COMMENTS ON NORMALIZED ELO

MICHEL VAN DEN BERGH

1. Introduction

Normalized Elo is introduced here [4]. The primary motivation for normalized Elo is that it is a measure for the amount of games it takes to prove that one engine is stronger than another, with a given level of significance. In other words it is an *objective measure* of strength difference.

In this document we make some cosmetic changes to the terminology introduced in loc. cit. In particular what was called "normalized Elo" will now be called "normalized t-value" and we redefine normalized Elo as the normalized t-value multiplied by an appropriate normalization constant. This is done to make the comparison with ordinary logistic Elo more intuitive.

2. Background

The *normalized t-value* for the strength difference of two engines is defined as

$$\boxed{t_n := \frac{\mu - 1/2}{\sigma_{\mathrm{pg}}}}$$

where μ is the expected score and σ_{pg} is the standard deviation of the expected score *per game*. In the trinomial case σ_{pg} is the standard deviation of the outcome distribution of a game, scored as $0, 1/2, 1$. In the pentanomial case, σ_{pg} is the standard deviation of the outcome distribution multiplied by $\sqrt{2}$, where we score the outcome of a game pair as $0, 1/4, 2/4, 3/4, 1$.

The justification for this convention is that, whatever testing system we use, the normalized t-value $\hat{t}_n$ of a test is defined to be the usual t-value divided by the square root of the number of games. Then t_n is the asymptotic expectation value of $\hat{t}_n$. More precisely, asymptotically we have

$$\boxed{\hat{t}_n \sim N(t_n, \frac{1}{N})}$$

where N is the number of games.

In the trinomial case or in the pentanomial case with a perfectly balanced book we have

$$\sigma_{\mathrm{pg}} = \frac{1}{2}\sqrt{1-d} \tag{2.1}$$

where d is the draw ratio.

3. Normalization

Below we put

$$S(x) = \frac{1}{1 + 10^{-x/400}}$$

This is the function which converts ordinary ("logistic") Elo into an expected score. It is convenient to write

$$S(x) = L(\beta x)$$

where $\beta = \log(10)/400$ and L is the usual logistic function

$$L(x) = \frac{1}{1 + e^{-x}}$$

L satisfies the functional equation

$$L'(x) = L(x)(1 - L(x))$$

Let

$$\boxed{C_{e/t} := \frac{2}{\beta} = \frac{800}{\log(10)} \cong 347.43}$$

We claim that for small Elo differences we have

$$t_n \cong \frac{1}{C_{e/t}} \frac{e_l}{2\sigma_{\text{pg}}} \tag{3.1}$$

where e_l is the logistic Elo difference between two engines. To see this note

$$e_l \cong (s - S(0))/S'(0) = (s - 1/2)/(\beta L'(0)) = (s - 1/2)/(1/2(1 - 1/2)\beta) = 4(s - 1/2)/\beta$$

Hence

$$\frac{1}{C_{e/t}} \frac{e_l}{2\sigma_{\text{pg}}} \cong \frac{\beta}{2} \frac{4}{\beta} \frac{s - 1/2}{2\sigma_{\text{pg}}} = \frac{s - 1/2}{\sigma_{\text{pg}}} = t_n$$

We define the *normalized Elo difference* between two engines as

$$\boxed{e_n := C_{e/t} t_n} \tag{3.2}$$

In case of a perfectly balanced book it follows from (2.1) and (3.1) that

$$\boxed{e_n \cong \frac{e_l}{\sqrt{1 - d}}} \tag{3.3}$$

This simple formula is the motivation for the normalization introduced in (3.2). We see in particular that for $d = 0$ normalized Elo and logistic Elo coincide. For other draw ratios we have the following conversion table.

Draw ratio	0.00	0.30	0.50	0.60	0.70	0.80	0.90
Normalized Elo	5.00	5.00	5.00	5.00	5.00	5.00	5.00
Logistic Elo	5.00	4.18	3.54	3.16	2.74	2.24	1.58

Let us now discuss the duration of an SPRT test for H0:$e_n = e_{n,0}$ versus $H1 : e_n = e_{n,1}$. Under the assumption that the the Type I/II error probabilities are given by $\alpha = \beta = 0.05$ we get that the worst case expectation duration (which corresponds to the actual Elo being half way between H0 and H1) of the test is given by

$$\boxed{T=\frac{D}{(e_{n,1}-e_{n,0})^2}} \tag{3.4}$$

where

$$\boxed{D:=C_{e/t}^2\log(19)^2\cong 1046535}$$

This leads to the following table

Normalized Elo difference	1	2	3	4	5	6
Expected duration	1046535	261634	116282	65408	41861	29070

Note that D is close to 1000000 which is sufficiently accurate for back of the envelope calculations.

Let us derive the formula (3.4). We may equivalently consider an SPRT of $t_n=t_{n,0}$ versus $t_n=t_{n,1}$. Let us suppose that σ_{pg} is known (see §4.2 below for a discussion) so it is sufficient to consider an SPRT test for $\mu=s_0$ versus $\mu=s_1$, for suitable s_0, s_1. According to [6] the expected duration of such a test, when the actual score is μ is equal to

$$T=\frac{T(h_\mu)}{w^2}$$

where

$$\begin{aligned}
w&=\frac{s_1-s_0}{\sigma_{\mathrm{pg}}}=t_{n,1}-t_{n,0}\\
h_\mu&=\frac{2\mu-(s_0+s_1)}{s_1-s_0}=\frac{2t_n-(t_{n,0}+t_{n,1})}{t_{n,1}-t_{n,0}}\\
T(h)&=\frac{2b}{h}\frac{1-e^{-hb}}{1+e^{-hb}}\\
b&=\log\left(\frac{1-\alpha}{\alpha}\right)
\end{aligned}$$

when the Type I/II error probabilities are both equal to α.

The worst case is given when $t_n=(t_{n,0}+t_{n,1})/2$. In that case $\mu=(s_0+s_1)/2$ and hence $h_\mu=0$. Applying l'Hôpital's rule we find

$$T=\frac{T(0)}{(t_{n,1}-t_{n,0})^2}=\frac{b^2}{(t_{n,1}-t_{n,0})^2}=\frac{C_{e/t}^2 b^2}{(e_{n,1}-e_{n,0})^2}$$

If $\alpha=0.05$ then $b=\log(19)$ and we obtain (3.4).

For completeness we note that in case $t_n=t_{n,0}$ or $t_n=t_{n,1}$ the expected duration is given by a similar formula as (3.4) where the numerator is replaced by

$$\boxed{D':=\frac{9C_{e/t}^2\log(19)}{5}\cong 639770}$$

4. LLR computation

What we call an SPRT is strictly speaking a GSPRT [7] which is based on monitoring the Generalized Log Likelihood Ratio (which we denote by LLR below) of H1 versus H0. See [2] for an introduction.

4.1. **The exact LLR.** Assume given real numbers

$$a_1 < a_2 < \cdots < a_l$$

and a discrete probability distribution

$$P : \{a_1, \ldots, a_l\} \to \mathbb{R} : a_i \mapsto p_i\,.$$

with mean μ and standard deviation σ. Assume a sample taken from $\{a_1, \ldots, a_N\}$ according to P has sample distribution $(\hat{p}_i)_{i=1,\ldots,N}$. Let μ_{ref} be some reference value. Put

$$t = \frac{\mu - \mu_{\text{ref}}}{\sigma}, \tag{4.1}$$

We will give a numerical procedure to compute the MLE for the $(p_i)_i$ given the empirical distribution $\hat{p}$, subject to the constraint $t = t_*$ for a given t_*. Although in practice this procedure appears to converge rapidly, as confirmed by simulation, we mention the following caveats:

- we have not proved that the MLE is unique;
- we have not proved convergence, rapidly or not.

Anyway, keeping this in mind, we explain the procedure. We must maximize

$$\sum_i \hat{p}_i \log p_i \tag{4.2}$$

subject to

$$\sum_i p_i = 1 \tag{4.3}$$

$$\mu - \mu_{\text{ref}} - t_*\sigma = 0 \tag{4.4}$$

Let $(m_i)_{i\geq 0}$ be the moments of P. We rewrite (4.4) as

$$\phi(P) := m_1 - \mu_{\text{ref}}\, m_0 - t_*\sqrt{m_2 m_0 - m_1^2} = 0 \tag{4.5}$$

The distinguishing feature of ϕ is that it is homogeneous of degree 1 in $(p_i)_i$. We will now assume that $\phi(P)$ is an arbitrary expression in $(p_i)_i$, homogeneous of degree $\kappa \neq 0$. To compute the extremal value(s?) of (4.2) subject to (4.3) and the condition $\phi(P) = 0$ we use Lagrange multipliers. That is we need to solve $\partial\tilde{\phi}/\partial p_i = 0$ where

$$\tilde{\phi}(P) = \sum_i \hat{p}_i \log p_i - \lambda\Big(\sum_i p_i - 1\Big) - \theta\phi(P)$$

where in addition (4.3) and $\phi(P) = 0$ are satisfied. Or in other words

$$\frac{\hat{p}_i}{p_i} = \lambda + \theta\phi_i(P) = 0, \qquad i = 1, \ldots, l \tag{4.6}$$

for $\phi_i = \partial\phi/\partial p_i$. For use below we note the Euler identity

$$\sum_i p_i\phi_i(P) = \kappa\phi(P) \tag{4.7}$$

Rewriting (4.6) as

$$\hat{p}_i = p_i(\lambda + \theta\phi_i(P))$$

and summing over i, using (4.3), $\phi(P) = 0$ and (4.7) we obtain $\lambda = 1$. Conversely if $\lambda = 1$ then $\sum_i p_i = 1$. In other words we are reduced to solving the following system:

$$\begin{aligned} p_i(1 + \theta\phi_i(P)) &= \hat{p}_i \qquad i = 1, \ldots, l \\ \sum_i p_i\phi_i(P) &= 0 \end{aligned} \tag{4.8}$$

where we have used (4.7) again (here we use $\kappa \neq 0$ to have that (4.8) implies $\phi(P) = 0$). This suggests the following numerical procedure for solving (4.8).

Assume we have an estimate P_n for the MLE distribution. Then determine θ_{n+1} such that

$$\boxed{\sum_i \frac{\hat{p}_i\phi_i(P_n)}{1 + \theta_{n+1}\phi_i(P_n)} = 0} \tag{4.9}$$

and put

$$\boxed{p_{n+1,i} = \frac{\hat{p}_i}{1 + \theta_{n+1}\phi_i(P_n)}, \qquad i = 1, \ldots, l} \tag{4.10}$$

Note that automatically $\sum_i p_{n+1,i} = 1$ (excercise). In order to have $0 \leq p_{n+1,i}$ we should only consider solutions of (4.9) satisfying.

$$-1/v \leq \theta_{n+1} \leq -1/u$$

where $u = \min_i \phi_i(P_n)$, $v = \max_i \phi_i(P_n)$ (such solutions are unique). Furthermore in order for (4.9) to have a solution we also should have $uv < 0$.

In the case where $\phi(P)$ is as in (4.5) then we obtain

$$\boxed{\phi_i(P) = a_i - \mu_{\mathrm{ref}} - \frac{1}{2}t_*\sigma\left(1 + \left(\frac{a_i - \mu}{\sigma}\right)^2\right)}$$

so we should use this expression in (4.9) and (4.10).

The question remains what we should take for P_0. An obvious choice is $P_0 = \hat{p}$ but then it sometimes happens that the condition $uv < 0$ is not satisfied. A much safer choice seems to be a uniform distribution. I.e. $\forall i : p_i = 1/l$.

4.2. **An approximation.** In [5] a relatively elegant method was given to compute the LLR for an SPRT for the mean of a multinomial distribution. In fact this can be obtained from the approach in §4.1 by taking $\phi(P) = m_1 - s_* m_0$.

In [5] an approximation for the LLR was derived and in [3] this was compared to the exact one. It seems a good strategy to use this approximation if in addition we estimate σ (necessary to convert the mean into a t-value) from the test itself.

Then by [5] the LLR for an SPRT of H0:$\mu = s_0$ versus H1:$\mu = s_1$ may be approximated by

$$\mathrm{LLR} \cong \frac{n}{2}\log\left(\frac{\sum_i \hat{p}_i(s_0 - a_i)^2}{\sum_i \hat{p}_i(s_1 - a_i)^2}\right) = \frac{n}{2}\log\left(\frac{(\hat{\mu} - s_0)^2 + \hat{\sigma}^2}{(\hat{\mu} - s_1)^2 + \hat{\sigma}^2}\right) \tag{4.11}$$

where n is the sample size. Put

$$\hat{t} = \frac{\hat{\mu} - \mu_{\text{ref}}}{\hat{\sigma}}, \qquad t_0 = \frac{s_0 - \mu_{\text{ref}}}{\sigma}, \qquad t_1 = \frac{s_1 - \mu_{\text{ref}}}{\sigma} \tag{4.12}$$

Remark 4.1. In the context of the pentanomial model described above, $\mu_{\text{ref}} = 1/2$, $n = N/2$ and the t-values in (4.12) are the normalized t-values multiplied by $\sqrt{2}$. In the trinomial case the t-values coincide with the normalized ones.

Assuming that $\hat{\sigma}$ is a good approximation for σ we find from (4.11)

$$\text{LLR} \cong \frac{n}{2} \log\left(\frac{1 + (\hat{t} - t_0)^2}{1 + (\hat{t} - t_1)^2}\right) \tag{4.13}$$

It is however a bit inelegant that (4.13) does not reduce to the corresponding trinomial approximation, even with a perfectly balanced book. Taking advantage of the fact that in practice $\hat{t}$, t_0 and t_1 will be small compared to 1 and combining Remark 4.1 with the fact that for small x we have $1 + 2x \cong (1 + x)^2$ we arrive at our final formula

$$\boxed{\text{LLR} \cong \frac{N}{2} \log\left(\frac{1 + (\hat{t}_n - t_{n,0})^2}{1 + (\hat{t}_n - t_{n,1})^2}\right)} \tag{4.14}$$

which is valid both in the trinomial and pentanomial case. That (4.14) performs entirely satisfactorily is confirmed by simulation. See [1].

Remark 4.2. As $\hat{t}$, t_0 and t_1 will be small compared to 1, (4.14) can be further approximated by

$$\text{LLR} \cong \frac{N}{2}((\hat{t}_n - t_{n,0})^2 - (\hat{t}_n - t_{n,1})^2) = \frac{N}{2}(t_{n,1} - t_{n,0})(2\hat{t}_n - t_{n,0} - t_{n,1}) \tag{4.15}$$

This formula works just as well but it needs to be regularized in some way. Indeed at the beginning of a test (say after a few game pairs with identical outcomes) $\hat{\sigma}$ will still be very small[1] and hence $\hat{t}$ may be spuriously large. Then the same will be true for (4.15).

It is easy to see that for small $t_{n,0}$, $t_{n,1}$ the extremal values of the the log-factor in (4.14) are $\cong \pm(t_{n,1} - t_{n,0})$. This suggests an easy to use regularization rule which applies to (4.15) but also to other approximations: LLR/N *should be clamped to the interval* $[-(t_{n,1} - t_{n,0})/2, (t_{n,1} - t_{n,1})/2]$.

References

1. Michel Van den Bergh, *A multi-threaded pentanomial simulator in C*, `https://github.com/vdbergh/simul`.
2. ———, *A practical introduction to the GSPRT*, `http://hardy.uhasselt.be/Fishtest/GSPRT_approximation.pdf`.
3. ———, *Comparing the approximations for the Generalized Log Likelihood Ratio of a multinomial distribution*, `http://hardy.uhasselt.be/Fishtest/comparing_approximations.pdf`.
4. ———, *Normalized Elo*, `http://hardy.uhasselt.be/Fishtest/normalized_elo.pdf`.
5. ———, *The Generalized Likelihood Ratio for the expectation value of a distribution*, `http://hardy.uhasselt.be/Fishtest/support_MLE_multinomial.pdf`.
6. ———, *The SPRT for Brownian motion*, `http://hardy.uhasselt.be/Fishtest/sprta.pdf`.
7. Xiaoou Li, Jingchen Liu, and Zhiliang Ying, *Generalized Sequential Probability Ratio Test for Separate Families of Hypotheses*, `http://stat.columbia.edu/~jcliu/paper/GSPRT_SQA3.pdf`.

[1]In fact $\hat{\sigma} = 0$. However for simplicity we replace zero outcome frequencies with a small $\epsilon > 0$ to avoid division by zero.