# The creative GAN: simulating creative behaviour in GANs with IoT

R.F.L. Beening

[1] First supervisor: dr. Victor de Boer
[2] Second supervisor: Joost de Boo
[3] Second assessor: dr. Karine Miras

**Abstract.** Generative Adversarial Networks (GAN) are frequently used in online art-generating tools because of their ability to produce realistic artefacts. The growing use of art generators raises questions about the ability of a computer program to create art. Making art requires a human skill that is challenging to find in computers, namely creativity. We argue that a GAN is predominantly good in making more of what is already known and that this contradicts being creative. Creativity comes from experiences and interactions with people and the environment. Therefore, we propose to combine the GAN with an IoT to make it behave more creative. We first review the literature on computational creativity and evaluate the original GAN based on the theories from the review. We identify a set of three requirements for a new creative GAN design. Based on the requirements, we design a system that extends the original GAN with an interactive IoT system to simulate a more creative process. The prototype of the design shows a successful implementation of creative behaviour that can react to the environment and gradually change the direction of the generated images. The generated art is evaluated based on their creativity by doing task-based interviews. The results show that the generated images are creative depending on the participant's view of creativity. We recognise that these results are indifferent due to creativity's cultural and social implications.

**Keywords:** Generative Adversarial Networks · Interactive Art · Computational Creativity · Internet of Things.

## 1 Introduction

Generative art gradually inserts itself into our lives as digital tools to create art become more accessible. They can appear in the form of music, the visual arts, video art, virtual reality, robotics, theatre, and many more [1, 2]. Generative art is art created using an autonomous system, often a computer program [3, 1, 4]. There is always a part in generative art that the artist has no control over, and in some cases, the artist is entirely absent. To some computer programmers and artists, generative art is evidence of computers becoming their own creative artists [5, 6]. However, there seems to be a line between art made through a computer and art made by a computer. The more creative autonomy the computer program has, the more doubts people express if the work can be called art [1, 7].

One of the computer programs that opened up the world of generative art is the Generative Adversarial Networks (GANs). GANs are often used in online art-generating tools, making it easy for anyone to generate a realistic piece of art. The original GAN, introduced by Goodfellow et al., is designed to excel at producing fake samples that cannot be distinguished from real samples [8]. They describe the process as analogous to counterfeiters trying to fool the police. The counterfeiters want to make fake currency that looks real. The police then try to detect the fake currency from the real currency. Both teams improve their skills after numerous attempts until the fake currency is indistinguishable from the real currency. The same explanation works when describing an art forger trying to replicate an artwork and a critic trying to spot the difference. The main difference between an art forger and a genuine artist is that an art forger is penalised for being creative, whereas an artist is celebrated for being creative. Therefore, we question whether a GAN is suitable as an art generator. If we value creativity in human artists, we should hold the GAN as an art generator to similar standards.

Creativity is described as surprising, new and valuable ideas that come from human experiences and interactions with people and the environment [9, 10, 2, 11]. Therefore, we propose that the GAN can become more creative by simulating human creativity through real-world interactions. One way to do this, is through the use of the Internet of Things (IoT). IoTs are integrated with people's lives and have a personalised view of the world around them. An extension on the GAN with an IoT could give the GAN its own unique experience that depends on environmental factors and usage within the IoT.

Simulation of creativity in GANs, or more general in computers, can help us understand how creativity happens and what it means to be creative [12, 9]. In this paper, we explore creativity around GANs. We investigate how a GAN can become an autonomous artist that generates creative art through IoT. We do this by reviewing the state-of-the-art within computational creativity and evaluating if the GAN is currently creative. From this examination, we derive requirements for a design that we will refer to as the creative GAN. The new design combines a basic GAN with an interactive IoT system to simulate a more creative process. We build a prototype of the design that can produce unique images. Ultimately, we review the creative GAN's creative behaviour and evaluate the output through task-based interviews.

## 2 Literature review

This section looks into computational creativity and IoT. We first cover notions within computational creativity, generative art and autonomy, and bias in evaluating creative artefacts. Following is a brief section on the IoT's general structure.

### 2.1 Computational creativity

Computational creativity is the study of building programs that can exhibit behaviour that would be deemed creative in humans. These creative programs can be used in various artistic tasks, maths and engineering. Next to these practical applications, computational creativity can help us understand human creativity. Margeret Boden is one of the most prominent contributors in the field of computational creativity. Her seminal contribution to computational creativity is not about what makes a computer creative but what behaviour can make them appear creative [9, 13, 2].

**Value and novelty** In Boden's [2004] explanations of creativity, she talks about value and novelty of creative artefacts. A considerable number of other authors discussed these same terms concerning creativity and are one of the most occurring concepts [9, 10, 2, 11]. Value is the measurement of how good the solution fits the given problem, and novelty describes how different artefacts are from existing samples of its sort [14]. An artefact with high value and low novelty makes a solution already exists. High novelty and low value create something new but will not be useful for the given problem. In other words, the value measures the information content of the artefact, and novelty measures the entropy. There is no creativity in only information, and there is no creativity in entropy [15, 16]. Only when both of the notions are in balance there is creativity.

**Creativity and surprise** A related term to value and novelty that appears in numerous studies is surprise [14, 9, 17]. Surprise describes the balance between value and novelty. Something is no longer surprising when it deviates considerably from its original form (high-novelty). Vice versa staying too close to the original is not surprising anymore (high-value). Three sorts of surprise correspond to three ways creativity can happen. These are combinational, explorative, and transformative creativity [9].

*Combinational creativity* involves connecting ideas in an unfamiliar way. Combining two concepts is simple, but making combinations that make sense and do not seem random is much more challenging. There must be an understanding of each concept and possible relations to each other to make meaningful combinations. Combinational creativity seeks links between the concepts that were previously not seen; this link is what makes the combination make sense and engaging.

*Explorative creativity* is not about going straight for the solution but for the unconventional solution that still solves the problem. To illustrate how this works, we introduce a rule-based problem; it is blue and square. There exist various ways to make a blue square. Some solutions are evident, while others are more interesting and surprising. What is the size of the square? Light or dark blue? Does it have a texture? Making this square a hairy baby blue square the size of the Atlantic ocean still solves the given problem but beyond expectations. Explorative creativity is a great way to develop new ideas that still fit the given problem space.

The most complex of the three forms of creativity is *transformative creativity*. Where combinational and exploratory creativity worked within a given search space, the goal of transformative creativity is to search beyond this given space. We previously explained explorative creativity as treading along the edges of the rules; transformative creativity wants to bend the rules. Practical examples are the creation of new genres or the reconstruction of an existing genre. Creative programmers often avoid transformative creativity because this type of creativity often produces high novelty but low-value artefacts [12]. Therefore transformational creativity is believed to be better approached as inspiration for a program rather than a goal to be achieved [18].

The notions of surprise carve a path toward achieving creativity and can be used as inspiration for computationally creative design. According to Boden [2004], creativity cannot be reduced to a set of formal conditions due to the many interpretations of creativity in different cultures. Therefore, these notions are not a definite set, but they may extend beyond the given conditions [9, 10].

### 2.2 Generative art and autonomy

Creative computers will have to fight the impression that it is limited as an artistic programming tool [16]. The challenge is to find a way to minimise human intervention in the creative process. If the program relies too heavily on human creativity to produce its artefacts, it becomes a question of whether the computer is making the art or the human is using a computer to make art. More creative autonomy within the generative process could switch the computer from being the tool to being the artist.

**Defining generative art** In the strict definition of generative art that Boden [2009] suggests, it is a prerequisite that a computer can run by itself with zero interference. Only when this is the case can the output of the process be called generative art. She excludes the artist entirely from the program, except when the artist creates their desired input and program. Nothing is left to do but let the computer do the work. She acknowledges that the definition is too strict and makes it easy to exclude certain types of art, such as interactive art. Therefore, she also suggests a less strict formulation where the process can have minimal or zero interference from a human. What minimal entails is open for interpretation. Regardless, it serves as a caution when considering a (fully) autonomous creative system. It is important to be aware of who intervenes in the generative process and what kind of intervention it is.

**The role if interaction and perception** Several studies suggest methods for the implementation of human-like experiences without the direct intervention of a human. Some of the proposed methods to achieve this could be by adding interaction, perception or using big data [19, 16, 20, 21]. Although initially, interaction sounds counterintuitive, Jennings claims autonomy does not mean there is no interaction possible. He suggests autonomy can require more interaction, seeking beyond the programmer. He remarks that this is not yet proven but explains his reasoning by taking it back to how humans' creative process works: outside influences from the social world are leading in the creation of art. Even more so, the interaction with other artists and knowledge of other styles make artists develop their own styles [16]. Similar ideas come from the use of perception in generative art. Perception enables humans to interpret ideas differently from others, offering insight into fresh perspectives and possibilities [21]. The idea is that people use their perceptions to facilitate their art. For example, one must know how to listen to compose music. Perception also plays a big part in the retrieving of inspiration. Art school encourages students to look at art from other artists and seek inspiration from the environment for the creation of art [22]. The exact methods of perceiving and creating for a human may differ from that of a computer, but some form of perceptual grounding is requisite for a genuinely creative system [23, 21].

Some questions regarding the use of interaction and perception in a creative computer remain to be addressed. As far as we know, no previous research has investigated this in practice. The theoretical grounds are promising given the goal of generative art; to decrease human intervention and increase experience.

### 2.3 Bias in evaluating creative artefacts

The term creativity originates within human society. Our whole notion of creativity tends to manifest various social and cultural prejudices that are affected by various environmental factors [2, 24, 25, 26, 9]. These prejudices form conscious and subconscious in humans and cannot be ignored in evaluating creativity and art. We address issues caused by context bias and discuss the role of process and human evaluation in computational creative art.

**Contextual bias around art and generative art** Context is an all-powerful tool in the perception of art. A great example is how the artist Banksy approaches sales of his artworks. At times his art is sold on the streets of New York for a price of $60, where people walking by the stand show no interest in buying the art. Later the same work is up for auction and sells for up to $80,000. Things changed in this scenario that made the $60 painting a $80,000 painting. People now recognise it as an authentic Banksy, presented in a renowned auction house. The work remains the same as before, but the context around the painting significantly increases the value. We can see a similar phenomenon when discussing computational art. People tend to value human-made art more than generative art made by algorithms when researchers disclose the contextual information [6, 27]. Withholding this context shows that there is initially no bias in their assessment [25, 26]. However, humans judge art not only on the aesthetics but also on effort, skill, and creativity [6, 27]. Withholding context makes the assessment of art less biased while, at the same time, more complicated to evaluate.

**Evaluation by process** Several authors have recognised that bias is a problem in evaluating computational creativity and propose evaluating the process instead of the output. The main advantage is that the human prejudices are removed from the evaluation, making it a more objective standard to use [12, 28, 29, 26]. The demand for an unbiased evaluation of the process has resulted in numerous proposed frameworks, for example, the FACE and IDEA model, SPECS, and the creative tripod [30, 31, 26, 6]. However, none of the proposed frameworks became the research standard, and human evaluation is often preferred over the use of a framework.

**Evaluation by humans** Not all authors express the same belief that bias is a reason to avoid human evaluation. Most of the research on computational art is done by conducting a human evaluation [25, 7]. The authors are aware of human bias in the evaluation and consistently mention this as a point of debate. Ritchie advocates the use of human evaluation for generated artefacts [2]. He explains that we must judge computer creativity as we usually judge human creativity. If we want our computers to achieve creativity in the way humans achieve creativity, we must assess them on observable factors. Ritchie dismisses the whole process of creation in the judgment of the art. The amount of information present when looking at an artwork from a human is supposed to be the amount of information present for the generative art. Considering more information for either artist in the judging process would be unfair. Ritchie's statement on only judging the observable factors contradicts the earlier discussed views of Colton [2008], who believes that effort, skill, and creativity also play a big part in the assessment of art. It seems possible to avoid bias by only judging the observable factors, but according to Colton [2008], this makes it impossible to judge the entirety of the art.

The underlying creative mechanisms in humans are hard to measure, making the evaluation of creative artefacts challenging. Conclusions in previous research are inconsistent. The literature reflects these inconsistencies as there is, as far as we know, not one approach adopted as the gold standard. Opting for a pure human-based evaluation does not consider the possibility that a system can behave creatively, but where the output does not resemble that of human-made art [19]. Bias against computer-made art could dismiss this as non-creative [6, 27]. Therefore, rather than choosing one approach over the other, we prefer to not only use a human evaluation but also use an objective approach in combination.

### 2.4 IoT

An IoT environment is a system of interrelated computing devices that can transfer data to a network without human interference. The computing devices are generally better known as smart devices. Smart devices can act on anything from temperature, movement, energy use, and the number of washing cycles. An example of an action by a group of smart devices would be turning on the lights and thermostat when one enters their home. Some smart devices can also suggest actions to the user based on the data it collects. New data can change the IoTs actions and help it learn more about the environment.

All possible actions, relations, functions and types of measurements are captured in an ontology. Without an ontology, the smart devices cannot communicate with each other. An example of an ontology in the smart device domain is the Smart Applications REFerence (SAREF). Figure 1 shows a fragment of a SAREF structure for a device. We can gather from the figure that saref:Device has a property and a measurement. saref:Property links to saref:FeatureOfInterest that relates to real-world environmental factors that the device can measure.

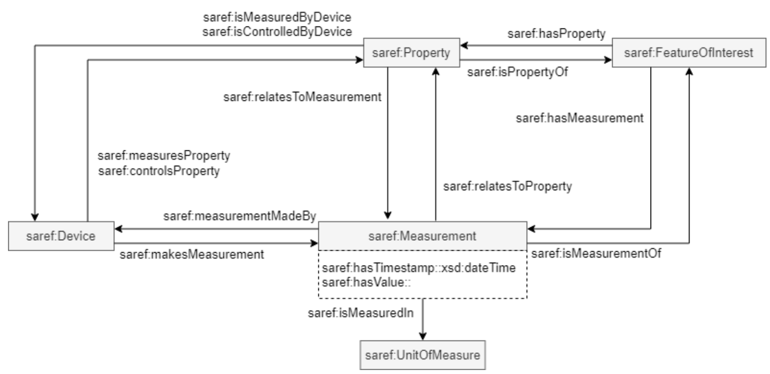


**Fig. 1.** Example of a SAREF ontology of a smart device.

The important note on IoT is that it not only consists of a group of sensors but is a community of devices interacting with each other and sharing their experiences. Therefore, IoT is a suitable digital substitute for human-like experience in computer creativity.

## 3 Review of the original GAN

This section introduces the original GAN and the inner processes, followed by a review of its creative abilities regarding the literature review in section 2. The conclusions from this review will be the basis of the requirements for the proposed design of the creative GAN.

### 3.1 Design

We briefly introduced the GAN in the introduction. We used the counterfeiter and the police analogy to demonstrate how the GAN works. The counterfeiter and the police are formally called the *generator* and the *discriminator*, which are the two main components of a GAN. Figure 2 illustrates the working pipeline of a simple GAN.

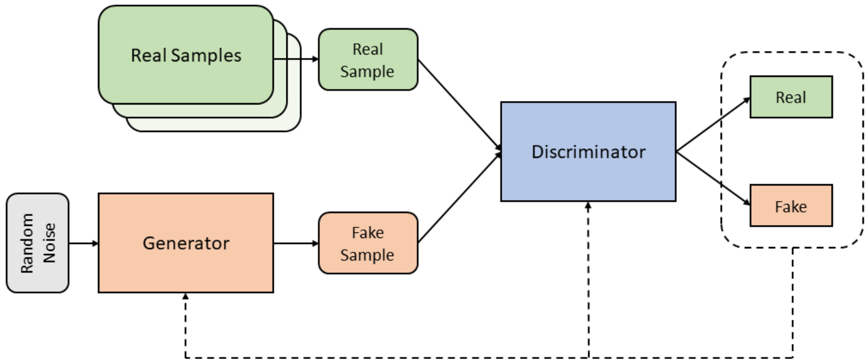


**Fig. 2.** Pipeline of the original GAN.

The pipeline illustrates that there are two inputs: *fake samples* and *real samples*. The generator, indicated by orange, generates the fake samples by taking *random noise* input. The random noise ensures that each sample created by the generator results in a new sample. The other input, the real samples input, can be anything from samples of handwritten notes to faces and photographs. The generator generates a batch of fake samples and, together with the real samples, are classified by the discriminator. The discriminator, indicated by green, is a typical classification model that discriminates the fake from the real samples. Both the discriminator and the generator are neural networks that get trained at the same time. The two models are updated based on the results of the discriminator. In turn, the discriminator gets better at classifying the samples, and the generator gets better at making real-passing fake samples. Eventually, the discriminator can no longer separate the fake samples from the real samples indicating that the generator makes good fake samples.

Figure 3 shows an example of GAN input and output. The training uses 25 images from the Color Field Painting style as input derived from the Wiki-art database [32].

The GAN copies several images from the real-data set in the generated dataset. The mostly blue and black paintings are examples of images that we both see in the real and fake set. Other images are relatively more unique in their design. We can ask ourselves if the GAN and its current output are creative. Or, as prior research suggests, is the GAN only good at making more of what is already known? [25, 7, 11]. We examine this question regarding value and novelty, surprise, and autonomy by relating these to the current circumstances.

### 3.2 Value and novelty

We first dissect the GAN in the two categories of value and novelty. Value and novelty are not behavioural aspects of creativity but concern the creative output. Looking again at figure 3, we can expect that there are a few generated samples

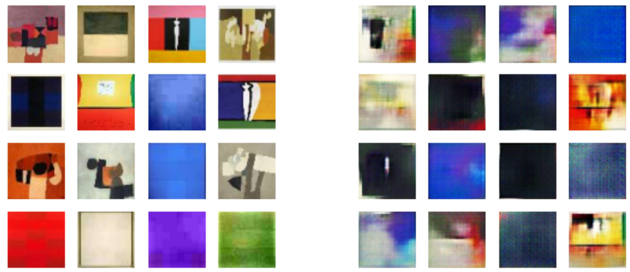

**Fig. 3.** Example of generated art by a GAN. On the left: 16 real Colour Field Painting samples, data from www.wikiart.org. On the right: 16 fake generated art samples.

that can go undetected between the real samples. That the generated images can not be distinguished indicates that the generated samples of the GAN are an excellent fit for the given problem, making more samples that fit within the real samples. Generating a good fit indicates that the GAN can produce high-value samples. The goal of the GAN to find a good fit contradicts the creative ideal of being both novel and valuable. We already mentioned that some images are not direct copies, but the images are in a similar style as the input images. The resemblances indicate that the generated images are of a low novelty. As for value and novelty, we consider the GAN unsuccessful in being a creative computer program. By making the output of the GAN more novel, creativity may be perceived as higher. One of the requirements for the new design will consider the creativity of the output.

### 3.3 Combinational, explorative and transformative creativity

The analyses from section 2.1 on the three types of creativity explained that creativity happens in three different ways. These are combinational, explorative and transformative creativity.

The original GAN primarily uses combinational creativity; each sample the GAN creates is an assemblage of the existing samples. Figure 4 exemplifies how the GAN processes two images into a new image. We can see that the GAN makes combinations but that these combinations are not necessarily compelling. When an interesting combination occurs it is rather based on chance than a thoughtful decision.

The GAN does not use explorative creativity or transformative creativity in the process. Explorative and transformative creativity are both based on a principle of threading along/outside the edges of the rules. They contradict the system's design to follow the rules. The game's rules are to match the real samples as best as possible, which contradicts the ideas of explorative and transformative creativity.

Currently, the system does not behave creative, considering how well the three types of creativity appear in the system. A second design requirement will

**Fig. 4.** GAN trained on 2 real samples. The left two images are the real input samples. The images on the right are the fake generated samples.

be improving in either combinational, explorative or transformative creativity or in combination.

### 3.4 Autonomy

At last, we look at the GAN as an autonomous creative system. The GAN fits well in some of the definitions discussed in section 2.2 on generative art. The only input a human has in this system is in the selection of the real data input. In the generating process, there is not one moment a human has to intervene. However, for autonomy to resemble human creativity, there should be more dynamics in the system than only the initial selection of the input data. Regarding creative autonomy, we also discussed how interaction and perception could play a role. The current GAN is not interactive or perceiving. The system's only experience is the initial experience expressed in the input data. Therefore, we consider adding interaction or perception as a requirement for the new design. Implementing more experience through interaction or perception can make the GAN behave more creative.

### 3.5 Previous research

Although many studies use GANs for creative purposes, the research on its self-creativity remains limited. A recent research paper was released reviewing the use of GAN to generate various forms of art [19]. The paper titled: "GAN computers generate arts-" lists numerous authors that used the GAN to execute their creative intentions. Only one paper actively addressed the relation between creativity and GAN. As far as we know, this is the only paper that adequately addresses the creativity in GAN. There is a need for additional studies that not only use the GAN in a creative context but also examine the GAN as a creative entity.

In the paper by Elgammel et al. [2017], they introduce the CAN: Creative Adversarial Networks. The authors build their CAN system on the theory suggested by Berlyne of "level of arousal", where human excitement on the aesthetics is central [33]. The creative improvement of the CAN is that the system trains on recognising styles and deviating from the style norms to increase the arousal potential. Their results showed that participants regularly could not distinguish their generated work from art generated by artists, and sometimes rated the generated art higher in various areas. The CAN provides an example that is already more successful in generating creative art than the original GAN. The CAN does not generate art based on experiences from the real world. Additionally, they do not address the creative behaviour of the GAN but focus on the creativity of the output. We propose a design where the emphasis is on the creative behaviour in the process instead of the creativity of the output.

So far in this chapter, we presented background information about how computers can become creative agents. In each section, we further elaborated on: what makes a computer appear creative, how is creativity autonomously simulated, and how is generated art evaluated. We evaluated the original GAN as a creative agent in the last section. Based on this related work, we can identify the creative opportunities of the GAN and define requirements. In the next section, we propose a design for a new creative GAN based on the related work.

## 4 The creative GAN

This section looks at the design for a self-generating art system using a GAN with an IoT extension. From the literature review we extract requirements and build a design and explain how the design works.

### 4.1 Requirements

We extract requirements from the literature review and the evaluation of the original GAN. Below is a list of the selected requirements.

(**R**1) Combinational, exploratory and/or transformative creativity.
(**R**2) Interaction and perception.
(**R**3) Creative output.

The first requirement (R1) is based on Bodens three ways that creativity can happen [9]. In the original GAN there is minimal combinational creativity present but this is more accidental than a given. The second requirement (R2) is based on implementation of human-like experiences. Jennings proposed a role for interaction in an autonomous system to simulate the outside influences from the social world that in artists are leading in the creation of art [16]. In combination with this requirement is the idea of perception in a creative system. Without perception, or an idea of the real world, it is not possible to interact and gain experience. The last requirement (R3) refers to how the original GAN generates images that are expected based on the given input data.

### 4.2 Pipeline

The design derived from the requirements is The creative GAN: an IoT art generator. We extend on the GAN with an interactive *IoT environment*. Section 2.4 explained the relevant characteristics of an IoT.

There are no changes made within the GAN. The IoT is an extension of the GAN that controls what real data the discriminator uses. The IoT environment becomes the decider on what real samples the system uses in training and in turn controls the output. Figure 5 shows the re-visioned pipeline.

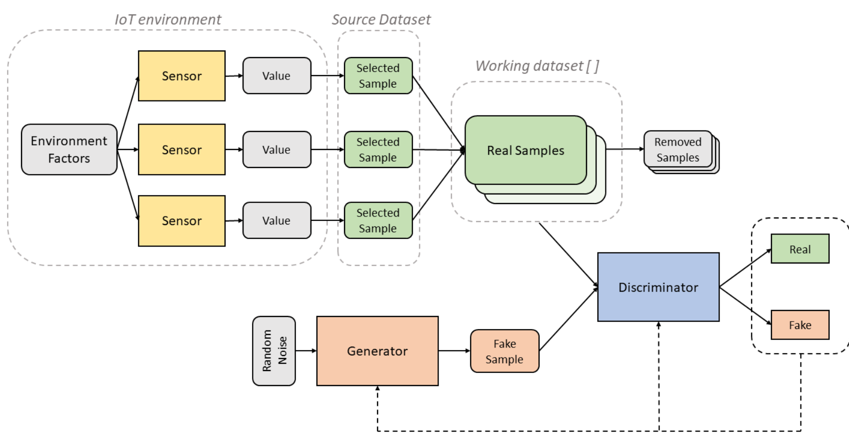


**Fig. 5.** Pipeline of the creative GAN. The GAN is extended with an IoT environment capturing live data and using this data to select samples for a working dataset.

**Data processing** The IoT gathers measurements from multiple sensors embedded in smart devices that are connected to the IoT. With the measurements, the system can select real data samples from the *source dataset*. The source dataset consists of numerous smaller sub-datasets with each its own theme or genre. Each sensor represents a different sub-dataset, and every raw value received from a sensor represents a specific file in that data set. For example; a thermostat contains a humidity sensor, the humidity sensor is linked to sub-dataset, the sensor measures a humidity of 39,0%. The specific file that is selected would then be 390.jpg in sub-folder humidity in folder thermostat. Depending on the size of the IoT system, the amount of data linked to the sensors can grow extensively as each sensor links to its own set. All the data linked to the sensors define the system's source dataset.

In order to use the IoT environment in an interactive setting, we need to review how the real data samples train in the original GAN. The original GAN continuously trains on one extensive real sample dataset that does not change during training. For the interaction to take effect in the output, we need to be able to change the real sample set continuously. The pseudo-code in Algorithm 1 illustrates how this would work in the new system.

**Algorithm 1** Pseudo code demonstrating the iterative loop and the epoch loop.

```
numberOfIterations
numberOfEpochs
sizeWorkingDataset = sizeLimit

for i in range(numberOfIterations) do                      ▷ start of iterative loop
    for sensor in environment do
        value ← sensorMeasurement
        img ← ./sourceDataset/style ∗ /value.jpg
        workingDataset.add(img)
        while sizeWorkingDataset ≥ sizeLimit do
            img.delete(./workingDataset/oldestFile)
        end while
    end for

    for j in range(numberOfEpochs) do                      ▷ start of epoch loop
        GAN ← train images workingDataset
        j++
    end for                                                ▷ end of epoch loop
    save generated images
    i++
end for                                                    ▷ end of iterative loop
```

We set up a *working dataset []* that contains the real samples selected through the measurements. A predetermined number of samples are inside the working dataset. The sensors measure new values and select new samples that get added to the working dataset. The program removes the oldest samples from the set to bring the number of samples back to the predetermined amount. The removed samples are not instantly lost because the system trains using a neural network. Traces of each sample are captured in the neural network's weights and do not suddenly disappear when the program removes the sample. The knowledge on the samples gradually becomes less up-front. After the program adds new samples, the network's weights can adjust to become a better fit for the updated working dataset. This way, the system explores new data samples and also "loses interest" in older data samples.

**Iterations and epochs** We call a pass through the working dataset an iteration. An iteration consists of selecting the new samples, removing the oldest samples, training on the updated working dataset, and saving the results. The training on the updated working dataset is referred to as an epoch. An epoch consists of generating fake samples, discriminating fake/real samples, and sending the feedback to both the generator and the discriminator. The system's number of iterations is endless, but the amount of epochs is predetermined. The amount depends on the size of the working dataset, which depends on the size of the IoT environment.

The new design can interact with its environment through the smart devices connected to the IoT. The system trains continuously on new datasets, enabling the system to generate diverse art under different real-world circumstances.

## 5 Prototype

The creative GAN is an interactive IoT-based art generator that explores art through perceiving the environment. Based on the design introduced in the previous section, we develop a prototype of the creative GAN. First, we present hardware, software and data, followed by fine-tuning and results.

All training occurs on the same system, with a 1660ti GPU and a single 16GB RAM. We use python 3.8 and TensorFlow 2.8 and run the program in the anaconda terminal in a virtual environment. Running the build directly in an IDE can cause problems regarding available memory, so Anaconda or similar is advised. The code is available at github.com/RFLBeening/creativeGAN.

### 5.1 Hardware

We set up the new GAN system starting with the hardware. The hardware consists of an Arduino UNO that represents an IoT. We use three sensors connected to the Arduino to act as the IoT environment. Appendix A shows the schematic of the hardware of the Arduino system. A detailed version of the creative GAN pipeline showing the specifics for this prototype is shown in Figure 6. This prototype uses three sensors, similar to the pipeline in section 4.2 (see Figure 5). These sensors correspond to two smart devices. The first sensor is a light dependent resistor (LDR); these sensors are in systems that, for instance, regulate the opening and closing of curtains. The second and third sensors are for temperature and humidity (DHT22). Both of these sensors are in a thermostat.

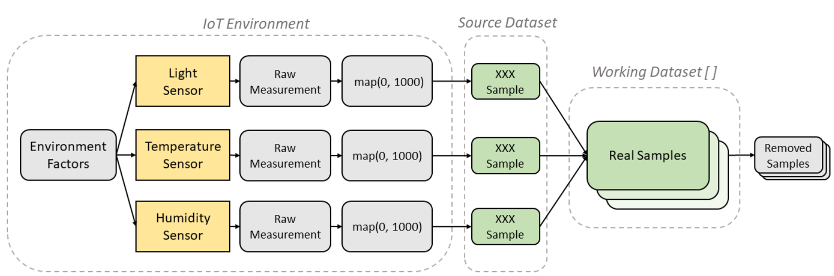


**Fig. 6.** Pipeline specific to the prototype.

### 5.2 Software

The original GAN has improved since its first introduction resulting in the WGAN-GP. The WGAN-GP is a GAN that uses the Wasserstein distance and gradient penalty instead of the original GAN classifier and gradient clipping. These improvements solve major problems the original GAN is prone to, for instance, non-convergence, mode collapse and instability. It seems appropriate to use the revised GAN that avoids these problems and makes training less sensitive to parameter changes.

The Wasserstein GAN with Gradient Penalty (WGAN-GP) from PyImageSearch University is the basis of the creative GAN[34]. We upgraded the image resolution from 64x64x3 to 128x128x3. By adding an extra convolutional layer to the generator and changing the input size of the discriminator, the system can process larger images. Replacing the WGAN-GP generator network from *128x1* $\rightarrow$ *8x8x512* $\rightarrow$ *16x16x256* $\rightarrow$ *32x32x128* $\rightarrow$ *64x64x64* $\rightarrow$ *64x64x3* to *128x1* $\rightarrow$ *4x4x1024* $\rightarrow$ *8x8x1024* $\rightarrow$ *16x16x512* $\rightarrow$ *32x32x256* $\rightarrow$ *64x64x128* $\rightarrow$ *128x128x64* $\rightarrow$ *128x128x3* and the critic input from *64x64x3* to *128x128x3*. Most of the other changes are additions to the code. We remove the real sample input and replace it with an image renew class. The class image renew handles the Arduino data and the selection and removal of the images in the working dataset. All new functions come together within the iterations. The pseudocode from section 4.2 illustrates how the iteration operates in the system (see Algorithm 1. An overview of the additions and changes in the code are in Table 1.

**Table 1.** WGAN-GP to creative GAN Summary of Changes.

| **Where** | **What** | **Original** | **Changes** |
|---|---|---|---|
| Main | Data input | Big data | Working dataset |
| Main | Added iterations | N/A | Iterations |
| Image renew | Added image renew class | N/A | Image renew |
| Generator | Convolutional layers | 4 layers | 5 layers |
| Discriminator | Input discriminator model | 64x64x3 | 128x128x3 |
| GANmonitor | Number of output images | 16 images | 9 images |

### 5.3 Data

To test the system, we trained the model using paintings from the publicly available WikiArt database [32]. This database contains 81.444 works of art in 27 different art styles. Before the system can process the images from the WikiArt dataset, they must be pre-processed. The pre-processing entails shaping the images into a 128x128 square and changing all names to numbered file names. For example, aaron-siskind_acolman-1-1955.jpg(1280x920) becomes 1.jpg(128x128).

**From measurement to image** The values of the sensors link to images of a specific art style from the WikiArt dataset. For the light sensors, this is Cubism; for the temperature sensor, Fauvism; and for the humidity sensor, Abstract Expressionism. For now, we choose three relatively low-detailed art styles better suited for learning in low resolution. Figure 7 shows samples of the three different art styles that the system uses for learning.

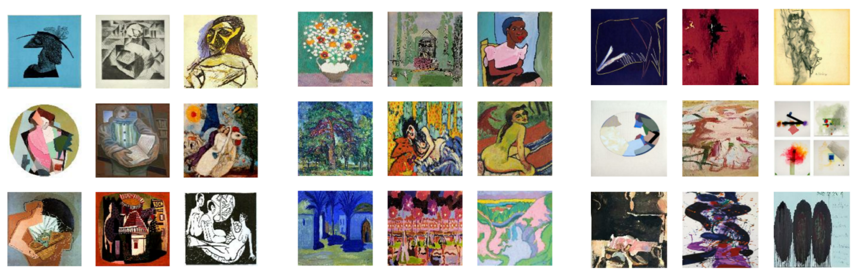

**Fig. 7.** Nine art samples randomly selected from the three linked art styles. From left to right: Cubism, Fauvism, Abstract Expressionism.

**Table 2.** Characteristics of the sensors in the prototype.

| Smart Device | Sensor | Sensor Range | Precision | Art style |
|---|---|---|---|---|
| Curtain | Light | (0, 1023) | 1 | Cubism |
| Thermostat | Temperature | (-40.0°C, 80.0°C) | 0.5°C | Fauvism |
| Thermostat | Humidity | (0%, 100%) | 1% | AE |

All sensors have diverse measurement ranges and sensitivity. As we can see in Table 2, the sensor ranges and precision are dis-proportioned between each sensor. These diversities lead to imbalances and errors in the working dataset. Errors occur when a measurement is a negative number, and imbalances occur when the change rate of specific sensors is higher or lower. To avoid the errors we map each sensor value to the same range of (0, 1000). A range of (0, 1000) means that one sensor can measure 1000 values that correspond to 1000 images, in practice this is not always true. Depending on the outer limits of the sensor and the precision it can not measure all 1000 values. The chances are lower that the environment reaches -40°C or 80°C than that it is entirely dark or light. The other major influence, the precision, causes another limitation to the number of active images. For example the humidity sensor has a range of (0%, 100%) that maps with precision of 1% to (0, 1000), this creates a situation where every 10th image is selected for training. These imbalances are not harmful in a system with many sensors or high-precision sensors; this will only reflect particular environmental activity more. However, in this prototype, we need to keep in mind that different precision of the sensors can cause imbalances and that certain art styles can become dominant or non-dominant in training.

**Working dataset** Finding the appropriate size of the working dataset is a complicated task. Many variables affect the number of images the working dataset should contain. The variables are the number and precision of the sensors, the system's specification, and the fluctuation in the environment. The prototype has three low-precision sensors, takes 3 seconds per epoch and has, on average, an environmental change every 3 minutes. For the prototype of this system, a working dataset of 20 images generated satisfactory results. Under different conditions, the size of the working dataset needs to be reconsidered.

### 5.4 Iterations and epochs

As we explained in section 4.2, the system consists of iterations and epochs. One iteration can contain multiple epochs. The amount of epochs eventually defines the length of an iteration and, in turn, the length between working dataset updates. With a working dataset set of 20 images, one epoch takes approximately 3 seconds, which means that every 3 seconds, the system checks the sensor values and saves the results. Ideally, we want to detect all environmental changes and set the number of epochs to 1. However, we need to consider the system specifications and saving the results of each iteration would cost a lot of internal storage. Additionally, because the prototype only contains three sensors, we do not need to measure every 3 seconds because the chance any of the measurements change is low. We decided upon 10 epochs per iteration, meaning every 30 seconds, there is an update on the working dataset and an output of the generated results. This system does not have an endpoint, meaning we do not have to define a training limit for the iterations. In theory, the system can run endless and keep generating new images. For testing, we limited the number of iterations to 4000 and epoch to 10 which takes 22-26 hours.

## 6 Results

In the previous sections, we demonstrated how we built the design in a working prototype. Now, we can look at the generated images. Figure 8 shows a sample of the generated images.

The model uses the images in the working dataset to make unique images in recognisable designs. We can see that the images produced are abstract, though if we were to assign a style to the images, it would be a combination of cubism and graffiti. That cubism is a dominant trait in the images is not surprising. The light sensor linked to the cubism has the highest activity of the three sensors. At times of many light changes, 18 out of the 20 images in the working dataset were cubism samples.

**Fig. 8.** Selection of generated art by the creative GAN.

We perceived many generated images outside the presented selection to be of high entropy and vague in their design. That this happens is not a unique problem to this system. GANs generally tend to produce more surreal-looking images. The flower paintings produced by GANGogh are very similar in nature to those produced by the creative GAN [35].

Another notable output of the system occurs in periods of low activity. The generated images start to look like a collage of existing images when the change rate of images is low, see Figure 9. Because there are only three sensors in the prototype, there may be extended periods that there is no or minimal change in the working dataset. This results in the over-training on the images currently in the data set and only stops when new images are added. If the number of sensors is higher than three, the change rate of the working dataset would be higher, and over-training is less likely.

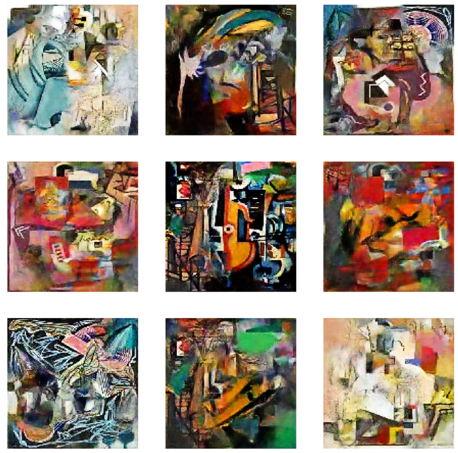

**Fig. 9.** Example of the generated images looking like a collage of the real data samples.

### 6.1 Interaction

In order to know if the prototype does as intended, we need to look better at the generated images. We want to look for changes in the generated images when changes occur in the environment. Visible changes would indicate that the system successfully incorporates new inspiration dictated by the environment and is successful regarding the second requirement. A moment in training is selected where the interaction occurred to investigate the process. We look at images surrounding sunset and sunrise, where changes in light interacted with the system. Figure 10 shows a range of 5 generated sets each 15 minutes apart for sunset. Figure 11 shows a range of 5 generated sets each 15 minutes apart for sunrise.

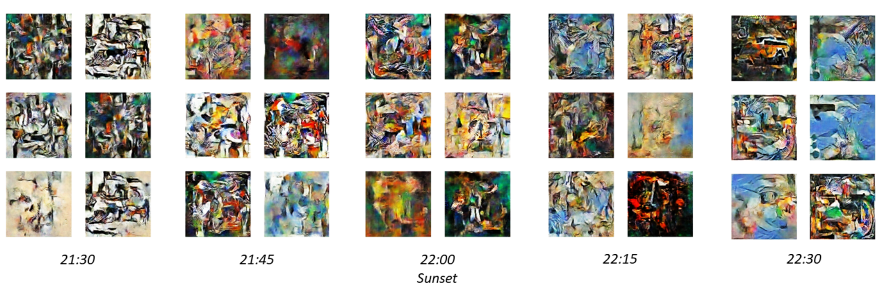


**Fig. 10.** Intermediate samples from the generator, before and after sunset. Iteration 2460: 21:30. Iteration 2510: 21:45. Iteration 2560: 22:00. Iteration 2610: 22:15. Iteration 2660: 20:30.


**Fig. 11.** Intermediate samples from the generator, before and after sunrise. Iteration 3860: 04:50. Iteration 3910: 05:05. Iteration 3960: 05:20. Iteration 4010: 05:35. Iteration 4060: 05:50.

The fundamental colour changes show that the system integrated the changes of images in the generated images. The sunset images start with bold lines and specs of colour; from left to right, this changes to more soft edges and prominent blue areas. In the images before sunrise, there are more cool tones at the start, which later turn to earthy colours.

The resulting images from the prototype give a good indication of the system's possibilities. The results show that the system can interact with the environment and produce diverse images according to the changes. The first step in further development would be extending the IoT system with more sensors to avoid a static working dataset.

## 7 Evaluation

We evaluate the creative GAN as an autonomous artist that generates creative art through IoT. In section 4, we set up the requirements for this system. These are, the system behaves creative according to Bodens' 3 ways of creativity (R1), the system uses interaction or perception to behave more creative (R2), and the system generates creative output (R3). As we suggested in section 2.3, we evaluate computational creativity by looking at the process and the generated output. First, we evaluate the process by reviewing the design regarding requirements R1 and R2. After that, we focus on the third requirement, where we evaluate the generated art of the GAN. This evaluation has a human-based approach where we conduct task-based interviews to verify the third requirement.

### 7.1 Requirement 1: combinational, explorative and/or transformative creativity

In this section, we review the process of the creative GAN regarding requirement R1. We ask ourselves, can the system behave creative according to Boden's suggested creativity types?

As discussed in section 3.1, the original GAN can make combinations but solely based on chance. In the creative GAN, the combinations are restricted to what samples are in the working data set. The current prototype is too incomplete to correctly investigate if this leads to more or maybe less attractive combinations. We did not directly address combinational creativity in the current implementation. With further development, the system can utilise the relations between devices in the IoT environment. A grouping of styles per room or device group can make related combinations possible. Still, many opportunities remain for combinational creativity using an IoT.

The second notion is explorative creativity. In the design of the creative GAN (see section 4.2), we explain how the system uses exploration to select samples for the working dataset. An extensive source dataset connects to the system, and the system explores different samples in the dataset based on the environment. Figures 10 and 11 from the results (see section 6) visualise how this appears in the output of the GAN. In both images we see a gradual change in terms of lines, shapes and colours.

The last notion is transformative creativity. Transformative creativity is described as "searching beyond the given space" [9]. The given space in an original GAN is the real data samples, as opposed to the creative GAN that also adds the environment as space. It depends on the interpretation of transformative creativity if the new creative GAN is considered transformative. If the interpretation relies on the system to be able to bend the rules, the system is not considered transformative. However, the system can be considered transformative if it is interpreted as searching beyond the GAN's bounds. Connecting the GAN to the real world exceeds the system's original bounds. Based on this assumption, we can call the new system transformative if we believe the second interpretation to be true.

In the creative GAN, several aspects of the three notions of creative surprise are present but certainly not yet fully exploited. Boden's [2004] three notions of creativity are free for any interpretation. This means that not everyone has similar ideas on using the three notions in a creative design. For the current work, it is sufficient to say that the design satisfies requirement R1 on explorative and transformative creativity based on the behaviour seen in the prototype. This assumption must be re-addressed when the prototype is further in its development.

### 7.2 Requirement 2: interaction and perception

This section covers the evaluation of the second requirement. For this, we try to answer the question: is interaction or perception possible in the system? The IoT system has multiple purposes in the proposed system. Firstly it acts as the system's interaction, and secondly, it is the system's perception. As suggested by Boden [2009], we carefully considered the type of interaction in the system. We find that implementing interaction in a way not directly steered by human intentions creates a subtle human interference that lets the system keep its autonomy. An additional advantage of the IoT is the perception of the real world through sensors. How the system uses the sensors to discover inspiration for new designs can be a computational equivalent to how humans use perception to find inspiration. The results from the prototype illustrate how the interaction with the environment changes the output of the creative GAN, suggesting that the system adheres to the description above (see Figures 10 & 11).

### 7.3 Requirement 3: output of the creative GAN

This section aims to address the third requirement. This requirement examines the question: is the output perceived as creative? To answer this question, we conducted a human evaluation using task-based interviews.

The participants in this study are knowledgeable in the art and creative media domain. We asked all participants to provide informed consent regarding data storage, audio recording and the freedom to answer or not answer questions.

We choose a task-based unstructured interview as the evaluation method. Unstructured interviews allow the freedom to ask relevant questions to the person at hand. As we discussed in section 2.3, people tend to have different opinions on art, so not all questions suit everyone. We like to keep the discussion open and let the interviewee free in their answers. These interviews aim to get insight into the new creative GAN as an interactive art generator. The interview proceeds in two stages. The first stage consists of questions about two specific tasks. For the tasks, we ask the participant to think aloud. The first task is to tell which and why images could be generative art. The second task is to tell which image they think is most or least creative. The second stage of the interview is to get to know the interviewee's stance on creativity and computer creativity; this can help explain their motives for specific choices. The interview format is found in Appendix B.

The first part of the interview consists of two tasks. The two tasks consist of:

- Identifying the generated art images.
- Describe the images based on creativity and how they compare to each other on the level of creativity.

For the two tasks, we use images from human-made art, online GAN-generated art, and the creative GAN-generated art. We collected the images from 4 different sources. Each image needs pre-processing according to the retrieved format. For the human-made art, we selected images from the WikiArt dataset in the genre Abstract [32]. For the other GAN-generated art, we used the online tool NightCafe and StaryAI to generate art with the prompt 'abstract'. We use the output from the results in section 6 for the creative GAN images. We acknowledge that there is selection bias in the samples by hand-selecting from both the creative GAN and the abstract WikiArt.

All creative GAN-generated images have a resolution of 128x128. We up-scaled the generated images from the 128x128 to a 512x512 resolution. For the images from the art database and the online GAN, we first sized down the images to fit the 128x128x resolution and up-scaled them to 512x512. This way, the participants will not be biased due to differences in qualities.

**Task on generative art identification** The first task concerns generative art identification. This task introduces the topic to the participants. We asked the participants to elaborate on why they thought an image was/or was not generated by a computer. Figure 12 shows the selection of images that were presented to the participants. The images 1, 3, 5 and 8 are creative GAN samples, and 2, 4, 6, 7 are human-made art samples from the WikiArt database. The data in Table 3 presents an overview of the answers to the question: which images do you think are computer-generated?

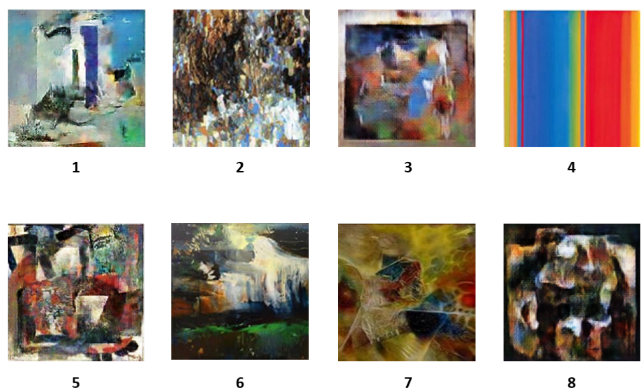


**Fig. 12.** Images shown to the participant during the generated art identifying task. Images 1, 3, 5 and 8 are creative GAN generated images. Images 2, 4, 6, 7 are human-made art.

**Table 3.** Answers to the question regarding if the presented image is computer made. Dark grey indicate a correct answer to the question, a light grey indicates an incorrect answer the question.

| **Participant** | **P1** | **P2** | **P3** | **P4** | **P5** |
|---|---|---|---|---|---|
| Image 1 | No | Yes | Yes | No | No |
| Image 2 | Yes | No | No | No | Yes |
| Image 3 | Yes | Yes | No | No | No |
| Image 4 | Yes | No | No | No | Yes |
| Image 5 | Yes | No | Yes | No | No |
| Image 6 | | Yes | Yes | Yes | No |
| Image 7 | No | Yes | Yes | Yes | Yes |
| Image 8 | Yes | Yes | No | Yes | Yes |

From the table, we can gather there is not one image continuously picked out as the generated art and vice versa as human-made art. We will look at the participant's explanations of their choices to make more sense of the data. We will first examine why they placed specific images in the generated art category. The participants mentioned glitchy, messy shapes and colouring, unfamiliarity, and the feeling that the image had a filter as the main reasons behind their decision. The other way around, the images they thought were human-made were considered unpredictable, emotional, or they could imagine seeing the work of art presented in a museum. Even though there were similarities in their argumentation, they did not choose the same images as the generated art. The similarities suggest they have a central vision about what would make artwork appear generated. Even though they had a similar vision, they did not choose the same images.

**Task on creativity level** In this task, we directly address the third requirement regarding the creativity of the output. We aim to compare the creative GAN between human-made art and other GAN-generated art in terms of creativity. We asked the participants whether they found the images creative and which images were the most and least creative. Figure 13 shows the selection of images that were presented to the participants. For this set of images, we used two creative GAN samples, two human-made art samples from the abstract WikiArt database, and two images generated by online GANs. The data in Table 4 presents an overview of the answers to the question: Do you think this image is creative?


**Fig. 13.** Images shown to the participant during the level of creativity task. Images 2 and 6 are creative GAN generated images. Images 1 and 5 are human-made art. Images 3 and 4 are online GAN generated art.

**Table 4.** Answers to the question: do you find this image creative. Dark grey indicates a positive answer to a creative GAN generated images, a light grey indicates a positive answer to either human-made art, or online GAN generated art.

| Participant | P1 | P2 | P3 | P4 | P5 |
|---|---|---|---|---|---|
| Image 1 | No | Yes | Yes | Yes | No |
| Image 2 | No | Yes | Yes | Yes | No |
| Image 3 | Yes | No | No | Yes | No |
| Image 4 | Yes | No | No | Yes | No |
| Image 5 | No | Yes | No | Yes | No |
| Image 6 | | Yes | Yes | Yes | No |

All participants mentioned originality as their primary criteria for deciding which image they perceived as more creative. Two out of the five participants placed all images in the same category. Participant 4 indicated that all of them were creative, and participant 5 indicated he could not indicate if they were creative or not without any context. The answers to this question became clearer in the second part of the interview. Participants 2 and 4 indicated that four or more images are creative and had a similar explanation of what they perceive as creative: a balance between logic and randomness. This description is very similar to the notions of value and novelty as explained in section 2.1. Participants 1 and 5, who chose two or fewer images to be creative, had a different view on creativity than the one mentioned by participants 2 and 4. Participant 1 looked at the art from an emotional perspective. Whereas Participant 5 was overall cautious with assigning creativity to only a still image. Similar to Colton [2008], they believed we could better judge creativity with more context on the process and knowledge of the skills involved. The answers to the question regarding creativity cannot tell us if the generated images are creative. However, it gives a good indication of how an image can be perceived as creative.

The mixed answers to the open questions in the second stage of the interview implied that it is probably not possible to make an image or art that everyone agrees on as being creative. The fact that several participants perceived the creative GANs images as creative indicates that the generated image is creative inside their view of creativity. As indicated in section 2.3, a person's view on creativity is bound by social and cultural prejudices, which are affected by various environmental factors. Therefore, we cannot make a definitive conclusion on the creativity of the output of the creative GAN. Only whether it is creative from that particular persons perspective.

## 8 Conclusion and Discussion

This paper further explores computational creativity, more specifically creativity in GANs. The original GAN does not behave creative when compared to computational creativity principles. Therefore, we proposed a new system to make the GAN behave more creative. Based on how humans use experience and interactions from the real-world in their creation of art, we proposed to extend the GAN with a matching system. We suggested an extension of the GAN with an IoT environment to simulate the human way of experiencing and interacting. We demonstrated a realisation of the system where sensors get to decide what styles the GAN learns.

The generated images showed that the sensors successfully influence the generative process. We further evaluated the new design by relating it to the computational creativity and creative autonomy literature. The evaluation showed that the system behaves according to Boden's [2004] ideas of exploration and transformation, and can perceive the environment and interact accordingly. Additionally, we conducted a human evaluation to gain insight into the creativity of the output. The human evaluation gave mixed answers on the creativity of the output. Some participants perceived the creative GAN art as more creative than the human-made and online GAN art, where others did not.

We leave open how to interpret the human participant's answers to the evaluation. Art and creativity are socially and culturally heavily influenced, causing bias in their judgement. Because of the nature of creativity, it is unlikely that everyone agrees on whether an artefact is creative or not. A different set of participants or a different set of images could have led to different answers. However, the answers provide an idea of what people value in creativity and gives us first hand experience in how this influences their outlook.

The next step for this project can be a real integration with an IoT system. Systems like the SAREF ontology are needed to interpret data from separate smart devices and test the proposed system more elaborately. The available data from SAREF was limited, and implementation of the system was beyond the scope of this research. Nevertheless, the current prototype already showed promising results when using sensors found in IoT devices. When the system is connected to a real IoT, different interactions can be explored. These interactions can be based on different sensor types or the existing relations between IoT devices.

## 9 Acknowledgements

This project would not have been possible without dr. Victor de Boer, who supervised this project and provided beneficial feedback throughout the process. Special thanks should go to Joost de Boo for serving as a second supervisor and providing me with great opportunities and creative inspiration. Thanks should also go to the participants of the interview. Not only did they take their time for the interview, but they also provided me with extra information and sources.

## References

[1] Margaret A Boden and Ernest A Edmonds. 2009. What is generative art? *Digital Creativity*, 20, 1-2, 21–46.

[2] Graeme Ritchie. 2007. Some empirical criteria for attributing creativity to a computer program. *Minds and Machines*, 17, 1, 67–99.
[3] Nick Collins. 2008. The analysis of generative music programs. *Organised sound*, 13, 3, 237–248.
[4] Mao Li, Jiancheng Lv, Jian Wang, and Yongsheng Sang. 2020. An abstract painting generation method based on deep generative model. *Neural Processing Letters*, 52, 2, 949–960.
[5] Gary R Greenfield. 2006. Art by computer program== programmer creativity. *Digital Creativity*, 17, 01, 25–35.
[6] Simon Colton. 2008. Creativity versus the perception of creativity in computational systems. In *AAAI spring symposium: creative intelligent systems*. Volume 8. Palo Alto, CA, 7.
[7] Marian Mazzone and Ahmed Elgammal. 2019. Art, creativity, and the potential of artificial intelligence. In *Arts* number 1. Volume 8. Multidisciplinary Digital Publishing Institute, 26.
[8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
[9] Margaret A Boden. 2004. *The creative mind: Myths and mechanisms*. Routledge.
[10] Geraint A Wiggins. 2006. Searching for computational creativity. *New Generation Computing*, 24, 3, 209–222.
[11] Sebastian Berns and Simon Colton. 2020. Bridging generative deep learning and computational creativity. In *ICCC*, 406–409.
[12] Margaret A Boden. 1998. Creativity and artificial intelligence. *Artificial intelligence*, 103, 1-2, 347–356.
[13] Simon Colton, Geraint A Wiggins, et al. 2012. Computational creativity: the final frontier? In *Ecai*. Volume 12. Montpelier, 21–26.
[14] Georgios N Yannakakis and Antonios Liapis. 2016. Searching for surprise. In ICCC.
[15] Michel Mendes France, Alain Hénaut, and Jacques Mandelbrojt. 1994. Art, therefore entropy. *Leonardo*, 27, 3, 219–221.
[16] Kyle E Jennings. 2010. Developing creativity: artificial barriers in artificial intelligence. *Minds and Machines*, 20, 4, 489–501.
[17] Michael Muller. 2014. Curiosity, creativity, and surprise as analytic tools: grounded theory method. In *Ways of Knowing in HCI*. Springer, 25–48.
[18] Graeme Ritchie. 2006. The transformational creativity hypothesis. *New Generation Computing*, 24, 3, 241–266.
[19] Sakib Shahriar. 2022. Gan computers generate arts? a survey on visual arts, music, and literary text generation using generative adversarial network. *Displays*, 102237.
[20] Lav R Varshney, Florian Pinel, Kush R Varshney, Debarun Bhattacharjya, Angela Schörgendorfer, and Yi-Min Chee. 2013. A big data approach to computational creativity. *arXiv preprint arXiv:1311.1213*.

[21] Derrall Heath and Dan Ventura. 2016. Before a computer can draw, it must first learn to see. In *Proceedings of the 7th international conference on computational creativity*. Citeseer, 172–179.

[22] Overleg Beeldende Kunsten. 2014. *Opleidingsprofielen beeldende kunst en vormgeving*. Volume 1. Mondriaan fonds.

[23] Edward Necka. 2011. Perception and creativity.

[24] Taras Kowaliw, Alan Dorin, and Jon McCormack. 2009. An empirical exploration of a definition of creative novelty for generative art. In *Australian Conference on Artificial Life*. Springer, 1–10.

[25] Ahmed Elgammal, Bingchen Liu, Mohamed Elhoseiny, and Marian Mazzone. 2017. Can: creative adversarial networks, generating" art" by learning about styles and deviating from style norms. *preprint arXiv:1706.07068*.

[26] Anna Jordanous. 2012. A standardised procedure for evaluating creative systems: computational creativity evaluation based on what it is to be creative. *Cognitive Computation*, 4, 3, 246–279.

[27] Jon McCormack and Mark d'Inverno. 2012. Computers and creativity: the road ahead. In *Computers and creativity*. Springer, 421–424.

[28] Geraint A Wiggins and Joydeep Bhattacharya. 2014. Mind the gap: an attempt to bridge computational and neuroscientific approaches to study creativity. *Frontiers in human neuroscience*, 8, 540.

[29] Alison Pease and Simon Colton. 2011. On impact and evaluation in computational creativity: a discussion of the turing test and an alternative proposal. In *Proceedings of the AISB symposium on AI and Philosophy*. Volume 39. Citeseer.

[30] Simon Colton, John William Charnley, and Alison Pease. 2011. Computational creativity theory: the face and idea descriptive models. In *ICCC*. Mexico City, 90–95.

[31] Anna Katerina Jordanous. 2012. *Evaluating computational creativity: a standardised procedure for evaluating creative systems and its application*. University of Kent (United Kingdom).

[32] [n. d.] Visual art encyclopedia. (). https://www.wikiart.org/.

[33] Howard Gardner. 1974. Aesthetics and psychobiology. (1974).

[34] 2022. Anime faces with wgan and wgan-gp. (2022). https://pyimagesearch.com/2022/02/07/anime-faces-with-wgan-and-wgan-gp/.

[35] Kenny Jones. 2017. Gangogh: creating art with gans. (2017). https://towardsdatascience.com/gangogh-creating-art-with-gans-8d087d8f74a1.

## Appendix A: Arduino schematic

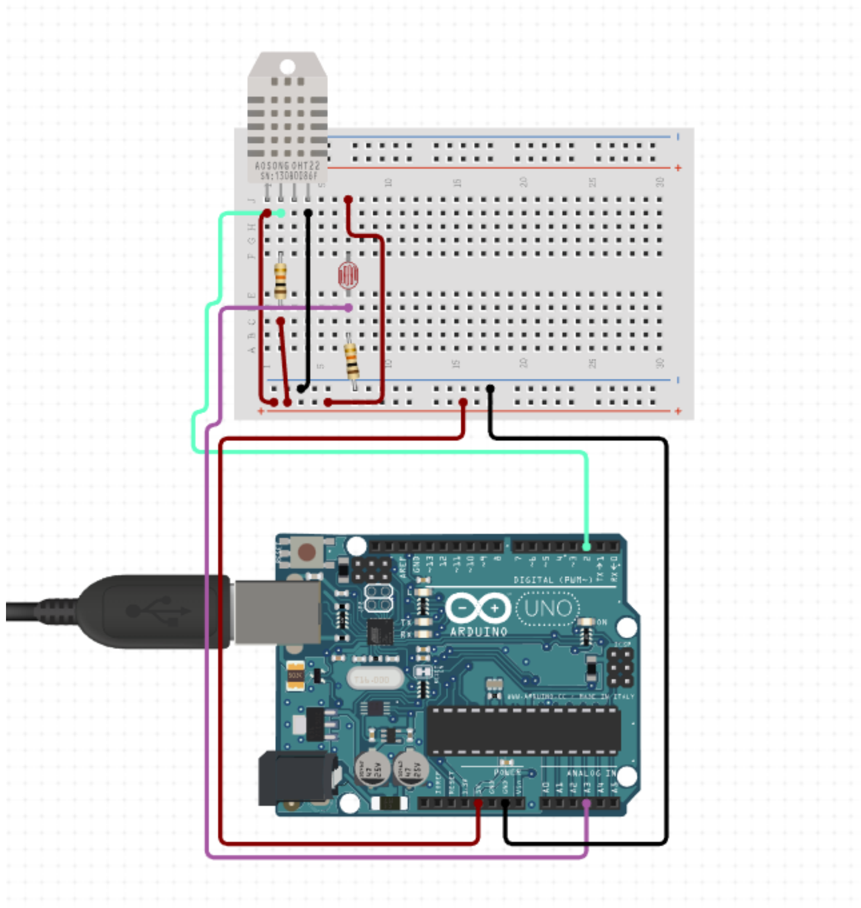

**Fig. A1.** Arduino schematic for creative GAN prototype

Appendix B: Interview format

# 1 Introduction

Welcome the participant.

Ask if they have read the consent form. If yes, then we proceed. If no, ask the participant to read it, and if there are no problems, we proceed.

Inform the participant on the structure of the interview: there are two stages. In the first stage we are going to show sets of images and ask questions surrounding the images. In the second stage we will be asking some open questions.

If participant is ready to begin, we proceed.

# 2 Tasks

The thinking aloud method: Ask participants to think aloud (say everything they are thinking) while performing a task.

## 2.1 Task 1

Show the participant the images. Ask them which image they think could be made by a computer.

## 2.2 Task 2

Show the participant the images. Ask them to indicate which images they perceive as creative. If the participant has difficulty doing this, ask them to order them based from most to least creative.

# 3 Open Questions

## 3.1 Question 1

Ask the participant what creativity is to them.

## 3.2 Question 2

Ask the participant if they believe computers can be creative.

# 4 End of the Interview

Ask if the participant has any questions regarding the interview or the project. Thank the participant for their time.