1}{\beta_t^2} \right)^{-\frac{1}{2}} = \sigma_{j,t}.$$

The remaining conclusions are consequences of three properties: memorylessness, incentive-compatibility (Theorem 5.5), and translation-invariance (ratings, skills, and performances are quantified on a common interval scale relative to one another).

Since the Elo-MM χ or Elo-MMR(∞) systems are memoryless, we may replace the initial prior and performance histories of players with any alternate histories of our choosing, as long as our choice is compatible with their current rating and uncertainty. For example, both i and j can be considered to have participated in the same set of rounds, with i always performing at $\mu_{i,t-1}$ and j always performing at $\mu_{j,t-1}$. Round t is unchanged.

Suppose $i > j$. Since i 's historical performances are all equal or stronger than j 's, Theorem 5.5 implies $\mu_{i,t} > \mu_{j,t}$.

Suppose $j > i$. By translation-invariance, if we shift each of j 's performances, up to round t and including the initial prior, upward by $\mu_{i,t-1} - \mu_{j,t-1}$, the rating changes between rounds will be unaffected. Players i and j now have identical histories, except that we still have $j > i$ at round t . Therefore, $\mu_{j,t-1} = \mu_{i,t-1}$ and, by Theorem 5.5, $\mu_{j,t} > \mu_{i,t}$. Subtracting the equation from the inequality proves the second conclusion. \square

5.2 Robust response

Another desirable property in many settings is robustness: a player's rating should not change too much in response to any one contest, no matter how extreme their performance. The Codeforces and TrueSkill systems lack this property, allowing for unbounded rating changes. Topcoder achieves robustness by clamping any changes that exceed a cap, which is initially high for new players but decreases with experience.

When $\rho > 0$, Elo-MMR(ρ) achieves robustness in a natural, smoother manner. To understand how, we look at the interplay between Gaussian and logistic factors in the posterior. Recall the notation in Equation (10), describing the loss function and weights.

THEOREM 5.7. *In the Elo-MMR(ρ) rating system, let*

$$\Delta_+ := \lim_{\rho_t \rightarrow +\infty} \mu_t - \mu_{t-1}, \quad \Delta_- := \lim_{\rho_t \rightarrow -\infty} \mu_{t-1} - \mu_t.$$

Then, for $\Delta_{\pm} \in \{\Delta_+, \Delta_-\}$,

$$\frac{\pi}{\beta_t \sqrt{3}} \left(w_0 + \frac{\pi^2}{6} \sum_{k \in \mathcal{H}_{t-1}} w_k \right)^{-1} \leq \Delta_{\pm} \leq \frac{\pi}{\beta_t \sqrt{3}} \frac{1}{w_0}.$$

PROOF. Using the fact that $0 < \frac{d}{dx} \tanh(x) \leq 1$, differentiating L' in Equation (10) yields

$$\forall s \in \mathbb{R}, w_0 \leq L''(s) \leq w_0 + \frac{\pi^2}{6} \sum_{k \in \mathcal{H}_{t-1}} w_k.$$

Now, the performance at round t adds a new term with multiplicity one to $L'(s)$: its value is $\frac{\pi}{\beta_k \sqrt{3}} \tanh \left(\frac{(s-p_k)\pi}{\beta_k \sqrt{12}} \right)$.

As a result, for every $s \in \mathbb{R}$, in the limit as $\rho_t \rightarrow \pm\infty$, $L'(s)$ increases by $\mp \frac{\pi}{\beta_t \sqrt{3}}$. Since μ_{t-1} was a zero of L' without this new term, we now have $L'(\mu_{t-1}) \rightarrow \mp \frac{\pi}{\beta_t \sqrt{3}}$. Dividing by the former inequalities yields the desired result. \square

The proof reveals that the magnitude of Δ_{\pm} depends inversely on that of L'' in the vicinity of the current rating, which in turn is related to the derivative of the tanh terms. If a player's performances vary wildly, the tanh terms will be widely dispersed, so any potential rating value will necessarily be in the tail ends of most of the terms. Tails contribute very small derivatives, enabling a larger rating change. Conversely, the tanh terms of a player with a very consistent performance history will contribute large derivatives, so the bound on their rating change will be small.

Thus, Elo-MMR naturally caps the rating changes of all players, and the cap is smaller for consistent performers. The cap will increase after an extreme performance, providing a similar “momentum” to the Topcoder and Glicko-2 systems, but without sacrificing incentive-compatibility (Theorem 5.5).

We can compare the lower and upper bound in Theorem 5.7: their ratio is on the same order as the fraction of the total weight that is held by the normal term. Recall that ρ is the weight transfer rate: larger ρ results in more weight being transferred into w_0 ; in this case, the lower and upper bound tend to stay close together. Conversely, the momentum effect is more pronounced when ρ is small. In the extreme case $\rho = 0$, w_0 vanishes for experienced players, so a sufficiently volatile player would be subject to correspondingly large rating updates. In the extended version of this paper,

we quantify an asymptotic steady state for the weights, and argue that $1/\rho$ can be thought of as a momentum parameter.

5.3 Runtime analysis and optimizations

Let’s look at the computation time needed to process a round with participant set \mathcal{P} , where we again omit the round subscript. Each player i has a participation history \mathcal{H}_i .

Estimating P_i entails finding the zero of a monotonic function with $O(|\mathcal{P}|)$ terms, and then obtaining the rating μ_i entails finding the zero of another monotonic function with $O(|\mathcal{H}_i|)$ terms. Using either of the Illinois or Newton methods, solving these equations to precision ϵ takes $O(\log \log \frac{1}{\epsilon})$ iterations. As a result, the total runtime needed to process one round of competition is

$$O\left(\sum_{i \in \mathcal{P}} (|\mathcal{P}| + |\mathcal{H}_i|) \log \log \frac{1}{\epsilon}\right).$$

This complexity is more than adequate for Codeforces-style competitions with thousands of contestants and history lengths up to a few hundred. Indeed, we were able to process the entire history of Codeforces on a small laptop in less than half an hour. Nonetheless, it may be cost-prohibitive in truly massive settings, where $|\mathcal{P}|$ or $|\mathcal{H}_i|$ number in the millions. Fortunately, it turns out that both functions may be compressed down to a bounded number of terms, with negligible loss of precision.

Adaptive subsampling. In Section 2, we used Doob’s consistency theorem to argue that our estimate for P_i is consistent. Specifically, we saw that $O(1/\epsilon^2)$ opponents are needed to get the typical error below ϵ . Thus, we can subsample the set of opponents to include in the estimation, omitting the rest. Random sampling is one approach. A more efficient approach chooses a fixed number of opponents whose ratings are closest to that of player i , as these are more likely to provide informative match-ups. On the other hand, if the setting requires incentive-compatibility to hold exactly, then one must avoid choosing different opponents for each player.

History compression. Similarly, it’s possible to bound the number of stored factors in the posterior. Our skill-evolution algorithm decays the weights of old performances at an exponential rate. Thus, the contributions of all but the most recent $O(\log \frac{1}{\epsilon})$ terms are negligible. Rather than erase the older logistic terms outright, we recommend replacing them with moment-matched Gaussian terms, similar to the transfers in Section 4 with $\kappa_t = 0$. Since Gaussians compose easily, a single term can then summarize an arbitrarily long prefix of the history.

Substituting $1/\epsilon^2$ and $\log \frac{1}{\epsilon}$ for $|\mathcal{P}|$ and $|\mathcal{H}_i|$, respectively, the runtime of Elo-MMR with both optimizations becomes

$$O\left(\frac{|\mathcal{P}|}{\epsilon^2} \log \log \frac{1}{\epsilon}\right).$$

If the contests are *extremely large*, so that $\Omega(1/\epsilon^2)$ opponents have a rating and uncertainty in the same ϵ -width bucket as player i , then it’s possible to do even better: up to the allowed precision ϵ , the corresponding terms can be treated as duplicates. Hence, their sum can be determined by counting how many of these opponents win, lose, or tie against player i . Given the pre-sorted list of ranks of players in the bucket, two binary searches would yield the answer.

Dataset	# contests	avg. # participants / contest
Codeforces	1087	2999
Topcoder	2023	403
Reddit	1000	20
Synthetic	50	2500

Table 1: Summary of test datasets.

In practice, a single bucket might not contain enough participants, so we sample enough buckets to yield the desired precision.

Simple parallelism. Since each player’s rating computation is independent, the algorithm is embarrassingly parallel. Threads can read the same global data structures, so each additional thread contributes only $O(1)$ memory overhead.

6 EXPERIMENTS

In this section, we compare various rating systems on real-world datasets, mined from several sources that will be described in Section 6.1. The metrics are runtime and predictive accuracy, as described in Section 6.2. Implementations of all rating systems, dataset mining, and additional processing used in our experiments can be found at <https://github.com/EbTech/Elo-MMR>.

We compare Elo-MM χ and Elo-MMR(ρ) against the industry-tested rating systems of Codeforces and Topcoder. For a fairer comparison, we hand-coded efficient versions of all four algorithms in the safe subset of Rust, parallelized using the Rayon crate; as such, the Rust compiler verifies that they contain no data races [32]. Our implementation of Elo-MMR(ρ) makes use of the optimizations in Section 5.3, bounding both the number of sampled opponents and the history length by 500. In addition, we test the improved TrueSkill algorithm of [30], basing our code on an open-source implementation of the same algorithm. The inherent sequentiality of its message-passing procedure prevented us from parallelizing it. All experiments were run on a 2.0 GHz 24-core Skylake machine with 24 GB of memory.

Hyperparameter search. To ensure fair comparisons, we ran a separate grid search for each triple of algorithm, dataset, and metric, over all of the algorithm’s hyperparameters. The hyperparameter set that performed best on the first 10% of the dataset, was then used to test the algorithm on the remaining 90% of the dataset.

6.1 Datasets

Due to the scarcity of public domain datasets for rating systems, we mined three datasets to analyze the effectiveness of our system. The datasets were mined using data from each source website’s inception up to October 9th, 2020. We also created a synthetic dataset to test our system’s performance when the data generating process matches our theoretical model. Summary statistics of the datasets are presented in Table 1.

Codeforces contest history. This dataset contains the current entire history of rated contests ever run on codeforces.com, the dominant platform for online programming competitions. The Codeforces platform has over 850K users, over 300K of whom are rated, and has hosted over 1000 contests to date. Each contest has a couple thousand participants on average. A typical contest takes 2 to 3 hours and contains 5 to 8 problems. Players are ranked by total

points, with more points typically awarded for tougher problems and for early solves. They may also attempt to “hack” one another’s submissions for bonus points, identifying test cases that break their solutions.

Topcoder contest history. This dataset contains the current entire history of algorithm contests ever run on the topcoder.com. Topcoder is a predecessor to Codeforces, with over 1.4 million total users and a long history as a pioneering platform for programming contests. It hosts a variety of contest types, including over 2000 algorithm contests to date. The scoring system is similar to Codeforces, but its rounds are shorter: typically 75 minutes with 3 problems.

SubredditSimulator threads. This dataset contains data scraped from the current top-1000 most upvoted threads on the website [reddit.com/r/SubredditSimulator/](https://www.reddit.com/r/SubredditSimulator/). Reddit is a social news aggregation website with over 300 million users. The site itself is broken down into sub-sites called subreddits. Users then post and comment to the subreddits, where the posts and comments receive votes from other users. In the subreddit SubredditSimulator, users are language generation bots trained on text from other subreddits. Automated posts are made by these bots to SubredditSimulator every 3 minutes, and real users of Reddit vote on the best bot. Each post (and its associated comments) can thus be interpreted as a round of competition between the bots who commented.

Synthetic data. This dataset contains 10K players, with skills and performances generated according to the Gaussian generative model in Section 2. Players’ initial skills are drawn i.i.d. with mean 1500 and variance 350^2 . Players compete in all rounds, and are ranked according to independent performances with variance 200^2 . Between rounds, we add i.i.d. Gaussian increments with variance 35^2 to each of their skills.

6.2 Evaluation metrics

To compare the different algorithms, we define two measures of predictive accuracy. Each metric will be defined on individual contestants in each round, and then averaged:

$$\text{aggregate}(\text{metric}) := \frac{\sum_t \sum_{i \in \mathcal{P}_t} \text{metric}(i, t)}{\sum_t |\mathcal{P}_t|}.$$

Pair inversion metric [24]. Our first metric computes the fraction of opponents against whom our ratings predict the correct pairwise result, defined as the higher-rated player either winning or tying:

$$\text{pair_inversion}(i, t) := \frac{\# \text{ correctly predicted matchups}}{|\mathcal{P}_t| - 1} \times 100\%.$$

This metric was used in the original evaluation of TrueSkill [24].

Rank deviation. Our second metric compares the rankings with the total ordering that would be obtained by sorting players according to their prior rating. The penalty is proportional to how much these ranks differ for player i :

$$\text{rank_deviation}(i, t) := \frac{|\text{actual rank} - \text{predicted rank}|}{|\mathcal{P}_t| - 1} \times 100\%.$$

In the event of ties, among the ranks within the tied range, we use the one that comes closest to the rating-based prediction.

6.3 Empirical results

Recall that Elo-MM χ has a Gaussian performance model, matching the modeling assumptions of Topcoder and TrueSkill. Elo-MMR(ρ), on the other hand, has a logistic performance model, matching the modeling assumptions of Codeforces and Glicko. While ρ was included in the hyperparameter search, in practice we found that all values between 0 and 1 produce very similar results.

To ensure that errors due to the unknown skills of new players don’t dominate our metrics, we excluded players who had competed in less than 5 total contests. In most of the datasets, this reduced the performance of our method relative to the others, as our method seems to converge more accurately. Despite this, we see in Table 2 that both versions of Elo-MMR outperform the other rating systems in both the pairwise inversion metric and the ranking deviation metric.

We highlight a few key observations. First, significant performance gains are observed on the Codeforces and Topcoder datasets, despite these platforms’ rating systems having been designed specifically for their needs. Our gains are smallest on the synthetic dataset, for which all algorithms perform similarly. This might be in part due to the close correspondence between the generative process and the assumptions of these rating systems. Furthermore, the synthetic players compete in all rounds, enabling the system to converge to near-optimal ratings for every player. Finally, the improved TrueSkill performed well below our expectations, despite our best efforts to improve it. We suspect that the message-passing numerics break down in contests with a large number of individual participants. The difficulties persisted in all TrueSkill implementations that we tried, including on Microsoft’s popular Infer.NET framework [29]. To our knowledge, we are the first to present experiments with TrueSkill on contests where the number of participants is in the hundreds or thousands. In preliminary experiments, TrueSkill and Elo-MMR score about equally when the number of ranks is less than about 60.

Now, we turn our attention to Table 3, which showcases the computational efficiency of Elo-MMR. On smaller datasets, it performs comparably to the Codeforces and Topcoder algorithms. However, the latter suffer from a quadratic time dependency on the number of contestants; as a result, Elo-MMR outperforms them by almost an order of magnitude on the larger Codeforces dataset.

Finally, in comparisons between the two Elo-MMR variants, we note that while Elo-MMR(ρ) is more accurate, Elo-MM χ is always faster. This has to do with the skill drift modeling described in Section 4, as every update in Elo-MMR(ρ) must process $O(\log \frac{1}{\epsilon})$ terms of a player’s competition history.

7 CONCLUSIONS

This paper introduces the Elo-MMR rating system, which is in part a generalization of the two-player Glicko system, allowing any number of players. By developing a Bayesian model and taking the limit as the number of participants goes to infinity, we obtained simple, human-interpretable rating update formulas. Furthermore, we saw that the algorithm is incentive-compatible, robust to extreme performances, asymptotically fast, and embarrassingly parallel. To our knowledge, our system is the first to rigorously prove all these properties in a setting with more than two individually

Dataset	Codeforces		Topcoder		TrueSkill		Elo-MM χ		Elo-MMR(ρ)	
	pair inv.	rank dev.	pair inv.	rank dev.	pair inv.	rank dev.	pair inv.	rank dev.	pair inv.	rank dev.
Codeforces	78.3%	14.9%	78.5%	15.1%	61.7%	25.4%	78.5%	14.8%	78.6%	14.7%
Topcoder	72.6%	18.5%	72.3%	18.7%	68.7%	20.9%	73.0%	18.3%	73.1%	18.2%
Reddit	61.5%	27.3%	61.4%	27.4%	61.5%	27.2%	61.6%	27.3%	61.6%	27.3%
Synthetic	81.7%	12.9%	81.7%	12.8%	81.3%	13.1%	81.7%	12.8%	81.7%	12.8%

Table 2: Performance of each rating system on the pairwise inversion and rank deviation metrics. Bolded entries denote the best performances (highest pair inv. or lowest rank dev.) on each metric and dataset.

Dataset	CF	TC	TS	Elo-MM χ	Elo-MMR(ρ)
Codeforces	212.9	72.5	67.2	31.4	35.4
Topcoder	9.60	4.25	16.8	7.00	7.52
Reddit	1.19	1.14	0.44	1.14	1.42
Synthetic	3.26	1.00	2.93	0.81	0.85

Table 3: Total compute time over entire dataset, in seconds.

ranked players. In terms of practical performance, we saw that it outperforms existing industry systems in both prediction accuracy and computation speed.

This work can be extended in several directions. First, the choices we made in modeling ties, pseudodiffusions, and opponent subsampling are by no means the only possibilities consistent with our Bayesian model of skills and performances. Second, it may be possible to further improve accuracy by fitting more flexible performance and skill evolution models to application-specific data.

Another useful extension would be to team competitions. Given a performance model for teams, Elo-MMR infers each team’s performance. To make this useful in settings where teams are frequently reassigned, we must model teams in terms of their individual members; unfortunately, it’s not possible to precisely infer an individual’s performance from team rankings alone. Therefore, it becomes necessary to condition an individual’s skill on their team’s performance. In the case where a team’s performance is modeled as the sum of its members’ independent Gaussian contributions, elementary facts about multivariate Gaussian distributions enable posterior skill inferences at the individual level. Generalizing this approach to other models remains an open challenge.

Over the past decade, online competition communities such as Codeforces have grown exponentially. As such, considerable work has gone into engineering scalable and reliable rating systems. Unfortunately, many of these systems have not been rigorously analyzed in the academic community. We hope that our paper and open-source release will open new explorations in this area.

ACKNOWLEDGEMENTS

The authors are indebted to Daniel Sleator and Danica J. Sutherland for initial discussions that helped inspire this work, and to Nikita Gaevoy for the open-source improved TrueSkill upon which our implementation is based. Experiments in this paper are funded by a Google Cloud Research Grant. The second author is supported by a VMware Fellowship and the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

- [1] CodeChef Rating Mechanism. codechef.com/ratings
- [2] Codeforces: Results of 2019. codeforces.com/blog/entry/73683
- [3] Farming Volatility: How a major flaw in a well-known rating system takes over the GBL leaderboard. reddit.com/r/TheSilphRoad/comments/hwff2d/farming_volatility_how_a_major_flaw_in_a/
- [4] Halo Xbox video game franchise: in numbers. telegraph.co.uk/technology/video-games/11223730/Halo-in-numbers.html
- [5] Kaggle milestone: 5 million registered users! kaggle.com/general/164795
- [6] Kaggle Progression System. kaggle.com/progression
- [7] LeetCode New Contest Rating Algorithm. [leetcode.com/discuss/general-discussion/468851/New-Contest-Rating-Algorithm-\(Coming-Soon\)](https://leetcode.com/discuss/general-discussion/468851/New-Contest-Rating-Algorithm-(Coming-Soon))
- [8] Open Codeforces Rating System. codeforces.com/blog/entry/20762
- [9] Topcoder Algorithm Competition Rating System. topcoder.com/community/competitive-programming/how-to-compete/ratings
- [10] Why Are Obstacle-Course Races So Popular? theatlantic.com/health/archive/2018/07/why-are-obstacle-course-races-so-popular/565130/
- [11] Sharad Agarwal and Jacob R. Lorch. 2009. Matchmaking for online games and other latency-sensitive P2P systems. In *SIGCOMM 2009*. 315–326.
- [12] Mark Yuying An. 1997. Log-concave probability distributions: Theory and statistical testing. (1997).
- [13] Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika* (1952), 324–345.
- [14] Shuo Chen and Thorsten Joachims. 2016. Modeling Intransitivity in Matchup and Comparison Data. In *WSDM 2016*. 227–236.
- [15] Rémi Coulom. [n.d.]. Whole-history rating: A Bayesian rating system for players of time-varying strength. In *CG 2008*. Springer, 113–124.
- [16] Pierre Dangauthier, Ralf Herbrich, Tom Minka, and Thore Graepel. 2007. TrueSkill Through Time: Revisiting the History of Chess. In *NeurIPS 2007*. 337–344.
- [17] Arpad E. Elo. 1961. New USCF rating system. *Chess Life* (1961), 160–161.
- [18] RNDr Michal Forišek. 2009. Theoretical and Practical Aspects of Programming Contest Ratings. (2009).
- [19] David A Freedman. 1963. On the asymptotic behavior of Bayes’ estimates in the discrete case. *The Annals of Mathematical Statistics* (1963), 1386–1403.
- [20] Mark E Glickman. 1995. A comprehensive guide to chess ratings. *American Chess Journal* (1995), 59–102.
- [21] Mark E Glickman. 1999. Parameter estimation in large dynamic paired comparison experiments. *Applied Statistics* (1999), 377–394.
- [22] Mark E Glickman. 2012. Example of the Glicko-2 system. *Boston University* (2012), 1–6.
- [23] Linxia Gong, Xiaochuan Feng, Dezhi Ye, Hao Li, Runze Wu, Jianrong Tao, Changjie Fan, and Peng Cui. 2020. OptMatch: Optimized Matchmaking via Modeling the High-Order Interactions on the Arena. In *KDD 2020*. 2300–2310.
- [24] Ralf Herbrich, Tom Minka, and Thore Graepel. 2006. TrueSkillTM: A Bayesian Skill Rating System. In *NeurIPS 2006*. 569–576.
- [25] Tzu-Kuo Huang, Chih-Jen Lin, and Ruby C. Weng. 2006. Ranking individuals by group comparisons. In *ICML 2006*. 425–432.
- [26] Stephanie Kovalchik. 2020. Extension of the Elo rating system to margin of victory. *Int. J. Forecast.* (2020).
- [27] Yao Li, Minhao Cheng, Kevin Fujii, Fushing Hsieh, and Cho-Jui Hsieh. 2018. Learning from Group Comparisons: Exploiting Higher Order Interactions. In *NeurIPS 2018*. 4986–4995.
- [28] Tom Minka, Ryan Clevon, and Yordan Zaykov. 2018. *TrueSkill 2: An improved Bayesian skill rating system*. Technical Report MSR-TR-2018-8. Microsoft.
- [29] T. Minka, J.M. Winn, J.P. Guiver, Y. Zaykov, D. Fabian, and J. Bronskill. /Infer.NET 0.3. Microsoft Research Cambridge. <http://dotnet.github.io/infer>.
- [30] Sergey I. Nikolenko, Alexander, and V. Sirotkin. 2010. Extensions of the TrueSkill TM rating system. In *In Proceedings of the 9th International Conference on Applications of Fuzzy Systems and Soft Computing*. 151–160.
- [31] Jerneja Premelč, Goran Vučković, Nic James, and Bojan Leskošek. 2019. Reliability of judging in DanceSport. *Front. Psychol.* (2019), 1001.
- [32] Josh Stone and Nicholas D Matsakis. The Rayon library (Rust Crate). crates.io/crates/rayon
- [33] Ruby C. Weng and Chih-Jen Lin. 2011. A Bayesian Approximation Method for Online Ranking. *J. Mach. Learn. Res.* (2011), 267–300.
- [34] John Michael Winn. 2019. *Model-based machine learning*.
- [35] Lin Yang, Stanko Dimitrov, and Benny Mantin. 2014. Forecasting sales of new virtual goods with the Elo rating system. *RPM* (2014), 457–469.