CREATIVITY AS DISRUPTIVE ADAPTATION – A COMPUTATIONAL CASE STUDY

TAPIO TAKALA
Helsinki University of Technology, Finland

Abstract. The article studies what creativity means when broken down to the very essential elements. Creative design is assumed to be a process that brings up new and useful products in a surprising way, i.e. against expectations but using a pattern instead of random search. These features are further analyzed with a case study of geometric packing problems. A computational problem solving agent is built that uses a small number of different strategies to place bottles in a case. Its adaptive control mechanism, based on reinforcement learning, leads to fixation of strategy and into a conflict when the strategy does not work. Relaxation of expectations allows it to (re)invent a novel strategy, which in favorable circumstances becomes a new fixation, an innovation. These circumstances are analyzed with respect to the problem domain and the order in which problems are presented to the agent, concluding that creative moments appear when entering the boundary of a subdomain where a different problem solving strategy must be applied.

1. Introduction

Creativity is a fundamental question in current AI research, but still far from being fully understood. Often it is considered to be a human trait alone, even a mystery.

In a review article, Buchanan (2001) treats the issue from psychological and computational points of view. Quoting him, *there is no consensus, just considerable ambiguity, about what we call creative behavior or what is involved in this behavior*. However, referring to Minsky, he concludes that *defining the criteria for acting creatively assures us that a computer can be creative.*

One reason why creativity evades definition is that it is contextual. In the larger context, referring to great inventors or artists, people talk about "big-C" Creativity, whereas similar behavior in mundane everyday activities is "small-c" creativity. In many of her articles, Boden (e.g. 1995) emphasizes the contextuality by distinguishing personal (P) creativity from historical (H) creativity. She also points out that creativity is to act against the rules by dropping or transforming the prevailing constraints. Mere novelty is not enough to define a product as being creative.

As long as we observe creativity in an open world, its definition remains contextual. In order to make a computational model we need a closed world that includes both a potentially creative subsystem (G) and the evaluating criteria (E) for it. Obviously the closed world would be smaller than the real society, or even a single person (figure 1).

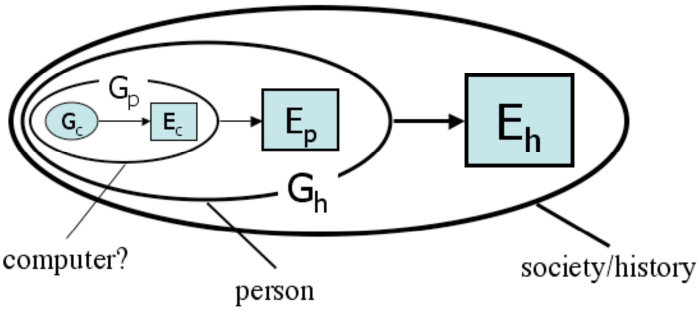

Figure 1. Context defines the expectations (E) against which the creativity of a generative process (G) is evaluated.

Because creativity is anyway contextual, there cannot be any objective definition based alone on the results or products of a creative process. We need to model the recognition of creativity as well. A major point in this article is, that the recognition process actually is more important than the generative part. If a person invents something by accident (i.e. starts using a novel solution), but doesn’t understand that it was a great invention, would he be called creative? Probably not. On the other hand, if an artist accidentally finds, for example, a beautiful piece of wood, and puts it in frames in an exhibition, is she then called creative? Yes, much more probably at least than the dumb inventor.

Researchers in artificial life (Langton 1989) have tried to extract the smallest set of common features to all living things (metabolism, homeostasis, reproduction, etc.), and to make a theoretical definition, based on which any artificial construction can be evaluated if it is alive (or more accurately “a-live”). Our approach is analogous to that, trying to find the minimal definition that applies to all creativity.

The aim of this article is to see what would be a minimal set of requirements for creative behavior, i.e. the smallest possible “$_c$” that still could be called creativity. Based on that it may be possible to start working with theoretical and computational models of creativity, that are free from human subjectivity.

1.1 ESSENTIAL ELEMENTS OF CREATIVITY

Creativity is often defined as a process producing something new and useful. As such this is not enough, because then even a slightest modification of an existing product would make it creative, which is not how we usually understand the word. Instead of just novel, a creative product should be unexpected, or surprising (Boden 1995).

Breaking the rules

But what makes it surprising? In order to know if something is creative, we ought to know the prevailing expectations, or the current habits of designing, against which a new approach can be compared. We need to model not only the agent that generates (more or less creative) products, but also the recognizing agent assessing when the product appears to be creative.

Often creativity is characterized as breaking the rules. However, there are different kinds of rules and different ways of breaking them. Simple ignorance of rules alone leads to anarchy. It may be useful for demonstrating social power (as Alexander the Great did when opening the Gordian knot with a sword), but otherwise it may be just absurd (for example, when breaking valuable objects while packing them in a small space by brute force). Instead of just breaking the existing rules, creativity brings up new rules that appear to be useful.

Creativity does not break such rules that keep the problem setting consistent. However, it breaks the illusory rules that are not actually given in the onset, but are only mentally built by the subject based on what has been useful in past situations (Takala 1999). Recognizing that these expectations are unnecessary leads to the “a-ha” moment of enlightment.

Relative novelty

If the context society does not know the proposed solution at the moment it is presented, it will be considered novel. This may mean that the invention is novel in the historical sense, but may also be caused by the society not having heard about it, or having known it but forgotten.

This brings us to the concept of situated (S) creativity (Gero 2002). Novelty is determined by comparing against the prevailing knowledge, actually against the expected rules only. This is emphasized by examples where a solution is immediately found to be self-evident but still creative: “Why didn’t I come to think about that?”

Gero (2002) also classifies designs as *routine*, if the same set of rules can be used throughout. By definition, if creativity is breaking the rules, this cannot be creative. Non-routine design is further divided into *innovative* and *creative*, the former only extending the ranges of variables in the rules, whereas the latter introduces new variables and thus makes more substantial changes in the design space. Although this classification was presented in the context of productive design act, it can be interpreted within recognition as well. When looking at a design, one may think of it in routine ways, or introduce innovative/creative viewpoints. They do not change the design itself but the way one thinks about it. With this interpretation, the classification conforms with the observations made above.

Illumination and other emotional reactions

Yet another important point in creativity is *reflection*: in order to be creative, a system must be able to observe its own behavior, and recognize the creative moment when rules are broken and a novelty appears. This feeling of illumination may be strong and it has a positive feedback, resulting in pleasure and *satisfaction*.

Satisfaction with past experiences may be the reason, why people tend to use familiar concepts rather than try out new ones, a behavior called *fixation* in this paper. Design fixation is a particular example of this, denoting a designer's premature commitment to a familiar solution, which then prevents considering other alternatives (Purcell and Gero 1996). The term is used here in a similar, though more general sense. Fixation is a necessary component in creativity, as it delimits the ways one will think, causing some ways to be more probable. These expected ways of thought are the rules to be broken in a creative act.

When a person is faced with a problem that does not get solved with the expected rules, a *conflict* appears. This is also important for creativity, as it triggers the need to think differently (Killander and Sushkov 1995). What typically happens in a creative process then is that the person, after trying to solve the problem persistently but in vain, runs into *frustration* and temporarily gives up. That starts the incubation period, during which the subconscious mind continues an opportunistic search with *relaxation* of constraints, and possibly comes out with an alternative, creative solution.

A creative mind does not freeze after a solution is found, but continues to face new situations, some of which may call for creative behavior. The satisfaction after a success may again be the underlying reason for *curiosity*, a behavior seeking for interesting problems to be solved, i.e. those potentially leading to novelties (Saunders 2001). Thus the creative personality continuously runs the cycle from fixation to conflict to illumination and to new fixation again.

The following list summarizes the set of features typical in creative behavior as described above. Tentatively these are proposed to be necessary requirements forming a definition of creative processes.

- rules/constraints for generating designs or interpreting them
- expectations, i.e. fixation to a set of rules
- problem triggering a conflict
- frustration, relaxation and opportunistic search
- reflection: recognition of novel rules
- positive emotional feedback (creative moment)
- curiosity

The practical research question now is: Can we computationally simulate the process and its essential features? If so, we have formulated them clearly enough to be represented in an artificial world. Experiments there may then confirm if the definition was right, i.e. if the chosen principles result in enough creative-like behaviors.

2. The case study

As a simulated world cannot contain many different aspects of real life, we need an illustrative case where the essential features of creativity can be demonstrated in an extremely simplified form. The core issues to be considered here are:

- a well-defined problem with a small number of parameters, in order to be simulated easily
- enough problem variability to span a domain of interesting, non-trivial problems

Geometrical packing problems appear to be a fruitful domain for our case study. They are important in practice, as even a small improvement in efficiency may mean huge savings in logistics industry. They are also non-trivial, as can be seen from the vast number of cases without a proof of optimality. They are easy to represent computationally, but often appear to be NP-complete, i.e. there is no better algorithm for finding a guaranteed optimum solution than an exhaustive search, or even worse: because the variables are continuous, a discrete exhaustive search is impossible. Of course, however, there are many heuristic strategies that work better in the average and that can find good approximate solutions.

Generally the task in a packing problem is to design a configuration of geometric objects, such that they fit inside a given constrained space.

Often they are expressed as optimization tasks, to minimize the container size for a given number of objects to be packed, for example. There are various different classes of problems, depending on the shapes of objects and the given constraints, see e.g. (Friedman 2005).

2.1 PROBLEM DEFINITION

In this study the particular problem was to pack a given number of bottles in a case. Assuming the bottles are upright in a box, the problem becomes two-dimensional packing of circles inside a rectangle. Placing one bottle into the case reduces to finding a point for the circle center that is at least the amount of its radius away from the box boundaries and from the other circles.

Depending on the number of bottles and size of the box, there may be an infinite number of solutions, a singular unique configuration, or the problem may be impossible to solve.

Figure 2a shows a typical commercial case, the six-pack, which appears to be the smallest configuration of six bottles. However, in large cases the rectangular placement is not optimal, as the best packing approaches the hexagonal (honeycomb) grid of figure 2b when the number of bottles increases. As will be seen below, neither of these well-organized patterns is always the best. Either one, or even a more irregular configuration may be optimal in a specific case.

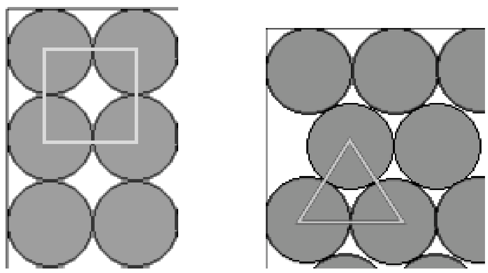

Figure 2. The (a) rectangular grid (six-pack) and the (b) hexagonal grid.

For this study, the bottle size and number of bottles for each case was fixed. On the other hand, size of the case was variable in order to generate a number of different problems. The overall problem setting thus was to make different six-packs, i.e. try to fit six bottles of same size into a rectangular case of given dimensions.

At first sight the problem may seem trivial. However, it is not easy to prove that a particular configuration is optimal, neither is it easy to find a rule for creating a satisficing design for a tight box. And if generalized to larger numbers of bottles, the problem becomes more complex with sometimes very irregular or even surprising optimal solutions.

2.2 THE METHOD

Our system is a simulated software agent that takes one bottle at a time and places it into the case of given size. If all six bottles fit in, the problem is successfully solved.

Potentially there is a huge number of possible approaches to the problem. To make it tractable, only three predefined strategies for placing a single bottle were used. Two of these are based on the rectangular and hexagonal grids of figure 2. The third strategy is simply a random placement.

Technically, the available places in a case are checked by drawing circles and box boundaries extruded by the circle radius (as in figure 4). This shows the occupied area where a new bottle center cannot be placed. Out of the remaining points inside the box, a position is selected using one of the three alternative strategies:

- seek a horizontally or vertically closest position to a previously placed circle. This always results in a rectangular grid, although not necessarily with an even length of rows.
- take the highest available point (and the leftmost if there are several equally high). This usually results in a hexagonal grid which is known to be optimal for infinite spaces, but also in specific cases generates many different variations.
- place the circle at a random free position. This does not usually make regular or dense arrays but, if we are lucky, it may in some cases result in better designs than either of the rules above..

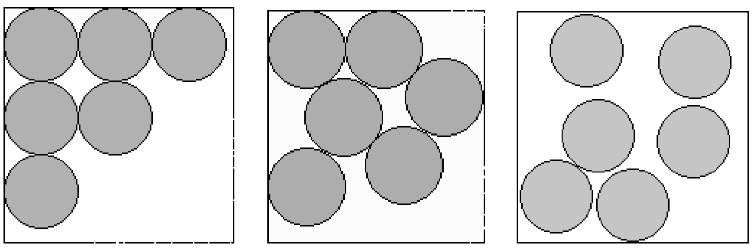

Figure 3: Sample configurations with each strategy:
(a) rectangular, (b) hexagonal and (c) random

If a placement strategy does not immediately succeed, a retrial is allowed for at most a given number of times (this particularly applies to random placement, where the trials are substantially different – the other algorithmic strategies result in the same configuration again).

In the retrials, another strategy may also be randomly chosen, based on probabilities assigned to each strategy. The probabilities are controlled by reinforcement learning, i.e. in case a strategy succeeds, its probability increases, whereas with a failure it decreases (the other probabilities are tuned accordingly to keep there sum as one). Every time a new problem is faced, the strategy with highest probability is tried first. This adaptive learning makes a positive feedback, causing the system to show persistence and to fixate in a recently successful strategy, even if others might do as well. If the dominant strategy does not work out, its probability decreases with retrials until another one becomes highest. Then the fixation ends and another strategy is tried out.

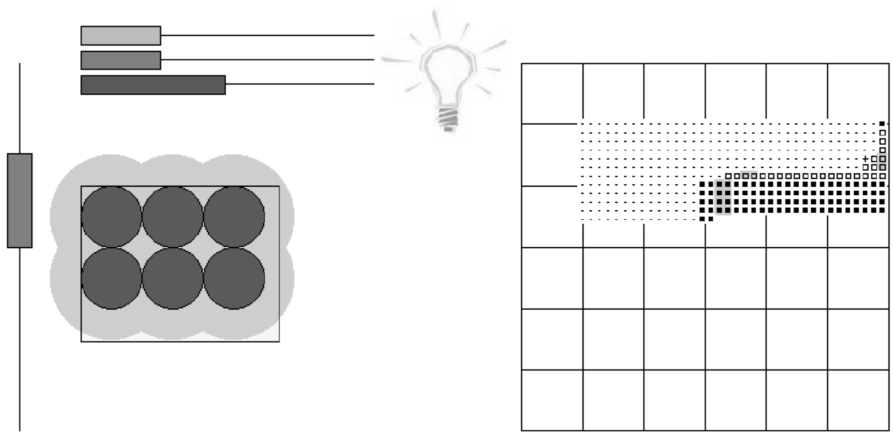

Figure 4. Snapshot of the program at the moment of enlightment, displaying the current problem (box with circles), probabilities for each strategy (bars above the box), the solution space (to the right), and a satisfaction indicator (to the left). The lighting bulb indicates that the system has just found a new applicable strategy.

Positive feedback from a successful solution is also separately cumulated as a “feeling of satisfaction”, which raises quickly, just like the probability of the successful strategy does. With a failing trial the satisfaction decreases, though slower than it raises. From changes in this emotional indicator one may recognize features belonging to a creative process: trials with a bad strategy fail and lead to dissatisfaction. But when a better strategy (accidentally) is found and used a few times, then satisfaction increases quickly. This sudden rise of satisfaction, corresponds to the “a-ha” experience. The system knows that it has found a new solution method, which consistently applies to those problems it has been trying to solve. In the figures 4-7 these cases are indicated with grey color. Figure 4 shows how the concepts discussed above are visualized for tghe observer.

Note that anthropomorphic terms, such as “fixation”, “frustration”, “creative moment” and “curiosity”, are purposely used throughout this article in order to point out the proposed analogies of the system’s behavior to creativity in real life. They are not to be taken literally, but in a metaphorical sense.

3. Results from test cases

When given a problem, the agent tries to solve it by placing circles in the case. It may try several times with different strategies, but eventually it either reports a solution or gives up after a limited number of trials. In order to study its dynamic behavior, another agent was programmed that serves the first one continuously with new problems. In the test runs I used different strategies to scan the problem space. Interesting phenomena were encountered, that are reported below in more detail.

The results are visualized with a diagram displaying the two-dimensional problem domain with width and height of the rectangular case as its axes. Each point represents a particular problem defined by the corresponding box dimensions. The origin is at the top-left corner and the grid markers are at units of one circle diameter. Thus, for example, the box fitting a conventional six-pack would be at points (2,3) and (3,2), as shown by the circles in figure 5. The marker symbol tells the strategy used to solve the problem at each point.

3.1 FIXATION OF STRATEGY

Generally the problem solving agent easily gets fixated to a particular strategy after a small number of successful trials. This is demonstrated in figure 5, where each horizontal row is a left-to-right scan of problems, i.e. the width of the box is increased in small increments while keeping the height constant. At the left end there is an impossible problem (too small box to fit six circles), shown with a (-) sign. During the scan the agent tries the three different strategies more or less randomly. When a problem becomes solved, the strategy successfully used is tried again for the next problem. As this is an easier one (a larger box), the agent succeeds again and then the same strategy is used throughout, resulting in a row with one repeating symbol only.

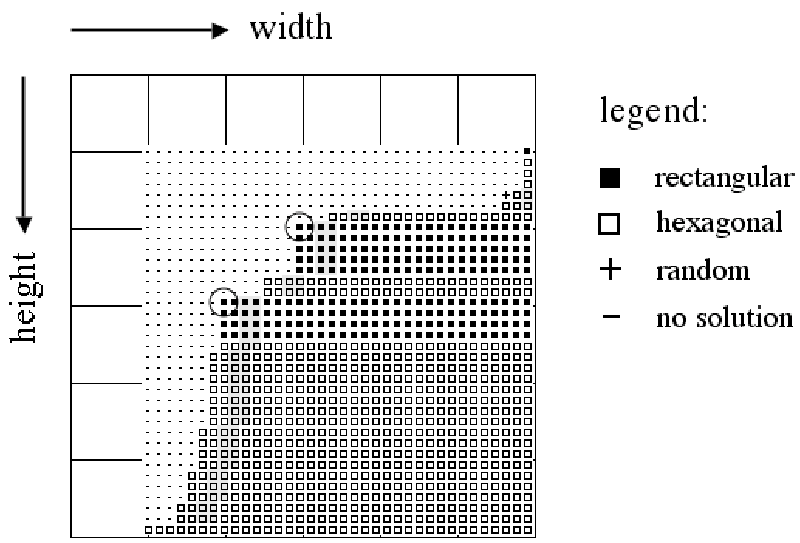

Figure 5. Fixation of strategy during a left-to-right scan of problems. Conventional six-box is marked with circles.

3.2 CREATIVE BEHAVIOR

A more interesting phenomenon happens when the agent surprises itself by reinventing a strategy it has forgotten (i.e. has not used for some time). This is demonstrated in figure 6, which is the same as figure 5, except that ordering of problems goes from right to left, i.e. from a trivial problem towards an impossible one. In the beginning of a row any strategy will do, and it becomes immediately fixated. Thus the right part of each line consists of similar symbols, which may randomly be any one of the three alternatives. Following the scan towards left, the same strategy may continue until the problem becomes intractable. However, in some places it appears that first the agent has given up (-) but then another symbol appears along the row, indicating that a new strategy has been chosen and used consistently until the problem becomes really impossible.

This is a simulated creative moment that can be distinguished as a disruptive behavior: after successfully solving a number of problems with the same strategy, the system runs into a situation where the usual strategy doesn’t work. With repeated unsuccessful trials it becomes frustrated, possibly giving up. However, deeming the problem impossible may be only seemingly right, and actually it is solvable by another strategy. This strategy is found after relaxation, i.e. by giving up the usual strategy and (more or less randomly) trying something else. When a new strategy succeeds, it is tried again in another situation and, if reinforced, becomes the new fixation. This overall behavior can be seen to resemble creativity as it contains the essential features listed in section 1.1.

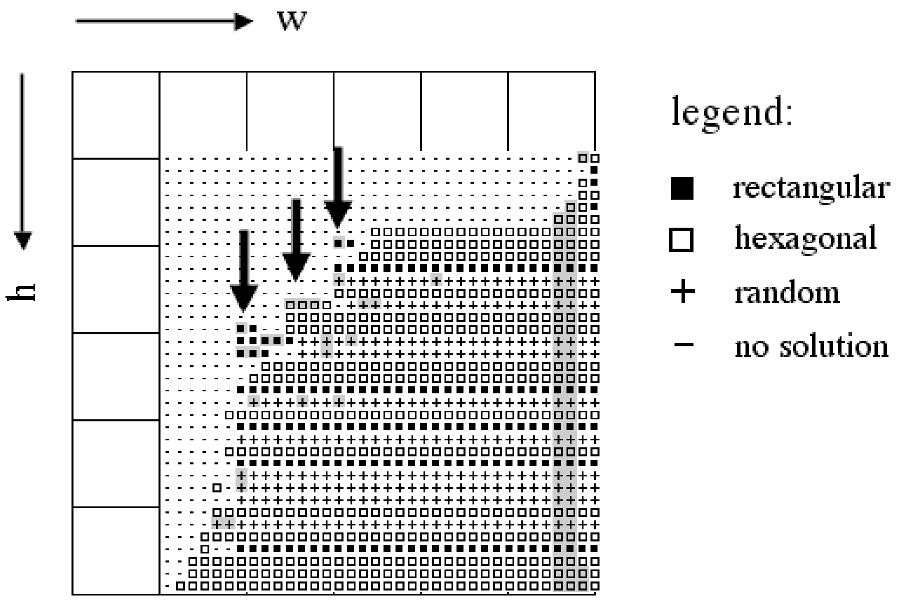

Figure 6. Creative moments appearing as grey spots in right-to-left scan of problems.

From this experiment it seems that the interesting creative behavior appears near the border of impossible problems. Thus a further question is: where are these cases exactly and why?

3.3 RANDOM WALK THROUGH THE DOMAIN OF INTERESTING PROBLEMS

An ordered scan of the problem domain gives an overall view of it. However, the results depend on the order of scans, as seen by comparing figures 5 and 6. Also a complete scan runs through the boring areas of impossible (towards top-left corner) and trivial (towards right-bottom) problems.

In order to explore the most interesting part inbetween, I made a feedback loop of both agents, or combined them into a double-agent. The problem domain is scanned with a random walk method controlled by curiosity that seeks interesting problems. Each time a new problem is given, it is taken from a point nearby the previous one, and on the opposite side to the point before that, but with small random variation. This makes a continuous randomly turning path. In order to keep it in the area of interest, a bias was added, turning the path towards the top-left corner (i.e. more difficult problems) whenever a case has been solved, and to the opposite direction if the trial fails. Some results are shown in figure 7. The path keeps in a rather narrow area, with the strategy symbol alternating here and there. Most consistently marked areas are the three “corners” near the center of the figure. In each of these only a single strategy seems to be applicable.

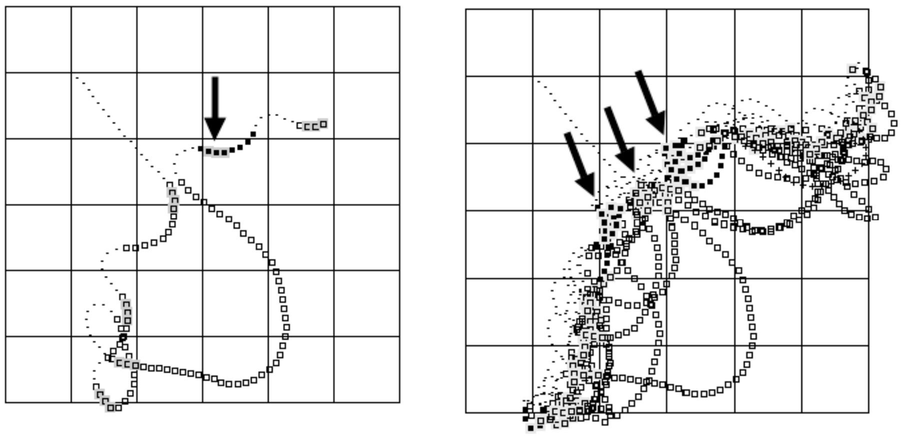

Figure 7. Controlled random walk in the area of interesting problems. Arrow in the left part indicates a point, where a new strategy was found. Arrows in the right part indicate areas where a single strategy seems to be applicable.

4. Discussion – analysis of the problem domain

To get a better understanding of the phenomenon, I explored the problem domain by regular scan with each single strategy in turn.

The subdomains where rectangular or hexagonal grids are applicable strategies are shown with dark grey in figures 8a and 8b, respectively. The trials not solved are depicted with light grey. Figure 8c shows the subdomain of rather trivial problems, as it was produced with the random strategy with at most five trials for each case.

Strictly impossible cases with width or height less than one diameter unit are not considered at all. Other impossible cases are difficult to find for sure, because we do not know all possible configurations, some of which might be better than those produced by simple strategies. Extensive statistical trials would give an impressions of this subdomain.

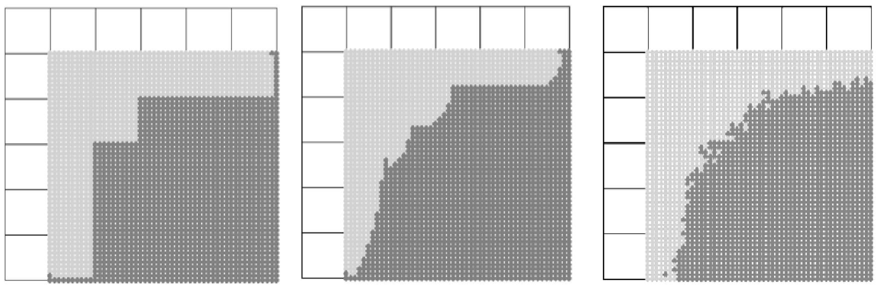

Figure 8. Subdomains for each strategy:
(a) rectangular, (b) hexagonal and (c) random

Most interesting are the non-trivial solvable cases, particularly those solvable by one strategy only. These are found by comparing the maps in figure 8. Most obviously the upper left corners of each subdomain are singular. Located in the points (2, 3) and (3, 2) for rectangular strategy are the conventional six-packs. At points (2.5, 2.73) and (1.86, 3.5) / (3.5 1.86) are the hexagonal dense packages, as depicted in figure 9.

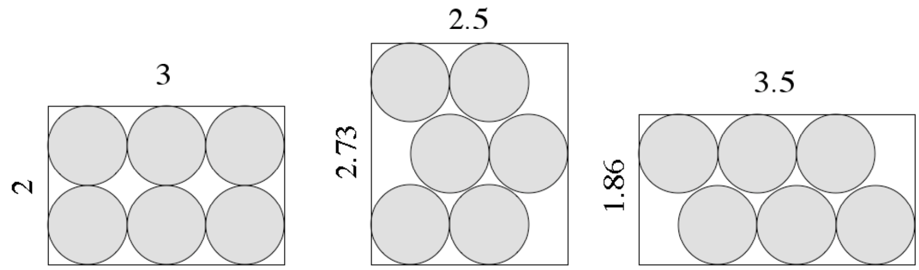

Figure 9. Optimally dense six-packs in rectangular and hexagonal form.

These are qualitatively different cases where another configuration would not work. A continuous deformation could be imagined leading from one to another, but none of our strategies would produce it. This incontinuity in solutions causes a disruptive leap in strategy if the succession of problems goes from one subdomain to another. The disruptive change can be understood as a creative moment.

4.1 IS THIS CREATIVITY?

One might ask how close to real creativity these experiments actually are. They demonstrate such features as producing novel and useful designs (though in a relatively small world), expectations based on experience rather than given constraints (fixation to a strategy), conflicting situations (the fixed strategy does not work), unexpected solutions with enlightment after relaxation of rules (based on random selection), and finally fixation of new rules (if application of the new strategy is successful).

What can be considered missing here are, for example, associative thinking for finding new solutions, sense of humor and aesthetics, recognition of patterns from random configurations, and inductive formation of rules from experiences. These are all found to be typical traits of creative people. Compared to human life, one has to keep in mind that the computer's capabilities are still extremely low. Taking that into account, I think that some new steps towards computational creativity have been taken despite of the missing parts.

One might also ask whether an extensive random search might work out better results than the strategies applied here. Even though this might be true, I think that it would be not as creative as a process. An individual solution alone lacks the pattern that can be applied more generally and that starts a new strategy. In a historical context someone might later recognize a pattern in it and then call the solution creative. In that case, creativity should be attributed to the later inventor.

5. Conclusions and future work

This work aims at demystifying the concept of creativity by demonstrating that it can appear even in very primitive situations in simulated environments, if we accept a minimalistic definition of creativity. Many characteristic features of creativity, as listed in the introduction, have been demonstrated in a simple setting.

Perhaps the most important point is that a model of creativity should include not only a generative element but also one with a partial selection of all potential rules, which forms the situational expectations, against which the generated constructs can be compared for novelty (as the "computer" part in figure 1). Other observations spawning from this are that fixation, conflicts and frustration, which normally are considered harmful for thinking, are also necessary elements in creativity. Yet another element is reflection about the process itself, whether it has broken its own rules. And in order to keep creativity running as a continuous process, positive feedback (satisfaction) and curiosity are also needed. The features listed above are suggested to form the minimalistic definition of creativity.

Potential for creative solutions depends on the problem domain, particularly on the structure of the subdomains of solutions with different strategies. If a single strategy (a pattern or a rule) covers all possible solutions, there is no place for surprises, everything happens as expected. That would be a domain of routine designing.

Creative moments appear at changes of subdomain, thus their realization depends on the order the problems are presented. Fixation to a default strategy is necessary in order to develop expectations, which are necessary for a surprise.

In case of continuous wandering in the problem space, creativity appears at subdomain boundaries. A typical situation is when approaching from more trivial (several strategies apply) towards more difficult area in problem space. If done in reverse order, the first applicable strategy becomes a fixation that is applied all over, and no conflict appears.

An important part missing here, as compared to real world creativity, is autonomous formation of strategic rules. In the case study all the strategies were predefined, and novelty is only relative but appropriate when situated: reinventing the wheel is a creative act, if one has already forgotten the existence of the wheel. In the future we should add inductive learning mechanisms for recognizing new patterns in random experiments and for representing them as formal configuration rules.

In a closed world like the case example here, we may take for granted a fixed number of patterns, each of which may be "reinvented" as new after being forgotten. In a more realistic situation, the inventions are cumulated as history, and new patterns compared to that ought to be invented.

Assuming that we still are in a closed world and there is no mystical creative source, the new patterns can only be formed by combinatorial processes. A future task is to study different possible forms of them. Known approaches are, for example, analogues, associative pattern matching, abductive logic, etc. Collectively these might be characterized as "layered search heuristics" (Buchanan 2001).

In this article I have challenged those previously widely accepted definitions of creativity that look at the generative process and its rule formation alone, without considering the rules as expectations to be broken. In doing so, I have tried to look from a different viewpoint the question, what creativity is. I have presented a set of features that are suggested to be essential to creativity. It remains open to discussion, which of these features are considered necessary, and if the set is sufficient to define creativity. To me, the process has contained several moments of personal creativity, bringing delight. If this approach will be wider accepted in the historical sense, I will be even happier.

In conclusion, the paper points out that creativity may be defined as a process adapting to new situations in a disruptive way, such that prevailing expected rules are suddenly changed. The detailed problems being solved, and the particular strategies to solve them, are not so important. As Buchanan (2001) puts it: *key to creativity is at the metalevel*.

Acknowledgements

I'd like to thank the reviewers for constructive comments that helped to make the point in this article clearer.

References

Boden, M: 1995, Creativity and Unpredictability, *Stanford [Electronic] Humanities Review* **4**(2): Constructions of the Mind. www.stanford.edu/group/SHR/4-2/text/boden.html

Buchanan, B: 2001, Creativity at the Metalevel – AAAI-2000 Presidential Address. *AI Magazine* **22**(3): 13-28.

Friedman E: 2005 (cited), Erich's Packing Center, www.stetson.edu/~efriedma/packing.html

Gero, JS: 2002, Computational Models of Creative Designing Based on Situated Cognition. *in* T Hewett and T Kavanagh (eds), *Creativity and Cognition 2002*, ACM Press, pp. 3-10.

Killander AJ and Sushkov VV: 1995, Conflict Oriented Model of Creative Design, in JS Gero, ML Maher and F Sudweeks (eds.), Preprints Computational Models of Creative Design (3rd Heron Island conference), pp.369-396.

Langton, C: 1989, Artificial Life, *in* CG Langton (ed.), *Artificial Life*, SFI Studies in the Sciences of Complexity VI, Addison-Wesley, pp. 1-47.

Purcell, TA. and Gero, JS: 1996, Design and other types of fixation, *Design Studies* **17**(4), pp. 363-383.

Saunders, R: 2001, *Curious Design Agents and Artificial Creativity*, Ph.D. Thesis, Faculty of Architecture, The University of Sydney.

Takala, T: 1999, Simplistic jokes and drawings as a guide towards computational creativity. In JS Gero and ML Maher (eds.*), Computational Models of Creative Design IV*. University of Sydney, 1999. pp. 33-46.