Introducing TRC Computer Race Ratings

By James Willoughby

Computers are very, very good at handicapping racehorses. If a computer is taught everything a human knows about the process, a human has no shot whatsoever of producing a set of ratings that finds more winners than a computer. None.

In the old world, people used to cling to the concept that human intuition and intelligence could defeat computers, even at tasks that had a finite state of possibilities like chess and Go. But it just turned out that the computers had not been programmed adequately.

Humans are good at handicapping too. We know that because the ratings that computers produce aren't all that different from humans. But computers are faster, more accurate and never get tired.

Until computer vision is refined, humans are still needed to collect the data about races germane to the task of handicapping. For example, a human might tell a computer which horses were hampered – and by how much – and which horses were eased at the end of a race. To really sharpen computer handicapping, a human might encode variables for the ease a horse was travelling or how hard-ridden it was.

Humans can also be useful for providing computers with so-called 'priors' – estimates about things the computer has not otherwise been taught and which it benefits from as a starting point.

The problem with computers is that they are just dumb robots. They carry out tasks according to their programmer's instructions and no more. Although some impressive tasks that computers can carry out in the real world are labelled as 'intelligent', they aren't intelligence as we possess it. Yet. But computers can mimic learning by iterative processes and by chaining vast networks of simple processing units together capable of understanding massively complex patterns in data.

Teaching the *TRC* computer to handicap

This article is not going to descend into mathematical details, which are abstruse to most followers of the sport. While there is no avoiding them if you are building a handicap model, it is relatively easy to understand the concept of computer handicapping explained in words.

Maths is, however, a concise, beautiful language all of its own and talking maths to a computer bridges the communication gap, so here is how we started the process of teaching the *TRC* computer to handicap horses:

$$\Psi\ (\text{distance beaten}) \sim N\ (\mu, \sigma)$$

$$\mu = \text{horse rating - race quality}$$

$$\sigma^2 = \sigma^2_{\text{course}} + \sigma^2_{\text{distance}} + \sigma^2_{\text{going}}$$

where Ψ is a scoring function normalising the result of a race such that $\boldsymbol{\Sigma}$ *scores = 0. Also,* ${}^1/_n\ \boldsymbol{\Sigma}$ *horse ratings = 100.*

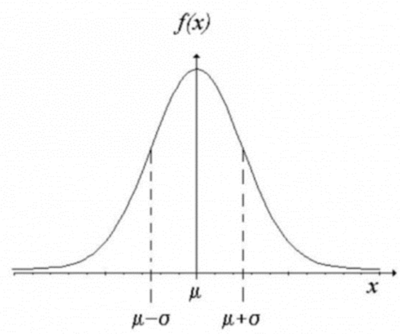

There were 156,906 performances in the *TRC* database, covering Group and Graded races from January 5, 2011, to February 9, 2021. First, the computer needs to know the obvious: the distance a horse is beaten is a function of the quality of a race and the racing merit of the horse. And it's the latter we are trying to rate.

We also need to tell it that horses don't always perform to the same level. Far from it. Their performance – measured in distance beaten where seconds or lengths are the units – actually follows a pattern called **the gamma distribution**. They come in at intervals behind the winner, with the greatest concentration coming just after the winner and the gaps tending to get wider as the back markers trail home.

It is interesting to ponder why this might be the case. A ten-runner horse race, for instance, isn't decided by ten individual time trials. If it were, we could describe the finishing times by a normal distribution, such as the one that describes adult male heights, for instance.

But a horserace is different because of tactics. In races where the winner goes for home in good time, lesser-talented rivals have to go hard in pursuit before it is really good for them. So, as you go down the field, each horse tends formerly prominent to run less efficiently if it were trying to win the race. It's a different matter with closers who run through beaten horses in their own time.

While the gamma distribution is well described in statistics, it can be a bit of a nightmare to deal with. The normal distribution is better behaved, and it is desirable for our purposes to take the distances behind the winner and somehow 'normalise them', that is, transform them mathematically so that when we stack up all the figures in the database, the pattern looks like adult male heights.

That's what the first line of the model above says to the computer. The fancy ψ is just the Greek letter Psi used here both to denote a 'normalising function'. It says: "Take the distances (or time lags) and manipulate them so that they follow a normal distribution. The tilde '~' is read "distributed as' and the N means the normal distribution.

But the model also says do something else. *Before* you normalise the distances, flip them about because our convention is that higher-rated horses are better and thus get beaten shorter distances. And turn lengths (or seconds) into rating points using a formula:

$$x = \frac{45\ .\ distance\ \ beaten}{distance\ \ beaten\ +12}$$

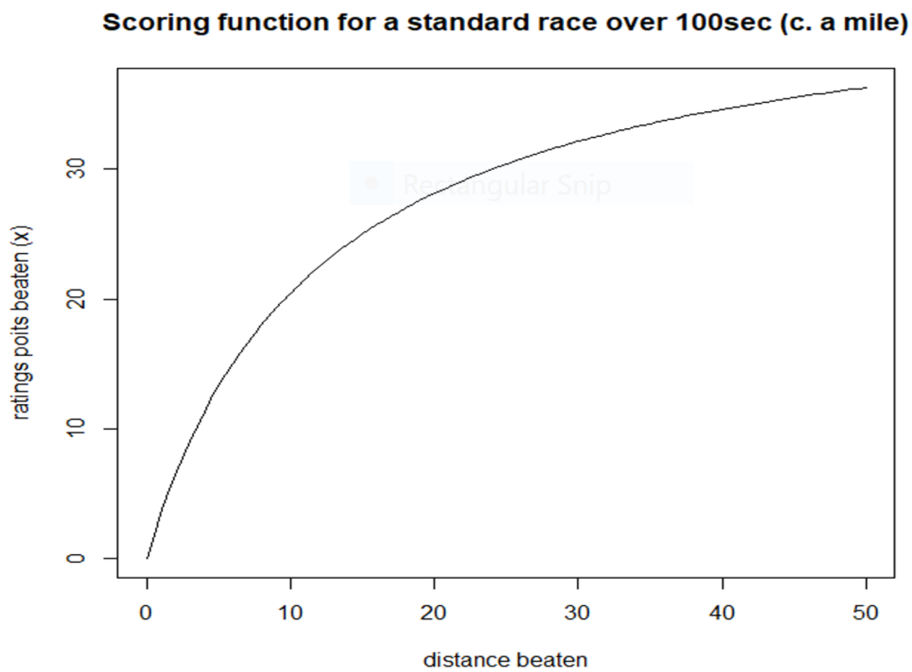

So, distance beaten is now a score which for the latest Breeders' Cup Mile looked like this:

Horse	DTW	PTW	SCORE
Order Of Australia IRE	0.00	0	17
Circus Maximus IRE	0.37	1	16
Lope Y Fernandez IRE	1.49	5	12
Ivar BRZ	2.99	9	8
Uni GBR	3.36	10	7
Halladay USA	4.11	11	6
Kameko USA	4.86	13	4
Factor This USA	5.98	15	2
Siskin USA	7.47	17	0
Raging Bull FRA	8.59	19	-2
Digital Age IRE	9.71	20	-3
Casa Creed USA	10.46	21	-4
March To The Arch USA	13.82	24	-7
Safe Voyage IRE	22.41	29	-12
	mean=>	14	

What we are doing is plugging the distance to the winner (**DTW**) into the above formula to produce rating points to the winner (**PTW**). Can you see the benefits of the formula? It expands the distances close behind the winner and shrinks those further away. Because the further behind a horse finishes, the less significant is a length or two. The 45 in the formula caps the maximum ratings points a horse can be defeated, while the ratio between 45 and 12 of just less than 4 determines how valuable the first length is (over a mile, remember.)

The final column takes the mean of **PTW**, which is 14, and subtracts each individual **PTW** from it. After this, the score is normalised (the process is beyond the scope of this article) and fixed in the database.

Leaning the effect of every course, distance and going on beaten lengths

So far, so good. Let's look at the way we taught the computer to handicap again:

$$\Psi \text{ (distance beaten)} \sim N(\mu, \sigma)$$
$$\mu = \text{horse rating - race quality}$$
$$\sigma^2 = \sigma^2_{\text{course}} + \sigma^2_{\text{distance}} + \sigma^2_{\text{going}}$$

where Ψ is a scoring function normalising the result of a race such that Σ *scores = 0. Also,* $^1/_n\ \Sigma$ *horse ratings = 100.*

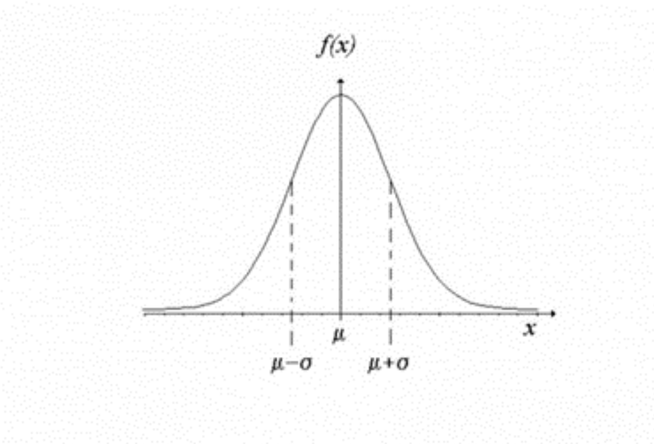

The picture to the right of the equations is a typical normal distribution. Its shape is described by two parameters, the mean μ (Greek letter Mu) which is the location, or middle, of the range of possible values, which itself is described by the standard deviation σ (Greek letter Sigma). The vertical dashed lines show the value on the horizontal axis corresponding to one standard deviation less (to the left of the mean) and more (to the right of the mean) than μ.

The computer first takes a stab at the ratings of all races and horses. It tries to find the set of ratings for all races and horses that has the highest likelihood, given the data. In other words, the ratings of all horses and races that are most consistent with the results.

Adjusting for conditions

For every performance, the difference between the horse rating and the race quality is the distance the computer expects the horse to be beaten. But, when the actual distance doesn't match, the computer has to figure out why. We tell it to expect that the variability from the distance it expects is a function of the course, the distance and the country-specific going ('Good' in France is very different going from 'Good' in Japan, for instance.)

We won't go over the importance of this consideration yet again. You can read about why we think it leads to horses in the Far East, in particular, being underrated in this article.

The computer thus has an allowance for every set of circumstances a horse can encounter. Here is a portion of the chart that we can extract from the model:

					-- Rating points beaten for lengths beaten --					
Course	Surface	Going	Distance	N	1	2	4	5	10	20
Ascot GBR	Turf	Soft	6	197	3.6	6.6	11.6	13.7	21.1	29.0
Ascot GBR	Turf	Soft	7	34	3.4	6.2	10.9	12.9	19.9	27.3
Ascot GBR	Turf	Soft	8	133	2.9	5.3	9.3	11.0	16.9	23.3
Ascot GBR	Turf	Soft	10	99	3.1	5.7	10.0	11.7	18.1	24.9
Ascot GBR	Turf	Soft	12	154	2.5	4.6	8.0	9.4	14.6	20.1
Belmont Park USA	Dirt	Fast	6	223	3.1	5.8	10.2	12.0	18.6	25.6
Belmont Park USA	Dirt	Fast	7	274	3.2	5.9	10.3	12.1	18.7	25.7
Belmont Park USA	Dirt	Fast	8	502	2.7	5.0	8.8	10.4	16.0	22.1
Belmont Park USA	Dirt	Fast	8.5	210	2.8	5.2	9.1	10.7	16.6	22.8
Belmont Park USA	Dirt	Fast	9	138	2.7	4.9	8.7	10.2	15.7	21.6
Belmont Park USA	Dirt	Fast	10	128	2.4	4.5	7.8	9.2	14.2	19.5
Belmont Park USA	Dirt	Fast	12	168	2.1	3.9	6.8	8.0	12.3	16.9
Tokyo JPN	Turf	Firm	7	265	5.1	9.4	16.5	19.4	30.0	41.2
Tokyo JPN	Turf	Firm	8	956	5.1	9.4	16.5	19.4	30.0	41.3
Tokyo JPN	Turf	Firm	9	521	4.8	8.9	15.5	18.2	28.2	38.8
Tokyo JPN	Turf	Firm	10	319	4.2	7.8	13.6	16.0	24.8	34.1
Tokyo JPN	Turf	Firm	12	673	3.4	6.4	11.2	13.2	20.4	28.0

These are not the exact allowances that we use in all circumstances, merely the mean of those allowances. If a race that is run under conditions that normally stretch finishers out has already resulted in wide margins between the runners, we won't extend the distances still further. But the allowances in the chart above give a decent means of comparison between conditions. The larger margins we allow for races in Japan and Hong Kong are only partly responsible for horses from those countries receiving standout ratings under our system; a larger part is their outstanding record internationally.

So, now the computer has more than 150,000 equations to solve simultaneously. Do you remember doing two of these at school? Because there are far fewer unknowns than equations - 46,000 horses and 16,000 races - we say the system of equations is 'overdetermined', which means there are many possible values for every horse in every country. The relative strength of racing in each country is partly determined by the performance of challengers from other countries:

Leaning not just the most likely ratings for all horses and races but also their range of values and the probabilities of these values requires software and some serious processing. All the while, the computer is learning how to handicap by tuning the values it comes up with according to how well they predict future races.

There are a few programming tricks required to find an optimal solution, but handicapping is really just an optimisation exercise for which methodology is extensive within academic literature.

The only arbitrary element is to centre the scale of final ratings so that they lay over the ratings used by the World's Best Racehorse committee and commonly used around the world for handicap ratings, one in which the world's best horses are around the 130 mark.

Retooling the rankings

All this, of course, is a powerful and interesting exercise in itself. But the point of it is to retool *TRC Global Rankings*. The *Racing Post* has always provided the ratings that feed into the system by determining the strength of each race, and thus its reward to the winners, but it can no longer guarantee a pan-global service, especially in some countries such as South America, South Africa and New Zealand, where we take the quality of racing seriously.

So, from this week (beginning Monday, February 15, 2021) *TRC Computer Race Ratings* replace *Racing Post Ratings*. We are making the change for all races retrospectively too, so for this week only there will be artificial changes to the competitors in each of the Horses, Jockeys, Owners and Sires categories as a result of us viewing some races differently in terms of quality. Count this is as *TRC Global Rankings* V2.

Before that, we better show you what the computer has produced. Here are the Top 30 performances globally since 2011, plus the Top 15 achieved in selected countries.

By the way, *TRC Computer Race Ratings* has found that ratings for female horses (who may switch from running against their own sex at level weights to receiving an allowance against males) are more accurate without considering the allowance. This is also consistent with our philosophy that female horses who defeat males should be ranked above those males in classifications. So, if you wish to compare *TRC Global Rankings* with other rating systems that do include allowances for females, remember to add the weight back in to their numbers.

Ranking by career – not performance

Similar to our ethos over *TRC Global Rankings*, we believe strongly that horses should not be ranked or classified in descending order of their best efforts. Instead, their careers are better assessed in the round, that is, by some weighted average of their best efforts in descending order.

We also regress their ratings towards the mean of our data, which is 100; horses with more starts have their ratings regressed less. This way, a *TRC* career rating expresses the length, breadth and depth of a career.

Bearing in mind that our data started in 2011 and we don't have a full career for horses around ten years or so ago (most notably, Frankel's amazing 2-year-old season), here are the leading 50 Thoroughbreds ranked by career rating with their best three performances and the circumstances of their career high:

Rank	Horse	Country	Category	Wins	Runs	G1	G2	G3	Career	*TRC*	BestRace
1	Frankel GBR	GBR	8T	10	10	9	0	1	**132.5**	*137 136 136*	1st, G1 Queen Anne Stakes (Ascot GBR, 2012-06-19)
2	Enable GBR	GBR	12T	13	16	11	0	2	**129.3**	*133 132 131*	1st, G1 King George VI And Queen Elizabeth Stakes (Ascot GBR, 2017-07-29)
3	Winx AUS	AUS	8T	35	41	25	9	1	**129.2**	*132 132 131*	1st, G1 William Hill Cox Plate (Moonee Valley AUS, 2016-10-22)
4	Black Caviar AUS	AUS	6T	17	17	14	3	0	**128.7**	*130 130 129*	1st, G1 Lexus Newmarket Handicap (Flemington AUS, 2011-03-12)
5	American Pharoah USA	USA	10D	9	10	8	1	0	**127.6**	*132 129 128*	1st, G1 Breeders' Cup Classic (Keeneland USA, 2015-10-31)
6	Treve FRA	FRA	12T	7	11	6	1	0	**127.4**	*132 129 129*	1st, G1 Qatar Prix de l'Arc de Triomphe (Longchamp FRA, 2013-10-06)
7	Almond Eye JPN	JPN	8T	10	13	9	0	1	**127.1**	*130 128 128*	1st, G1 Tenno Sho (Tokyo JPN, 2019-10-27)
8	Gun Runner USA	USA	9D	10	17	6	2	2	**126.6**	*132 128 127*	1st, G1 Pegasus World Cup Invitational Stakes (Gulfstream Park USA, 2018-01-27)
9	Cracksman GBR	GBR	12T	6	9	4	2	0	**126.4**	*132 129 126*	1st, G1 Qipco Champion Stakes (Ascot GBR, 2017-10-21)
10	Golden Horn GBR	GBR	12T	5	7	4	1	0	**126.2**	*131 129 127*	1st, G1 Qatar Prix de l'Arc de Triomphe (Longchamp FRA, 2015-10-04)
11	Kingman GBR	GBR	8T	6	7	4	0	2	**125.8**	*130 129 126*	1st, G3 Aon Greenham Stakes (Newbury GBR, 2014-04-12)
12	Arrogate USA	USA	10D	4	7	4	0	0	**125.8**	*131 130 129*	1st, G1 Travers Stakes (Saratoga USA, 2016-08-27)
13	Wise Dan USA	USA	8T	18	23	11	6	1	**125.3**	*128 127 125*	1st, G3 Ben Ali Stakes (Keeneland USA, 2012-04-22)
14	Battaash IRE	GBR	5T	11	18	4	6	1	**125.3**	*129 128 127*	1st, G1 Coolmore Nunthorpe Stakes (York GBR, 2019-08-23)
15	Orfevre JPN	JPN	12T	11	18	6	5	0	**125.1**	*129 128 126*	1st, G1 Arima Kinen (Nakayama JPN, 2013-12-22)
16	California Chrome USA	USA	9D	10	18	7	3	0	**124.8**	*129 128 125*	2nd, G1 Breeders' Cup Classic (Santa Anita USA, 2016-11-06)
17	Excelebration IRE	GBR	8T	6	12	3	2	1	**124.7**	*129 127 125*	1st, G2 CGA Hungerford Stakes (Newbury GBR, 2011-08-13)
18	Magical IRE	IRE	10T	11	26	7	3	1	**124.7**	*131 126 124*	2nd, G1 Longines Breeders' Cup Turf (Churchill Downs USA, 2018-11-03)
19	So You Think NZL	IRE	10T	6	11	5	0	1	**124.6**	*128 127 127*	1st, G1 Tattersalls Gold Cup (Curragh IRE, 2011-05-22)
20	Ghaiyyath IRE	GBR	10T	8	11	4	1	3	**124.6**	*128 127 127*	1st, G1 Hurworth Bloodstock Coronation Cup Stakes (Newmarket GBR, 2020-06-05)
21	Stradivarius IRE	GBR	16T	14	21	7	7	0	**124.3**	*128 125 123*	1st, G1 Gold Cup (Ascot GBR, 2020-06-18)
22	Minding IRE	IRE	8T	8	11	7	1	0	**124.2**	*128 125 124*	1st, G1 Queen Elizabeth II Stakes (Ascot GBR, 2016-10-15)
23	Accelerate USA	USA	9D	8	17	5	3	0	**124.1**	*128 127 125*	1st, G1 $1 Million TVG Pacific Classic (Del Mar USA, 2018-08-19)
24	Maximum Security USA	USA	9D	8	10	6	1	1	**124.1**	*126 125 125*	1st, G1 Saudi Cup (Riyadh KSA, 2020-02-29)
25	Farhh GBR	GBR	10T	2	7	2	0	0	**123.9**	*128 126 125*	1st, G1 JLT Lockinge Stakes (Newbury GBR, 2013-05-18)
26	Shared Belief USA	USA	9D	8	9	5	2	1	**123.9**	*127 126 123*	1st, G2 San Antonio Invitational Stakes (Santa Anita USA, 2015-02-07)
27	Gran Alegria JPN	JPN	8T	6	9	4	1	1	**123.7**	*127 126 125*	1st, G1 Yasuda Kinen (Tokyo JPN, 2020-06-07)
28	Authentic USA	USA	9D	5	7	3	1	1	**123.7**	*128 124 123*	1st, G1 Longines Breeders' Cup Classic (Keeneland USA, 2020-11-07)
29	Ribchester IRE	GBR	8T	6	15	4	1	1	**123.6**	*127 127 126*	2nd, G1 Queen Elizabeth II Stakes (Ascot GBR, 2016-10-15)
30	Songbird USA	USA	8.5D	12	14	9	2	1	**123.6**	*127 123 123*	1st, G1 Cotillion Stakes (Parx USA, 2016-09-24)
31	City Of Light USA	USA	7D	5	7	4	1	0	**123.6**	*128 126 125*	1st, G1 Pegasus World Cup Invitational Stakes (Gulfstream Park USA, 2019-01-26)
32	Atlantic Jewel AUS	AUS	8T	7	8	4	3	0	**123.4**	*125 125 124*	1st, G1 Schweppes Thousand Guineas (Caulfield AUS, 2011-10-12)
33	Order Of St George IRE	IRE	14T	8	17	3	1	4	**123.4**	*126 126 124*	1st, G1 Palmerstown House Estate Irish St. Leger (Curragh IRE, 2015-09-13)
34	Cirrus Des Aigles FRA	FRA	10T	15	37	7	4	4	**123.3**	*126 126 124*	1st, G1 Prix Ganay (Longchamp FRA, 2012-04-29)
35	Moonlight Cloud GBR	FRA	7T	10	17	6	0	4	**123.3**	*127 125 125*	1st, G1 Prix Maurice De Gheest - Goldikova (Deauville FRA, 2011-08-07)
36	Australia GBR	IRE	8T	4	6	3	0	1	**123.3**	*128 123 122*	1st, G1 Investec Derby (Epsom GBR, 2014-06-07)
37	Crystal Ocean GBR	GBR	12T	7	15	1	1	5	**123.3**	*129 127 125*	2nd, G1 King George VI And Queen Elizabeth Qipco Stakes (Ascot GBR, 2019-07-27)
38	Blue Point IRE	GBR	6T	9	18	4	2	3	**123.1**	*127 126 123*	1st, G1 King's Stand Stakes (Ascot GBR, 2018-06-19)
39	Al Kazeem GBR	GBR	10T	8	18	4	2	2	**123.0**	*126 125 124*	1st, G1 Coral-Eclipse (Sandown GBR, 2013-07-06)
40	Gold Ship JPN	JPN	12T	11	26	6	4	1	**122.8**	*125 123 123*	1st, G1 Takarazuka Kinen (Hanshin JPN, 2013-06-23)
41	Flintshire GBR	FRA	12T	7	21	5	1	1	**122.8**	*126 125 124*	2nd, G1 Qatar Prix de l'Arc de Triomphe (Longchamp FRA, 2015-10-04)
42	Postponed IRE	GBR	12T	7	15	4	3	0	**122.8**	*126 125 124*	1st, G1 Juddmonte International Stakes (York GBR, 2016-08-17)
43	Highland Reel IRE	GBR	12T	9	25	7	1	1	**122.8**	*126 124 124*	1st, G1 Longines Breeders' Cup Turf (Santa Anita USA, 2016-11-05)
44	Found IRE	IRE	10T	5	19	3	0	2	**122.8**	*128 124 123*	1st, G1 Qatar Prix de l'Arc de Triomphe (Chantilly FRA, 2016-10-02)
45	West Coast USA	USA	9D	3	9	2	0	1	**122.8**	*127 127 125*	1st, G1 Travers Stakes presented by NYRA Bets (Saratoga USA, 2017-08-26)
46	Nathaniel IRE	GBR	10T	3	8	2	1	0	**122.7**	*128 124 122*	1st, G1 Coral-Eclipse (Sandown GBR, 2012-07-07)
47	Beauty Generation NZL	HKO	8T	15	31	8	4	3	**122.7**	*124 124 122*	1st, G2 Bochk Wealth Management Jockey Club Mile (Sha Tin HKO, 2018-11-18)
48	Sharp Azteca USA	USA	8D	5	11	1	2	2	**122.7**	*127 125 124*	1st, G1 Cigar Mile Handicap presented by NYRA Bets (Aqueduct USA, 2017-12-02)
49	Midnight Bisou USA	USA	8.5D	13	20	5	5	3	**122.7**	*125 124 123*	2nd, G1 Saudi Cup (Riyadh KSA, 2020-02-29)
50	Chrono Genesis JPN	JPN	8T	5	11	3	1	1	**122.7**	*128 125 120*	1st, G1 Takarazuka Kinen (Hanshin JPN, 2020-06-28)

Top 50 racing careers since 2011 according to TRC Computer Race Ratings

Next time

There is a ton of insight that we have derived from this exercise, and having developed software to produce ratings like this we are excited to share it with you. *TRC Computer Race Ratings* will soon be available as a click-through on all rankings, but they are already powering the rankings you see on the site today.