AISB Convention 2015, 20-22nd April, Canterbury

The Society for the Study of Artificial Intelligence and Simulation of Behaviour

AISB Convention 2015
University of Kent, Canterbury

Proceedings of the AISB 2015 Symposium on Computational Creativity

Edited by Mohammad Majid al-Rifaie, Jeremy Gow and Stephen McGregor

ISBN: 978-1-908187-50-5

Introduction to the Convention

The AISB Convention 2015—the latest in a series of events that have been happening since 1964—was held at the University of Kent, Canterbury, UK in April 2015. Over 120 delegates attended and enjoyed three days of interesting talks and discussions covering a wide range of topics across artificial intelligence and the simulation of behaviour. This proceedings volume contains the papers from the *Symposium on Computational Creativity*, one of eight symposia held as part of the conference. Many thanks to the convention organisers, the AISB committee, convention delegates, and the many Kent staff and students whose hard work went into making this event a success.

—Colin Johnson, Convention Chair

Introduction to the Symposium

Over the last few decades, computational creativity has attracted an increasing number of researchers from both arts and science backgrounds. Philosophers, cognitive psychologists, computer scientists and artists have all contributed to and enriched the literature.

Many argue a machine is creative if it simulates or replicates human creativity (e.g. evaluation of AI systems via a Turing-style test), while others have conceived of computational creativity as an inherently different discipline, where computer generated (art)work should not be judged on the same terms, i.e. as being necessarily producible by a human artist, or having similar attributes, etc.

This symposium aimed at bringing together researchers to discuss recent technical and philosophical developments in the field, and the impact of this research on the future of our relationship with computers and the way we perceive them: at the individual level where we interact with the machines, the social level where we interact with each other via computers, or even with machines interacting with each other.

This 2nd International Symposium on Computational Creativity (CC2015) featured a number of presentations covering a range of topics in the evolving field of Computational Creativity. Issues addressed will include practical work in the area, theoretical approaches to creativity, and philosophical questions raised on the potential of non-human "creative" agents.

Topics of interest for this symposium included, but were not limited to: novel systems and theories in computational creativity, in any domain (e.g. drawing and painting, music, story telling, poetry, games, etc); the evaluation of computational creative systems, processes and artefacts; theory of computational aesthetics; representational issues in creativity, including visual and perceptual representations; social aspects of computational creativity, and intellectual property issues; creative autonomy and constraint; computational appreciation of artefacts, including human artworks.

We would like to thank all the members of the Programme Committee for the generous support and excellent work in evaluating the submissions.

—Mohammad Majid al-Rifaie, Jeremy Gow

Organising Committee

- Mohammad Majid al-Rifaie (Goldsmiths, University of London, UK)
- Jeremy Gow, (Goldsmiths, University of London, UK)
- Stephen McGregor (Queen Mary, University of London, UK) – Publicity

Programme Committee

- Mark Bishop (Goldsmiths, University of London, UK)
- Simon Colton (Goldsmiths, University of London, UK)
- Mark d'Inverno (Goldsmiths, University of London, UK)
- Pablo Gervás (Universidad Complutense de Madrid, Spain)
- Bipin Indurkhya (AGH University of Science and Technology, Kraków, Poland)
- Mohammad Ali Javaheri Javid (Goldsmiths, University of London, UK)
- Anna Jordanous (Kent University, UK)
- Francois Pachet (SONY Computer Science Laboratory Paris, France)
- Alison Pease (University of Dundee, UK)
- Georgi Stojanov (American University of Paris, France)
- Dan Ventura (Brigham Young University, USA)
- Geraint Wiggins (Queen Mary, University of London, UK)

Contents

Pablo Gervás, Tightening the Constraints on Form and Content for an Existing Computer Poet	1
Mohammad Ali Javaheri Javid, Mohammad Majid al-Rifaie and Robert Zimmer, An Informational Model for Cellular Automata Aesthetic Measure	9
Anna Jordanous, Four PPPPerspectives on Computational Creativity	16
Stephen McGregor and Mariano Mora McGinity and Sascha Griffiths, How Many Robots Does It Take? Creativity, Robots and Multi-Agent Systems	23
David C. Moffat, The Creativity of Computers at Play	30
Jiří Wiedermann and Jan van Leeuwen, Towards a Computational Theory of Epistemic Creativity	35

Tightening the Constraints on Form and Content for an Existing Computer Poet

Pablo Gervás [1]

Abstract. Existing systems for the automated generation of poetry often attempt to simplify the task by taking advantage of free-form poetry - to avoid the need to achieve rigorous poetic form - and poetic licence - to avoid the need of conveying a specific message at a semantic level. This is acceptable as an initial step, but once acceptable solutions have been found for the simplified version of the problem, progress can be made towards higher goals by enriching the initial problem statement. The present paper describes an attempt where an existing computer poet, originally developed to produce poems in a given form but with no specific constraints on their content, is put to the task of producing a set of poems with tighter restrictions on both form and content. Alternative generation methods are devised to overcome the difficulties, and the various insights arising from these new methods and the impact they have on the set of resulting poems are discussed in terms of their potential contribution to better poetry generation systems.

1 Introduction

Computer generation of poetry is a flourishing area of research in the context of computational creativity. In the last few years there has been a significant increase in the number of approaches to the task, and an extension to work in languages previously untried. By its nature, the task of generating a poem, when addressed by either a computer or a human, has to satisfy constraints at two very different levels. One level concerns the sequence in which the words appear in the poem. For a draft to be acceptable there has to be some way in which the words in it appear to link to one another, to make sense as a linguistic message. This constraint is applicable to the whole poem but essentially it operates at a local level, based on how each word can be seen to follow on from the previous one. A different level concerns certain macro-structural features that may be desirable in a poem, such as being distributed over a number of lines of specific lengths in terms of syllables, or having rhyming words occur at the end of particular lines. This corresponds to the poem satisfying some form of poetic stanza.

The problem of poetry generation is in fact rather more complex, because these two levels of constraints are just formulations of the overall specification at the extremes of a continuum. In truth, the way in which the sequence of words builds up is also expected to satisfy constraints on form - usually based on the relative positions of stressed syllables within a line, sometimes expressed in terms of feet - and there must also be some sense to be made between the different parts of the poem at a linguistic level. This is why human quality poetry is a tall order that few computer programs can tackle to the satisfaction of their critics. However, two higher level characteristics of poetry can be exploited to simplify the problem from an engineering point of view. First, poetry can also exist in free form, where constraints on line length, stress patterns or rhyme may be waived in favour of a more expressive poem at a semantic level. Second, the concept of poetic licence allows poets to sometimes violate linguistic expectations in favour of a more pleasing poem in terms of form. Traditionally, these two characteristics are applied in opposition to one another: if free-form is chosen for a poem, it is usually so that its linguistic expression does not have to be forced in any way to express the poet's meaning; if poetic licence is applied, it is usually to fit the poet's meaning into a particular poetic form where conventional phrasings might not work. Computer generated poetry often operates at the confluence of these two approaches relying on one to avoid the need to achieve rigorous poetic form and on the other to avoid the need of conveying a specific message at a semantic level. As the full problem is so complex, it is acceptable to apply a certain degree of simplification so that progress can be made in spite of the difficulties. However, the original goal must be kept in mind, so that once acceptable solutions have been found for the simplified version of the problem, progress can be made towards it by enriching the initial problem statement.

The present paper describes such an attempt. An existing computer poet, originally developed to take advantage of the characteristics of poetry described above, is set to the task of producing a set of poems with tighter restrictions on both form and content. The approach previously followed to poetry generation is shown to have limitations when the task is rephrased in this way. These limitations are analysed in terms of the current theoretical descriptions of computational creativity, and alternative generation methods are explored.

2 Previous Work

The work presented in this paper brings together some of the existing theoretical accounts of computational creativity and a number of efforts for computer generation of poetry. Both of these separate topics are reviewed in the present section.

2.1 Computational Creativity

Much of the work done on computational creativity over the past few years has been informed by Margaret Boden's seminal work describing creativity in terms of search over a conceptual space [3]. Boden formulated the search of ideas in terms of search over a conceptual space. Such a conceptual space would be defined by a set of constructive rules. The strategies for traversing this conceptual space in search of ideas would also be encoded as a set of rules. This view of

[1] Instituto de Tecnología del Conocimiento, Universidad Complutense de Madrid, email: pgervas@ucm.es

computational creativity was taken a step further in [25] by specifying formally the different elements involved (the universe of possible concepts, the rules that define a particular subset of that universe as a conceptual space, the rules for traversing that conceptual space, and a function for evaluating points in the conceptual space reached by these means).

In his pioneering work on the evaluation the creativity of computer programs, Ritchie [20] outlined a set of empirical criteria to measure the creativity of the program in terms of its output. Ritchie's criteria are defined in terms of two observable properties of the results produced by the program: novelty (to what extent is the produced item dissimilar to existing examples of that genre) and quality (to what extent is the produced item a high-quality example of that genre). He also put forward the concept of *inspiring set*, the set of (usually highly valued) artefacts that the programmer is guided by when designing a creative program. Ritchie's criteria are phrased in terms of: what proportion of the results rates well according to each rating scheme, ratios between various subsets of the result (defined in terms of their ratings), and whether the elements in these sets were already present or not in the inspiring set.

This idea of the inspiring set was taken a step further in [10], where the issue of how systems might take their prior output into account when evaluating the novelty of subsequent artifacts. This lead to the introduction of the concept of a *dynamic inspiring set*, one where system outputs are progressively updated into the inspiring set so they can inform later generative processes.

Colton and Wiggins [6] introduced the term *curation coefficient* to identify the proportion of system results that an impartial observer of system output would be happy to present to third parties. When estimated for a system addressing creative tasks it provides a reasonable measure of how much of the merit of presented system output can be attributed to the system itself and how much to the person actually selecting which particular outputs to present.

2.2 Computer Generated Poetry

Computer generation of poetry has traditionally addressed the constraints outlined in section 1 in terms of two different strategies: one is to reuse large fragments of text already formatted into poem-like structures of lines [7], and the other is to generate a stream of text by some procedure that ensures word-to-word continuity and then establish a distribution of the resulting text into lines by some additional procedure.

The reuse of text fragments already distributed into poetic lines was poineered by [17, 16] and it has more recently been used by [23, 12, 5, 24, 21, 4, 19]. In all these cases, either lines or larger poem fragments from exiting poems are subjected to modifications - usually replacement of some of the words with new ones - to produce new poems. A refinement on this method the selected fragment is stripped down to a skeleton consisting only of the POS tags of each line, and words corresponding to the desired content are used to fill this skeleton in. This procedure is followed in [8, 1, 22].

Alternative procedures rely on building a stream of text from scratch, and resort to various techniques to ensure the continuity of the textual sequence. One early approach was to rely on linguistic grammars to drive the construction. This was the approach followed in [13, 14], where TAG grammars were employed. A more popular alternative is the use of n-grams to model the probability of certain words following on from others. This corresponds to reusing fragments of the corpus of size n, and combining them into larger fragments based on the probability of the resulting sequence. This is the main approach for ensuring text coherence used in [2, 11, 9, 7]. All these different computer poets rely on various additional methods for establishing constraints on the resulting poem drafts.

To ensure that resulting poems satisfy constraints on poem structure in terms of lines, systems that build a stream of text from scratch rely on either building each line separately [7] or applying a separate procedure for distributing the resulting text into poetic lines [11, 9].

2.3 The WASP System

The development described in this paper was carried out over an existing version of the WASP system [11, 9].

Combining ngram modelling and evolutionary approaches, the WASP poetry generator had been built using an evolutionary approach to model a poet's ability to iterate over a draft applying successive modifications in search of a best fit, and the ability to measure metric forms. It operates as a set of families of automatic experts: one family of content generators or *babblers* – which generate a flow of text that is taken as a starting point by the poets –, one family of *poets* – which try to convert flows of text into poems in given strophic forms –, one family of *judges* – which evaluate different aspects that are considered important –, and one family of *revisers* – which apply modifications to the drafts they receive, each one oriented to correct a type of problem, or to modify the draft in a specific way. These families work in a coordinated manner like a cooperative society of readers/critics/editors/writers. All together they generate a population of drafts over which they all operate, modifying it and pruning it in an evolutionary manner over a number of generations of drafts, until a final version, the best valued effort of the lot, is chosen. In this version, the overall style of the resulting poems is strongly determined by the accumulated sources used to train the content generators, which are mostly n-gram based. Several versions have been developed, covering poetry generation from different inspirational sources as different sets of training corpora are used: from a collection of classic Spanish poems [11] and a collection of news paper articles mined from the online edition of a Spanish daily newspaper [9]. Readers interested in a full description are referred to the relevant papers. However, two specific aspects of this implementation are relevant for the present paper. First, the various judges assign scores on specific parameters – on poem length, on verse length, on rhyme, on stress patterns of each line, on similarity to the sources, fitness against particular strophic forms... – and an overall score for each draft is obtained by combining all individual scores received by the draft. A specific judge is in charge of penalising instances of excessive similarity with the sources, which then get pushed down in the ranking and tend not to emerge as final solutions. Second, poets operate mainly by deciding on the introduction of line breaks over the text they receive as input.

3 Can a Computer Poet Undertake a Commission for a Set of Themed Poems?

The work reported in this paper arose in response to a request received by the author to provide a set of poems generated by the WASP poetry system to be included in a book chapter about computational creativity. The request explicitly indicated that these poems should never have been published anywhere else, to avoid possible problems with copyright. Additionally, the author decided that the poems should aim to achieve a certain thematic unity, somehow relating to the circumstances in which they were commissioned. Finally, the author wanted to include data on the curation coefficient applicable, and to maximise its value to highlight the relative merit of the system itself in the achievement.

These conditions posed a challenge to the existing implementation of the WASP system. First, because the system as it stood had no means for driving the resulting poems towards particular themes. Second, because the procedures already in place for ensuring originality were inefficient. Third, because prior versions of the system had relied on low values of the curation coefficient: only a very small subset of actual system output was worthy of presentation to a wider audience.

The final set of poems was achieved by a recombination of some of the existing modules with new modules specifically designed for the ocasion, and by a new procedure for generating poems that abandoned the original generate and test approach underlying the evolutionary version of the system for a more informed generative approach that applied backtracking in search of solutions that better fulfilled the driving constraints.

3.1 Developing Text Babblers for the Themed Commission

As the book for which the poems were commissioned was to be published in Mexico, it was decided that the poems should have a Mexican theme. As the babbler modules rely on an ngram model of language to produce sequences of text that are word to word coherent, the overall style of the resulting poems is strongly determined by the accumulated sources used to train the content generators. For this initiative, a corpus of training texts was constructed by combining an anthology of poems by Mexican poets compiled from the Internet, and a set of news articles mined from the web pages of an online Mexican daily newspaper.

Earlier attempts to generate based on the simpler model trained only over the set of news items resulted in a candidate texts that were very difficult to adjust to any given poetic form. This related to the fact that the sequences of words contemplated in the ngram model resulting from news items only did not include enough combinations with a potential for poetic form. When the training set was expanded with an additional set of poetic texts, the resulting set of candidate texts showed a greater potential for composition into poetic forms.

This observation corroborates the intuition that the set of training texts used to train the ngram model imposes a certain overall style on the texts that can be produced. But it also raises the question of whether the desired poetic form is obtained at the price of replicating fragments of the poems being used as part of the inspiring set. This issue is addressed below.

3.2 Limitations of the Original Evolutionary Approach

The original WASP evolutionary system was designed to produce an initial large population of drafts - based on its ngram-based babbler modules -, to compose these into poem drafts by inserting line breaks at appropriate places - relying on its poet modules -, and to select as output a quality subset from those candidate drafts by applying the fitness functions implemented in its judge modules. This procedure was effective because it allowed the system to zoom in towards the regions of the overall conceptual space - as defined by the ngram model of language being used - that held potentially valuable text fragments from the point of view of poetic form - as defined by the fitness functions. This procedure was reasonable when the only constraint on the result was that it satisfy a certain poetic form. Specific poet modules and fitness functions would be designed for the particular poetic form, say, for a *cuarteto*, and the system would explore all the possible poems of this form arising from the given ngram model. This approach had two disadvantages for the present initiative: one related to form and one related to theme.

The existing solution was devised to drive the system towards poems of a particular type. When giving priority to theme, a certain flexibility in form could be introduced, allowing for poems with different poetic forms as long as they were consistent with the theme. To achieve this in terms of an evolutionary approach required the development of a confusing set of composition modules - capable of generating drafts in several poetic forms - and complex fitness functions - allowing for different fitness according to which particular poetic form was being considered. This lead to the consideration of alternative implementations.

The existing solution also had no obvious way of constraining results to particular themes. The word content of the results is constrained by the ngram model used, but an ngram model small enough to ensure that particular themes are present in the result would be too small to allow sufficient word recombinations to achieve valuable poetic forms. Additional elements could be added to the fitness function to rule out candidate drafts diverging from the desired themes, but this solution clashed with the decision above to consider alternative implementations.

3.3 Redeploying WASP Modules with a Different Purpose

A first attempt was carried out to simply redeploy the existing WASP modules - babblers, poets and judges - with the new purpose in mind. Under the new circumstances, judgements on candidate drafts could become more radical: if drafts were not related to the desired theme, they could be ruled outright. This had another consequence on the overall design: the reviser modules, which allowed exploration of the conceptual space by replacing certain words with others at random were seen to have little positive effect. Given the accumulated set of constraints on the results, random changes had a high probability of reducing fitness rather than improving it.

A formative evaluation was carried out over the existing prototype, configured so that a very large population of drafts was built, composed into a number of possible poetic forms, and evaluated using judges that combined fitness functions for theme, the various poetic forms considered, and originality. The revision modules were switched off for this test.

Fitness functions for theme relied on a set of input words to characterise the desired theme, penalising the drafts that did not include any of them, and reinforcing the drafts that did.

Fitness functions for poetic forms were already available as judge modules, and a simple combination of judges for different poetic forms was employed.

The fitness function for originality was addressed by developing a specific judge module that held the complete set of texts in the training set as a master file. Every line appearing in a candidate draft was searched for in the master file, and the candidate draft was rejected if the particular sequence of words in any of its lines appeared as a continuous unit anywere in the master file. This ensured that only lines that combined elements from different parts of the training set in innovative ways were considered by the system.

This approach generated a very large set of results but with very low average quality. This might have been acceptable if the set of results was mined for valuable drafts, but this would imply a very low curation coefficient for the final set.

3.4 Revising the Constructive Procedure to Match the New Circumstances

It was clear from the experiment described above that at least two improvements were required to fulfill the goals we had set out to fulfil. One was to improve the fitness functions overall so that only results of a higher quality survived the evaluation stage. Another was to somehow improve the construction procedure itself so that better quality results were produced. The evolutionary paradigm of the original approach required mostly random procedures for generation and revision, with quality to be achieved by means of evolutionary operators combined with selection in terms of the fitness function. But this approach clashes with the fact that the conceptual space that we want to explore is constrained to the set of texts that can be derived from the ngram model under consideration. For the evolutionary operators to guarantee that mutation and cross over produce results that are still within the desired conceptual space, they would have to be restricted to operations that take into account the ngram model during mutation and/or cross over.

The option of refining the revisers by enriching them with knowledge so that the changes they introduced were more informed was seen as impractical, and it was preferred to overhaul completely the generation procedure so as to take advantage of the available information to only generate valuable results in the first place.

The revised version of the construction procedure expanded the initial solution for babblers, which was based on extending a candidate sequence of words with further words that have a non-zero probability of appearing after the last word of the sequence, according to the ngram model. In both versions, at each choice point, the system is faced with a number of possible continuations. In the earlier version, this choice was taken randomly. In the new version, the choice is made taking into account additional criteria, covering the following issues: relation to theme, plausibility of sentence ending, control over repetition of sentences already generated, and restriction to overall length of sentences.

The first criterion to consider involves the initial constraints on theme, giving preference to options related to the desired theme.

The second criterion is designed to rule out cases where a draft is ended at a point where the word sequence under consideration does not allow the ending of the sentence.

The third criterion aims to avoid having the system repeat itself. A model of short term memory for sentences has been added, so that continuations of sentence drafts that replicate sentences constructed recently are avoided.

The final criterion ensures that text candidates are restricted to single sentences, and the overall length is restricted by introducing a check on the accumulated length of the word sequence that starts giving priority to continuations that close off the sentence after a given threshold length has been achieved.

The set of judges is revised so that drafts in any one of the following situations are ruled out directly:

- candidate drafts with line lengths beyond 14
- candidates drafts that have lines of different lengths

Additional judges have been developed that reinforce drafts were a certain pattern of rhyme can be spotted:

The procedure for composing candidate texts into valid poetic forms is revised in the following way. For any given candidate text the poetic composition module:

- finds the set of line lengths that have a potential to give an exact break down of the total number of syllables in the text
- composes a number of candidate draft poems based on the input text, each one distributing the text into lines of the corresponding length as worked out above
- evaluates the resulting set of poem drafts
- returns only those that are positively evaluate in terms of the judges for metric form

The described adaptations result in an exploratory software that takes a long time to run - as it explores exhaustively the portions of the conceptual space established by the given ngram model that include words from the desired theme - and produces a much smaller set of candidate drafts. These candidate drafts are of high quality in terms of poetic form - they correspond to stanzas of lines of the same length in syllables - but are surprisingly short in length - they very rarely exceed two lines. This restriction on length is a result of the interplay between the configuration that limits texts to single sentences and the restriction that the system start trying to close sentence as soon as a minimally valid length has been reached. In spite of the fact that judges have been included to prioritise poem drafts that exhibit rhyme, the set of results very rarely does.

This set of results is not in itself a convincing set of poems with which to satisfy the received commission. But it constitutes a treasure trove of valuable material generated by the system: it is by construction innovative - in terms of p-creativity as described by Boden, given that the originality judges check each line against the master file built from the training corpus and rule out any replications - and it is remarkable in its poetic form - as guaranteed by the remaining judges. It is a small set, but large enough to allow a further step of recombination of these poem snippets with one another.

The construction procedure was therefore extended with a further stage that considered these poem drafts as possible ingredients to combine into larger poems. The heuristics considered to drive this recombination process were as follows:

- the set of poem snippets was classified into groups according to the length of their lines in syllables
- poem snippets of the same length of line were further grouped together into sets related by shared rhymes
- larger poem drafts were built by combining together the sets of snippets of the same line length that had shared rhymes

The initial set of small poem drafts was produced in 6 separate runs with the same configuration, designed to carry out 1000 attempts to build poem drafts fulfilling the constraints as described above. The data on number of valid poem drafts found in each of these runs is presented in Table 1. Runs 2, 3 and 6 had to be aborted without finishing for practical reasons unrelated to system operation.

Run #	Valid drafts found
1	149
2	46
3	106
4	150
5	8
6	10

Table 1. Rates of success in the runs for collecting an initial set of poem snippets.

The average rate of success over this limited set of data - excluding the data for aborted runs on the grounds that no record is available of the number of attempts they had carried out before being stopped - is 13.5 %. Given the complexity of the conceptual space that is being searched, this rate is considered very acceptable.

The total number of snippets obtained in this way that was used as input for the procedure for composing larger poem drafts was 469.

The procedure for recombining the generated poem snippets into larger poem drafts produced 42 poem drafts, as described in Table 2. Overall these poems have used 18 different rhymes, irregularly spread over the set of resulting poems. The numbers provided for the complete set of poems do not correspond to the addition of the specific values for different line lengths because poem lengths and rhyme schemes are sometimes repeated for different line lengths.

Line lengths	Poems	Poem lengths	Rhyme schemes
6	1	1	1
7	9	6	7
8	6	3	6
9	10	7	10
10	14	8	14
12	2	1	-
All	42	11	31

Table 2. Description in terms of line lengths of the set of poem drafts obtained by recombination of snippets

The poems that resulted from this process were of different size, and for each particular poem size a rhyme schema results from the way in which snippets sharing rhyming lines have been combined. The analysis of the resulting set in terms of these emerging stanzas and rhyme schemes is presented in Tables 3 and 4.

Stanza size	Rhyme schemes
10	ABABCAABDA
14	BAABACADAEFBAA
	ACABABADAEAFCA
15	EACAFGBADADAHCA
20	AEACABABFAGAADADACAC
21	BACBAADAAAEAFAGHAAIJA

Table 4. Description of the longer stanzas in the set of poem drafts obtained by recombination of snippets

Of the 42 poems generated, 13 poems were deemed to be unusable as a result of problems in the generation process. The type of problems that were identified included issues of incorrect scanning of line lengths due to the appearance of punctuation signs not covered by the parsing procedures (2), undesirable repetition of subsets of lines (5), occurrence of unknown words (4), inclusion of unacceptable rude words (2).

The issue with incorrect scanning of line lengths has now been corrected.

Repetition of fragments of poems of more than one line is discouraged. The ones appearing in the result set have been tracked down to a small bug in the recombination process that should be easy to fix.

Some of the unknown words appear because the corpus of news items is mined directly from the web and the pre-processing procedures applied to clean up the html code sometimes miss non-words that end up in the training set. Improvements on the clean up procedure already under way should avoid this problem in the future.

Another source of problematic words is the use of foreign languages proper names, also frequent in news items. These words are acceptable in terms of their semantics contribution but their spelling confuses the metric analysis module of the system, which computes an incorrect number of syllables for them. This in its turn affects the composition processes that convert the resulting text into poetic form.

Rude words seem to have been used in some of the news items in the corpus, or possibly in some of the poems. But they are not considered desirable for the commissioned set of poems.

Of the remaining 29 poems, 7 were selected to be included in the book chapter that gave rise to the commission. This selection was based on general quality, but also on how well the selected poems fitted the desired theme. The 22 poems that were not selected show acceptable quality, and they were excluded from the selection for one of the following reasons:

- they shared some lines with the poems already selected
- their relation to the desired theme was not clear
- they included mentions of entities too specific to Mexican current news to be easily identified by a general public
- they included proper names of individuals featuring in the Mexican news
- they were overlong

Example results of the poems produced in this way are presented in Tables 6, 7,8,9,10 and11. These examples correspond to a second stage of selection out of the 22 poems that had not been chosen for inclusion in the set of poems commissioned for the book chapter.

The poems presented in Tables 5, 6, 7 and 8 correspond to four-line poems of different number of syllables per line (7, 7, 8 and 9 respectively), and showing different rhyme schemes (BACA, ABCA, ABAC). Together they illustrate the ability of the system to find the most metrically appropriate form for presenting a given text, using different lengths of line in syllables as required. They also illustrate the ability of the system to operate with different rhyme schemes to make the most of a given text.

Toda era una ave larga	Everything was a long bird
que cuando se conforman.	that when they conform.
Admitió que se tienen	He admitted that they have
registradas personas.	registered persons.

Table 5. Example of a poem of 4 lines of 7 syllables with rhyme scheme BACA, with an approximate English translation.

Muestra también. Esta	Shows as well. This
noche adonde yo soy.	night were I am.
Subraya que para ellos	Underlines that for them
ya no salgas. Estrellas.	come out no more. Stars.

Table 6. Example of a poem of 4 lines of 7 syllables with rhyme scheme ABCA, with an approximate English translation.

Aspecto que se encontraban	Aspect that they were finding
ejemplares. Nuevamente.	exemplars. Again.
Señalaron que no haya	They pointed out that there should not be
más daños como los niños.	more harm like the children.

Table 7. Example of a poem of 4 lines of 8 syllables with rhyme scheme ABAC, with an approximate English translation.

Stanza size	4	5	6	7	8	9
Rhyme schemes	ABBA	BACDA	BAAAAB	CBADABA	BACADAAE	ABABCBBDB
	ABAB	ABCDA	ABABAB	BACDAAE		ABCDEAFBA
	AAAA		ABCABC			ABACDEAFA
	BACA		ABABCA			BACADAEAF
	ABAC		ABABAC			ABACDAAEF
	ABCA		BAACAA			
			ABCADA			
			ABCAAD			
			ABACAD			

Table 3. Description in terms of the shorter stanzas in the set of poem drafts obtained by recombination of snippets

Zona militar. Qué delicia	Military zone. What a delight
delgada incomprensible. Amiga.	thin incomprehensible. Girl friend.
Agueda era luto pupilas	Agueda was mourning green
verdes. Sobrepasa. Guerrero.	pupils. Overshoots. Warrior.

Table 8. Example of a poem of 4 lines of 9 syllables with rhyme scheme ABAC, with an approximate English translation.

The poem presented in Table 8 is made of 4 *eneasílabos* of 9 syllable lines. Lines 2 and 3 share an asonant rhyme in *i-a*. The restriction on early closure of sentences has produced here a certain *staccato* feeling that is in line with the topic being addressed. Serendipity has led to a marked contrast between "military" and "delight", followed up with a surprisingly appropriate "incomprehensible". In spite of the choppy phrasing, as "girl friend" and "delight" agree in gender in Spanish, there is an implicit thread to the first two lines that is quite evocative. The third line mentions the female proper name "Agueda", rounding up this impression. This is again serendipitous. But it poses the question of whether similar criteria might not be used to derive selection heuristics so that future versions of the system can attempt to achieve similar effects. The final word "warrior" is ambiguous, and may originally have been intended as a reference to the Mexican state of Guerrero, but also links up with the military theme.

Juegan el largo recorrido	They play the long tour
desde su muerte ya no salgas.	from his death come out no more.
Séptimo. Cordero tranquilo	Seventh. Peaceful lamb
cordero que paces tu grama.	lamb that grazes its grass.
Silencios. Cordero tranquilo	Silences. Peaceful lamb
cordero que paces tu grama.	lamb that grazes its grass.

Table 9. Example of a poem of 6 lines of 9 syllables with rhyme scheme ABABAB, with an approximate English translation.

The poem presented in Table 9 is composed of 6 *eneasílabos* of 9 syllable lines. Lines 1, 3 and 5 rhyme together, and so do lines 2, 4 and 6. The rhyming is poor because it basically involves some the line endings being repeated twice. However, this arises from a parallelism trope - same linguistic structure used repeatedly with slight variations of content - and this makes the repeated rhyme somewhat more acceptable. The repetition is serendipitous and arises from the fact that particular sequences of words that match well a given poetic form tend to be reused to fill in certain stanzas ("Cordero tranquilo // cordero que paces tu grama. "), relying on different fragments of similar length to cover the initial first few syllables ("Séptimo.", Silencios."). Remember the constituent snippets were originally built separately, and they are only combined by application of the described composition heuristics. The apparent rhetorical effect is a consequence of the interaction between the limitation in the constructive procedure for poem snippets and the composition heuristics. Having noticed this interaction, we hope to include it as a system feature in future releases. In this particular case, the sequence in which the different fragments appear also achieves a significant effect, with the neighbouring mention of "death" and "lamb" evoking a certain hint of Christian symbolism. The effect of the early closing policy for sentences is also apparent in this poem.

Engalanados por los derechos	Garlanded by the rights
del niño indígena. Apago soles.	of the indigenous child. I switch off suns.
Concluido el objetivo que exista	Having achieved the goal that it exists
todo el mes para que ya sin nombre.	the whole month so that now nameless.
Dichosa puerta que nos transforman.	Happy gate that they transform for us.
Solidaridad vocación. Hombres.	Solidarity vocation. Men.
Acción nacional tiene un enorme	National action has an enormous
pez que se ilumina. Guatemala.	fish that lights up. Guatemala.

Table 10. Example of a poem of 8 lines of 10 syllables with rhyme scheme BACADAAE, with an approximate English translation.

The poem presented in Table 10 is composed of 8 *decasílabos* of 10 syllable lines. Lines 2, 4, 6 and 7 rhyme together. It presents interesting features that arise from the fact that sentences in the news items corpus are not generally well suited for partition over several valid metric lines, which lead to them being cut off abruptly at points where the closure makes syntactic sense. The texts in the poetic part of the corpus perform better in this sense, possibly as a result of being composed with metric form in mind. Drafts where the system alternates fragments from the two different parts of the corpus tend to achieve greater sentence lengths, as well as interesting contrasts between day to day pragmatic topics arising from news items and grander and more abstract topics obtained from the poems in the corpus. The Mexican theme is hinted at by the mention of the indigenous child.

The poem presented in Table 11 is included as an example of a longer poem. It has 15 *heptasílabos* or 7 syllable lines. Lines 2, 3, 5, 7, 9, 13, and 14 share asonant rhyme in *a-o* and lines 4 and 12 share asonant rhyme in *a-a*. This results in a rhyme scheme of the form CAABADAEAFGBAA. The Mexican theme is apparent in the mentions of citizens of two different Mexican states ("michoacanas", women from Michoacán; and "Queretanos", men from the town of Querétaro).

Tus ojos. Vinos tintos	Your eyes. Red wines
blancos rosados. Nardo.	whites rosés. Tube rose
Amo tus ríos claros.	I love your clear rivers.
Tal vez esta medalla.	Maybe this medal.
Antes que este hachazo	Before this axblow
nos sacude. Imagínate.	shakes us. Imagine.
Séptimo. Pinté el tallo	Seventh. I painted the stem
luego el cáliz después.	then the calyx afterwards.
Ganar. Solicitamos.	Winning. We request.
Sólo soy un prisionero.	I am only a prisoner.
Admitió que se estima	He admitted that it is estimated
que mil michoacanas	that a thousand Michoacans
acudan. Queretanos.	turn up. Queretans.
Valor. Acompañado	Valour. Accompanied
por Margarita Flores.	by Margarita Flores.

Table 11. Example of a poem of 15 lines of 9 syllables with rhyme scheme CAABADAEAFGBAA, with an approximate English translation.

4 Discussion

Over the complete run, 29 out of 42 poems were considered acceptable, and 13 of those have been submitted for publication in different media. The remaining 16 poems are less impressive but acceptable overall - they are not included here for lack of space -, although they do have the disadvantage of sharing some lines with the preferred poems. This, however, should not be considered as a demerit of the poems themselves. Instead, it should be thought of as an issue of incompatibility between possible system outputs in terms of originality. Once one particular line has been included in a poem submitted to the public, the system should refrain from including it in further output. This issue had already been described in [10], and more attention should be paid to it in poetry generators in the future.

These numbers lead to a curation coefficient for the described system run of around 69 %. It is important to note that higher curation coefficients are desirable. This is contrary to intuition, which suggests that high values of the curation coefficient imply a need for significant mediation between system output and publishable results. The contrary is in fact the case, as a high curation coefficient implies that a large percentage of system output can be passed to the public directly.

An interesting feature of the system described in this paper is that instead of establishing as configuration parameters values for features such as number of lines, number of syllables per line, or rhyme scheme to use, it relies on an exploratory procedure that allows the system to find optimal values for these features depending on the text that it has to convey. This leads to the variety of line and poem lengths, and the broad range of rhyme schemes that appear in the result set.

This variation in the range of rhyme schemes might be presented as an argument in favour of the perceived creativity of the system. The exploratory procedure in place relies on a fitness function that assigns higher value to poems that exhibit rhyming lines, but it does not prescribe any particular patterns for the rhymes. This results in output that satisfies rhyme schemes not traditionally used by human poets. This could be interpreted as a shortcoming, but it can also be considered as a creative feature.

The reliance on a corpus of training texts to produce candidate texts to compose into poetry introduces a number of dependencies between the particular training set chosen and the range of output text that can be generated. In the examples above this has been shown to lead to poems satisfying certain thematic constraints, not necessarily arising from explicit theme related constraints but simply as a result of having constrained the corpus to text somehow related to the theme. The issue of explicit constraining on theme needs to be explored further.

The influence of the training corpus has also been shown to affect the plasticity of the resulting texts when trying to compose them into poetic metrical forms. Certain styles of prose, such as that used in news items, are less conducive to composing into metrically acceptable forms than those custom-composed for such a form of expression. This should be taken into account when building training corpora for this type of system. On the other hand, the combination of corpus elements coming from different domains can lead to interesting contrasts that may result in a perception of originality in the final results.

One interesting point is the role of punctuation. As a result of the way the ngram models are constructed, most punctuation sign are stripped away from the texts before training. Question and exclamation marks are left in because they impact the syntax of the sentences they appear in. The output candidate texts are therefore generally devoid of punctuation. This introduces a degree of freedom that provides some leeway for human readers to find possible valid interpretation of the resulting poems. Readers should consider the possibility of revising the poems to consider whether simple punctuation, like the insertion of commas or semi-colons at certain points might improve them. It is after all, a task that editors of poetry sometimes do take out of the hands of their poets, even when they are human. In any case, having noticed the possible significance of this issue, the development of a system module to address such a refinement task is being considered as future work.

A final question to be considered is that of the originality of the output set in contrast to the inspiring set, here understood to correspond to the training corpus of texts. This question features prominently in Ritchie's set of criteria for evaluating the output of creative systems [20]. The system presented in this paper includes by construction a filter on candidate poem drafts that rejects them if they include a line that can be found as a continuous sequence anywhere within the training set. This should ensure that no line in any of the resulting poems correspond to lines in the poems in the training set, and it should also reduce significantly the chance that sentences in the training corpus are replicated verbatim.

Although many of the points outlined above deal with features that are specific to poetry, some of them can clearly be considered as valuable insight for computational creativity beyond poetry generation. First, the idea that creative systems should evolve towards versions where the role of a human observer curating a a subset of system outputs as valid for publication is reduced to a minimum. Second, the need to consider not only originality with respect to the inspiring set but also with respect to other elements in the result set. Third, the observation that, once the desired target is sufficiently specified, the introduction of randomness in the constructive procedure can have a negative impact. Fourth, that tightening the constraints on the desired target is likely to lead to increases in the time taken to produce results, and to decreases in the amount of results produced. However, the results obtained in this way are more likely to be of high quality. This point can be related to the stated view of Douglas Hofstadter that constraints are crucial to creativity.

5 Conclusions

The evolutionary solutions attempted in the past for poetry generation in the WASP system worked very well for the unconstrained exploration of broad conceptual spaces, where all parts of the space from a thematic point of view where equally valid as solutions, and constraints could be specified only in terms of metrical form. When constraints on theme are taken into consideration, it pays to relax the constraints on form, so that the system may look for the optimal poetical form covering a given theme. This has lead to the development of an exploratory procedure that sets its own values at run time for features such as poem length, line length, and rhyme scheme.

The refinement of the procedure for generating sentences to certain types of candidate - sentences of acceptable length and that can be understood as acceptably closed - had the consequence of restricting the possible outputs of the initial poem composition procedure to very short poem drafts. To compensate, a second stage of poem draft recombination has been added that builds larger poems from the set of initial candidate drafts. This recombination procedure is based on line length and shared rhymes, which leads to a result set that emulates reasonably well the composition of poems in terms of stanzas shaped together by rhyme.

The ratio of acceptable system outputs over total system outputs is reported, and it is argued to result in a very positive value of the curation coefficient.

The analysis of system outputs has lead to the identification of a number of positive features that have been included by serendipity, but which hold a very high potential for inclusion in future releases of the described system as quality-enhancing improvements. To handle these features might require an elaboration of the construction procedure as an interaction between a number of cooperating experts, in the way described in [15].

ACKNOWLEDGEMENTS

This paper has been partially supported by the projects ConCreTe 611733 and PROSECCO 600653 funded by the European Commission, Framework Program 7, the ICT theme, and the Future and Emerging Technologies FET program.

REFERENCES

[1] Manex Agirrezabal, Bertol Arrieta, Mans Hulden, and Aitzol Astigarraga, 'Pos-tag based poetry generation with wordnet', in *Workshop on Natural Language Generation (ACL 2013)*, (2013).

[2] Gabriele Barbieri, Franois Pachet, Pierre Roy, and Mirko Degli Esposti, 'Markov constraints for generating lyrics with style.', In Raedt et al. [18], pp. 115–120.

[3] M. Boden, *Creative Mind: Myths and Mechanisms*, Weidenfeld & Nicholson, London, 1990.

[4] John Charnley, Simon Colton, and Maria Teresa Llano, 'The flowr framework: Automated flowchart construction, optimisation and alteration for creative systems', in *5th International Conference on Computational Creativity, ICCC 2014*, Ljubljana, Slovenia, (06/2014 2014).

[5] Simon Colton, Jacob Goodwin, and Tony Veale, 'Full-FACE poetry generation', in *Proceedings of the International Conference on Computational Creativity 2012*, pp. 95–102, (2012).

[6] Simon Colton and Geraint A. Wiggins, 'Computational creativity: The final frontier?', In Raedt et al. [18], pp. 21–26.

[7] Amitava Das and Björn Gambäck, 'Poetic machine: Computational creativity for automatic poetry generation in bengali', in *5th International Conference on Computational Creativity, ICCC 2014*, Ljubljana, Slovenia, (06/2014 2014).

[8] P. Gervás, 'WASP: Evaluation of different strategies for the automatic generation of spanish verse', in *Proceedings of the AISB-00 Symposium on Creative & Cultural Aspects of AI*, pp. 93–100, (2000).

[9] P. Gervás, 'Evolutionary elaboration of daily news as a poetic stanza', in *Proceedings of the IX Congreso Español de Metaheurísticas, Algoritmos Evolutivos y Bioinspirados - MAEB 2013*, (2013).

[10] Pablo Gervás, 'Dynamic inspiring sets for sustained novelty in poetry generation', in *Second International Conference on Computational Creativity*, México City, México, (04/2011 2011).

[11] Pablo Gervás, 'Computational modelling of poetry generation', in *Proceedings of the AISB13 Symposium on Artificial Intelligence and Poetry*, (2013).

[12] Hugo Gonçalo Oliveira, 'PoeTryMe: a versatile platform for poetry generation', in *Proceedings of the ECAI 2012 Workshop on Computational Creativity, Concept Invention, and General Intelligence*, C3GI 2012, Montpellier, France, (August 2012).

[13] H. M. Manurung, 'Chart generation of rhythm-patterned text', in *Proc. of the First International Workshop on Literature in Cognition and Computers*, (1999).

[14] H. M. Manurung, *An evolutionary algorithm approach to poetry generation*, Ph.D. dissertation, University of Edimburgh, Edimburgh, UK, 2003.

[15] Joanna Misztal and Bipin Indurkhya, 'Poetry generation system with an emotional personality', in *5th International Conference on Computational Creativity, ICCC 2014*, Ljubljana, Slovenia, (06/2014 2014).

[16] Oulipo, *Atlas de littérature potentielle*, number vol. 1 in Collection Idées, Gallimard, 1981.

[17] R. Queneau, *100.000.000.000.000 de poèmes*, Gallimard Series, Schoenhof's Foreign Books, Incorporated, 1961.

[18] Luc De Raedt, Christian Bessière, Didier Dubois, Patrick Doherty, Paolo Frasconi, Fredrik Heintz, and Peter J. F. Lucas, eds. *ECAI 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31 , 2012*, volume 242 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2012.

[19] Fam Rashel and Ruli Manurung, 'Pemuisi: A constraint satisfaction-based generator of topical indonesian poetry', in *5th International Conference on Computational Creativity, ICCC 2014*, Ljubljana, Slovenia, (06/2014 2014).

[20] G. Ritchie, 'Some empirical criteria for attributing creativity to a computer program', *Minds & Machines*, **17**, 67–99, (2007).

[21] Jukka M. Toivanen, Oskar Gross, and Hannu Toivonen, 'The officer is taller than you, who race yourself! using document specific word associations in poetry generation', in *5th International Conference on Computational Creativity, ICCC 2014*, Ljubljana, Slovenia, (06/2014 2014).

[22] Jukka M. Toivanen, Matti Järvisalo, and Hannu Toivonen, 'Harnessing constraint programming for poetry composition', in *Proceedings of the International Conference on Computational Creativity 2013*, pp. 160–167, (2013).

[23] Jukka M. Toivanen, Hannu Toivonen, Alessandro Valitutti, and Oskar Gross, 'Corpus-based generation of content and form in poetry', in *Proceedings of the International Conference on Computational Creativity 2012*, pp. 175–179, (2012).

[24] Tony Veale, 'Less rhyme, more reason: Knowledge-based poetry generation with feeling, insight and wit', in *Proceedings of the International Conference on Computational Creativity 2013*, pp. 152–159, (2013).

[25] G. Wiggins, 'Searching for Computational Creativity', *New Generation Computing, Computational Paradigms and Computational Intelligence. Special Issue: Computational Creativity*, **24**(3), 209–222, (2006).

An Informational Model for Cellular Automata Aesthetic Measure

Mohammad Ali Javaheri Javid ✉ [1] and **Mohammad Majid al-Rifaie** [2] and **Robert Zimmer** [3]

Abstract. This paper addresses aesthetic problem in cellular automata, taking a quantitative approach for aesthetic evaluation. Although the Shannon's entropy is dominant in computational methods of aesthetics, it fails to discriminate accurately structurally different patterns in two-dimensions. We have adapted an informational measure to overcome the shortcomings of entropic measure by using information gain measure. This measure is customised to robustly quantify the complexity of multi-state cellular automata patterns. Experiments are set up with different initial configurations in a two-dimensional multi-state cellular whose corresponding structural measures at global level are analysed. Preliminary outcomes on the resulting automata are promising, as they suggest the possibility of predicting the structural characteristics, symmetry and orientation of cellular automata generated patterns.

1 INTRODUCTION

Cellular Automata (CA) initially invented by von Neumann in the late 1940s as material independent systems to investigate the possibility self-reproduction. His initial cellular automaton to study the possibility of self-reproduction was a two-dimensional (2D) cellular automaton with 29 states and 5-cell neighbourhood. A cellular automaton consists of a lattice of uniformly arranged finite state automata each of which taking input from the neighbouring automata; they in turn compute their next states by utilising a state transition function. A synchronous or asynchronous interactive application of state transition function (also known as a *rule*) over the states of automata (also referred to as *cells*) generates the global behaviour of a cellular automaton.

The formation of complex patterns from simple rules sometimes with high aesthetic quality has been contributed to the creation of many digital art works since the 1960s. The most notable works are "*Pixillation*", one of the early computer generated animations [32], the digital art works of Peter Struycken [31, 36], Paul Brown [5, 12] and evolutionary architecture of John Frazer [18]. Although classical one-dimensional CA with binary states can generate complex behaviours, experiments with 2D multi-state CA have shown that adding more states significantly increases the complexity of behaviour, therefore, generating very complex symmetrical patterns with high aesthetic qualities [21, 22]. These observations have led to the quest of developing a quantitative model to evaluate the aesthetic quality of multi-state CA patterns.

This work follows Birkhoff's tradition in studying mathematical bases of aesthetics, especially the association of aesthetic judgement with the degree of complexity of a stimulus. Shannon's information theory provided an objective measure of complexity. It led to emergence of various informational theories of aesthetics. However due to its nature, the entropic measure fails to take into account spacial characteristics of 2D patterns which is fundamental in addressing aesthetic problem for CA generated patterns.

[1] Goldsmiths, University of London, email: m.javaheri@gold.ac.uk
[2] Goldsmiths, University of London, email: m.majid@gold.ac.uk
[3] Goldsmiths, University of London, email: r.zimmer@gold.ac.uk

2 CELLULAR AUTOMATA ART

The property of CA that makes them particularly interesting to digital artists is their ability to produce interesting and logically deep patterns on the basis of very simply stated preconditions. Iterating the steps of a CA computation can produce fabulously rich output. The significance of CA approach in producing digital art was outlined by Wolfram in his classical studies on CA behaviours in [39]. Traditional scientific intuition, and early computer art, might lead one to assume that simple programs would always produce pictures too simple and rigid to be of artistic interest. But extrapolating from Wolfram's work on CA, "it becomes clear that even a program that may have extremely simple rules will often be able to generate pictures that have striking aesthetic qualities-sometimes reminiscent of nature, but often unlike anything ever seen before" [39, p.11].

Knowlton developed "*Explor*" system for generating 2D patterns, designs and pictures from explicitly provided 2D patterns, local operations and randomness. It aimed not only to provide the computer novice with graphic output; but also a vehicle for depicting results of simulations in natural (i.e. crystal growth) and hypothetical (e.g. cellular automata) situations, and for the production of a wide variety of designs [23]. Together with Schwartz and using *Explor*'s CA models, they generated "*Pixillation*", one of the early computer generated animations [32]. They contested in the *Eighth Annual Computer Art Contest* in 1970 with two entries, "*Tapestry I*" and "*Tapestry II*" (two frames from *Pixillation*). The "*Tapestry I*" won the first prize for "*new, creative use of the computer as an artist's tool*" as noted by selecting committee and covered the front page of *Computers & Automation* on Aug. 1970.

Meertens and Geurts also submitted an entry to the *Eighth Annual Computer Art Contest* with "*Crystalization*" as an experimental computer graphics generated by a asynchronous cellular automaton. Their entries were four drawings intended to generated patterns that combine regularity and irregularity in a natural way [20]. Peter Struycken, the Dutch contemporary digital artist has created many of his works "*Computer Structures*" (1969), "*Four Random Drawings for Lien and Ad*" (1972), "*Fields*" (1979-1980) with binary and multi-state CA [31, 36]. Paul Brown, the British contemporary digital artists also applied various CA rules in his static and kinematic computer arts. "*Neighbourhood Count*" (1991), "*Infinite Permutations V1*" (1993-94), "*Infinite Permutations V2*" (1994-95), "*Sand*

Lines" (1998), "*My Gasket*"(1998) "*Chromos*" (199-2000) [5, 12] are some of his CA generated works.

John F. Simon Jr created a series of art projects called "*Art Appliances*" using a CA based software and LCD panels to exhibit CA pattern formations. "*Every Icon*" (1996), "*ComplexCity*" (2000) and "*Automata Studies*" (2002) are examples of his art appliances. Driessen and Verstappen have produced "*Ima Traveler*" (1996) and "*Breed*"(1995-2007) digital arts in a three-dimensional CA space. Dorin's "*Meniscus*" [13] and McCormack's "*Eden*" [27] are further examples of interactive artworks built on the bases of CA rules. In addition, a combination of CA with other Alife techniques (e.g. evolutionary computing or L-systems) has been used to explore a set of rules generating patterns with aesthetic qualities [9, 34].

Fig. 1 shows some experimental patterns generated by the authors to demonstrate the generative capabilities of CA in creating appealing complex patterns from various initial configurations.

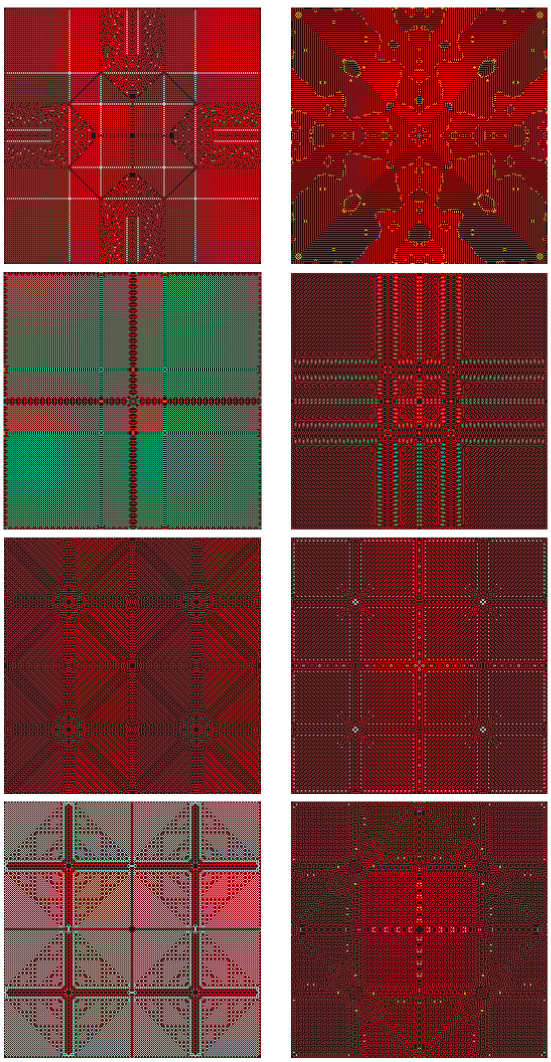

Figure 1. Sample 2D CA generated complex symmetrical patters

3 DEFINITION OF CELLULAR AUTOMATA

In this section, formal notions of 2D CA are explained and later referred to in the rest of the paper.

Definition 1: A cellular automaton is a regular tiling of a lattice with uniform deterministic finite state automata as a quadruple of $\mathcal{A} = \langle L, S, N, f \rangle$ such that:

1. L is an infinite regular *lattice* in $\mathbb{Z}$,
2. $S \subseteq \mathbb{N}^0$ is a finite set of integers as *states*,
3. $N \subseteq \mathbb{N}^+$ is a finite set of integers as *neighbourhood*,
4. $f : S^{|N|} \mapsto S$ is the *state transition function*.

The state transition function f maps from the set of neighbourhood states $S^{|N|}$ where $|N|$ is the cardinality of neighbourhood set, to the set of states $\{s_0, .., s_{n-1}\}$ synchronously in *discrete time* intervals of $t = \{0, 1, 2, 3, ..., n\}$ where t_0 is the *initial time* of a cellular automaton with *initial configuration*. A mapping that satisfies $f(s_0, ..., s_0) = s_0$ where $(s_0 \in S)$, is called a *quiescent state*.

In a 2D square lattice ($\mathbb{Z}^2$) if the opposite sides of the lattice (up and down with left and right) are connected, the resulting *finite* lattice forms a torus shape (Fig.2) which is referred as a lattice with *periodic boundary conditions*.

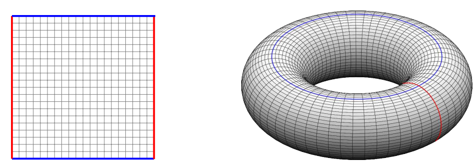

Figure 2. Connecting the opposite sides of a lattice forms a torus

The state of each cell at time $(t+1)$ is determined by the states of immediate surrounding neighbouring cells (nearest neighbourhood) at time (t) given a neighbourhood template. There are two commonly used neighbourhood templates considered for 2D CA. A five-cell mapping $f : S^5 \mapsto S$ known as *von Neumann neighbourhood* (Eq. 1) and a nine-cell mapping $f : S^9 \mapsto S$ known as *Moor neighbourhood* (Eq. 2).

$$s_{i,j}^{t+1} = f\begin{pmatrix} & s^t_{(i,j+1)} & \\ s^t_{(i-1,j)} & s^t_{(i,j)} & s^t_{(i+1,j)} \\ & s^t_{(i,j-1)} & \end{pmatrix} \quad (1)$$

$$s_{i,j}^{t+1} = f\begin{pmatrix} s^t_{(i-1,j+1)} & s^t_{(i,j+1)} & s^t_{(i+1,j+1)} \\ s^t_{(i-1,j)} & s^t_{(i,j)} & s^t_{(i+1,j)} \\ s^t_{(i-1,j-1)} & s^t_{(i,j-1)} & s^t_{(i+1,j-1)} \end{pmatrix} \quad (2)$$

Since the elements of the S are non-negative integers and discrete instances of time are considered, the resulting cellular automaton is a *discrete time-space* cellular automaton. These type of CA can be considered as *discrete dynamical systems*.

4 INFORMATIONAL AESTHETICS

The topic of determining aesthetics or aesthetic measures have been a heated debate for centuries. There is a great variety of computational approaches to aesthetics in visual and auditory forms including mathematical, communicative, structural, psychological and neuroscience. A thorough examination of these methodologies from different perspective has been provided in [19]. In this section, some informational aesthetic measures are presented. Our review is focused on informational theories of aesthetics as these are the ones that conform with this work directly.

Birkhoff suggested an early aesthetic measure by arguing that the measure of aesthetic (M) is in direct relation with the degree of *order* (O) and in reverse relation with the *complexity* (C) of an object [11]. Given that order and complexity are measurable parameters the aesthetic measure of (M) is:

$$M = \frac{O}{C} \tag{3}$$

Even though the validity of Birkhoff's approach to the relationship and definition of order and complexity has been challenged [38, 15, 16, 14], the notion of *complexity* and objective methods to quantify it remains a prominent parameter in aesthetic evaluation functions.

Shannon's introduction of *information theory* provided a mathematical model to measure the degree of uncertainty (entropy) associated with a random variable [33]. The entropy H of a discrete random variable X is a measure of the average amount of uncertainty associated with the value of X. So $H(X)$ as the entropy of X is:

$$H(X) = -\sum_{x \in \mathcal{X}} P(x) \log_2 P(x) \tag{4}$$

The definition of entropy for X has a logarithm in the base of 2 so the unit of measure of entropy is in $bits$.

Moles [28], Bense [7, 6, 8] and Arnheim [2, 3, 4] were pioneers of the application of Shannon's entropy to quantify order and complexity in Birkhoff's formula by adapting statistical measure of information in aesthetic objects. Berlyne used informational approach in his psychological experiments to determine humans perceptual curiosity of visual figures [10]. Bense argued that aesthetic objects are "vehicles of aesthetical information" where statistical information can quantify the aesthetical information of objects [7]. For Bense order is a process of artistic selection of elements from a determined repertoire of elements. The aesthetic measure (M_B) is a the relative redundancy (R) of the reduction of uncertainty because of selecting elements from a repertoire $(H_{max} - H)$ to the absolute redundancy (H_{Max}).

$$M_B = \frac{R}{H_{max}} = \frac{H_{max} - H}{H_{max}} \tag{5}$$

where H quantifies entropy of the selection process from a determined repertoire of elements in $bits$ and H_{max} is maximum entropy of predefined repertoire of elements [8]. His informational aesthetics has three basic assumptions. (1) Objects are material carriers of aesthetic state, and such aesthetic states are independent of subjective observers. (2) A particular kind of information is conveyed by the aesthetic state of the object (or process) as *aesthetic information* and (3) objective measure of aesthetic objects is in relation with degree of order and complexity in an object [29].

Herbert Franke put forward an *aesthetic perception* theory on the ground of *cybernetic aesthetics*. He made a distinction between the amount of information being stored and the rate of information flowing through a channel as *information flow* measured in $bits/sec$ [17]. His theory is based on psychological experiments which suggested that conscious working memory can not take more than 16 $bits/sec$ of visual information. Then he argued that artists should provide a flow of information of about 16 $bits/sec$ for works of art to be perceived as beautiful and harmonious.

Staudek in his multi criteria approach (informational and structural) as *exact aesthetics* to Birkhoff's measure applied information flow I' by defining it as a measure assessing principal information transmission qualities in time. He used 16 $bits/sec$ reference as channel capacity $C_r = 16\ bits/sec$ and a time reference of 8 seconds ($t_r = 8s$) to argue that artefacts with $I > 128\ bits$ will not fit into the conscious working memory for absorbing the whole aesthetic message [35].

Adapting Bense's informational aesthetics to different approaches of the concepts of order and complexity in an image in [30], three measures based on Kolmogorov complexity [25], Shannon entropy (for RGB channels) and Zurek's physical entropy [40] were introduced. Then the measures were are applied to analyse aesthetic values of several paintings (Mondrian, Pollock, and van Gogh). Machado and Cardoso [26] proposed a model based on Birkhoff's approach as the ratio of *image complexity* to *processing complexity* by arguing that images with high visual complexity, are processed easily so they have highest aesthetic value.

5 INFORMATION GAIN MODEL

Despite the domination of entropic measures to aesthetic evaluation functions, it has a major shortcoming in terms of reflecting structural characteristics of 2D patterns. Examples in Fig.3 illustrate this shortcoming by showing the calculations of entropy for 2D patterns with the same density but different structural regularities and complexities. Fig.3a is a uniformly distributed patterns (a highly ordered pattern), Fig.3b and Fig.3c are patterns with identical structures but in vertical and horizontal orientations. Fig.3d is randomly arranged pattern (a random pattern). As it is evident from the comparison of the patterns and their corresponding entropy value, all of the patterns have the same entropy value. This clearly demonstrates that Shannon's entropy fails to differentiate structural differences among these patterns. In the case of measuring complexity of CA generated patterns especially with multi-state structures, it would be problematic if only entropy used as a measure of complexity for the purpose of aesthetic evaluation.

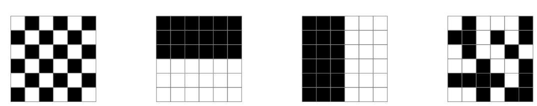

(a) $H = 1$ (b) $H = 1$ (c) $H = 1$ (d) $H = 1$

Figure 3. The measure of entropy H for structurally different patterns with the same density of 50%

In order to overcome this problem we have adapted *information gain* model introduced as a method of characterising the complexity of dynamical systems [37]. It has been applied to describe quantitatively the complexity of geometric ornaments and patterns arising in random sequential adsorption of discs on a plane [1]. The information gain G, also known as Kullback-Leibler divergence [24], measures the amount of information required to select a discrete random variable X with state j if prior information about variable X is known at the state of i.

$$G_{x_{ij}} = -\log P_{(x_i|x_j)} \tag{6}$$

where $P_{(x_i|x_j)}$ the conditional probability of the discrete random variable x at state i given its state j. Then from Eq. 6 *mean information gain* $\overline{G}$ would be the average information gain from possible states $(i|j)$:

$$\overline{G} = \sum_{i,j} P(i,j) G_{ij} = -\sum_{i,j} P_{i,j} \log P(i|j) \tag{7}$$

where $P_{(i,j)}$ is the joint probability of the variable x at state i and variable x at state j. Considering Eq. 7, we can define a structural complexity measure for a multi-state 2D cellular automaton as follows:

Definition 2: *A structural complexity measure is the mean information gain of a cell having a heterogeneous neighbouring cell in a multi-state two-dimensional cellular automaton pattern.*

$$\overline{G} = -\sum_{i,j} P_{(i,j)} \log_2 P_{(i|j)} \tag{8}$$

where $P_{(i,j)}$ is the joint probability of a cell having the i state (colour) and the neighbouring cell has the state (colour) j in a given neighbouring cell. And $P_{(i|j)}$ is the conditional probability of the state (colour) i given that its neighbouring cell has state (colour) j in one of four directions of up, low, left or right. The quantity $\overline{G}$ measures average information gain about other elements of the structure (e.g. the state of the neighbouring cell in one of the four directions), when some properties of the structure are known (e.g. the state of a cell). It can be noted that the combined probabilities of $P_{i,j}$ and $P_{i|j}$ describe spatial correlations in a pattern so that $\overline{G}$ can detect inherent correlations of patterns. Considering neighbourhood templates of a 2D CA in Eq.1 and Eq. 2, following variations of $\overline{G}$ can be defined where for each cell in i state given its neighbouring cell in j state in any of directions.

$$\overline{G}_u = -\sum_{i,j_{(x,y+1)}} P_{(i,j_{(x,y+1)})} \log_2 P_{(i|j_{(x,y+1)})} \tag{9}$$

$$\overline{G}_d = -\sum_{i,j_{(x,y-1)}} P_{(i,j_{(x,y-1)})} \log_2 P_{(i|j_{(x,y-1)})} \tag{10}$$

$$\overline{G}_l = -\sum_{i,j_{(x-1,y)}} P_{(i,j_{(x-1,y)})} \log_2 P_{(i|j_{(x-1,y)})} \tag{11}$$

$$\overline{G}_r = -\sum_{i,j_{(x+1,y)}} P_{(i,j_{(x+1,y)})} \log_2 P_{(i|j_{(x+1,y)})} \tag{12}$$

The measure is applied to calculate structural complexly of sample patterns in Fig 4 to demonstrates the ability of $\overline{G}$ in discriminating structurally different 2D patterns. The calculations have been performed for each elements of the patterns having a heterogamous colour in one of the four directions from two possible colours.

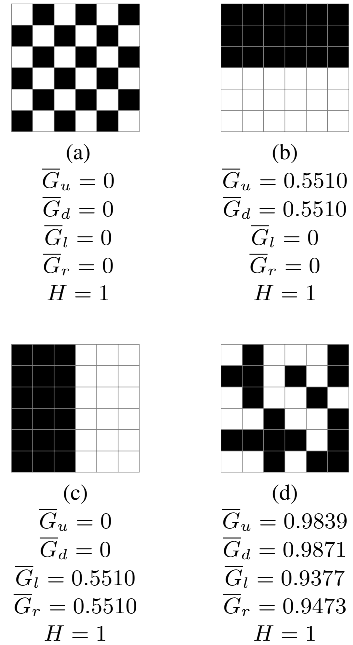

Figure 4. The comparison of entropy H and $\overline{G}$ for structurally different patterns but with the same density of 50%

6 EXPERIMENTS AND RESULTS

A set of experiments were designed to examine the effectiveness of $\overline{G}$ in discriminating structurally different patterns generated by multi-state 2D CA. The chosen experimental cellular automaton maps three states represented by *green*, *red* and *blue* colour cells. The quiescent state cells represented with *white* colours. The size of the lattice is set to 129×129 cells. Two set of experiments are conducted: (1) a single cell as initial configuration and (2) a randomly seeded initial configuration with 50% destiny of three states (*green*, *red*, *blue*). Both of the experiments are conducted for 300 successive time steps. The $\overline{G}$ for four directions along with their corresponding entropy H are measured in $bits$.

Fig. 7 and Fig. 8 illustrate the formation of 2D patterns for a sample of 12 time steps $\{0, 10, 20, 30, 40, 50, 60, 70, 80, 100, 200, 300\}$ starting from two different initial configurations and their corresponding $\overline{G}$ and H. Figs. 7 and 6 shows the plots of $\overline{G}$ and H for 300 time steps. The $\overline{G}$ measures in Fig. 7 which shows the formation of 2D patterns from a single cell are conforming in directional calculations; it means that each cell in the patterns has exactly the same amount of information regarding their neighbouring cell in one of the four directions. Therefore it indicates that the development of the patterns are symmetrical in the four directions. In other words, the cellular automaton with a single cell as its initial configuration has created 2D pattens with four-fold rotational symmetry. The measure in Fig. 8 starts with $\overline{G} \approx 1.7$ for the random initial configuration and with $H \approx 1.5$ (maximum entropy for a three-state patterns since $\log_2 3 = 1.5848$). The formation of patterns with local structures reduces the value of $\overline{G}$. The values of $\overline{G}$ are not conforming in any directional calculations which indicates the development of less ordered ("chaotic") patterns. From the comparison of H with $\overline{G}$ in the set of experiments, it is clear that it would be very unlikely to discriminate the structural differences of patterns with a single measure

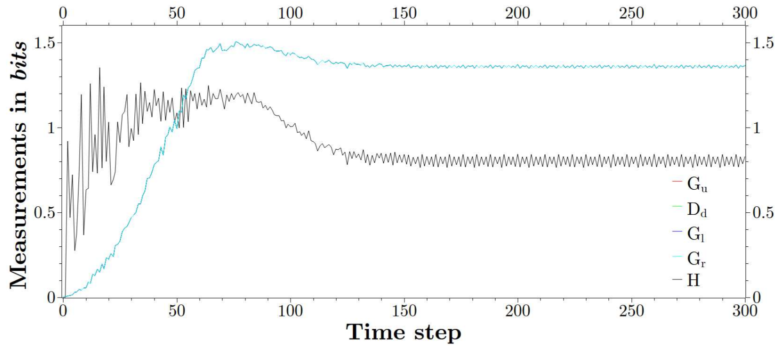

Figure 5. The plot of $\overline{G}$ and H for 300 time steps starting from single cell

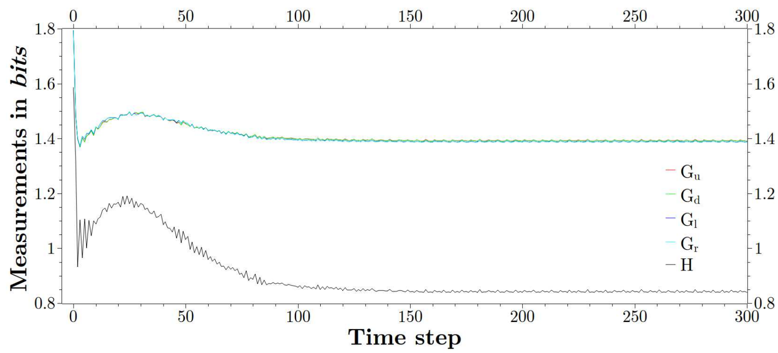

Figure 6. The plot of $\overline{G}$ and H for 300 time steps starting from random initial configuration

of H given the diversity of patterns that can be generated by various 2D CA state transition functions. Computing directional measures of $\overline{G}$ and comparing their values provides a more subtle measure of structural order and complexity of a 2D pattern. The conformity or non-conformity of $\overline{G}$ measure in up, down, left and right neighbouring cells clearly gives us not only an accurate measure of structural characteristics of 2D patterns but they also provide us with information about the orientation of the patterns as well.

7 CONCLUSION

Cellular automata (CA), which are fundamental to the study of self-replicating systems, are powerful tools in generating computer art. The multi-state 2D CA rule space is a vast set of possible rules which can generate interesting patterns with high aesthetic qualities. The application of CA in digital art has been reviewed; and the concepts of order and complexity from Shannon's information entropy perspective in the CA framework has been analysed concluding that existing informational aesthetic measures do not capture structural differences in 2D patterns. In order to address the shortcomings of informational approaches to computational aesthetics, a mean information gain model was adapted to measure both structural complexity and distinguish symmetrical orientation of 2D CA patterns. The measure takes into account conditional and joint probabilities of the information gain value that a cell offers, given a particular position of its neighbouring cells. The effectiveness of the measure is shown in a series of experiments for multi-state 2D patterns generated by a cellular automaton. The results of the experiments show that the mean information gain model is capable of distinguishing the structural complexity of 2D CA patterns as well as their symmetrical orientation. Having a model to evaluate the aesthetic qualities of CA generated patterns could potentially have a substantial contribution towards further automation of the evaluative component in the CA based computer generated art. This could also enable us to have an integrated process of generation-evaluation which is a subject of on going research.

8 ACKNOWLEDGEMENTS

We would like to offer our special thanks to anonymous reviewers for their constructive suggestions to improve the quality of this paper.

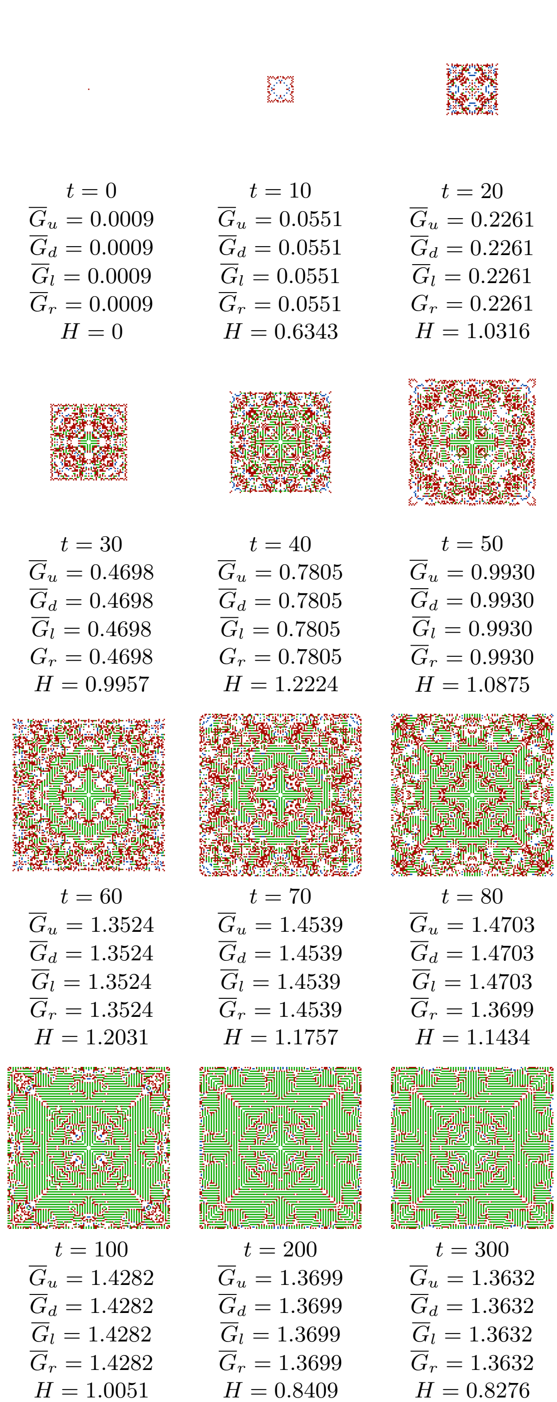

Figure 7. Patterns generated from a single cell as initial configuration and their corresponding $\overline{G}$ and H values

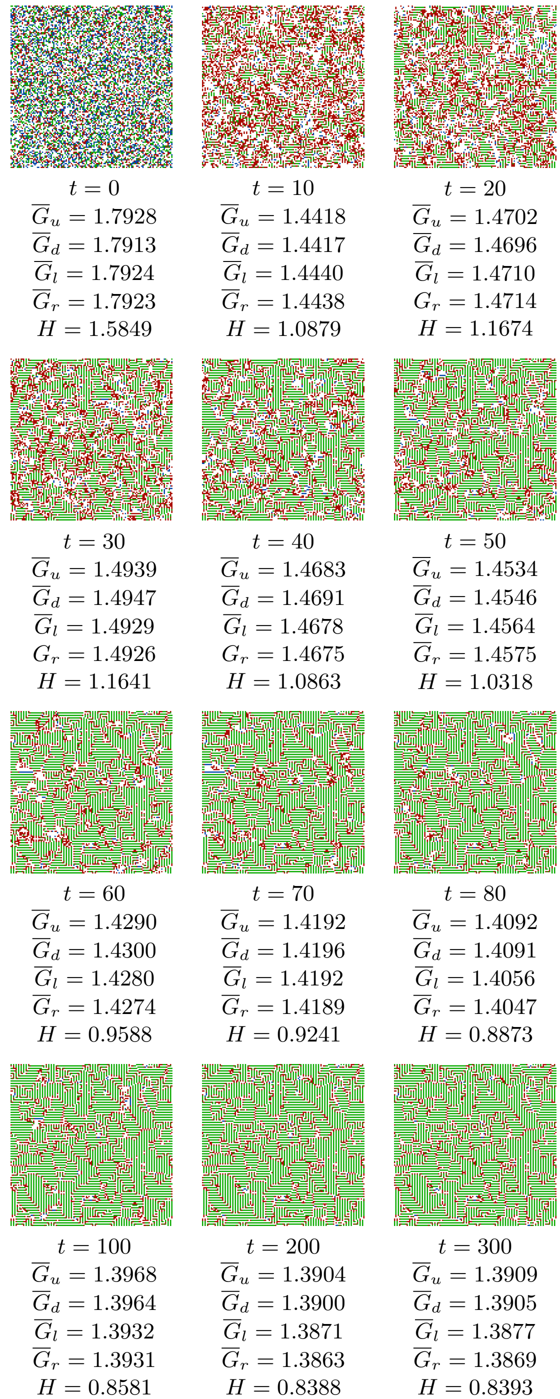

Figure 8. Patterns generated from a 50% seeded density as initial configuration and their corresponding $\overline{G}$ and H values

REFERENCES

[1] Andrienko, Yu. A., Brilliantov, N. V., and Kurths, J., 'Complexity of two-dimensional patterns', *Eur. Phys. J. B*, **15**(3), 539–546, (2000).

[2] Rudolf Arnheim, *Art and visual perception: A psychology of the creative eye*, Univ of California Press, 1954.

[3] Rudolf Arnheim, 'Towards a psychology of art/entropy and art an essay on disorder and order', *The Regents of the University of California*, (1966).

[4] Rudolf Arnheim, *Visual thinking*, Univ of California Press, 1969.

[5] Honor Beddard and Dodds Dodds, *Digital Pioneers*, V&A Pattern, V&A Publishing, 2009.

[6] M. Bense and G. Nee, 'Computer grafik', in *edition rot*, eds., Max Bense and Elisabeth Walther, volume 19, Walther, Stuttgart, (1965).

[7] Max Bense, *Aestetica: Programmierung des Schönen, allgemeine Texttheorie und Textästhetik...*, Agis-Verlag, 1960.

[8] Max Bense, *Kleine abstrakte ästhetik*, E. Walther, 1969.

[9] Katie A Bentley, 'Exploring aesthetic pattern formation', in *Generative Art 2002 conference proceedings*, (2002).

[10] Daniel Ellis Berlyne, 'Conflict and information-theory variables as determinants of human perceptual curiosity.', *Journal of experimental psychology*, **53**(6), 399, (1957).

[11] G.D. Birkhoff, *Aesthetic Measure*, Harvard University Press, 1933.

[12] Paul Brown, 'Stepping stones in the mist', in *Creative evolutionary systems*, pp. 387–407. Morgan Kaufmann Publishers Inc., (2001).

[13] Alan Dorin, 'The virtual ecosystem as generative electronic art', in *Applications of Evolutionary Computing*, 467–476, Springer, (2004).

[14] Hans J Eysenck, 'An experimental study of aesthetic preference for polygonal figures', *The Journal of General Psychology*, **79**(1), 3–17, (1968).

[15] Hans Jürgen Eysenck, 'The empirical determination of an aesthetic formula.', *Psychological Review*, **48**(1), 83, (1941).

[16] Hans Jurgen Eysenck, 'The experimental study of the 'good gestalt' –a new approach.', *Psychological Review*, **49**(4), 344, (1942).

[17] Herbert W. Franke, 'A cybernetic approach to aesthetics', *Leonardo*, **10**(3), 203–206, (1977).

[18] John Frazer, *An evolutionary architecture*, Architectural Association Publications, Themes VII, 1995.

[19] Philip Galanter, 'Computational aesthetic evaluation: past and future', in *Computers and Creativity*, eds., Jon McCormack and Mark d'IInverno, pp. 255–293. Springer, (2012).

[20] Leo Geurts and Lambert Meertens, 'Crystallization', *Computers and Automation*, **19**(8), 22, (1970).

[21] Mohammad Ali Javaheri Javid, Mohammad Majid al Rifaie, and Robert Zimmer, 'Detecting Symmetry in Cellular Automata Generated Patterns Using Swarm Intelligence', in *Theory and Practice of Natural Computing*, eds., Adrian-Horia Dediu, Manuel Lozano, and Carlos Martín-Vide, volume 8890 of *Lecture Notes in Computer Science*, pp. 83–94. Springer International Publishing, (2014).

[22] Mohammad Ali Javaheri Javid and Rene te Boekhorst, 'Cell Dormancy in Cellular Automata', in *International Conference on Computational Science (3)*, eds., Vassil N. Alexandrov, G. Dick van Albada, Peter M. A. Sloot, and Jack Dongarra, volume 3993 of *Lecture Notes in Computer Science*, pp. 367–374. Springer, (2006).

[23] Kenneth C Knowlton, 'Explor-a generator of images from explicit patterns, local operations, and randomness', in *Proceedings of 9th Meeting of UAIDE*, pp. 544–583, (1970).

[24] S. Kullback and R. A. Leibler, 'On Information and Sufficiency', *The Annals of Mathematical Statistics*, **22**(1), pp. 79–86, (1951).

[25] Ming Li, *An introduction to Kolmogorov complexity and its applications*, Springer, 1997.

[26] Penousal Machado and Amílcar Cardoso, 'Computing aesthetics', in *Advances in Artificial Intelligence*, 219–228, Springer, (1998).

[27] Jon McCormack, 'Evolving sonic ecosystems', *Kybernetes*, **32**(1/2), 184–202, (2003).

[28] Abraham Moles, *Information theory and esthetic perception. Trans. JE Cohen.*, U. Illinois Press, 1968.

[29] Frieder Nake, 'Information aesthetics: An heroic experiment', *Journal of Mathematics and the Arts*, **6**(2-3), 65–75, (2012).

[30] Jaume Rigau, Miquel Feixas, and Mateu Sbert, 'Conceptualizing birkhoff's aesthetic measure using shannon entropy and kolmogorov complexity', in *Workshop on Computational Aesthetics*, eds., Douglas W. Cunningham, Gary Meyer, and Laszlo Neumann, pp. 105–112, Banff, Alberta, Canada, (2007). Eurographics Association.

[31] Ir Remko Scha, 'Kunstmatige Kunst', *De Commectie*, **2**(1), 4–7, (2006).

[32] L.F. Schwartz and L.R. Schwartz, *The Computer Artist's Handbook: Concepts, Techniques, and Applications*, W W Norton & Company Incorporated, 1992.

[33] Claude Shannon, 'A mathematical theory of communication', *The Bell System Technical Journal*, **27**, 379–423 & 623–656, (October 1948).

[34] Karl Sims, 'Interactive evolution of equations for procedural models', *The Visual Computer*, **9**(8), 466–476, (1993).

[35] Tomáš Staudek, *Exact Aesthetics. Object and Scene to Message*, Ph.D. dissertation, Faculty of Informatics, Masaryk University of Brno, 2002.

[36] Peter Struycken, 'Splash 1972/1974', in *Artist and computer*, ed., Ruth Leavitt, 30–31, Harmony Books, (1976).

[37] Renate Wackerbauer, Annette Witt, Harald Atmanspacher, Jürgen Kurths, and Herbert Scheingraber, 'A comparative classification of complexity measures', *Chaos, Solitons & Fractals*, **4**(1), 133–173, (1994).

[38] D. J. Wilson, 'An experimental investigation of Birkhoff's aesthetic measure', *The Journal of Abnormal and Social Psychology*, **34**(3), 390, (July 1939).

[39] Stephen Wolfram, *A New Kind of Science*, Wolfram Media Inc., 2002.

[40] Wojciech H Zurek, 'Algorithmic randomness and physical entropy', *Physical Review A*, **40**(8), 4731, (1989).

Four PPPPerspectives on Computational Creativity

Anna Jordanous [1]

Abstract. From what perspective should creativity of a system be considered? Are we interested in the creativity of the system's output? The creativity of the system itself? Or of its creative processes? Creativity as measured by internal features or by external feedback? Traditionally within computational creativity the focus had been on the creativity of the system's Products or of its Processes, though this focus has widened recently regarding the role of the audience or the field surrounding the creative system. In the wider creativity research community a broader take is prevalent: the creative Person is considered as well as the environment or Press within which the creative entity operates in. Here we have the Four Ps of creativity: Person, Product, Process and Press. This paper presents the Four Ps, explaining each of the Four Ps in the context of creativity research and how it relates to computational creativity. To illustrate how useful the Four Ps can be in taking a fuller perspective on creativity, the concepts of novelty and value explored from each of the Four P perspectives, uncovering aspects that may otherwise be overlooked. This paper argues that the broader view of creativity afforded by the Four Ps is vital in guiding us towards more encompassing and comprehensive computational investigations of creativity.

1 Introduction

A practical issue arises when considering the evaluation of a computational creativity system: from what perspective should creativity of a system be considered? Are we interested in the creativity of the system's output? The creativity of the system itself? Or of its creative processes? Creativity as measured by internal features or by external feedback?

The computational creativity community has traditionally considered creativity from the perspective of the creative output produced by a system, or the processes employed within creative systems (with notable exceptions, such as Saunders [48]). The call for this ICCC 2014 conference invites papers addressing the 'Process vs. product: addressing the issue of evaluating/estimating creativity (or progress towards it) in computational systems through study of what they produce, what they do and combinations thereof.'

This paper argues that to consider process and product is not enough; computational creativity should be considered and explored from four different perspectives, known as the Four Ps: the creative Person, Product, Process and Press (or environment) [43, 26].

The Four Ps have long been prevalent in creativity research relating to humans[2] and enable a more inclusive and encompassing approach to the study of creativity and accommodating multiple relevant perspectives. Here the Four Ps are presented and considered in the light of how they are relevant to computational creativity researchers.

1.1 The product/process debate in computational creativity evaluation

> 'As a research community, we have largely focussed on assessment of creativity via assessment of the artefacts produced.' [8, p. 1]

As illustrated by the ICCC 2014 call for papers, one important debate in computational creativity is about whether evaluation of a creative system should focus exclusively on the output produced by the system, or whether the processes built into the system should also be taken into account. Should both product and process should be included in evaluation [39, 8, 20], or should evaluation concentrate solely on the product of systems [45]? Ritchie [45] stated that examining the process is unimportant for creativity, arguing that humans normally judge the creativity of others by what they produce, because one cannot easily observe the underlying process of human creativity. Ritchie therefore advocated a black-box testing approach, where the inner program workings are treated as unknown and evaluation concentrates on the system's results. Later, however, Ritchie [46] conceded that it can be important to consider a system's 'mechnisms' in the case of 'more theoretical research'[46, p. 147].

While it is true that we can only use the material we have available to form an evaluation, evaluation experiments [36, 19] show that people often make assumptions about process in their judgements on product. As Hofstadter pointedout, '*covert mechanisms* can be deeply probed and eventually revealed merely by means of watching *overt behaviour* ... [this approach] lies at the very heart of modern science.' [15, quoted in p. 10, [39]]. Pearce & Wiggins [36] discussed how our interpretation of how something was produced is important, even if the actual method is unknown, and that such an interpretation can be derived if people are repeatedly exposed to the compositional systems (human or computational) that they are evaluating. Collins [6] discussed how making reasonable assumptions can assist the reverse-engineering[3] of program code from output, in scenarios where white-box testing (evaluation with access to the program code) is not possible.

Colton [8] acknowledged Ritchie's arguments but quotes examples from art to demonstrate that process is as important as the end product when evaluating creativity, at least in the artistic domain. As evidence, Colton cites conceptual art for details on conceptual art in the context of this debate, where the concepts and motivations behind the artistic process are a significant contribution of the artwork. Sol LeWitt defined Conceptual Art [25] as an art form where 'the

[1] University of Kent, UK, email: a.k.jordanous@kent.ac.uk

[2] Variants of these *Ps* also arise in slightly different guises in non-related areas, such as software project management [16] or education [2].

[3] Reverse-engineering is the process of identifying and perhaps replicating how a product is made, through analysis of that product.

idea or concept is the most important aspect of the work. ... The idea becomes a machine that makes the art.' Two examples are Tracey Emin's controversial exhibit *My Bed* (1999) and Duchamp's Fountain (1917). Jordanous [20] makes similar arguments for creativity in musical improvisation, finding that the process of improvisation is often seen as more relevant for creativity than the end result.

If assessing how creative a piece of conceptual art or a musical improvisation is, solely by evaluating the product, then there are two negative consequences:

1. The primary intentions of the artist/musician are ignored (their focus is on how the creative work is made rather than the end result).
2. The level of creativity presented will probably be underestimated, especially if the creative process results in producing something that might seem commonplace outside the context of that art installation/musical performance.

Colton [8] also posed a thought experiment that considers two near-identical paintings presented at an exhibition. In the first painting, the dots are placed randomly, whereas in the second, the dots' locations represent the artist's friendships with various people. Colton argued that the second painting would be more appealing to purchase than the first, though the end product is very similar, due to the process by which it was created. Colton's thought experiment illustrates how process can impact on our judgement of creative artefacts, though one could question if the experiment explores perception of creativity, or of quality/appeal.

The thought experiment described by Ventura [54] gives further evidence (perhaps unintentionally) on how knowledge of the creative process affects how we evaluate creativity. Two creative systems, the RASTER and iRASTER systems, were designed by Ventura to be decidedly non-creative. If these systems were implemented and their generated images were given to people to evaluate without telling the evaluators how they were produced, the evaluators may well rate the creativity of the system highly. Supplying the evaluators with details of how a program works, though, could have a detrimental impact on the subsequent evaluations [11, 8].

One issue with creativity is analogous to the adage that a magician never reveals their secrets. This adage is based on the fact that tricks do not appear so impressive once you have found out how the magician performed the trick. Similarly things can appear to be less creative when you know how they were produced:[4]

> 'it is not unknown for critics of AI to refuse to accept programs as creative (or intelligent) once the mundane mechanistic nature of the inner workings are revealed' [44, p. 4]

Colton [8] intentionally sidestepped this issue by reporting on his artistic system in high-level terms only, rather than giving details of the program [8, p. 8].

Until recently, computational creativity evaluation methodologies mainly looked solely at a system's *products* [45, for example] or at a combination of the *products* and the *process* [39]. Recently it has been acknowledged that there is more to creativity than process and product, with the Creative Tripod [8], whose evaluative framework is influenced by how an audience perceives the creativity of a system, SPECS [20] which requires the researcher to investigate what creativity means in the context of their system, and the FACE/IDEA models [9] which consider various aesthetic features and interactions between audience and system. Work on computationally creative societies has also developed in the last few years [48, is a significant example].

Along a similar broadening of perspectives, the next section brings in work from the wider creativity research community, examining further viewpoints - the creative *person* operating in a *press*/environment - and relating these viewpoints to a computational creativity standpoint.

2 The *Four Ps* of creativity

One major approach in creativity research is to break down creativity into four perspectives, commonly referred to as the *Four Ps* [43, 51, 34, 26, 49, 53, 35]:

- Person: The individual that is creative.
- Process: What the creative individual does to be creative.
- Product: What is produced as a result of the creative process.
- Press: The environment in which the creativity is situated.

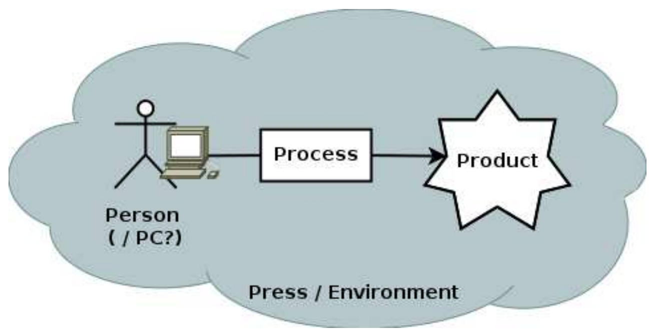

Figure 1. A simplified view of how the Four Ps fit together in creativity

Rhodes [43] was perhaps first to identify the four P perspectives. Rhodes collected 40 definitions of creativity and 16 definitions of imagination. The '*Four P*' dimensions of creativity emerged from analysis of these definitions.[5] Several people seem to have independently identified four similar themes of creativity [26, 51, 34, 35], boosting the credibility of the Four Ps.

Plucker, Beghetto & Dow [41] conducted a literature survey investigating the use (or absence) of creativity definitions in creativity research. As part of this review, Plucker et al. used their analysis to derive their own definition by identifying reoccurring themes and forming these into an inclusive definition which happens to account for each of the Four Ps:

> 'Creativity is the interaction among *aptitude, process, and environment* by which an individual or group produces a *perceptible product* that is both *novel and useful* as defined within a *social context*' [41, p. 90]

In reviewing Four Ps research, Kaufman [23] described addendums that have been suggested for the Four Ps: persuasion [49]

[4] If the inner workings of a program are very impressive, complex or novel, then we may still be impressed by the program, but this is a different perspective to whether or not we think the program is creative.

[5] As Rhodes' work appeared in a relatively unknown journal, many later advocates of a 'Four Ps'-style approach to creativity seem unaware of Rhodes' contribution (e.g. Odena, 2009, personal communications), so fail to cite him.

and potential [47]. In general, however, the Four Ps have been adopted as they were originally conceived by various researchers [43, 51, 34, 26].

2.1 The Four Ps: Person

This perspective addresses human characteristics associated with creative individuals or groups of people. Encouraged by Guilford's call in 1950 for studying the creative person, an abundance of different personal characteristics have been associated with creativity [43, 51, 24, 53, 35], ranging from personality traits, attitudes, intelligence and temperament to habits and behaviours (for example curiosity, persistence, independence and openness). Some of these are closely related; others are contradictory. Rhodes mentioned the relevance to creativity of people's personality traits, attitudes and habits, physique and intelligence and the identifiable features of creative people, as well as referring to people's temperament, habits, self-concept, value systems, defence mechanisms, and behaviour [43, p. 307].

Empirical studies up until 1968 were summarised by Stein [52] and were combined into a list of 18 distinct personality characteristics of a creative person, including aspects such as curiosity, persistence, independence and openness. Stein used these characteristics to identify creative individuals for study. There is a risk of circularity here, as the selection criteria for creative individuals chooses people to be studied, then the study involves examining those characteristics and criteria. Stein's work has not stood the test of time, with few current citations.

Several researchers subdivide the 'Person' category into finer-grained groups. Three sub categories of the 'pupil' perspective emerged during Odena & Welch's work [35]: personal characteristics of the pupil, their individual learning style (either adapting to new information or deriving new information themselves) and the influence of the pupil's background. Koestler [24] described three types of creative person: the *Artist*, the *Sage* and the *Jester*. Through Tardif & Sternberg's review of definitions of creativity, three main categories were identified with which to describe creative people: cognitive attributes, personality attributes/motivation and developmental influences. Tardif & Sternberg suggested three resultant modes of study of human creativity: cognitive psychology; psychometric testing; and study of human development.

These discrepancies and the sheer quantity of attributes together place an obstacle in the way of compiling a definitive list of attributes of a creative person and instead provoke disagreements on exactly which cognitive characteristics should be attributed to creative people. Tardif & Sternberg's review showed that as of 1988, different authors highlighted a variety of characteristics, with no general consensus and no characteristics common to all reports [53, Table 17.1, p. 434].

2.1.1 *The Person in Computational Creativity*

In computational creativity, the creative person could be analogous to the computer, or perhaps more accurately, to the computer program, software, or to a creative agent within a multi-agent system. Here the machine is the hardware hosting the creative agent, much as we might distinguish between physical and functional characteristics of a 'Person'.

Interesting work has been done on modelling creative agents, for example by Saunders [48], although the emphasis in computational creativity software tends to be on product generation and to some extent, process modelling, rather than on the modelling of characteristics of a creative Person in computational format. This is because computational creativity systems tend to be oriented towards a particular goal or domain, rather than being generally creative, as we can see by the plethora of domain-specific systems (as opposed to modelling of creative personal characteristics) in the various proceedings of ICCC conferences (International Conference on Computational Creativity). As argued in [20], different types of creativity require domain specific skills to some extent, so domain-specific computational creativity systems tend to be built around the most prominent necessary skills for that domain.

In terms of evaluating creative systems, Colton's Creative Tripod [8] emphasises the need for systems to demonstrate skill, imagination and appreciation before they can be considered as a candidate creative system, all three of which are alluding to personal characteristics.

Features, traits and aspects of the creative system can be studied, and it would be fascinating to explore how general creative personal characteristics could be specifically modelled within creative systems (see the *Process* section, next). Computational modelling of characteristics that encourage creativity could help us progress our systems to be able to be creative in more than one system which they were originally designed for; this would be significant progress in our pursuit of modelling creativity as a phenomenon which transcends different types of creative activity.

The 'Person' could also entail the individual(s) interacting with a creativity support system or co-creative system which interacts with people[27, 22]. Another possible interpretation of the 'Person' in computational creativity would be to acknowledge the role of the programmer(s), tester(s), researcher(s) and peers involved in shaping the project.

2.2 The Four Ps: Process

The creative process has been broken down into a series of sequential or cyclic stages occurring over time [42, 55] or subtasks [35].

In their work on student creativity in school music lessons, Odena & Welch [35] broke down the creative process into subtasks, identifying various types of process (e.g. different activities, group process, the structuredness or otherwise of a process and composition by improvisation) rather than tracing a linear progression of subprocesses.

It is often stressed that creativity is not just the first flash of inspiration, but is also the activity that validates, develops, and refines that first idea; rather than occurring at one point in time, creativity develops over a period of time [55, 42, 53]. Tardif & Sternberg [53] questioned whether creativity is a social or an individual process. The social view of creativity has notably been promoted by Csikszentmihalyi [12].

2.2.1 *The Process in Computational Creativity*

In computational creativity, the creative process might be that employed by a single piece of software, or the interactions between multiple machines or programs, or the interactions between machine and human users. As described above, the computational creativity community has given some attention to the concept of creative processes employed within computational creativity, with growing attention paid to this aspect in recent years. For example, the FlowR framework [5] is designed to facilitate creative computational workflows by chaining together processes in a linear pattern, and from personal communications with members of the project team, there are plans to consider non-linear chains of processes as well. Additionally, the work by Joanna Misztal on poetry generation [31] specifically focusses on the processes required to generate poetry, at various levels of abstraction.

The generate-and-test [30, 38] or engagement-reflection approach [40] specifically models the creative process as a cycle of generating artefacts then improving the generation process via evaluating the generation phase. This is an approach which deserves broader adoption within computational creativity; evaluation is a critical part of the creative process [42, 12]. In terms of post-implementation evaluation, the FACE model for evaluation of creative systems [9] places importance on computational systems being able to report on the creative process (this report is referred to in the FACE framework as a *Frame*).

There are multiple theories about how human creativity processes are structured (see for example [42, 12, 23, 14]). Computational creativity research can provide a test-bed for these psychological theories and allow us to explore if implementing the theories result in creative behaviour. Conferences such as the Creativity and Cognition series showcase work that links between theory and practice to some extent, but further activity along these lines would emphasise the validity of computational creativity research, allowing computational work to contribute to human creativity research and vice-versa.

2.3 The Four Ps: Product

Many authors advocate that *proof* of creativity is necessary to be considered creative [21, 53, 41, 44]. The product-centric view adopted by computational creativity researchers such as Ritchie [45], that creative products are both necessary and sufficient for creativity, was present in earlier human creativity research [21]. But, inspired by Guilford's seminal 1950 address on creativity research, emphasis in human creativity research shifted from identifying creative individuals post-production of creative work, to predicting future potential for creativity in individuals. This change in emphasis is illustrated in the proliferation of psychometric tests [23, 19] within creativity research.

Tardif & Sternberg [53] considered the creative product more briefly than the other three 'Ps' in their review, deciding that while a creative product is essential for creativity, it is not enough merely to generate a product; the product should also be considered in a domain-specific context.

Computational creativity research has long acknowledged the importance of the output or artefacts generated by creative systems, as described above. To borrow a metaphor from human creativity research, it has been common (until recently) for computational creativity to follow the product-centric approach to creativity as advocated by Kagan: '*Creativity* refers to a product, and if made by a man, we give him the honor of the adjective' [21, p. viii].

2.3.1 The Product in Computational Creativity

Generating creative products has been an area of significant success for computational creativity. To see examples, one just needs to consult any year's proceedings of the International Conference on Computational Creativity where there are multiple examples to be found of systems which are reported in terms of the products they generate. The success of systems is often reported in terms of what kind of artefacts they generate, as noted in [18]. Some systems have been evaluated using Graeme Ritchie's empirical criteria [44, 45], which exclusively focuses on evaluating the products of computational systems without considering any of the other three Ps.[6]

2.4 The Four Ps: Press/Environment

The Press perspective encompasses a bidirectional perspective between the environment which influences the creator and receives the creative work, and the creator who publicises their work and is given feedback on what they produce. Tardif & Sternberg [53] considered both creative domains themselves and the social environments in which creative people are influenced as they employ creative process, advertise their creative products and receive feedback. Rhodes [43] concentrated on the role that the environment plays on a person during the creative process, rather than how the creative produce is judged by the external world after being created. Rhodes reflected on how everyone is different, so everyone perceives the world in a unique way and processes ideas according to their own contexts.

Of the Four Ps, this is the perspective that is often neglected when one takes an individualistic view of creativity. In general creativity theorists do however acknowledge the influence of the environment in which creativity is situated [49, 13]. If one concentrates on an individual's creativity, however, the Press perspective is often neglected, even if unintentionally. For example, although stating that '[t]o be appreciated as creative. a work of art or a scientific theory has to be understood in a specific relation to what preceded it' [3, p. 74], Boden's treatment of creativity mainly focused on different cognitive processes of creativity, rather than a detailed examination of social or environmental influences.

2.4.1 The Press in Computational Creativity

Some computational creativity researchers are starting to highlight the importance of the environment in which a creative system is situated [50, 17, 37, 48], with some of this work influenced by the DIFI (Domain-Individual-Field-Interaction) framework [12]. Social interaction between creative agents and their audience is an area which has been neglected by all but a few groups of researchers: for example nearly 75% of papers in the 2014 International Conference on Computational Creativity failed to make any reference to social or interactive aspects of creativity. But creativity cannot exist in a vacuum. A recent increase in development of the interactivity of creative systems (especially where this affects the way these systems works) is pleasing to see and deserves further attention [10].

There is a separate point to acknowledge regarding Press in computational creativity. As computational creativity researchers, we should stay aware of any potential biases that may be introduced, should an audience be aware that the creative agent of interest is computational rather than human [32, 19].[7]

2.5 Interaction between the Four Ps

Simonton [49] saw discrepancies between combining the Four Ps in theory and in practice:

> 'Now, in an ideal state of affairs, it should not matter which one of the four p's our investigations target, for they all will converge on the same underlying phenomenon. ... But reality is not so simple, needless to say. The creative process need not

[6] Recently proposed evaluation methods such as [8, 9, 19] place more emphasis on the other three 'Ps'.

[7] Many thanks to the anonymous reviewer who noted this point.

> arrive at a creative product, nor must all creative products ensue from the same process or personality type; and others may ignore the process, discredit the product, or reject the personality when making attributions about creativity.' [49, p. 387]

From this, one conclusion which seems to follow naturally is that an accurate and comprehensive definition of creativity must account for the (potential) presence of all four aspects, in order to be complete. Simonton, however, concluded that '[i]f we cannot assume that all four aspects cohesively hang together, then it may be best to select one single definition and subordinate the others to that orientation' [49, p. 387], with his natural research inclination leading him to focus his work on *persuasion*, his term for the Press/Environment aspect.

The mysterious impression often associated with creativity [56, 3, 23] can be explained to some extent when one or more of the Four Ps are not accounted for:

> 'Each strand [of the Four Ps] has unique identity academically, but only in unity do the four strands operate functionally. It is this very fact of synthesis that causes fog in talk about creativity and this may be the basis for the semblance of a "cult".' [43, p. 307]

Rhodes argued that creativity research should follow a specific path: 'from product to person and thence to process and to press.' [43, p. 309]

> 'Objective investigation into the nature of the creative process can proceed in only one direction, i.e. from product to person and thence to process and to press.' [43, p. 309]

Such a statement makes Rhodes's contribution less useful. For example, the Press (environment) in which one is creative has some influence on the creative Process, so one may prefer to study how Press and Person interact before looking at Process issues. Simonton viewed creativity as how a person's ideas emerge as influential when that person, by chance, has new ideas and promotes them to influence others. Creative people would not be equivalent to lucky people, by this interpretation, but chance would intervene in their success. Simonton refers to this as the 'chance-configuration theory' that 'outlines the general conditions that favor creativity' [49, p. 422].

Tardif & Sternberg [53] treated each of the Four Ps individually, 'as these really are separate levels of analysis, and it is from comparisons within levels that coherent statements about our knowledge of creativity can be made' [53, p. 429]. Tardif & Sternberg's summary is weakened somewhat by this as it does not make comparisons across the Four Ps, despite highlighting Simonton's emphasis on the interactions and relations between these four views [49]. In contrast Mooney [34] argued that the four approaches should be integrated in a model of creativity, proposing a model that 'puts together the four approaches by showing them to be aspects of one unifying idea' [34, p. 333]. While Mooney's claims become rather grandiose at points, Mooney's more specific contributions on creativity match neatly with the four Ps approach identified elsewhere at that date [43, 51]

2.5.1 *Interaction between Four Ps in Computational Creativity*

This paper argues that we can make significant progress in computational creativity by considering all four Ps in our computational creativity work. Tony Veale's tagline for the ICCC'2012 conference sums up current aspirations of computational creativity well; Veale characterises computational creativity research as 'scoffing at mere generation for more than a decade'. Generation of creative products is only a quarter of the full picture of creativity, only one of the Four 'Ps'. Granted, we have achieved much success in product generation, as exemplified by exhibitions, concerts and other demonstrations of creative products reported in various papers on computational creativity systems [18]. However, the more mature work and exciting potential comes from the incorporation of the other three Ps, at least to some extent, such as in [40, 48, 31].

3 Applying the Four Ps: examples of *novelty* and *value*

Novelty (originality, newness) and value (usefulness, appropriateness) form key parts of creativity [28, 3, 45, 20], often being identified as the two main aspects of computational creativity [39, 45, 4, for example].[8] Work in computational creativity illustrates both novelty and utility from each of the Four P perspectives, although some perspectives are represented more plentifully within computational creativity than others. To illustrate the discussions above, we can discuss novelty and value in computational creativity from each of the Four P perspectives. Considering novelty from each of the Four Ps:

Product Novelty is well associated with system outputs and products: how novel are the generated artefact(s)? The novelty of artefacts generated by computational creativity systems is a key consideration in Ritchie's empirical criteria for evaluating creative systems [45].

Process A creative process can take a novel approach or be implemented in a novel way, perhaps employing new algorithms or techniques or different approaches. Efforts at trying new processes and combinations thereof are being encouraged by systems such as the FlowR framework [5], which focuses specifically on enabling us to chain different processes together for creative purposes.

Person Creativity can be performed by a new creative entity, which demonstrates or uses novel characteristics relevant to that creativity. As is often encountered in computational creativity work, implementing or running a creative system on new hardware or in different software may also impact upon the system's performance and may have unexpected results. The number of new systems presented each year at the International Conference on Computational Creativity exemplifies how novel creative entities continually arise in computational creativity research.[9] (Also, the novelty of unexpected results is often unintentionally exemplified when live demos of these systems are attempted in unfamiliar computing setups.)

Press The creativity demonstrated by a system can be noted as being novel in a particular environment, even though it may be commonplace in other environments. The system may also exploit the surrounding press in previously unexplored ways. This was demonstrated neatly by the combination of two systems in [33], where a textual annotation system interacted with a system that generates emotion-driven music. The combination resulted in novel interpretations of fairy tales; such results would not have arisen were the systems operating in isolation.

Considering value from each of the Four Ps:

[8] It should be clarified that for this author, creativity consists of considerably more than novelty and value, though these are two key components of creativity. See [20].

[9] See `http://www.computationalcreativity.net/conferences`.

Product Value is also well associated with system outputs and products: how valuable or good are the generated artefact(s)? This is a highly current area of concern within computational creativity, with much evaluation concentrating on the quality of output [18].

Process The creative processes being incorporated within creativity can be useful in themselves for learning or studying how certain approaches and techniques work or for cross-application to new areas. Systems with an emphasis on modelling process, such as Misztal and Indurkhya's poetry generator [31] bring added utility by what they reveal about the processes being modelled.

Person Some creators become more valuable than others as a contributor in their field, based on their personal characteristics, experience and influence.[10] The same can be noted for creative systems to some extent; some are cited more often than others, for example Simon Colton's HR mathematical discovery system [7] (which provides a useful example of creativity in a non-artistic domain).

Press If creative activities benefit the external world in some way, then they have value to the press. As example, Harold Cohen's AARON colouring system has received much external attention, from media discussions [29] through to inspiring a screensaver for personal computers via `http://www.kurzweilcyberart.com`.

These above lists are not intended to be a full and conclusive portrait of novelty and value within computational creativity. What these lists illustrate is the different viewpoints that can be uncovered using the Four Ps as *signposts* with which to guide our thinking around computational creativity. The breadth of issues mentioned above shows aspects of novelty and value within computational creativity which may not always be accounted for if taking a product/process-oriented viewpoint; however it is argued here that those perhaps-overlooked aspects give us a closer rendition of creativity, guiding us away from incomplete viewpoints of creativity in the context of our computational work.

4 Summary

The difficulty of understanding what creativity is should not discourage us from such an attempt [43, 41, 8]. In creativity research, the *Four Ps* construct ensures we pay attention to four key aspects of creativity: the creative Person, the generated Products, the creative Process and the Press/Environment hosting and influencing the creativity. This framework helps us to consider creativity more broadly.

For example, if viewing *novelty* and *value* from the perspectives of *product*, *process*, *person* and *press*, we uncover various interpretations of these two key concepts within computational creativity which may otherwise have been overlooked. The *Four Ps* framework helps to highlight different perspectives on creativity, to portray creativity in a fuller context.

ACKNOWLEDGEMENTS

Many thanks to Carly Lassig for originally making me aware of the Four Ps and pointing me towards Rhodes [43]. Thanks also to various computational creativity researchers who provided feedback on earlier versions of these thoughts while in development, and the anonymous reviewers of this paper.

[10] This has been found, for example, in the recent Valuing Electronic Music project `http://valuingelectronicmusic.org` [1], where some people's endorsements can have a greater influence on the perceived value of an electronic musician and their work.

REFERENCES

[1] Daniel Allington, Byron Dueck, and Anna Jordanous. Networks of value in Electronic Dance Music: SoundCloud, London, and the importance of place, submitted.

[2] J. B. Biggs, 'Constructing learning by aligning teaching: constructive alignment', in *Teaching for quality learning at university: what the student does*, 11–33, SRHE & Open University Press, Buckingham, (2003).

[3] Margaret A. Boden, *The creative mind: Myths and mechanisms*, Routledge, London, UK, 2nd edn., 2004.

[4] David Brown, 'Computational artistic creativity and its evaluation', in *Computational Creativity: An Interdisciplinary Approach*, number 09291 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany, (2009).

[5] John Charnley, Simon Colton, and Maria Teresa Llano, 'The flowr framework: Automated flowchart construction, optimisation and alteration for creative systems', in *Proceedings of the 5th International Conference on Computational Creativity*, (2014).

[6] Nick Collins, 'The analysis of generative music programs', *Organised Sound*, **13**(3), 237–248, (2008).

[7] Simon Colton, *Automated theory formation in pure mathematics*, Distinguished dissertations, Springer, London, UK, 2002.

[8] Simon Colton, 'Creativity versus the perception of creativity in computational systems', in *Proceedings of AAAI Symposium on Creative Systems*, pp. 14–20, (2008).

[9] Simon Colton, John Charnley, and Alison Pease, 'Computational Creativity Theory: The FACE and IDEA descriptive models', in *Proceedings of the 2nd International Conference on Computational Creativity*, pp. 90–95, Mexico City, Mexico, (2011).

[10] Simon Colton and Geraint A. Wiggins, 'Computational creativity: The final frontier?', in *Proceedings of 20th European Conference on Artificial Intelligence (ECAI)*, pp. 21–26, Montpellier, France, (2012).

[11] David Cope, *Computer Models of Musical Creativity*, MIT Press, Cambridge, MA, 2005.

[12] Mihaly Csikszentmihalyi, 'Society, culture, and person: a systems view of creativity', in *The Nature of Creativity*, ed., Robert J. Sternberg, chapter 13, 325–339, Cambridge University Press, Cambridge, UK, (1988).

[13] Beth A. Hennessey and Teresa M. Amabile, 'The conditions of creativity', in *The Nature of Creativity*, ed., Robert J. Sternberg, chapter 1, 11–38, Cambridge University Press, Cambridge, UK, (1988).

[14] Beth A. Hennessey and Teresa M. Amabile, 'Creativity', *Annual Review of Psychology*, **61**, 569–598, (2010).

[15] Douglas Hofstadter, 'Creativity, brain mechanisms and the Turing test', in *Fluid concepts and creative analogies*, ed., Douglas Hofstadter, 467–491, Harper Collins, New York, (1994).

[16] Ivar Jacobson, Grady Booch, James Rumbaugh, James Rumbaugh, and Grady Booch, *The unified software development process*, volume 1, Addison-Wesley Reading, 1999.

[17] K. E. Jennings, 'Developing creativity: Artificial barriers in artificial intelligence', *Minds and Machines*, **20**(4), 489–501, (2010).

[18] Anna Jordanous, 'Evaluating evaluation: Assessing progress in computational creativity research', in *Proceedings of the Second International Conference on Computational Creativity (ICCC-11)*, Mexico City, Mexico, (2011).

[19] Anna Jordanous, *Evaluating Computational Creativity: A Standardised Procedure for Evaluating Creative Systems and its Application*, Ph.D. dissertation, University of Sussex, Brighton, UK, 2012.

[20] Anna Jordanous, 'A Standardised Procedure for Evaluating Creative Systems: Computational creativity evaluation based on what it is to be creative', *Cognitive Computation*, **4**(3), 246–279, (2012).

[21] *Creativity and learning*, ed., Jerome Kagan, Beacon Press Press, Boston, MA, 1967.

[22] Anna Kantosalo, Jukka Toivanen, Hannu Toivonen, and Ping Xiao, 'From isolation to involvement: Adapting machine creativity software to support human-computer co-creation', in *Proceedings of 5th International Conference on Computational Creativity, ICCC*, Ljubljana, Slovenia, (2014).

[23] James C. Kaufman, *Creativity 101*, The Psych 101 series, Springer, New York, 2009.

[24] Arthur Koestler, *The act of creation*, Danube Books, New York, 1964.

[25] Sol LeWitt, 'Paragraphs on conceptual art', *Artforum International Magazine*, **June 1967**, (1967).

[26] Douglas W. MacKinnon, 'Creativity: a multi-faceted phenomenon', in

Creativity: A Discussion at the Nobel Conference, ed., J. D. Roslansky, 17–32, North-Holland Publishing Company, Amsterdam, The Netherlands, (1970).
[27] Mary Lou Maher, 'Computational and collective creativity: Who's being creative?', in *Proceedings of the 3rd International Conference on Computer Creativity*, Dublin, Ireland, (2012).
[28] Richard E. Mayer, 'Fifty years of creativity research', in *Handbook of Creativity*, ed., Robert J. Sternberg, chapter 22, 449–460, Cambridge University Press, Cambridge, UK, (1999).
[29] Pamela McCorduck, *Aaron's code: Meta-art, artificial intelligence, and the work of Harold Cohen*, WH Freeman, New York, NY, 1991.
[30] G. McGraw and Douglas Hofstadter, 'Perception and creation of diverse alphabetic styles', *AISB Quarterly*, **85**, 42–49, (1993).
[31] Joanna Misztal and Bipin Indurkhya, 'Poetry generation system with an emotional personality', in *Proceedings of 5th International Conference on Computational Creativity, ICCC*, (2014).
[32] David C. Moffat and Martin Kelly, 'An investigation into people's bias against computational creativity in music composition', in *Proceedings of the 3rd International Joint Workshop on Computational Creativity (ECAI06 Workshop)*, Riva del Garda, Italy, (2006).
[33] Kristine Monteith, Virginia Francisco, Tony Martinez, Pablo Gervás, and Dan Ventura, 'Automatic generation of emotionally-targeted soundtracks', in *Proceedings of the 2nd International Conference on Computational Creativity*, pp. 60–62, Mexico City, Mexico, (2011).
[34] Ross L. Mooney, 'A conceptual model for integrating four approaches to the identification of creative talent', in *Scientific Creativity: Its Recognition and Development*, eds., Calvin W. Taylor and Frank Barron, chapter 27, 331–340, John Wiley & Sons, New York, (1963).
[35] Oscar Odena and Graham Welch, 'A generative model of teachers' thinking on musical creativity', *Psychology of Music*, **37**(4), 416–442, (2009).
[36] Marcus Pearce and Geraint Wiggins, 'Towards a framework for the evaluation of machine compositions', in *Proceedings of the AISB Symposium on AI and Creativity in Arts and Science*, York, UK, (2001).
[37] Alison Pease and Simon Colton, 'On impact and evaluation in computational creativity: A discussion of the Turing Test and an alternative proposal', in *Proceedings of the AISB'11 Convention*, York, UK, (2011). AISB.
[38] Alison Pease, Markus Guhe, and Alan Smaill, 'Some aspects of analogical reasoning in mathematical creativity', in *Proceedings of the International Conference on Computational Creativity*, pp. 60–64, Lisbon, Portugal, (2010).
[39] Alison Pease, Daniel Winterstein, and Simon Colton, 'Evaluating machine creativity', in *Proceedings of Workshop Program of ICCBR-Creative Systems: Approaches to Creativity in AI and Cognitive Science*, pp. 129–137, (2001).
[40] R. Pérez y Pérez, A. Aguilar, and S. Negrete, 'The ERI-Designer: A computer model for the arrangement of furniture', *Minds and Machines*, **20**(4), 533–564, (2010).
[41] Jonathan A. Plucker, Ronald A. Beghetto, and Gayle T. Dow, 'Why isn't creativity more important to educational psychologists? Potentials, pitfalls, and future directions in creativity research', *Educational Psychologist*, **39**(2), 83–96, (2004).
[42] Henri Poincaré, 'Mathematical creation', in *The Foundations of Science: Science and Hypothesis, The Value of Science, Science and Method.*, volume Science and Method [Original French version published 1908, Authorized translation by George Bruce Halsted], chapter III of Book I. Science and the Scientist, 383–394, The Science Press, New York, (1929).
[43] Mel Rhodes, 'An analysis of creativity', *Phi Delta Kappan*, **42**(7), 305–310, (1961).
[44] Graeme Ritchie, 'Assessing creativity', in *Proceedings of the AISB Symposium on AI and Creativity in Arts and Science*, pp. 3–11, York, UK, (2001).
[45] Graeme Ritchie, 'Some empirical criteria for attributing creativity to a computer program', *Minds and Machines*, **17**, 67–99, (2007).
[46] Graeme Ritchie, 'Uninformed resource creation for humour simulation', in *Proceedings of the 5th International Joint Workshop on Computational Creativity*, pp. 147–150, Madrid, Spain, (2008).
[47] M.A. Runco, 'Creativity, cognition, and their educational implications', in *The educational psychology of creativity*, ed., John Houtz, 25–56, Hampton Press, New York, NY, (2003).
[48] R. Saunders, 'Towards autonomous creative systems: A computational approach', *Cognitive Computation*, **4**(3), 216–225, (2012).
[49] Dean Keith Simonton, 'Creativity, leadership, and chance', in *The Nature of Creativity*, ed., Robert J. Sternberg, chapter 16, 386–426, Cambridge University Press, Cambridge, UK, (1988).
[50] Ricardo Sosa, John Gero, and Kyle Jennings, 'Growing and destroying the worth of ideas', in *Proceedings of the 7th ACM Creativity and Cognition conference*, pp. 295–304, Berkeley, California, (2009).
[51] Morris I. Stein, 'A transactional approach to creativity', in *Scientific Creativity: Its Recognition and Development*, eds., Calvin W. Taylor and Frank Barron, chapter 18, 217–227, John Wiley & Sons, New York, (1963).
[52] Morris I. Stein, 'Creativity', in *Handbook of personality theory and research*, eds., Edgar F. Borgotta and William W. Lambert, 900–942, Rand McNally, Chicago, IL, (1968).
[53] Twila Z. Tardif and Robert J. Sternberg, 'What do we know about creativity?', in *The Nature of Creativity*, ed., Robert J. Sternberg, chapter 17, 429–440, Cambridge University Press, Cambridge, UK, (1988).
[54] D. Ventura, 'A reductio ad absurdum experiment in sufficiency for evaluating (computational) creative systems', in *Proceedings of the 5th International Joint Workshop on Computational Creativity*, pp. 11–19, Madrid, Spain, (2008).
[55] Graham Wallas, *The Art of Thought*, C. A. Watts & Co, London, UK, abridged edn., 1945.
[56] Raymond Williams, *Keywords: a vocabulary of culture and society*, Fontana/Croom Helm, Glasgow, UK, 1976.

How Many Robots Does It Take? Creativity, Robots and Multi-Agent Systems

Stephen McGregor and **Mariano Mora McGinity** and **Sascha Griffiths**[1]

Abstract. This paper seeks to situate computational creativity within the context of ongoing theoretical and practical investigations of environmentally situated and dynamic systems. Beginning with a consideration of the evidently goal directed nature of creativity, the problem of how teleological behaviour emerges in a fundamentally physical world. Creativity is reassessed as a search for goals in a dynamic environment rather than as a pursuit of a fixed goal in a stable and finite space of possible actions. A significant consequence of this evaluative shift is the impossibility of considering truly creative systems as anything other than embodied agents deeply entangled in an environmental situation. Two fields are discussed as potential habitats for such systems: robotics and multi-agent systems. Creativity from the perspective of ongoing research in these areas is considered, and some preliminary thoughts for future directions of enquiry are offered.

1 Introduction

This paper will address the question of the relationship between goals and creativity. Notions of purpose are so deeply ingrained in the standard view of creation that creativity itself is often defined in terms of the accomplishment of some expressive objective. Implicit in the problem of modelling creativity, however, is the emergence of end directed action in a reductionist world: how can something that is not in a physical sense present nonetheless contribute to the operation of a physically supervenient system?

Having posed the question of how a creative agent views its own objectives, the paper will turn to an exploration of the related problem of causality. In particular, the emergence of absent causes – which is to say, the influence of possible worlds, both historical and futuristic, removed from present reality – is addressed. This etiological inquiry is couched in terms of evolution by natural selection, with a brief consideration of this well researched process as a model of evidently goal directed and therefore potentially creative behaviour. A general hypothesis regarding the viability of explaining goals as emergent properties of complex systems, grounded in contemporary theoretical investigations of dynamic systems, will be put forward. *Contra* the idea that computationally creative agents must necessarily be handed a well defined goal by an external designer, dynamic processes are proposed as a basis for models that can discover their own goals through collaboration and environmental interaction.

This theoretical consideration is followed by a preliminary exploration of two compelling areas of research that move beyond what has been the *de rigueur* constraint satisfaction approach to computational creativity. First the topic of robotics will be considered from the perspective of the modelling of creativity, with particular attention to the problem of how a robot obtains, represents, and adapts its own goals. Robots are importantly embedded in a physical environment, and this situation opens the door to the possibility of the emergence of dynamic attractors that might be construed as new and unexpected goals outside any representation of an objective built into a robot's programming. The conclusion of this investigation will be that it seems reasonable to at least consider the possibility of an adequately flexible robot formulating goals that can be considered as evidence of its own creativity.

Next multi-agent models are considered, with particular attention to the ways in which complex patterns of activity with the trappings of intentionality can emerge from interactions within a population of agents individually following very basic sets of predetermined rules. As with robots in their environmental entanglements, swarms of agents have some prospect of generating collective behaviour that can be interpreted as being directed towards ends outside of the simple constraint satisfaction requirements programmed into the functioning of each independent agent. In the case of multi-agent models, the model becomes the environment, with the attractors that arise in the course of interactions becoming the handles for assessing the system in terms of the formulation and pursuit of goals. On the one hand, interpretation of action in a simulation of such a system presumably still falls back on the analysis of an external observer. On the other hand, the emergent properties of such systems potentially offer models of the parallel emergence of cognitive phenomena such as creativity in a physically grounded universe. Again, there seems to be scope for considering the implementation of multi-agent systems as a form of computational creativity that begins with a traditional programming task but that subsequently goes beyond mere constraint satisfaction.

The ideas proposed in this paper are at this stage indications of direction for future research. Exciting work is currently being done in the fields of both robotics and multi-agent systems, with some applications specifically towards modelling creative systems [1, 17]. This paper is intended to serve as a bellwether for further research in this direction, with the objective of moving beyond a constraint satisfaction approach to computational creativity. Traditional rule based implementations of creative agents have accomplished much in recent years, but reconsidering the emergence of goals within highly dynamic environments offers the grounding of a more robust argument for creative autonomy arising within the systems themselves. This reconsideration of the relationship between creativity, goals, and the environmental situation of creative agents can furthermore become a platform for extended discussions of interesting philosophical questions about causation and cognition.

[1] Queen Mary University of London, UK, email: s.e.mcgregor, m.mora-mcginity, sascha.griffiths@qmul.ac.uk

2 Creativity and Goals

> What redeems it is the idea only... and an unselfish belief in the idea – something you can set up, and bow down before, and offer a sacrifice to...
>
> JOSEPH CONRAD, *The Heart of Darkness*

If an agent cannot choose its own goals, can its "creative" behaviour really be considered creative? This question is familiar to anyone who has attempted to construct a creative agents, especially agents that produce works of art. "*Why* did the computer choose this word, this note, this color, and not another?" is generally asked by people confronted with an artistic piece produced by a computer. If the same question were put to an artist the answer would most likely be as unsatisfactory as that provided by the computer, yet it is the computer's response that is most troubling. The assumption is that art made by computers is something that needs to justify itself, whereas it is accepted that artists produce according to their inspiration, perhaps because works of art are seen as reflections of the mind or the personality of the artist who made them, and it is troubling to think of a computer as having a "personality". However, the fuzziness of the artist's response seems to reflect something essential in the creative process, something which can and should be exploited in the design of artificial creative systems: the goal is not fixed; it shifts and drifts and wanders; it might not even exist *a priori* but rather will emerge from the creative process itself.

2.1 Causality and Teleology

Aristotle's "Doctrine of the Four Causes" presents a framework for understanding the relationship between actions and motivations in a world of physical events [2]. Starting with the essentially reductive premise that the material nature of entities is at the root of actions in the world, the Doctrine outlines a hierarchy of relationships, culminating in the theory of how teleological – which is to say, goal directed – behaviour stands as the "final cause" that explains the regularity with which functional effects are produced in the natural world. Aristotle's four causes can be enumerated:

1. Material Cause - the behaviour indicated by the physical properties of matter
2. Efficient Cause - the consequences of the manipulation of physical material
3. Functional Cause - the reasoning regarding efficient causes that informs actions on materials
4. Final Cause - the goal that motivates functional planning

Conventional approaches to creativity generally descend Aristotle's causal ladder: there is a goal, a plan for achieving this goal, a set of actions carried out to realise that plan, and a world of physical relationships in which those actions have consequences. Indeed, a fundamental principle of a certain approach to aesthetics is that the perception of beauty involves the recognition of a function that defines an artefact and an appreciation of the creative process employed in the achievement of that functionality [46]. An alternative theory, rooted in the philosophy of Kant, considers aesthetic experience to unfold in a perceptual domain of its own, involving a detachment from any practical consideration of an object of beauty [20]. Even in this latter case, though, beauty, from the perspective of a creator, becomes an objective unto itself, with the elicitation of an aesthetic response in principle indicating the achievement of this goal. So, regardless of the theoretical grounding adopted by an analyst, creativity seems to be bound up in an end directed process.

Computational creativity has tended to adopt a similar line. Ritchie has characterised the creative behaviour of an information processing machine in terms of the identification of a class of existing artefacts that qualifies as a target domain, subsequent generation of artefacts that are expected to fall within this domain, and then evaluation on the part of the system of whether the creative goal has been achieved [32]. Output produced without some sort of goal criteria has been described as "mere generation", a ramble through a state space that, regardless of its consequences, cannot be properly considered as creativity [11]. This lines up well with Boden's description of levels of computational creativity, with high level transformations of state spaces trumping lower level recombinations of elements within a predefined space [10]. It is in this transformational degree of symbol manipulation, involving the delineation of a state space above a traversal of a known space, where the complexity of the goal directed aspect of creativity becomes evident. A fundamental challenge for a computer scientist interested in designing autonomously creative systems is therefore to understand what it would mean for computers to make decisions about the definition of their own search spaces.

But it is not even clear how teleological processes arise in the material world, reducible, as it is, to the interactions of physical fields on a very small scale. Deacon has taken Aristotle's Doctrine as a starting point for his own exploration of the emergence of goal oriented behaviour in material reality, beginning with the premise that modern philosophy has sometimes tended to use dualism and homunculi to obscure the hard question of final cause [13]. For Deacon the first step up from the tumult of pure physics is a consideration of thermodynamic processes, by which a tendency that is so reliable it has a nomic aspect emerges from the random interaction of particles. In fact, despite the regularity implicit in the terminology "laws of thermodynamics", there is no principle that requires systems to move towards entropic arrangements; it is just the overwhelmingly likely outcome of a stochastic process. The kernel of teleology might be discovered in the apparent lawfulness of entropy that arises in systems that are actually just complex and unpredictable.

Like Deacon, Kauffman recognises the seeds of emergence in the way that order can spontaneously come about in a dynamic system, giving rise to interpretable attractors [19]. The contemporary case for emergence maintains that nested hierarchies of interacting attractors can be extrapolated into apparently teleological behaviour. At the higher end of the scale exist cases like evolution by natural selection, which, while it has been grasped through an astounding act of reductionist interpretation, can nonetheless only really be understood as a process directed towards the goal of fitness—and in fact it has been argued that evolution itself should be treated as a creative process. To put it simply, an evolved organism is a confluence of functions that result in their own perpetuation. Taking an example offered by Millikan, the biological operation of an animal can only be understood in terms of the functional role that the creature's various organs play in sustaining life, and these functions have been determined through an assiduous process of evolutionary trial and error: a lion's heart exists in order to pump blood through the lion's body, even though the genetic and developmental process that resulted in the existence of the organ cannot have been somehow aware of that outcome [24].

But, Millikan asks, what happens if a fully developed lion comes into existence spontaneously? While the lion might be considered an operational organism, it is tempting to conclude that its organic components have no function in the sense of having been selected because of a goal they accomplish. An evolved lion inherits properties of goal directedness from the generational history of organisms that has contributed to its fitness. This extension of the lion's emergent identity into the past corresponds to a converse projection of the functional properties of its components towards the accomplishment of future goals, specifically the goals of the lion surviving and reproducing. The spontaneous lion, on the other hand, while it also has some hope of coincidentally surviving and replicating, has simply happened: it cannot be interpreted as the fulfillment of a goal that has emerged in the unfolding of events in a complex and unpredictable environment. In terms of Aristotle's efficient cause, the lions are identical, but in terms of final cause they seem to be completely different.

Bickhard has responded to Millikan's case for a connection between causal history and functionality, however, by arguing that the history of a system cannot be a part of its ongoing operation [9]—history is, presumably, a contextualised interpretation of a present situation. Instead, Bickhard proposes, function should be understood in terms of the contribution a functional component makes to its system's persistence in a state that defies the entropic tendencies of the universe [8]. It is the case that the dynamics of complex and chaotic systems result in the emergence of processes that, in their regularity, seem to have a sense of following some kind of rule. This shift away from the basic laws of physics begins with processes such as thermodynamics, where the regularity lies precisely in the predictable breakdown of structure in systems, and moves out towards the further from equilibrium states that characterise the process of evolution, or more explicitly cognitive apparati such as representational symbols.

So by the emergentist account, causation is understood in terms of nested layers of dynamically coupled, intricately entangled processes, with each emerging attractor becoming an element in a higher level of interactions. This view escapes the paradoxes of trying to incorporate some representation of the system's past into its current operation, and at the same time seeks to explain the evident gravity of future outcomes in the workings of higher order complexes. The upshot of this is that teleological processes are necessarily associated with systems that are highly non-linear on several levels, an insight that sits well with the enactivist world view of Varela, Thompson and Rosch, who suggest that a mindful agent – which is to say, one capable of the planning and execution inherent in creativity – must be situated in a deeply interactive relationship with a dynamic and unpredictable environment [42].

There is a gravely concerning ramification to this conclusion from the perspective of a computer scientist interested in designing autonomously creative agents, however: if teleological processes only emerge in the context of complex interaction with a chaotic environment, it is difficult to imagine how a symbol manipulating machine could hope to creatively flourish in its rule based domain. Considering that even computational processes modelled non-deterministically can be reduced to linear operations, the case for a strictly algorithmic system producing output that would be judged even basically creative seems doomed. Two possibilities immediately present themselves as the beginning of a solution to this challenge: the modelling of dynamic interactions between rule following agents, and the physical construction of environmentally situated robots.

3 Robots

The classic intelligent agent concept [34] entails that an agent should be able to use actuators to manipulate its environment, which it monitors via perception. The agent has goals which it is trying to satisfy via its actions. Whether these goals have been reached is subject to an evaluation which the agent achieves by applying a metric. In classic AI, the agent's environment is understood much less literally than it is in robotics. Robots exist in a physical world that they actively manipulate and that directly affects their actions. Also, they share this world with humans. What follows offers a brief survey of contemporary approaches to robotics as they relate to creativity, followed by some thoughts on the future exploration of robots as creative agents.

From the perspective of the description of creative systems, the great appeal of robots lies in their situation in the same highly non-linear environment from which human creators have emerged. As a first approximation, robots might be considered to have goals that are handed to them by a designer, grounded in external observations: in this case, the robot becomes an expression of its own creator's stance towards the world, and even in this basic instance a dynamic emerges where the robot's behaviour can become an element in a larger creative system, with the designer responding to the robot's successes and failures through subsequent design decisions. In what follows, this scenario will be case in terms of robots as a form of creative expression. More complexly, robots might be modelled as adaptive agents involved in a feedback loop with their own environments. In this case, while there may be overarching goals handed to a robot by a designer, it is the behaviour that emerges in the pursuit of this goal that may be considered creative. Ultimately, it is conceivable that robots or perhaps even more compellingly networks of robots might be involved in processes with unpredictable outcomes that can be interpreted as the emergence of truly goal oriented causation.

It has recently been argued [27] that real progress in natural language processing will depend on a more human-like machine which has a situated knowledge embodied in its own physical form. This presence [26] is necessary for a cognitive architecture which is more human-like and therefore capable of a human-like command of natural language. This may just as well be just as true for other cognitive abilities.

Feldman [16] sees two possible ways in which a robot can fully understand human subjective experience. One is a full simulation of the human body to gain insights into human experience. The other would be a new type of grounding that builds up an understanding of the world through the robot's own sensors and bodily experience.

Creative automata and machines which exist in the physical world have been built for centuries. There is, for instance, the case of von Kempelen's speaking machine [43, 15], which was a hybrid between a research project on the human vocal apparatus and an entertainment tool similar to a musical instrument.

Creativity in the domain of robotics can be conceptualised in terms of creative activities that are performed by intelligent agents capable of performing a full action-perception loop which takes the environment into consideration. Within this action-perception loop the, agent must have some "creative goal".

3.1 Agents and Embodiments

In order to understand what it means for a robot to be creative we will now describe a few systems which do in some way fulfill the criterion of being "deemed creative" if they were "performed by a human" [44]. Creative robots come in two flavours currently: they are either presented as being creative themselves or they are used as tools for expressing a human's creativity. We will fist deal with the later kind of robot for creative tasks.

Robots as a form of creative expression are teleoperated, which is to say their actions are determined by the perceptions and decisions of a human performer. Ogawa et al. [30] report on a teleoperated robot called the “Geminoid” [36] being used for a task in which the android and an actor performed a play live on stage together. This robotic agent had the following properties:

- The android takes the shape of a physical body which is modelled on an actual female human. The body has 12 degrees of freedom (DoF). These are mainly used for its facial expressions which closely copy the operator’s facial movements. It also has loudspeakers which transmit the operator’s voice to the audience.
- Perception is accomplished through a camera system which lets the operator see the machine’s view of its environment.
- The machine’s processing of the environment is realised by feeding the video back to the operator, and its actions are hence based on receiving “commands from the human operator”.

So the robot’s body itself is used for artistic expression. The authors conclude, based on experiment, that the robot actually improved the audience’s sense of immersion in the performance. This is a surprising result but shows that the embodiment through the artificial agent can actually generate a different level of “meaning”, as the authors suggest. It is actually the human-like but not-human body that generates this added meaning.

Robots as creative agents are autonomous to a certain extent. Tresset and Deussen [40] report on a robot, named e-David, which creates visual art through painting on a canvas. This agent had the following properties:

- e-David is not anthropomorphic (human-like). It is an industrial robot that only consists of an arm. The arm is also its actuator, with which it manipulates a pencil or brush.
- The perceptive apparatus is a camera system.
- The system performs the action-perception loop by creating an image it intends to paint and then monitoring its output by perceiving the painting as it emerges through its own actions applied to a canvas.

Embodiment is crucial in the case of e-David. The authors list thirteen ways in which e-David’s embodiment has a direct impact on the final result of the visual art it produces. These include the velocity at which the arm moves, the pressure it applies to the painting, and control of the amount of paint on the brush. All of the factors have a direct effect on the visual appeal of the product which e-David produces. Thus, this robot demonstrates the importance of considering the physical presence of an artificial creative system in the creation of visual art.

Both the Geminoid and e-David illustrate how important the actual physicality of an intelligent agent it is and how their individual embodiments shape their creative output. However, the processing system in each case is actually quite different. Whereas e-David is autonomous in its actions to a large extent, the Geminoid is operated by a human. Thus, these two specific robots have different levels of autonomy and one needs to debate what “responsibilities”, in the sense of Colton and Wiggins [11], they take on within the creative process.

3.2 Goals

As already illustrated, robotic agents that use their physical appearance and structure to pursue creative objectives can differ in their goals. Whereas the Geminoid in the study discussed above tries to evoke emotional response in an audience, e-David monitors its own output on a canvas via a visual feedback system.

Similarly, musical robots have goals which they pursue. In this case, the environment is typically the musical instrument with which the robots interact physically.

A robot coordinating its own body in a creative process

Batula and Kim [6] present a system which plays the score of a two-finger piece on a piano. The robot in this case is a small humanoid. Its environment is a keyboard. Its perception relates to the monitoring of its own motions and audio-feedback.

The robot’s goal is to play the piece it has been assigned correctly. The authors frame their research as an investigation into the motorics required for musicianship. The robot’s goals are simple: it detects mistakes in its own playing. This is very much in line with our argumentation. The system’s physicality comes from the control of its own limbs in relation to the velocity of its playing. The robot controls its own motion, and the decisions of how to play rely solely on its own bodily control.

A robot coordinating with another body in a creative process

A contrasting approach is presented by Mizumoto et al [25]. In their approach the focus is on ensemble performance. The goal is for the agent to ask: “Am I creating the same output as another agent?”

This is a different question because the machine is no longer in control of the speed at which the product is created. The robot plays a theremin while the human plays drums. The robot’s perception is used to actually calculate the action of the actuators, in contrast to the actuators acting independently to exert force on the physical environment. The required processing relies on a coupled-oscillator model.

3.3 Environments

What kind of environments do robots encounter in the course of creative processes? The comedic robot is one recent concept which has been implemented. Thus far, these robots are the only agents which actually treat an audience as their environment. They do exactly what an intelligent agent does by monitoring what effects their actions have on the environment.

Audience Monitoring

Other agents with which robots interact may be artificial (see section 4) or human audiences. Knight [21] analyses the impact of embodiment on performances in robot theatre. Knight et al [22] present a system which tells jokes to an audience.

In the system described a small humanoid robot is the comedic agent. Its goal is to make the audience laugh. It monitors levels of audience interest and attention (more precise methods are further described in [23]). The robot presents jokes and will choose the sequence of jokes in accordance with the audience’s reaction. This is a direct application of the action-perception loop. The quality of the creative output is measurable in the sense that the audience reaction is the operationalisation of what the output should achieve.

Interacting with the Audience

Katevas at al [18] also use a humanoid robot as a stand-up comedian. In their performance, however, the robot actively engages with the audience and directly addresses individual members of the audience. In this way, the robot influences the outcome of the creative process. The goal is an active audience reaction, so the robot tries to improve the outcome and generate more laughs by engaging the audience.

As such the robot is not only relying on its output in the form of jokes, but also actively and preemptively shapes the audience’s reaction and hence its environment’s reaction to the jokes. This can be considered a different approach. If joke telling is considered an artistic and creative process, then the audience's reaction is the measure by which one can tell whether the result of the activity is of good quality. The robot here imitates the practices of human stand-up comedians by actively inducing a reaction in the audience. It does not just rely on the humorous value of the verbal stimuli it presents to the audience.

3.4 Creative Robots

In line with the theoretical points raised above, entertainment robotics is a growing market [7, 33]. The potential here is vast. A robot can use the principles outlined above to become an active companion [12, 3], giving itself an advantage over static media such as television broadcasting or film.

The three principles addressed here, embodiment, goals, and environments, will play a crucial goal in developing systems that can be deemed creative. This section has illustrated differing approaches to all three of these topics. In designing creative robotic systems, the human designer will have to think carefully about how the agent will pursue its goals within the given environment.

In line with the argument in this paper, for a robot to be truly creative it must be able to show adaptive behaviour. Embodiment will obviously be given from the outset in a robotic system, influencing the system's actions, perception, and interaction with the environment in a non-trivial way. However, real adaptivity for creativity will arise only if the robotic agent is able to define its own goals. An approach to robotics which includes this kind of behavioural autonomy is evolutionary robotics [29]. This approach assumes that the agent has some kind of overall goal such as playing a musical piece or amusing an audience via comedic practices. The sub-goals upon which the system operates would have to be adaptable. One way of implementing such a strategy would be to devise methods that allow the robot to choose between the goals outlined above (see section 3.2), or, with respect to interacting with the environment, choosing between the two strategies of, for instance, interaction with an audience as described above (see section 3.3).

4 Multi-Agent Systems

It is sometimes easy to forget that artists are not totally isolated from their environment: they come into contact with other artists who are tackling problems and trying to reach goals very similar to their own. Artists, scientists, chess players, normal people trying to make ends meet—creative people are influenced by other people, and they themselves influence other people, very often people with whom they are in no direct contact. Think of the generations of musicians influenced by Beethoven or of mathematicians working on problems formulated by Gauss.

In fact, one would be justified in thinking that creative processes are never the work of one individual alone, no matter how visionary and illuminating her thinking might be: every creator stands on the shoulders of giants. The intention here is to discuss how this interaction might be modelled through artificial agents, and how such an interaction might influence the behaviour of the agents towards, ultimately, determining the goal of the creative process itself.

As they relate to the imperative of creative goals as behavioural causes, the appeal of multi-agent systems is their potential for producing emergent attractors which cannot be understood as components of any single agent's behaviour. Agents themselves may be goal oriented – indeed, their processes are typically modelled in terms of the satisfaction of very basic criteria – but these goals are simplistic, whereas the operation of the overall system is nuanced. The power of simple agents collaborating to develop and realise complex goals can be observed in various real-world contexts, from the swarm behaviour of certain insects to the efficacious productivity of financial markets and indeed the homeostatic condition of entire ecological systems. This paper considers the question of how computers might be used to model multi-agent systems and then to analyse the potential for considering these systems as generators and executors of creative objectives.

4.1 Interacting agents

Multi-agent systems have been used extensively to model the origin and evolution of an impressive array of different social constructs [14, 31, 35], from ant or termite colonies to computer networks to economic markets. Agents are assigned a more or less rigorous set of beliefs, desires and intentions which determines their interaction. Agents are goal-oriented: their actions are determined by a desire to maximise a reward function, and it is through their interaction that the system evolves.

Most interestingly, multi-agent systems can show emergent properties: interaction between the agents allows the self-organisation of system properties that were not originally part of the system. Self-organisation, i.e. the lack of a centralised element imposing structure on the emergent property, is an important characteristic of such systems, revealing how organised properties can arise from simple interactions alone. These kinds of systems have been used, for instance, to model the self-assembling of biological complex structures[28], or to model the origin and evolution of language [37, 5, 4, 38]. In Steels' work, agents create and agree on a lexicon to name a series of objects in their environment. Their interaction follows a protocol specified in a "language game", similar to the language games described by Wittgenstein [45]. Van Trijp [41] shows that the "Naming Game" will converge towards a stable lexicon if certain requirements are met.

According to Tomasello [39]:

> The current hypothesis is that it is only within the context of collaborative activities in which participants share intention and attention, coordinated by natural forms of gestural communication, that arbitrary linguistic conventions could have come into existence evolutionarily...

This hypothesis seems to validate the modelling approach. An effort to understand creative processes as an attempt at collaborative behaviour by intelligent agents might prove to be very fruitful.

4.2 Creative processes as collaboration: a thought experiment

We propose a thought experiment which could help to illuminate the relationship between goal-seeking behaviour and creativity. Agents of different physical or cognitive characteristics are placed in an environment and forced to collaborate in order to achieve a series of tasks. To simplify things, we propose the following interaction rules:

1. All interactions are one-to-one: two agents are chosen and made to interact.
2. Agents are chosen at random: the system does not show a topology, i.e. it is a mean-field system.

3. One of the agents adopts the role of the demonstrator; the other is the observer.

Both agents have a clear idea of the task that is to be carried out. However, their different physical and cognitive skills require them to adapt their own actions to the task: some agents are better equipped to carry out the task in one way, whereas others must find efficient ways to carry out the task. During the interaction, the demonstrator performs the task in the most efficient way it can. Obviously, this way depends on all the previous experience of the agent. More particularly, it depends on what it has learned from all its previous interactions with other agents.

Following this demonstration, the observer must decide whether it is fit to perform the action in the same way. It does this by attempting to imitate the demonstrator. If it cannot, it must try to find a way to perform the action in a way that will resemble the demonstrator's actions, only adapted to its own abilities. If it succeeds in carrying out the action, then the observer will include this action into the set of actions it is capable of carrying out to perform the assigned task; the game is successful and two new agents are chosen to play new game. If, on the other hand, after a fixed number of attempts the agent is incapable of performing the action in a satisfactory way, then the game starts again, only now a new task is chosen: the goal changes. The new task should be similar to the previous one, if possible, so that agents might be able to identify properties of the task that are difficult for them, and perhaps learn to avoid them or find a way around them.

The hypothesis offered here is that such a system would become stable, i.e. it would reach a point after which all interactions would be successful. At this point all agents would have learned how to behave when forced to carry out a collaborative task. Every agent would have learned to adapt its own goal, according to its capabilities, to fulfill the task in a cooperative manner. Every agent would have learned how to work around what it cannot do.

5 Conclusion

This paper has examined the idea that creativity can be understood in terms of a process of adaptation on the part of agents attempting to accomplish a set of goals in complex and unpredictable environments. The hypothesis presented here is that agents dynamically coupled with their environments might become involved in the instigation of higher level emergent features that can be interpreted as potentially surprising and valuable new goal directed behaviours. There is scope for hoping that a network of multiple environmentally situated agents, each independently working towards their own micro-goals, will remit a systemic shift that in turn can become a target for discovery of new possible goals available to agents. From an external perspective, such a system offers the overall impression of being directed towards goals that are not in any way present in the programming that defines the behaviours of its components. In the physical universe, definable as it is in terms of a few simple rules of interaction, has nonetheless become a cauldron for such complex emergent systems evolution and cognition. In the same sense, a system of simple, interactive, environmentally oriented computational agents might have a chance of developing patterns of behaviour that can collectively be considered creative.

Existing work in the pertinent fields of robotics and multi-agent systems has been briefly discussed. The embodied situation of robots invites a consideration of the development of goal directed behvaiour in an unpredictable environment. And the dynamics of multi-agent systems present a platform for investigating the possibility of treating the attractors that emerge unexpectedly in the course of interaction as unanticipated creative objectives. The juxtaposition of these two topics in the context of computational creative naturally suggests an amalgamation: a potential project developing swarms of individually adaptive robots, treating their own community of robotic co-agents as an environment embedded in the physical world, with each robot adapting its behaviour based on its interaction with its peers. On an individual level, the robots would update their procedures based on observations of other robots and with the pre-programmed objective of accomplishing simple goals. On a collective level, the robotic system as a whole might very well take on an emergent aspect, with unexpected intimations of higher level organisation. The question raised by such a model is whether the system's proclivity for organising itself in a surprising and potentially effective way can be considered the creative discovery of a new objective.

ACKNOWLEDGEMENTS

Griffiths is supported by ConCreTe: the project ConCreTe acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 611733. McGregor is supported by EPSRC grant EP/L50483X/1.

REFERENCES

[1] Mohammad Majid al Rifaie, John Mark Bishop, and Suzanne Caines, 'Creativity and Autonomy in Swarm Intelligence Systems', *Cognitive Computation*, **4**(3), 320–331, (2012).

[2] Aristotle, *The Works of Aristotle Volume II: Physica*, The Clarendon Press, Oxford, 1930. Translated by R. P. Hardie and R. K. Gaye.

[3] Ruth S Aylett, Ginevra Castellano, Bogdan Raducanu, Ana Paiva, and Mark Hanheide, 'Long-term socially perceptive and interactive robot companions: challenges and future perspectives', in *Proceedings of the 13th international conference on multimodal interfaces*, pp. 323–326. ACM, (2011).

[4] A. Baronchelli, V. Loreto, L. Dall'Asta, and A. Barrat, 'Bootstrapping communication in language games: Strategy, topology and all that', in *The Evolution of Language, Proceedings of the 6th International Conference (EVOLANG6), edited by A. Cangelosi, A. D. M. Smith & K. Smith, World Scientific Publishing Company*, (2006).

[5] J Batali, 'Computational simulations of the emergence of grammar', in *Approaches to the Evolution of Language: Social and Cognitive Bases*, eds., J R Hurford, M Studdert-Kennedy, and Knight C., 405–426, Cambridge University Press, Cambridge, (1998).

[6] Alyssa M Batula and Youngmoo E Kim, 'Development of a mini-humanoid pianist', in *10th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2010)*, pp. 192–197. IEEE, (2010).

[7] George A Bekey, *Autonomous Robots*, MIT Press, Cambridge, MA, 2005.

[8] Mark Bickhard, 'The Emergence of Contentful Experience', in *What Should Be Computed to Understand and Model Brain Function?*, ed., Tadashi Kitamura, World Scientific, (2001).

[9] Mark Bickhard, 'The Dynamic Emergence of Representation', in *Representation in Mind, Volume 1: New Approaches to Mental Representation (Perspectives on Cognitive Science)*, eds., Hugh Clapin, Phillip Staines, and Peter Slezak, Elsevier, (2004).

[10] Margaret A. Boden, *The Creative Mind: Myths and Mechanisms*, Weidenfeld and Nicolson, London, 1990.

[11] Simon Colton and Geraint Anthony Wiggins, 'Computational creativity: The final frontier?', in *Proceedings of 20th European Conference on Artificial Intelligence (ECAI)*, eds., L. De Raedt, C. Bessiere, D. Dubois, P. Doherty, P. Frasconi, F. Heintz, and P. Lucas, pp. 21–26, Montpellier, France, (2012). IOS Press.

[12] Kerstin Dautenhahn, Sarah Woods, Christina Kaouri, Michael L Walters, Kheng Lee Koay, and Iain Werry, 'What is a robot companion-friend, assistant or butler?', in *Intelligent Robots and Systems,*

2005.(IROS 2005). 2005 IEEE/RSJ International Conference on, pp. 1192–1197. IEEE, (2005).

[13] Terrence W Deacon, *Incomplete nature: How mind emerged from matter*, WW Norton & Company, New York, NY, 2011.

[14] Mark D'Inverno and Michael Luck, *Understanding Agent Systems*, Springer Verlag, 2nd edn., 2004.

[15] Homer Dudley and Thomas H Tarnoczy, 'The speaking machine of Wolfgang von Kempelen', *The Journal of the Acoustical Society of America*, **22**(2), 151–166, (1950).

[16] Jerome A Feldman, *From Molecule to Metaphor: A Neural Theory of Language*, MIT Press, Cambridge, MA, 2006.

[17] Petra Gemeinboeck and Rob Saunders, 'Creative Machine Performance: Computational Creativity and Robotic Art', in *Proceedings of the Fourth International Conference on Computational Creativity*, (2013).

[18] Kleomenis Katevas, Patrick G.T. Healey, and Matthew Tobias Harris, 'Robot stand-up: Engineering a comic performance', in *Proceedings of the 2014 Workshop on Humanoid Robots and Creativity at the 2014 IEEE-RAS International Conference on Humanoid Robots (Humanoids 2014)*, Madird, Spain, (2014). Available: $http://cogsci.eecs.qmul.ac.uk/humanoids/Katevasetal\_2014.pdf$.

[19] Stuart Kauffman, *At Home in the Universe: The Search for Laws of Complexity*, Oxford University Press, 1995.

[20] Gary Kemp, 'The Aesthetic Attitude', *British Journal of Aesthetics*, **39**(4), 392–399, (1999).

[21] Heather Knight, 'Eight lessons learned about non-verbal interactions through robot theater', in *Social Robotics*, eds., B. Mutlu, C. Bartneck, J. Ham, V. Evers, and T. Kanda, 42–51, Springer, (2011).

[22] Heather Knight, Scott Satkin, Varun Ramakrishna, and Santosh Divvala, 'A savvy robot standup comic: Online learning through audience tracking', in *International Conference on Tangible and Embedded Interaction*, Funchal, Portugal, (2010).

[23] Heather Knight and Reid Simmons, 'Estimating human interest and attention via gaze analysis', in *2013 IEEE International Conference on-Robotics and Automation (ICRA)*, pp. 4350–4355. IEEE, (2013).

[24] Ruth Millikan, *Language, Thought, and Other Bilogical Categories*, MIT Press, Cambridge, MA, 1984.

[25] Takeshi Mizumoto, Takuma Otsuka, Kazuhiro Nakadai, Toru Takahashi, Kazunori Komatani, Tetsuya Ogata, and Hiroshi G Okuno, 'Human-robot ensemble between robot thereminist and human percussionist using coupled oscillator model', in *, 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1957–1963. IEEE, (2010).

[26] Roger K Moore, 'Presence: A human-inspired architecture for speech-based human-machine interaction', *Computers, IEEE Transactions on*, **56**(9), 1176–1188, (2007).

[27] Roger K Moore, 'From talking and listening robots to intelligent communicative machines', in *Robots that Talk and Listen – Technology and Social Impact*, 317 – 336, De Gruyter, Boston, MA, (2014).

[28] Radhika Nagpal, Attila Kondacs, and Catherine Chang. Programming methodology for biologically-inspired self-assembling systems, 2002.

[29] Stefano Nolfi and Dario Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-organizing Machines*, Cambridge, MA, 2000.

[30] Kohei Ogawa, Koichi Taura, and Hiroshi Ishiguro, 'Possibilities of androids as poetry-reciting agent', in *RO-MAN, 2012 IEEE*, pp. 565–570. IEEE, (2012).

[31] Liviu Panait and Sean Luke, 'Cooperative Multi-Agent Learning: The State of the Art', *Autonomous Agents and Multi-Agent Systems*, **11**(3), 387–434, (2005).

[32] Graeme Ritchie, 'Some Empirical Criteria for Attributing Creativity to a Computer Program', *Minds and Machines*, **17**, 67–99, (2007).

[33] Florian Röhrbein, Sascha Griffiths, and Laura Voss, 'On Industry-Academia Collaborations in Robotics', Technical Report TUM-I1338, Technische Universität München, (2013).

[34] Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, Prentice Hall International, Harlow, third edn., 2013.

[35] Yoav Shoham and Kevin Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge University Press, 2008.

[36] S.Nishio, H.Ishiguro, and N.Hagita, *Geminoid:Teleoperated android of an existing person*, chapter 20, 343–352, I-Tech, Vienna, Austria, 2007.

[37] L. Steels, 'A Self-Organizing Spatial Vocabulary', *Artificial Life*, **2**(3), 319–332, (January 1995).

[38] L. Steels, 'Self-organization and Selection in Cultural Language Evolution', in *Experiments in Cultural Language Evolution*, ed., L Steels, John Benjamins, Amsterdam, (2012).

[39] M. Tomasello, *Origins of Human Communication*, MIT Press, Cambridge, MA, 2008.

[40] Patrick Tresset and Oliver Deussen, 'Artistically skilled embodied agents', in *Proceedings of AISB2014*, Goldsmiths, University of London, (1st - 4th April 2014).

[41] R van Trijp, 'The Evolution of Case Systems for Marking Event Structure', in *Experiments in Cultural Language Evolution*, ed., L Steels, 169–205, John Benjamins, Amsterdam, (2012).

[42] Francisco J. Varela, Evan Thompson, and Eleanor Rosch, *The Embodied Mind*, MIT Press, Cambridge, MA, 1991.

[43] Wolfgang von Kempelen, *Über den Mechanismus der menschlichen Sprache nebst Beschreibung einer sprechenden Maschine*, J.V. Degen, Vienna, 1791. Facsimile re-print from 1970, Stuttgart: Frommann-Holzboog.

[44] Geraint A Wiggins, 'A preliminary framework for description, analysis and comparison of creative systems', *Knowledge-Based Systems*, **19**(7), 449–458, (2006).

[45] L. Wittgenstein, *Philosophische Untersuchungen*, Joachim Schulte Wissenschaftliche Buchgesellschaft, Frankfurt, 2001.

[46] Nick Zangwill, 'Aesthetic Functionalism', in *Aesthetic Concepts: Sibley and After*, eds., Emily Brand and Jerrold Levinson, Oxford University Press, Berlin, (2001).

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.

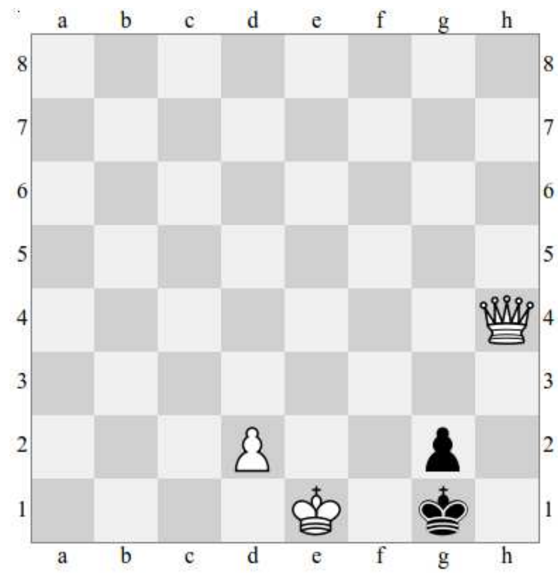

Figure 1. White to play and mate in two moves. From: [8].

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.

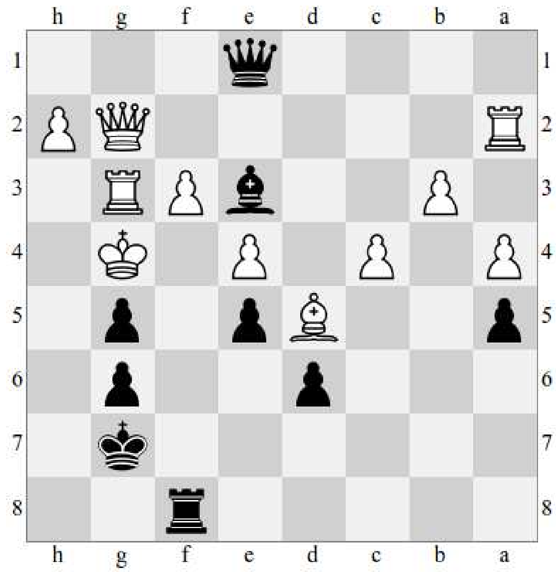

Figure 2. Black to play and mate in three. From: [11].

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).

Towards a Computational Theory of Epistemic Creativity

Jiří Wiedermann[1] and **Jan van Leeuwen**[2]

"The creative act is not an act of creation in the sense of the Old Testament. It does not create something out of nothing: it uncovers, selects, re-shuffles, combines, synthesizes already existing facts, idea, faculties, skills. The more familiar the parts, the more striking the new whole.
A. Koestler [9]

Abstract. We investigate the computational process of creativity from the viewpoint of our recent thesis stating that computation is a process of knowledge generation. Rather than considering the creativity process in its full generality, we restrict ourselves to so-called epistemic creativity which deals with the processes that create knowledge. Within this domain we mainly concentrate on elementary acts of creativity — viz. drawing analogies. In order to do so using the epistemic framework, we define analogies as certain relationships among linguistic expressions and we state what knowledge must be discovered in order to resolve a given incompletely specified analogy. We assume analogies are formed in a natural language and also require that a solution of each analogy must contain an explanation why the resulting analogy holds. Finally, the difference between non-creative and creative computational processes is discussed. Our approach differs from the majority of previous approaches in stressing the knowledge discovery aspects of computational creativity, in requiring explanations in analogy solving and, last but not least, in including theory-less domains serving as knowledge base for knowledge discovery process.

1 INTRODUCTION

Creativity is an activity producing knowledge in the form of ideas, artifacts or behavior that is new for its creator and in some way valuable or important for him or her. Without creativity, no artificial system can aspire to be on par with human intelligence. In its most developed form creativity permeates all human activities. It has been subject of studies in many academic disciplines, among them in psychology, cognitive science, education, philosophy (particularly philosophy of science), technology, theology, sociology, linguistics, economics, and in arts. While all of these disciplines have defined creativity according to their own paradigms and needs, hardly any of them made a serious effort to reveal the underlying mental mechanisms supporting and enabling the process of creativity. This is perhaps due to the fact that the anticipated nature of these mechanisms has been assumed to lay outside of the disciplines at hand. But there is one exception to this rule, and this is the field of artificial intelligence, and especially artificial general intelligence (AGI). Mechanisms of artificial creativity have been intensively studied in cognitive science as well. Due to its omnipresence in many fields of study, the literature concerning creativity is immensely rich and too extensive to be discussed, summarized or referenced fully here.

When inspecting definitions of creativity in whatever discipline, AGI included, two things strike the eye: first, the definitions are very informal, given in a natural language, and second, the definitions hardly ever mention the term knowledge. Especially the latter fact is quite surprising since, perhaps with the exception of artistic creativity, the ability to create new knowledge permeates all domains of creativity. In such domains the primary purpose of creativity is to generate or to demonstrate new knowledge in whatever form — be it conventional knowledge used in everyday life, or scientific knowledge, or a skill, behavior, or a "materialized knowledge" (i.e., knowledge embedded into objects, their functioning, shape or appearance). This kind of creativity is called *epistemic creativity*. Mokyr (cf. [11]) describes it as "actually creating new knowledge or combining existing fragments of knowledge in altogether new ways", as part of his more general view of productive creativity. How can the functioning of epistemic creativity effectively be understood?

It is true that the research field called "knowledge discovery" has become quite popular since the 1990s. Knowledge discovery describes the process of automatically searching large volumes of structured (databases, XML) and unstructured (text, documents, images, multimedia) data for patterns that can be considered knowledge about the data. When compared to what one expects from epistemic creativity, the field of knowledge discovery, despite its name, merely extracts knowledge about the data without having the ambition to create new knowledge other than that which can be straightforwardly extracted from data. This, by the way, can be illustrated by the fact that in the research papers in this field, the word "creativity" is used quite rarely.

It is also true that at the intersection of the fields of artificial intelligence, cognitive psychology, philosophy, and the arts there is a flourishing multidisciplinary endeavor called "computational creativity" (also known as artificial creativity or creative computation). Its roots go back to the nineteen sixties. The field is concerned with theoretical and practical issues in the study of creativity. Here the situation seems to be fairly opposite to the previous case: while the field teems with the word "creative" and all its derivatives, the notion of "knowledge" is much less frequent here. In part this could be due to the fact that the field very often seeks its inspiration in artistic creativity. The field is looking for its theoretical foundations.

In our opinion this frequent overlooking of the connection between (computational) creativity and knowledge generation - where the latter obviously is the main sense of epistemic creativity - may have been caused by an insufficient understanding of what compu-

[1] Institute of Computer Science of CAS and Czech Institute of Informatics, Robotics and Cybernetics of CTU, Prague, Czech Republic email: jiri.wiedermann@cs.cas.cz
[2] Center for Philosophy of Computer Science, Utrecht University, the Netherlands email: J.vanLeeuwen1@uu.nl

tation is. In our recent works [17],[18], [19], [15] we coined the idea that the classical view of computation, based on the ways information is processed by all sorts of machine models (typically by Turing machines), prevents us from clearly seeing the main purpose of computations. The classical view favors the view of HOW computations are performed, instead of WHAT they are doing, i.e. of what is their sense. We hold the view that computation is any process of *knowledge generation*, as we have demonstrated in our previous works. Note that the notion of knowledge generation is machine-independent: we are not interested how, by what means, knowledge is generated, be it in a serial, parallel, interactive, or any other way. What counts is what knowledge is generated.

Changing the view of what computation is may have dramatic consequences. For instance, in the past, various authors have argued that cognition is not computation (cf. [3], [14]), where they have viewed computation in its classical sense, through classical models and scenarios of computations. Under the new view, cognition becomes knowledge generation, and thus, computational, independently of the underlying machine model and computational scenarios. The previous problem vanishes thanks to a new apprehension of computation.

Seeing computations as knowledge generation processes does not automatically turn every computation into a creative process. Intuitively, epistemic creativity requires more than producing knowledge according to some rigid schema (program), counting with some fixed number of alternatives each of which corresponding to a certain prespecified circumstance. For creativity, we require more: new, original alternatives (pieces of knowledge) satisfying as many required constraints as possible must be discovered within the existing knowledge and combined in a novel way under whatever circumstance that cannot be known beforehand. From the candidate alternatives, the one best fitting the constraints must get chosen. This leads to a computational view of epistemic creativity.

The ideas described in the last paragraph answer the often posed question why people have ideas and computers don't. The reason why computers are not creative can have two reasons. The first one is that in the majority of cases when an average person is using a computer, creativity is not required by the application (e.g., in looking for a train schedule). The second answer concerns the so-far quite rare cases where creativity is required — e.g., when consulting symptoms of a disease, or asking for a nice analogy. In such situations a computer will probably not be as creative as we would like to see because it is programmed without understanding how creativity works and what its prerequisites are. Nonetheless, the essence of epistemic creativity has been described in the last two sentences of the previous paragraph. Can we say more about the respective creative processes? Can we be more specific in describing which knowledge generating processes can be seen as creative processes? What are the prerequisites for computational creativity? (Note that we are using the term "computational creativity" in a new, broader sense than mostly used in the eponymous research field.)

In this paper we will answer the last three questions from the epistemic viewpoint of computations. As it turns out, answering the last three questions in their full generality is not easy. Therefore, in what follows we will first investigate but a specific case of creativity. We will concentrate on one of the simplest cases of creativity, and this is analogy solving. Solving an analogy can be seen as an elementary creativity act that calls for discovering and displaying new relations between known pieces of knowledge. Then we will extend our study to a general case of new knowledge discovery.

The structure of the paper is as follows. In Section 2 we present our view of computation as knowledge generation that will offer a unified framework for our further consideration of computations. Special attention is paid to computations in theory-less domains corresponding to natural languages. Section 3 contains the main contribution of the paper. After some preliminaries in Subsection 3.1. analogies and their formal definition in the epistemic framework is presented in Subsection 3.2. The "hard to vary" principle is described, enabling a "quality" judgment of explanatory analogies. In Subsection 3.3. metaphors and allegories as variants of analogies are considered. Subsection 3.4. deals with the efficiency issues in analogy solving. The entire Section 4 is devoted to the general problem of knowledge discovery. Finally, Section 5 contains a general discussion, also paying attention to the difference between creative and non-creative knowledge generation. Conclusions are given in Section 6.

The contribution of the paper to the present state of the art of the theory of computational creativity can be seen in several planes. First, the epistemological view of computations offers a natural unified framework for studying problems related to epistemic creativity. Second, this framework, being machine independent, allows the consideration of theory-less knowledge domains. Third, pertaining to analogy solving, the requirement for a computation to be accompanied by evidence that it works as expected is mirrored in the definition of analogy by a similar demand for analogy explanations. Fourth, explanations attached to each solution of explanatory analogies allows one to judge their explanatory power via the "hard to vary" principle. Finally, our considerations shed further light on the general problem when a computational process is a creative process.

2 COMPUTATION AS KNOWLEDGE GENERATION

Viewing computation as knowledge generation as described in [17], [18] and [19], requires certain ingredients that we first describe informally.

Knowledge in our framework is knowledge in the usual sense of this word. This, of course, does not look like a definition of knowledge, but we need not be very specific. For illustration purposes only, we cite the following definition from Wikipedia: *Knowledge is a familiarity with someone or something, which can include facts, information, descriptions, or skills acquired through experience or education. It can refer to the theoretical or practical understanding of a subject. It can be implicit (as with practical skill or expertise) or explicit (as with the theoretical understanding of a subject); it can be more or less formal or systematic.* Obviously, knowledge according to this definition is observer–dependent.

Any knowledge is a part of a so-called *epistemic domain,* or *domain of discourse,* corresponding to the kind of knowledge we are interested in. Such a domain can be given formally — as in mathematical or logical theories (e.g., theory of recursive functions) or entirely informally, in a natural language, as all sentences describing phenomena in a real world. Intermediate cases (like physical, chemical or biological theories) described in part formally and in part informally are also acceptable. In any case, we must have means to describe the so-called *pieces of knowledge* (e.g., axioms, sentences or formulae in formal theories, or words and linguistic expressions in informal theories described in a natural language).

The final ingredient we require are so-called *inference rules* applicable to the pieces of knowledge in a given domain allowing constructing, generating new pieces of knowledge that will still belong to the domain at hand. Again, in the case of formal theories these rules are also formal rules (like deductive rules in logic), but we also allow entirely informal ones, corresponding to "rational thinking" in the case of informal theories.

The epistemic domain together with the corresponding inference rules form the *epistemic theory*.

Each computation we will consider will generate knowledge from some epistemic domain with the help of the corresponding computational process. We will say that such a computation will be *rooted* in this domain. Starting from the so-called *initial knowledge* the computational process will generate *output knowledge* within the given epistemic domain. Depending on the epistemic domain, initial knowledge is given in the form of axioms, definitions, observations, facts, perceptions, etc. The output knowledge may take the form of propositions, theorems or proofs in the case of formal theories, and statements, hypotheses, scientific laws, or predictions in the case of natural sciences. In the case of informal theories (like theory of mind, arts, etc.) the generated knowledge takes the form of conceptualization, behavior, communication, utterances in a natural language, thinking, and knowledge about the world formed mostly in a natural language or in a form of scientific theories and other writings.

From what has been said above one can see that the epistemic domains range from so-called *theory-full* domains corresponding to formal, abstract theories to *theory-less* domains that admit no formal descriptions for capturing e.g. behavior in common life situations (cf. [13]).

In order for a computation to generate knowledge there must be evidence (e.g., a proof) that explains that the computational process works as expected. Such an evidence must ascertain two facts: (i) that the generated knowledge can be derived within the underlying epistemic theory, and (ii) that the computational process generates the desired knowledge.

The latter is the key to the following more formal definition (cf. [18]). In this definition we assume that the input to a computation is part of both the underlying epistemic domain (and thus of the theory) and the initial data of the computational process. Do not forget that although the notation used in the definition formally resembles the notation used in the formal theories, we will also be using it in the case of informal epistemic domains.

Definition 1 Let T be a theory, let ω be a piece of knowledge serving as the input to a computation, and let $\kappa \in T$ be a piece of knowledge from T denoting the output of a computation. Let Π be a computational process and let E be an explanation. Then we say that process Π acting on input ω generates the piece of knowledge κ if and only if the following two conditions hold:

- $(T, \omega) \vdash \kappa$, i.e., κ is provable within T from ω, and
- E is the (causal) explanation that Π generates κ on input ω.

We say that the 5-tuple $C = (T, \omega, \kappa, \Pi, E)$ is a computation rooted in theory T which on input ω generates knowledge κ using computational process Π with explanation E.

When considering epistemic creativity in the sense of human mental ability, one usually thinks of it in the context of a natural language. How could the corresponding computation (seen as knowledge generation) be captured by the above definitions?

First of all, one must bear in mind that the underlying knowledge domain is a domain comprising, in principle, all human knowledge. This knowledge can be seen as a union of various specific knowledge domains which vary from theory-full to theory-less domains. The respective knowledge is thus *heterogeneous knowledge* and natural language serves as an important, and in fact, the only one known mediator among the respective theories. The less formal the knowledge is the more it relies on the natural language. The "inference rules" for heterogeneous domains are a mix of informal and formal rules. That is, when one speaks within theory-less domains, the informal rules of "rational thinking" are used. Otherwise, speaking within theory-full domains one makes use of the rules corresponding to that domain. Natural language provides not only a tool for initial forming and describing a theory, it also provides a unified tool for understanding all theories and "moving" among them. Last but not least, natural language and its semantics provide a link between a theory and the physical world. Only due to natural language and only within a theory one can explicate meaning of the expressions of a natural language, i.e., their semantics. Namely, in our framework the meaning of any expression of a natural language is given by knowledge pertinent to this expression within a certain domain of discourse. This knowledge comes again in the form of a theory stating all contexts and relationships among them in which the expressions at hand can be used. That is, this theory captures the ways in which usage of an expression makes sense in various contexts. Semantics is knowledge and therefore it can be generated by a computation. From this viewpoint all computations, including the computations that generate knowledge based on understanding natural language, bear a homogeneous structure despite the fact that the underlying knowledge as a whole covers many epistemic domains.

The knowledge framework behind a computation over the domain of a natural language will normally be based on *cooperating theories.* This is an extremely complex system since in principle to each word a theory (in our general sense) is attached, controlling the proper use of this word. In general, such a theory depends not only on the word at hand, but also on the context in which the word is being used. In the case of embodied cognitive systems the context does not only refer to the grammatical context, but also to the entire perceptual situation (cf. [16]). All this leads to a complex intertwining of the respective theories working of the internal models of the world. If realized along the lines sketched above, the underlying cooperative theories should display understanding. The problem of understanding is the central problem of AGI and our approach to computation seems to offer a versatile tool for capturing the related issues. This is because it concentrates on the specification of WHAT the sense of understanding is, while postponing the questions HOW this can be realized. Nevertheless, it is fair to state that so far we do not know much about cooperating theories leading to computational understanding.

Second, what computational process is behind a natural language? It is the process running in our heads. Although we do not know the details of how it works, we do know that it generates knowledge that we can describe by natural language as indicated above. And finally, what corresponds to the explanation? Again, it is an explanation in a natural language.

To summarize, we see that natural language is used here as a means for describing the underlying theory-less domain and the inferences over such domain, as well as for explaining the respective computations as performed by the human brain. Note the analogous situation in classical computing where, for example, λ-calculus is used both as a programming language and as the underlying model of computation.

3 COMPUTATIONAL CREATIVITY

3.1 Preliminaries

Any computation as defined in the previous section generates knowledge. Nonetheless, as remarked in the Introduction, this does not necessarily mean that any computation should be seen as a creative process, as a process that generates something new, original, unexpected, surprising, deserving a special interest or having some worthy value as required in epistemic creativity. This "surprise effect" does not happen when an output of a computation can routinely be produced in a straightforward way, following pre-programmed paths corresponding to a priori envisaged circumstances. The majority of current computer programs works in this way. Typical examples include the computation of a function. Such a process can be seen as generating explicit knowledge (i.e., a function value corresponding to the input value) from implicitly described knowledge that is given in the form of an algorithm. There is no room for creativity in such a process. Note that, e.g., various editors and spreadsheets belong to the category of such computations. Operating systems can serve as an example of an interactive non-creative computational process. What they do can be subsumed as an iteration of the following activity: *"if so and so happens, do so and so"*. In computations of this kind no creativity is assumed, since it is not required by the applications at hand.

What about database searches? Here, pieces of knowledge are sought by searching a finite amount of data ("knowledge items") using a specified criterion. Is here some room for creativity? Now the answer is not so simple as in the previous case. In "old fashioned" databases as used in the early days of computing that used to seek an item satisfying a certain condition within the set of structured data, the situation was similar to the previous case. But think about the following case: a "database" (or rather: a knowledge base) containing all knowledge possessed by an average person (whatever it might mean), i.e., knowledge contained in the mind of that person. The query would be as follows: *"name me an animal living in a desert having the same relation to its living environment as has a shark to the ocean"* (the example taken from [16]). In this case, we can obtain an answer "I don't know" (e.g., from a child), or "a camel" (from an average educated person), or "a desert lion" (from an informed animal rights activist), or even "Cataglyphis bicolor" (a desert-dwelling ant also called "the Sahara desert ant"), from some joking entomologist. Now, were there some aspects of creativity in delivering any od these answers? Which of these answers is the best? And, last but not least, what was the mechanism enabling the answering of such a query?

3.2 Analogies

The last example has been an example of analogy solving. Discussions and studies of analogies go back to the ancient philosophers, since analogies have always played in important role in reasoning in logics, science, law and elsewhere. The role of analogy has been intensively studied for years in cognitive science (cf. [8], [10]). The notion of analogy is rarely formally defined. What one can find in the literature, vocabularies and on the web are informal definitions serving to the purpose of the underlying field. Thus, one can find definitions like *"analogy denotes a similarity between like features of two things, on which a comparison may be based;* or *"a comparison between one thing and another, typically for the purpose of explanation or clarification"*, or *"analogy is a figure of language that expresses a set of like relations among two sets of terms"*. In logic, *"analogy is a form of reasoning in which one thing is inferred to be similar to another thing in a certain respect, on the basis of the known similarity between the things in other respects"*.

There are many variants of analogies. For the purpose of knowledge generation we will be especially interested in so-call *explanatory analogies*. Such analogies create understanding between something unknown by relating it to something known. They provide insight or understanding by relating what one does not know with what one knows. Thus, these analogies may be seen as providing elementary creativity steps in deriving new knowledge. This approach where knowledge is not obtained by simply composing pieces of old knowledge has to be contrasted with the classical epistemological procedures of knowledge generation. Such procedures are usually described as extrapolations of repeated observations, or of known facts, as some variants of an induction process. In this process, there is no creativity aspect: knowledge is merely transformed from one form to an other. However, it is reasonable to expect that the ability to create new knowledge must also include the ability to create new explanations, not merely extrapolating or generalizing the past experience.

In order to better understand explanatory analogies, we will need a more formal definition of analogy that will enable us to see the finer details of the envisaged computational process of creating knowledge leading to analogy solution. Therefore, for our purposes the desired definition should fit into the framework of epistemic computations.

The starting point will be to choose a suitable theory in which the respective computations will be rooted. In this respect, note that all informal definitions of analogies involve direct or indirect reference to natural language. Moreover, they are using linguistic expressions like features, relations, similarity, comparison, or explanations. Therefore, a natural choice for such a theory would be a natural language $\mathcal{NL}$ possessing the richness of linguistic expressions needed to understand and resolve analogies. The (informal) rules corresponding to $\mathcal{NL}$ would be those of "rational thinking", and the corresponding computational process will be that produced by the human brain (cf. Definition 1) and the discussion thereafter.

In the following definition (taken from [16]) the adjective *linguistic* will mean that the corresponding expressions, predicates or relations are not described in any formal logical calculus or theory — rather they are described by expressions of a natural language $\mathcal{NL}$ corresponding to the respective pieces of knowledge. These pieces of knowledge form the *knowledge base* of $\mathcal{NL}$. Their validity usually cannot be proved formally but can be known from experience, empirically or from hearsay.

Definition 2 *Let $S = (s_1, \ldots, s_k)$ and $T = (t_1, \ldots, t_k)$ be two sequences of linguistic expressions from $\mathcal{NL}$. If there exists a linguistic k-ary predicate $P \in \mathcal{NL}$ such that both $P(S)$ and $P(T)$ hold and linguistic relations $R_1, \ldots, R_k \in \mathcal{NL}$ such that $R_i(s_i, t_i)$, for $i = 1, \ldots, k$ holds, then we say that S* is analogous to *T w.r.t. predicate P and relations $R_1, \ldots, R_k$.*

Parameters $s_1, \ldots, s_k$ and $t_1, \ldots, t_k$ are called attributes *of S and T, respectively. Relations R_i's are called* similarity relations*.*

Note that the linguistic expressions, predicates and relations are all described as expressions of a chosen natural language $\mathcal{NL}$.

Definition 3 *Using the notation from Definition 1, given S and T,* analogy solving *is a knowledge generating process whose purpose is to find linguistic predicate P and linguistic relations $R_1, \ldots, R_k$ such that S is analogous to T w.r.t. predicate P and relations $R_1, \ldots, R_k$.*

We say that P is a conjecture *and $P(S)$, $P(T)$ and $R_i(s_i, t_i)$ are the* explanation *of this conjecture.*

To illustrate the use of the introduced formalism, consider again the example from Subsection 3.1. Excerpting from [16]: If $S =$

$(shark, ocean)$ and $T = (camel, desert)$, then we may define predicate $P(x, y)$ as"x lives in y" and R_1 as "both *shark* and *camel* are animals", R_2 as "both *ocean* and *desert* are living environments". Then the claim "x lives in y" is the conjecture and the facts that "*camel* lives in a *desert*" , "*shark* lives in *ocean*" , "both *shark* and *camel* are animals" and "both *ocean* and *desert* are living environments" are its explanation. The previous task is often described as "the relation of *shark* to *ocean* is like the relation of *camel* to *desert*" and abbreviated as *shark : ocean :: camel : desert*.

If all expressions in S are known and only some expressions from T are missing, then S is called the *source* and T is called a *target* of the analogy. Then the whole analogy inclusively of its explanation can be seen as an *explanatory analogy.* The task of finding both the conjecture and its explanation is an act of knowledge discovery. This is because in general the predicates corresponding to the conjecture and the explanations must be discovered among the pieces of knowledge that are at one's disposal.

We have already noted that an explanatory analogy might admit more than one solution. For instance, the solution of analogy *shark : ocean :: ? : desert* could have been either a camel, or a desert lion, or a Sahara desert ant. Under some circumstance, the answer "I don't know" could also be correct. In order to judge the quality and validity of an answer, we must also know the respective explanation. If all explanations are evaluated by an observer as valid, then what answer is the best? In such a case, the best answer would be the one which maximizes the number of relations between the source and the target (i.e., maximizes number k in Definition 2). For instance, in our case, the answer "desert lion" is to be preferred, because in addition to similarity relations R_1 and R_2 it also satisfies relation R_3 "both $shark$ and $desert\ lion$ are predators". The more similarity relations the candidate solution of the incomplete analogy satisfies, the harder it is to come with a different solution. We say that the solution at hand is "hard to vary". According to Deutsch [6], such a solution has a better "explanatory power" than the other competing solutions.

The multitude of answers points to the fact that the answer is observer dependent. The "less knowledgeable" observer might not know about the existence of desert lions and therefore the answer "camel" would sound more plausible to him or her. An observer not knowing any animal living in a desert obviously must answer "I don't know".

3.3 Variants of analogies

Analogies also occur in a number of different forms which can be seen as generalizations or specific cases of our definition of analogy. Let us mention but a few of such instances of analogy.

A more general case is the case of so-called *incomplete analogies*, in which one has to find an analogy between two (or even more) linguistic notions S and T but not all (possible none of the) attributes of neither notion are given. That is, a part of a solution must also be the discovery of the respective attributes of T and S whose pairs correspond to the similarity relations, and the maximization of the number of such pairs. Such problems occur, e.g., in taxonomy dealing with classification of things or concepts based on sharing similar features. In such cases the degree of creativity seems to be higher than in the cases described by Definition 2.

In the opposite direction, a metaphor is a special type of analogy. A *metaphor* is an expression of language that describes a subject by comparing it with another unrelated subject resembling the original subject only in some semantic aspects, on some points of comparison. Both subjects then share the same semantic property which is not immediately apparent from the names of both subjects (cf. the metaphor *"time is money"*) (cf. [10]).

An extended metaphor is *allegory,* in its most general sense. Allegory has been used widely throughout the histories of all forms of art, largely because it readily illustrates complex ideas and concepts in ways that are comprehensible to its viewers, readers, or listeners. Allegories are typically used as literary devices or rhetorical devices that convey hidden meanings through symbolic figures, actions, imagery, and/or events, which together create the moral, spiritual, or political meaning the author wishes to convey (cf. Wikipedia). Re-casting allegory into our framework, allegory usually establishes similarity relations between the narrative story and its possible interpretations in a real or imaginary world. Discovering such relations is a task for allegory creation as well as their projection into the solution of the allegory at hand. The idea is that this projection is not usually obvious at the first sight and its discovery is a task for the observer. In this sense, the similarity relations are "indirectly defined" and depend on the individual taste and knowledge of the observer. Aesthetics and emotions can play an important role in this process. In this way, both creating an allegory by its creator as well as its "deciphering" by an observer are creative acts.

Finally, let us mention the most general and the most important case that plays a crucial role in scientific discovery, and this is the case of the resolution of a "flaw" in a theory. In our framework (cf. Definition 1), the scenario of such a situation is as follows: consider theory T working well over some epistemic domain until one day an input ω to T is found delivering output κ_1. This output is different from output κ_2 which for some reason was expected (e.g. κ_1 disagrees with observations or with experiments). Now the question is, what is the minimal adjustment of theory T such that it would predict output κ_2 on input ω while retaining its ability to work correctly for all other inputs? Clearly, this is another variation on the theme of analogy. This time however, the epistemic theory T itself has become the source and the new theory T' the target of the analogy, and one has to invent new attributes of the target theory preserving as much of the old theory as possible while repairing its flow w.r.t. input ω. Of course, it may happen that theory T is "irreparable" and T' will be completely different from T. History of science knows a lot of such examples (cf. the clash of Darwinism and creationism).

3.4 Efficiency issues in analogy solving

In the framework of epistemic computations one cannot speak about complexity of computations in the classical sense. This is because, in this case, no concrete computational model is used. What can be done for a computation generating complex knowledge, is to describe what partial knowledge or pieces of information are needed in order to generate the knowledge.

Consider the case of solving an explanatory analogy in the form as described by Definition 2 and 3.

The input knowledge for our computation consists of two linguistic predicates $S, T \in \mathcal{NL}$ from a natural language $\mathcal{NL}$, respectively, with $S = (s_1, \ldots, s_k)$ and $T = (t_1, \ldots, t_k)$. Since we are dealing with explanatory analogy we will assume that some (but not all) attributes t_i's in T are left unspecified. Let $\mathbb{K}$ be the knowledge base that is at the disposal of our computation.

In order to resolve such analogy, we need to discover the following knowledge:

1. we have to check whether an object corresponding to predicate S does exist in $\mathbb{K}$. If not, the answer would be *"I don't know"* and we are done.

2. for each object $T' \in \mathbb{K}$ satisfying predicate T in specified attributes, we check whether in $\mathbb{K}$ there exists
   (a) a k-ary predicate P satisfying $P(S) = P(T')$ (i.e., we are looking for a conjecture). If there is no such P the answer would again be *"I don't know"* and we are done.
   (b) next, we look for linguistic relations of similarity $R_1, \ldots, R_k \in \mathbb{K}$ such that $R_i(s_i, t_i)$, for $i = 1, \ldots, k$ holds. If such relations are found then the answer would return object T', conjecture P and explanations R_i's. Otherwise the answer would again be *"I don't know"* and in either case, we are done.

   If no object T' is found then the answer is *"I don't know"* and we are done.

A more involved procedure would be needed in case the necessary knowledge is not found and we don't want to "give it up". If this happens then it is possible to consult "external sources" such as the web, encyclopedias, monographs, experts, etc. In any case, one can see that resolving explanatory analogies is a quite demanding task, requiring in the worst case knowledge of all items in the underlying knowledge base.

Can we say at least something about the computational complexity of solving explanatory analogies? Well, in any case, when we are dealing with analogies over finite knowledge bases, the previous "algorithm" of finding a solution (if it exists) solves in fact a combinatorial problem over a finite domain and therefore can be solved in finite time.

Obviously, the solution of an analogy problem, and in general, of any creativity task depends on all items in the underlying knowledge base. In order to address the essence of the problem of knowledge discovery in terms of the size of the underlying knowledge base we also use a metaphor, viz. the *metaphor of a mosaic.* Namely, a simplistic view of knowledge discovery is that we seek a piece or pieces of knowledge that fit into a certain unfinished mosaic composed from pieces of knowledge possibly from various domains. Here, "fit" means that the new pieces of knowledge are related to the existing pieces by a certain set of known eligible relations that can be either of a syntactic or a semantic nature. (Note that this was also the case of analogies and metaphors.) Then the *creativity problem* is the task of composing a solution of a problem from finitely many pieces of knowledge that have to be related in a logical way in order to come up with the desired solution. It is interesting to observe that a mosaic where only few pieces are missing can be seen as a hypothesis, or a conjecture. In the case of explanatory analogy solving the size of the mosaic is bounded by the number of attributes of both source and target predicate (parameter k in Definition 2). If n is the size of the knowledge base then solving the mosaic problem requires inspection of at most $\binom{n}{k} = O(n^k)$ subsets of the knowledge base. This means that for sufficiently large n and a fixed k the mosaic problem is of polynomial complexity and thus fixed parameter tractable in k.

A problem similar to the creativity problem — the so-called *domino problem* — has been studied in classical complexity theory (based on Turing machines). In 1966, Berger [2] proved that the domino problem is (classically) undecidable if the pieces of knowledge can be used an arbitrarily number of times. The basic idea of the proof is to have a mosaic to encode a halting computation of a Turing machine.

On the one hand, this explains the difficulty of finding new knowledge in general: there is no (Turing machine) algorithm solving such a task. On the other hand, solutions with a small number of pieces are relatively easy to find by a combinatorial search. It is interesting to note that the unrestricted creativity problem seems to be one of the few known undecidable problems of practical significance.

4 Discovering Knowledge

In [5] Barry Cooper asked, whether information can increase in a computation. Indeed, how could a computation produce information which has not already been somehow encoded in the initial data? This does not seem to be possible. An exhaustive answer to this problem has been given by S. Abramsky in [1]. He concludes that, while information is presumably conserved in a total (closed) system, there can be information flow between, and information increase in, subsystems. Note that in our definition of computation we have considered computational processes rooted in the underlying epistemic domains. This can be viewed as though computations are "observing" their "environments" as captured in their knowledge bases, and indeed, some of them even update the underlying knowledge or gain information from cooperating theories under an interactive scenario. In this case it is possible for such a computation to discover new knowledge.

More precisely, it is possible to go beyond the current knowledge explicitly represented in a knowledge base. This can be done by discovering new relationships among the elements of knowledge, or to discover an element or elements of knowledge that satisfy a required relationship to the existing pieces of knowledge, or to gain a new piece of knowledge from "external sources". By "discover" we readily mean to make something explicitly known, i.e., to obtain explicit knowledge of something for the first time. As an example of new knowledge one can take the resolution of a given analogy.

When speaking about creativity in the sense of knowledge generation one must take into account that knowledge can only be generated from knowledge — this is in fact the essence of our definition of computation. Thus, there exist two opposite processes related to knowledge processing: knowledge acquisition, and knowledge generation.

There are many ways of knowledge acquisition: by reason and logic, by scientific method, by trial and error, by algorithm, by experience, by intuition, from authority, by listening to testimony and witness, by observation, by reading, from language, culture, tradition, conversation, etc.

The purpose of *acquisition processes* is to let the information enter into a system and to order it — via computation — into the existing theory or theories over the pertinent knowledge domains (and represent it in a knowledge base). Such domains take various forms of conceptualizations which are part of the respective theory. A *conceptualization* is a simplified, abstract view of the world representing the given knowledge domain. It captures the objects, concepts and other entities and their relationships existing within the knowledge domain at hand (cf. Wikipedia). Obviously, any knowledge acquisition process builds and updates the existing theories.

The purpose of the *knowledge generation process* is to produce knowledge in reaction to the external or internal requests. One can distinguish two basic principles of knowledge generation: syntactic and semantic knowledge mining. Both methods make use of specific inference mechanisms whose purpose is to discover hidden patterns in the data.

Syntactic knowledge mining works solely over the data representing knowledge. It takes into account only the syntax of the respective data, not their meaning, and also the syntactic inference mechanisms of the underlying theory. Syntactic knowledge mining is the computational process of discovering patterns mainly in large data sets involving methods at the intersection of artificial intelligence, machine learning, statistics, and database systems. Finding a pattern corresponding to a certain relationships among data not previously known certainly counts as knowledge discovery.

Semantic knowledge mining is the main engine of creativity. It also looks for hidden patterns in data (knowledge representations) which are semantically, rather than syntactically, related. Usually, based on the semantics of one pattern (the base) and a semantic relation an other pattern (the target), possibly in a different knowledge domain, with a similar semantic structure as the base pattern, is sought satisfying the required relation.

Very often the task of semantic knowledge mining is formed in a natural language. That is, the items to be sought and the relations to be satisfied are described in linguistic terms (as was the case of analogy solving). This complicates the searches since the meaning of linguistic terms must be known. The meaning of each term is, in fact, a piece of knowledge — a theory describing (the properties of) the notion at hand in various contexts. Thus, semantic knowledge mining calls for detecting similarities between theories usually related to different knowledge domains.

Discovering a, in a sense, "parallel" theory to a given one contributes to a better understanding of either theory since it enables to expect relations holding in one theory to also hold in its pendant theory. This is an important element of insight, explanation and understanding. Insight, understanding and explanation make only sense within a theory. They must follow from known facts and rational thoughts.

Unfortunately, general mechanisms of semantic knowledge discovery, explanations and understanding are largely unknown. What we have described here are but the first steps along the respective road.

5 DISCUSSION

We have already remarked at several occasions that, although we see every computation as a knowledge generating process, we cannot automatically consider any computational process to be a creative process. Stating this, when will a computation become a creative process? This seems to be a "million dollar question" of the entire field.

As an example, consider computing x^2, given x. Is this computation a creative process? Compare this with solving the incomplete analogy *shark : ocean :: ? : desert*. Where is the difference? We even know for both cases an algorithm leading to an answer. So where is a difference? Why is finding a solution of the former task considered to be a "non-creative" task whereas finding a solution of the latter is considered to be a creative task?

Well, there seem to be two important differences. In solving the first task one can compute directly with x and manipulate it as the computation requires. In the second case, the computation requires to discover other notions contained in the knowledge base, and the answer depends on what knowledge is stored in the knowledge database at that time. As we have seen, the solution of the second task need not be unique. And the second difference is in the complexity of explanations. While in the first case we must offer an explanation as required by Definition 1, in the second case, in principle, we must offer two explanations: one as asked by Definition 1 related to the correctness of the computation, and the second one required by Definition 2 concerning the correctness of analogy drawing. In general, is there a clear cut between non-creative and creative computational knowledge generation? Nevertheless, the extreme cases can be distinguished.

One might think that there is one more difference. Namely, that in the first case we do not need to know the semantics of x whereas in the second case it is necessary to know the semantics of the "parameters" of the analogy. But this is not true — both computations proceed without knowing the semantics of the respective notions.

Thus, as it appears, in creative knowledge generation (i.e., in computational creativity) the resulting knowledge depends, in addition to the discovery algorithm, on the contents of the underlying knowledge base. The result need not be uniquely defined and in some case need not be defined at all. The respective "creative computation" must work over whatever complete or incomplete knowledge base over the domain of the natural language at hand.

An aspect that seems implicit in "epistemic creativity" is that it isn't driven by the search for a pre-determined answer. In other words, creativity seems be synonymous with "unanticipated solution". In this context the underlying computational process is divergent since it leads to many answers, solutions, knowledge items even from domains that are not internal already but may be imported from elsewhere. Thus creativity seems to involve the generation of options that do not follow by mere deterministic reasoning. (If an artist has found a style that he can repeat, the question is whether it remains a creative process after the first time.)

The bottom line seems that a creative process is not a special type of computation to begin with but a whole collection of computations as also seen from the schema of resolving explanatory analogies in Subsection 3.4. This process is guided at best by some overall triggering process. This latter process may also hold a criterion for judging the computations or rather, the knowledge they come up with, a kind of objective function (which may not be well determined nor a "function" either). In the case of analogies we opted for the "hard to vary" criterion. In any case, this would mean that the question "when is a computation creative" is perhaps not the proper one to ask, if we accept that it is rather a complex process of "divergent computations".

Interestingly, in [12] the author attributes the difference between creative and non-creative *mental processes* not to the underlying computational/functional mechanisms, but rather to the way in which the mental process is experienced. This, however, throws no light on the nature of the underlying mechanisms.

In the context of computational creativity our analysis of analogy solving has revealed that the larger the knowledge base the greater the potential for discovering new knowledge. In order to have the knowledge base as large as possible it must potentially involve all the existing human knowledge and the creative agent must have a command of natural language in order to be able to navigate among various knowledge domains. Along these lines it appears that among the main obstacles of the progress in AGI is our insufficient knowledge of natural language processes concerned with the interactions between computers and human (natural) languages and representation of knowledge accessible to natural languages. Automatic procedures building the respective knowledge bases must be sought (cf. [16]).

Finally, one remark regarding the series of recent writings and interviews of one of the world top thinkers, the prominent British physicist David Deutsch (cf. [7]). At these occasions he has repeatedly stressed that *"The field of artificial (general) intelligence has made no progress because there is an unsolved philosophical problem at its heart: we do not understand how creativity works."*

In spite of what is known about computational creativity (cf. [4]) and despite of the enormous activities in this field, there is something in Deutsch's statement. What is still missing in all known approaches, are the phenomenal issues related to creativity. The “phenomenal component” of creativity seems to be required for a genuine understanding and realization of creative acts. In our approach we have covered up this problem by the requirement of a full mastering of the natural language. This appears to be impossible without engagement with issues around consciousness and free will, and this is why we have stressed the central role of natural language in epistemic creativity processes.

6 CONCLUSION

Our approach is consistent with the modern philosophical view accepted since ancient times that creativity is a form of discovery of new knowledge rather than some kind of inspired guessing. In this discovery process the role of natural language is indispensable since it serves as a universal language bridging various theory-less knowledge domains serving as knowledge base for a knowledge discovery process. Our approach to the problems of computational creativity via the epistemic view of computations offers a natural and uniform framework for the investigation of such problems. Under this view, computational creativity is simply seen as a specific kind of computational knowledge discovery in the underlying knowledge base. The richer the knowledge base the higher the potential for creativity is possessed by the corresponding computations. From this viewpoint, the classical, “non-creative” computational processes are but a special, in a sense “degenerated” kind of computations that do not make use of epistemic theories corresponding to knowledge domains described by explicit knowledge. The epistemic view of computations points to the full capability of computations by revealing their creative potential already in their very definition.

ACKNOWLEDGEMENTS

The research of the first author was partially supported by ICS AS CR fund RVO 67985807 and the Czech National Foundation Grant No. 15-04960S. We thank the anonymous reviewers for their insightful comments that have helped to improve the final version of our paper.

REFERENCES

[1] S. Abramsky. Information, processes and games. In J. van Benthem and P. Adriaans, editors, Handbook of the Philosophy of Information, pages 483549. Elsevier Science Publishers, Amsterdam, 2008

[2] Berger, R.: The undecidability of the Domino Problem, Memoirs of the American Mathematical Society 66, 1966.

[3] Bishop, J.M., (2009), A Cognitive Computing fallacy? Cognition, computations and panpsychism, Cognitive Computing 1:3, pp. 221-233

[4] Boden, M.: The Creative Mind: Myths and Mechanisms (Weidenfeld/Abacus & Basic Books, 1990; 2nd edn. Routldge, 2004)

[5] Cooper, B.: Turing centenary: The incomputable reality Nature 482,465, 23 February 2012

[6] D. Deutsch, *The Beginning of Infinity. Explanations That Transform the World.* Penguin, 2011, 496 pages

[7] D. Deutsch, “Creative Blocks”, AEON Magazine, 02 October 2012

[8] G. Edelman, *Second Nature: Brain Science and Human Knowledge.* Yale University Press, 2006, 224 p.

[9] Koestler, A.: The Act of Creation, Penguin Books, New York, 1964.

[10] G. Lakoff, and M. Johnson, *Metaphors We Live By.* Chicago, IL: The University of Chicago Press, 1980

[11] Mokyr, J.: Mobility, Creativity, and Technological Development: David Hume, Immanuel Kant and the Economic Development of Europe. Session on Creativity and the Economy, German Association of Philosophy, Berlin, Sept. 18, 2005. In G. Abel, ed., Kolloquiumsband of the XX. Deutschen Kongresses fr Philosophie, Berlin 2006, pp. 1131-1161.

[12] Nanay, B.: An experiential account of creativity. In: Elliot Paul and Scott Barry Kaufman (eds.): The Philosophy of Creativity. Oxford: Oxford University Press, 2014, pp. 17-35.

[13] Valiant, L.: Probably Approximately Correct: Nature’s Algorithms for Learning and Prospering in a Complex World, Basic Books, (2013)

[14] van Gelder, T.: What might cognition be, if not computation? The Journal of Philosophy, Vol. 92, No. 7. (Jul., 1995), 345-381.

[15] J. van Leeuwen, J. Wiedermann: Knowledge, representation and the dynamics of computation. To appear in: G. Dodig-Crnkovic, R. Giovagnoli (Eds): Representation and Reality: Humans, Animals and Machines, 2015, Berlin: Springer

[16] Wiedermann, J.: The creativity mechanisms in embodied agents: An explanatory model. In: 2013 IEEE Symposium Series on Computational Intelligence (SSCI), IEEE, 2013, pp. 41-45.

[17] Wiedermann, J. van Leeuwen, J.: Rethinking computation. In: *Proc. 6th AISB Symp. on Computing and Philosophy: The Scandal of Computation - What is Computation?*, AISB Convention 2013 (Exeter, UK), AISB, 2013, pp. 6-10

[18] Wiedermann, J., van Leeuwen, J.: Computation as knowledge generation, with application to the observer-relativity problem. In: *Proc. 7th AISB Symposium on Computing and Philosophy: Is Computation Observer-Relative?*, AISB Convention 2014 (Goldsmiths, University of London), AISB, 2014

[19] Wiedermann, J.,van Leeuwen, J.: What is Computation: An Epistemic Approach. In: G. Italiano *et al.* (Eds), *SOFSEM 2015: Theory and Practice of Computer Science*, Proceedings, Lecture Notes in Computer Science Vol 8939, Springer, 2015, pp. 1-13.