

# On the Notion of Framing in Computational Creativity

**John Charnley, Alison Pease and Simon Colton**

Computational Creativity Group  
Department of Computing, Imperial College  
180 Queens Gate, London SW7 2RH, United Kingdom.  
ccg.doc.ic.ac.uk

## Abstract

In most domains, artefacts and the creativity that went into their production is judged within a context; where a context may include background information on how the creator feels about their work, what they think it expresses, how it fits in with other work done within their community, their mood before, during and after creation, and so on. We identify areas of framing information, such as motivation, intention, or the processes involved in creating a work, and consider how these areas might be applicable to the context of Computational Creativity. We suggest examples of how such framing information may be derived in existing creative systems and propose a novel dually-creative approach to framing, whereby an automated story generation system is employed, in tandem with the artefact generator, to produce suitable framing information. We outline how this method might be developed and some longer term goals.

## Introduction

Michael Craig-Martin's 1973 work, *An Oak Tree*, comprises a glass of water on a shelf and an accompanying text, in which Craig-Martin claims that the object which appears to be a glass of water is really an oak tree. The text takes the form of a question and answer session written by Craig-Martin about how he has changed the water into a tree:

A. [...] I've changed the physical substance of the glass of water into that of an oak tree.

Q. It looks like a glass of water.

A. Of course it does. I didn't change its appearance. But it's not a glass of water, it's an oak tree.

...

Q. Haven't you simply called this glass of water an oak tree?

A. Absolutely not.

Craig-Martin is rather mysterious as to how he has accomplished the change:

Q. Was it difficult to effect the change?

A. No effort at all. But it took me years of work before I realised I could do it.

Q. When precisely did the glass of water become an oak tree?

A. When I put the water in the glass.

Q. Does this happen every time you fill a glass with water?

A. No, of course not. Only when I intend to change it into an oak tree.

The status of the piece as a work of art is then raised:

Q. Do you consider that changing the glass of water into an oak tree constitutes an art work?

A. Yes.

Q. What precisely is the art work? The glass of water?

A. There is no glass of water anymore.

Q. The process of change?

A. There is no process involved in the change.

Q. The oak tree?

A. Yes. The oak tree.

This is an example of human creativity which is taken seriously in its field. First shown in 1974, it was bought by the National Gallery of Australia in Canberra in 1977, and has been exhibited all over the world with the text translated into at least twenty languages. Like many important works, opinion is divided: artist Michael Daley referred to it as "self-deluding" and "pretentious" (Daley August 31 2002), while art critic Richard Cork wrote:

"I realise that one of the most challenging moments occurred in 1974 when the Rowan Gallery mounted an exhibition of Michael Craig-Martin's work." (Cork October 9 2006).

Researchers in Computational Creativity (CC) can learn from this work. The main point that we consider in this paper is that *the artefact* (the glass of water) *has no creative value without the title and accompanying text*. The value, the creativity associated with the piece, lies in the narrative surrounding the glass of water. This point has clear implications for CC, which has traditionally focused on artefact generation, to the extent that the degree of creativity judged to be in the system is often considered to be entirely dependent on characteristics of the set of artefacts it produces (for instance, see (Ritchie 2007)). Very few systems in CC currently generate their own narrative, or framing information. Artefacts are judged either in isolation or in conjunction with a human-produced narrative, such as the name of the system and any scientific papers which describe how it works. Researchers in CC will be familiar with what Bedworth and Norwood call "carbon fascism" (Bedworth and Norwood 1999) (the bias that only biological creativity can produce valuable artefacts), and, for the most part, computer-generated creative artefacts are not taken seriously

by experts in the domain in which the artefacts belong. We believe that enabling creative software to produce its own framing information will help to gain acceptance from these experts. While *An Oak Tree* may be a rather extreme example of the importance of framing information, we hold that such information almost always plays some role in creative acts, and is a fundamental aspect of human creativity. We consider here which types of framing information we could feasibly expect a piece of software to produce, and begin to propose ways in which we could formalise this. Specifically, we consider three areas in computational terms: *motivation* (why did you do X?), *intention* (what did you mean when you did X?), and *processes* (how did you do X?). We make the following contributions:

1. We highlight the importance of framing information in human creativity.
2. We propose an approach to automatically generating framing information, in which a separate creative act of automated story generation is performed alongside traditional artefact generation.

### Framing in human creativity

Sir John Tusa is a British arts administrator, radio and television journalist, known for his BBC Radio 3 series *The John Tusa Interview*, in which he interviews contemporary artists. These interviews have been reproduced as two books in which he explores the processes of creativity (Tusa 2003; 2006). We have analysed all thirteen interviews in his most recent collection, in order to provide the starting point for a taxonomy of framing information. In the following discussion, unless otherwise specified, all page numbers refer to this collection of interviews (Tusa 2006). We identified two categories which artists spoke about: INTERNAL, or inward looking, in which the artist talks about their own *Work*, *Career* and *Life*, and EXTERNAL, or outward looking, in which the artist talks about their view of their *Audience* and *Field*.

### The artist's work

Discussion about artists' work is very common in the Tusa interviews. This might concern a specific piece; such as how an artist feels about a piece, what they think it expresses, or how it relates to everyday concepts; or this might concern details of the generative process, such as how the work is created, how processes involved in its creation fit together, or whether a new technique or material changed the way that something was done. As an example below, Cunningham (MC) relates his work to scientific and religious concepts:

*MC*: ...it was the statement of Einstein's which I read at that time, where he said, 'There are no fixed points in space.' And it was like a flash of lightning; I felt, Well, that's marvellous for the stage. Instead of thinking it's front and centre, to allow any point, very Buddhist, any point in the space to be as important as any other. (p. 66)

### The artist's career

A picture of the structure of an artist's career, in terms of his or her past, present and possible future directions, can aid

understanding of current work. Questions about previous work include asking how two pieces differ; what category work from certain periods falls into; classification of a career into different stages, such as early and late, or pre-work X and post-work X.

Examples from (Tusa 2006) include the questions: "So you think you are recognisably the same person, creatively the same person as you would have been if you'd stayed in New York?" (to Forsythe, p. 93); "What's the next stage of your evolution as a maker of ballets?" (also to Forsythe, p. 105), and "When you look back over the last twenty years, would you ever have guessed that the work that you do would have travelled so far .... I mean this is an extraordinary journey. How aware have you been of the evolution as you've been through it?" (to McBurney, pp. 181-2).

### The artist's life

Audiences are interested in the personalities and influences behind society's "creative heroes". Topics of interest include political, intellectual, personal, cultural and religious influences; value systems; reasons for working in a particular area; important events in the life of the artist, and so on.

John Tusa asks many questions in this vein. For instance, he asks: "Are you an optimist or a pessimist as a person?" (to Rovner); "When did you discover that you had this condition called Dysgraphia, where I think the brain wants to write words as pictures?" (to Viola); "What do you feel, as you're coming in to work?" (to McBurney); and "What music do you like?" (to Piano), and has some rather poignant exchanges, such as one with Viola in which he asks about a near-death experience (pp. 221-3); and this exchange with Rovner (MR):

*JT*: Are you lonely as an artist?

*MR*: You mean as an artist or as a person?

*JT*: Well as a person who is an artist.

*MR*: I'm alone. I don't know if I'm lonely. I am single, you know I'm a single person, I'm a single person. (p. 213)

### The artist's view of their audience

The perception that an artist has of his or her audience may influence their work. Queries in this topic included questions about effect of a particular field on audiences, and what the effect of certain pieces of work are on the collective subconscious. Egoyan, for instance, discusses responsibility to one's audience with Tusa (p75).

### The artist's view of the field

Embedding a particular artist's work into the context of a body of work is one of the purposes of framing information. Queries include definitional questions about particular fields, and their relationship to other fields; how a piece fits into a field; in which field an artist sees themselves; the influence of external characteristics such as politics, or how modern advancements such as new techniques have affected a field; the history of a field and directions in which it could go, and so on. For instance, Egoyan discusses how he thinks video compares to film (p76), and Forsythe talks about how his work fits into great classical ballets (p. 106).

## Framing for Computational Creativity

Analysis of the interview responses suggests a new direction for CC: *enabling creative software to generate some of its own framing information*. As with human artworks, the appeal of computer creativity will be enhanced by the presence of framing. However, there are obvious restrictions on the scope to which the various forms of framing apply in the computer generated context. Here we consider three areas in computational terms: motivation (why did you do X?), intention (what did you mean when you did X?), and processes (how did you do X?).

### Motivation

Many creative systems currently rely upon human intervention to begin, or guide, a creative session and the extent to which the systems themselves act autonomously varies widely. In some sense, the level to which these systems could be considered self-motivating is inversely proportional to the amount of guidance they receive. However, it is possible to foresee situations where this reliance has been removed to such an extent – and the human input rendered so remote – that it is considered inconsequential to the creative process. For instance, the field of Genetic Programming (Koza 1992) has resulted in software which can, itself, develop software. In the CC domain, software may eventually produce its own creative software which, in turn, produces further creative software, and so forth. In such a scenario, there could be several generations in an overall genealogy of creative software. As the distance between the original human creator and the software that directly creates the artefact increases, the notion of self-motivation becomes blurred.

Beyond this, the scope for a system's motivation towards a particular generative act is broad. For example, a suitably configured system may be able to perform creative acts in numerous fields and be able to muster its effort in directions of its own choosing. With this in mind, we can make a distinction between *motivation to perform creative acts in general*, *motivation to create in a particular field* and *motivation to create specific instances*.

In the human context, the motivation towards a specific field may be variously influenced by the life of the artist, their career and their attitudes, in particular towards their field and audience. Several of these are distinctly human in nature and it currently makes limited sense to speak of the *life* or *attitudes* of software in any real sense. By contrast, we *can* speak of the *career* of a software artist, as in the corpus of its previous output. This may be used as part of a process by which a computer system decides which area to operate within. For example, we can imagine software that chooses its field of operation based upon how successful it has previously been in that area. For instance, it could refer to external assessments of its historic output to rate how well-received it has been, focusing its future effort accordingly.

The fact that a computer has no *life* from which to draw motivation does not preclude its use as part of framing information. All those aspects missing from a computer could, alternatively, be simulated. For example, we have seen music

software that aims to exhibit characteristics of well-known composers in attempts to capture their compositional style (Cope 2006). The extent to which the simulation of human motivation enhances the appeal of computer generated artefacts is, however, still unquantified. The motivation of a software creator may come from a bespoke process which has no basis in how humans are motivated. The details of such a process, and how it is executed for a given instance, would form valid framing information, specific to that software approach.

### Intention

The aims for a particular piece are closely related to motivation, described above. A human creator will often undertake an endeavour because of a desire to achieve a particular outcome. Factors such as attitudes to the field contribute to this desire. Certainly, by the fact that some output is produced, every computer generative act displays intent. The aims of the process exist and they can, therefore, be described as part of the framing. In the context of a computer generative act, we might distinguish between *a priori* intent and intentions that arise as part of the generative process. That is, the software may be pre-configured to achieve a particular goal although with some discretion regarding details of the final outcome, which will be decided during the generative process. The details of the underlying intent will depend upon the creative process applied. For example, as above, software creators might simulate aspects of human intent.

Intent has been investigated in collage-generation systems (Krzeczkowska et al. 2010). Here, the software based its collage upon events from the news of that day with the aim of inviting the audience to consider the artwork in the context of the wider world around them. This method was later generalised to consider wider combinations of creative systems and more-closely analyse the point in the creative process at which intentionality arose (Cook and Colton 2011).

### Processes

In an act of human creativity, information about the creative process may be lost due to human fallibility, memory, awareness, and so on. However, in a computational context there is an inherent ability to perfectly store and retrieve information. The majority of creative systems would have the ability to produce an audit trail, indicating the results of key decisions in the generative process. For example, an evolutionary art system might be able to provide details of the ancestry of a finished piece, showing each of the generations in between. The extent to which the generative process can be fully recounted in CC is, nevertheless, limited by the ability to fully recreate the sources of information that played into the generative process. Software may, for instance, use information from a dynamic data source in producing an artefact, and it may not be possible to recreate the whole of this source in retrospect.

One system that produces its own framing is an automated poetry generator currently being developed (Colton, Goodwin, and Veale 2012). In addition to creating a poem, this system produces text which describes particular aspects of

its poetry that it found appealing and aspects of how it generated its output. In order to fully engage with a human audience, creative systems will need to adopt some or all of the creative responsibility in generating framing information.

Details of the creative process are valid aspects of framing information, which are relevant to both computational and human creative contexts. There is a notion of an appropriate level of detail: extensive detail may be dull and the appreciation of artefacts is sometimes enhanced by the absence of information about the generative process.

### Examples of framing for Computational Creativity

There are many ways in which creative systems might generate their own framing information. For example, an automated art system, such as AARON (McCorduck 1991), could store details of all its previous artworks and provide an assessment of how a new piece differs, in various respects, from its past output. A poetry system, such as (Colton, Goodwin, and Veale 2012), might reveal the general mood of the inspiring source it used as a basis for an affective poem. Mathematical software, such as HR (Colton 2002), could be given the ability to compare the conjectures it finds against on-line mathematical databases and report on how its output relates to known theorems. By corollary, an art system could appeal to image databases to suggest similarities to other artists. A simple enhancement to the collage generation program of (Krzeczowska et al. 2010) could see it provide the text of the news story that formed the inspiration for the collage. In this mode, the framing information would become as important an aspect of the overall presentation as the collage itself. The artwork would be a combination of both the collage and underlying story, rather than the collage alone. This list is by no means exhaustive. The varied nature of framing information that we have been describing shows that the opportunities for enhancing works with framing are extensive.

### A dually-creative approach to framing

Framing information has the potential to greatly impact an audience's assessment of an artefact. In some instances, framing is arguably as much a part of the overall creative presentation as the artefact itself: this was seen in Craig-Martin's *An Oak Tree*, described above, as well as, for example, elements of Marcel Duchamp's *readymades* series such as *Fountain*. The information can be as simple as a title for the artefact, or might encompass much of the type of framing indicated in our analysis. Framing can add to the mystique and mystery surrounding an artefact, as we have described.

Framing information need not be factually accurate. Information surrounding human creativity can be lost, deliberately falsified or made vague for artistic impact. Thus, the generation of framing information can itself be seen as a creative act. The overall impact of the *package* – namely the artefact and the associated framing information – will depend on both the assessed quality of the artefact, together with the impression given by the framing information. We propose one approach to artefact-with-framing generation,

where the two are produced simultaneously, by a dually-creative process. Under this approach, the most appropriate creative paradigm for the framing information would be a form of automated storytelling. One part of a combined system would create the artefact itself and a storytelling aspect would generate a framing story. The framing story could be as simple or as complex as those which accompany human creations. Tools which were able to perform tasks such as metaphor and analogy (see (Gentner, Holyoak, and Kokinov 2001; Gibbs Jr. 2008)) might be integrated into the storytelling aspect.

In the previous section, we discussed aspects of framing which might be relevant to the CC setting. This information could form much of the input to the story generation system, becoming part of the basis of the story. For example, purely factual information about how the software arrived at the final product could be retained. Given that there is no requirement for the framing story to be factually correct, some or all of the story might be fictional and there is no prescription for the extent to which the framing story should directly correspond with the artefact. Consider, for example, a framing story which describes all aspects of the creative process in full detail compared with a framing story consisting of a random seemingly-unrelated word. Both have artistic value, but in entirely different ways.

An initial approach to the fact that no configuration of a particular automated story generation system would be able to generate the variety of framing stories that we have witnessed in human creativity, might be to develop a small number of story-telling paradigms, each based upon a particular story template. One challenge might then be to achieve an appropriate balance between fact and fiction in the generated stories. In future, we might hand such decisions over to the software. For example, a sufficiently-able software suite might decide which story-telling paradigm is most appropriate for a particular effect, the balance between fact and fiction and how extensive the framing should be. In a more complex manifestation, the story might form an interactive dialogue, providing answers to audience queries in a manner akin to an interview. As with human creativity, the answers to those questions may be entirely at the whim of the generating system. Going further, software might employ story generation approaches to simulate aspects of the framing information which might otherwise be absent, such as a religious belief or other motivation. This could, in turn, feed back into the generation of the creative artefact itself. Storytelling for framing information represents an interesting challenge for our existing and future automated story generation systems.

### Related work

In (Colton, Pease, and Charnley 2011; Pease and Colton 2011), two generalisations were introduced with the aim of enabling more precise discussion of the kinds of behaviour exhibited by software when undertaking creative tasks. The first generalisation places the notion of a generative act, wherein an artefact such as a theorem, melody, artwork or poem is produced, into the broader notion of a *creative act*. During a creative act, multiple types of generative acts are

undertaken which might produce framing information, *F*, aesthetic considerations, *A*, concepts, *C*, and exemplars, *E*; in addition to generative acts which lead to the invention of novel generative processes for the invention of information of types *F*, *A*, *C* and/or *E*.

The second generalisation places the notion of assessment of the aesthetic and/or utilitarian value of a generated artefact into the broader notion of the impact of a creative act, *X*. In particular, an assumption was introduced that in assessing the artefacts resulting from a creative act, we actually celebrate the entire creative act, which naturally includes information about the methods underlying the generation of the new material, and the framing information, which may put *X* into various contexts or explain motivations, etc., generally adding value to the generated artefacts over and above their intrinsic value.

The introduction of these two generalisations enabled the FACE and IDEA descriptive models to be introduced as the first in the fledgling formalisation known as *Computational Creativity Theory*. In this paper we have extended this model by exploring the notion of *framing*.

### Future work and conclusions

Creativity is not performed in a vacuum and the human context gives an artefact meaning and value. Implicit in the Computational Creativity Theory models so far developed is the notion that the FACE information/artefacts resulting from creative acts can be seen as invitations to a dialogue. For instance, when a person appreciates a painting, they are encouraged to ask questions of it, and look for answers, either explicitly from the artist or some perceived notion of how artists work, via visual interrogation of the piece itself, or through certain cultural contexts; for example, by understanding the culture in the time and place when the painting was produced.

Despite the importance of framing information as part of the overall artistic endeavour, we are only aware of a very small number of systems that generate framing information to accompany their creative output. We have proposed one approach to this, whereby automated story generation is used to generate framing information. There are no real bounds to what information such framing can contain, its basis in fact versus fiction, or the format in which it is presented. Consequently, we suggest that initial attempts be restricted to a small number of simplified paradigms, taking their basis from a more complete investigation into how humanly-produced framing information relates to CC. Expanding upon this starting point, we imagine software taking over some of the creative responsibility for the framing information, such as determining the story-telling paradigm and the story's emphasis or level of detail.

Craig-Martin, via his narrative of *An Oak Tree*, opens up a dialogue with the viewer on the nature of essence, proof, faith, matter, reality, art, and so on. The viewer engages with this narrative, which includes the manner of presentation of the piece, Craig-Martin's background as an artist and a person, critics' and artists' responses to the piece, stories surrounding the work and effects that it has on everyday life

(for instance, there is a myth that Australian customs officials barred it from entering the country since it was classified as "vegetation", and in February 2012 the first three hits from google images on the search term "an oak tree" are images of Craig-Martin's work). We anticipate that the direction outlined in this paper will form an important axis of development for CC systems. Our long-term goal is to help to develop CC to such an extent that one day a piece of creative software will appear in the table of contents of a collection of Tusa-style interviews, to discuss its work and itself, alongside other contemporary artists.

### Acknowledgements

This work has been funded by EPSRC grant EP/J004049. We are grateful to the three reviewers, who raised interesting points.

### References

- Bedworth, J., and Norwood, J. 1999. The Turing test is dead. In *Proceedings of the 3rd conference on creativity and cognition*.
- Colton, S.; Goodwin, J.; and Veale, T. 2012. Full face poetry generation. In *Proceedings of the Third International Conference on Computational Creativity*.
- Colton, S.; Pease, A.; and Charnley, J. 2011. Computational creativity theory: The FACE and IDEA descriptive models. In *Proceedings of the 2nd International Conference on Computational Creativity*.
- Colton, S. 2002. *Automated Theory Formation in Pure Mathematics*. Springer-Verlag.
- Cook, M., and Colton, S. 2011. Automated collage generation – with more intent. In *Proceedings of the Second International Conference on Computational Creativity*.
- Cope, D. 2006. *Computer Models of Musical Creativity*. Cambridge, MA: MIT Press.
- Cork, R. October 9, 2006. Losing our vision. *New Statesman*.
- Daley, M. August 31, 2002. Tracey left on the shelf. *The Guardian*.
- Gentner, D.; Holyoak, K.; and Kokinov, B. 2001. *The Analogical Mind: Perspectives from Cognitive Science*. Cambridge, MA: MIT Press.
- Gibbs Jr., R. W., ed. 2008. Cambridge, UK: Cambridge University Press.
- Koza, J. R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA, USA: MIT Press.
- Krzeczkowska, A.; El-Hage, J.; Colton, S.; and Clark, S. 2010. Automated collage generation – with intent. In *Proceedings of the 1st International Conference on Computational Creativity*.
- McCorduck, P. 1991. *AARON's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen*. W.H. Freeman and Company.
- Pease, A., and Colton, S. 2011. Computational creativity theory: Inspirations behind the FACE and the IDEA models. In *Proceedings of the 2nd International Conference on Computational Creativity*.
- Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17:67–99.
- Tusa, J. 2003. *On Creativity*. London: Methuen.
- Tusa, J. 2006. *The Janus aspect: artists in the twenty-first century*. London: Methuen.