

Application Domains Considered in Computational Creativity

Róisín Loughran and Michael O’Neill

Natural Computing Research and Applications Group (NCRA)
University College Dublin
Ireland
roisin.loughran@ucd.ie

Abstract

We present a review of papers presented at IJWCC and ICCC, specifically considering what applications these papers are engaged with, either directly in generative systems or indirectly in evaluation or framework proposals. The primary focus of this work was to ascertain if there are any trends in the applications considered over the years, any topics that are becoming more dominant or any that have been neglected. Our initial classification among 16 specific categories indicated that Music was the most popular application domain; when we reconsidered seven broader categories we determined that papers involving variations of language processing were most popular. We considered the trend among application domains over the past 12 years and noted that contrary to early discussions on creativity, problems based on logic, science or mathematics do not appear often. We consider the implications of this research as to what information it may convey both to the computational creativity community and to a general computer science audience.

Introduction

A Computationally Creative system is defined as one that can be shown to exhibit behaviour deemed to be creative (Colton, Wiggins, and others 2012). A concise, generalised and context-free meaning of the term *creative* has yet to be defined, however. As such, many scientific studies in the field of Computational Creativity (CC) develop and describe systems that exhibit creativity in a specific application domain. Discussion and evaluation of such systems is then dependent on their ability to function within the given domain. Although some studies within the CC field consider creativity in a more generalised sense with no domain in mind, the majority of papers — even those that are not specifically describing a system designed to produce a single artefact — discuss the merits of the work undertaken in relation to one or more specified applications. This paper takes a quantitative examination of the application domains considered in CC research, specifically from those papers published by the CC community at the main annual events from 2004 to 2016. For this study we consider each paper individually and make a subjective categorisation, rather than using any autonomous, lexical classification techniques.




Figure 1: Number of registrations, papers submitted and papers accepted at IJWCC and ICCC from 2004-2017

CC is a young but expanding field that has been gaining momentum over the last decade. After the International Joint Workshops in Computational Creativity (IJWCC) held as part of larger conferences in 2006 and 2007, the first stand alone workshop was held in 2008. This was further developed into the first International Computational Creativity Conference (ICCC) in 2010 which grew steadily in the following years; ICCC16 had 51 published papers with over 100 registered attendees for the first time. The growth of the event in terms of number of registrations (where available) and number of papers submitted and accepted is illustrated in Figure 1¹. This indicates a general increase in participation and interest in the event over the years. The ICCC conference series is the only scientific conference devoted entirely to all aspects of CC. As such, the proceedings from these conferences offer a comprehensive insight into the work that has been undertaken by leaders of this field. The work submitted, reviewed and presented at these events shapes the field of CC and the direction in which it is going.

This paper considers application domains investigated in this field, not to find the best domain or negate the value of any specific domain but rather to establish if there is a trend

¹Numbers kindly supplied through personal correspondence from members of the Association for Computational Creativity.

in the application domains under consideration by the community, and what — if any — implications this may have for the field as it progresses. CC is still a relatively young field, and yet it encompasses an extremely broad range of topics. As such, it is important to regularly take stock and review any direction it may be taking. The following section offers further discussion on our motivation for this study and what we hope to achieve in undertaking it. The remainder of the paper describes the method by which we categorised the reviewed papers and discusses the results obtained and what conclusions we may draw from them.

Motivation

We conducted the proposed study to review the application domains within which contemporary research on CC is being considered. The term *creative* is one which is colloquially understood and yet inherently difficult to define in a generalised context. In discussions on creativity, either human or machine, there is a natural tendency to use the generation of examples of creative artefacts to demonstrate creative behaviour. Such examples will belong to a specific application domain e.g art, music or literature, but while creativity may be exhibited as such, creativity in general is not specific to any given application, artefact or domain. In humans, achievements in such domains are generally attributed to creative ability — people that are artistic, musical or poetic are often described as creative. Such ideas lend to the notion that any creativity requires special ability, that being creative is only exhibited as an aesthetic talent to be honed and nurtured by the few that display it, rather than an innate ability possessed by us all. Creative behaviour is not limited to remarkable achievements in aesthetic domains however. Personal or P-creativity (as opposed to Historical or H-creativity) is a personal creative ability all people possess, displayed in the generation of an idea that may already exist in history but is new to the individual (Boden 1998). Most people use P-creativity regularly, any time they solve a problem or have a new idea. In studying creativity in a scientific manner, it is this P-creativity that is of most interest, regardless of the domain in which it is demonstrated.

CC is still a young inter-disciplinary field, but as can be seen from an increasing number of publications, it is a field that is growing. Within any growing field, new studies build on research previously undertaken, and early studies are often considered pioneering work in the area. Preliminary studies are continued, developed and used as the basis for more mature studies. As the field develops, ideally we would wish for a balanced approach to the domains being used for applications. If we assume that the presence of creativity is not dependent on the application domain, we must consider that any continued focus towards one domain over another may introduce a bias within the field in general. Furthermore, to encourage development of the field and to attract new researchers and talent to the area we may wish to ensure that new applications are continuously being explored and that the domains considered do not become stagnant.

It is important to note that papers are discussed here *only* in terms of which application they use — even if this application is not the primary focus of the work. Some papers

are application-focussed while others merely use the application in their discussion of a broader creativity issue. We wish to help avoid the automatic use of applications without any more justification than they have been ‘used before’ and thus gather inertia to remain dominant within the field. Such problems have been noted to emerge in evaluation within other applied computing areas, for instance, using ill-suited benchmarks that have been noted to persist in the field of genetic programming (McDermott et al. 2012).

There have been a number of previous studies concerned with the direction and development of the field of CC. A review of the history (and predicted future) of the field was offered in 2009 in the interim between the IJWCC and the start of the ICCC conference (Cardoso, Veale, and Wiggins 2009). This paper considers creativity in general and reviews a number of approaches that have been undertaken in CC along with challenges and progress in the field. It concludes with an optimistic outlook for development within field — one that has been seen to come to fruition with the increasing success of ICCC in the intervening years. Many other survey style papers focus on one aspect of CC, such as problems in evaluation. The difficulty in defining creativity naturally leads to a resultant difficulty in evaluating a creative system. This has led to a number of authors doing self evaluation, minimal evaluation or no evaluations at all on their systems. The lack of evaluation in CC systems has been noted throughout the development of the field (Boden 1998; Cardoso, Veale, and Wiggins 2009; Jordanous 2011). Such studies highlight the need for a clear definition of what can be considered creative.

A method of semi-automated domain conceptualisation on papers from the last six years of ICCC was proposed in (Pollak et al. 2016). The current paper differs in that it offers no automation in the categorisation of applications. We consider each paper individually and determine the domain from the discussion given by the author rather than any information extraction or analysis of the syntax in the papers. The purpose of this work is to consider the domains in which the academic discussion on CC has been undertaken in recent years, specifically by looking at publications at ICCC, and to determine if there are any trends worth noting and what such trends might mean for the overall CC community. This review would be of use to CC researchers concerned with the progress of the field, but also to those new to the area that wish to know what problems have been addressed so far by established researchers. In addition to this methodical review, we wish to reflect on how a focus on domain-based results may be influencing current and future methods of evaluating creativity.

Analysis

In total, 353 papers from 12 years were considered. The number of papers submitted and accepted each year is shown in Figure 1. Table 1 gives the names of sessions that were incorporated in each event. While many of the session names share similarities across years, it is clear there are no ‘standard’ categories that all papers must fall into. For the purpose of this study, we considered each paper individually and assigned it to a specific category.

Table 1: Overview of organisation of papers in each year

Event	Sessions
IJWCC04 Part of ECCBR04	unspecified
IJWCC05 Part of IJCAI04	Mathematical and analogical creativity, Theoretical issues in computational creativity (x2), Creativity in the literary domain, Creativity in the music domain, Creativity in other human activities
IJWCC06 Part of ECAI04	Visual creativity, Musical creativity (x2), Frameworks, Linguistic creativity (x2)
IJWCC07	Creativity in narrative, Analogy and language, Musical creativity, Applied creative systems, Frameworks for creativity
IJWCC08	Theory of creativity, Techniques to get creativity, Storytelling, Music, Platforms and experimental frameworks
ICCC10	Music: patterns and harmony, Visual art, Analogy and metaphor, Stories, Social aspects, Foundations, Music: creation/generation, Creativity support: tools
ICCC11	The Applied, The Social, The Narrative, The Cybernetic, The Foundational, The Helpful, The Cognitive, The Exploratory
ICCC12	Conceptual blending, Analogy, Search, Reflections, Generative systems, Evaluation (x2), Computers being creative, Cognition and computation, Creativity and language
ICCC13	Metaphor in computational creativity, Creativity via computational evolution, Creative processes, Music, Visual art, Computational processes for creativity, Evaluating computational creativity, Poetry, Narrative, Collective and social creativity, Embodied creativity
ICCC14	Co-creation, Visual arts, Videogames, Poetry, Music, Evaluation, Evaluation/Data, Language/Narrative (x2), High level issues
ICCC15	Creative autonomy, Evaluation in the arts, Creative mechanisms, Language, Evaluation of creativity, Musical interaction, Conceptual blending, Visual arts, Games music and cocktails, Creativity support, Imagination and curiosity, Co-creativity, Language
ICCC16	Search, Evaluation, Interaction, Models of creativity, Visual arts, Narratives, Language, Generating structure, Beyond the fence, Blending, Software platforms

Table 2: Description of the 16 initial categories and the higher-level grouping each category was assigned to

Category Name	Description	Higher-level Grouping
Story	story-telling, plot development, character development	NLP
Language	general language syntax, lexicology, translation	NLP
Analogy	analogy and metaphor (text-based)	NLP
Literature	poetry, haiku, sonnet generation or analysis	NLP
Humour	language systems based on understanding or generating humour	NLP
Design	design implementation, description or augmentation	Other
Coding	programming and generating coding solutions	Other
Games	generating, augmenting or playing computer games	Other
Other	any specifically named system not in one of the named categories	Other
Sound	sound generation and analysis, sound effects	Music
Music	music generation, analysis or composition	Music
Maths (and Science)	mathematical formulae, scientific problems, numerical problems, theorems	Logic
Logic	logical problems, general problem-solving	Logic
Image	image generation, analysis or composition	Image
Concept	general high-level concepts (not text-based)	Concept
None	papers that do not discuss any application	None

Categorisation of papers

Categorisation was conducted subjectively by the authors through a review of each paper. This categorisation was performed personally rather than using autonomous, lexical classification to ensure we encapsulated the intended application domain of the author. Using an autonomous, statistical analysis of the papers would likely produce different results, but it was author intent that we found to be more interesting, particularly in view of what this may say about the direction of the the field, as discussed later in the paper.

Many CC papers are based on generative systems. Such systems are trivial to categorise from an application domain perspective; a system that generates paintings is clearly in the domain of visual art or images. A large number of systems are not so specific however. The Call for Papers for these events have always supported the submission of pa-

pers on general creativity, high level concepts or position papers that discuss developments within the field. Such studies often do not specify any application domain. For the purpose of this study we assign these papers to the category of ‘None’. Likewise, some papers mention multiple application domains. Papers based on evaluation of creativity may present results in a number of different domains. In these cases the paper is placed in multiple categories. For example, (Kantosalo et al. 2014) is considered to be in the categories Humour, Choreography, and Design. Of course, this is only the categorisation for the proposed work; the primary focus in this paper is in investigating human-computer co-creation. In this way, the categorisations proposed here do not necessarily correlate with the session organisation as detailed in Table 1. We are purely considering papers from application domains explicitly stated by the authors.

From inspecting each paper in the catalogue we identi-

fied 16 initial individual categories: Logic, Story, Language, Analogy, Sound, Design, Maths, Image, Music, Literature, Concept, Humour, Coding, Games, Other and None. These were chosen as the main topics described explicitly by authors in numerous works. A brief explanation of each category and which papers were included in each is given in Table 2. While most of these are self explanatory as practical applications, one notable exception is the category of ‘Concept’. This category described papers that were not purely positional — they described experiments and offered results — but focussing on higher-level concepts or ideas, rather than a specific physical object as the artefact associated with their work. A number of papers focussed on conceptual blending, such as (Martins et al. 2016), were best categorised as ‘Concept’.

As expected some papers were easily categorised, but many were found to be more difficult to attribute to one individual category. Only papers that explicitly stated more than one application were given multiple categorisations; papers whose application domain was ill-defined or appeared to span multiple domains in one study were subjected to a judgement on our part and assigned to one category. This category was always chosen as that which appeared to be in focus from the authors perspective in discussing their work, rather than making a judgement based on the title, abstract or which session it was included in. For example, (Ventura 2008) could be considered a theoretical or concept-based paper, yet it discusses hypothetical images. Although no images were created by this system, it has been categorised as an ‘Image’ paper as this is the way the paper has been discussed. Other papers arguably could be categorised as either concept or analogy, or possibly story or analogy. Again in such circumstances we categorised in favour of the discussion presented in the individual papers. Certain papers raised severe difficulties in categorisation. (Johnson 2012) mentions nearly all aesthetic fields — music, art etc. yet the overall discussion is mostly concerned with the definition of CC. Arguably, such a paper could be considered to have almost all applications or none. In this case, we have categorised it as ‘None’. A small number of papers required such a subjective categorisation. For complete transparency, a full list of each paper and which category we attributed it to is available in the accompanying appendix: <http://tinyurl.com/lg2aqq4>.

Throughout this discussion, no distinction has been made between long and short papers or those that were presented orally or as posters. However, Demonstrations and Show and Tell sessions were not included in this paper.

Category reduction As described above, these 16 categories were chosen according to the application domain specified by the authors. Many of these initial categories share similar properties and could be amalgamated into broader groupings; there is no one ideal number of categories in such a study. For an alternative level of analysis we reduced the number of categories by grouping together those that could be considered similar. We reduced the 16 sub-categories into the 7 higher-level categories as detailed in the third column of Table 2. As evident from this table many of

the subcategories can be re-categorised as Natural Language Processing (NLP). This covers any application that directly involves text analysis and understanding. Notably we did not consider ‘Concept’ to be part of this grouping as those papers categorised as Concept were not text-based but considered the notion of a concept as a higher level or abstract idea. This greatly reduced number of application categories enables a clearer analysis of the results reported in the following section.

‘Other’ expansion Conversely, the single category of ‘Other’ clearly can refer to a large number of subcategories of applications. Any application that is specifically named but does not belong to one of the categories defined above is considered to belong to Other. It is arguably possible to again consider some of these as sub-categories of more generalised applications described above. For instance, Archaeology is given as the application in one early paper (Cos et al. 2007) and while on inspection this does amount to image analysis, the authors have framed and written the paper from the perspective of Archaeology. Again in cases such as this, where a novel application has been explicitly mentioned by the author, we have chosen this as the given application domain. While presently, these are all categorised as ‘Other’ we consider individual applications and the increase in the use of specific topics in recent years in the results below.




Figure 2: Number of papers in each category, 2004-2016

Results

An overview of the total numbers of papers submitted in each of the original 16 categories summed over all 12 years is shown in Figure 2. This indicates that Music is the most popular single category across all years, followed by Other and Image. The number of papers categorised in to the reduced number of categories, as specified in Table 2 is displayed in Figure 3. It is clear from this figure that when considered with this categorisation, papers based on NLP are actually more popular over the years. This may be unsurprising as the topic of NLP encompasses numerous smaller




Figure 3: Number of papers in each (reduced) category, 2004-2016

yet always popular categories such as Storytelling, Literature (including poetry) and Humour. It is not just papers based on Language whose main application is considered NLP, but those whose application inherently require a semantic understanding such as those in Analogy or Humour.

The percentage of papers in a given category in each year is shown in Figure 4. This shows the trend in applications considered across all years. One point to note from this graph is that the ‘Other’ category has become more prevalent in recent years; this category has had the highest, or second highest, percentage of papers in every year since 2013. We can see from Table 1 that there have been more papers accepted since 2014; as the field has grown and more papers are being written, there are more papers considering new applications. This growing diversity can only be beneficial to the field of CC in general as it indicates new areas of interest, new ideas and new problems being considered.

One surprising result evident from Figures 2 through 4 is the lack of papers based on Logic, Mathematical or Scientific problems. Over all years, studies based on scientific problems or applications have not been popular among CC papers. This is quite surprising when we consider that early discussions on Creativity were often illustrated with scientific, logical or mathematical problems. Much discussion on creativity by Boden is on the scientific and mathematical works of Poincaré, Kekulé and Einstein (Boden 2004) — a point reiterated in the discussion on the development of CC (Cardoso, Veale, and Wiggins 2009). Despite this, papers written by the CC community have focussed on the more traditionally creative or aesthetic applications such as music, art and literature.

Other applications

We have noted that the category of ‘Other’ has become increasingly popular in recent years. This covers the generalisation of topics that have only appeared once such as Cocktail preparation (Pagnutti and Whitehead 2015), to others such as Choreography that may have originally been




Figure 4: Percentage of papers in each category, 2004-2016

considered obscure but have now gathered a following. Choreography in particular has been the main application for six individual papers — four of which were in 2016. We have not explicitly defined a threshold number which must be reached for a topic to become a category, but if choreography remains this popular in the next few years it would rightfully be considered a category of its own. Often it does not take many publications on a topic for it to be considered typical within the field. For instance there were just three individual papers published on recipe creation (Butnariu and Veale 2006; Morris et al. 2012; Shao, Murali, and Sheopuri 2014) before this was chosen as a topic for discussion within a framework paper (Grace and Maher 2015). Many of these Other applications draw from variations or combinations of the categories defined above, but place them in a specific context. For example Internet Memes (Costa, Oliveira, and Pinto 2015) are a specific combination of Humour and Image, or Computer Icons (Confalonieri et al. 2015) are a combination of Image and Design. The most ambitious singular application undertaken to date is surely the Beyond the Fence musical (Colton et al. 2016). The creation and production of this musical was the result of a collaboration from researchers with experience in a wide range of application domains. Large scale projects such as this that may be considered a unique (or ‘Other’) application will undoubtedly always result from a combination of other established application domains.

Papers were considered to be within the Other category when the authors stated the application domain explicitly. A complete list of Other topics that have occurred at least once in this literature includes: role-playing games, map generation (2), puzzles, evolutionary robotics (3), archaeology, cinematography, improvisational theatre, furniture arrangement, chat communication, advertising (2), exploratory gene analytics, play, web-comics, choreography (2), identity structures (2), visual narratives, recipes (3??), subvertising, dementia care, flowcharts (2), modelling, animation (2), interior design, agents in 3D environments, travel, fashion, computer icons (2), 3D vases, cocktails,

kinematics, authoring, scientific discovery, internet meme (2), maze navigation, internet movie database, 3D objects from 2D objects, Beyond the Fence (2)².

Session organisation

The names of the individual sessions taken from Programs and Proceedings at each event is given in Table 1. It is interesting to note that the naming of the organised paper sessions differ from year to year. In some years, the session names have been very application focussed whereas in others they were not. Naturally this is at the discretion of the organisers but this variety in session naming indicates a fluidity within the development and progression of focus in the field. Interestingly, the number of papers related to a given application did not necessarily directly influence the choice of session names. For example, there is a session name focussed on music in each year except 2011, 2012 and 2016 (there are in fact two sessions on music in 2010). However, these years did not lack in papers focussed on musical applications. From Figure 4 we can see that in fact in 2012 and 2016, Music was the second highest application domain represented in accepted papers.

Discussion

We have presented a study focussed on application domains within CC research, but we are not attempting to establish one ‘most creative’ domain. Contrary to colloquial sentiment there is no one domain that is more creative than another; it would be very difficult to determine if an autonomously generated recipe for curry displayed more or less creativity than an autonomously generated piano melody. Absurd as this comparison may seem, if we measure the creativity exhibited by a system purely on the output produced, this type of comparison would become inevitable. Such situations can only be avoided through domain-independent evaluation of the system. Merit should be assigned to progress in the creativity exhibited by a system, rather than to superficial adjustments that merely change the output of a system in a given domain. Systems that produce artefacts requiring more domain knowledge or more complex representation can appear to be more impressive than those that create in simpler domains. An increase in complexity does not necessarily imply an increase in creativity however. It is vitally important when evaluating creative systems that it is the system being evaluated — the processes it undertakes to create, given the domain knowledge that has been presented to it. For evaluation to be domain independent, it must take into consideration all domain knowledge learned by or available to the system, and somehow measure the leap that the system made from this knowledge to what it was able to produce. Given the current state of a system, the representation it uses and the training data, grammars or other a priori information it has access to — what intuitive leap does it make in creating its output? In asking this we may first wish to consider if an autonomous system can actually make an ‘intuitive’ leap, or if it can only

²A graphical display of the spread of these topics was not possible, but those that had multiple instances are shown in parentheses.

be considered intuitive once a person acknowledges it to be so? Furthering this we may have to ask: Can a computer be creative if there is no-one there to call it creative?

The current definition of a CC system is one which ‘exhibits behaviour deemed to be creative...’ thus it is the *behaviour of a system* that needs to be evaluated; the application domain is merely the setting for the experiment. There is a circular, self-referential issue in that CC is defined in terms of creative behaviour, which is often displayed in the creation of artefacts in a given domain before evaluation (of said creativity) inevitably happens in this domain. Hence it can be very problematic to evaluate without considering the application, even though we state that the presence (or level) of creativity is not dependent on the given domain. This entanglement of evaluation and domain knowledge results from the definition of CC, and the definition of creativity in general. As long as CC is defined in terms of ‘behaviour deemed to be creative’ we are relying on an adjudication of actions (behaviour) in comparison to an ill-defined concept (creativity). Without any further specifics we automatically create and evaluate systems in our preferred domain. Should a definition also make some reference to an intuitive leap, or creative step in terms of the abstraction or emergence of a new idea from knowledge already obtained? Even if we did consider incorporating this into a definition — how would one measure such a creative step?

CC research is undertaken within a broad range of subject areas. Attempting to limit this by, for instance, suggesting all research should be conducted only in certain domains would be counter-productive to progress. For an individual researcher to switch application domain may involve a steep learning curve in developing new expertise before any experimental progress could be made. If we state that the presence of creativity is not dependent on working within one application domain then this would appear to be a waste of time, a dismissal of many bodies of work and it would strangle the work of many prominent researchers within the field. Furthermore, as creativity itself remains an ill-defined concept, restricting the areas in which it is studied could hinder development in some unknown way. It is still possible that we could learn more about creativity through one application over another. Hence, we consider it to be beneficial to keep considering more applications rather than less — while focussing evaluation on the system, rather than the product created through representation within the given domain. In this study we only considered papers published at ICCG; many relevant papers have been published elsewhere. Multiple journals have had special issues on CC and there are many other conferences, workshops and journals that look at computational aspects of specific domains such as music, art or design. While a fully comprehensive review of the field would consider all such events, such a review would be infeasible. We have chosen ICCG papers as a representation of the field as a whole.

The field of Artificial Intelligence (AI) has witnessed a similar focus on application-based systems leading to the development of Weak (or narrow) AI instead of Strong (or general) AI. Many high profile successful instances of Weak AI have made headlines in recent years such as Deep Blue

(Campbell, Hoane, and Hsu 2002) or AlphaGo (Silver et al. 2016). These systems are highly impressive in beating world class humans at a specific task and gain media attention and prestige to their programmers. This may benefit the reputation and status of the field in general, yet these systems do not exhibit or possess a General AI that can tackle multiple different problems as a human would. Domain-specific CC systems are falling into the same single-application trap. A music system hailed as ‘creative’ will not recognise the creativity in a joke unless it is also programmed to recognise the humour representation *and* possesses an ability to recognise general creativity — a term we are still struggling with. An ideal general creative system would be able to generate or appreciate a creative act regardless of domain, yet as with general AI, there is no such system at the moment. The comparison between AI and CC is a natural one if we consider creativity as a feature of human intelligence, but CC should not be labelled as mere application within AI. CC remains a field in its own right as long as we ensure that the questions considered are not limited to the description of creative applications. Such a topic is better described as Creative AI, and does tend to involve more aesthetic endeavours such as the generation of art or images. The focus of the field of CC has always been on developing an understanding of what it means to be creative and how we can emulate this creativity autonomously in computational systems.

Future Steps

It was noted above that there is a lack of papers on scientific and logical problems. As so many early studies in creativity did consider such problems to be relevant (or even fundamental) to creative thinking, it would appear that this indicates a potential gap in the field. Addressing this would require some act to entice researchers to undertake research in one of these areas. The proposal of an annual problem-solving competition to coincide with the annual conference could potentially address such an issue. Similar open competitions have been incorporated into other conferences, for instance the ‘Humies’, a human-competitive event held annually at GECCO (Humies 2017). Such a competition should require the development of an autonomous system that solves a specific logical or scientific problem. As the application domain is set, each system would be adjudicated on the creativity it displayed in its approach to the problem. While such a competition may require some organisation, it would encourage development within this domain.

Much early work in creativity was based on symbolic AI, with representations that could be grasped and understood by the user. In more recent years, applied computing research (including CC) has moved towards a more statistical, machine learning approach. The lack of explanation offered by such systems may have influenced the move away from logical problems towards more subjective, aesthetic problems. Even so, we have a responsibility to communicate the possibilities of what can be achieved through CC more clearly to a general computer science and research audience. Personal experience has indicated that many researchers familiar with machine learning or data science are under the assumption that research in CC always involves either music

or art. It is unfortunate that even among experienced, applied computer scientists, the use of the term *Creative* still translates to aesthetic or artistic. The more fundamental meaning of Creativity and the possibilities that can be reached through proper understanding and research needs to be better portrayed as the field develops.

Developing systems that tackle real-world problems could attract more funding for our own areas of research, either through industrial relations or for academic funding proposals. As an applied computer science area we should always bear in mind that good problem solving requires creativity. Real-world problems such as those proposed in (McCaffrey and Spector 2011) or the propositions for managing Dementia Care (Zachos and Maiden 2013) would be of great interest to the public in general. Acquiring knowledge and developing systems to benefit society and the world around us is surely the ultimate goal of any scientific research; arguably we have a moral responsibility to encourage the development of solutions in such areas in any manner possible.

Conclusion

We have presented a review of application domains considered throughout annual CC events over the past 12 years. By concentrating on the applications considered, rather than the overall purpose of the papers, we hoped to gain some insight as to which topical domains are typically used in discussing the subject of CC. This paper focussed entirely on the domains discussed in publications in the field, while simultaneously stating that the presence of creativity is not dependent on any given domain. Although we would like to state that this is because creativity is ‘domain-independent’, at the moment we would just state that this is because the presence of creativity is not determined by the given application domain. We note that papers based on NLP are continuously well-represented across years. Conversely we note a lack of studies based on logical or scientific problems. Tackling scientific, logical or realistic issues could help bring the reputation of CC away from a purely aesthetic domain towards developing solutions for real world problems.

It is difficult at this time to predict how the field will progress in the coming years, but if the current level of growth is to continue, one can assume CC will become increasingly important field within applied computer science. It is imperative that the field remains balanced as it grows and that we remember to reflect on all areas of growth. As a computational field, a number of autonomous systems for analysing papers in the field are emerging such as Dr Inventor (O’Donoghue et al. 2015) that considers the relationships between studies and the system proposed in (Pollak et al. 2016). As such analytical systems are developed, we must ensure to take a step back to consider the implications of the results obtained, what this may tell us about the field and how we can use this information to shape the development of the field as it progresses.

Acknowledgments

We would like to thank members of the ACC for personal correspondence in collecting this data, in particular Raphael

Pérez y Pérez, Dan Ventura, Amilcar Cardoso, Geraint Wiggins, Pablo Gervás, Dragana Miljkovic and Vincent Corruble. This work is part of the App'Ed project funded by Science Foundation Ireland under grant 13/IA/1850.

References

- Boden, M. A. 1998. Creativity and artificial intelligence. *Artificial Intelligence* 103(1):347–356.
- Boden, M. A. 2004. *The creative mind: Myths and mechanisms*. Psychology Press.
- Butnariu, C., and Veale, T. 2006. Lexical combinatorial creativity with “gastonaut”. In *The Third Joint Workshop on Computational Creativity, 17th European Conference on Artificial Intelligence, ECAI. Riva del Garda: Università di Trento*.
- Campbell, M.; Hoane, A. J.; and Hsu, F.-h. 2002. Deep blue. *Artificial intelligence* 134(1):57–83.
- Cardoso, A.; Veale, T.; and Wiggins, G. A. 2009. Converging on the divergent: The history (and future) of the international joint workshops in computational creativity. *AI Magazine* 30(3):15.
- Colton, S.; Llano, M.; Hepworth, R.; Charnley, J.; Gale, C.; Baron, A.; Pachet, F.; Roy, P.; Gervás, P.; Collins, N.; et al. 2016. The beyond the fence musical and computer says show documentary. In *Proceedings of the International Conference on Computational Creativity*.
- Colton, S.; Wiggins, G. A.; et al. 2012. Computational creativity: the final frontier? In *ECAI*, volume 12, 21–26.
- Confalonieri, R.; Corneli, J.; Pease, A.; Plaza, E.; and Schorlemmer, M. 2015. Using argumentation to evaluate concept blends in combinatorial creativity. In *Proceedings of the Sixth International Conference on Computational Creativity*. Brigham Young University.
- Cos, M. A.; y Pérez, R. P.; Altas, C. L.; Hidalgo, M.; Llera, A. A.; Filosóficas, I.; and de México, A. 2007. A generative grammar for pre-columbian artistic production: The case of el tajín style. In *Proceedings of the 4th International Joint Workshop on Computational Creativity*, 39.
- Costa, D.; Oliveira, H. G.; and Pinto, A. M. 2015. ‘In reality there are as many religions as there are papers’ - First steps towards the generation of internet memes. In *Proceedings of the Sixth International Conference on Computational Creativity*, 300.
- Grace, K., and Maher, M. L. 2015. Specific curiosity as a cause and consequence of transformational creativity. In *Proceedings of the International Conference on Computational Creativity June*.
- Humies. 2017. 14th Annual “Humies” Awards for human-competitive results produced by genetic and evolutionary computation. <http://gecco-2017.sigev.org/index.html/Humies>.
- Johnson, C. G. 2012. The creative computer as romantic hero? computational creativity systems and creative personae. In *Proceedings of the third international conference on computational creativity*, 57–61.
- Jordanous, A. 2011. Evaluating evaluation: Assessing progress in computational creativity research. In *Proceedings of the second international conference on computational creativity (ICCC-11). Mexico City, Mexico*, 102–107.
- Kantosalo, A.; Toivanen, J. M.; Xiao, P.; and Toivonen, H. 2014. From isolation to involvement: Adapting machine creativity software to support human-computer co-creation. In *Proceedings of the Fifth International Conference on Computational Creativity*, 1–8.
- Martins, P.; Pollak, S.; Urbancic, T.; and Cardoso, A. 2016. Optimality principles in computational approaches to conceptual blending: Do we need them (at) all? In *Proceedings of the Seventh International Conference on Computational Creativity*.
- McCaffrey, T., and Spector, L. 2011. How the obscure features hypothesis leads to innovation assistant software. In *Proceedings of the second international conference on computational creativity*, volume 1, 120–122.
- McDermott, J.; White, D. R.; Luke, S.; Manzoni, L.; Castelli, M.; Vanneschi, L.; Jaskowski, W.; Krawiec, K.; Harper, R.; De Jong, K.; et al. 2012. Genetic programming needs better benchmarks. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, 791–798. ACM.
- Morris, R. G.; Burton, S. H.; Bodily, P. M.; and Ventura, D. 2012. Soup over bean of pure joy: Culinary ruminations of an artificial chef. In *Proceedings of the 3rd International Conference on Computational Creativity*, 119–125.
- O’Donoghue, D. P.; Abgaz, Y.; Hurley, D.; and Ronzano, F. 2015. Stimulating and simulating creativity with dr inventor. In *Proceedings of the Sixth International Conference on Computational Creativity*. Brigham Young University.
- Pagnutti, J., and Whitehead, J. 2015. Generative mixology: An engine for creating cocktails. In *Proceedings of the Sixth International Conference on Computational Creativity*, 212.
- Pollak, S.; Boshkoska, B. M.; Miljkovic, D.; Wiggins, G. A.; and Lavrac, N. 2016. Computational creativity conceptualisation grounded on iccc papers. In *Proceedings of the Seventh International Conference on Computational Creativity*.
- Shao, N.; Murali, P.; and Sheopuri, A. 2014. New developments in culinary computational creativity. In *Proceedings of the Fifth International Conference on Computational Creativity*.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of go with deep neural networks and tree search. *Nature* 529(7587):484–489.
- Ventura, D. A. 2008. A reductio ad absurdum experiment in sufficiency for evaluating (computational) creative systems. In *Proceedings of the 5th international joint workshop on computational creativity*. Computational Creativity.
- Zachos, K., and Maiden, N. 2013. A computational model of analogical reasoning in dementia care. In *Proceedings of the Fourth International Conference on Computational Creativity*, volume 48.