Towards Ontology Use, Re-use and Abuse in a Computational Creativity Collective (A Position Statement)

Simon Colton
Computational Creativity Group
Department of Computing, Imperial College, London
http://www.doc.ic.ac.uk/ccg

Computational creativity is broadly defined as the study of building software which exhibits behaviour that would be deemed creative if exhibited by a person. In more practical terms, we investigate how to engineer software that takes on some of the creative responsibility in arts and science projects which produce culturally interesting *artefacts* such as poems, theorems, paintings, melodies, etc. To this end, there are numerous examples of creative software being employed in musical composition, visual arts, pure mathematics, natural language generation, scientific discovery, video game design, and many more areas of discourse. Moreover, the computational creativity community is beginning to come to consensus on some of the thorny research questions that have arisen, such as: which AI processes are more suited to generative applications; how can we measure levels of creativity in software; and what roles can software have in creative acts? Our contributions to computational creativity research have revolved around our two pieces of research software: the HR system [2] and The Painting Fool (www.thepaintingfool.com). The former is mathematical theory formation software which can start with the bare minimum about a domain of pure mathematics, such as how to divide one number by another, and end with a rich theory of concepts, conjectures, theorems and proofs. The latter is an automated painter which we hope will one day be taken seriously as a creative artist in its own right.

The majority of software developed by computational creativity researchers – including our own – is given domain knowledge only about its specific area of application. For instance, our HR software is given enough background information about domains of pure mathematics to enable it to invent concepts in those particular domains, but it is not given wider mathematical knowledge and is certainly not provided with information outside the sphere of pure mathematics. This is largely acceptable in domains where there are objective measures of value with which we can assess the artefacts produced by the creative systems. However, we argue in [3] that in certain domains (most noticeably the visual arts), the creativity and intelligence of the creator is taken into account when assessing the value of the artefacts that he/she/it produces. In particular, in such domains, the cultural awareness of the artist may well be questioned when people assess the value of their work. In these situations, there is much need for the kind of knowledge stored in ontologies, where creative software can use these knowledge bases to enhance its intelligence, in order to produce more interesting and culturally valuable artefacts. As

we argue later, in addition to the *use* of such ontologies, we will also need to *re-use* the knowledge in one ontology for situations that it was not conceived for, and also to *abuse* the ontologies, by using them for rather more imaginative purposes.

We are currently engaged in a project to build a *computational creativity collective* of software, knowledge and resources (www.doc.ic.ac.uk/ccg/collective). The purpose of the collective will be to bring together software which is able to generate different types of artefacts (e.g., poems, pictures, melodies, theorems, etc.) in such a way that the output from one process can become the stimulus for another process. In this way, we hope to show that such software – when used in combinations – can produce culturally relevant artistic and scientific artefacts of real value to various communities. There will be (at least) four arms to the collective. The first arm will contain the generative and analytical algorithms for producing artefacts. Currently, the collective has more than 100 processes involving image and text manipulation, and the retrieval of web material. In addition, it has more than 50 *mashups* which combine two or more of these processes. For instance, one mashup in the collective retrieves the most current tweets from a twitter user, extracts keywords from the tweet, finds images from Google Images which match the keywords, and puts the images together into a visual Collage. This kind of text-to-art mashup follows on from the project discussed in [6]. The second arm of the collective will be a set of reasoning methods which are able to fruitfully combine the processes in the collective automatically. The first instantiation of such a combined reasoning process will be based on global workspace architectures, in particular, the work described in [1]. The third arm of the collective will be source materials (images, sounds, music, text, etc.) which can be accessed by the processes in the production of the artefacts. Naturally, any artefacts which are produced by the collective will be added to this source material.

The final arm of the collective will be a knowledge base of information which can be queried by processes in the collective, and a set of techniques for querying, combining and possibly abusing the knowledge base in order to add more intelligence and imagination to the way the processes in the collective operate. We envisage the use of ontologies to represent the information in the knowledge base, and we hope that the creative usage of ontologies will lead to interesting developments in the theory of modular ontology design. Researchers interested in computational creativity are already turning to ontologies both to test their methods on, and as sources of information. For instance, the work of Veale et. al. investigates analogical reasoning for ontology construction in [9], and their work regularly uses ontologies such as WordNet in creative fashions, e.g., for the creation of neologisms [8]. We can also claim that the automated theory formation routine in HR is a method for ontology construction, whereby the concepts in the ontology are invented by the software. A third example of ontology construction methods might be found in conceptual blending [5], which is also regularly used in computational creativity settings, for instance in the projects described in [7]. We envisage that these kinds of knowledge generation engines will form an important part of the way in which the collective is used for creative artefact generation.

We adhere to the maxim that *Good Art Makes You Think*, as espoused in [4]. Unfortunately, the majority of artworks generated entirely by computer are rather decorative in nature, hence are unlikely to be classed as culturally important works. In contrast, we aim for the artefacts produced by the collective to be designed to encourage the consumer of them to engage their mental faculties when interpreting the pieces. For instance, in the work described in [6], entirely independently of any input from ourselves, the software

downloaded a news story about the war in Afghanistan from the UK's Guardian newspaper's website, and in response, produced a collage of paintings including a bomber plane, an explosion, a family and a girl in ethnic headwear, which forces the viewer to think about the issues raised in the news story. This highlights a rather direct way in which computer generated art can make people think. There are a number of other ways in which this can be achieved, e.g., through metaphor, disguise, shock, humour, satire, etc. Each of these methods will need information about how the world fits together, in order to present a coherent artefact of value to the consumers. This is where we envision the usage of ontological knowledge bases to be paramount in order for the software in the collective to appreciate the nature of the artefacts it produces.

We also believe that the abuse of ontologies will be fruitful for creative purposes. In particular, we are interested in applying ontologies in the exact opposite way to how they are normally used: in order to *create* rather than resolve ambiguity and to suggest how the world *doesn't* work, rather than explain how it does. Creating ambiguity is an essential part of the artistic process, where artefacts are designed to be interpreted by people, hence cannot be too literal (naïve) in conception. The fact that a piece of art is interpreted in often very different ways by different people often adds much to the value of the artwork. We believe that looking at how ontologies have been designed to minimise ambiguity may help to guide artefact consumers to – eventually – interpret artworks in appropriate and enlightening ways. In addition, we envisage working on ways in which to use common sense knowledge, and possibly more focused domain knowledge, to determine what isn't true, but could be, and what is entirely fantastical, via the analysis of ontologies of real world knowledge. Such imaginative behaviour is central to creative processes in people and computers. We plan to further research the possibilities for the productive **use** of ontologies to make our creative software more culturally aware; the productive **re-use** of ontologies for new domains, in the pursuit of analogical reasoning, the construction of metaphors, etc.; and the productive **abuse** of ontologies for imaginative and playful purposes. To this end, we look forward to working closely with the various communities working in ontology design.

References

[1] J Charnley. *A Global Workspaces Approach to Combining Reasoning Systems*. PhD thesis, Department of Computing, Imperial College, London, 2010.

[2] S Colton. *Automated Theory Formation in Pure Mathematics*. Springer-Verlag, 2002.

[3] S Colton. Creativity versus the perception of creativity in computational systems. In *Proceedings of the AAAI Spring Symposium on Creative Systems*, 2008.

[4] S Colton. Seven catchy phrases for computational creativity research. In *Proceedings of the Dagstuhl Symposium on Computational Creativity*, 2009.

[5] G Fauconnier and M Turner. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. Basic Books, 2002.

[6] A Krzeczkowska, J El-Hage, S Colton, and S Clark. Automated collage generation – with intent. In *Proceedings of the First International Joint Conference on Computational Creativity*, 2010.

[7] F Pereira. *Creativity and Artificial Intelligence: A Conceptual Blending Approach*. Mouton de Gruyer, 2007.

[8] T Veale. Tracking the lexical zeitgeist with wordnet and wikipedia. In *Proceedings of the 17th European Conference on Artificial Intelligence*, 2006.

[9] T Veale and Guofu Li. Ontological cliques - analogy as an organizing principle in ontology construction. In *Proceedings of The International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management*, 2009.