

University of Dundee

The importance of applying computational creativity to scientific and mathematical domains

Pease, Alison; Colton, Simon; Warburton, Chris; Nathanail, Athanasios; Preda, Irina; Arnold, Daniel

Published in:

Proceedings of the 10th International Conference on Computational Creativity, ICCC 2019

Publication date:

2019

Licence:

CC BY

Document Version

Publisher's PDF, also known as Version of record

[Link to publication in Discovery Research Portal](#)

Citation for published version (APA):

Pease, A., Colton, S., Warburton, C., Nathanail, A., Preda, I., Arnold, D., Winterstein, D., & Cook, M. (2019). The importance of applying computational creativity to scientific and mathematical domains. In K. Grace, M. Cook, D. Ventura, & M. L. Maher (Eds.), *Proceedings of the 10th International Conference on Computational Creativity, ICCC 2019* (pp. 250-257). (Proceedings of the 10th International Conference on Computational Creativity, ICCC 2019). Association for Computational Creativity.


General rights

Copyright and moral rights for the publications made accessible in Discovery Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Proceedings of the 10th International Conference on Computational Creativity



June 17 – 21
Charlotte, North Carolina
USA

Kazjon Grace
Michael Cook
Dan Ventura
Mary Lou Maher



Proceedings of the Tenth International Conference on
Computational Creativity

ICCC 2019

Charlotte, North Carolina, USA | 17 - 21 June

Kazjon Grace, Michael Cook, Dan Ventura, Mary Lou
Maher (Editors)

Published by the Association for Computational Creativity
(ACC)



University of North Carolina at Charlotte
USA

<http://computationalcreativity.net/iccc2019/>

First published 2019

TITLE: Proceedings of the 10th International Conference on Computational Creativity

EDITORS: Kazjon Grace, Michael Cook, Dan Ventura, Mary Lou Maher

ISBN: 978-989-54160-1-1

Published by the Association for Computational Creativity (ACC)

Copyright notice: all contents of these proceedings by ACC, the Association for Computational Creativity, published under a Creative Commons Attribution (CC BY) license <https://creativecommons.org/licenses/by/4.0/>, which allows unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.



Attribution CC BY



The Importance of Applying Computational Creativity to Scientific and Mathematical Domains

Alison Pease¹, Simon Colton², Chris Warburton¹, Athanasios Nathanail³,
Irina Preda¹, Daniel Arnold³, Daniel Winterstein¹ and Mike Cook²

¹University of Dundee, UK

²Queen Mary University of London, UK

³Heriot-Watt, University of Dundee, UK

Abstract

Science and mathematics are currently under-represented in the computational creativity (CC) community. We discuss why the CC community should apply their work to mathematical and scientific domains, and argue that this would be mutually beneficial for the domains in question. We identify a key challenge in Automated Reasoning – that it has not achieved widespread adoption by mathematicians; and one in Automated Scientific Discovery – the need for communicability of automatically generated scientific knowledge. We recommend that CC researchers help to address these two challenges by: (i) applying systems based on cognitive mechanisms to scientific and mathematical domains; (ii) employing experience in building and evaluating interactive systems to this context; and (iii) using expertise in automatically producing framing functionality to enhance the communicability of automatically generated scientific knowledge.

Introduction

Despite the best efforts of successive ICCO organising committees and the computational creativity (CC) community, CC has always attracted substantially more interest from researchers in artistic than scientific and mathematical domains. In their 2017 study of application domains in CC, (Loughran and O'Neill 2017) found that, of 16 categories, papers on Maths, Science and Logic accounted for only 3% of the 353 papers on CC across 12 years. Of course, some work is domain independent, or at least not easily assigned to an academic discipline, such as the body of work on CC and curiosity (for instance (Grace et al. 2017b)). Even taking this into account, it is clear that science and mathematics are vastly under-represented in our community.

There are many reasons why this may be the case. Firstly, AI researchers in scientific domains may well be doing creativity-related work in other contexts but not engaging with the CC community. Automated reasoning (usually deductive reasoning in mathematics) and automated scientific discovery (usually inductive reasoning in a scientific domain) are both thriving subfields of AI, with internationally recognised journals as outlets for publication and engagement; certainly these will contain work relevant to our field but couched in different terminology with different methodologies. Secondly, other priorities in scientific domains may

have led to a focus on techniques such as search, data-mining and automated deduction. Since these generate results of interest to domain experts, the more difficult, fluid and tenuous concept of creativity may be seen as unnecessary, risky or simply not a priority. This may particularly be the case given the various “AI winters” in the twentieth century (the second of which ended in 1993, just six years before the first workshop on CC), and the need for AI to “prove itself” (Crevier 1993). Thirdly, it may be easier to be a hobbyist game designer or artist or composer (many CC researchers are deeply involved in the domains in which they work), than an AI researcher and also an occasional physicist. Fourthly, CC researchers may consider that even if generation is possible within scientific domains, evaluation is too difficult. How we should evaluate our work and our systems has always been a contentious, albeit important, issue in CC, with few proposed evaluation metrics and the majority of researchers still arguing for value along the lines of “we/people liked the system’s output” or “we/people couldn’t distinguish the system’s output from human produced work” (Jordanous 2014). It might be the case that in science, the main evaluation metric – “is it true?”, or “does it work?” – is considered simply too expensive or difficult to demonstrate. Even if evaluation is possible, we may be more prone to dismissing initial results as uninteresting in science than in artistic domains. For instance, we may get a greater sense of progress from a system working in game design which outputs a new (rather basic) game, than one working in geology which outputs a new (rather basic) result.

In this position paper we argue that neglecting scientific and mathematics domains in CC is at best a wasted opportunity, and at worst a significant problem for the field. Deep learning and ML are making inroads everywhere: generative arts, processors, Go, machine vision, and so on, and we need to consider as a community where this leaves us. We believe that it is essential to the health of our field that we reach out as a community at this stage, both to domain experts in science and mathematics and to those in related AI areas. The benefits of doing so will go both ways: we argue that research in CC can help to address key challenges in both Automated Reasoning and Automated Scientific Discovery. As AI is used more and more in science, there is greater dependence on blackbox machine learning systems. While providing greater predictive power, this often comes

at the cost of understanding. We call this the *Understandability Problem* and argue that it will become a big issue in science, which we will have to address. Twenty years of thinking about computational creativity has provided us with valuable tools for addressing these problems. This paper is a call to arms to CC researchers to apply their work to science and mathematics.

A deeper look at the 3% of papers in Maths, Science and Logic (including, for instance, (Juršič et al. 2012; O’Donoghue et al. 2015)) is outwith the scope of this paper. Similarly, we leave aside the question of what creativity or non-creativity looks like in the arts or sciences; for now we simply assume that creative endeavours take place in both domains.

Loughran and O’Neill argue that “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.” (Loughran and O’Neill 2017, p. 7) and that “It is imperative that the field remains balanced as it grows and that we remember to reflect on all areas of growth.” (*Ibid.*). In this paper we support and present further arguments for this position, alongside practical recommendations for doing so.

What are the sciences and arts?

The concept of science is not a straightforward one. The division of the origins of learning and systematic production of new knowledge into disciplines as we know them tends to take into account at least some of the following: methodologies, objects of study (which can be shared with other disciplines), a body of accumulated knowledge (which is generally not shared with other disciplines), theories and concepts, terminology and an institutional manifesto (so that it can reproduce itself) (Krishnan 2009, p. 9). Sciences include *Natural sciences*, which are subdivided into *physical sciences* (chemistry, physics, astronomy), *life sciences*, or *biology* (zoology, botany, ecology, genetics) and *earth science* (geology, oceanography, meteorology, palaeontology); *Social sciences* (psychology, sociology, economics, law, political science); *Formal sciences* (mathematics, logic, theoretical computer science, statistics); and *Applied sciences*, which are subdivided into *engineering* (computer science, civil engineering, electrical engineering, mechanical engineering); *health sciences* (medicine, dentistry, pharmacy) and *agriculture*. The number and variety of sciences makes generalisations difficult, and core values vary accordingly. However, values commonly associated with the (rather unhelpfully named) “hard sciences” include repeatability, reproducibility, predictability, generality and understandability. This last value is particularly cherished: for instance, Roger G. Newton sums it up as “The primary aim of most physical scientists is to understand and explain the workings of Nature.” (Newton 2000, p. 4).

The arts are possibly even harder to define. Indeed, Gallie specifically uses “Art” as an example of an essentially contested concept. This is a concept, the definition of which is “not resolvable by argument of any kind” (Gallie 1955 1956, p. 169). Julie Van Camp, writing in the context of United

	Science	Arts
<i>Aesthetic:</i>	truth	beauty
<i>Approach:</i>	problem-driven	artefact-driven
<i>Task:</i>	analytic	generative
<i>Terminology:</i>	discover	create
<i>Status:</i>	objective	subjective
<i>Goal:</i>	knowledge	self-expression

Table 1: Possible perspectives on scientific and artistic endeavours.

States Congressional policy on arts education, provides the following extensional definition:

The term ‘the arts’ includes, but is not limited to, music (instrumental and vocal), dance, drama, folk art, creative writing, architecture and allied fields, painting, sculpture, photography, graphic and craft arts, industrial design, costume and fashion design, motion pictures, television, radio, film, video, tape and sound recording, the arts related to the presentation, performance, execution, and exhibition of such major art forms, all those traditional arts practiced by the diverse peoples of this country. (*sic*) and the study and application of the arts to the human environment.¹

As a starting point, we could suggest (generalising, controversially) some of the differences between the sciences and the arts as shown in Table 1. In particular, the terminological difference between discovering and creating may explain our field’s current focus on the arts. Of course, the real-world everyday lived experience of *doing* science or *doing* art is far more complex than Table 1 would suggest. Studies of interpretations of seismic data in geology, for instance, show the large number of different expert interpretations of the same seismic section, highlighting the subjectivity involved (Bond et al. 2007). These interpretations are used to analyse subsurface geology, and form the basis for many exploration and extraction decisions. Even in cases where interpreters report that an interpretation is relatively straightforward, there are significant differences in interpretation, leading to significantly different predictions, for instance about gross pore volume or gross rock volume (Rankey 2003). While objectivity may be the goal here, such studies suggest that this aspect of geological practice is closer to visual art interpretation than it is to some other scientific domains.

Similarly, studies of the backstage production of mathematics show that beauty is often a guiding value (Inglis and Aberdeen 2015); there is a high level of disagreement amongst experts about the validity of certain proofs (Inglis et al. 2013); and proofs and theories are often considered to be socially constructed rather than discovered (Lakatos 1976). Less structured knowledge such as our ability to reason logically has been shown to be highly context dependent (for instance, participants in the Wason Selection Task were unable to solve a logical problem at an abstract level

¹http://web.csulb.edu/~jvancamp/361_r8.html

but could solve it correctly when it was framed in a familiar context (Wason and Shapiro 1971)); constructing grounding metaphors to the physical world and abstract linking metaphors argued to be fundamental to our understanding and construction of mathematical knowledge (Lakoff and Núñez 2000); and even the language in which reasoning occurs affecting our preconceptions, perceptions and assumptions (Barton 2009). An analogous story could be told in the arts; for instance, in some contexts paintings are criticised for being beautiful, with the goal being truth, or knowledge (Derrida 1987).

Dibbets expresses the relationship between arts and sciences as follows:

But in the end, we all do very much of the same. All scientists, artists, composers and writers are intensively occupied imagining something that does not yet exist. They find themselves at the borders of areas where up to then hardly anyone found himself, trying to solve problems that are incomprehensible to others, trying to answer questions no one has ever asked. Here, they share a vision on things that are not yet real. (Dibbets 2002, p. 1)

Some of these interdisciplinary features are recognised in curriculum design and teaching featuring transferable skills, in which one skill may be learned within a scientific context and developed or employed in an arts context, or vice versa (see for instance (Gaff and Ratcliff 1996)). Of course, the need for so many interdisciplinary initiatives (and related concepts such as transdisciplinarity, pluridisciplinarity, and multidisciplinary) may suggest that some traditional discipline boundaries are no longer drawn in a helpful way. The evolving role and functionality of AI systems further complicates things. The focus of AI researchers, particularly in machine learning, is often on the skills they hope to simulate rather than a particular domain in which they are usually employed. This may be a more productive approach than the typical CC focus on domain over skill.

Automated Reasoning

Brief history

Automated Reasoning (AR) is a flourishing academic and industry community, with a range of publication venues, including the Journal of Automated Reasoning, the International Joint Conference on Automated Reasoning (IJCAR) and Conference on Automated Deduction (CADE). It has a relatively long history in Artificial Intelligence research: experiments were conducted as early as 1955, with Newell, Shaw and Simons Logic Theorist, which searched forward from axioms to look for proofs of results. Theorem provers HOL, NuPrl and Nqthm, and a variety of other approaches and software tools were in development in the mid-1980's for practical reasoning about programs: (Jones 2003) gives an account of the early history of AR. Notable recent successes include Tom Hales and his team's formalisation of their proof of the Kepler conjecture, using several theorem provers to confirm Hales' major 1998 paper (Hales et al. 2010); and Georges Gonthier and team's 2012 formalisation of the 255 page odd-order theorem (Gonthier 2013) (one of

the most important and longest proofs of 20th century algebra) in the Coq theorem prover.

Key Challenge

While the simulation of mathematical reasoning has been a driving force throughout the history of AI, *it has not achieved widespread adoption by mathematicians*. This is now seen as one of the key challenges in the field. The 2017 and 2019 Big Proof I and II Programmes² included under the Programme Theme description:

The programme is directed at the challenges of bringing proof technology into mainstream mathematical practice.³

and

The scale and sophistication of proof technology is approaching a point where it can effectively aid human mathematical creativity at all levels of expertise. (*Ibid.*)

We can hypothesise many reasons as to why there remains a disconnect between automated and human reasoning. There may be cultural reasons: mathematicians are typically not trained to use Automated Theorem Provers (ATPs), it is not usually part of the undergraduate course or subsequent training and practice. It may simply be the case that perhaps mathematicians that use AR become known as computer scientists (definitions of both of these professions are fluid and somewhat overlapping). Lastly – and this oft-cited reason is our focus here – it may be because current systems cannot do mathematics in the ways that humans do: machine proofs are often considered by mathematicians to be unclear, uninspiring and untrustworthy, as opposed to human proofs which can be deep, elegant and explanatory.

Opportunities for CC

Traditionally there have been two barriers to developing systems which produce “human-like” mathematics: firstly, it is difficult to know what this is; and secondly, it is difficult to automate (Bundy 2011; Gowers 2000). The growing interdisciplinary study of mathematical practice, started by Polya (Polya 1945) and Lakatos (Lakatos 1976), can shed light on the first of these problems. They were early advocates of the (as yet unarticulated) view that it is fruitful to look at what Hersh later termed “backstage mathematics” – the informal workings and conversations about “mathematics in the making” (as opposed to “frontstage mathematics” – textbook or publication-style “finished mathematics”) (Hersh 1991). This rapidly growing body of work is interdisciplinary to varying degrees, bridging mathematics, history, sociology, philosophy, education and cognitive science of mathematics.

²These were hosted at the Isaac Newton Institute (INI) for Mathematical Sciences (2017) and, as a follow on INI satellite event (2019) at the International Centre for Mathematical Sciences in Edinburgh, and organised by some of the most influential people in automated reasoning today: <https://www.newton.ac.uk/event/bpr>, <https://www.newton.ac.uk/event/bprw02>

³<https://www.newton.ac.uk/event/bpr>

Automated Reasoning is largely based on the traditional model of mathematics as a solitary, logic-based endeavour, largely comprising of constructing mathematical proofs. This contrasts with work in the study of mathematical practice, which recognises that mathematics largely takes place in a social context; that it involves “soft” aspects such as creativity, informal argument, error and analogy; and that mathematical knowledge comprises far more than mere proof, including definitions, examples, conjectures explanations, and so on.

Developments in the study of mathematical practice include work on visualisation, such as diagrammatic reasoning in mathematics; analogies, such as between mathematical theories and axiom sets; and mathematical concept development, such as ways to determine potential fruitfulness of rival definitions. Lakoff and colleagues (Lakoff and Núñez 2000) and Barton (Barton 2009) have explored the close connection between language and thought, and shown that images, metaphors and concept-blends used in ordinary language shape mathematical (and all other types of) thinking. At the heart of many of these analyses lies the question of what proof is for, and the recognition that it plays multiple roles; explaining, convincing, evaluating, aiding memory, and so on, complementing or replacing traditional notions of proof as a guarantee of truth). This in turn gives an alternative picture of machines as members of a mathematical community.

These developments present opportunities for researchers in CC which would help to address the second barrier in the “human-like” computing movement – that of difficulty in automation.

Recommendation 1: CC researchers who have developed systems based on cognitive mechanisms, such as concept-blending, analogies and metaphors (eg. (Veale 2012; Li et al. 2012; Baydin, Lopez De Mantaras, and Ontanon 2012; O’Donoghue and Keane 2012)) may consider applying these systems to mathematical domains.

Recommendation 2: CC researchers who have experience in building and evaluating interactive systems which enhance an expert user’s creativity (eg. (Bray and Bown 2016; A. et al. 2014; Karimi et al. 2018)) may consider conducting their work with expert mathematicians. This might, for instance, follow user-centred design, development and testing, and perhaps bridge work between AR and user mathematicians.

Automated Scientific Discovery

Brief history

Whereas AR traditionally has deduction at its heart, Automated Scientific Discovery (ASD) uses induction and abduction to make new taxonomies, laws, theories, models, predictions and explanations. Again, this endeavour started early in the history of AI, with Herbert Simon’s work in 1966 on scientific discovery and problem solving (Simon 1966) and DENDRAL, which used heuristic search to systematically evaluate all of the topologically distinct arrangements of a set of atoms within the rules of chemistry (Lindsay et

al. 1993). The BACON set of programmes were also early examples of ASD, which used rule-based induction to re-discover empirical rules in history of physics and chemistry (Langley 1979). The field became more active in the late 1970s and early 1980s, starting to use machine learning data, and moved from replicating historical events to discovering new ones, including results in astronomy, biology, chemistry, geology, graph theory, and metallurgy (see (Langley 2000) for a review). Today, widespread use of sophisticated machine learning techniques, alongside an explosion of data has led to discoveries such as metallic glass, cost-effective plastic solar cells, and drug discovery.

The use of computational systems to find patterns in scientific data is not without critics. For instance, Genevra Allen highlighted accuracy and reproducibility issues with scientific discoveries made by machine-learning techniques in her recent talk at the 2019 Annual Meeting of the American Association for the Advancement of Science (AAAS).⁴

Key Challenge

Pat Langley – one of the most influential figures in ASD – highlights the need for *communicability of automatically generated scientific knowledge*. In (Langley 2002) he reviews successful and unsuccessful ASD systems, concluding that scientists want interactive, mixed-initiative, rather than fully automated systems. Results must be communicated in language and notation which is familiar to the scientist for collaboration to be successful. He also argues this point in (Džeroski, Langley, and Todorovski 2007):

This emphasis on exchanging results makes it essential that scientific knowledge be *communicable*. (*Ibid.* p2).

Opportunities for CC

The importance and value of narratives that explain, contextualise, comment, and frame generated artefacts for public or expert consumption has been recognised in CC (Charnley, Pease, and Colton 2012), and in a human-only context is known to affect perception of creativity in artistic domains. The CC project to enable creative software to produce its own framing information is in its infancy, but it forms a fundamental part of one of the main evaluation frameworks, as the “F” part of the FACE model (Colton, Pease, and Charnley 2011), and initial work has emerged. This approach goes beyond the explainable AI movement (Došilović, Brčić, and Hlupić 2018), as it aims to show motivation, aesthetic judgements, and so on; telling the story behind the creation of an artefact. We foresee this being an increasingly important area of research in CC, with an increasing level of sophistication: from explanation to justification to argument and dialogue with a user about the value, method of production, motivation etc. behind output. How framing information should be developed is a research programme in its own right. Whilst much attention has been focused on making the outputs of ML systems more accurate and robust, there

⁴https://eurekalert.org/pub_releases/2019-02/ru-cwt021119.php

is also a need for framing information which is more explanatory, more understandable to users and less prone to misinterpretation.

Recommendation 3: CC researchers who are developing systems which can automatically produce framing information (eg. (Grace et al. 2017a; Tomašič, Žnidaršič, and Papa 2014; Colton, Goodwin, and Veale 2012)), may consider applying them to ASD. This may perhaps in collaboration with researchers and existing systems in that field, with a focus on producing useful framing information in a scientific context.

The Understandability Problem in Science and Mathematics

The key challenges that we have identified in both AR and ASD both concern understandability in science and mathematics, and present two different approaches to the problem. Our distinction by domain is, to a certain extent, artificial, and our suggested approaches and recommendations could be used in either domain. In this section we discuss more generally the issue of understanding in science and mathematics, and what that might mean in an AI world.

Understanding in human science and mathematics

Roger Newton’s quote above about the primary aim of physical scientists being to understand and explain nature is uncontroversial, but difficult to unpack. Ever since the entirety of our collective scientific knowledge became too large for a single polymath to comprehend, we have had to outsource our understanding to others. The institutionalised ways in which trust of others’ understanding and progress is handled started with the early universities, and developed with the invention of the printing press, academic journals, the peer review process and so on. Knowledge and understanding is a social process, as argued by (Martin and Pease 2013), but even in the human-only case, this gets complicated. The longest proof in history, of the Classification of Finite Simple Groups (CFSG), is over 10,000 pages, spread across 500 or so journal articles, by over 100 different authors, and took 110 years to complete. What does understanding mean here? Perhaps a handful of people understand the proof in its entirety, and when they die it is not obvious that any one person will ever again understand the entire proof.

In the example of the CFSG, it is considered sufficient that someone once understood the proof. However in the ongoing case of the *abc* conjecture, this is not the case. This conjecture – proposed in 1985, on relationships between prime numbers – is considered to be one of the most important conjectures in number theory (more significant than Fermat’s Last Theorem; in fact Fermat’s Last Theorem would be a corollary of the proof). A proof would be “one of the most astounding achievements of mathematics of the twenty-first century.” (Goldfeld, in (Ball 2012)). In 2012 Shinichi Mochizuki – a mathematician with a good track record, having proved “extremely deep” theorems in the past (Conrad in (Ball 2012)) – produced a 500-page proof. The problem is that the techniques and mathematical objects which Mochizuki has developed to use in his

proof are so new and strange that it would take a reviewer or mathematical colleague most of their career to understand them, before they were able to understand and verify the proof. Despite some efforts from Mochizuki and a handful of his followers to make his work accessible, currently his proof has neither been published nor accepted by mainstream mathematicians, for the simple reason that they don’t understand it.

Crowd-sourced mathematics, in which open conjectures are solved collaboratively via online communities, has been used for around ten years now by a subset of the mathematical community as a new way of producing mathematics through collaboration and sharing (Gowers and Nielsen 2009). Nielsen argues that this has resulted in “amplifying collective intelligence” in his book *Reinventing Discovery* (Nielsen 2011). It has certainly resulted in some original and significant new proofs (for instance, the proofs of the Bounded Gaps Between Primes and the Bounded intervals with many primes, in the 2014 Polymath8 project (eg. (Polymath 2014))). Here it is perfectly possible for a person to be a co-author but not fully understand the proof in their own paper.

Understanding in mixed-initiative science and mathematics

Adding computers to the social process, to form a combination of people, computers, and mathematical archives to create and apply mathematics – a “mathematics social machine” (Martin and Pease 2013) – further complicates matters. Take *automated theorem proving*; the task of deciding whether a given formal statement follows from a given set of premises (Sutcliffe and Suttner 2001). The least informative approach would be to produce merely a “yes”, “no” or “unknown” response. Not only is this devoid of *explanation*, but it also hides the effects of any bugs; requiring the user to either trust the results, or verify the implementation.

This can be mitigated by having the system instead generate a *proof object*: a formal argument for *why* a given statement follows or does not follow. Once generated, a proof object’s validity can be checked without requiring any knowledge of how it was created, thus avoiding the need to trust or verify the (potentially complicated) search and generation procedures. Theorem provers which produce proof objects that are trivial to check by independent *proof checker* programs (which are themselves easily verified, due to their simplicity) satisfy the *de Bruijn criterion* (Barendregt and Wiedijk 2005); examples are Coq (Barras et al. 1997) and Isabelle/HOL (Nipkow, Paulson, and Wenzel 2002).

Proof objects are not a complete solution to understandability, since they can still be quite inscrutable to human users. This often depends on how closely the chosen formal system is able to encode the user’s ideas: for example, the formal proof of the Kepler Conjecture was performed using a system of Higher Order Logic (HOL) (Hales et al. 2015) whose proof objects (natural-deduction style derivations), whilst tedious, are in principle understandable to a user experienced with both the software and problem domain. The same cannot be said of the Boolean Pythagorean Triples problem, a statement of Ramsey theory involving the

structure of the natural numbers. Rather than taking a high-level approach like HOL, (2016) analysed sets $\{0 \dots n\}$ for larger and larger n , encoding these restricted versions of the problem into the language of boolean satisfiability (SAT), and found that the problem is unsatisfiable for $n = 7825$, and hence for the natural numbers as a whole. In this case, the proof object demonstrates this unsatisfiability using 200 terabytes of propositional logic clauses (compressable to 68 gigabytes). Not only is this far too much for any human to comprehend, but the concepts used in the proof (boolean formulae) are several layers removed from the actual problem statement (natural numbers, subsets and pythagorean triples).

Whilst “low level” formalisms like SAT are less understandable or explanatory for users, they are far more amenable to automation than more expressive logics. Despite the proof for the Boolean Pythagorean Triples problem being many orders of magnitude larger than that of the Kepler Conjecture, the latter is well beyond the ability of today’s automated theorem provers due to its encoding in HOL. Instead, it took 22 collaborators 9 years just to formalise the proof (Hales had previously produced a less formal proof, hundreds of pages long and accompanied by unverified software; yet another reminder that human-generated artefacts are not necessarily understandable either).

Forgoing understandability

It may be the case that, given the increase in power, generality and predictiveness that ML approaches give, and the increasing complexity and amount of scientific knowledge, we decide to forgo understandability in science. As a community we would be in a unique position to develop thinking on this, and to answer questions such as whether we should try to replace understandability with something else. We suggest identifying and engaging with stakeholder groups in science and mathematics to ensure that we develop in directions which will be fruitful and useful to society.

Another possible solution to the problems described here would be to forego understandability in the current sense, or rather to change our notion of what *kind* of thing we are aiming to understand. For instance, could a neural network itself be considered to be a scientific discovery, analogous to the discovery of a new plant? It may be that AI systems become objects of study in the same way as the human brain is currently an object of study, with methods and approaches from neural science, psychology, cognitive science and so on employed to understand an AI system and its behaviour and interactions. There is an interesting analogy between ways in which we can “interrogate” a neural network, for instance via generating inputs aligned to deep features (by specifying a deep-level state, then “training” the input to get close), and how we use introspection and analysis to understand human learning. We’re gradually becoming cognitive scientists and psychologists for the robots.⁵ This is already

⁵The term “robopsychologist” was coined by Isaac Asimov in (Asimov 1950) to describe the study of the personalities and behavior of intelligent machines.

an active research area, with (Jonas and K.P. 2017) offering a cautionary tale. Again, as a community we would be in a position to provide a unique perspective on this, having reflected on what constitutes an artefact and how they might be evaluated as novel or significant discoveries.

Concluding Remarks

Recent developments in other areas of AI – principally machine learning (ML) – have led to astonishingly rapid progress in generative processes. Research in Constructive Machine Learning has led to impressive generative results in both the arts and sciences, including painting, music, poetry, gaming, drug design, and gene design – usually in collaboration with domain experts. Our concern is that the sheer size and combined resources of the ML community may render generative work in CC untenable, potentially leading to an identity crisis in the field.

CC has long been seen as more than “mere generation”.⁶ Celebrating and automating other aspects of creative acts in addition to generation – such as making aesthetic judgements, producing framing information (background information about the work) and finding new meta-level processes – is partly what distinguishes us from other AI fields. As generative results in neighbouring areas of AI become more sophisticated, we may wish to focus on these other aspects of the creative act. Extending our repertoire to include more scientific domains will further strengthen our communal identity and enhance our value to other AI researchers and to domain experts in science.

There is also the question of whether CC output might run up against natural boundaries in some areas of the arts. For instance, it is possible that in highly expressive domains, such as poetry, computationally produced poems will not be taken seriously, given the lack of authenticity of life experiences they have. This was discussed in (Colton, Pease, and Saunders 2018), in which the authors argue that a lack of authenticity is a looming issue in CC. Authenticity is not so inherently valuable in the sciences.

Furthermore, it *may* be the case that as the novelty and backstory of computer-generated art wears off, society questions whether we want more computer-produced paintings or poems. The question as to whether we still want more computer-generated science or mathematics, however, seems less likely to be asked: we *always* want more science and mathematics. We suggest this only hesitantly. At the turn of the 20th century, photography caused an explosion in the productivity of art. Flooding the market with images forced artists to redefine their value and led to the creation of modern art, transforming individual self-expression. Subareas of photography have themselves developed as unique art forms, such as wildlife photography and photojournalism. Art has further been transformed through digital technology by filters and editing. We can take inspiration from this: advances in AI can saturate old ways of thinking, but naturally open up new ones. If high quality computationally produced

⁶The informal, tongue in cheek slogan at the 2012 ICCA conference was “scoffing at mere generation for more than a decade.” - although this has been challenged, for instance by (Ventura 2016).

art becomes common-place, art as we know it will be transformed forever: a lot of concepts in art might collapse, but at the same time new concepts which are currently unpredictable might emerge. Of course, applying CC to science may equally saturate certain fields or kinds of work. We raise this here to begin a conversation on where CC in the arts may eventually lead, and as a further potential concern about focusing all our energy on the arts.

People are not naturally good at science. The history of science and scientific methodology, the length of time it takes to train a scientist and the high number of published research findings in science which are considered to be false or sub-standard⁷ all hint at the difficulty of the scientific enterprise. This is partially due to political and institutional factors such as pressure to publish, conflicts of interest and a culture which is often more competitive than collaborative; but also partially due to the constant battle to avoid the large number of cognitive biases that adversely affect our reasoning and judgements (Haselton, Nettle, and Andrews 2005; Sutherland 2013). On the other hand, the arts – while also difficult to do well – do not usually go against our natural way of thinking, and can be seen as a celebration of our humanity. In many ways science should be an obvious application domain for computational creativity. This paper is a call to arms for the whole CC community: to apply systems based on cognitive mechanisms to scientific and mathematical domains; to employ experience in building and evaluating interactive systems to this context; and to use expertise in automatically producing framing functionality to enhance the communicability of automatically generated scientific knowledge.

Acknowledgements

We are very grateful to the thoughtful reviews and meta-review. We gratefully acknowledge financial support from EPSRC Grant EP/P017320/1. The fourth author is supported by the NERC National Productivity Investment Fund (NPIF) grant, ref. no. NE/R01051X/1 and the James Watt Scholarship scheme at Heriot-Watt University. The eighth author is supported by the Royal Academy of Engineering under the Research Fellowship scheme.

References

A., K.; Toivanen, J. M.; Xiao, P.; and Toivonen, H. 2014. From isolation to involvement: Adapting machine creativity software to support human-computer co-creation. In *Proc. of the 5th ICCS*.

Asimov, I. 1950. *I, Robot*. Gnome Press.

Baker, M. 2016. 1,500 scientists lift the lid on reproducibility. *Nature* (533):452–454.

Ball, P. 2012. Proof claimed for deep connection between primes. *Nature* 489(7414).

Barendregt, H., and Wiedijk, F. 2005. The challenge of computer mathematics. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 363(1835):2351–2375.

Barras, B.; Boutin, S.; Cornes, C.; Courant, J.; Filliatre, J.-C.; Gimenez, E.; Herbelin, H.; Huet, G.; Munoz, C.; Murthy, C.; et al. 1997. *The Coq proof assistant reference manual: Version 6.1*. Ph.D. Dissertation, INRIA.

Barton, B. 2009. *The Language of Mathematics: Telling Mathematical Tales*. Mathematics Education Library, Vol. 46. Springer.

Baydin, A. G.; Lopez De Mantaras, R.; and Ontanon, S. 2012. Automated generation of cross-domain analogies via evolutionary computation. In *Proc. of the 3rd ICCS*.

Bond, C. E.; Gibbs, A.; Shipton, Z.; and Jones, S. 2007. What do you think this is? “Conceptual uncertainty” in geoscience interpretation. *GSA Today* 17(11):4–10.

Bray, L., and Bown, O. 2016. Applying core interaction design principles to computational creativity. In *Proc. of the 7th ICCS*.

Bundy, A. 2011. Automated theorem provers: a practical tool? *Ann Math Artif Intell* 61:3–14.

Charnley, J.; Pease, A.; and Colton, S. 2012. On the notion of framing in computational creativity. In *Proc. of the 3rd ICCS*.

Colton, S.; Goodwin, J.; and Veale, T. 2012. Full-face poetry generation. In *Proc. of the 3rd ICCS*.

Colton, S.; Pease, A.; and Charnley, J. 2011. Computational creativity theory: The FACE and IDEA descriptive models. In *Proc. of the 2nd ICCS*.

Colton, S.; Pease, A.; and Saunders, R. 2018. Issues of authenticity in autonomously creative systems. In *Proc. of the 9th ICCS*.

Crevier, D. 1993. *AI: The Tumultuous Search for Artificial Intelligence*. New York, NY: BasicBooks.

Derrida, J. 1987. *The Truth in Painting*. Univ. of Chicago Press.

Dibbets, J. 2002. Interactions between science and art. *Cardiovascular Research* 56(3):330–331.

Došilović, F. K.; Brčić, M.; and Hlupić, N. 2018. Explainable artificial intelligence: A survey. In *41st International convention on information and communication technology, electronics and microelectronics (MIPRO)*, 0210–0215. IEEE.

Džeroski, S.; Langley, P.; and Todorovski, L. 2007. Computational discovery of scientific knowledge. *Computational Discovery: LNAI* 4660:1–14.

Gaff, J. G., and Ratcliff, J. L. E. 1996. *Handbook of the Undergraduate Curriculum A Comprehensive Guide to Purposes, Structures, Practices, and Change - The Jossey-Bass Higher and Adult Education Series*. Wiley.

Gallie, W. B. 1955 – 1956. Essentially contested concepts. *Proceedings of the Aristotelian Society* 56:167–198.

Gonthier, G. 2013. Engineering mathematics: the odd order theorem proof. *POPL* Roberto Giacobazzi and Radhia Cousot (eds):pp. 1–2.

Gowers, T., and Nielsen, M. 2009. Massively collaborative mathematics. *Nature* 461(7266):879–881.

Gowers, T. 2000. Rough structure and classification. *GAFA (Geometric And Functional Analysis)* Special volume – GAFA2000.

Grace, K.; Maher, M. L.; Mohseni, M.; and Perez Y Perez, R. 2017a. Encouraging p-creative behaviour with computational curiosity. In *Proc. of the 8th ICCS*.

Grace, K.; Maher, M.; Mohseni, M.; and Pérez y Pérez, R. 2017b. Encouraging p-creative behaviour with computational curiosity. In *Proc. of the 8th ICCS*.

⁷Meta-scientific studies suggest that 85% of biomedical research efforts are wasted (Macleod et al. 2014) and 90% of respondents to a recent survey in Nature agreed that there is a ‘reproducibility crisis’ (Baker 2016) (see (Munafò et al. 2017; Ioannidis 2005) for further details).

- Hales, T. C.; Harrison, J.; McLaughlin, S.; Nipkow, T.; Obua, S.; and Zumdeller, R. 2010. A revision of the proof of the Kepler conjecture. *Discrete & Computational Geometry* 44(1):1–34.
- Hales, T.; Adams, M.; Bauer, G.; Dang, D. T.; Harrison, J.; Hoang, T. L.; Kaliszyk, C.; Magron, V.; McLaughlin, S.; Nguyen, T. T.; et al. 2015. A formal proof of the Kepler conjecture. *arXiv preprint arXiv:1501.02155*.
- Haselton, M. G.; Nettle, D.; and Andrews, P. W. 2005. The evolution of cognitive bias. In Buss, D. M., ed., *The Handbook of Evolutionary Psychology*. Hoboken, NJ, US: John Wiley and Sons Inc. 724–746.
- Hersh, R. 1991. Mathematics has a front and a back. *Synthese* 88(2):127–133.
- Heule, M. J.; Kullmann, O.; and Marek, V. W. 2016. Solving and verifying the boolean pythagorean triples problem via cube-and-conquer. In *International Conference on Theory and Applications of Satisfiability Testing*, 228–245. Springer.
- Inglis, M., and Aberdein, A. 2015. Beauty is not simplicity: An analysis of mathematicians’ proof appraisals. *Philosophia Mathematica* 23(1):87–109.
- Inglis, M.; Pablo, J.; Mejia-Ramos; Weber, K.; and Alcock, L. 2013. On mathematicians’ different standards when evaluating elementary proofs. *Topics in Cognitive Science* 5(2):270–282.
- Ioannidis, J. P. A. 2005. Why most published research findings are false. *PLOS Medicine* 2(8).
- Jonas, E., and K.P., K. 2017. Could a neuroscientist understand a microprocessor? *PLoS Comput Biol* 13(1).
- Jones, C. B. 2003. The early search for tractable ways of reasoning about programs. *IEEE Annals of the History of Computing* 25(2):26–49.
- Jordanous, A. 2014. Stepping back to progress forwards: Setting standards for meta-evaluation of computational creativity. In *Proc. of the 5th ICCS*.
- Juršič, M.; Cestnik, B.; Urbančič, T.; and Lavrač, N. 2012. Cross-domain literature mining: Finding bridging concepts with cross-bee. In *Proc. of the 3rd ICCS*, pp. 33–40.
- Karimi, P.; Grace, K.; Maher, M.; and Davis, N. 2018. Evaluating creativity in computational co-creative systems. In *Proc. of the 9th ICCS*.
- Krishnan, A. 2009. What are academic disciplines? Some observations on the disciplinarity vs. interdisciplinarity debate. Technical Report 03/09, National Centre for Research Methods.
- Lakatos, I. 1976. *Proofs and Refutations*. Cambridge, UK: CUP.
- Lakoff, G., and Núñez, R. 2000. *Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being*. New York: Basic Books.
- Langley, P. 1979. Rediscovering physics with bacon. In *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pp. 505–507. Tokyo: Morgan Kaufmann.
- Langley, P. 2000. The computational support of scientific discovery. *International Journal of Human-Computer Studies* 53:1149–1164.
- Langley, P. 2002. Lessons for the computational discovery of scientific knowledge. In *Proceedings of First International Workshop on Data Mining Lessons Learned*, 9–12. Sydney: University of New South Wales.
- Li, B.; Zook, A.; Davis, N.; and Riedl, M. 2012. Goal-driven conceptual blending: A computational approach for creativity. In *Proc. of the 3rd ICCS*.
- Lindsay, R. K.; Buchanan, B. G.; Feigenbaum, E. A.; and Lederberg, J. 1993. Dendral: A case study of the first expert system for scientific hypothesis formation. *Artificial Intelligence* 61(2):209–261.
- Loughran, R., and O’Neill, M. 2017. Application domains considered in computational creativity. In *Proc. of the 8th ICCS*, 197–204.
- Macleod, M. R.; Michie, S.; Roberts, I.; Dirnagl, U.; Chalmers, I.; Ioannidis, J.; Al-Shahi Salman, R.; Chan, A.; and Glasziou, P. 2014. Biomedical research: increasing value, reducing waste. *Lancet* (383):101–104.
- Martin, U., and Pease, A. 2013. Mathematical practice, crowd-sourcing, and social machines. In Carette, J.; Aspinall, D.; Lange, C.; Sojka, P.; and Windsteiger, W., eds., *Intelligent Computer Mathematics*, volume 7961 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg. 98–119.
- Munafò, M. R.; Nosek, B. A.; Bishop, D. V. M.; Button, K. S.; Chambers, C. D.; Percie du Sert, N.; Simonsohn, U.; Wagenmakers, E.-J.; Ware, J. J.; and Ioannidis, J. P. A. 2017. A manifesto for reproducible science. *Nature Human Behaviour* 1(1):0021+.
- Newton, R. G. 2000. *The Truth of Science: Physical Theories and Reality*. Cambridge, Massachusetts: Harvard University Press.
- Nielsen, M. 2011. *Reinventing Discovery: The New Era of Networked Science*. USA: Princeton University Press.
- Nipkow, T.; Paulson, L. C.; and Wenzel, M. 2002. *Isabelle/HOL: a proof assistant for higher-order logic*, volume 2283. Springer Science & Business Media.
- O’Donoghue, D., and Keane, M. T. 2012. A creative analogy machine: Results and challenges. In *Proc. of the 3rd ICCS*.
- O’Donoghue, D. P.; Abgaz, Y.; Hurley, D.; Ronzano, F.; and Sagion, H. 2015. Stimulating and simulating creativity with dr invento. In *Proc. of the 6th ICCS*, pp. 220–227.
- Polya, G. 1945. *How to solve it*. Princeton University Press.
- Polymath, D. 2014. Variants of the selberg sieve, and bounded intervals containing many primes. *Research in the Mathematical Sciences* 1(12).
- Rankey, E. C. 2003. Interpreter’s corner – that’s why it’s called interpretation: Impact of horizon uncertainty on seismic attribute analysis. *The Leading Edge* 22(9).
- Simon, H. A. 1966. Scientific discovery and the psychology of problem solving. In (Ed.), R. C. C., ed., *Mind and cosmos: Essays in contemporary science and philosophy*. University of Pittsburgh Press.
- Sutcliffe, G., and Suttner, C. 2001. Evaluating general purpose automated theorem proving systems. *Artificial intelligence* 131(1):39–54.
- Sutherland, S. 2013. *Irrationality: The enemy within*. Pinter and Martin Ltd.
- Tomašič, P.; Žnidaršič, M.; and Papa, G. 2014. Arts, news, and poetry – the art of framing. In *Proc. of the 5th ICCS*.
- Veale, T. 2012. From conceptual “mash-ups” to “bad-ass” blends: A robust computational model of conceptual blending. In *Proc. of the 3rd ICCS*, pp. 1–8.
- Ventura, D. 2016. Mere generation: Essential barometer or dated concept? In *Proc. of the 7th ICCS*.
- Wason, P. C., and Shapiro, D. 1971. Natural and contrived experience in a reasoning problem. *Quarterly Journal of Experimental Psychology* 23:63–71.