
BUSINESS SYSTEMS

VISIT TO MARSAILLE

Dr Andrew Flitman

The travel grant was issued to enable me to present a refereed paper at Neurap'96. The paper is to be published in full proceedings of the conference. Because of the French strikes at the time, the conference was postponed. I was informed of this through a telephone conversation with Diane West after I arrived in London.

I therefore made alternative arrangements, and worked with Dr Paul Rafenes at London Business School. Dr Rafenes runs the Decision Sciences Research Group (generally regarded as the world's leading group in my area of finance modelling). Our joint work resulted in our setting up closer ties between our groups resulting in a 10 week visit by one of my PhD students in April/May 1996. We have also been encouraged by Dr Rafenes to submit a paper to a special edition of the Journal of Forecasting (possibly the world's most prestigious journal in that area) of which Dr Rafenes is editor. We discussed funding of the PhD students visit to London, and his program of activity. We therefore collaborated in an application to the British Council for funding.

In the time allocated to the conference trip, I also visited the Business School at Warwick University and further explored joint research interests, as well as learning from their methods of obtaining external funding (Warwick University is the most successful of UK universities in obtaining outside funding, to the extent that this funding now exceeds 50% of their total revenue). It is my intention (when time permits) to emulate their approach and develop a research forum on financial modelling research.

Other academic activities undertaken during my trip included access to the Bodlean Library at Oxford, and visiting Vienna Technical University (extending my German to it's utmost!).

COMPUTER SCIENCE

AN INFORMATION-THEORETIC AND A GAUSSIAN FOOTY-TIPPING COMPETITION

David L. Dowe, Kevin L. Lentin, Jon J. Oliver and A. John Hurst

"The Master said, Yu, shall I teach you what knowledge is? When you know a thing, to recognise that you know it, and when you do not know a thing, to recognise that you do not know it. That is knowledge."

- Analects of Confucius (transl. by Arthur Waley), Book II, No. 17.

"And, yea, having single-handedly out-scored an unnamed team in the opening game of the season, did Gary ascend and come to sit on the right shoulder of the full back, from whence verily did his nephew later attend a party at the dwellings of one of the FCIT general staff."

- source unknown, due apologies made.

It seems to be a shortcoming in many statistical and "machine learning" methods that, while they are prepared to predict, they are not always willing to associate a probability or degree of certainty with the prediction. Before taking a risk in life, such as going to the shops in winter without an umbrella or making a risky financial investment, we are not only interested in whether it is more likely to rain, or not to rain, and in whether the prices of our investment are more likely to go up than down, but we would also like some idea of the probability of the outcome. We really would like to know whether it's 51% likely that it won't rain (whereupon we might take our umbrella) or whether it's 99.9% likely not to rain (whereupon we might not bother taking our umbrella).

The FCIT Faculty footy-tipping competition rewards people for correctly tipping the winning team in AFL football matches, and the mainstream football tipping competition running in the Computer Science department has been of a very similar nature for a number of years.

Send items for next Newsletter to Sheila Pullin at:

*spullin@fcit-m1.fcit.monash.edu.au
by 30 August*

"Right"/"wrong" is a very obvious and sensible criterion which is used in multiple-choice exams and, indeed, most tipping competitions.

However, there is a time and place in life for knowing how sure we are of ourselves. Confucius seems (above, top quotation) to think that this is important, as undoubtedly do the TAB, a lot of bookmakers and many casinos. The ability of bookmakers to make a livelihood is dependent to some degree upon how closely they can pick the true probability of the outcome of a horse race or of a footy team winning, or whatever event. Bypassing the mathematics, it turns out that the correct reward to give someone whose asserted probability was p for an event that ended up actually happening is $\log(p)$. We take logarithms to the base 2, so that tippers get a score in bits. Tippers are given 1 bit per game, since a tip of 0.50 on each team is guaranteed to cost $\log_2(0.50) = -1$ either way, and so an ultra-conservative tipper can always guarantee a score of at least $1 + (-1) = 0.0$ per game. We say above that the logarithm function is the "correct" reward because it turns out in this case that the optimal expected long-term strategy is to tip the "true" probability of the result of the match. An over-confident tip can see a tipper gain as much as $1 + \log_2(1.0) = 1 + 0 = 1$ bit for a game, but if the tipper were wrong, a score of $1 + \log_2(1-p)$ will be devastating when p is close to 1.0 and $1-p$ is close to 0.0. The boldest tip anyone had taken in the early part of the 1995 season was 0.95, with at least one such tip having been unsuccessful. The reader might care to see (bottom lines of table) that in Round 4 of the 1996 season, some successful 0.99 tips were ventured, as was (ouch!) an unsuccessful 0.95 tip. (Right towards the end of the 1995 season, there were some desperate tips made, some of which went desparately wrong.)

The Computer Science department's first probabilistic tips were for the Round 3 matches of the 1995 AFL season, which began with a match on Saturday 15 April, 1995. Examples of this logarithmic reward function discussed above being used in research work date back to Professor Phil Dawid of University College, London, and more recently to Wallace and Patrick (1993, p20) and Dowe and Krusel (1993, p3).

We didn't begin until Round 3 (1995) because it wasn't until a few days prior to Round 3 that Jon Oliver suggested to David Dowe in passing in the

corridor that this might be a good idea. David e-mailed the department's football-tipping mailing list with the scoring system (discussed above), and Kevin Lentin did all the hard work of writing the code and setting things up. In the middle of the season, after Round 11 (1995), we opted to have two parallel competitions - the continuation of the original competition and a new competition which re-started after Round 11. Graham Farr has made his observation that the sense of gambling is part of the appeal in this (since an extreme tip of 0.95 or so can seem a tempting way of getting ahead fast, but it runs the risk of hurting when one is wrong). I (dld) dare say that this competition is largely responsible for re-juvenating my interest in a game that I thought I had stopped following decades ago. Graham also contributed to discussions on the "draw" rule (which we had to use in 1996 Round 4).

I had a large probability that the Guinness Book of Records would not express great interest in this, but I'd never written to the Guinness Book of Records (or the FCIT Newsletter) before, and this seemed like the time. My confident probability that the Guinness Book of Records would not show sufficient interest was vindicated. I hope the FCIT Newsletter is more sympathetic - if you're reading this, it is.

Believe it or not, there is serious academic research in this. If you don't agree, just ask a bookmaker. One might also look at the theory of the "Weighted" and "Average" tippers, which are definitely of academic interest. Yes, information theory and MML (Minimum Message Length) (Wallace and Boulton, 1968; etc.) is relevant here.

Also, as Confucius points out and as at least Chris Wallace and yours truly (dld) seem to agree, in a competition like this, in the long run, the way to make the most of whatever knowledge you have is to tip the probability that you genuinely believe. For example, one pretending to be more sure of themselves than they are will surely pay the price. On this point, you could also ask Kevin Korb his thoughts about "calibrating" the various tippers.

Some tipping strategies include Kevin Korb's modification of the chess-based Elo rating system. At least one other person is now experimenting with this idea. And John Hurst is one of at least three people whose tips are largely based on the content of the Friday Football Guide. Yours truly believes that the former idea in particular, and

possibly also a variant of the latter, should be very relevant in the ranking of cricket teams.

In case the ideas above behind the info-theory tipping competition might almost seem clear, as of Round 1, 1996 (which commenced on Fri 29 March with the result Ablett 9-6-60 def. Melbourne 8-8-56, not to mention Geelong's 27-21-183 - but I digress), we've got a new, even more complicated, competition. This time, the idea is to pick the margin (μ) of the game, and also an estimate of the expected error (σ) in your tip. This tip defines a "Normal" or "Gaussian" distribution over the margin. By Round 4 (1996), the designer (dld), the programmer (ajh) and most of the participants (or those that haven't dropped out) almost seem to have the hang of this. Suffice it to say that both competitions are about minimising a Kullback-Leibler distance and that they are, quite seriously, both relevant to problems of scientific prediction and combining opinions of different "experts".

Two of many possible ways of combining our "expert" tips are (i) to average them and (ii) to somehow weight them, where the better tippers are given larger weights than the less successful tippers. The "Average" tipper and a "Weighted" tipper were both entered in the 1995 information-theoretic competition. The weightings used in the "Weighted" tipper were based on Ray Solomonoff's (1964) ideas. The "Average" tipper performed very well in 1995, leading 11 human tippers and trailing only 3 human tippers after Round 21. The "Weighted" tipper was a clear winner in 1995, but it is not yet clear whether or not this was because it was inadvertently given a slight but nonetheless unfair advantage. It is desired that we have the "Average" and (correct) "Weighted" tippers for the 1996 information-theoretic competition, as the writer is relatively sure we probably soon will. The mathematics has also been done to put "Average" and "Weighted" tippers in the Gaussian competition.

If any of this interests you, the "old", 1995, WWW site

<http://www.cs.monash.edu.au/~kevinl/footy.html>
and the new (1996) WWW site is at
<http://www.cs.monash.edu.au/~ajh/footy/results/>.

We'll almost conclude with an MML-based statistical analysis of the results of the info-theory footy-tipping competition after its first round, Round 3 (1995).

All eight competitors had scores greater than zero (suggesting from early data that, as we already knew, football is indeed not random) and the four competitors who had been to MML meetings were all doing worse than the four who hadn't.

Let's conclude with some acknowledgments, a list of references, some results from after Round 4 of the 1996 Info-theory competition and some results after Round 4 of the inaugural, "centenary year", 1996 Gaussian tipping competition. Tippers were given the option of being anonymous in this publication. Thanks to ajh for reducing the results of the emergingly popular Info Theory competition (with entrants from two states, three faculties and three companies) to fit onto one page. (Please see the WWW pages if anything is a little illegible.) Thanks also to gfarr for feedback, and thanks to the tippers for making the competition (and lending their names here).

The info-theory and Gaussian results after 1996 Round 4 should both be and probably both are displayed near this article.

References:

D. L. Dowe and N. Krusel, "A decision tree model of bushfire activity", Technical Report #93/190, Department of Computer Science, Monash University, Clayton 3168, Australia.

D. L. Dowe and K. L. Lentin, "Information-theoretic footy-tipping competition - Monash", Computer Science Association Newsletter (Australia) pp55-57, 1995.

A. E. Elo, "The USCF Rating System", U.S. Chess Federation Publication, 1966. Described also in: Harkness, K., 'Official chess handbook', New York: David McKay, 1967, pp183-205.

K. B. Korb, "The Collapse of Collective Defeat", 1992.

R. J. Solomonoff, "A formal theory of inductive inference I, II", Information and Control, 7, 1964, pp1-22 and pp224-254.

C. S. Wallace and D. M. Boulton, "An Information Measure for Classification", Computer Journal, Vol.11, No.2, 1968, pp185-194.

C. S. Wallace and J.D. Patrick, "Coding Decision Trees", Machine Learning, 11, pp7-22, 1993.

Footy Sweep 1996

Info Theory Results after Round 4

Department of Computer Science

Name	Bri(30) Nor	Car(47) Mel	Col(37) Fre	St.K(85) Fit	Foo Haw	Ade(96) Ess	Wes Gee(18)	Ric Syd(1)	This Week	Season Total
james-JamiesScuglia	Nor(0.700) -0.737	Car(0.800) 0.678	Fre(0.650) -0.515	St.K(0.600) 0.263	Foo(0.600) -0.029	Ess(0.650) -0.515	Wes(0.650) -0.515	Ric(0.000) 1.000	-0.369	5.557
ajh-JohnHurst	Nor(0.579) -0.248	Car(0.816) 0.707	Col(0.621) 0.313	St.K(0.734) 0.554	Foo(0.645) -0.063	Ade(0.538) 0.106	Gee(0.538) 0.106	Ric(0.674) -0.617	0.856	5.205
dld-DavidDowe	Nor(0.539) -0.117	Car(0.817) 0.708	Col(0.623) 0.317	St.K(0.735) 0.556	Foo(0.646) -0.064	Ade(0.539) 0.108	Gee(0.539) 0.108	Ric(0.675) -0.621	0.995	5.014
kevinl-Kevin	Nor(0.700) -0.737	Car(0.700) 0.485	Col(0.650) 0.379	St.K(0.650) 0.379	Foo(0.680) -0.100	Ade(0.630) 0.333	Gee(0.670) 0.422	Ric(0.750) -1.000	0.161	4.988
jono-jono	Nor(0.550) -0.152	Car(0.820) 0.714	Col(0.630) 0.333	St.K(0.750) 0.585	Foo(0.550) -0.007	Ade(0.500) 0.000	Gee(0.680) 0.000	Ric(0.680) -0.644	0.829	3.001
gfarr-GrahamFarr	Bri(0.510) 0.029	Car(0.780) 0.642	Col(0.530) 0.084	St.K(0.680) 0.444	Foo(0.650) -0.068	Ade(0.530) 0.084	Wes(0.550) -0.152	Ric(0.700) -0.737	0.325	2.836
chatz-chatz	Nor(0.550) -0.152	Car(0.950) 0.926	Col(0.550) 0.138	St.K(0.600) 0.263	Foo(0.500) 0.000	Ess(0.600) -0.322	Gee(0.600) 0.263	Ric(0.600) -0.322	0.794	2.504
@omen.melbourne.sgi.com	Nor(0.731) -0.894	Car(0.942) 0.914	Col(0.514) 0.040	St.K(0.746) 0.577	Foo(0.652) -0.070	Ess(0.600) -0.322	Gee(0.564) 0.174	Ric(0.776) -1.158	-0.740	2.456
Ian.Coulter-DropKick	Bri(0.700) 0.485	Car(0.900) 0.848	Col(0.700) 0.485	St.K(0.900) 0.848	Foo(0.530) -0.003	Ess(0.520) -0.059	Gee(0.520) 0.057	Ric(0.800) -1.322	1.340	1.340
@arts.monash.edu.au	Nor(0.600) -0.322	Car(0.600) 0.263	Col(0.600) 0.263	St.K(0.600) 0.263	Foo(0.600) -0.029	Ess(0.600) -0.322	Gee(0.600) 0.263	Ric(0.600) -0.322	0.057	0.350
karen-kazza	Nor(0.600) -0.322	Mel(0.550) -0.152	Fre(0.550) -0.152	St.K(0.600) 0.263	Foo(0.600) -0.029	Ess(0.550) -0.152	Gee(0.550) 0.138	Ric(0.600) -0.322	-0.729	-0.422
annn-annn	Bri(0.600) 0.263	Car(0.900) 0.848	Col(0.700) 0.485	St.K(0.700) 0.485	Haw(0.600) -0.029	Ess(0.800) -1.322	Wes(0.700) -0.737	Ric(0.700) -0.737	-0.743	-1.607
scottt-fnord	0.000	0.000	0.000	0.000	0.000	Ade(0.700) 0.485	0.000	0.000	0.485	-2.004
@melb.convergent.com.au	Nor(0.600) -0.322	Car(0.990) 0.986	Col(0.700) 0.485	St.K(0.600) 0.263	Haw(0.550) -0.007	Ess(0.600) -0.322	Wes(0.650) -0.515	Ric(0.700) -0.737	-0.169	-2.051
anitaj-Anita	Nor(0.891) -2.198	Mel(0.671) -0.604	Fre(0.682) -0.653	Fit(0.601) -0.326	Haw(0.531) -0.003	Ess(0.720) -0.837	Wes(0.719) -0.831	Ric(0.619) -0.392	-5.843	-4.167
jefsharp-JeffreySharp	Nor(0.600) -0.322	Car(0.900) 0.848	Col(0.550) 0.138	St.K(0.900) 0.848	Haw(0.700) -0.126	Ess(0.600) -0.322	Wes(0.700) -0.737	Ric(0.900) -2.322	-1.995	-5.235
@adelaide.DIALix.oz.au	Nor(0.850) -1.737	Car(0.990) 0.986	Col(0.900) 0.848	St.K(0.900) 0.848	Haw(0.900) -0.737	Ess(0.650) -0.515	Gee(0.850) 0.766	Ric(0.950) -3.322	-2.863	-6.816
Debbie.Hedger-Debbie										
@lib.monash.edu.au										
kelly-AndrewKelly										
jamesf-James										

Hmmm! Should we worry about jamies's root privilege? Does he know something I don't (apart from the root password)? Be that as it may, inspired/deranged tipping on his part for the last game of the round saw him keep his lead over ajh. But it would have been season's curtains for him, had that umpire not given the free kick against Richardson!!!

These results have had the correct formula for drawn games applied:

$$\text{score} = 1 + \frac{1}{2}(\log_2(p) + \log_2(1-p))$$

Footy Sweep 1996

Department of Computer Science

Gaussian Tipping Results after Round 4

Name	Bri(30) Nor	Car(47) Mel	Col(37) Fre	St.K(85) Fit	Foo Haw	Ade(96) Ess	Wes Gee(18)	Ric Syd(1)	This Week	Season Total
kevinl-Kevin	Nor(20) 45.0/1.29	Car(37) 30.0/2.69	Fre(10) 45.0/1.40	St.K(20) 40.0/0.45	Foo(10) 40.0/2.31	Ade(23) 60.0/0.70	Gee(15) 50.0/2.03	Ric(25) 35.0/2.15	13.004	44.319
dld-DavidDowe	Nor(18)	Car(30)	Col(15)	St.K(34)	Foo(17)	Ade(7)	Gee(6)	Ric(30)	12.333	40.381
ajh-JohnHurst	Nor(24)	Car(64)	Col(12)	St.K(18)	Foo(43)	Ade(30)	Gee(30)	Ric(20)	11.411	33.641
kelly-AndrewKelly	Nor(10)	Car(60)	Col(1)	St.K(60)	Haw(10)	Ess(10)	Wes(10)	Ric(60)	9.962	30.138
karen-kazza	Nor(45)	Car(45)	Col(45)	St.K(45)	Foo(45)	Ess(45)	Gee(45)	Syd(20)	6.058	22.849
Ian.Coulter@arts.monash.edu.au-DropKick	100.0/0.62	100.0/1.03	100.0/1.03	100.0/0.91	100.0/0.88	100.0/-0.40	100.0/0.98	100.0/1.00	11.772	22.440
jamies-JamieScuglia	Bri(30)	Car(70)	Col(19)	St.K(62)	Foo(7)	Ess(5)	Gee(6)	Ric(60)	8.053	21.658
gfarr-Graham FarrV2	50.0/2.03	50.0/1.88	50.0/1.94	50.0/1.88	50.0/2.02	50.0/-0.91	50.0/1.99	50.0/0.96	11.380	21.644
jono-jono	Nor(15)	Car(25)	Fre(10)	St.K(20)	Foo(15)	Ess(20)	Wes(15)	Ric(15)	1.569	18.980
gfarr-GrahamFarr	30.0/1.14	40.0/2.13	30.0/1.00	40.0/0.45	30.0/2.59	40.0/-3.71	30.0/1.89	30.0/2.56	11.380	3.245
(anonymous)	Bri(1)	Car(25)	Col(5)	St.K(20)	Foo(15)	Ade(3)	Wes(5)	Ric(20)	8.156	-63.317
	60.0/1.60	55.0/1.78	60.0/1.56	55.0/0.89	50.0/1.97	60.0/0.03	50.0/1.88	60.0/1.68		
	Nor(10)	0.0/0.00	0.0/0.00	0.0/0.00	0.0/0.00	0.0/0.00	0.0/0.00	0.0/0.00		
	Bri(1)	Car(25)	Col(5)	St.K(20)	Foo(15)	Ade(3)	Wes(5)	Ric(20)		
	60.0/1.60	55.0/1.78	60.0/1.56	55.0/0.89	50.0/1.97	60.0/0.03	50.0/1.88	60.0/1.68		
	Nor(30)	Car(60)	Col(30)	St.K(60)	Haw(30)	Ess(20)	Gee(40)	Ric(50)		
	30.0/-0.12	50.0/1.98	40.0/2.33	40.0/2.07	30.0/2.05	-40.0/-3.71	30.0/2.38	40.0/1.18		