

Towards Goal-aware Collaboration in Artistic Agent Societies

Otto Hantula and Simo Linkola

Department of Computer Science

University of Helsinki

{otto.hantula, simo.linkola}@helsinki.fi

Abstract

We study the effects of goal-awareness in artistic agent societies creating evolutionary art. Particularly, we examine how goal-awareness may be utilized in modeling an agent’s peers when the aesthetic goals of the agent and its peers are subject to change. The agents use the learned peer models to choose their collaboration partners, and may alter their own aesthetic goal for the duration of the collaboration in order to enhance the potential of the collaboration outcomes. In addition, we demonstrate how goal-awareness can be used to guide the aesthetic goal change. The empirical evaluation indicates that agents which can adapt to their collaboration partners are more likely to reach favorable collaboration outcomes, even when their partners perceive fundamentally different properties from the artifacts.

Introduction

An agent seeking to select suitable collaboration partners in a creative society where the agent’s and its peers’ *aesthetic goals* are subject to change raises the need for dynamic peer models. We study how *goal-awareness* (Linkola et al. 2017), the ability to monitor and control one’s own goals, can be utilized in peer modeling and collaboration partner selection, and to facilitate favorable collaboration outcomes. Further, we demonstrate how an agent can use goal-awareness in conjunction with novelty-seeking (curious) behavior to strategically change its own aesthetic goals. In empirical evaluation, we observe that goal-aware agents are more likely to reach favorable collaboration outcomes and strategic aesthetic goal change causes emergent phenomena encompassing the whole society.

We build upon our earlier work (Linkola and Hantula 2018), where we investigated how artistic agents creating evolutionary art (Sims 1991; Romero and Machado 2007) could find feasible collaboration partners in a society consisting of agents with different skills and aesthetic preferences. Each agent aims to produce both valuable and novel outputs (Boden 1992) using aesthetic measure for value and a memory of previously seen artifacts for novelty. The agents interact with each other through artifact exchanges and pairwise collaboration, in which the agents aim to jointly create an artifact. To distinguish favorable collaboration partners, each agent learns a private model of its peers,

which it then utilizes in collaboration partner selection.

In our previous work the agents had static aesthetic preferences which they couldn’t change. In this paper, we expand the problem setting by allowing each agent’s aesthetic goals to change, resulting in a dynamic and more complex situation. To handle the increased complexity, we provision the agents with different traits of goal-awareness.

Drawing from self-adaptive (see, e.g. Salehie and Tahvil-dari (2009)) and self-aware systems (Lewis et al. 2015), our main focus is in metacreativity (Linkola et al. 2017). First, we are interested in how goal-awareness may benefit an agent in selecting its collaboration partners in dynamic situations where its own and its peers’ preferences are subject to change. Second, the goal-aware agents have an ability to adapt their aesthetic goals to a given collaboration partner for the duration of the collaboration. Third, we demonstrate how an agent can use its memory and goal-aware peer models in order to make strategic changes to its aesthetic goal satisfying both its collaboration- and novelty-seeking goals.

The rest of the paper is organized as follows. In the next section, we give motivation and background for our work. Then, we describe our agent society in general and different components of individual agents. We then move to our contributions. First, we go through on a conceptual level the distinct ways we utilize goal-awareness in this paper. Next, we define a new goal-aware peer modeling scheme which exploits the linear nature of our selected aesthetic functions, and describe how an agent can make the strategic aesthetic goal changes using the goal-aware peer model. Then, we outline our empirical experiment setup and present the main results from the experiments. We finish with discussion and conclusions.

Background

Our paper studies social behavior of artistic agents, and we are interested in emergent phenomena during the agent society’s lifetime, in the context of computational social creativity (see, e.g. Saunders and Bown (2015)). A prominent conceptualization of social creativity is the system’s view of creativity (Csikszentmihalyi 1988). It describes how the accumulated cultural artifacts, i.e. *the domain*, the experts of a given *field* and (each) *individual* are in constant interaction and affect each other. The major claim of the system’s view is that creativity is not in any single component but in

the interaction between all three components. In this paper, we focus on the individuals' ability to find suitable collaboration partners within the changing field, and how strategic behavior of individuals may cause emergent macro-level phenomena in the society. The domain is modeled implicitly as the collection of artifacts in the agents' memories.

Agent-based simulations have been extensively utilized to study social aspects of creative phenomena. For example, Saunders and Gero (2001) report emergence of communication cliques between agents with matching hedonic functions in a society of curious agents producing evolutionary art; Sosa and Gero (2005) reveal emerging social roles, such as gatekeepers, and other social phenomena when simulating designers and their societies; and Gabora and Tseng (2014) show that self-regulation of new ideas may have a positive effect on the mean fitness of ideas present in an agent society.

However, in agent-based simulations, interaction between the agents is typically defined using simple rules. The agents may be (directly) affected by the choices and actions of their neighbors or the society as a whole, but they do not model distinct peers in order to make strategic decisions about their own behavior involving those peers.

On the other hand, the skills, preferences and other properties of the collaboration participants have a direct result on the collaboration outcomes (Uzzi and Spiro 2005). To be able to distinguish favorable collaboration partners or otherwise act with social intent, an agent has to have a model of its peers' "minds" (Castelfranchi 1998).

Collaboration is essential for (computational) creativity allowing the participants to produce artifacts they could not by themselves (Paulus and Nijstad 2003; Uzzi and Spiro 2005; Pérez y Pérez et al. 2010). In computational creativity, collaboration of independent creative agents has gathered the most attention in musical domain. However, even in the musical domain, the set of collaborating agents is typically fixed, e.g. to ensembles where each agent plays a different instrument (Eigenfeldt et al. 2017).

Overall, there is a prominent lack of research considering how independent creative agents should model their peers' collaboration potential and utilize the peer models in their decision making, e.g. when selecting collaboration partners.

Peer modeling becomes a dynamic problem if an agent or its peers are subject to change as time elapses. However, this is the standard situation in creative societies: agents evolve in their style, aesthetic preferences and other properties.

In computational creativity, being aware of one's own creative process and being able to adjust it is often called metacreativity (Linkola et al. 2017). A particularly eminent aspect of metacreativity is goal-awareness. In conjunction with interaction-awareness, a goal-aware agent is provided with tools to adapt to its collaboration partners and change how it perceives its peers in a significant manner. Particularly, an agent may envision how it would observe artifacts if its goals would be different. The agent may then utilize this knowledge and temporarily adapt its goals to a new collaboration partner.

In this work, we hope to take the first steps to address the different concerns mentioned above. Building upon our pre-

vious work with interaction-aware agents, we study social behavior of creative agents which interact with their peers intentionally. We aim to add to the understanding of how goal-awareness may aid the agents in their peer modeling and collaboration partner selection, during the collaboration process and in strategically changing their aesthetic goals with respect to how they see their peers.

Agent Society

The agent society consists of a diverse set of artistic agents creating images that are novel and valuable to them. The agents differ in their image creation skills, aesthetic goals and what they are able to perceive in an image. In particular, we are focusing on the effects of changing aesthetic goals through goal-awareness.

The agent society is simulated iteratively. At odd time steps each agent creates a new image individually (we call these *solitary artifacts*). At even time steps the agents pair up and collaborate with their partner, aiming to produce a jointly created artifact. An individual agent can create and evaluate artifacts, as well as interact with its peers by sending them artifacts and through *collaboration*. To guide its interaction with other agents, and possibly other behavior, an agent learns a *peer model*.

An agent creates evolutionary art (Sims 1991) using an *evolutionary engine*. The artifact *evaluation* utilized in the evolutionary engine is based on perceived *value* and *novelty*. An agent has one *aesthetic measure* it uses to compute value and a limited *memory* of seen artifacts it uses to compute novelty. As its aesthetic goal, an agent has a target value for the aesthetic measure. The ability to adjust this target value is the key feature introduced in this paper. When an agent changes its current target value, we call it *movement*.

Next, we move on to describe these abilities and components on a general level. For the full details of the evolutionary engine's configuration and the collaboration process, we refer the reader to Linkola and Hantula (2018).

Evolutionary engine An agent creates a new image using an evolutionary engine, initializing the engine's population partly using the images it has previously made during the simulation. The evolutionary engine uses genetic programming to evolve an expression tree, which is used to calculate the value for each (x, y) coordinate in an image. The tree consists of terminals (leafs) and functions (inner nodes). An agent's image creation skills are determined by the subset of functions it has for creating expression trees.

Aesthetic measure and value For the purposes of this paper we use two aesthetic measures present in our earlier work: entropy and fractal dimension (Linkola and Hantula 2018). Entropy is defined by the color distribution in an image and fractal dimension measures an image's structural properties. Each agent has only one of these two measures, but the actual target within the aesthetic measure's bounds is different for each agent. For the complete descriptions of how the objective values of the aesthetic measures are computed, we guide the reader to den Heijer and Eiben (2014).

The value of an artifact I is calculated based on the evaluating agent's aesthetic measure v and target value, i.e. *aes-*



Figure 1: Example of collaboration between two agents. On the left is a solitary image by a fractal dimension agent, on the right by an entropy agent. On the center is an artifact the agents made in collaboration, showing traits from both agents.

thetic goal g . The closer the objective aesthetic measure calculated from the artifact is to the goal, the more valuable the artifact is. The value is a linear mapping of the distance from the aesthetic goal, calculated with the following formula:

$$\text{value}(I) = \begin{cases} 1 - \frac{|g-v|}{v_{max}-v_{min}}, & \text{if } |g-v| < v_{max} - v_{min} \\ 0, & \text{otherwise,} \end{cases}$$

where v_{min} and v_{max} are the minimum and the maximum values for the aesthetic measurement, respectively.

We chose entropy and fractal dimension as aesthetic measures because of their potential to complement each other. Further, their asymmetrical relationship provides an interesting case for analysis: agents with entropy tend to create images of high complexity regardless of their exact target value, but complexity's target value does not have a strong relation to entropy in the images produced by our agents.

Memory and novelty An agent has memory for up to 500 artifacts, where the agent can store artifacts it has seen. The artifacts can be created by itself or other agents. If the memory is full when storing a new artifact, the oldest artifact in the memory is forgotten.

The novelty of an artifact is evaluated with the function $\text{novelty}(I) = \min_m \text{ed}(I, m)$, where I is the artifact being evaluated, m is an artifact in the agent's memory, and $\text{ed}(\cdot)$ is the normalized Euclidean distance between the artifacts. In other words novelty is the euclidian distance to the closest artifact in the agent's memory.

Evaluation Using the value and novelty calculated from an artifact, an agent uses the following function to get the final evaluation: $\text{eval}(I) = \frac{1}{2}\text{value}(I) + \frac{1}{2}\text{novelty}(I)$.

Movement An agent can change its aesthetic target value, or aesthetic goal, which is used for creating artifacts and selecting collaboration partners. The movement changes what kind of artifacts an agent creates (what they see valuable) and with whom it collaborates. We run tests with two different types of movement. First is completely random, where the new goal is drawn from a uniform distribution. Second utilizes goal-awareness and curiosity in determining the new goal. These are explained in detail later.

Peer model An agent learns a peer model of the other agents from the artifacts they create. The peer model is used to select collaboration partners, to change one's aesthetic goal for collaboration and for goal-aware movement. We use two Q-learning based learning schemes for the peer models, which are described in their own section.

Collaboration In collaboration, a pair of agents merge their artifact creation skills and aesthetic goals aiming to produce an artifact jointly. The collaboration follows an alternating co-creation process (Kantosalo and Toivonen 2016), where the collaboration partners evolve the same artifact set in turns iteratively (see Linkola and Hantula (2018) for details of the collaboration process). Figure 1 shows an example of collaboration between two agents.

If the agents agree on an artifact to be produced as a collaboration's result, we call the collaboration *successful*. If they can't agree on an artifact, no artifact is produced. To negotiate about the collaboration artifact, both agents keep a hall-of-fame of the best artifacts seen during the collaboration process (sorted to an increasing rank, the best artifact having the first rank). At the end of the collaboration, agents compare their hall-of-fames and pick an artifact which has the smallest combined rank as the collaboration result, i.e. they agree on it, if there exists an artifact which is in both hall-of-fames.

Manifestations of Goal-awareness

A goal-aware agent is able to observe how it reaches its current goals and adjust these goals if it sees fit (Linkola et al. 2017). In essence, goal-awareness facilitates creative autonomy (Jennings 2010) in an agent, aiding the agent to change its creative process and produce potentially previously unreachable artifacts.

There are three ways in which our agents can utilize goal-awareness. First, if the agent is aware of the goals of its peers, it can model those goals and their changes and use that information to select feasible collaboration partners with respect to its own current goals. Second, if the agent models its peers' goals, it may adjust its own aesthetic goal for the

duration of the collaboration, possibly enhancing the collaboration potential. Third, the agent can use the learned peer models to make strategic changes to its aesthetic goal.

Next, we describe on a conceptual level the different ways in which goal-awareness is implemented in this work.

Peer modeling Changing aesthetic preferences of the agent and its peers imposes new challenges on peer modeling. The peer model has to contain sufficiently accurate and topical information about peers for it to have any value as an asset in an agent’s decision making.

Goal-awareness provides capabilities to handle changing aesthetic preferences. If an agent is able to imagine how it would perceive a certain artifact if its aesthetic goal would be different, it can keep an alternative peer model for each of its goals. The agent can adjust each of these peer models when it perceives an artifact from another agent. Then, when the agent changes its own aesthetic goal, it can assimilate the alternative peer model most suitable for the current goal without the need to build the peer model from scratch.

Exploiting the alternative peer models is central for the peer modeling scheme proposed in this paper, ga-Q, which is described in detail in the next section.

Adaptation during collaboration Our agents utilize goal-awareness in collaboration by changing their own aesthetic goal to align with the partner’s goal. When the collaboration begins, the agent in the collaboration pair that got to select its partner chooses a *temporary goal*, which it uses during the collaboration. The selected partner doesn’t change its goal. Selecting the partner perceived as best and then selecting the temporary goal to suit the partner can be seen as a combination of the selfish and altruistic approaches (Linkola and Hantula 2018). First the agent selfishly selects the partner it personally likes most. Then it altruistically adapts its own goal to be the best possible collaboration partner for the other agent.

Strategic movement Our agents use curiosity to guide when and where to move. There are three factors that affect when an agent decides to move: how long the agent has had its current aesthetic goal, how good artifacts it is producing with respect to its current goal and what the other agents are creating. An agent doesn’t want to stay in the same place for too long. It wants to produce valuable artifacts, moving if it fails to do so. A place currently being explored by the society is also seen as less interesting.

When the agents move, they use their memory to guide their movement to less explored areas, exhibiting curiosity. The agents utilize goal-awareness by considering their peers’ aesthetic goals, trying to move to areas currently unoccupied by other agents, but still having sufficient collaboration potential. The strategic movement is described in detail in its own section.

Peer Modeling

Peer modeling is the basis for intentional interaction between the agents in our experiments. The learned peer model

is used to select good collaboration partners, change the agent’s aesthetic goal for collaboration and guide movement. Because the agents have dynamic aesthetic goals, the model has to be able to quickly adapt to the changes in the learning agent itself and the peers. To enable some of the aspects of goal-awareness, we describe a peer model that models all of the agent’s possible aesthetic goals simultaneously, even though the agent can only have one of them at a time.

The learning scheme for peer modeling we propose here is an extension to the learning scheme called *hedonic-Q* in Linkola and Hantula (2018). Hedonic-Q is based on Q-learning (Watkins and Dayan 1992), which is a common reinforcement learning method, that maintains Q-values for state-action pairs based on received reward. The Q-value for a state-action pair is the expected utility of choosing the action while in the state. Hedonic-Q uses a simplified, stateless version of Q-learning, with update rule $Q(a_i) \leftarrow Q(a_i) + \lambda(r - Q(a_i))$ (Claus and Boutilier 1998), where a_i is the action of selecting peer i as collaboration partner, r is the reward and λ is the learning rate (we use $\lambda = 0.9$). It would be natural to use the evaluation of the collaboration artifact created with peer i as the reward. Instead, an agent uses its evaluations of i ’s solitary artifacts as an approximation, learning how much it likes its peers artifacts. In our experience this works well, because the agent gets information about the peers from all created solitary artifacts and not just from its collaborations (Linkola and Hantula 2018).

The new learning scheme used in this paper, *ga-Q* (goal-aware-Q), extends hedonic-Q with goals. The update rule for ga-Q is $Q(g, a_i) \leftarrow Q(g, a_i) + \lambda(r - Q(g, a_i))$, where g is an aesthetic goal. For the reward ga-Q uses its own evaluation of its peers’ artifacts just like hedonic-Q. Ga-Q learns how much it likes its peers’ artifacts relative to g .

Ga-Q requires discrete goals, but the agent’s aesthetic goal is a continuous value. We discretize the continuous value by dividing the range of possible goal values to B equal sized bins, resulting in B goals for ga-Q. So if the aesthetic goal is bounded in $[v_{min}, v_{max})$, this interval is divided into B , $(v_{max} - v_{min})/B$ sized subintervals, which represent the goals for ga-Q. From now on we refer to the discretized goals with g_b . We use $B = 20$.

One of the greatest benefits of ga-Q is, that the learning agent can update all possible goals simultaneously based on a single artifact. For each goal g_b in the ga-Q model, the artifact is evaluated using the middle point of the bin as the aesthetic goal. Then $Q(g_b, a_i)$ is updated using this evaluation as the reward. This way the agent already knows how to act, when it changes its aesthetic goal to a new one, even if it has never had that goal before.

An agent uses the peer model learned using hedonic-Q by sorting its peers into a preference order using their corresponding $Q(a_i)$ values. The ordering is done similarly with ga-Q, by first mapping the agent’s current aesthetic goal to g_b and then using the $Q(g_b, a_i)$ values.

Adaptation during collaboration During collaboration, the agent that got to choose its partner in the collaboration pair uses the Q-values to choose a new aesthetic goal in the following way. If the selected partner is peer i , the temporary goal is the middle point of the bin that corresponds to

the goal $\max_{g_b} Q(g_b, a_i)$. This means choosing the goal that maximizes the agent’s appreciation of its collaboration partner’s artifacts, maximizing the chance that the collaborating agents have some common ground, i.e. the agents appreciate similar artifacts. After the collaboration, the temporary goal is changed back to the goal the agent had before collaboration. The selected partner does not change its goal.

Strategic Movement

In this section, we describe how an agent strategically changes its aesthetic goal, i.e. moves. First, we describe how the movement is triggered, and then we describe how the new aesthetic goal is decided.

Choosing to move We model an agent’s desire to move as a value c , accumulating whenever the agent observes an artifact. When c exceeds a fixed threshold c_t , an agent chooses a new aesthetic goal, i.e. moves. When an agent changes its goal, c is reset to 0.

An agent A accumulates c when observing artifact I as follows:

$$c = \begin{cases} c + \frac{1}{\text{value}(I)^n}, & \text{if } A \text{ is a (co-)creator of } I \\ c + \max\left\{0, 1 - \frac{|g-v|}{2s}\right\}, & \text{for other images,} \end{cases}$$

where n is the number of simulation steps since the agent’s last goal change, g is the agent’s current aesthetic goal, v is the aesthetic measure value of the received artifact and s is the bin size used for ga-Q.

The formulation for (co-)created artifacts accumulates c exponentially faster the longer the agent fails to create value with its current goal. On the other hand, the accumulation is less pronounced when the agent has just moved, giving the agent time to adjust itself to its new aesthetic goal. The accumulation of c for other artifacts is larger when its peers are creating artifacts close to the agent’s aesthetic goal, making the agent move in shorter intervals if its close to its peers.

The threshold c_t is designed so that an agent accumulates enough curiosity to trigger movement with every 10th time step under two assumptions. First, an agent is able to completely satisfy its own aesthetic goals (produces solitary and collaborated artifacts with value 1.0). Second, each observed peer artifact’s objective aesthetic value is drawn from uniform distribution within the aesthetic bounds.

Moving Once the movement is triggered, an agent decides its new aesthetic goal based on its memory, the Q-values and potentially the agent’s current aesthetic goal.

We describe two different ways an agent can strategically change its aesthetic goal: static and dynamic, both utilizing ga-Q. Static movement is as likely to move to any place within its aesthetic bounds. Dynamic movement prefers aesthetic goals closer to its current goal, decreasing the desire to move linearly with distance to a new aesthetic goal.

The new aesthetic goal is chosen as follows:

1. Agent calculates how many artifacts in its memory fall into each ga-Q goal bin, and filters out the four (20%) most crowded bins.
2. Agent filters out any remaining bins which are perceived to contain a peer, i.e. have an agent which has maximum

Q-value in that bin, and estimates the collaboration potential of each remaining bin as the sum of the four highest Q-values in it.

3. If the movement type is dynamic, the agent scales each remaining bin’s value according to its closeness to its current aesthetic goal.
4. Agent selects the bin with the highest value and chooses a new aesthetic target using a uniform distribution defined by the bin’s borders.

The decision process is designed to satisfy an agent’s novelty-seeking and collaboration goals. By filtering out the areas where most artifacts have been observed and the areas which are expected to contain peers, an agent tries to find a novel place within the society. The collaboration goal is satisfied by choosing a place which has a good collaboration potential.

Experiment Setup

With our experiments we aim to investigate intentional collaboration partner selection in a dynamic agent society, using goal-awareness to benefit the partner selection and the collaboration process, and finally guiding the aesthetic goal change with curiosity and goal-awareness. Our main research questions are:

1. How does the new learning scheme ga-Q perform? Especially compared to hedonic-Q.
2. How does adapting to the collaboration partner with a temporary aesthetic goal affect the results?
3. Does using curiosity and goal-awareness in changing one’s aesthetic goal benefit collaboration?
4. What kind of emergent behavior arises on the society’s level, when curiosity and goal-awareness are used in changing one’s aesthetic goal?

In our experiments, we have 16 agents and 2 aesthetic measures. Half of the agents have entropy (ENT) as their aesthetic measure and the other half have fractal dimension (FRD). Each agent is initialized with a goal for the aesthetic measure. The aesthetic goal g for an agent A is initialized uniformly from the following bounds: $g \in [0.5, 4.5]$ if A uses entropy and $g \in [0.5, 1.8]$ if A uses fractal dimension.

The nucleus of our experiment setup is a simulation run consisting of 200 iterative time steps ($S = (s_1, s_2, \dots, s_{200})$). The agents start each simulation with empty memories, creating the first images using only their aesthetics and evolutionary engine. The simulation is run 30 times for each learning scheme and movement configuration present in our experiments, resulting in a total of 180 runs. For the results we report the experiment setup run averages.

At odd time steps each agent creates a solitary artifact. At even time steps the agents select their collaboration partners and create artifacts in pairs. All of the created artifacts are sent to all agents for evaluation. An agent memorizes all of the artifacts it has created. An artifact created by another agent is memorized, if it exceeds the agent’s thresholds for novelty (0.4) and value (0.5). If an agent changes its aesthetic goal, it does so at the start of an odd time step, before it starts creating a solitary artifact.

Table 1: Collaboration success ratios and various value measures for learning schemes. The statistics are averages of 30 simulation runs for each experiment configuration (column), displayed with 99% confidence interval.

Measurement	Collaborator selection and adaptation				Strategic movement	
	Hedonic-Q	Ga-Q _{fxd}	Ga-Q _{ada}	Random	Ga-Q _{st}	Ga-Q _{dyn}
Collab. success (CS) %	83.4 ± 0.8	83.5 ± 0.8	93.4 ± 0.5	67.4 ± 1.3	95.4 ± 0.6	94.7 ± 0.6
Value, own solitary	.962 ± .002	.962 ± .002	.962 ± .002	.960 ± .001	.981 ± .001	.969 ± .005
Value, collab. selector	.934 ± .002	.936 ± .002	.939 ± .001*	.892 ± .003	.968 ± .001*	.942 ± .004*
Value, collab. partner	.941 ± .002	.941 ± .002	.946 ± .002	.920 ± .002	.974 ± .001	.958 ± .007
Novelty, own solitary	.536 ± .004	.531 ± .007	.532 ± .004	.535 ± .006	.595 ± .006	.539 ± .009
Novelty, collab. selector	.535 ± .004	.528 ± .007	.534 ± .005	.523 ± .006	.594 ± .006	.541 ± .009
Novelty, collab. partner	.534 ± .004	.527 ± .007	.533 ± .005	.521 ± .006	.594 ± .006	.540 ± .009
-Collaboration between different aesthetics (selector → selected), collaboration attempt counts in brackets						
ENT → FRD CS%	83.5 ± 1.6 (4252)	83.9 ± 1.4 (4243)	86.8 ± 1.5 (4116)	70.6 ± 2.0 (6471)	91.8 ± 1.3 (4253)	89.0 ± 1.6 (3135)
FRD → ENT CS%	80.7 ± 1.6 (2680)	79.0 ± 2.2 (2703)	84.1 ± 2.5 (2632)	71.5 ± 2.8 (6381)	86.0 ± 2.3 (2523)	86.0 ± 3.0 (2599)
Total attempts	6932	6946	6748	12852	6776	5734

* To make the values comparative between the experiments, the value is computed using the agent’s real aesthetic goal, not the adapted goal for the duration of the collaboration.

At the start of the collaboration time steps, the agents are arranged into a random order. Then one by one the agents use their preference list (defined by the learning scheme) to select a partner. The partner is the first agent in their preference list, which does not yet have a collaboration partner on this time step.

Creating a solitary and collaboration artifact takes roughly the same amount of resources. In our experiments, 10 evolutionary iterations are used for both solitary and collaboration artifacts. In the collaboration process, the agents do 5 iterations each.

Next, we describe our two different experiment setups.

Collaborator selection and adaptation With this setup we aim to investigate using goal-awareness for selecting collaboration partners in a dynamic situation and for adapting to the collaboration partner. These experiments also serve as a baseline for the strategic movement setup.

We run the setup for hedonic-Q and ga-Q (with and without adaptation), and compare the results to a baseline where the collaboration partners are selected randomly. We refer to these runs as hedonic-Q, ga-Q_{fxd} (without adaptation), ga-Q_{ada} (with adaptation) and random.

In this setup all learning schemes change their aesthetic goal randomly. The agents have 0.2 probability to change their goal in the beginning of each solitary time step, making them change their goal on average on every 10th time step. The new goal is drawn from uniform distribution in the bounds of the agent’s aesthetic measurement (see above).

Strategic movement With this experiment we aim to see how using curiosity and goal-awareness in changing one’s aesthetic goal affects the collaboration results and does strategic movement give rise to emergent phenomena on the macro-level.

We experiment with the two different strategic movements using ga-Q: static (ga-Q_{st}) and dynamic (ga-Q_{dyn}).

Results

We now proceed to present the results from our experiments.

Collaborator selection and adaptation We see from Table 1 that all peer modeling schemes are able to produce more collaboration artifacts (higher CS%) with higher value than random collaboration. Especially ga-Q_{ada} is able to collaborate successfully, which means that adaptation to the partner’s aesthetic goal is beneficial for our collaboration process. Surprisingly, there isn’t much difference in CS% between hedonic-Q and ga-Q_{fxd}, although hedonic-Q has to relearn all Q-values every time the agent changes its aesthetic goal. Ga-Q_{fxd} should cope with the goal changes better, because it maintains Q-values for all goals simultaneously. We return to this in discussion.

Collaboration value Interestingly, adapting to the collaboration partner does not decrease the adapting agent’s value for collaboration artifacts, as seen in row 3 of Table 1, comparing ga-Q_{fxd} to ga-Q_{ada} (the value is calculated using the agent’s real aesthetic goal, instead of the temporary goal). This is probably caused by the selfish selection of partner by the selector agent, which is the one adapting, choosing peers who are close to its goal.

Collaboration between aesthetics From the last three rows of Table 1, it can be seen that hedonic-Q, ga-Q_{fxd} and ga-Q_{ada} all have significantly less collaboration between the aesthetics compared to random. Overall the entropy agents are able to select partners leading to more successful collaborations from the fractal dimension agents than vice versa. On rows 8 and 9 ga-Q_{ada} has statistically significantly higher CS% than ga-Q_{fxd} (Welch’s t-test, p-values 4.5e-05 and 1.3e-04 respectively). This shows that adapting to the collaboration partner is beneficial for collaboration between the aesthetics.

Strategic movement In Table 1 the two rightmost columns show a general improvement in collaboration suc-

Table 2: Average moving distance, clustering and rate.

Measurement	Aesthetic	Random	Ga-Q _{st}	Ga-Q _{dyn}
Average area covered in 10 steps (normalized)	ENT	0.424	0.232	0.059
	FRD	0.416	0.198	0.048
Average number of agents in the same bin	ENT	1.217	1.925	1.104
	FRD	1.201	4.333	1.292
Average number of aesthetic goal changes	ENT	20.279	21.971	15.025
	FRD	19.533	27.313	17.217

cess and value for the strategic movement, compared to random movement.

Movement From Table 2 we observe that ga-Q_{st} agents change their aesthetic goal more and operate in a much larger aesthetic range within 10 steps than ga-Q_{dyn} agents. Still ga-Q_{st} has less overlap in the bins than ga-Q_{dyn}, indicating a less spread out society. This combined with the high values for ga-Q_{st} in Table 1, it seems that the ga-Q_{st} society is very opportunistic, always jumping to the most promising place together. Ga-Q_{dyn} is more conservative and spread out in its movement. We reflect on this more in discussion.

These differences between ga-Q_{st} and ga-Q_{dyn} can be seen in Figure 2, too. In the ga-Q_{st} runs the whole society moves tightly together, even when the target is oscillating intensely, as happens with entropy. With fractal dimension the society tends to stay in the high end of the aesthetic bounds. In the ga-Q_{dyn} runs the society also moves together, but in a more spread out manner. The collective targets of the society do not oscillate, but rather move steadily.

Novelty As seen in Table 1 rows 5-7, novelty is quite similar between all the schemes, except for Ga-Q_{st}. Ga-Q_{st} finds more novelty than the others due to its curiosity and ability to move in the whole aesthetic range. The high novelty is probably also partially caused by the FRD agents favoring complex artifacts, which tend to be more novel. However, we observed in our experiments, that entropy agents also produced notably more novelty with Ga-Q_{st} than the other schemes. Ga-Q_{dyn} is also guided by curiosity, but it mostly operates in a small range around its current target, making finding novelty more difficult.

Discussion and Conclusions

We have presented a new goal-aware peer model, ga-Q. The peer model enables an agent to envision alternative aesthetic goals, allowing the agent to temporarily adapt to its collaboration partner, and position its own aesthetic goals in relation to its peers' aesthetic goals. Our experiments indicate that the goal-aware selection of temporary goal for the collaboration is beneficial to our collaboration process and that the curious and goal-aware movement is beneficial for both collaboration and solitary artifact creation. Goal-awareness can also facilitate collaboration between the aesthetics.

Ga-Q Overall, ga-Q shows potential as a straightforward way to provide agents with a goal-aware peer modeling technique. It is easily generalizable to societies where new peers are introduced and old ones may leave. When a new peer




Figure 2: Heat maps of the whole society's typical aesthetic goal movement during a single run for static (upper two) and dynamic (lower two) strategic movement.

enters, a new Q-value $Q(g, a_j)$ can be created for each goal with a default value. When a peer leaves, all Q-values related to it can simply be dropped. Similarly, new goals can be created and old goals can be dropped. However, the number of goals in ga-Q grows exponentially with the number of an agent's aesthetic goals, making it impractical in situations where the number of aesthetic goals an agent has is high.

The way in which we use the ga-Q's Q-values to calculate the collaboration potential for a goal is quite unconventional. In our case, using the sum of the top Q-values makes sense, because an agent might not be able to select its favorite peer in the partner selection process, as that peer might already be in a collaboration pair. Therefore the agents should aim to select goals for which good collaboration partners exist, even if they don't get to select the best one.

Adapting to collaboration partner In the results, we observed that ga-Q without adapting to the collaboration partner is close to hedonic-Q. The reason is that the change of one's aesthetic goal happens before an agent receives the solitary artifacts from its peers. Q-learning's learning speed is fast enough to adapt using one step's worth of information. If the agent wouldn't get new information between its own aesthetic goal change and partner selection, or if the learning rate was lower, hedonic-Q wouldn't be able to make informed choices, while ga-Q should be relatively unaffected.

Strategic movement Our results for strategic movement are dividing. Even though ga-Q_{st} has the highest collaboration success, value and novelty, it might not be the most desirable way of implementing a society. The rapid nature of static movement's collective aesthetic goal changes (see Figure 2) renders the whole society unstable.

Ga-Q_{st} is also heavily affected by the asymmetric nature of the two aesthetic measures. The ENT agents produce nearly only artifacts which the FRD agents observe to belong to a couple of bins near the higher end of their aesthetic bounds. This causes the FRD agents to swarm around these bins, unable to move away from them.

Further, the ability to change one's aesthetic goal arbitrarily far might not be preferable, e.g. agents drastically changing their aesthetic goal might not be able to make full use of their accumulated expertise. For a more spread out and conservative search of the domain, ga-Q_{dyn} seems preferable. However, the two different strategic movements can be seen as different points on the same scale: how much the agent prefers new aesthetic goals close to its current aesthetic goal.

Lastly, memorylessness of our strategic movement implementation makes it undesirable for long processes. The agent does not accumulate information of the aesthetic goals it has previously possessed, and thus the swarming behavior of the static FRD agents may emerge. For more sustained processes, time-awareness has to accompany strategic movement in order for the agent to understand the history of its own aesthetic goals and utilize that knowledge in its decision making.

To conclude, we believe that for true social intent, the agents need to model their peers and their interaction. By experiments we hope to have gathered some insight towards such intent. In the future, we aim to study more closely how time-awareness can be used in strategic movement in conjunction with goal-awareness and how societies with diverse strategic movement behaviors evolve over time.

Acknowledgments. This work has been supported by the Academy of Finland under grant 313973 (CACs).

References

- Boden, M. 1992. *The Creative Mind*. London: Abacus.
- Castelfranchi, C. 1998. Modelling social action for AI agents. *Artificial Intelligence* 103(1):157–182.
- Claus, C., and Boutilier, C. 1998. The dynamics of reinforcement learning in cooperative multiagent systems. In *Proceedings of the Fifteenth National/Tenth Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence*, 746–752.
- Csikszentmihalyi, M. 1988. Society, culture, and person: A systems view of creativity. In Sternberg, R. J., ed., *The Nature of Creativity: Contemporary Psychological Perspectives*. Cambridge University Press. 325–339.
- den Heijer, E., and Eiben, A. 2014. Investigating aesthetic measures for unsupervised evolutionary art. *Swarm and Evolutionary Computation* 16:52–68.
- Eigenfeldt, A.; Bown, O.; Brown, A. R.; and Gifford, T. 2017. Distributed musical decision-making in an ensemble of musebots: Dramatic changes and endings. In *Proceedings of the Eight International Conference on Computational Creativity*, 88–95.
- Gabora, L., and Tseng, S. 2014. The social impact of self-regulated creativity on the evolution of simple versus complex creative ideas. In *Proceedings of the Fifth International Conference on Computational Creativity*, 8–15. Ljubljana, Slovenia: Josef Stefan Institute, Ljubljana, Slovenia.
- Jennings, K. E. 2010. Developing creativity: Artificial barriers in artificial intelligence. *Minds and Machines* 20(4):489–501.
- Kantosalo, A., and Toivonen, H. 2016. Modes for creative human-computer collaboration: Alternating and task-divided co-creativity. In *Proceedings of the Seventh International Conference on Computational Creativity*, 77–84.
- Lewis, P. R.; Chandra, A.; Faniyi, F.; Glette, K.; Chen, T.; Bahsoon, R.; Torresen, J.; and Yao, X. 2015. Architectural aspects of self-aware and self-expressive computing systems: From psychology to engineering. *Computer* 48(8):62–70.
- Linkola, S., and Hantula, O. 2018. On collaborator selection in creative agent societies: An evolutionary art case study. In *Proceedings for the Seventh International Conference on Computational Intelligence in Music, Sound, Art and Design*, 206–222.
- Linkola, S.; Kantosalo, A.; Männistö, T.; and Toivonen, H. 2017. Aspects of self-awareness: An anatomy of metacreative systems. In *Proceedings of the Eight International Conference on Computational Creativity*, 189–196.
- Paulus, P., and Nijstad, B. 2003. *Group Creativity: Innovation through Collaboration*. Oxford University Press.
- Pérez y Pérez, R.; Negrete, S.; Nalosa, E. P.; Ávila, R.; Castellanos, V.; and Lemaitre, C. 2010. Mexica-impro: A computational model for narrative improvisation. In *Proceedings of the International Conference on Computational Creativity*, 90–99.
- Romero, J., and Machado, P., eds. 2007. *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*. Natural Computing Series. Springer Berlin Heidelberg.
- Salehie, M., and Tahvildari, L. 2009. Self-adaptive software: Landscape and research challenges. *ACM Transactions on Autonomous and Adaptive Systems* 4(2):14:1–14:42.
- Saunders, R., and Bown, O. 2015. Computational social creativity. *Artificial Life* 21(3):366–378.
- Saunders, R., and Gero, J. S. 2001. Artificial creativity: A synthetic approach to the study of creative behaviour. In *Computational and Cognitive Models of Creative Design V*, 113–139. Sydney, Australia: University of Sydney.
- Sims, K. 1991. Artificial evolution for computer graphics. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques*, 319–328.
- Sosa, R., and Gero, J. S. 2005. Social models of creativity. In *Proceedings of the International Conference of Computational and Cognitive Models of Creative Design VI*, 19–44. Heron Island, Australia: Key Centre of Design Computing and Cognition, University of Sydney, Australia.
- Uzzi, B., and Spiro, J. 2005. Collaboration and creativity: The small world problem. *American Journal of Sociology* 111(2):447–504.
- Watkins, C. J., and Dayan, P. 1992. Q-learning. *Machine learning* 8(3-4):279–292.