The Creativity of Computers at Play

David C. Moffat
Department of Computing
Glasgow Caledonian University, UK.
D.C.Moffat@gcu.ac.uk

Abstract. There are many domains where creative software is being energetically developed, from writing and art to music and mathematics. These domains are open, without clear measures of value, and usually depend on humans to judge the creativity. While such research is obviously relevant to the nature of creativity, it may be that another creative domain is relatively overlooked; namely, that of puzzles.

This paper proposes the game of chess as a good domain in which to demonstrate, investigate and develop computational creativity. It shows some initial comparisons on two chess puzzles, one of which novices or even non-players could follow. The results support the case for computational creativity of programs that play in this domain. In conclusion, all puzzle or strategy games are suitable research testbeds for creativity, both natural and artificial.

1 Introduction — Creative domains

There are many domains where software has been tested for creativity, and is being energetically developed, from writing and art to music and mathematics. These domains are open, without clear measures of value, and typically depend on humans to judge the creativity. While such research is obviously relevant to the nature of creativity, it may be that another creative domain is relatively overlooked; namely, that of puzzles and play.

1.1 Games as a domain for computational creativity

Within the subject of games, AI has been able to make several large contributions. Most of them are general AI techniques, but one or two belong more specifically to the sub-field of computational creativity. First, let us recall that solving problems can be a creative activity, even if the solution is already known to somebody else.

Some researchers take the position that video games are highly relevant for the field of computational creativity. Liapis et al [7] go as far as to call games the "killer app for computational creativity." I certainly agree with their promotion of this perspective; but even they limit themselves in this recent position paper to matters which are generally forms of procedural content generation. My argument here pushes into the different role of computer as *player*.

1.1.1 Solving problems can be creative.

It is often said, at least in passing, that it takes creativity to solve (hard) problems. Engineering and design are creative endeavours, after all; and they consist largely in solving problems. They are not considered to be part of the "creative industries" however: they are not *called* "creative" (in the English-speaking world), and so they tend to get passed over in favour of the more overtly artistic domains. Even engineers themselves (such as AI researchers) tend to have this bias, as is evident in the field of computational creativity.

That is unfortunate, it seems to me, because the arts are in some ways still too challenging for the research field of computational creativity. In particular, to assess the quality of the supposedly creative products (computer generated art, music, jokes and poetry) requires human judgement; and that is extremely slow compared to computer speeds. Research could progress very much faster if only computers were set to work in a creative domain that did not depend on human reaction (at least not in real-time).

The suggestion of this paper is that we do have such a creative domain, and that it is relatively overlooked so far. The domain is that of games; and in particular the *playing* of them. Games are often puzzles in their own right, or they include puzzles within them, as modern video games do. In a typical story based video game, the player is expected to make decisions without having enough information to be sure, and without being able to foresee all the consequences. That is in essence a form of puzzle. There are puzzles placed throughout such games in their "levels" or areas within the virtual world where part of the story takes place. The player has to solve these puzzles before being able to move on through a door, or to the next level.

1.1.2 Games in computational creativity today.

Games are in fact a domain for the field of computational creativity, in the form of video games, and that is because it takes a great deal of labour to make the content for such games with their virtual worlds for player's characters to wander around in.

In order to save costs, video game programmers naturally make specialist software tools to help the designers generate the so-called "levels" of the game. The levels are virtual spaces filled with objects like: trees and houses, roads and walkways, obstacles and vehicles, and computer-controlled "non-player" characters, and the instructions they need to help them navigate around the space in an apparently intelligent way. In the bigger games there are many levels or areas with whole farms, fields and forests, and the virtual towns and cities have to be planned out just as real cities have town planners. To generate so much content for games is only feasible because of the specialist software that takes up much of the burden.

These software tools are increasingly automated, and able to make more appropriate design decisions, to better help the human designers. What the tools do is called "procedural content generation" (or PGC).

1.1.3 PGC is not play.

PGC is an increasingly important part of the industry, as well as an active area of academic research in computational creativity (or AI). Because it helps in the creative process of game design, PGC is obviously a part of the field of computational creativity. But PGC is AI for the making of games, not the playing of them; and it is *play* that is the focus of this paper.

There are other common AI contributions to games, including the use of finite state machines, fuzzy logics, decision trees, search algorithms, and occasionally even neural networks and genetic algorithms. These are AI, but are not part of the field of computational creativity. Neither are they uniquely applied to games, but are rather general techniques developed for and applied to other domains.

The work on search algorithms for games is a healthy and exciting research area these days, especially with the recent developments in Monte Carlo Tree Search (MCTS). Search algorithms like this are used to plan moves in puzzles and adversarial games, usually, like chess. In other words, search algorithms are used to make computers *play* games, but are seen as a mainstream AI technique that is useful for games, rather than as belonging to the sub-field of computational creativity. If that is an oversight, then it is the aim of this paper to correct it.

As other authors have recently noted then, PGC is an active and rich area for computational creativity [7] and [3]. However it is the computer as player that is of interest to me here, and is the area that is still treated relatively lightly, in my view.

1.2 Games and puzzles in AI history

While games have some overlap with computational creativity, they have been far more important to AI in general. It could be asserted that no other domain has been more important to AI, in fact. Let us first consider why that might be so, and then go on to reconsider creativity in that context.

1.2.1 AI has been at play since it began.

In a curious parallel to human development, the field of AI began playfully, before turning to more serious matters as it matured.

Even before modern digital computers existed, thinkers like Turing [10] and Shannon [9] were designing chess playing algorithms, and speculating that computers would one day play chess well enough to beat human players. If only they could have seen how right they were!

Rather like a child, AI in the early days was fed on challenges that led its development, including games like chess and checkers, and puzzles like trying to plan how to put childrens' toy wooden blocks on top of each other in a certain order. These tasks are usually called "toy problems" but they surely count as puzzles as well.

Games and puzzles were chosen as development challenges because they are formally and concisely specifiable, with clear goal conditions, and yet only humans could play them. Being thus characteristic of human intelligence, they were naturally seen as natural aims for computers (AI) to tackle. In the very name of AI, the early preoccupation with intelligence is clear to see. However, the related concept of *creativity* was mentioned much less often than intelligence. It still is, to this day, and indeed the research effort that declares its interest in creativity is tiny compared to the world's AI research.

On the other hand, when humans play, they are often said to be creative, in the way they develop interesting strategies or styles of play, or in finding novel but useful solutions to problems. Before we dismiss the possibility that computers might be creative in the way they play games, or solve problems, we should examine how humans are creative in play, if they are.

2 Creativity and play :

2.1 Play is creative for humans

Children and young animals are naturally playful. They play as part of growing up, in order to learn about their world. Humans are especially busy with play of all kinds, as first recognised by the Dutch historian Huizinga in his classic book asserting the layful nature of man, *Homo Ludens* [4]. Especially for humans, games are used to structure interactions and provide a context in which children (and adults) can play. This leads their cognitive and social development.

2.1.1 Play also encourages creativity.

This is partly because of the nature of the playground, which is a place of safety, but where different roles can be acted at the same time. Players can pretend to perform actions that in real life would be dangerous or impossible. For example, little boys often love to play with toy guns, and pretend they are shooting at each other. Later on, they may play first-person shooter video games like "Medal of Honor". Although they are bigger boys by then, or even full grown men, and the game has more "adult content", they are nevertheless still essentially playing as they did when they were little boys, with pretend guns. It is the safety of the game situation, and the pretence of it, that encourages a creative approach. Because there can be no serious consequences, and the danger is only pretend, and not real, it allows experimentation with different acts, from the illegal to the lethal and from tabu to terrorism.

Experimental thinking is necessary to creativity, as is taking the chance of being wrong. Making poor decisions in real life can have grave consequences, but in games failure is an opportunity to learn by trying again. Trying more risky actions, or a wider variety of actions, means that a there is more chance of discovering actions or decisions that lead to success eventually, even if initially they did not seem to. The style of thinking or problem solving in games or puzzles is thus ideally suited to finding new ways to achieve the desired goals. After more playing, more and better ways to win may be found. Eventually the player or puzzle-solver can discover the best and most elegant solutions: and these can properly be called "creative."

2.1.2 When is a puzzle solution creative?

The two most typical characteristics of creative products are commonly held to be *novelty*, and *quality* (or value). That is by now approaching a consensus [12]. We may question the novelty and the quality of a solution then, but is the "solution" the answer to a puzzle, is it something else, like the way that the answer was found?

To simplify the discussion at this point, let us consider the creativity of the *product* of thought, and not of the *process*, nor of the *producer*. The thinker of thoughts (the producer) is either a human or a computer, but we do not want our assessment of creativity to fall into a confusion about the nature of the thinker, such as whether it is warm to the touch, or as cold as metal. A definition of creativity that depends on body temperature has clearly gone wrong somewhere.

The way that thoughts are produced may be called creative with more legitimacy; but as some other authors do [12] I shall exclude this matter from the discussion, at least for this paper. That leaves the question of whether the *product* of the thought processes (or calculations or algorithms) can be creative.

In the case of the solving of puzzles then, and of the playing of games which are often sequences of problems, we wish to know whether any solutions that are found can be called creative. If they are, then we should call those solutions creative, no matter who or what found them (e.g. human or computer).

2.1.3 On the novelty *of solutions*

Certainly for games and puzzles, the notion of creativity is immediately under threat here, because the solution must already be known by the person who sets the puzzle. Any game must have a way to win, and there must be a way to solve any puzzle, and there must be a way to check when the players have solved it correctly. Otherwise, they will get frustrated with wasting their time if there is no solution for them to find.

Following Boden's distinction between H-novelty (historical novelty) and P-novelty, we note simply that puzzle solutions are not H-creative, because the solution was already known [1, 2]. However, as the puzzle solver did not know it yet, the solution is new to him or her or it, so it is P-novel (for psychological novelty).

In a research strategy where we wish to study the psychological processes of creativity, this P-novelty is the ideal notion for us. It means that we can evaluate how well different algorithms perform in finding solutions that we already know about. To study algorithms that are aimed at H-novelty would be to apply our knowledge of creative processes, excitingly but would be appropriate only *after* we have gained the knowledge; and that can be arrived at best by studying P-novelty first.

Note that the creative *process* has just returned, uninvited but naturally enough, in that last point.

2.1.4 On the quality *of solutions*

As well as P-novelty, we need our problem-solving algorithms to produce *good* solutions, before we can call them creative. Here again, it is an advantage to research into games and puzzles as problem-solving domains. The evaluation of solution quality is typically built into the game or puzzle as part of its specification, usually in the form of a points score.

2.2 Is AI at play creative?

Although we left the issue of *process* behind, and attempted to make the final *product* bear the test of creativity alone, consideration of the extra criterion of *surprise* brought the *process* issue in again through the back door. It might be that the character of the *process* is what will ultimately determine whether we think that an algorithm is creative.

The source of creativity is still disputed in the field, with some researchers such as Indurkhya [5] including the audience or culture and society at large as co-contributors. That is an interesting view, but here we focus on the cognitive process as a determinant of creativity.

First let us consider playful algorithms as candidates for computational creativity. If people can be creative in the way they play games, then when AI plays games, and solves puzzles, is it being creative as well? Let us take the game of chess as an example.

3 Chess for (creative?) computer play

There is a deep history of chess in AI, which makes it a potentially rich domain for the field of computational creativity if it can be shown to be relevant in that regard. The world of chess is itself rich, and includes many forms of chess play, and other playfulness. Let us focus here on chess puzzles, or "compositions."

Iqbal and Yaacob [6] reported an extensive study on chess puzzles, and their aesthetics for human observers. They showed some of the major components of a chess puzzle that people would see as beautiful. This is interesting and innovative work on the *beauty* of chess, and related to, but not the same as, my concern here; which is the potential for *creative play* in chess. Let us turn to a couple of example chess puzzles or "compositions" that are beautiful, but also can be called creative.

In a composition, a strong player (such as a chess Grandmaster) sets up a position on the chessboard and challenges us to find the winning play. An example is shown here, in Fig. 1, with "white to play and mate in two moves." The composition is by the famous chess player Susan Polgar, who was a child prodigy and the first ever female player to become a full Grandmaster in her own right.

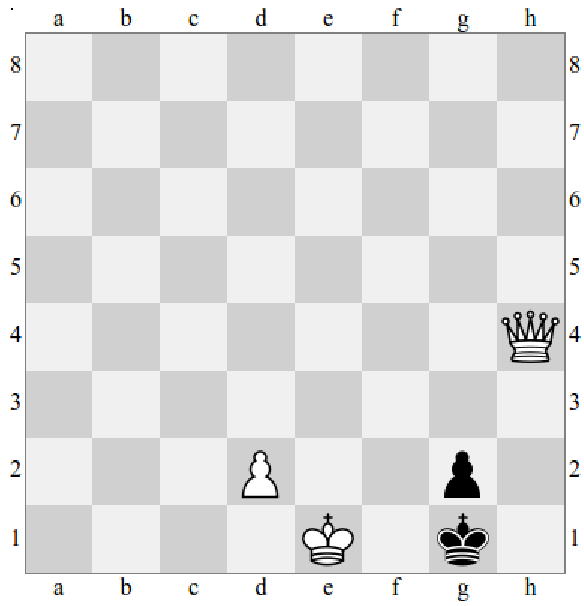

Figure 1. White to play and mate in two moves. From: [8].

A more complex composition, in Fig. 2, (from [11]) is also by Susan Polgar.

This is quite a difficult puzzle, which Polgar has specifically asked people to try to solve themselves, without using the help of a computer. The author of the article is a chess columnist, who loves chess compositions, but took a whole evening to solve this one. The solutions to both of these puzzles are in the next section, in case readers wish to try to solve them on their own first. That will help to give a sense of any creativity needed or involved in solving the compositions.

In both cases, the common characteristics of good chess compositions are on show. The puzzles are difficult to solve, intriguing because the obvious attempts are not correct, and therefore contain an element of misdirection. It is as if the composer anticipates the thought processes of the solver and baffles them. To solve such puzzles quickly is therefore an impressive feat, and shows some deeper understanding of the chess positions.

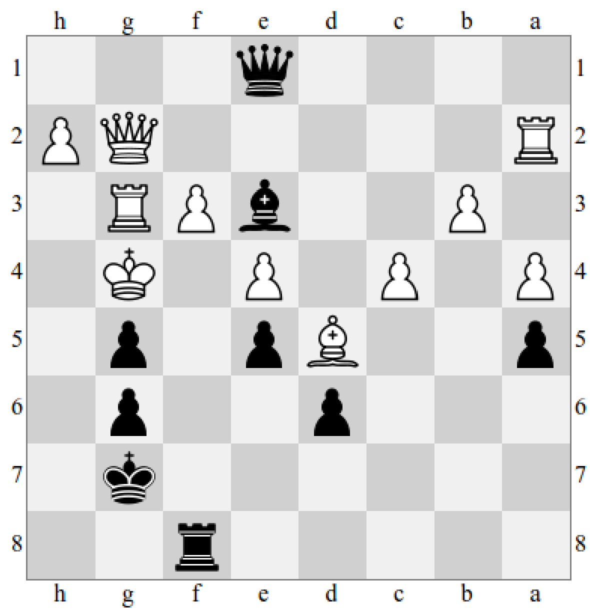

Figure 2. Black to play and mate in three. From: [11].

The upshot is to create a feeling of surprise in the solver, when the solution is finally shown; or else if the solver finds the solution himself, there is a feeling of satisfaction, and appreciation of the artistry in the composition if it is a good one.

3.1 Computer performance on the puzzles

While it takes a human player some time to solve the puzzles, computer programs can solve them much quicker. To illustrate this, a modest but convenient computer player was tested with both puzzles (available at http://www.apronus.com/). It runs in a Javascript browser, and was timed on a small notebook computer with only 1GB of RAM memory and a 1.6GHz Intel Atom CPU.

The first puzzle is relatively easy, and a fair player might find the solution in well under a minute. The computer found the solution in 200ms. (It is to move the white king away from the black king, giving him space to move out, which is then his only option; but luring him into a trap. The queen swoops down next to him and it's checkmate. `1. Kd1, Kf1. 2. Qe1 #`).

The second puzzle is more serious, and even most Grandmasters would probably take at least five to ten minutes to solve it. The same computer took only 700ms. Weisenthal gives a nice walk-through of the thought processes of a typical player trying to solve the puzzle, which even a novice player could follow. He shows how such compositions are constructed to mislead and tease the solver [11]. (The trick is to see the second move, which is a relatively quiet one, not suggesting itself to the typical chessplayer; and that white is then oddly helpless against the quiet threat. `1. ..., Rf4. 2. K x g5, Bb6. 3. ..., Bd8 #`).

3.2 Assessment of the computer's creativity

Can we say that the computer algorithm that solved the two compositions is creative? Well it finds the correct solution, which it did not know beforehand, so its product is both novel (to itself) and valuable. Indeed the computer is exactly as creative as any human solver by this reckoning; but as the computer is so much faster, it is that much more "creative", in the terms given above.

What about the extra criteria of creativity mentioned earlier, namely that os *surprise*? The surprise is built into the puzzle by the composer, in the sense that it was designed to have a non-obvious solution that would thus be hard to find. This property is again equal for both computer and human solver; but again the computer's great speed tells in its favour.

Objectively then, by the criteria of creativity laid out in this paper, and on the results of this limited test of two puzzles, the computer is more creative than any human expert player.

That may be an astonishing and unwelcome conclusion for some readers, especially given that the chess algorithms were never written in order to specifically address the question of computational creativity in the first place.

3.3 Possible objections and resolution

One common objection to this claim of computational creativity will be to complain that computers and only calculating their way to a solution. In this case they are executing a "brute force" search. This is an appropriate term for chess algorithms, and indeed it is exactly how it was envisaged from the beginning of AI by founders like Shannon and Turing that computers would come to play chess. The ironic wit in the term is deliberate — the computer is displaying only a brute form of intelligence, and yet with such power that it gives an uncanny impression of genuine intelligence.

This objection of brute force, or of mere calculation, is a classic objection to AI in all its forms, and is immediately persuasive to ordinary people, as well as many experts. However, it is not quite fair as a supposedly unfavourable comparison with human cognition, for the following reasons at least:-

1. computer "cognition" is apparently very different, but that does not make it necessarily inferior or worse. To assume that anything different from us must be inferior is characteristic of racism and xenophobia, and is outwith science.
2. human cognition is itself not well understood in any case. This makes it too tempting to overstate any claim that other cognition is different from it, without having any solid basis.

While it is true that we feel that our human thought processes are often intuitive, and not to be explained, they are also successful at the same time. This gives our own creativity a mystique that we cannot attribute to algorithms once we understand how they work. But again, to rest on a vague concept like "intuition" as the key distinction between two supposedly different kinds of cognition, seems too hasty and unsound.

4 Conclusion

Starting from a commonly shared notion of what creativity is, we have taken a tour through some chess puzzle territory, to explore the possibility that chess algorithms might be good models for computational creativity. We found that computer performance in this respect is high, and that we are thus bound to accept that computers are creative, or else we have to re-examine our conceptions and definitions of creativity.

Computers in this domain can easily exceed human performance, which is already a contribution to the field of computational creativity. However the main intention of this paper is to establish the viability and even suitability of computer games, with chess as an example, as a research domain for the field. It appears in conclusion that this potential may have been generally underestimated to date. Reasons for this might include a general prejudice against rational reasoning as being creative; or against computers especially. But whatever reasons for it there may be, the point remains that computers and algorithms, as game players and puzzle solvers (not only composers), are not yet fully appreciated by the field, which continues to devote more attention to the arts. As the area of games and puzzles is more tractable however, for evaluation especially, we should expect better progress with this as a research domain.

5 Acknowledgements

I would like to thank the two anonymnous reviewers for their comments and suggesting a couple of references.

REFERENCES

[1] Margaret A. Boden, 'Creativity and artificial intelligence', *Artificial Intelligence*, **103**, 347–356, (1998).
[2] Margaret A. Boden, 'Computer models of creativity', *AI Magazine*, **30**(3), 23–34, (Fall 2009).
[3] Michael Cook and Simon Colton, 'ANGELINA – coevolution in automated game design', in *Proceedings of the Third International Conference on Computational Creativity*, p. 228, Dublin, Ireland, (may 2012).
[4] J. Huizinga, *Homo Ludens: A Study of the Play-element in Culture*, Beacon paperbacks ; 15 : Sociology, Beacon Press, 1955.
[5] Bipin Indurkhya, 'Whence is creativity?', in *Proceedings of the Third International Conference on Computational Creativity*, p. 62–66, Dublin, Ireland, (May 2012).
[6] Azlan Iqbal and Mashkuri Yaacob, 'Advanced computer recognition of aesthetics in the game of chess', in *WSEAS Transactions on Computers*, p. 497–510, (May 2008).
[7] Antonios Liapis, Héctor P. Martínez, and Georgios N. Julian Togelius, 'Computational game creativity', in *Proceedings of the Fifth International Conference on Computational Creativity*, (2014).
[8] "Susan Polgar", Wikipedia, 2015, January 12.
[9] Claude E. Shannon, 'Programming a computer for playing chess', *Philosophical magazine*, **41**(314), 256–275, (1950).
[10] "Alan Turing", Pearson; and the A.M.Turing Trust, 1953.
[11] Joe Weisenthal. Here's a delicious chess problem that had me scratching my head for hours. Business Insider website. http://www.businessinsider.com/black-to-move-and-mate-in-3-problem-2013-8, August 14 2013.
[12] Geraint Wiggins, 'Searching for computational creativity', *New Generation Computing*, **24**(3), 209–222, (2006).