Quickly Detecting Skill Trace in Games

Cameron Browne
Department of Advanced Computing Sciences
Maastricht University
Maastricht, The Netherlands
cameron.browne@maastrichtuniversity.nl

***Abstract*—This short paper introduces the notion of the *skill trace* of a game, which indicates its potential for tactical and/or strategic interest. A simple method is presented for quickly detecting the skill trace of a given game.**

***Index Terms*—General game playing (GGP), automated game evaluation, game metrics, skill trace, Monte Carlo tree search (MCTS).**

I. Introduction

When discussing the character of a game, it is useful to distinguish the degree of skill versus chance that the game involves. Uncertainty plays a vital role in games [4] but a degree of skill is typically required for the game to have lasting merit. Throughout human history, games have changed and evolved as they are passed on orally from individual to individual and generation to generation [11], and it is those games that achieve a good balance of skill and chance that tend to survive [6].

While there have been some studies into measuring the balance of skill versus chance in games, e.g. [1] and [5], there is no standard method for achieving this. This paper describes a simple method for quickly detecting indicators of the potential for skill in a given game.

A. Strategic and Tactical Potential

A game's potential for skill essentially comes down to its potential for *strategic* and/or *tactical* play, where *strategic play* refers to the high-level goals that the player aims to achieve while *tactical play* refers to the low-level actions that achieve those goals. Strategic depth, especially, is an important quality for a game to possess if it is to have lasting merit [14].

Lantz *et al.* describe the useful notion of the *strategy ladder* [9] as shown in Fig. 1. The three plots each show a game, where the dots on each plot indicate points at which players learn increasingly effective strategies through their experience with the game. The leftmost plot (white dots) shows a game with easily obtained strategies that quickly yield perfect play; such a game would be easy to learn and master but of little interest thereafter. The rightmost plot (white dots) shows a game with difficult strategies that would take a long time to learn; such a game would be difficult and frustrating for players and would most likely die out. The central plot (black dots) shows a game with many strategies that can be learnt in a linear fashion; such a game would be easy to learn and encourage further play as deeper strategies are discovered. The early strategies of such a game (circled) are of especial interest, as games with such strategies (and tactics) that beginners could obtain would be easier to teach and learn, and would be more likely to be played; these are the games that would flourish and survive.

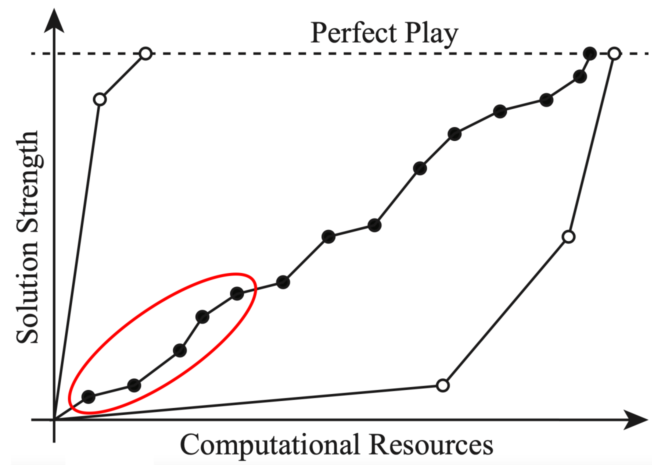

Fig. 1. Strategy ladder [9] showing early strategies in a well-formed game.

B. Motivation

The Digital Ludeme Project [2] aims to model the evolution and spread of games throughout human history. This involves evaluating candidate reconstructions of the rules for ancient games based on the available (partial) evidence [3], and a method for quickly detecting the skill trace of a given rule set could go a long way to correctly evaluating potentially vast numbers of candidate reconstructions. As we are concerned with the historical transmission of games [11], we are primarily interested in detecting rule sets that allow simple strategies and tactics at the lower end that beginners might learn and pass on to others, such as those circled in Fig. 1. This makes the task easier as strong "oracle" players with more complete understanding of the games are not required.

The strategic depth of a game is often gauged by the number of distinct skill levels that the game supports [14]. However, this approach typically assumes a large dataset of existing games between known players – which is not the case for newly reconstructed rule sets – and can be susceptible to the ranking scheme used [12]. We now describe a simple alternative measurement.

II. Skill Trace

We define the *skill trace* of a game as an indication of the degree to which the game rewards strategic and/or tactical play as opposed to random play. Note that we do not need to know exactly *what* those strategies and/or tactics are, or how many there are, just that the potential for them exists.

The approach involves running a sequence of M matches with T trials per match between standard UCT [8] agents at low iteration counts, one weak and one strong, successively doubling the iteration counts with each match and observing the resulting win rates. The assumption is that any difference in playing strength between the two AIs in each match-up is largely due to tactical/strategic play that the the strong AI is able to stumble upon with its superior search budget that the the weak AI with its shallower search fails to achieve.

The iteration counts are based on the *branching factor* (i.e. number of actions) at the base state for each move. That is, if the current state has BF moves, and UCT_n indicates a standard UCT agent with a search budget of n iterations, then matches are run between:

- UCT_{BF} vs UCT_{2BF}
- UCT_{2BF} vs UCT_{4BF}
- UCT_{4BF} vs UCT_{8BF}
- UCT_{8BF} vs UCT_{16BF}
- ...

until a given time limit is reached.

For each match, the average win rate for the strong AI (with the higher iteration count) is recorded, where 1=win, 0=draw and -1=loss. A regression line is then calculated through the strong AI win rates (shown in Fig. 2) from the second match onwards. The first match result is discarded due to the anomalous behaviour of UCT at such low iteration counts; the first BF iterations constitute random move choices as each root node child is tried in turn, then the next BF move choices tend to favour high-reward move choices until the UCB1 exploration factor kicks in from iteration $2BF$ onwards. The result is that UCT_{2BF} tends to unduly outperform UCT_{BF} in most cases, as can be seen in the obvious spike that defies the trend at $BF = 1$ in Fig. 2.

The y intercept of the regression line at $T + 1$ is then calculated (white dot). This is a projection of the next expected win rate, which slightly rewards upward-trending regressions and slightly penalises downward-trending regressions:

$$y = f(M + 1) \ in \ [0..1] \tag{1}$$

An estimate of the area A between the 0 line and positive part of the plot is then obtained by adding the win rate scores (clamped to the range [0..1] and squared). The final skill trace value ST is given by:

$$ST = y + (1 - y)A \tag{2}$$

This calculation provides a linear interpolation between the next expected win rate (i.e. y intercept) and the area covered

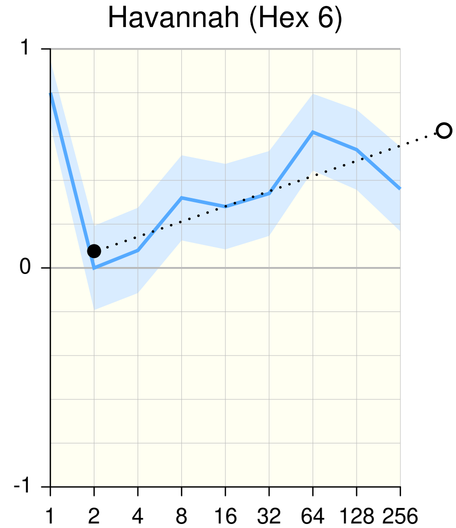

Fig. 2. Skill trace estimate of Havannah (size 6) over 20s of self-play. The x axis shows i (Weak AI BF multiplier) while the y axis shows the Strong AI's mean win rate per match. The error band shows the 95% confidence interval.

by the plot in cases in which the y intercept tends to zero. The approach is summarised in Algorithm 1.

Algorithm 1 Skill Trace (ST) Estimation

SkillTrace(**Game** G)
  return Calculate(Generate(G))

Generate(**Game** G)
  double $scores[] \leftarrow 0$
  int $m \leftarrow 1$
  while $time < Budget$
    for T $trials$
      $scores[m]$ += SelfPlayUCT(G, 2^{m-1}, 2^m)
    $m \leftarrow 2m$
  return $scores[]$

Calculate(**double** $scores[]$)
  $f \leftarrow$RegressionLine($scores[]$)
  $y \leftarrow f(M+1)$ // in range [0..1]
  $A \leftarrow \sum_{m=1}^{M} max(0, scores[t])^2$
  return $y + (1 - y)A$

The function SelfPlayUCT(G, 2^{m-1}, 2^m) runs a single self-play trial for game G between UCT agents with search budgets $2^{m-1}BF$ and $2^m BF$ respectively, where m is the current match count. Note that play order should alternate with each trial to alleviate any inherent first or second move advantage. For multiplayer games with $P > 2$ players, one strong agent is played against $P - 1$ weak agents, and the weak agent results collapsed to a single average value.

A. Inspiration

The inspiration for the regression line approach goes back to Hausdorff's 1919 topological dimension metric [7]. This

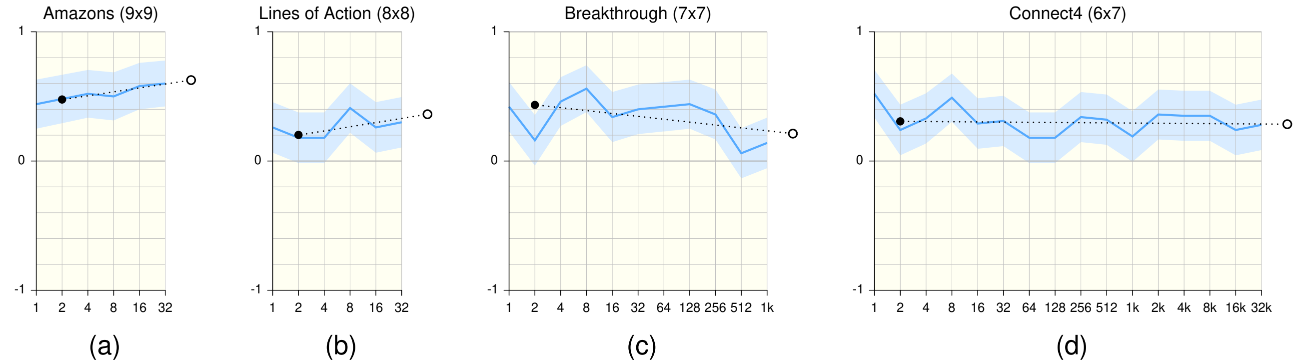

Fig. 3. Skill trace estimates of some well-known games: (a) Amazons, (b) Lines of Action, (c) Breakthrough and (d) Connect4.

was the precursor to Richardson's linear regressions through log/log plots of coastline lengths measured at successively doubled intervals, which inspired Mandelbrot's fractal dimension metric [10, p.33]. The ST metric is similarly estimated from a series of successively doubled samples, for which taking the linear regression serves to smooth local fluctuations in win rates to reveal an overall trend (ideally upwards).

III. Examples

The method described above was implemented for the Ai Ai general game player [13] and applied to a number of games. The number of trials per match was set to $T = 100$ and the UCT exploration constant set to a default of $C = \sqrt{2}$. All tests were run on a single thread of an $i7$ processor with a time setting of 20 seconds per game (unless otherwise stated).

Fig. 2 shows the result for Havannah on a size-6 board, indicating a strong skill trace with an upward trend, as would be expected from a game of such strategic depth. Fig. 3 shows the results for some other well-known and well-studied games. Amazons (a) and Lines of Action (b) show strong skill traces with upward trends while Breakthrough (c) and Connect4 (d) show lower relatively flat skill traces, all as would be expected.

Fig. 4 shows the results for some simple games played on small boards. Tic Tac Toe (a) shows minimal skill trace that quickly converges to zero, while Three Mens Morris (b) and Mu Torere (c) also converge quickly to zero but show stronger traces of skill than Tic Tac Toe early on, as might be expected. This is the reason for linearly interpolating between the y intercept and the area A below the plot, in order to detect early skill traces in games whose win rates converge to zero.

Conversely, Fig. 5 (left) shows the results for a pure chance game called Last Chance Saloon,[1] in which players add pieces to a board until it is full and then roll a dice to decide the winner. This game produces a score of $ST = 0$ as expected. The game Can't Stop (Fig. 5, right) involves a strong chance element but also involves some skill, which is detected.

Fig. 6 shows skill trace measurements for two more complex games of notable strategic depth. The plots shown took much longer than 20 seconds to compute – 1 hour and 4 hours, respectively – but show a strong upward trend as expected. These upwards trends can be found in less time if needed.

Based on these results, the ST score appears to be useful in distinguishing games that exhibit:

- No skill (flat line at $y = 0$).
- Some early skill (non-zero area sloping down to $y = 0$).
- Some constant skill (flat line above $y = 0$).
- Significant skill (upward slope).

Note that any skill trace detected by UCT agents running such low iteration counts is more likely to be due to the stronger agent stumbling upon effective low-level combinatorial (i.e. tactical) plays rather than high-level (i.e. strategic) plays which would require much stronger levels of play.

IV. Validation

It is difficult to validate the ST calculation objectively as the balance of skill and chance in any given game has to date been more of a subjective intuition with no convenient yardstick to measure against. In lieu of such a yardstick, ST estimates were correlated against BoardGameGeek (BGG)[2] user ratings for 31 games implemented in Ai Ai, as shown in Fig. 7. BGG user ratings indicate the average user's satisfaction with each game, and are remapped from [0..10] to [-1..1] in the figure.

The Pearson correlation coefficient for the two sets of measurements (i.e. BGG score and ST score for each game) is $r = 0.6213$, which is significant at the 99.9% level. This suggests a general correlation between the ST score of a game and the average player's satisfaction with that game.

V. Conclusion

Skill trace is a new metric aimed to detect indicators of strategic and/or tactical potential in games. It is simple to implement and fast to run, typically providing a result for most games within seconds.

Future work will include evaluating the technique across a wider range of games; unfortunately general game players do not typically provide obviously flawed games for counterexamples. The Backgammon family of games will be especially

[1] This game was invented for this experiment as a counterexample; such flawed games are difficult to find in practice.

[2] https://boardgamegeek.com

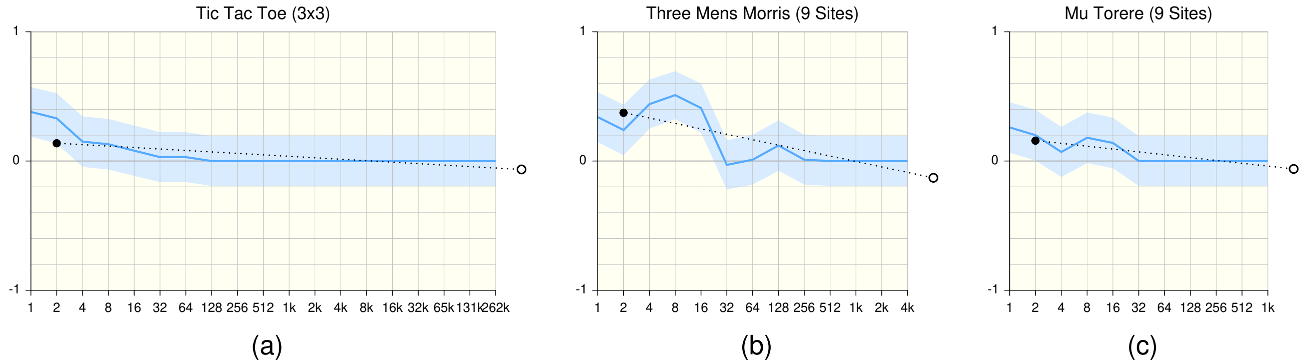

Fig. 4. Skill trace estimates of three small games on similar boards: (a) Tic Tac Toe, (b) Three Mens Morris and (c) Mu Torere.

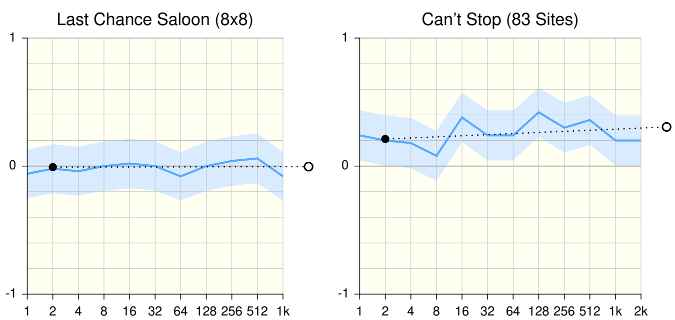

Fig. 5. Skill trace estimates of games with strong chance elements.

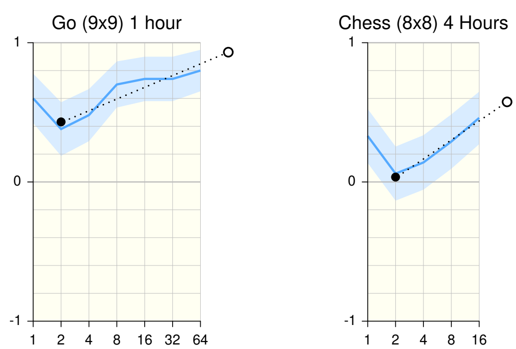

Fig. 6. Skill trace estimates of more complex games: Go (9×9) and Chess.

interesting test cases, as these games have evolved over the centuries to achieve a fine balance between skill and luck.

ACKNOWLEDGEMENTS

Thanks to Stephen Tavener for testing these ideas in Ai Ai. This work was funded by the European Research Council as part of the Digital Ludeme Project (ERC CoG #771292).

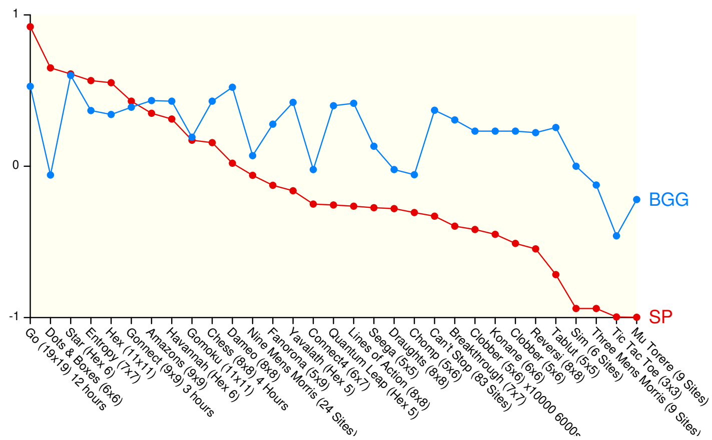

Fig. 7. Plot of BGG and ST scores for 31 games implemented in Ai Ai.

REFERENCES

[1] P. Borm and B. Genugten, "On a relative measure of skill for games with chance elements," *TOP: An Official Journal of the Spanish Society of Statistics and Operations Research*, vol. 9, pp. 91–114, 2001.

[2] C. Browne, "Modern techniques for ancient games," in *IEEE Conference on Computational Intelligence and Games*. Maastricht: IEEE Press, 2018, pp. 490–497.

[3] C. Browne, D. Soemers, É. Piette, M. Stephenson, M. Conrad, W. Crist, T. Depaulis, E. Duggan, F. Horn, S. Kelk, S. Lucas, J. Neto, D. Parlett, A. Saffidine, U. Schädler, J. Silva, A. de Voogt, and M. Winands, "Foundations of Digital Archæoludology," http://arxiv.org/abs/1905.13516, Schloss Dagstuhl Research Meeting, Germany, Tech. Rep., 2019.

[4] G. Costikyan, *Uncertainty in Games*. Massachusetts: MIT Press, 2013.

[5] P. Duersch, M. Lambrecht, and J. Oechssler, "Measuring skill and chance in games," Univ. Heidelberg, Department of Economics, Discussion Paper Series 643, Tech. Rep., 2017.

[6] I. Finkel, "Discussing Ancient Board Games," https://www.youtube.com/watch?v=wFoWj7BAMo8, 12 May 2017.

[7] F. Hausdorff, "Dimension und äußeres maß," *Mathematische Annalen*, vol. 79, no. 1–2, pp. 157–179, 1919.

[8] L. Kocsis and C. Szepesvári, "Bandit based Monte-Carlo planning," in *Machine Learning: ECML 2006*, ser. Lecture Notes in Computer Science (LNCS), J. Fürnkranz, T. Scheffer, and M. Spiliopoulou, Eds. Springer, Berlin, Heidelberg, 2006, vol. 4212, pp. 282–293.

[9] F. Lantz, A. Isaksen, A. Jaffe, A. Nealen, and J. Togelius, "Depth in strategic games," in *AAAI '17 Workshop on What's Next for AI?* San Francisco: AAAI Press, 2017.

[10] B. B. Mandelbrot, *The Fractal Geometry of Nature*. New York: W. H. Freeman and Company, 1977.

[11] H. J. R. Murray, *A History of Board-Games Other Than Chess*. Oxford: Clarendon Press, 1952.

[12] S. Tavener, "Uct skill ladders for small games," *Ai Ai Pages*, 2020. [Online]. Available: http://mrraow.com/uploads/AiAiReports/uct_skill_ladders.html

[13] ——, "Ai Ai Home Page," http://mrraow.com/index.php/aiai-home/, accessed: 30-11-2021.

[14] M. J. Thompson, "Defining the Abstract," *The Games Journal*, 2000. [Online]. Available: http://thegamesjournal.com