

Supportive and Antagonistic Behaviour in Distributed Computational Creativity via Coupled Empowerment Maximisation

Christian Guckelsberger¹, Christoph Salge², Rob Saunders^{3,4} and Simon Colton^{1,4}

¹Computational Creativity Group, Goldsmiths, University of London, UK

²Adaptive Systems Research Group, University of Hertfordshire, UK

³University of Sydney, Australia; ⁴The Metamakers Institute, Falmouth University, UK
c.guckelsberger@gold.ac.uk

Abstract

There has been a strong tendency in distributed computational creativity systems to embrace embodied and situated agents for their flexible and adaptive behaviour. Intrinsically motivated agents are particularly successful in this respect, because they do not rely on externally specified goals, and can thus react flexibly to changes in open-ended environments. While supportive and antagonistic behaviour is omnipresent when people interact in creative tasks, existing implementations cannot establish such behaviour without constraining their agents' flexibility by means of explicitly specified interaction rules. More open approaches in contrast cannot guarantee that support or antagonistic behaviour ever comes about. We define the information-theoretic principle of coupled empowerment maximisation as an intrinsically motivated frame for supportive and antagonistic behaviour within which agents can interact with maximum flexibility. We provide an intuition and a formalisation for an arbitrary number of agents. We then draw on several case-studies of co-creative and social creativity systems to make detailed predictions of the potential effect the underlying empowerment maximisation principle might have on the behaviour of creative agents.

Introduction

If we peek into a painting class, we might observe teachers prescribing certain techniques to tackle a task, and students suggesting each other different brushes or materials to achieve their goals. While the teachers' behaviour can be considered as positive but antagonistic, we understand the students as supportive. Both forms of behaviour are closely related to constraints. Constraints limit the space of possible creative trajectories, and thereby make the exploration of an initially vast set of creative possibilities feasible. This focus also allows an agent to achieve mastery of some techniques before approaching others (Stokes, 2005). Constraints also present the challenge to achieve similar results by different means. Overcoming constraints is the key to transforming a space of possibilities, one crucial aspect of creativity (Boden, 1995). We consider *positive, but antagonistic behaviour* as imposing constraints on another agent temporarily, so they can learn to master them. When maintaining constraints permanently, this can turn into harm, and lead to the sabotage of other's creative endeavours. *Support* in contrast is present if an agent actively helps another in e.g. learning a particular technique to overcome constraints.

We suggest that embracing such supportive and both positive and negative antagonistic behaviour could advance research in distributed computational creativity (CC). Reproducing such behaviour could not only improve our understanding of human creativity, but it could also prove to be essential in the construction of genuinely autonomous creative systems. If we want artificial agents to be taken seriously as partners in creative activities, we require them to challenge us. In other words, we want them to constrain the actions we can undertake, so we can practice the mastery of the remaining and discover alternative routes. Likewise, we want such systems to help us escape from situations where we are very limited in our potential interactions with a certain medium. Negative antagonistic behaviour might allow for the emergence of cliques and diverging creative paths.

The challenge of realising such behaviour in a distributed CC system lies in finding a way to formalise and operationalise support and antagonism in an interactive and dynamic context, and preferably allow for the flexible and seamless transition between these two modes. We need to define these behaviours in a generic way, that allows agents to act in open-ended environments without clear goals, which are commonplace in CC. Such a situation cannot be mastered with predetermined behaviour, and we will address this issue by means of intrinsically motivated agents.

In this paper, we analyse existing co-creativity and social creativity systems as representatives of distributed CC, and conclude that there is no means yet to foster supportive and antagonistic behaviour in such agents without prescribing specific interactions, and thereby limiting the agents' flexibility. We suggest to use the information-theoretic, intrinsic motivation of *empowerment* (cf. Salge, Glackin, and Polani, 2014) to formalise the degree to which an agent is constrained in a creative activity. As our main contribution, we define the principle of coupled empowerment maximisation (CEM) as a generic mechanism to enable the emergence of supportive and antagonistic behaviour in distributed CC systems, without putting explicit constraints on the types of interactions. Empowerment corresponds to an agent's potential influence on the environment at a certain time. Coupled empowerment maximisation consequently motivates an agent to act in a way which maximises or minimises this capacity for other agents. Importantly, it allows to seamlessly shift between supportive and antagonistic behaviour.

Background

Co-creative and social creativity systems are only meaningful if each agent has a different perspective on a shared world, allowing them to complement each other, and for creativity to emerge from their interaction. Only embodied, situated and intrinsically motivated agents afford such a genuinely personal perspective (Guckelsberger and Salge, 2016). We will briefly describe and motivate these notions, and relate to existing projects in the field.

Embodied and Situated Agents

There is a common notion that creativity does not occur in a vacuum (cf. Jordanous (2015)). It is a *situated* activity, in that it relates to a cultural, social and personal context. Moreover, and in line with Saunders et al. (2010), we suggest that a large portion of creative behaviour, just like other processes constituting intelligence (cf. Rosch, Thompson, and Varela, 1992), is conditioned on an agent's *embodiment*. Put differently, we suggest that creativity is structured by how an agent's morphology, sensors and actuators enable its interaction with the world.

Robots are becoming increasingly popular in CC research (cf. Saunders et al., 2010; Saunders and Gemeinboeck, 2014; Brodbeck, Hauser, and Iida, 2015). Nevertheless, being embedded in the physical environment is neither sufficient nor necessary for an agent to be deemed embodied and situated. It is not sufficient because a robot could be governed by a central controller alone, following a classic computationalist approach. In contrast, embodied and situated agents must implement a tight interplay between physical and information-theoretic aspects of the agent, i.e. between the sensors, actuators, limbs and the controller. Pfeifer, Iida, and Bongard (2005) note that embodiment is only given if changes to one component can affect every other; moving from a greyscale to a color camera sensor might allow an agent to differentiate the consequences of its actions further, potentially leading to more diverse behaviour.

We take the stance that embodiment does not require a physical environment, so long as a virtual environment gives rise to the same effects. Nevertheless, many studies employ robots, because their situatedness makes the simulation of a rich environment obsolete. Furthermore, a physical environment affords a more natural interaction between humans and artificial agents. It also allows for morphological computation, where part of an agent's computational burden is taken over by its morphology, e.g. by constraining its joints. Saunders et al. (2010) argue that taking advantage of the physical world can expand an agent's behavioural range.

Being embodied and situated comes with a restricted access to the world, i.e. an agent can only perceive and affect parts of it. This leads to the emergence of an *Umwelt* (Von Uexküll, 1982), i.e. an agent's world of significance, which shapes its intrinsically motivated goals or the way that extrinsically motivated goals are perceived. Changing an agent's embodiment can change its *Umwelt*, and therefore also the way it interacts with the world and other agents. Pickering (2005) argues that this embodied and situated perspective leads to creativity when exploiting opportunities,

and overcoming embodiment-relative constraints in an environment. We believe that this systemic view represents the main motivation for distributed CC, over any mere engineering concerns. Here, creativity emerges from the interaction of multiple agents, both human and artificial, with different perspectives on the world, and on potentially shared tasks.

Embodied and situated agents also challenge the mini-me problem in CC, i.e. the problem that creativity is often attributed to the designer instead of the artificial agent. The behavioural complexity of embodied and situated agents is to a large extent determined by their interaction with the environment. Instead of explicitly programming, we have to engineer for emergence, leading to more robust behaviour which might be novel and surprising even for the designer.

Intrinsic Motivation

Pickering (2005) argues that human creativity cannot be properly understood, or modelled, without an account of how it emerges from the encounter between the world and intrinsically active, exploratory and productively playful agents. *Intrinsic motivation* was first named by White (1959) while observing animals engaging in such behaviours in the absence of an obvious reinforcement or reward. Ryan and Deci define the term from a psychological point of view as "Performing an activity for its inherent satisfactions rather than for some separable consequence" (Ryan and Deci, 2000). Being *extrinsically motivated* in contrast means to perform an activity for an externally prompted, instrumental value. Oudeyer and Kaplan (2008) complement this view with a definition informed by robotics and AI. The converging point is the reliance on the sensorimotor flow and agent-internal experience alone, independent of the involved channels' semantics. Intrinsically motivated agents are not dependent on externally defined goals, but can still form goals intrinsically. This allows for higher flexibility and adaptivity especially in open-ended environments which are commonplace in creative activities. Intrinsic motivation was identified in philosophy (Kieran, 2014) and in psychological experiments as an important factor in producing more creative artefacts (Amabile, 1985), by driving the exploration of creative options.

Related work in CC focusses mainly on the notions of curiosity and novelty, but also on surprise (Maher, Brady, and Fisher, 2013) and expectation (Grace and Maher, 2014, 2015). In co-creativity and social creativity, models of curiosity is particularly popular: Saunders (2007) developed a system of curious design agents which evolve abstract art. In *Curious Whispers*, intrinsically motivated robots generate and play music to each other (Saunders et al., 2010). Merrick and Maher (2009) employ curiosity to support the learning of tasks in adaptive characters in multiuser games. In *Accomplice*, Saunders and Gemeinboeck (2014) establish a playful interaction of curious robots with a human audience.

Co-Creativity and Social Creativity

In creativity studies, co-creativity refers to several people contributing to the creative process in a blended manner (Candy and Edmonds, 2002). In this paper, we will use the term for the more specific *human-computer co-creativity*

(Davis, 2013), describing the interaction of one person or multiple people with one or more artificial agents to generate a creative product. There are many subcategories such as mixed initiative systems, live algorithms and collaborative AI for artistic tasks, each stressing different aspects such as the order of interaction, time constraints, or the task concerned. Much research has been done on robotic live music improvisation, e.g. *Ja'maa*, a modification of the percussion robot *Haile* (Weinberg, Driscoll, and Thatcher, 2006), and the interactive Marimba player *Shimon* (Hoffman and Weinberg, 2010). Other researchers look at co-creativity in sketching: In the *Drawing Apprentice* system, a person and a software agent take turns to add to a virtual canvas (Davis et al., 2014). Jacob and Magerko (2015) investigate human-computer co-creativity in dance and interactive art by means of the *Viewpoint AI* system. Here, a human performer and virtual agent collaborate to improvise movements in real-time. A co-creative system which is less about artistic tasks is the ongoing *Computational Play Project*, where robots will eventually engage with children and toys in pretend play, i.e. the “subsequent enactment of a narrative experience using physical objects” (Magerko et al., 2014).

Within CC, the notion of *social creativity*, comprising creative cultures, creative societies, and computer social creativity, refers to computer-computer interaction, i.e. groups of artificial agents which produce and share artefacts. Co-creativity thus represents the fundamental mechanism in social creativity systems, if understood as creative interactions between purely artificial agents. There are two overlapping perspectives on the use of social creativity systems: One is inspired from research in artificial life and sociology, and employs systems as testbeds for investigating social factors in human creativity. Here, the produced artefacts are of minor interest (cf. Saunders and Bown, 2015). For instance, Steels (1995) as well as Saunders and Grace (2008) study the emergence of shared vocabularies and the formation of agent cliques engaging in “language games”.

The second perspective is directed towards the development of autonomous creative systems, and considers the mechanisms inherent in social creativity, e.g. dialogue, reflections and multiple perspectives, as means to achieve this goal (cf. Corneli et al., 2015). Such systems could produce valuable artefacts, but their value and novelty might be intrinsic to the system, i.e. only meaningful to the artificial agents themselves. They often draw on concepts from cognitive science such as the *Blackboard* architecture which is inspired by the *Global Workspace* model (Baars, 2005). The latter expresses the idea that distributed sources of knowledge or different roles, represented by competing mental processes, can be leveraged to cooperatively solve problems that no single constituent could solve alone.

The Blackboard architecture is particularly popular in poetry and narrative generation, but these systems struggle to incorporate competition and cooperation in a flexible and adaptive way. Only few implement a strong coupling between an agent’s body and its environment: The story generator by Laclaustra et al. (2014) creates stories by recording the interaction of multiple agents with different roles. Eigenfeldt investigated music in social creativity, both by means

of the software agent ensembles *Drum Circle* (Eigenfeldt, 2007) and *Musebots* (Eigenfeldt, Bown, and Carey, 2015), which other researchers can modify to produce a collective composition. A physical realisation of an ensemble is given by the 12 arm drum robot *MahaDeviBot*, where each limb is controlled by one agent (Eigenfeldt and Kapur, 2008).

Social creativity systems usually employ at least one agent to direct the flow of actions. For instance, Laclaustra et al. (2014) define the role of a “director” in their story generation system, which sets new goals for the acting agents. In the *Virtual Storyteller*, Theune et al. (2003) use a director to introduce new characters and objects, give characters specific goals, and deny them to perform certain actions. Similarly, Eigenfeldt’s ensembles employ a conductor agent to control the composition (Eigenfeldt, Bown, and Carey, 2015). In co-creativity, the human is usually, but not always in control and introduces goals into the system: *Curious Whispers* encourages people to interact with the robots via a synthesiser. While some take a traditional “master” role in trying to teach the robots tunes, others act more passively and try to learn from- and copy the robots. Our formalism is designed to foster sensible agent behaviour in both cases.

Case-Studies

We conducted three case-studies to analyse if and how present co-creativity and social creativity systems realise supportive and antagonistic behaviour. We evaluated systems which situate intrinsically or extrinsically motivated agents in a physical or virtual environment.

Curious Whispers Developed by Saunders et al. (2010), *Curious Whispers* represents a society of intrinsically motivated robots which generate and listen to tunes. The main goal of this social creativity system is to investigate the effects of embodiment on creativity in a physical environment. Each robot is equipped with a pair of microphones, touch sensors to avoid collisions, a speaker and four wheels. The robots are driven by an intrinsic measure of interest- ingness, which is quantified by mapping the novelty of the current sensor input on a Wundt curve. Novelty is quantified by comparing how new percepts are encoded in a self-organising map serving as the robot’s long-term memory. The robots can listen to two sources of sound at a time, and move closer to the one which is considered more interesting.

Saunders et al. (2010) suggest that engaging in social relations represents one crucial means of how embodiment can foster creativity. In their system, such relations remain shallow: robots play their tunes when getting bored of listening to others, and might consequently be engaged by other robots which find their tune interesting. The programmed behaviours of the individual agents in *Curious Whispers* are deliberately minimal, so any supportive and antagonistic behaviour would be an emergent property of the system. Exposure to tunes biases the robots in the generation of new instances. They thus appear to engage in mutual support to explore the space of potential tunes as they are passed between and modified by each other. Nevertheless, there is no apparent antagonistic behaviour, and at no point do the agents act to directly influence the performance of others.




Figure 1: Causal Bayesian network representing the interaction of a memoryless agent with the world as a perception-action loop unrolled in time. Arrows denote causal dependencies between its sensors S , actuators A and the rest of the world R .

Shimon *Shimon* (Hoffman and Weinberg, 2010) is an extrinsically motivated, robot marimba player which improvises in real-time to a human pianist’s performance. It consists of four arms which can be moved separately on a shared rail in parallel to the marimba’s front side. Each arm has one mallet for the bottom-, and one for the top-row keys. A MIDI-listener attached to the electric piano and an adjustable metronome are used as sensors. The project implements embodied cognition by understanding music not as a sequence of notes, but as a choreography of movements constrained by the robots morphology. The performance is based on different interaction modules which analyse and react to the sensory input. Each responds to a different challenge, e.g. to react in time with the right tempo, or to play beat-matched, synchronised and chord-adaptive patterns.

The description of a live performance sheds light on the nature of interaction. Each interaction module matches a phase in performance, and is addressed independently through the human player who provides the notes, beat and tempo. It lasts until a certain condition, e.g. a limit of played bars, is met. This rigid setup, together with a preset of rhythms and a pre-programmed crescendo finale, shows that provided a time and stimulus, the robot performs a prescribed set of hard-coded interactions to establish support. At no time does it challenge the player.

Drawing Apprentice Davis et al. (2015) introduce a co-creative system in which an extrinsically motivated software agent and a person take turns to improvise line drawings on a virtual canvas. The *Drawing Apprentice* is based on a cognitive architecture inspired by Enactivism, which considers creativity as emerging from an improvised interaction with the environment and other agents. It differentiates three types of awareness, which are associated with different layers of perceptual logic. The system receives a line input from the user, analyses and adopts the perceptual layer the user is currently in, and generates an improvised response. Each layer focusses on a different scale of the drawing, determines how and over which timespan the system will analyse the user’s input, and puts constraints on the possible responses. For instance, the local logic only takes the user’s last input into account, and complements it by mirroring, scaling or translation. The regional logic in contrast analyses a series of past strokes and employs gestalt principles to group them, while the global logic analyses the whole composition.

The *Drawing Apprentice* reflects *Shimon*’s system architecture to some extent, as it constrains potential responses by means of dedicated modules. Nevertheless, it autonomously selects which module to perform. The system arguably supports and inspires the user in their activity, by complementing their drawing in interesting ways. Nevertheless, the sys-

tem is explicitly grounded in such supportive behaviour by design. Some responses might feel like a constraint to the user, but the design does not seem to embrace such antagonistic behaviour explicitly. Being extrinsically motivated, the system can only react in previously anticipated ways.

Summary The case-studies show us that present systems with intrinsically motivated agents exhibit emergent and thus highly flexible and adaptive behaviour, but do not have a means to establish a truly supportive or antagonistic mode of interaction. Systems with extrinsically motivated agents prescribe such interactions rigidly, but are limited to what the system designer anticipates as supportive in a certain situation beforehand, which is particularly difficult in a physical environment without a clearly defined interaction interface or fixed goals. Importantly, no project realises antagonistic behaviour. Next, we will introduce the CEM principle to overcome this situation. We later recall the case-studies and show how the principle could apply.

Formal Model

We propose the CEM principle as a candidate mechanism to enable the emergence of supportive and antagonistic behaviour in co-creative and social creativity systems, without putting explicit constraints on the interactions. We first provide an intuition and a formal definition of empowerment and the empowerment maximisation (EM) principle, followed by a formalisation of CEM and an algorithmic description for a scenario with two agents.

Empowerment and Empowerment Maximisation

Empowerment, the quantity underlying the CEM principle, is defined over the relationship between an agent’s actuators and sensors, and as such is sensitive to the agent’s embodiment and *Umwelt*. In a deterministic environment, empowerment quantifies an agent’s options in terms of availability and visibility. In a stochastic setting, this generalises to the potential influence of an agent’s actions on its environment, and to the extent to which it can perceive this influence afterwards. Empowerment is measured in bits of information (Shannon, 1948). It is zero when the agent has no control over its sensors, i.e. when all actions lead to the same perception, and it grows when different actions lead to different perceivable outcomes. For simplicity, the interaction presented here is discrete in time and space. Nevertheless, continuous implementations exist and were evaluated both in virtual environments and in robotics. An introduction to empowerment with a survey of motivations, intuitions and past research can be found in (Salge, Glackin, and Polani, 2014).

At the centre of the empowerment definition is the interpretation of an agent’s embodiment as an information-




Figure 2: Perception-action loop for two agents $(S^1, A^1), (S^2, A^2)$ interacting in turnwise order. The first agent (S^1, A^1) is coupled to the second (S^2, A^2) . Dotted (orange) lines indicate the estimation of future sensor states and dashed (purple) lines represent the calculation of empowerment, which comprises further estimation steps.

theoretic communication channel. For any arbitrary separation between an agent and a world we can define sensor variables S and actuator variables A , as those states that allow for the in- and outflow of information to the agent, respectively. This interaction with the world is usually described as a perception-action loop (Touchette and Lloyd, 2004), which can be analysed by means of a causal Bayesian network as in Fig. 1 and Pearl’s interventional calculus (Pearl, 2000). In the figure, arrows imply causation between random variables: the agent’s actions A only depend on its sensor S , which in turn is determined by the rest of the system R . The latter is affected both by the preceding system state and the agent’s actions. The interventional causal probability distribution $p(S_{t+1}|S_t, A_t)$ therefore represents the (potentially noisy) communication channel between actions and future sensor states.

Empowerment is then defined as the maximum potential information flow (Ay and Polani, 2008) that could possibly be induced by a suitable choice of actions, in a particular state s_t . This can be formalised as the channel capacity:

$$\begin{aligned} \mathfrak{E}_{s_t} &= \max_{p(a_t)} I(S_{t+1}; A_t) \\ &= \max_{p(a_t)} H(S_{t+1}) - H(S_{t+1}|A_t) \\ &= \max_{p(a_t)} \sum_{A, S} p(s_{t+1}|s_t, a_t) p(a_t) \log \frac{p(s_{t+1}|s_t, a_t)}{\sum_A p(s_{t+1}|s_t, \hat{a}_t) p(\hat{a}_t)} \end{aligned}$$

Here, $I(S_{t+1}; A_t)$ represents the mutual information between sensors and actuators, which is based on the difference of regular $H(S_{t+1})$ and conditional entropy $H(S_{t+1}|A_t)$. For details on these information-theoretic notions, see (Cover and Thomas, 2006).

Empowerment is *local*, i.e. the agent’s knowledge of the local dynamics $p(S_{t+1}|S_t, A_t)$ are sufficient to calculate the quantity. The information-theoretic grounding makes it domain-independent and universal, i.e. it can be applied to every possible agent-world interaction, as long as this interaction can be modelled as a probabilistic perception-action loop. This implies that it can be computed on arbitrary agent morphologies, and can cope with changes being made to it. Finally, empowerment as presented here is task-independent, i.e. it is not evaluated relative to a specific goal.

Crucially, empowerment does not measure an agent’s actual, but their potential influence on the environment. The EM principle suggests that an agent should, in absence of any explicit goals, choose actions which are likely to lead to

states with higher influence on the environment, i.e. more options. A greedy agent would thus choose the action with the highest expected empowerment, i.e. which most likely leads to future sensor states with maximum empowerment. Based on the properties outlined in the previous paragraph, EM satisfies the criteria for intrinsic motivation by Oudeyer and Kaplan (2008), which were outlined earlier.

Coupled Empowerment Maximisation

Coupled Empowerment Maximisation is defined on the embodiment of multiple agents, and represents an extension of the general maximisation principle. It is based on the observation that an agent’s actions might not only affect its own, but also the empowerment of other agents. CEM represents an action policy which explicitly considers this relationship. While we assume that an agent would always maximise its own empowerment over the long term, we suggest to look at how both maximising and minimising other agents’ empowerment shapes the behaviour of individual agents and groups. Given the general intuition of empowerment from the previous section, we hypothesise that minimising or maximising another agent’s empowerment establishes a general frame for supportive and both positive and negative antagonistic behaviour, respectively. A positive, antagonistic agent constrains others temporarily in order to benefit them over the long term. A negative, antagonistic agent in contrast maintains these constraints permanently.

We will define the coupled empowerment for an arbitrary number of agents, but limit our examples to two. For simplicity, we will also assume agents to interact in a turn-wise order. Fig. 2 shows the extended perception-action loop for two interacting agents (S^1, A^1) and (S^2, A^2) . Here, agents are considered as distinct from the rest of the world R . Due to the turnwise interaction, they do not have to account for their sensor states at intermediate stages where they are not permitted to act, e.g. the second agent at $t + 2$, and we consequently omitted these variables. The diagram shows that by performing a certain action a_t at time t , the first agent potentially affects both the second agent’s sensor state at $t + 1$, and its own sensor at $t + 2$, which in turn also depends on the second agent’s action choice at $t + 1$.

CEM suggests that the active agent chooses its actions in order to both maximise its own expected empowerment and to maximise or minimise the empowerment of the coupled agents. This is formalised by Eq. 1, the general action selection policy. Here, parameters α_i determine the influ-

$$\pi(s_t) = \arg \max_{a_t} [\alpha_n \cdot \pm E[\mathfrak{E}^n]_{a_t} + \alpha_{n-1} \cdot \pm E[\mathfrak{E}^{n-1}]_{a_t} + \dots + \alpha_1 \cdot E[\mathfrak{E}^1]_{a_t}] \quad (1)$$

ence of individual couplings. We use the notion of *expected* empowerment here, because the active agent cannot be sure about how the other agents might behave. The calculation of coupled empowerment therefore involves several estimation steps, which are illustrated by means of Alg. 1 and Fig. 2. For the supportive case and two agents, the active, first agent has to calculate the expected coupled empowerment of each of its actions a_t . As a first step, the agent has to estimate which potential follow-up sensor states of the second agent S_{t+1}^2 can be reached via a_t . The agent then has to take into account how the second agent could potentially act, in order to estimate its own future sensor state S_{t+2}^1 . This is indicated by the dotted (orange) lines in Fig. 2. From there, the agent has to perform another round of estimations in order to infer the local dynamics $P(S_{t+3}^2|a_{t+2}, s_{t+2})$ and $P(S_{t+4}^1|a_{t+3}, s_{t+3})$, which eventually enable the calculation of both agents' empowerment (dashed, purple lines in Fig. 2). Finally, the agent has to calculate the expected coupled empowerment $E[\mathfrak{E}^C]_{a_t}$, by combining its own $E[\mathfrak{E}^1]_{a_t}$ and the second agent's expected empowerment $E[\mathfrak{E}^2]_{a_t}$, given the current action a_t .

CEM is not constrained to a particular number of agents; nevertheless, the computational complexity grows exponentially the more agents are involved. Note that this is not problematic if we employ several empowerment maximising agents e.g. in a social creativity system, as long as each agent is only coupled to a small number of others. Different means of optimisation exist, e.g. based on monte-carlo techniques (Salge, Glackin, and Polani, 2014), the information-bottleneck method (Anthony, Polani, and Nehaniv, 2014) and deep neural networks (Mohamed and Rezende, 2015).

Algorithm 1 Calculating the action policy of the first agent in a two-agent scenario, based on supportive CEM.

```

function  $\pi(s_t, \alpha)$ 
  for all  $a_t \in A_t^1$  do
    Estimate  $P(S_{t+1}^2|a_t, s_t)$ 
    for all  $s_{t+1} \in S_{t+1}^2$  do
       $\mathfrak{E}_{s_{t+1}}^2 \leftarrow \text{CALCEMPowerment}(s_{t+1})$ 
      for all  $a_{t+1} \in A_{t+1}^2$  do
        Estimate  $P(S_{t+2}^1|a_{t+1}, s_{t+1})$ 
        for all  $s_{t+2} \in S_{t+2}^1$  do
           $\mathfrak{E}_{s_{t+2}}^1 \leftarrow \text{CALCEMPowerment}(s_{t+2})$ 
        end for
      end for
    end for
     $E[\mathfrak{E}^2]_{a_t} \leftarrow \sum_{s_{t+1}} P(s_{t+1}|a_t, s_t) \times \mathfrak{E}_{s_{t+1}}^2$ 
     $E[\mathfrak{E}^1]_{a_t} \leftarrow \sum_{s_{t+1}} P(s_{t+1}|a_t, s_t) \times \sum_{a_{t+1} s_{t+2}} P(s_{t+2}|a_{t+1}, s_{t+1}) \times \mathfrak{E}_{s_{t+2}}^1$ 
     $E[\mathfrak{E}^C]_{a_t} \leftarrow \alpha \times E[\mathfrak{E}^2]_{a_t} + (1 - \alpha) \times E[\mathfrak{E}^1]_{a_t}$ 
  end for
  Perform  $a_t : a_t = \arg \max_{A_t} E[\mathfrak{E}^C]_{a_t}$ 
end function

```

Coupled Empowerment Maximisation in Computational Creativity

Saunders and Gemeinboeck stress that intrinsic motivation is at the core of the creative process, when agents engage in “a reflective exploration of possibilities” (Saunders and Gemeinboeck, 2014). Empowerment quantifies the possibilities available to an agent in a certain situation in a very generic way. Klyubin, Polani, and Nehaniv (2008) argue that empowerment maximisation could be realised by, or even help constituting specialised motivations, such as curiosity and novelty which CC research focused on in the past. The goal of this section is to provide an intuition of CEM and motivate its potential in embodied and situated distributed CC by recalling the previous examples and case-studies.

Supportive and Antagonistic Behaviour Revisited

We suggest that CEM establishes a generic frame for supportive and antagonistic behaviour to emerge in distributed CC. We already motivated the potential benefits of negative antagonistic behaviour in a social creativity scenario, but not for human-computer co-creativity: Since a positive antagonistic agent might struggle with determining the time-frame for imposing constraints, e.g. in the presence of uncertainty, we suggest to use empowerment minimisation in co-creativity as a shortcut for positive antagonism.

Davis et al. (2015) note that from an enactivist perspective, expertise in a field is not only about knowledge, but to a large part about the mastery of an agent's sensorimotor contingencies (O'Regan and Noe, 2001). Maximising empowerment, either in respect to the own or another agent's embodiment, translates to developing more nuanced action-percept couplings, and can thus be interpreted as maximising an agent's sensorimotor expertise.

Empowerment maximisation is not goal-directed, and Klyubin, Polani, and Nehaniv (2008) hypothesise it to be a good policy in the absence of any explicit goals. Nevertheless, by coupling the maximising agent to other, goal-driven agents such as the human collaborator in a co-creativity system, or to “director” and “conductor” agents in social creativity, we can induce the goals of the coupled agent into the active agent's behaviour. We would consequently expect a CEM-driven agent to either support or sabotage the current goal of the agent it is currently coupled to. At the same time, the maximising agent takes on some control, by influencing the empowerment of the other.

Recalling the Case-Studies

This section illustrates how CEM can foster supportive and antagonistic behaviour in the earlier case studies. One way to affect empowerment is by constraining or widening an agent's options directly. An agent's empowerment is maximum, if all potential actions are available and lead to clearly separable outcomes. A positive but antagonistic *Drawing Apprentice* could challenge the human co-creator by limiting its toolbox to thick brushes only, or by restricting the

colour palette to cold tones. *Shimon* could maximise its own empowerment by moving its mallets into a position which allow it to react most flexibly to the pianist in the next time step. It could support the pianist vice versa by playing a tune which allows for many potential responses.

Empowerment can also be affected by limiting the aspects which an agent's sensor can differentiate in the environment: The positive, but antagonistic *Drawing Apprentice* might switch the output of the virtual canvas to greyscale, while maintaining the internal colour scheme. The human partner would consequently perceive many colours alike. The software agent would challenge them to practice *Grisaille*, a technique of painting exclusively in shades of grey, which was particularly popular among the old masters.

Klyubin, Polani, and Nehaniv (2008) demonstrate that empowerment can serve as an immediate guide for sensor and actuator evolution during an agent's lifetime. This requires the modifying agent to have access to the actuators and sensors. In the *Drawing Apprentice*, the virtual canvas serves as proxy to the human's perceptions and actions; *Shimon* in contrast cannot directly access the perceptual apparatus of its human collaborator. It could increase its own empowerment by evolving its actuator to apply more force to its mallets, allowing for a wider range of distinct sounds.

In *Curious Whispers*, the sensors and actuators of other robots are also inaccessible. This scenario illustrates how empowerment can be affected by modifying an agent's environment. Here, it becomes most obvious how an embodied agent's behaviour is connected to its morphology and the external environment, and how important different roles and abilities are to distributed CC. In *Curious Whispers*, a negative antagonistic agent could disturb other agents listening to a performance, by playing a noise which makes it impossible to differentiate between the different sounds that were originally played. Consequently, the listening agents will not be able to pick up the exact tune for their own performances. If the other agents were able to express a different spectrum of tones, they would maximise their empowerment by switching to the part which the antagonistic agent could not disturb, eventually leading to the differentiation of artefacts. We can think of even more complex interventions: If there were movable parts in the environment, a supportive agent might improve another agent's rehearsal by moving these parts into a position where they block noise out.

Conclusion and Future Work

We suggest to understand antagonistic and supportive behaviour in distributed CC as imposing constraints on other agents, and helping them to overcome them. We translated the information-theoretic notion of empowerment to the creative domain to formalise the degree to which embodied, situated and intrinsically motivated agents are constrained in their creative activity. We then defined the principle of coupled empowerment maximisation as a means to enable support and antagonistic behaviour in distributed CC systems. We used CEM as an *intuition pump* to demonstrate which behaviours CEM-driven agents might exhibit in three existing co-creativity and social creativity systems. Although this is one possible application, the strength of CEM lies in

its capacity to allow for the implementation and subsequent emergence of supportive or antagonistic behaviour *online*.

The emergent behaviour of embodied, situated and intrinsically motivated agents is often surprising, and hard to predict. The most important next step is therefore to evaluate the formalism in an actual co-creativity or social creativity system, and to investigate the effects under experimental conditions. The examples in this paper focus on agents in a common sense. Nevertheless, we are also interested in applying CEM to scenarios where agency is attributed to objects. This might allow us to evolve artefacts to compete with others, or to maximise synergistic effects. Our research yielded a strong correspondence between the environment that an agent can act upon, and the notion of conceptual spaces. As part of future work, we want to investigate more thoroughly how (coupled) empowerment maximisation relates to the exploration and transformation of conceptual spaces. We only considered "constraints" in terms of which actions are *possible* for an agent in a certain situation. Nevertheless, in many creative processes, options can also be limited by what is *desireable*, e.g. from an aesthetics point of view. Integrating such "soft constraints" into the formalism represents another, promising avenue of research.

Acknowledgments

CG is funded by EPSRC grant EP/L015846/1 (IGGI). CS is funded by the H2020-641321 socSMCs FET Proactive project. SC's work was supported by EPSRC grant EP/J004049 (Computational Creativity Theory).

References

- Amabile, T. M. 1985. Motivation and Creativity: Effects of Motivational Orientation on Creative Writers. *Personality and Social Psychology* 48(2):393–399.
- Anthony, T.; Polani, D.; and Nehaniv, C. L. 2014. General Self-Motivation and Strategy Identification: Case Studies Based on Sokoban and Pac-Man. *TCIAIG* 6(1):1–17.
- Ay, N., and Polani, D. 2008. Information Flows in Causal Networks. *Advances in Complex Systems* 11(1):17–41.
- Baars, B. J. 2005. Global Workspace Theory of Consciousness: Toward a Cognitive Neuroscience of Human Experience. *Progress in Brain Research* 150:45–53.
- Boden, M. 1995. Creativity and Unpredictability. *Stanford Humanities Review* 4(2):1–18.
- Brodbeck, L.; Hauser, S.; and Iida, F. 2015. Morphological Evolution of Physical Robots through Model-Free Phenotype Development. *PLoS one* 10(6).
- Candy, L., and Edmonds, E. A. 2002. Modeling Co-Creativity in Art and Technology. In *Proc. 4th Conf. Creativity and Cognition*, 134–141.
- Corneli, J.; Jordanous, A.; Shepperd, R.; Llano, M. T.; Misztal, J.; Colton, S.; and Guckelsberger, C. 2015. Computational Poetry Workshop: Making Sense of Work in Progress. In *Proc. 6th Int. Conf. Computational Creativity*, 268–275.
- Cover, T. M., and Thomas, J. A. 2006. *Elements of Information Theory*. Wiley-Interscience, 2nd edition.

- Davis, N.; Popova, Y.; Sysoev, I.; Hsiao, C.-P.; Zhang, D.; and Magerko, B. 2014. Building Artistic Computer Colleagues with an Enactive Model of Creativity. In *Proc. 5th Int. Conf. Computational Creativity*, 38–45.
- Davis, N.; Hsiao, C.-P.; Popova, Y.; and Magerko, B. 2015. An Enactive Model of Creativity for Computational Collaboration and Co-creation. In Zagalo, N., ed., *Creativity in the Digital Age*. Springer. 223–243.
- Davis, N. 2013. Human-Computer Co-Creativity: Blending Human and Computational Creativity. In *Proc. 9th Conf. AIIDE*, volume 2013, 9–12.
- Eigenfeldt, A., and Kapur, A. 2008. An Agent-Based System for Robotic Musical Performance. In *Proc. Int. Conf. New Interfaces for Musical Expression*, 144–149.
- Eigenfeldt, A.; Bown, O.; and Carey, B. 2015. Collaborative Composition with Creative Systems : Reflections on the First Musebot Ensemble. In *Proc. 6th Int. Conf. Computational Creativity*, 134–141.
- Eigenfeldt, A. 2007. Drum Circle: Intelligent Agents in Max/MSP. In *Proc. Int. Conf. Computer Music*, 2–5.
- Grace, K., and Maher, M. L. 2014. What to Expect When You're Expecting: The Role of Unexpectedness in Computationally Evaluating Creativity. In *Proc. 5th Int. Conf. Computational Creativity*.
- Grace, K., and Maher, M. L. 2015. Specific Curiosity as a Cause and Consequence of Transformational Creativity. In *Proc. 6th Int. Conf. Computational Creativity*, 260–267.
- Guckelsberger, C., and Salge, C. 2016. Does Empowerment Maximisation Allow for Enactive Artificial Agents? In *Proc. 15th Conf. ALIFE (to appear)*.
- Hoffman, G., and Weinberg, G. 2010. Gesture-based human-robot Jazz improvisation. In *2010 IEEE Int. Conf. Robotics and Automation*, 582–587.
- Jacob, M., and Magerko, B. 2015. Interaction-based Authoring for Scalable Co-creative Agents. In *Proc. 6th Int. Conf. Computational Creativity*, 236–243.
- Jordanous, A. 2015. Four PPP Perspectives on Computational Creativity. In *AISB 2015 Symposium on Computational Creativity*, 16–22.
- Kieran, M. 2014. Creativity as a Virtue of Character. In Paul, E. S., and Kaufmann, S. B., eds., *The Philosophy of Creativity: New Essays*. Oxford Scholarship. 125–144.
- Klyubin, A. S.; Polani, D.; and Nehaniv, C. L. 2008. Keep Your Options Open: An Information-Based Driving Principle for Sensorimotor Systems. *PLoS one* 3(12):1–14.
- Laclaustra, I. M.; Ledesma, J. L.; Méndez, G.; and Gervás, P. 2014. Kill the Dragon and Rescue the Princess: Designing a Plan-based Multi-agent Story Generator. In *Proc. 5th Int. Conf. Computational Creativity*.
- Magerko, B.; Permar, J.; Jacob, M.; Comerford, M.; and Smith, J. 2014. An Overview of Computational Co-creative Pretend Play with a Human. In *Proc. 1st Workshop Playful Virtual Characters, 14th Conf. Intelligent Virtual Agents*.
- Maher, M. L.; Brady, K.; and Fisher, D. H. 2013. Computational Models of Surprise in Evaluating Creative Design. In *Proc. 4th Int. Conf. Computational Creativity*, 147–151.
- Merrick, K. E., and Maher, M. L. 2009. *Motivated Reinforcement Learning. Curious Characters for Multiuser Games*. Springer.
- Mohamed, S., and Rezende, D. 2015. Stochastic Variational Information Maximisation. In *Proc. 29th Conf. NIPS*, 1–9.
- O'Regan, K., and Noe, A. 2001. A Sensorimotor Account of Vision and Visual Consciousness. *Behavioral and Brain Sciences* 24:939–1031.
- Oudeyer, P.-Y., and Kaplan, F. 2008. How Can We Define Intrinsic Motivation? In *Proc. 8th Int. Conf. Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, 93–101.
- Pearl, J. 2000. *Causality: Models, Reasoning and Inference*. Cambridge University Press.
- Pfeifer, R.; Iida, F.; and Bongard, J. 2005. New Robotics: Design Principles for Intelligent Systems. *Artificial Life* 11(1-2):99–120.
- Pickering, J. 2005. Embodiment, Constraint and the Creative Use of Technology. In *Proc. Freedom and Constraint in the Creative Process in Digital Fine Art*.
- Rosch, E.; Thompson, E.; and Varela, F. J. 1992. *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.
- Ryan, R., and Deci, E. 2000. Intrinsic and Extrinsic Motivations: Classic Definitions and New Directions. *Contemporary Educational Psychology* 25(1):54–67.
- Salge, C.; Glackin, C.; and Polani, D. 2014. Empowerment – an Introduction. *Guided Self-Organization: Inception* 67–114.
- Saunders, R., and Bown, O. 2015. Computational Social Creativity. *Artificial Life* 21(3):366–378.
- Saunders, R., and Gemeinboeck, P. 2014. Accomplice: Creative Robotics and Embodied Computational Creativity. In *Proc. 50th Anniversary Convention of the AISB Symposium*.
- Saunders, R., and Grace, K. 2008. Towards a Computational Model of Creative Cultures. In *AAAI Spring Symposium: Creative Intelligent Systems*.
- Saunders, R.; Gemeinboeck, P.; Lombard, A.; Bourke, D.; and Kocaballi, B. 2010. Curious Whispers: An Embodied Artificial Creative System. In *Proc. 1st Int. Conf. Computational Creativity*, 100–109.
- Saunders, R. 2007. Towards a Computational Model of Creative Societies Using Curious Design Agents. *Engineering Societies in the Agents World VII*:340–353.
- Shannon, C. E. 1948. A Mathematical Theory of Communication. *Bell Systems Technical Journal* 27:379–423.
- Steels, L. 1995. A Self-Organizing Spatial Vocabulary. *Artificial Life* 332(1995):319–332.
- Stokes, P. D. 2005. *Creativity From Constraints: The Psychology of Breakthrough*. Springer.
- Theune, M.; Faas, S.; Faas, E.; Nijholt, A.; and Heylen, D. 2003. The Virtual Storyteller: Story Creation by Intelligent Agents. In *Proc. Conf. Techn. for Interactive Digital Storytelling and Entertainment*, 204–215.
- Touchette, H., and Lloyd, S. 2004. Information-Theoretic Approach to the Study of Control Systems. *Physica A: Statistical Mechanics and its Applications* 331:140–172.
- Von Uexküll, J. 1982. The Theory of Meaning. *Semiotica* 42(1):25–78.
- Weinberg, G.; Driscoll, S.; and Thatcher, T. 2006. Jamaa A Middle Eastern Percussion Ensemble for Human and Robotic Players. In *Proc. Int. Conf. Computer Music*, 464–467.
- White, R. W. 1959. Motivation Reconsidered. *Psychological review* 66(5):297–333.