Artistically Skilled Embodied Agents

Patrick Tresset[1] and **Oliver Deussen** [2]

Abstract. In this paper we report on our research into the development of computational embodied systems dedicated to the production of artworks. We present a conceptual framework that introduces the notion of style-space in relation to the visual arts. This framework underpins and guides our general approach to the development of autonomous agents capable of producing objects that have artistic value. In addition we introduce the importance of intentionality and embodiment, two interrelated elements that are fundamental in the appreciation of visual artworks such as paintings and drawings. Our practical research has at this stage only touched upon the development of agents that are adapted and sufficiently skilled to display the lowest level of artistic autonomy. Having such agents is essential to conduct further research into the autonomous invention of artistically valuable styles. Two of these skilled agents and their production are presented in this document.

1 Introduction

The work presented in this paper is concerned with applied computational creativity in the context of the visual arts, with a focus on static pictorial expression, specifically observational drawing and painting.

One of the particular and essential characteristics of art is that artists develop–find–adopt–invent their own specific sets of rules, strategies and techniques. As a consequence, art is a vast conceptual space with a large number of valid paradigms, theories and expressions. In this context, our approaches and views do not pretend at universality, but describe possible and viable research and development strategies to progress towards achieving our aims. The term *artwork* encompasses a wide range of medium, such as videos, sculptures, installations, softwares, paintings, drawings, prints, in the context of this paper it exclusively means observational drawings and paintings. Although it constitutes an important and integral part of our art practice and research, in this paper we leave aside the role of the robotic systems as artworks in their own right.

Our aim is to develop autonomous systems that are capable of conceiving and producing artifacts that have a range of qualities and characteristics that enable their status as a work of art. Objects, to be considered as having such status, must be exhibited–evaluated–appreciated–acquired in a contemporary art context, and in the same manner as artist-made artworks.

The research described in this paper is artist-led and as such it is at present primarily focused on developing computational systems that have the skills required to produce objects that are considered as artworks. Although increasing the systems' level of autonomous artistic creativity will raise the artistic value of both the robots and the artifacts produced, this research can only be pursued using artistically-skilled agents. The development, exploration and exploitation of the style-spaces described in subsection 3.1 occur essentially through practice, and can only be envisaged after viable embodied skilled agents have been developed. This approach could be considered as an analogy to one of education paradigm in the fine arts where, until the end of the nineteenth century and for a large part of the twentieth century (until the 1960's) art education was skill-led. Art students would first acquire skills and only when the skill-set was mastered would they be expected to attempt to create their own work. Generally, it would take a few years[3] to reach this level. Progressing from the ground up and with a focus on practical results, is a necessary strategy as our research is more prominently tested, reviewed and disseminated through exhibitions than it is through the publication of research papers.

At this stage, our research has been focused on developing skilled embodied agents. Although the entities we have developed produce objects that are arguably considered as artworks, objectively they only implement very low-level creative behaviours. However, when using measures of creativity such as those described in [19], the systems can be evaluated as being more creative than we may have initially considered, we will clarify this in the conclusions when commenting on the systems' output.

The paper is structured in the following manner: in section 2 we give a brief overview of some existing computational systems that aim to produce artworks. In section 3 we present a formalisation of what we name *style-space* with subsections on the importance of intentionality and embodiment in art. We then describe two embodied systems, in section 4, we present Paul, a robot dedicated to drawing faces, in section 5, we describe the system developed using e-David, a painting robot conceived at the University of Konstanz, Germany.

2 In the computational arts context

For centuries we have been manufacturing images automatically, early examples include Vaucanson's automated loom, which is often put forth as a predecessor to the computer, and Maillardet's draftsman-writer automaton. In the twentieth century Tinguely comes to mind with his drawing machines. In this paper we are not looking at systems that are considered as artworks in themselves, such as softwares, robots, interactive installations, net-art, we are focusing on systems that generate two-dimensional static images.

We can trace the beginning of the computer art movement to the 1960's which saw artists writing software with the intention of producing works of art. The work of artists such as Manfred Mohr, Frieder Nake, Roman Verostko, Paul Brown and Ernest Edmond are notable examples of this early strand [27]. This early period coincided with and should be seen in the context of significant movements in the arts [10], where the removal of traces of the artist from

[1] Goldsmiths, University of London, United Kingdom, email: p.tresset@gold.ac.uk
[2] University of Konstanz, Germany, email: oliver.deussen@uni-konstanz.de

[3] In the Beaux-Arts schools 3 to 4 years

Figure 1. Paintings from the series Paul's Memories, 2013

the artwork was an important feature[4]. Later, in the 1980's, the advent of evolutionary art with the work of Karl Sims [25] and William Latham [28], and later works linked to artificial life such as those of Penousal Machado [15], Leonel Moura [18] or Rui Antunes [1].

Another branch of computational art is driven by an investigation into how an artist produces artworks, with two main strands: *a)* work driven by individuals who are originally art practitioners: the pioneering work of Harold Cohen with AAron [17], in the late 1980's Ed Burton's work with ROSE [5], more recently with the work of Steve DiPaola [8] and also our work [31]; *b)* work driven by researchers who are originally computational scientists: such as Simon Colton's work on the Painting Fool [7], Oliver Deussen and Thomas Lindemeier's work on e-David [14].

The computer art movement is arguably a significant art strand that has since the beginning attracted collectors, held important founding exhibitions, and is now the subject of art history publications. We see dedicated collections in major museums such as the Victoria and Albert museum and yet still computer art is apart from academic contemporary art history. This might naturally evolve as the digital natives[5] take progressively more importance in the art-world.

3 Art, style, intentionality and embodiment

3.1 Art and style-spaces

Frameworks such as the one described in [32] which formalise and detail Boden's views on creativity [2], offer a perspective where creativity is considered related to a conceptual space, the set of all concepts, a notion that covers everything that is the result of creativity, concrete or abstract. To formalise our perspective on visual art we consider what we name *style-spaces* as subspaces of the *art-space*.

Works of art display and have embedded in them an ensemble of characteristics. The combination, the expression of these characteristics is what we call style. Great and minor artists produce artworks that have an individual identifiable style. Styles are not isolated or autonomous, they contain traces of influences from past and contemporary art history. These influences are not expressed directly as a kind of patchwork or collage, but rather they are combined and in a certain manner "digested" by the artist. Artists make these influences their own, transformed by their "personality" or what we could consider the artists charateristics, such as their strengths and weaknesses (psychological, motor, perceptual, cognitive). This process enables the artwork produced to display rich coherent individual style, in which art history and contextual influences are embedded. This facilitates the perception and appreciation of artworks as works of art.

We can describe the artist's *major creative event* as the establishment of their novel–original–personal style-space. Typically, this type of creative event occurs a limited amount of times in an artists career. Generally, the first of such events occurs only after years of learning, practice and exploration. For example, an artist such as Pablo Picasso who has the reputation of being highly creative, had 6 distinct major creative events over a career that spanned 70 years[6], whilst an artist such as Francis Bacon had only one significant such event.

The found–developed–created artist-style-space can been seen as a subspace of a movement-style-space, itself a subspace of the period-style-space, which is a subspace of the art-space. The artist-style-space contains subspaces or regions related to periods, that include subspaces related to series which are populated by individual artworks. Each style-space and subspace is characterised by a number of dimensions, and governed by specific rules that define elements such as conceptual stances, colour harmonies, type of composition. Associated to each of these style-spaces and subspaces are strategies and techniques used for the development and production of artworks. Meta-rules, strategies and techniques associated to parent spaces would also apply to subspaces. From this perspective, artworks can be seen as constructed systems that are functional in the style-space.

Artists develop–create–invent their personal–original–identifiable style-space, these style-spaces have commonalities with existing style-spaces, this occurs through the understanding of rules, techniques and strategies related to style-spaces established by other artists, achieved for example by copying existing artworks. These sets of inherited strategies, rules and knowledge (schemata) are developed–adapted–personalised through practice, exploration and experimentation. As in other domains of expertise, schemata can be described as the sets of knowledge and strategies for: action, information gathering, evaluation, planning required to achieve a task, with the schemata for high-level tasks relying on lower-level schemata.

Works of art can be considered as systems with a visual appearance that has been invented–developed–executed to act on the ob-

[4] In the modernist period: the suprematists, De Stiljlt and in contemporary art: minimalism, conceptual, op-art, process art. Motivations for the removal of the human are multiple such as spiritual, purity, the removal of sentimentalism, anti-individualism, art for art and political motivation.

[5] People born after the democratisation of the personal computer

[6] The number varies depending on the historians and sources but in our view: Blue Period, Pink Period, African Period, Cubism, Surrealism, Classicism

server's perceptual/cognitive system to produce a certain aesthetic–artistic–emotional experience. We can consider the establishment of a style-space (artist, period, series) as the result of transformatory creativity, and the development/execution of artworks as the result of exploratory creativity [2].

Although level of creativity is generally considered a measure of artistic value, it is difficult to examine it without considering mastery and expertise, and in the context of the arts these elements are intimately linked, one being useless without the other. Without a level of expertise, major creative events in an artists career would be unlikely to occur or be recognised as such by the artist, and without mastery, style-space exploration would not be possible.

The loose formalisation described above enables the attribution of levels of creative/artistic autonomy to computational systems that produce artworks. The highest level of autonomy being attributed to systems that are capable of creating–developing–finding their own style-space, with the level of autonomy decreasing when the system is only capable of establishing period-style-space, decreasing further still for the ability to establish series-style-space, which is considered at a higher level of artistic autonomy than the production of individual artworks.

The establishment and existence of these rich coherent style-spaces is essential for the consideration of the produced artifacts as works of art, as they provide recoverable traces of historical context and artistic coherence. Developing systems that only mimic the superficial appearance of artworks, would be like looking at a mathematician's blackboard and based only on superficial observation, developing a system to produce images using the same signs arranged in the same manner, and then expecting this image to express a mathematical truth.

The systems described in this paper display only the lowest level of artistic autonomy, that is to say the systems are not yet able to develop their own style-spaces. As artist–researchers–developers we have been responsible for the establishment of the artist, period and series style-spaces. At present the agents are responsible for the creation of artworks that are functional in these style-spaces. It is important to note that the stylistic appearance of the artworks produced is not that of a pastiche. The artworks are interpretations naturally influenced by the agent's characteristics and capabilities (physical and computational).

3.2 Intentionality

Although it is difficult to define what art is, we can define the artist as an expert who conceives and produces artifacts that are exhibited in galleries, museums and other public spaces, with these artworks being exhibited, appreciated, acquired and collected by actors or institutions belonging to the art-world[7].

Our goal is not for the systems to mimic precisely the appearance of human-made artworks so as to create pastiches, but to produce artworks that trigger similar aesthetic emotions in the viewer as a human-made drawings or paintings. With this in mind we should identify and define a range of characteristics that enable the perception and categorisation of human-made objects as being works of art by the art-world.

There are two perspectives on visual arts perception/appreciation:

[7] In the same manner as scientist's research publications are reviewed by specialists/experts we believe that it should be the same for art. Experts in the arts include: curators, collectors, art critics, amateurs, patrons, art historians and practitioners. In this paper the term art-world can be defined as a range of people that are able to form an informed opinion about artworks.

Figure 2. Paul sketching Amy, 2012, photo by Steph Horak

a) the cognitive approach that generally postulates that artworks have intrinsic artistic value which can be understood through studying the effect of the appreciation of artworks at the neurological level; *b)* the humanities approach that argues that an artwork's appreciation is strongly influenced by contextual knowledge related to its production.

In the cognitive approach, scientists from empirical aesthetics [21, 16, 24], neuroaesthetics [6, 33, 26] use tools and methods from psychology and neurosciences to further the understanding of the mental and neural processes involved in art appreciation. This focus on perception and observable cognitive processes does not take into account the humanities approach and even often rejects it, looking for artistic universals that are sufficient in accounting for artistic value.

The humanities approach, without rejecting the intrinsic aesthetic value of artworks, considers the role of context (artistic, historical, sociological), as well as the artist's intentions as both having an effect on the production and appreciation of artifacts as works of art.

In the context of computationally-produced artworks, it is easier to take into account the cognitive approach rather than the humanities approach, therefore developing systems that produce images that superficially look like artworks, and focusing on this aim. But considering that actors of the art-world in the majority take the humanities perspective, we have to accept this framework if we want the artifacts produced by our systems to be appreciated as work of arts. As a consequence we must consider that intention-rich artifacts that can be associated to coherent artistic and historical contexts are more likely to be considered works of art.

Pignocchi in [23] unifies the cognitive and humanities approaches using a model of the experience of artworks based on the mechanisms of intention attribution. Pignocchi claims that the traditional notion of "intention" is too restrictive when considered in the humanities approach, and argues that an observer can not only recover the artist's high-level overt intentions, but also *"all the mental states—conscious or not, propositional or not—that have played a causal role during the production of the work"*. Pignocchi's framework postulates that:

> "the perception of an artwork activates rich intention at-

tributions, implicit and explicit, conscious and unconscious, where the intentions attributed are potentially any kind of mental states that could have played a causal role during the production of the work."

An artwork contains the history of the intentions that led to its creation, and some of this history is recoverable by the observer and contributes to the artwork's appreciation. At low level, the brushstrokes, marks or pen's traces on the paper, their overlapping and interactions, are a direct memory of the artist's decision and actions over time and are recoverable by the observer.

Based on Pignocchi's model we can hypothesise that when produced using the rules and strategies related to coherent styles-spaces, traces and cues are embedded in the artworks that enable their association to rich and complex artistic and historical contexts.

3.3 Embodiment

There are a number of fundamental differences between a computational system designed to produce images that look superficially like drawings or paintings and an embodied computational system that physically produces paintings or drawings. These differences are essentially due to: *a)* the characteristics of the robot and the characteristics of the medium; *b)* the physicality of the produced output; *c)* the perceived agency of the system.

The characteristics of the robot, the movements and tasks it can achieve, are highly constrained and require careful consideration and complex control in an embodied system. In the case of painting:

- the control of the pressure of the brush onto the canvas
- the angle at which the brush touches the canvas
- the path of the brushstroke
- the velocity of the movements
- controlling the quantity of paint loaded onto the brush
- the control of the paint viscosity, texture, transparency
- mixing paints to obtain specific colour
- the time it takes to move from one area to another area
- the time it takes to reload the paint
- time taken to clean the brushes
- the number of brushes it can use
- the characteristics of paints and brushes
- the range of pigments available

These elements impose strong constraints or are too complex to control when developing an embodied system that acts in the real world. In a purely computational system, they are controlled effortlessly and without constraints, by simply changing some parameters. In this situation one could think that it is more logical to develop a non-embodied system capable of producing works of art, but this is not the case. The complexity and limits encountered when developing an embodied system forces the developer to simplify and adapt to the constraints, and as a consequence bring a form of stylisation that defines and charaterises with precision a local style-space, based on both the robots and the mediums' characteristics.

With regards to the physicality of the output, paintings and drawings are physical objects that have specific qualities that can only be appreciated when in their presence. Artworks are objects with aesthetic qualities that are valued, not only for what they represent and how they depict but also for their material properties such as scale, textures, how the surfaces reflect light. Reproductions in books or printed posters in no way enable the experience of appreciating the artwork in direct physical contact; a digital print of an image produced by a non-embodied system, even of the best possible quality, is not capable of encapsulating the aesthetic richness of a drawing or painting. More prosaically, a painting or a drawing has more value than a print commercially, due to its uniqueness and its material qualities. Furthermore, physically produced objects encapsulate recoverable memories of the succession of processes that have created them [22, 13, 9], enabling in turn a perception of the succession of the artist's intentions over time.

Although not often considered or known by the public at large, most artists of a certain status had/have assistants and collaborators who take charge of some of the tasks involved in the production and commercialisation of artworks. In the past (pre-nineteenth century) young artists would learn their craft and acquire the necessary skills by working for Masters, some of them becoming Masters in their own rights, others pursuing roles as workers in the Master's studio/workshop. In regards to artist's assistants, the contemporary situation is not totally dissimilar in that art students and young artists often work for practicing artists as a way to earn an income and learn certain aspects of their profession. This use of assistants to aid the conception and fabrication of artworks does not affect the works capability of being perceived and considered as art. With this in mind, the use of computational systems should not prevent their production being considered as art as long as they are perceived as such by the audience.

As for the perceived agency of the system, studies have shown that humans have a tendency to consider robots as social agents, as having personalities and agency [3, 4]. It is likely that with the knowledge that an artwork has been produced by a robot, the observer will be able to consider the robot as the originator of the intentions that led to the artworks' creation, as the artist.

Figure 3. Sketch by Paul, 2011

4 Paul

This section gives a brief description of a robot named *Paul*, an artwork and obsessive artificial drawing entity that was created using technologies and ideas developed in the context of AIkon II, a project conducted at Goldsmiths, University of London by the first author in collaboration with Prof. Frederic Fol Leymarie. A detailed description of Paul and the algorithms driving it can be found in [31].

The AIkon II systems we initially developed to investigate the observational drawing practice, were not embodied in a robot. Yet it rapidly became apparent that embodiment was necessary. There are a number of factors that motivate this decision: *a)* the physicality of drawing, having a system that draws by physically moving a pen in contact with a medium: we have seen previously that the way artworks are made has a strong influence on the way they are perceived; *b)* the manner in which robots are controlled: contemporary software framework architecture to control robots is composed of concurrently running processes that communicate. This architecture forces one to approach problems from an angle which is more appropriate for understanding how to implement complex systems, such as drawing from observation; *c)* the advantages for dissemination: the appeal that robots have on an audience attracts interest and attention.

The strategies, techniques and aesthetics implemented in Paul are derived and in accord with those characterising Tresset's drawing-style-space. This space was developed and influenced by years of practice and by studying in particular the drawings of Leonardo Da Vinci, Egon Schiele and Alberto Giacometti.

Paul is designed to limit complexity, and only to fulfill its function, drawing. As such it is composed of a three-jointed planar arm with an extra joint to raise and lower the pen in contact with the paper and an actuated pan and tilt camera used as its eye. Both arm and eye are bolted to a single school desk where the drawing paper is attached using pins (fig.2).

Figure 4. Painting from the series: Paul's Memories, 2013

In our experience of exhibiting Paul's production to a wide audience, we have noticed that one of the interesting properties of the produced portraits is that they are perceived, considered and appreciated as drawings. When observing a series, Paul's drawings are recognised as drawn by the same author meaning that they display an autographic style. Contrary to other computational systems that produce drawings from photographs such as AIkon-I [30, 29], drawings produced by Paul do not display the same serial uniformity of treatment (fig.5).

Figure 5. Sketches by Paul, 2011

A number of factors can account for the perceived quality of Paul's sketches, such as the choice of paper, layout and composition. When Paul draws lines, their paths are extracted from the responses of Gabor filters [20]. Such filters are known to be good models of simple cells in the early visual cortex (V1) [12] and as a consequence accentuate what would be perceived as salient features [11]. The use of simulated visual feedback to constrain and evaluate the random exploration during the shading process is sufficient to produce patterns that are perceived as being the result of an intentional process. Paul's drawings are the result of a sequence of movements and as such they are the record of a process. Evidence that the traces forming part of a drawing by Paul are the results of movements can be found in the lines irregularities. Although these irregularities are not akin to the imperfections a human might produce, they have characteristics that could only be the result of a pen in motion driven by an articulated arm. Furthermore, the layering of successive lines and of successive shading patterns adds to the drawing being perceived as the result or consequence of a sequence of intentional movements/processes (fig.3).

Paul and its drawings have been widely exhibited including in major art museums. Since 2011, *the Pauls*[8] have drawn thousands of people and hundreds of drawings have been acquired by the public in galleries, museums and art-fairs. One of Paul's drawings is in the Victoria and Albert Museum's collections.

5 Painting Paul's Memories

This section reports on the first author's research and the artistic outcome of a 9-month residency with the research team in Konstanz led

[8] There are to date ten Pauls in existence.

by Prof. Oliver Deussen. The aim of the residency was for Tresset to experiment with the robot *e-David* [14] and develop algorithms to drive it with the view of creating series of paintings. The aims of the team in Konstanz for the e-David system are: *i)* approximating the manual painting processes by a machine; *ii)* finding out to what extent the system is able to produce artistic looking paintings; *iii)* looking for new forms of visual representation that are especially suited to painting machines; *iv)* finding out how to introduce high-level semantic information into the process.

In recent years, methods for image understanding have developed a lot, so it is plausible that painting machines of the future could 'know' what they draw and automatically adapt their painting strategy [14]. In the subsection 5.1 e-David's technical setup is described, followed in subsection 5.2 by a description of the ideas and the techniques used to paint the *Paul's Memories* series.

5.1 e-David

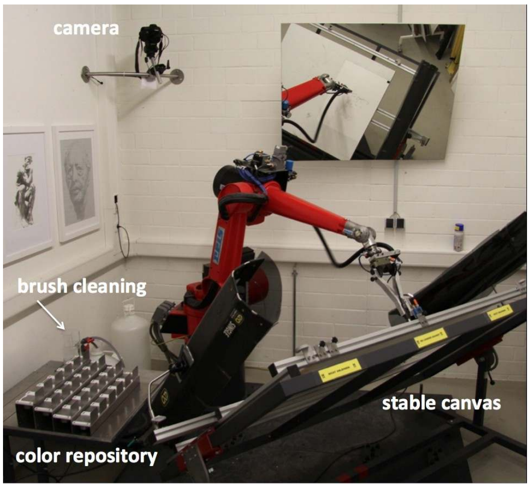

Figure 6. e-David

e-David consists of an industrial robot (Reis RV20-6) equipped with a specialized picking device for grasping up to five different brushes that can be used during the painting process (fig. 6). Paint is stored in a repository; so far the system allows for the use of up to 24 different colors. The robot accesses a paint container by mechanically opening the cover plate and dipping the brush into the color. The setup is amended by a cleaning and drying facility. Together with the brush grasping tool, a distance sensor is mounted on the robot arm which allows the system to compensate for curved canvases and even the mechanical tolerances of the robot while moving the brush over the canvas. Such tolerances are in the sub-millimeter range but still visible for image styles with long strokes and fine brushes.

The visual feedback of the system is created using a Canon EOS 5D Mark II SLR camera with a 21MPixel sensor and a fixed 50mm lens. This provides a resolution of about 1mm on the canvas. Two specialized fluorescent tubes with polarisation filters illuminate the canvas, and the camera also has a polarisation filter.

The robot is controlled by an assembler-like language, and a server application that controls the machine and accepts XML commands was built to ease its use. Most commands are plotter-like instructions such as pen selection and drawing, but specialized commands for measuring the canvas and the handling of brushes and colors have also been implemented. A second server is used for the camera. The camera allows for the control of all necessary functions from the computer. Images are created in XYZ color space and are calibrated using geometric calibration patterns and color sheets. The geometric calibration is within the range of a single pixel, and to date the color calibration is in the range of 5%, which is sufficient for current applications but must be enhanced in the future. The rationale of building an industrial robot rather than just using a pen plotter, was to remove the limitation of the use of only one type of physical instrument when creating paintings in this manner.

The current implementation of e-David can successfully use any brush or pencil. It can paint with a number of painting mediums such as ink, acrylic and oil paints. The inaccuracies of brush-stroke rendering and the unpredictability of color interactions on a canvas are reduced with the use of the feedback loop. Tests have shown that this behavior has proved to be very important since it not only allows the use of relatively simple simulation techniques, but can also easily be adapted to different strategies for placing strokes.

5.2 Paul's Memories

The aim of the research and work with e-David conducted during the first author's residency was to produce series of paintings that could potentially be appreciated as work of arts. As our approach postulates, there is a need for artifacts to be located in a series-style-space, period-style-space and the individual/original artist-style-space, and at present it is not possible for our systems to develop their own style-spaces. As a consequence it is necessary that the researcher/developer takes responsibility for establishing the rules, parameters, strategies and techniques related to these spaces in order for the produced artifacts to be perceived as works of art.

When artists are in the research-experimentation phase that will lead to the establishment of the artist-style-space, a number of possibilities are explored, techniques developed to be later abandoned due to various factors such as technical limitation or more promising avenues being found. For the work developed here, such an abandoned space is explored and reused. After a couple of years of exploration this potential style-space was abandoned essentially due to the artist's skills limitations.

The style-space explored here was established as the result of experimentations by the first author with techniques and styles of pre-impressionist periods, combining the use of chiaroscuro, impasto and glazes. This research was aimed at finding ways to depict human faces and was influenced by Rembrandt, Francisco Goya and also by twentieth century painters such as Leon Kossoff, Howard Hodgkin and Yan Pei-Ming.

The style-space is defined and characterised by: a lighting akin to chiaroscuro, achieved by relatively few thick,well-defined brush-strokes (impasto), with tonal variations achieved from a succession of thin monochromatic uniform transparent layers (glazes).

The series-style-space established for Paul's Memories series is defined and characterised by: a stylisation reminiscent of early computational images achieved by using only straight brushstrokes, with a limited number of orientations, a limited discretised tonal range (8 levels), and three distinct brushstroke widths. The choice of the series subject matter: the individuals portrayed in this series are taken from Paul's memories (Paul stores a digital image of each person it has drawn).

5.2.1 The painting process

We provide here a brief and superficial description of the painting process and of the algorithms used.

- The robot only paints with thick white acrylic paint on a pre-prepared canvas board of a mid-grey colour, with three different sized brushes.
- Nine paintings are painted at a time, arranged on a 3 x 3 grid. The full painting process for each set takes between 24 and 36 hours[9].
- The uniform monochrome glazes are applied by a skilled human operator. Although this operation is seemingly simple when executed by a human with a brush, it requires constant adaptation of speed and pressure, and removal of excess paint based on visual feedback. This adaptation is necessary to achieve extremely thin and uniform layers. It would have been a complex research project in itself to get a robot to perform this phase of the painting process. It could perhaps have been achieved automatically by having the robot spray the glazes, but we did not have the resources, competences and time to develop and implement this technology. Furthermore, the delegation of a non-creative repetitive low-level task to a human adds an interesting element to the context in which the paintings were created.
- At each iteration the glazes' colouration alternates between a dark bluish purple and a dark reddish purple.
- At each iteration the glazes' transparency is increased by 50%.
- There are 8 cycles, one for each of the tonal values.
- Each cycle:
  - the robot takes a picture of the canvas (PCI)
  - brushstrokes, length, orientations and locations are computed
  - the robot paints a set of white brushstrokes with a large brush
  - the robot takes a picture of the canvas (PCL)
  - the brushstrokes' length, orientations and locations are computed
  - the robot paints a set of white brushstrokes with a medium brush
  - the robot takes a picture of the canvas (PCM)
  - the brushstrokes', length, orientations and locations are computed
  - the robot paints a set of white brushstrokes with a small brush
  - when the paint is dry the human operator applies a glaze
  - For each cycle a level of the graylevel image, discretised using a k-means clustering algorithm, is used as a binary map (GM)
  - For each of the brush sizes, a brushstroke map (BM) is computed. BM is a low-resolution texture-flow map of the subject image, the resolution of each map being related to the brush size (larger brush size, lower resolution). The orientation of each cell of the map is limited to 8 angles
  - For the larger brushstrokes, GM is used as a mask to select from BM, the set of brushstrokes to be painted. To select the medium and small brushstrokes set, using PCI and PCL or PCM as visual feedback, a new binary map is computed that represents which area of GM has not yet been painted

The paintings produced in the Paul's Memories series are arguably in a coherent identifiable personal style, and have been evaluated by a number of artists, a curator from a major contemporary art museum and other amateurs as being of artistic interest (fig.1, 4). From discussions with people who have seen the paintings it is clear that they are evaluated with the same standards as if they were painted by a human artist. Although these paintings do not resemble work we have painted ourselves, we would be very satisfied to achieve such work.

[9] 5 sets were painted, 2 of them were failures

6 Conclusions

In this paper we have described views, considerations and strategies that shape and motivate our development of artistically-skilled embodied agents used to produce artworks. The numerous discussions we have had with actors of the art-world and the art-loving public, around the drawings by Paul and the Paul's Memories series of paintings, have convinced us that these artifacts are considered as works of art.

The question of evaluating the level of creativity of a computational system is not a simple one. If we consider the perceived creativity in both of the systems that we have presented in this paper, they are both perceived as being creative and artistic by the public at large and art specialists. Even we, the creators of the systems' are very often surprised by the artifacts produced by Paul and e-David. There is a level of originality in each artifact and yet they can be perceived as displaying an autographic style. We are aware that even with the use of visual feedback and other low-level evaluations, both systems can not at present be evaluated as being creative. Although shaped by our personal style-spaces, the systems have different strengths and limitations to our own and as a consequence the work they produce does not belong to us. It is an interpretation of our style shaped by their characteristics, and for a number of reasons, objective or subjective, we evaluate their artworks as more interesting and artistically valuable than what we have produced by hand.

Now that we have developed skilled agents, the next stages will aim at progressively increasing the level of autonomous creativity without decreasing the quality of the artworks produced. The framework presented in this paper encourages a development approach from the ground up. In the context of drawing and painting, the lowest level behavior is mark-making. The next step in our research is to use reinforcement learning (RL) to allow the robot to learn to gain more control when tracing lines without losing its apparent spontaneity and without removing traces of its intrinsic characteristics. RL is a technique using trials, reward and punishment to learn, or to modify behaviours.

As we progress in enabling the systems to gain more creative autonomy by developing their own style-spaces, we will have to consider the difficult issue of the system's ability to assess the quality of its own work. The objective is for the Paul and e-David systems to be able to be influenced by their artistic environment, art history, contemporary arts and their own practice, to invent their own style-space where our influence is only one amongst many others.

ACKNOWLEDGEMENTS

Paul is based on technologies developed in the context of the AIkon-II project which was co-directed with Prof. Frederic Fol Leymarie and was in part supported by a Leverhulme Trust grant. Patrick Tresset's nine month residency in the informatik department of the University of Konstanz was financed through a senior fellowship at the University's Zukunftskolleg.

REFERENCES

[1] R. F. Antunes and F. Fol Leymarie, 'xtnz-an evolutionary three-dimensional ecosystem', in *Proceedings of the 4th International Conference on Digital Arts, Artech*, pp. 201–204, (2008).

[2] M. A. Boden, *The creative mind: Myths and mechanisms*, Routledge, 2003.

[3] C. Breazeal and B. Scassellati, 'How to build robots that make friends and influence people', in *Intelligent Robots and Systems, 1999. IROS'99. Proceedings. 1999 IEEE/RSJ International Conference on*, volume 2, pp. 858–863. IEEE, (1999).

[4] C. L. Breazeal, *Designing Sociable Robots with CDROM*, MIT press, 2004.

[5] E. Burton, 'Thoughtful drawings: A computational model of the cognitive nature of children's drawing', in *Computer Graphics Forum*, volume 14, pp. 159–170. Wiley Online Library, (1995).

[6] A. Chatterjee, 'Neuroaesthetics: a coming of age story', *Journal of Cognitive Neuroscience*, **23**(1), 53–62, (2011).

[7] S. Colton, 'The painting fool: Stories from building an automated painter', in *Computers and creativity*, 3–38, Springer, (2012).

[8] S. DiPaola, 'Exploring a parameterised portrait painting space', *International Journal of Arts and Technology*, **2**(1), 82–93, (2009).

[9] D. Freedberg and V. Gallese, 'Motion, emotion and empathy in esthetic experience', *Trends in cognitive sciences*, **11**(5), 197–203, (2007).

[10] E. J. Ilfeld, 'Contemporary art and cybernetics: Waves of cybernetic discourse within conceptual, video and new media art', *Leonardo*, **45**(1), 57–63, (2012).

[11] L. Itti, C. Koch, and E. Niebur, 'A model of saliency-based visual attention for rapid scene analysis', *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, **20**(11), 1254–1259, (1998).

[12] J. P. Jones and L. A. Palmer, 'An evaluation of the two-dimensional gabor filter model of simple receptive fields in cat striate cortex', *Journal of Neurophysiology*, **58**(6), 1233–1258, (1987).

[13] M. Leyton, *The structure of paintings*, Springer, 2006.

[14] T. Lindemeier, S. Pirk, and O. Deussen, 'Image stylization with a painting machine using semantic hints', *Computers & Graphics*, (2013).

[15] P. Machado and A. Cardoso, 'All the truth about nevar', *Applied Intelligence*, **16**(2), 101–118, (2002).

[16] C. Martindale, *The clockwork muse: The predictability of artistic change.*, Basic Books, 1990.

[17] P. McCorduck, *Aaron's code*, WH Freeman & Co., 1991.

[18] L. Moura, 'A new kind of art: The robotic action painter', *X Generative Art Conferece, Politecnico di Milano University*, (2007).

[19] A. Pease, D. Winterstein, and S. Colton, 'Evaluating machine creativity', in *Workshop on Creative Systems, 4th International Conference on Case Based Reasoning*, pp. 129–137, (2001).

[20] N. Petkov, 'Biologically motivated computationally intensive approaches to image pattern recognition', *Future Generation Computer Systems*, **11**(4), 451–465, (1995).

[21] R. W. Pickford, 'Psychology and visual aesthetics.', (1972).

[22] A. Pignocchi, 'How the intentions of the draftsman shape perception of a drawing', *Consciousness and cognition*, **19**(4), 887–898, (2010).

[23] A. Pignocchi, 'History and intentions in the experience of artworks', *Topoi*, 1–10, (2012).

[24] A. P. Shimamura and S. E. Palmer, *Aesthetic science: Connecting minds, brains, and experience*, Oxford University Press, 2012.

[25] K. Sims, 'Evolving 3d morphology and behavior by competition', *Artificial life*, **1**(4), 353–372, (1994).

[26] R. L. Solso, *Cognition and the visual arts*, MIT press, 1996.

[27] A. M. Spalter, *The computer in the visual arts*, Addison-Wesley Longman Publishing Co., Inc., 1998.

[28] S. Todd and W. Latham, 'Evolutionary art and computers', (1994).

[29] P. Tresset and F. Fol Leymarie, 'Generative portrait sketching', in *Proceedings of VSMM*, ed., Hal Twaites, 11th International Conference on Virtual Systems and MultiMedia, pp. 739–748, Ghent, Belgium, (2005). Archaeolingua publishers.

[30] P. Tresset and F. Fol Leymarie, 'Aikon: the artistic/automatic ikonograph', in *ACM SIGGRAPH 2006 Research posters*, p. 37. ACM, (2006).

[31] P. Tresset and F. Fol Leymarie, 'Portrait drawing by paul the robot', *Computers & Graphics*, (2013).

[32] G. A. Wiggins, 'A preliminary framework for description, analysis and comparison of creative systems', *Knowledge-Based Systems*, **19**(7), 449–458, (2006).

[33] S. Zeki, 'Art and the brain', *Journal of Consciousness Studies*, **6**(6-7), 6–7, (1999).