**Kevin McGuire**

*Paper:* **A Turing Test for Generative Art**

***Abstract:***
When we produce generative algorithms whose output closely matches the style of a particular architect [1] or artist, what exactly is it that we think we are doing? Do we believe we are simply mimicking the architect, producing shallow results that have the surface qualities of the architect but no substance, much as a parrot might speak words [2] but there is no meaning? Is it "mannered"? Or do we believe we are somehow capturing something deeper, a kind of partial working copy of the architect's brain and their creative essence?

How can we even begin to have this discussion?

In the area of Artificial Intelligence, there is the concept of the Turing Test [3] which allows us to discuss when we believe a computer program is behaving "intelligently". The Turing Test is not a real test, but rather a rhetorical framework in which one can disentangle the discussion around artificial intelligence.

We can apply a similar philosophical framework to the problem of generative art. By applying an equivalent to the Turing Test, we can discuss topics like mimicry, forgery, functional correctness, topological similarity, branding, signature style, whether a machine can be creative, and whether a machine can create art.

***Topic: Philosophy of Mind***

***Author:***
***Kevin McGuire***
Independent
Canada

***References:***
[1] G.Matteo Codignola, "Generative Recognizable Codes from Gaudi'", Generative Art International Conference, Milan, 2000
[2] Joseph Weizenbaum, "ELIZA — A Computer Program For the Study of Natural Language Communication Between Man And Machine", Communications of the ACM 9 (1), USA, 1966
[3] Alan Turing, "Computing Machinery and Intelligence", Mind, Oxford University Press, England, 1950

***Contact:***
***kevin@mcguire.name***

***Keywords:***
Artificial intelligence, imitation, creativity

# A Turing Test for Generative Art

**Kevin McGuire**
*Independent, Canada*
*www.mcgui.com*
*e-mail: kevin@mcguire.name*

## Abstract

When we produce generative algorithms whose output closely matches the style of a particular architect or artist, what exactly is it that we think we are doing? Do we believe we are simply mimicking the architect, producing shallow results that have the surface qualities of the architect but no substance? Or do we believe we are somehow capturing something deeper, a kind of partial working copy of the architect's brain and their creative essence?

How can we even begin to have this discussion?

The Turing Test is a rhetorical framework helpful in disentangling the discussion around artificial intelligence. By applying an equivalent to the Turing Test to the questions of Generative Art, we can discuss topics like mimicry, topological similarity, branding, signature style, whether a machine can be creative, and whether a machine and create art.

## 1. Introduction

When we speak of generative art in terms of genetic codes [1], it is sometimes not clear how literally we should interpret the notion of genetic content. We understand that by using a generative approach we produce a family that conforms to the rules of the algorithm, a family whose form has origin in the genetic material. But what do we believe we've actually captured in that algorithm?

Codignola [2] demonstrated a system that could produce results that remarkably resembled the work of Antoni Gaudi. What exactly do we believe he did? Was this just a kind of shallow mimicry [3], a surface resemblance to some of the obvious qualities of Gaudi? Or might we consider him to have captured something deeper, something that by so closely mimicking the subject, could be thought to be equivalent to it?

## 2. The Turing Test

This is a difficult subject to unravel because of all the elements involved: the metaphor of genetics, encoding of form versus encoding of origin of form, perception, objective notions of morphological similarity, subjective notions of

aesthetics, the hand of the human who created the generative system, and finally, the role of the artist versus the role of the audience.

These problems are similar in many ways to the problem of discussing whether machines can think. In order to discuss whether a machine can think, one must first define what it means to think; to discuss what it means for a machine to be intelligent, one must define what it means to be intelligent. But these subjects are just as intractable.

In 1950, Alan Turing proposed a thought experiment which he referred to as “the imitation game”, now referred to as the “Turing Test”. There are three elements: the machine, a human subject, and an experimenter. The experimenter is physically separated from the machine and the human subject by means of a keyboard and screen, so that his only interaction is via a set of typed questions and answers. The machine is constructed in such a way as to try to mimic the behaviour of a human. The experimenter does not know whether they are interacting with the machine or the human. The goal is to see whether the experimenter can tell the difference between the two, detect the fact that they’re interacting with the machine. The machine “wins” the imitation game if the experimenter cannot tell the difference.

In so doing, Turing replaces the intractable question, “What is intelligence” by a more tractable one, “Are there imaginable digital computers which would do well in the imitation game?” [4]. Essentially, it merely claims that one can consider the machine to be intelligent if it behaves in such a way as to be indistinguishable from an intelligent being, if it convincingly mimics a human.

## 3. The Turing Test applied to Generative Art

A direct application of the Turing Test to Generative Art would be to replace “intelligence” with “art”. Thus instead of asking “Can machines think?” we’d ask, “Can machines create art?” But this line of inquiry won’t get us far. While as we believe we can describe intelligent behaviour in terms of it being indistinguishable from that of a human, and in so doing, define intelligence as “that which humans exhibit”, no such definition for art is at the ready, except for the trivial and circular definition “it is the thing that artists create”. That is, we cannot talk about when a machine has sufficiently mimicked the creation of art because we can’t even answer the question for humans.

Instead we’re going to narrow the imitation game to creative output and make the question more tractable. It becomes:

> “Can machines produce artefacts which are indistinguishable from the creative output of a human?”

Indeed, Codignola said with respect to his work on generating form according to Gaudi,

> *The next step was to verify the exact working of the new generative codes by means of 3d scenaries, that are recognizable as "Antoni Gaudì specie's architecture"*

Where “recognizable” is the key for our purpose. To make the discussion more concrete, we’ll pick some specific outputs and humans.

## 4. Spot the BMW

For this thought experiment, let us suppose that a very clever industrial designer has programmed a generative algorithm to produce different styles of cars. The program consists of a single setting: a dial for different brands of cars. The output is a set of drawings of cars of that brand, minus tell-tale giveaways such as hood ornaments or other emblems.

We set the dial to “BMW”. Do we believe the output can be considered “drawings of BMWs”? Has the program mimicked what a BWM designer does? Who better to ask than the man on the street!

Imagine we take the drawings and mix them in with real studio drawings taken from the design centres of other car manufacturers, drawings for cars not in production. We do this to level the playing field and ensure all the cars are unfamiliar to the subject. We ask passer-by’s which car is the BMW. One might reasonably argue that any inability to spot the BWM may be the fault of the subject (they’re just not good at it), not of the generative machine. Therefore we’ll create a control group which contains real studio drawings from BWM, again of cars not in production. Since the generative machine may happen to produce drawings which match existing cars and it would be trivial for people to spot those, we will similarly remove from the generated set any which are easily identifiable as existing models.

If passer-by’s pick the generated BWM with about the same frequency as the studio BWM from the control group, we can reasonably claim that we’ve mimicked what BMW designers do.

We’ve also captured “What it means for something to look like a BMW”. We claim this because, after all, the concrete purpose of branding is to exactly achieve that reaction: immediate recognition on the part of the public that that artefact is of that brand. Therefore if the general public can’t distinguish BWM designer drawings from generated drawings, we can conclude that the generative system is mimicking BWM designers. This of course is bad news for the designers who work at BMW, since their jobs are no longer required.

### 4.1 Argument: Such rules cannot be written down

This argument says that we could never construct the generative algorithm in the first place, that what it means for something to look like a BWM can’t be written down. I’ve never worked at BMW, or any other car manufacturer for that matter, but I’m going to hazard to guess that during one’s career at BMW one gets better at

designing BWMs. Thus there's some externalized transfer of knowledge from experienced designers to new designers. Did no one think to write it down?

In any event, the same argument might be made for the Turing test, that the rules for intelligent behaviour can't be known. Turing avoids this completely by simply defining the criteria under which we might agree that a machine behaves intelligently without specifying how such a machine might be constructed, although he does conjecture that given advances, such a machine could plausibly be built someday. The same can be argued for the BWM generator. The thought experiment is one of mimicry: if the machine can produce an outward result that is indistinguishable from the human's, we can objectively ascribe to the machine the label of being capable of that behaviour in the same way we ascribed it to the human. In the Turing test it's about intelligence, in our case about design and art.

### 4.2 Argument: Being a good BWM designer is about more than just producing cars that look like BMW's

We didn't say that we'd produce *nice* BWMs, or ones that could be sold successfully, simply that they'd be recognizable. In Turing terms this might be phrased as, "While we might produce a machine that appears to be intelligent, we wouldn't necessarily want to be friends with it". So we concede that recognisability is a necessary but not sufficient condition. It seems the BMW designer jobs are safe, at least for the time being.

Or are they?

## 5. Discovered, Gaudi's lost sketches!

For the next thought experiment, we will claim that we've discovered a previously unknown sketch book of Antoni Gaudi (say, buried away in his mother's attic). In truth no such book has been discovered; instead we have a generative system which produces drawings which resemble the work of Gaudi.

In this version of the test, we will send the images from "the sketch book" to an esteemed panel of architectural historians, architects, and other experts. What if they cannot tell the difference? What if they believe these to be images produced by Gaudi himself? In that even we would claim that we've successfully mimicked Gaudi. Who could say otherwise?

We are so pleased to have our Gaudi generator, the first thing we might do (after apologizing to the experts for duping them), is to apply it to the job of completing the Sagrada Familia. How wonderful it will be, that we can finally complete Gaudi's work, *as if he had been there to complete it himself!* [5]

Now surely the work of Gaudi is the culmination of his life, his experiences, his knowledge, his passion. Thus, if the results are Gaudi'esque, is not the machine too? Haven't we encoded something of what it means to be Gaudi, the human, his life, his

experiences, etc.? Have we created a kind of partial working copy of Gaudi, in a box?

### 5.1 Argument: Art is the sole domain of a conscious entity, the sole domain of humans

Turing addresses a similar argument in terms of consciousness as a requirement for intelligent behaviour. He cites Professor Jefferson's Lister Oration for 1949 as an example of such an argument:

> Not until a machine can write a sonnet or compose a concerto because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that machine equals brain—that is, not only write it but know that it had written it. No mechanism could feel (and not merely artificially signal, an easy contrivance) pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes, be charmed by sex, be angry or depressed when it cannot get what it wants. (445/6)

Turing argues that the only way to know if an entity is conscious or not is to be that entity. He also provides a little “viva voce” dialog to illustrate the kinds of responses that one would naturally interpret as pleasure, anger, etc., and from which one would reasonably conclude that the entity has emotion. Any entity that behaved such would be considered to have those emotions, machines included.

We can extrapolate from this that any machine that can mimic the full gamut of human emotions is theoretically capable of producing art.

## 6. Conclusion

We’ve posed the question, “Can machines produce artefacts which are indistinguishable from the creative output of a human?” and applied a variation of the Turing test as a way to frame the discussion.

While we’ve avoided defining art, we’ve argued that we cannot discount the possibility of machines producing it. We’ve also argued that the more we are capable of mimicking the output of a specific artist or architect, the more we have captured something about what it means to be that artist or architect.

This line of reasoning is likely emotionally bothersome for some. The notion that we may replace specific artist/architects with computerized versions has a kind of “Invasion of the Body Snatchers” [6] feel to it that somehow reduces our very notions of humanity. However, if the goal of Generative Art is to codify creative output, to “generate” “art”, we must question what qualities we’re willing to ascribe to the generator, and the extent to which they should or should not differ from those we ascribe to creative acts of humans.

Is generating Gaudi to know him? Or on more practical terms, who do you assign the output copyright to, the programmer, the computer, or Gaudi?

## 7. References

[1] “Meta-Code, Identity’s Borders”, Celestino Soddu, Generative Art Conference GA2004 proceedings, Milan 2004

[2] G.Matteo Codignola, “Generative Recognizable Codes from Gaudi’”, Generative Art 2000, Proceedings of the 3rd International Conference, Milan 2000

[3] Joseph Weizenbaum, "ELIZA — A Computer Program For the Study of Natural Language Communication Between Man And Machine", Communications of the ACM 9 (1), USA, 1966

[4] Alan Turing, “Computing Machinery and Intelligence”, Mind, Oxford University Press, England, 1950

[5] Schumacher, Edward (1 January 1991). "Gaudi's Church Still Divides Barcelona". The New York Times.

[6] “Invasion of the Body Snatchers”, Allied Artists Pictures, USA, 1956