# An institutional approach to computational social creativity

**Joseph Corneli**
Computational Creativity Group
Goldsmiths College
London, UK
j.corneli@gold.ac.uk

### Abstract

Modelling the creativity that takes place in social settings presents a range of theoretical challenges. Mel Rhodes's classic "4Ps" of creativity, the "Person, Process, Product, and Press," offer an initial typology. Here, Rhodes's ideas are connected with Elinor Ostrom's work on the analysis of economic governance to generate several "creativity design principles." These principles frame a survey of the shared concepts that structure the contexts that support creative work. The concepts are connected to the idea of computational "tests" to foreground the relationship with standard computing practice, and to draw out specific recommendations for the further development of computational creativity culture.

## Introduction

One two-part claim is advanced and defended herein: *Elinor Ostrom's theory of institutions can be used to design systems that exhibit computational social creativity, and a culture supports this work*. The contribution takes the form of several candidate "design principles," a literature survey that elaborates them, and an analysis that connects these ideas to common programming practice.

The paper is structured as follows. The "Background" section describes Ostrom's (1990) *Institutional Analysis and Development* (IAD) framework, focusing on her proposed design principles for commons management. To connect these ideas to social creativity, the paper draws on the 4Ps (Person/Process/Press/Product), a model for thinking about creative contexts (Rhodes 1961) that has been brought to bear in theorising computational creativity (Jordanous 2016). This is summarised and slightly adapted. In the subsequent main section of the paper, "Testing for Creativity", Ostrom's design principles are transposed from the world of commons management to the world of computational social creativity. This section looks for ways to connect the propsed creativity design principles to computational methods, and also draws on contemporary thinking in the philosophy of technology, with examples from familar social computing settings like Wikipedia. A two-part example dealing with both the "soft" culture of the computational creativity community and potential software-based interventions is presented in the "Example" section. Finally, the "Discussion and Conclusions" highlight the relevance of this work for computational creativity culture, systems, and evaluation.

## Background

This section summarises the motivation for the paper, introduces Elinor Ostrom's work, and reviews Rhodes's 4P framework. The central parts of this section are Table 1 and 2, which list Ostrom's design principles for managing a commons, and transpose them to creative domains.

**Motivation** The current investigation is motivated, in part, by the idea of *Ecologically Grounded Creative Practice* (Keller, Lazzarini, and Pimenta 2014). Within a given ecological niche, agents and objects interact; niches can also be brought into relationship in creative ways. The current work has in mind relatively sophisticated agents with their own "contextual maps" and the ability to participate in "reading and writing computational ecosystems" (Antunes, Leymarie, and Latham 2015). Such agents will use, view, critique, and evaluate the work and workflow of other agents. Although computational agents with all of these features do not exist yet in any robust form, we can reason about them, and in so doing, help design the future of *computational social creativity* (Saunders and Bown 2015) – an emerging research area at the nexus of artificial life, social simulation, and computational creativity.

**Elinor Ostrom's "design principles"** To contextualise this effort, we must begin with a short excursus into economics. Ostrom's work is typically applied to study the management of natural resources. In economics jargon, the specific resources considered are *rivalrous* and *non-excludable*. This means that consumption by one party precludes consumption by a rival, and that it is not directly possible to for anyone to block others' access to the resource. Economic goods with these two properties are referred to as *common pool resources* (CPRs); see (Ostrom 2008). Fisheries and forests are important examples. Economic actors have incentives to exploit these resources, however, there are natural limits on total consumption. In principle, a CPR might be gobbled up due to individual greed: this is the so-called tragedy of the commons, and one does not have to look too far for examples. However, in practice, the tragic outcome does not always transpire. Ostrom's theoretical and empirical perspective helps understand why, and emphasises:

> (1) the importance of group attributes and institutional arrangements in relation to the structure of incentives and utilities for individual decision making; and (2) the likelihood of a broader set of possible outcomes, including user-group institutional solutions (McCay and Acheson 1990, p. 23)

Ostrom's ideas have recently been applied to analyse Wikipedia, considered as an "expressive commons" (Safner 2016). Wikipedia is *non-rivalrous* in consumption, if we accept the metaphor "to read is to consume." However, *contribution* to Wikipedia presents a range of salient social dilemmas, and efforts to manage them are reflected, for example, in the *Neutral Point Of View* (NPOV) policy, which helps produce "articles that document and explain major points of view, giving due weight with respect to their prominence in an impartial tone."[1]

IAD focuses on *action situations*, framed in three phases: *context*, *action*, and *outcome*. Importantly, this part of the theory is not linked to the particular details of CPRs. Ostrom uses the term *institution* to refer to the "shared concepts used by humans in repetitive situations organized by rules, norms, and strategies" (Ostrom 2010). We will return to these concept categories later and consider them further from a computational perspective. For now, our way into thinking in terms of IAD will be by way of several *design principles* for the successful management of CPRs that Ostrom described; see Table 1. These principles work together to support institutions that maintain the integrity of the commons – for example, by ensuring that behaviour is monitored, that knowledgeable and concerned parties are the ones who make specific rules, and that conflicts do not get out of hand (Ostrom et al. 2012, p. 79).

**The four Ps** We can bootstrap our contextual understanding of creativity with the help of an existing model. Rhodes (1961, pp. 307-309) intends "the four Ps" to refer to the following facets of creativity, which are familiar from everyday experiences of creativity in society.

> Person – *personality, intellect, temperament, physique, traits, habits, attitudes, self-concept, value-systems, defense mechanisms, and behavior.*
>
> Process – *motivation, perception, thinking, and communication.*
>
> Product – *an idea embodied into a tangible form.*
>
> Press – *the relationship between human beings and their environment.*

We will shortly use these concepts to rewrite the items in Table 1, replacing the focus on *appropriation* with a focus on *contribution* that befits a theory of social creativity.

Jordanous makes a case for thinking about computational creativity using Rhodes's 4P's, starting with a critique of the strategies used in the evaluation of computational creativity, which, she emphasises, is "traditionally considered … from the perspective of the creative output produced by a system" (Jordanous 2016).

Ecological thinking suggests that that it is quite limited to take the final product as the sole term of analysis. At least we might like to introduce the "embedded evaluation" of creative products into the creative process, and build agents that are aware of some contextual features of their environment. For example, these agents might ask: *How similar or how different is my generated artwork to an existing artwork, or*

**Ostrom's design principles**

**1A. User boundaries**
"Clear boundaries between legitimate users and nonusers must be clearly defined."[2]

**1B. Resource boundaries**
"Clear boundaries are present that define a resource system and separate it from the larger biophysical environment."

**2A. Congruence with local conditions**
"Appropriation and provision rules are congruent with local social and environmental conditions."

**2B. Appropriation and provision**
"The benefits obtained by users from a common-pool resource (CPR), as determined by appropriation rules, are proportional to the amount of inputs required in the form of labor, material, or money, as determined by provision rules."

**3. Collective-choice arrangements**
"Most individuals affected by the operational rules can participate in modifying the operational rules."

**4A. Monitoring users**
"Monitors who are accountable to the users monitor the appropriation and provision levels of the users."

**4B. Monitoring the resource**
"Monitors who are accountable to the users monitor the condition of the resource."

**5. Graduated sanctions**
"Appropriators who violate operational rules are likely to be assessed graduated sanctions (depending on the seriousness and context of the offense) by other appropriators, by officials accountable to these appropriators, or both."

**6. Conflict-resolution mechanisms**
"Appropriators and their officials have rapid access to low-cost local arenas to resolve conflicts among appropriators or between appropriators and officials."

**7. Minimal recognition of rights to organise**
"The rights of appropriators to devise their own institutions are not challenged by external governmental authorities."

**8. Nested enterprises**
"Appropriation, provision, monitoring, enforcement, conflict resolution, and governance activities are organised in multiple layers of nested enterprises."

Table 1: Ostrom's design principles as expressed in the meta-review carried out by Cox, Arnold, and Tomás (2010)

[1] `https://en.wikipedia.org/wiki/Wikipedia:Five_pillars`

[2] Repetition is *sic*, the point being that the boundaries must be both distinct and explicitly defined.

**Proposed creativity design principles**

**1A.** The population of Producers who can add to or alter the resource is clearly defined.

**1B.** The boundaries of the Place must be well defined.

**2A.** The Process is related to local conditions.

**2B.** Contributing to the Product has benefits for the Producer that are proportional to the efforts expended.

**3.** Most Producers who are affected by the rules governing contribution can participate in modifying the operational rules.

**4A.** Tests document the interaction of Producers and Place.

**4B.** Tests can be modified by Producers or their representatives.

**5.** Producers who violate operational rules in the domain will be assessed sanctions by other Producers.

**6.** Producers have rapid access to low-cost local arenas to resolve conflicts.

**7.** The rights of Producers to devise institutions governing their contributions are not challenged by external authorities.

**8.** Contribution, testing, enforcement, conflict resolution, and governance and are organised in multiple layers of nested Places and agencies.

Table 2: "Creativity design principles" formed by switching the polarity of entries in Table 1 to emphasise contribution rather than appropriation, and using the concept of "tests" to connect to computing practice

*to the components thereof, or to the initial conception for the work?* This route is quite close to Ritchie's (2007) empirical criteria for judging a final product against an "inspiring set" – but now makes evaluation an explicit part of the creative process. Some recent work in computational creativity emphasises embedded evaluation (Gervás and León 2014). However, as Jordanous argues, creative products are just one part of the overall creative process – and the 4Ps help expose the other features.

Unfortunately, however, Rhodes's thinking and terminology is too anthropocentric for our current purpose. As Ostrom describes it, action situations are to be understood using seven clusters of variables: *participants*, *positions*, *potential outcomes*, *action-outcome linkages*, *participant control*, *types of information generated*, and *costs and benefits assigned to actions and outcomes* (Ostrom 2009, p. 14). Nowhere does this mention a "Person". Continuing the adaptations begun by Jordanous (2016), the four Ps will be rendered here as Producer/Process/Product/Place. It is important to emphasise that these labels are strictly more inclusive than Rhodes's, and more abstract. In particular, the Place corresponds to Ostrom's action situation, structured in advance by contextual features. This adapted 4P model is reminiscent of the *Domain-Individual-Field Interaction* (DIFI) model due to Feldman, Csikszentmihalyi, and Gardner (1994), if we understand Domain $\approx$ Place, Individual $\approx$ Producer, and Field $\approx$ (a collection of) established Processes. Note that contextual theories, broadly construed, pose a long-standing challenge for computing, partly because "what context is changes with its context" (Gundersen 2014, p. 343). One possible working definition is that: "Context is what contrains a problem solving [scenario] without intervening in it explicitly" (Brézillon 1999). Another relevant remark is that context is "defined solely in terms of effects in a given situation" (Hirst 2000).

In developing an *institutional approach to computational social creativity*, we will look for the rules, norms, and strategies that can be used to establish suitable and effective contextual relationships between Process(es), Place(s), Producer(s), and Product(s).

**Transposing the design principles into "creativity design principles" and translating them into technical terms**

*Software testing* is embodied in the formal ideas of *assertions*, *advice*, and *contracts*. Related programming methodologies aim to build *executable specifications* and may make use of *test-driven development* (TDD). These techniques provide various ways for (evolving) programs to interact with their context. These ideas can help us translate Table 1 into technical terms. To get started, Table 2 uses the 4P terminology and the generic notion of a test to transpose Ostrom's design principles into "creativity design principles."

## Testing for creativity

The current section elaborates the candidate creativity design principles outlined above, expanding each with relevant literature and examples, and seeking the ways in which each principle could be applied within a software system.

### 1A. User boundaries

"The population of Producers who can add to or alter the resource is clearly defined."

In user-oriented computing, this principle is often addressed using *Access Control Lists* (ACLs) or other permissions mechanisms. The corresponding tests are relatively simple: either each modifiable object in the system has a piece of metadata about it that says who can modify it, or each user has a piece of metadata attached to his or her user account that says which resources they can modify.

Before granting access to a resource, we may require that a Producer implements certain protocols. In a client-server architecture, the client generally communicates using an existing API and may need to implement a certain set of callback functions or adhere to other restrictions. Noncompliant user behaviour after access has been granted may result in access being revoked. Thus, for example, even though Wikipedia is "the encyclopedia anyone can edit," violating the site's principles may lead to a IP-based block, or a username-based ban.

### 1B. Resource boundaries

"The boundaries of the Place must be well defined."

The source of this well-definedness may come from "either side." That is, the Place may advertise its definition in terms of its APIs and other criteria (as above) together with guarantees on output behaviour in the style of "Design by Contract" (Mitchell and McKim 2002); alternatively, Producers may implement tests that restrict the Places that they will engage with.

In a simple example of the latter sort, a game-playing agent might resign if it estimates that its position is unwinnable. The fact that different participants can have different perspectives points to an interesting special case in which the (shared) definition of the Place arises in an emergent manner. This phenomenon is especially important if we "[take] a broad view of creativity as any process in which novel outcomes emerge" (Saunders and Bown 2015).

**2A. Congruence with local conditions**

"The Process is related to local conditions."

In Ostrom's original formulation, local conditions were broken down along axes of "time, place, technology, and/or quantity of resource units" (Ostrom 1990). With respect to theorising the local conditions of creativity, we can gain a useful perspective by turning briefly to the psychoanalyst Winnicott's treatment of "the exciting interweave of subjectivity and objective observation" which takes place in an "area that is intermediate between the inner reality of the individual and the shared reality of the world" (Winnicott 2002, p. 86). We are then led to consider those local conditions that exist in the "interweave" of Place and Producer. For example, roboticist Andy Clark proposes a theory of extended cognition, in which enminded beings "use" the environment to self-program and are not just programmed by the environment (Clark 1998). However as Clark points out elsewhere, "it becomes harder and harder to say where the world stops and the person begins" (Clark 2001). In short, the mind is not separated from the body or environment but grounded in perception (Ingold 2000).

A corresponding computational test is found in the earlier example of embedded evaluation, in which existing artefacts are employed as a virtual sensorium. More broadly, this principle concerns making sense of, or "parsing", the Place (and the other P's). This Process is well described by Steigler's notion of *grammatisation*: "processes by which a material, sensory, or symbolic flux becomes a gramme," or, more simply, "the production and discretisation of structures" (Tinnell 2015). Remember that while a given agent is trying to make sense of the world, others are likely trying to make sense of that agent as well. Framed as dilemma, the last word would likely be: "program or be programmed" (Rushkoff 2010) – but reflecting on Clark's comment above, we see that this can become somewhat complex.

**2B. Appropriation and provision**

"Contributing to the Product has benefits for the Producer that are proportional to the efforts expended."

The usual way of thinking about computers – as non-agentive machines – would render the above-stated principle perfectly meaningless. In connection with principle 2A, we should here remark: "That an object is more profitable or effective is only a secondary consequence of its refinement" (Chabot 2013, p. 12). In any case, before we can think about "benefits" in the case of a non-human (and non-living) Producer, the phrasing of the current principle leads us to ponder the cost of their "efforts."

It may be best to change tack, and ask, with Terrance Deacon, "In what sense could a machine be alive?" (Deacon 2014). If a machine were responsible for maintaining its own energy supply, its features of outward-orientation might give cause to say that the machine has a "self" (Deacon, Haag, and Ogilvy 2011). Consider for example the Ethereum project, which provides protocols for distributed computing and the creation of "decentralized autonomous organisations" – whose organisation relative to the outside world is mediated by cryptocurrency, referred to as "fuel" (Wood 2014).

From a testing standpoint, the key requirements are: an ability to judge whether a given option can be (tentatively) thought of as beneficial, and, ideally, a memory that can compare these judgements with iterations of similar situations later on. In this way we would recover the foundations of reinforcement learning, and, as Ostrom points out, the core logic behind the development of new institutions:

> "How about if you do $A$ in the future, and I will do $B$, and before we ever make a decision about $C$ again, we both discuss it and make a joint decision?" (Ostrom 2009, p. 19)

**3. Collective-choice arrangements**

"Most Producers who are affected by the rules governing contribution can participate in modifying the operational rules."

Let us reflect in more detail on the *rules* that comprise – along with biophysical and material conditions and community attributes – the locally-contextual variables which determine or constrain an action situation (Ostrom 2009, p. 15). At their simplest, these rules are "if-then" statements giving instructions that determine the behaviour of persons in certain roles. As such, each rule contains a logical test, and changing the rules means writing new tests.

Ostrom develops a grammar around this idea, and defines *regulatory rules* with the following formula:

> ATTRIBUTES of participants who are OBLIGED, FORBIDDEN, OR PERMITTED to ACT (or AFFECT an outcome) under specified CONDITIONS, OR ELSE. (Ostrom 2009, p. 187)

*Norms* and *strategies* are defined using a simplified formula, also cast in terms of *attributes*, *deontics*,[3] *aim*, and *conditions* (Ostrom 2009, p. 140). The prescriptive terms may be assigned a particular weight, and actions and consequences may also be assigned a particular cost or value (Ostrom 2009, p. 142). Some relevant actions are: *be* in a position, *cross* a boundary, *effect* a choice, *jointly exercise* partial control together with others, *send or receive* information, *pay out or receive* costs or benefits, and *take place* (for outcomes) (Ostrom 2009, p. 191).

[3] I.e., the presciptive valence – obliged, forbidden, or permitted, as above – for norms, not for strategies.

Something more needs to be said about the assertion that Producers "can" participate in changing (or creating) rules, norms, and strategies. In practice, participatory systems tend to be lossy. Changes to rules and structures will tend to be carried out by those Producers who are *most* affected – and *thus* most knowledgeable; cf. Ostrom et al. (2012, p. 79). The structure of new rules is predicted by Conway's Law:

> [T]here is a very close relationship between the structure of a system and the structure of the organization which designed it. (Conway 1968)

Specifically, the proposed relationship is "homomorphism": following Conway, any Product will mirror the hyper-local conditions that describe the Producers' social context. Furthermore, it seems likely that Products will mirror local environmental conditions in the Place. This points to importance of a broad class of tests that would be described as environmental "sensors". This theme will be developed more fully below.

**4A. Monitoring users**

"Tests document the interaction of Producers and Place."

The straightforward view suggested by the idea of "monitoring" is to deploy some global functionality that keeps track of the actions of all participating Producers within a Place. But this function can be broken up and distributed out among the Producers themselves. In the first instance, what a Producer produces is sensory data. Sensors are generally deployed along with effectors or (more broadly) transducers that translate the sensory information into action. So, monitoring is important for modelling any action or interaction whatsover. For example, The Painting Fool compares an initial altered snapshot (sensory data) to the painted image that it generates in response to that snapshot, and judges the quality of its output on that basis (Colton and Ventura 2014). This example could be extended to theorise "proprioceptive" sensing and judgement about effected actions more broadly. Filtering upstream data is another simple application of sensors, which Keller (2012) describes as an "ecocompositional technique." In short, an ecological view on monitoring suggests that it can be distributed out among participants and that this is vital for social creativity.

**4B. Monitoring the resource**

"Tests can be modified by Producers or their representatives."

The environment itself also filters and selects (Kockelman 2011). Some of these conditions are fatal for living beings in the environment, and more broadly may provide terminating conditions for the constituent Processes in a Place. It would be too much to say that *all* tests can be modified by Producers. Rather, Producers may have programmatic access to those tests which transform potentially fatal (or at least fateful) features of the Place and participating Producers into *data*. This opens up the possibility of directly modifying decision making processes on the one hand, or of passing along information about the fitness landscape to future generations of Producers in a (co-) evolutionary framework on the other (DeLanda 2011).

Simply put, *data* is lack of uniformity within some context (Floridi 2016). In the case of monitoring the extractive use of CPRs, direct and compelling feedback about instances of non-uniform or otherwise aberrant resource usage define critical (i.e., decisive) points within a resource management structure. In creative contexts "critique" is no less important.

**5. Graduated sanctions**

"Producers who violate operational rules in the domain will be assessed sanctions by other Producers."

Economic sanctions are generally punishments, which are presumed to have a clear meaning or a direct impact on behaviour. However, there are other cases in which Producers' interactions with other Producers will not be punitive so much as, for example, educative or otherwise formative.

In an artistic context, "sanctions" may range from constructively critical reviews to outright condemnation to no response at all. The *Iterative Development-Execution-Appreciation* (IDEA) cycle (Colton, Charnley, and Pease 2011) introduces *well-being* and *cognitive effort* ratings from which several derived measures of audience response can be computed (e.g., by averaging across audience members). This can readily be extended to a developmental or peer production context. "Audience" might be re-thought as a "public," or as Rhodes's "press" (as originally formulated) to capture the idea that its response has a direct effect on the Producer. Inasmuch as the Producer is produced, feedback from the "parent" Producer(s) is especially important to this formative Process.

**6. Conflict-resolution mechanisms**

"Producers have rapid access to low-cost local arenas to resolve conflicts."

Wikipedia's edit wars provide a familiar example (Viegas et al. 2007; Yasseri et al. 2012). These are carried out on the pages of the encyclopedia itself, and resolved using supplementary pages. Machine-generated metadata is relied upon throughout. These mechanisms are low cost: the "stigmergic" self-organisation patterns exemplified by open online communities make fairly minimal demands on participating agents (Heylighen 2015). Nevertheless, structure matters: cases of direct and unresolvable conflict must usually be referred a higher authority, e.g., sitewide guidelines and policies, or available arbitration committees. Opportunities to jointly exercise partial control are, again, often Products, and the creation of a communication channel – a Place within a Place – is another formative Process, which Jakobson (1960, p. 355) calls the "phatic function."

The theme of local scale suggests more and less representative examples. For instance, academic research is currently organised in a much more segmented and localised format than Wikipedia. *Modularity* is one of three features that are hypothesised to support *commons based peer production* (CBPP) (Benkler 2002). However, CBPP requires not just decomposability into modules but relatively fine *granularity* of these modules, and as well as a *low cost of integration* to bring disparate pieces of work together once they are completed – possibly "subsidised" by an assistive technology, like Wikipedia's metadata systems. Creative and scientific writing, at the level of individual papers or books, tends to miss features that would allow this work to scale up (Kim, Cheng, and Bernstein 2014) – even though science and literature represent impressively huge "virtual" collaborations.

The most straightforward test related to this theme is that a Producer needs to be able to *detect* conflict, either between itself and other Producers, or between incompatible goals. In order to resolve a conflict – or to organise work on a project to avoid conflicts in the first place – a Producer will probably need to reason about the project's structure.

### 7. Minimal recognition of rights to organise

"The rights of Producers to devise institutions governing their contributions are not challenged by external authorities."

The foremost external authority to be concerned about in a computational creativity setting is the programmer. A "mini me" critique can readily be levelled by CC sceptics (Colton 2012). We are still in early days for autonomous creative systems and general AI, and involvement of programmers and others in teaching systems how to devise institutions is at least as relevant as teaching them how to conform to pre-given instructions.

Keeping in mind the earlier reflections on Winnicott, a relevant set of tests would compare the frequency of user- or programmer-generated changes in the system, with the frequency of changes coming from the system itself. This is the thrust of the diagrammatic formalism of creative acts developed by Colton et al. (2014): with considerable further work we could expand the ability of computer systems to participate in, or fully automate, such modelling activity. A basic challenge in applying the formalism from Colton et al. is to identify the individual "creative acts" that a given Producer has made. The tests that would reveal these acts in a given stream of Products tend to be domain-specific.

### 8. Nested enterprises

"Contribution, testing, enforcement, conflict resolution, and governance and are organised in multiple layers of nested Places and agencies."

That the Place or the Producer would be layered isn't a surprise; many systems have a hierarchical aspect. What is perhaps more surprising is that many of the features that make up a "creative ecosystem" must themselves be *produced*, which points to the inherent multiplicity of Producers. Here, Producers are seen as self-organising the structure of their interrelationships and interconnections at various levels. Developing a computational treatment of such a system divorced from real world applications would be a thankless and ultimately futile task. Effort may be better spent on developing programs that model and participate in existing creative ecosystems. In such cases, there would be real-world empirical tests of success, coming from users.

## Example

This section uses the creativity design principles discussed above to describe some of the creativity-supporting institutions in place at the Seventh International Conference on Computational Creativity (ICCC 2016), and to explore potential additions and adaptations for future ICCCs.

**I.** The crucible for the current paper was a unique set of ongoing discussions (see "Acknowledgements") (2A). At first, the hope was to co-author the paper with one of these discussants, but due to time constraints this was not possible, so it became a single-author paper (1A). The ICCC call helped motivate writing up the ideas (2B), partly because the conference is open to papers that are informed by and contribute to various disciplines at varying degrees of formality. However, ICCC enforces rigourous academic standards, using slightly different evaluation criteria for papers submitted to each of five "tracks" (1B). Reviewers used the Easychair website to bid for papers to review, and to share discussions and debate about these papers in case of disagreement (6). Papers that were seen as less relevant were rejected outright, or potentially (as a norm) allocated briefer slots in the conference schedule (5). The current paper was conditionally accepted, which meant that it entered into a "shepherding" process, whereby a senior programme committee member could check (4A) whether the author followed through on specific reviewer requirements (4B). By and large authors are given free rein to write papers about any topic relevant to computational creativity, if they do so in a rigorous academic style (7). This entails reflecting on certain themes-held-in-common – but the conference seems to lose some opportunities for structuring engagement more deeply, e.g., around common tools or challenge problems (8?). Presumably only the conference steering committee can change the conference's overall rules; however, it should be noted that reviewer requirements constitute fine-tuned rule-setting at the level of individual papers (3?).

**II.** The reflections above begin to suggest ways in which we might make better use of software systems in creative partnership. One realistic idea would be to use computer programs to help with paper review tasks. Essay grading software is now mainstream, and services like WriteLab can help authors simplify their writing and catch grammar and logic errors.[4] Agent-based reviews or a shift to *post-publication review*, in which reviews are offered "after an article is published, much like commentary on a blog post" (Ford 2013, p. 316) would change the population of reviewers (1A). Moving beyond blogs to wikis, lists of open problems from prior publications could be collected, compared, and explicitly referenced with semantic links (1B) (Tomlinson and others 2012). This could begin to make explicit the ways in which a given paper constitutes an advance (2A, 2B). The development, use, and maintenance of shared tools (APIs, open source software) and design patterns for computational creativity could be encouraged (3). A standardised testing approach based on challenge problems, as in the

[4] `https://writelab.com` offers a freemium service for students, but is "always free for instructors."

recently announced OpenAI Gym,[5] with worked examples, explicit evaluation metrics, and variant versions (4A, 4B) could help the community move towards, and enforce, standards of *replicability* and *generalisability* (5). Partial "wikification" and semantisation of the research area is already underway with systems like FloWr (Charnley, Colton, and Llano 2014) and ConCreTeFlows (Znidaršič et al. 2014), but it is unclear whether these systems will merge, or diverge, or if a new standard will come along (6?). Once shared technologies and datasets are in common use, computational agents will be better able to contribute to the field (7). It is to the advantage of computational creativity researchers to develop applications and application environments that we – and others – agree are useful (8).

## Discussion and Conclusions

This section reviews the contribution above, beginning with a link to related work. Specifically, Ostrom's high-level Institutional Analysis and Development (IAD) framework can be regarded in parallel with the high-level outline of the *Standardised Procedure for Evaluation of Creative Systems* (SPECS) (Jordanous 2012). SPECS suggests that, in order to evaluate creativity, it is necessary to put forth a definition of "what it is to be creative," and then to specify criteria by which creativity will be measured before formulating the evaluation. IAD suggests that institutions operate within a certain context, which afford certain kinds of actions, and that these lead to certain observable outcomes. To wit:

| IAD | SPECS |
|---|---|
| Context | Definition |
| Action | Criteria |
| Outcome | Evaluation |

In IAD, context can be thought of as a collection of "exogenous variables," (Ostrom 2009, p. 13, esp. Figure 1.1) including pre-defined rules, that shape what happens in the action situation at the heart of the analysis. We have described several candidate design principles that outline potential rules for guiding action in creative settings. This suggests the possibility of recording a definition and set of criteria for evaluating social creativity in a general domain. Pragmatically, this definition might unpack the 4Ps in terms of Ostrom's variables (participants, positions, etc.).

SPECS could be criticised for being overly abstract: in other words, for simply describing good practice in any empirical investigation. IAD adds many more specifics, which have necessarily been presented in a compressed form here. It is hoped that this first attempt to use IAD to theorise computational social creativity will motivate future explorations that further unpack social creativity using Ostrom's ideas.

The creativity design principles offer guidelines (and with minor changes, hypotheses) for members of the computational creativity community to test out in practice. More empirical work is needed to validate (or improve) these principles. On the cultural side, more attention should be given to the fact that our institutions – including institutions for building institutions – are analysable in programmatic terms.

[5] `https://gym.openai.com/`


## Acknowledgements

This work was supported by the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET-Open Grant number 611553 (COINVENT).

Ongoing conversations with Anna Jordanous, Steve Corneli, and Rasmus Rebane about computational creativity, economics, and semiotics (respectively) were key drivers in this paper's evolution. The anonymous reviewers made many useful comments on matters of presentation.


## References


Antunes, R. F.; Leymarie, F. F.; and Latham, W. 2015. On writing and reading artistic computational ecosystems. *Artificial life* 21(3).

Benkler, Y. 2002. Coase's Penguin, or Linux and the Nature of the Firm. *Yale Law Journal* 112:369.

Brézillon, P. 1999. Context in problem solving: a survey. *The Knowledge Engineering Review* 14(1):47–80.

Chabot, P. 2013. *The philosophy of Simondon: Between technology and individuation*. Bloomsbury.

Charnley, J.; Colton, S.; and Llano, M. T. 2014. The FloWr Framework: Automated Flowchart Construction, Optimisation and Alteration for Creative Systems. In Ventura et al. (2014).

Clark, A. 1998. *Being there: Putting brain, body, and world together again*. MIT press.

Clark, A. 2001. *Natural-born cyborgs?* Springer.

Colton, S., and Ventura, D. 2014. You Can't Know my Mind: A Festival of Computational Creativity. In Ventura et al. (2014).

Colton, S.; Pease, A.; Corneli, J.; Cook, M.; and Llano, T. 2014. Assessing Progress in Building Autonomously Creative Systems. In Ventura et al. (2014).

Colton, S.; Charnley, J.; and Pease, A. 2011. Computational creativity theory: The FACE and IDEA descriptive models. In *Proceedings of the Second International Conference on Computational Creativity*, 90–95.

Colton, S. 2012. The Painting Fool: Stories from Building an Automated Painter. In *Computers and creativity*. Springer. 3–38.

Conway, M. E. 1968. How do committees invent? *Datamation* 14(4):28–31.

Cox, M.; Arnold, G.; and Tomás, S. V. 2010. A review of design principles for community-based natural resource management. *Ecology and Society* 15(4):38.

Deacon, T.; Haag, J.; and Ogilvy, J. 2011. The emergence of self. In Wentzel Van Huyssteen, J., and Wiebe, E. P., eds., *In Search of Self: Interdisciplinary Perspectives on Personhood*. Wm. B. Eerdmans Publishing Co.

Deacon, T. W. 2014. In what sense could a machine be alive? Thursday 3rd April, 9:30-10:30, AISB'50, Convention Plenary, Goldsmiths College, University of London.

DeLanda, M. 2011. *Philosophy and simulation: the emergence of synthetic reason*. Continuum.

Feldman, D. H.; Csikszentmihalyi, M.; and Gardner, H. 1994. *Changing the world: A framework for the study of creativity.* Praeger Publishers/Greenwood Publishing Group.

Floridi, L. 2016. Semantic conceptions of information. In Zalta, E. N., ed., *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, Center for the Study of Language and Information (CSLI), Stanford University.

Ford, E. 2013. Defining and characterizing open peer review: A review of the literature. *Journal of Scholarly Publishing* 44(4):311–326.

Gervás, P., and León, C. 2014. Reading and Writing as a Creative Cycle: the Need for a Computational Model. In Ventura et al. (2014).

Gundersen, O. E. 2014. The role of context and its elements in situation assessment. In Brézillon, P., and Gonzalez, A. J., eds., *Context in Computing*. Springer. 343–357.

Heylighen, F. 2015. Stigmergy as a Universal Coordination Mechanism: components, varieties and applications. *Human Stigmergy: Theoretical Developments and New Applications*.

Hirst, G. 2000. Context as a spurious concept. In Gelbukh, A., ed., *International Conference CICLing-2000: Conference on Intelligent Text Processing and Computational Linguistics (Proceedings)*, 273–287.

Ingold, T. 2000. *The Perception of the Environment: Essays on Livelihood, Dwelling and Skill*. Routledge.

Jakobson, R. 1960. Linguistics and Poetics. In Sebeok, J., ed., *Style in Language*. MIT Press. 350–377.

Jordanous, A. 2012. A Standardised Procedure for Evaluating Creative Systems: Computational Creativity Evaluation Based on What it is to be Creative. *Cognitive Computation* 4(3):246–279.

Jordanous, A. 2016. Four PPPPerspectives on Computational Creativity in theory and in practice. *Connection Science* 28:194–216.

Keller, D.; Lazzarini, V.; and Pimenta, M. S. 2014. Ubimus Through the Lens of Creativity Theories. In *Ubiquitous Music*. Springer. 3–23.

Keller, D. 2012. Sonic ecologies. In Brown, A. R., ed., *Sound musicianship: Understanding the crafts of music*. Cambridge Scholars Publishing. 213–227.

Kim, J.; Cheng, J.; and Bernstein, M. S. 2014. Ensemble: exploring complementary strengths of leaders and crowds in creative collaboration. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, 745–755. ACM.

Kockelman, P. 2011. Biosemiosis, technocognition, and sociogenesis: Selection and significance in a multiverse of sieving and serendipity. *Current Anthropology* 52(5):711–739.

McCay, B. J., and Acheson, J. M. 1990. *The question of the commons: The culture and ecology of communal resources*. University of Arizona Press.

Mitchell, R., and McKim, J. 2002. *Design by Contract, by Example*. Addison-Wesley.

Ostrom, E.; Chang, C.; Pennington, M.; and Tarko, V. 2012. *The Future of the Commons: Beyond Market Failure and Government Regulation*. Institute of Economic Affairs.

Ostrom, E. 1990. *Governing the commons: The evolution of institutions for collective action*. Cambridge: Cambridge University Press.

Ostrom, E. 2008. The challenge of common-pool resources. *Environment: Science and Policy for Sustainable Development* 50(4):8–21.

Ostrom, E. 2009. *Understanding institutional diversity*. Princeton University Press.

Ostrom, E. 2010. Institutional analysis and development: Elements of the framework in historical perspective. In *Historical Developments and Theoretical Approaches in Sociology*, volume 2. EOLSS. 261–288.

Rhodes, M. 1961. An analysis of creativity. *The Phi Delta Kappan* 42(7):305–310.

Ritchie, G. D. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* 17:67–99.

Rushkoff, D. 2010. *Program or be programmed: Ten commands for a digital age*. Or Books.

Safner, R. 2016. Institutional Entrepreneurship, Wikipedia, and the Opportunity of the Commons. `http://ssrn.com/abstract=2564230`.

Saunders, R., and Bown, O. 2015. Computational social creativity. *Artificial life*.

Tinnell, J. 2015. Grammatization: Bernard Stiegler's theory of writing and technology. *Computers and Composition* 37:132–146.

Tomlinson, B., et al. 2012. Massively distributed authorship of academic papers. In *CHI'12 Extended Abstracts on Human Factors in Computing Systems*, 11–20. ACM.

Ventura, D.; Colton, S.; Lavrač, N.; and Cook, M., eds. 2014. *Fifth International Conference on Computational Creativity, ICCC 2014*. Association for Computational Creativity.

Viegas, F. B.; Wattenberg, M.; Kriss, J.; and Van Ham, F. 2007. Talk before you type: Coordination in Wikipedia. In *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*, 78–78. IEEE.

Winnicott, D. W. 2002. *Playing and reality*. Routledge. (Original work published 1971.)

Wood, G. 2014. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum Project Yellow Paper*.

Yasseri, T.; Sumi, R.; Rung, A.; Kornai, A.; and Kertész, J. 2012. Dynamics of conflicts in Wikipedia. *PLOS ONE* 7(6):e38869.

Znidaršič, M.; Miljković, D.; Perovšek, M.; Pollak, S.; Kranjc, J.; Cherepnalkoski, D.; and Lavrač, N. 2014. First report on framework and data. Technical Report D6.2, ConCrete: Concept Creation Technology, Project Number 611733. ICT – Future and Emerging Technologies (FET).