

Big Data and Creativity

PALLE DAHLSTEDT

Division of Interaction Design, Department of Computer Science and Engineering, Chalmers University of Technology & University of Gothenburg, SE-412 96 Göteborg, Sweden. Email: palle.dahlstedt@gu.se. Also Department of Communication & Psychology, Aalborg University, Rendsburggade 14, DK-9000 Aalborg, Denmark

Big data and machine learning techniques are increasingly applied to creative tasks, often with strong reactions of both awe and concern. But we have to be careful about where to attribute the creative agency. Is it really the machine that paints like van Gogh, or is it a human that uses a high-level tool to impart one pattern upon another, based on her aesthetic preferences? In this paper, the author analyses the problem of machine creativity, focusing on four central themes: the inherent convergence of machine learning and big data techniques, their dependence on assumptions and incomplete data, the possibility of explorative search as a new creative paradigm, and the related problem of the opacity of results from such methods. The Google Deep Dream project is brought in as an example to illustrate the discussion. Information and complexity are brought into the discussion as central concepts for both creative processes and the resulting artefacts, concluding that the complexity of the interaction between the creative agent and the environment during the creative process is a crucial parameter for meaningful creative output. Based on the exposed limitations in current technologies, the author concludes that the principal creative agency still lies in the developers and users of the tools, not in the data processing itself. Human effort and input still matters. But we can take a constructive approach, regarding big data techniques as tools one order of magnitude more complex than what was available before, allowing artists to work with abstractions previously unfeasible for computational work.

Introduction

As an artist I explore the unknown, and as a creativity researcher I study how the unknown can be structured and explored, and I develop tools and technologies to do it. The unknown is an infinite space, full of data to be discovered – or generated, depending on your ontological standpoint. There are many parallels between my art and my research on one hand, and the problems and issues connected with Big Data on the other. In this article, I will apply some of the theoretical reasoning I have previously developed in the field of artistic and computational creativity to the

discussion about Big Data in general, and to Big Data as potentially used by artists. I will draw upon personal experience from my artistic practice, research into the mechanisms of creativity, a number of examples, and experience from the development of tools to manage such huge or infinite data spaces.

In some contexts, big data is regarded as the solution to almost every conceivable problem: artificial intelligence, business optimization, customized content creation, crisis management, epidemics prevention, and much more. At a symposium on the topic, the vice president of a leading data analysis company talked about big data as ‘the new oil’.¹ These optimistic attitudes towards big data are partly justified. It can certainly lead to new insights, research results, and economic and societal change. However, we have to be cautious. In a discussion at the same symposium, the moderator asked an important question: can we find correlations and patterns in big data if we don’t know what we’re looking for? The panellists answered unanimously ‘yes’ to this, and briefly described the inference models they use to find such correlations. But I do not think we really can be that optimistic, and this goes back to some fundamental problems of creativity and artificial intelligence.

I will only give a short answer here, and then return to the problem later. These models look for correlations in unknown data, but only for correlations we know how to describe. They will not be able to find new kinds of correlations or new types of patterns in the data. In a sense, we will find what we are looking for. To find the unexpected is related to creativity – it is not about finding more of the same, but finding something that is truly new, not only in content, but in form and structure.

The main themes I will discuss are:

- The use of typical Big Data techniques has a conforming effect, i.e. it is inherently non-creative.
- Explorative search as the creative process offers a new paradigm of artistic creativity.
- Material and results obtained through large space search or bottom-up methods are often opaque and difficult to understand, which may make them less useful.
- Big Data techniques and the way we use them are based on assumptions and design choices, and these have aesthetic implications.

The text is structured as follows. After an introduction to big data, and to the limits of our perception of it, two essential problems of big data approaches are discussed – inherent convergence, and incomplete data and context. Potential work-arounds are discussed, followed by a discussion about style and novelty and their role in creative processes. Spatial exploration is introduced as a metaphor for creative processes, which together with generative techniques provide for a new creative paradigm: sow, cultivate and harvest. Implications of this new paradigm are discussed, as is the problem of the opacity of machine-generated solutions. This is followed by some thoughts on the role of human skill and effort in creative process, and a section on the role of complexity in computer-generated art. Finally, Google

Deep Dream is analysed as an example of a big data system with supposed creative abilities, in the light of the previous discussions. The paper ends with some conclusions about using big data techniques in creative applications, and the role of the human in such applications.

What is Big Data?

So what do we mean by ‘big data’, and why is it important to artists? There are many definitions of big data. Here is a typical example from a leading big data company:

Big data is high-volume, high-velocity and/or high-variety information assets that demand cost-effective, innovative forms of information processing *that enable enhanced insight, decision making, and process automation.*² (Emphasis added)

Often, the very ‘big-ness’ of the data is loosely defined by three Vs: large *volumes* of data, data collected at high *velocity*, and data of big *variety*. Another leading big data company prefers to add two more characteristics to that: *variability* (here meaning big variation in speed of acquisition) and *complexity*, but those could be considered to be included in the original third V – variety.³ The way the term big data is most often used, however, it refers not only to the data itself but to the technologies needed to manage it, and to an approach towards what data can be used for, since it opens up data applications that were hitherto unthinkable.

So, when talking about big data in this text, I do not only mean the huge datasets, but also the surrounding infrastructures and technologies, and applications of them. For example, I include the following:

- The use of probabilistic machine learning models as a representation of the properties of a very large set of data or the behaviour of a very large set of agents, gathered through data.
- The idea that one can extract and define generative processes and behaviour from a large amount of data through analysis.
- Methods to quickly search through very large datasets, through for example advanced indexing and filtering.

Big data is in itself an abstraction, but acquires meaning and purpose from its use, through various kinds of analysis and data extraction, and integration of such analyses or the results thereof in products and services. Such processing invariably includes data reduction, i.e. making data more manageable through filtering, visualization or as data for training machine-learning models. All data processing is based on assumptions that constrain the future use of the data. Data collection is based on assumptions about what is important in the data (since you can rarely save *all* data), in what form it should be stored, what is really represented in the data and how it is structured. In data analysis, we introduce assumptions about what kinds of patterns we are looking for, how we think they can be detected, along which dimensions will data be reduced, etc.

Big Data Applications

The field of big data has emerged naturally, as an answer to the question: how can we make better use of all this information that is gathered, for different purposes? Data are gathered for a variety of reasons, in different contexts such as business logistics on websites, logging of people's interactions with online services, communication flows, and libraries, archives and databases of information that have been built up over decades but are now more easily consolidated.

The purpose of big data processing and analysis is often one of optimization of cost and time. It can also be to detect design needs and as a basis for decision making. In research it opens up new methods in, for example, DNA analysis. As with all new tools, it also leads to new kinds of services, such as search engines, including Google's image search,⁴ which would not be possible without huge amounts of training data for its image analysis models (images together with text containing descriptions of image contents). Other common uses are recommendation systems for media services and web shops, and the tailoring of media contents (entertainment, but potentially also art) to the recipients' assumed needs and preferences.⁵ There are already simple systems for media generation, e.g. software for word suggestions as you type on your mobile, which is based on big datasets from many users as well as continuous analysis of your own typing. There is software that can imitate your musical playing in real time and continue where you stopped.⁶ From there, it is but a short step to software generation of aesthetic material and intelligent creative tools, and such systems already exist in experimental form. From there, the step is not so big to fully autonomous creative systems. Or is it?

Big Data and Art

There is already some direct usage of big data techniques in the arts. For example, in the field of generative arts, and more specifically within the area of computational creativity, machine learning and AI techniques are used to generate new material from large training sets, and datasets of different forms are used as raw material for artworks. There is a whole genre of art based on or relating to visualization of data, or where significant portions of raw material are derived from large datasets, either from scientific observations or output from scientific models (the well-known composer Iannis Xenakis used the latter approach in his orchestral work *Pithoprakta* in the 1950s), or from databases or real-time data streams from the social or natural sciences, humanities (e.g. linguistics) or medicine.⁷

In the discussion, I will also include generative artworks where very large datasets are derived from generative systems, because they have similar properties – you have more data points than you can overview, and you need to find interesting details or subsets therein, derive new similar material from properties of the dataset, translate the data to a new medium, or draw overall conclusions from the properties of the dataset.

Experience Big Data

The human mind is limited. We can only hold a small amount of information in working memory, but over time we are able to form mental maps of larger sets of information,^{8–10} and mental models of the underlying processes,¹¹ such as geographical landscapes we move in, timbre spaces of musical instruments, experimental research data, or parameter spaces of machines we use in our work. We can to some extent, based on such mental maps and models, predict what specific sound will be produced given a certain fingering on the instrument, or the behaviour given certain parameters of a machine. Still, such mental maps and models are not perfect. And our abilities soon fail us when presented with extremely large datasets, if nothing else because they take a very long time to perceive and process. Also, we have a very limited bandwidth for how much data we can take in per time unit. If we are exposed to more, we cannot process it.

To illustrate this, I have prepared a few very short sound snippets, which are available in the online version of this article (Media Example 1). Please listen to those before you read on.

Media example 1. Seconds of data (<https://clyp.it/d44ggnqb>)

What you can hear may have appeared as five short noise bursts. But it was the following: a complete reading of George R.R. Martin's series of novels *The Song of Ice and Fire* (well-known as the basis of the TV series *The Game of Thrones*), a complete reading of the Bible, the complete music produced by the Beatles, a recitation of the whole of the Koran, and a reading of Shakespeare's complete works. Each one lasted exactly one second, but they were not time-compressed. Instead, they have been divided into short snippets of one second, and then all of these snippets were played at the same time, panned from left to right in the order they appear.¹² In the case of the Bible, you heard *Genesis* to the far left, and *Revelations* to the far right. So, in theory, no information was lost – but we cannot digest it in such a short time. Too much data becomes a grey mass. Regarding amounts of raw data, *The Song of Ice and Fire* is the largest (about 4300 pages, and 200 hours of reading aloud), and the Beatles' complete production the shortest at about 10 hours of music.

Humans can perceive things of a certain scale, in space, in magnitude, and in time. But we have invented scientific instruments to translate phenomena which fall outside of that scale, so that we can observe them anyway. Measurements of imperceptible quantities are visualized with indicators or graphs of an adequate size. A microscope translates very small things to something that fits our retina, and a telescope does the same with very large but very far things, and also gathers large amounts of light to produce images adjusted to the sensitivity of the human eye. In the same way, we can speed up or slow down moving images to be able to perceive and understand very fast or very slow processes, which are otherwise hard to grasp.

Big Data is a grey, formless mass to the naked mind. When datasets are too big, we need translations that afford an overview, or at least tell us something about the

data in the form of its distribution and properties, its overall characteristics. Such a translation is by definition a simplification and reduction, but it can help us to form relevant mental maps or models of the dataset. We can use various data analysis and visualization tools to translate immense datasets to something we can make sense of. We can look at subsets of the data through filtering, reducing the number of dimensions, to be able to form useful mental maps and models. We can look at statistical properties such as mean values, standard deviations and correlations between parameters. Or we can create data models, e.g. we can train machine learning models so that they embody the essential properties of the data.

All these techniques can help us to interact directly with data, in search for the interesting specific. But they all rely on reduction and simplification, so we gain something, and something is lost.

Inherent Non-creativity of Machine Learning: What is in the Box stays in the Box

Since the raw data is big, and much too extensive for brute re-search each time you are looking for something, it must be indexed, structured, compacted, or modelled through machine learning models. As such models are designed to find or produce the most likely outcome, they converge (metaphorically) towards the mean. They stay inside the box by definition. Let us look a bit closer at this problem.

First, another direct experience of big data, showing the convergence towards realistic text output, when using a gradually more complex underlying generative model, see Media Example 2 in the online version of this paper.

Media example 2. A video showing – in a stylized graphical layout – gradually converging texts of my co-authors, one word at the time for a selection of researchers, driven by a probabilistic model trained on a corpus of academic publications of each researcher. In the middle, a text generated from a model trained on all authors is shown (<https://vimeo.com/345663627>).

The problem here is both related to the convergence of probabilistic models and to the use of an incomplete representation of the real world, based on a number of design decisions. The model learns from the training set, i.e. the data fed to it, selected according to what somebody wants it to learn. The data are fed into a structure that somebody designed to capture what they consider essential properties. The data are provided in the medium chosen by somebody for the representation (text, 2D image bitmaps, sound files, etc), sometimes because it was the only practically available format, and sometimes because that particular format is easy or computationally cheap to process. If all this is done right, the model will be able to generate something that has similar statistical properties as the objects in the training set – in the chosen medium, according to the representation and data selection made.



Figure 1. The first few bars of each of J.S. Bach's 15 *Two-Part Inventions*¹⁶.

An example is David Cope's work with musical style studies.¹³ He trained a probabilistic model on existing music of, for example, Debussy, and then generated material with similar properties, finally arranging this material into a new piece sounding as if it were composed by Debussy. But since this new material has been generated to have the same properties as authentic music by Debussy (regarding sequences and probabilities of notes, intervals and rhythm values, for example), it will not add anything that Debussy had not already composed. It will be completely inside the box. In addition, the model will be unable to catch qualities of the works that are excluded from the chosen representation, such as Debussy being inspired by Indonesian gamelan music at the world exhibition in Paris in 1889. A statistical model based on such reduced data may capture musical traces of such gamelan influences, but will not be able to add new influences of the same kind, from other music he could have heard, interpreted in a way that would be characteristic for Debussy.¹⁴ It captures low-level properties of the music, but is disconnected from the semantic dimension of the music, from its link to the outside world. It emulates symptoms, not causes.

Such imitational systems are quite common in the field of artificial intelligence and computational creativity research. Let us look at another example illustrating why they are inherently uncreative (this argument was first presented in a paper on algorithmic composition in 2001).¹⁵

Consider Johann Sebastian Bach's collection of *Two-Part Inventions* (Figure 1).

Each of these inventions is different from the others, but they also share some common traits. If we metaphorically visualize their conceptual content as a diagram of a subset of ideas in a conceptual space (of all possible musical ideas), each one contains some unique ideas and some shared ideas – in relation to the collection of all 15 inventions. Each invention adds to the surface area, and together they form a large area in the space of possible musical ideas, here depicted as a circle, with the conceptual content of each invention marked as a subset within it (see Figure 2).



Figure 2. A hypothetical diagram of the set of ideas constituting (from left to right): Bach's First Invention, Inventions 1 and 2, the five first Inventions, and finally, all 15 Inventions, with a circle showing the complete set of ideas of the whole collection.

A machine learning model learns the properties of this set of material and will generate material with similar properties. This means that all material coming out of it will be inside this area. But remember that each invention added to the combined area in the conceptual space – it contributed something unique. That means that if Bach would have composed one more invention, it would most certainly have stretched outside of the marked area, i.e. it would have contained some conceptual material that did not already exist in any of the other inventions. There is no way we can know in which direction he would have expanded his ideas. But we can be sure he would have included something new also in the 16th Invention. Still, a computer model will stay inside the area of the 15 Inventions, no matter how many new inventions it composes, simply because there is no way to know in which direction Bach would go next. To model that, the creative model would have to include his whole situation, environment, and his complete musical knowledge. It would have to model the causes of the music, the conditions for its coming into being – not the results. A model based on learning from outcomes cannot extrapolate into the unknown – only interpolate within the known. Patterns within outcomes do not reveal the causes, since there is no one-to-one mapping back from cause to effect.

The Problem of Incomplete Data and Context: We Get the Answers We are Looking For

In the previous section, I discussed the scope of data models, preventing them from breaking out of the training set. But there are also other principles keeping them inside the box. Typical big data analysis methods look for correlations and patterns using various probabilistic inference methods,¹⁷ resulting in networks of dependencies between variables in the data. But they are only able to look for patterns we know how to describe. The search for correlations relies on assumptions based on previous knowledge. It is not possible to just pose the question 'find all dependencies in this dataset' unless it is very small or contains a very limited set of variables. Thus, we essentially ask well-formed questions to the data of the form: 'We think A depends on B and C, B depends on D, and C depends on D – what are the strengths of these dependencies given this data?' The result is a set of probabilities for each

variable, given the state of the other variables, which can be used as a generative model to mimic the behaviour of the process that generated the original data. It is of course a little bit more complex than that, but I simplify slightly here to make it easier to understand.

The result is a model that is based on correlations, not causality. This is a classic problem in empirical science, but it is especially troublesome in the field of computational creativity, because it does not help us understand what generates the data, nor the causal relationships behind the data, only the strength of the covariance between different variables. The trained model and the detected correlations capture some patterns in the training data, but the causal dependencies of the outside world are too complex and are difficult to include in the model. And the model only has access to the data, not to the rest of the outside world. Hence, we can only talk about properties of existing entities and relationships in the dataset, but not how they came to be, and not deduce procedures to generate them.

In addition, nonlinear dependencies in the data can be formulated in infinitely many different ways, and are therefore hard or impossible to find, without knowledge of the underlying system.

This is a severe problem limiting creative use of such models. Creativity is about producing new patterns, while most machine techniques for analysing large data detect, classify or reproduce known patterns. Generative models that could involve causal relationships might be able to generate more creative results. But not without including substantial parts of the real world in the model.

The problem of missing causal relationships is related to the problem of data reduction, mentioned earlier. Many dimensions of data are missing, and the original context of the data is lost. This is a devaluation of the specific, since the models try to generalize all data and fit it into a pattern, while it lacks most information that caused it in the first place. To illustrate why this is a problem in creative applications, let us once again consider the *Two-Part Inventions* of J.S. Bach. Each invention is in a different key, and each invention also has a different character. To better understand these pieces, you need to know about the doctrine of the affections, a Baroque theory stating that music and the arts could evoke certain affects through certain movements and tonal characteristics (see Mattheson for historical sources,¹⁸ with a summary and partial translation in Mattheson and Lennberg¹⁹). According to this doctrine, each key represented a certain mood, which was partly caused by the non-equal tunings that were used in the Baroque times. For example, D major represented royalty and fanfares, possibly because that was the most common tonality of trumpets and because it sounded quite pure in the most common tuning systems of the time. G major was a very talkative key, while E major symbolized pain, sorrow and agony – since it sounded very awkward and skewed in the tuning systems used. These ideas were widespread in Bach's times, and most definitely influenced his works, including his *Inventions*. These characters are clearly audible in those pieces. Another quality of the *Inventions* is that they were explicitly composed by Bach as pedagogic examples of how to compose, to be used when teaching his sons, and others. We can assume that he then wanted to give a broad range of examples of

how to construct melodies, harmonies and counterpoint, and different affects, and this pedagogical aim also had aesthetic implications (this idea is further elaborated by Lowrance in relation to another of Bach's collections of keyboard works, but also applies to the Inventions²⁰). Any generative model that tries to imitate Bach without taking the aesthetic, affective and semantic contexts of these pieces and of his time into account is doomed to fail. Again, it will imitate symptoms, not causes.

Another example of how the context is essential is Picasso's periods (blue period, rose period, cubism, neo-classical period, and so on). A data model might be able to detect them and cluster his works into periods, and this could, for example, be done with an unsupervised learning method such as a Kohonen self-organizing map,²¹ given the right kind of input. In this way, software systems can be said to be able to distinguish between the periods, referring each of his works to a specific period. But this is a passive, after-the-fact ability. These models cannot capture the reasons for his change of style, which may have been caused by incidents in his life, by significant events, encounters with other art (e.g. African art), a feeling of having exhausted the expressive potential of the previous period, or exchanges of ideas with friends and colleagues. And they cannot generate paintings from a hypothetical period Picasso could have entered if he had lived longer, because they only take into account the properties of the *outcome* of Picasso's creative processes, which in a larger context should also include his life in all its details.

Probabilistic machine learning methods are passive and analytical, and once again are not able to capture the real complexity behind the creative processes of these works of art. While data reduction may be a meaningful simplification, enhancing important dimensions in the data by discarding insignificant ones, every data reduction is still based on assumptions, and those have implications.

Possible Solutions

Above, I have described some of the main problems with generative models based on data: their inherent convergence, and the lack of context and causality. Are there ways around these problems?

If we look at the problem of generating novel artistic material based on learning from an existing corpus, there may be ways around it. Maybe not ways to generate the music Bach would have written if he would have continued, but ways to use probabilistic machine learning models to generate unheard music. But it requires creativity by the users, and the system itself will remain uncreative. For example, a generative system trained on a specific learning set of music could generate novel material which is organic and complex (features the user may want to inherit or retain from the training data, i.e. from the real world), by altering the weights or connections inside the probabilistic model through some kind of mutations, inversions or transpositions, thus diverting from the properties of the training set. Still, this would probably have to be done blindly or through trial and error, for reasons that are explained in the next section. To my knowledge, this has not been done yet.

One could also combine different training sets, either through the merging of two or more already trained models, or by training a model on two quite disparate sets of data, to achieve an advanced form of hybrid results. This is an approach I used myself when generating the lyrics for a work for vocal quartet and electronics, *Var fjärilsam / Be Butterflyish* (2000). It was a commissioned work, and the theme was *passion*. I chose to train a probabilistic model with text material from two related but very different categories: romantic amateur teenage poetry, and erotic short stories, both harvested from the internet. The generated output was a strange kind of hybrid text material, centred on love and eroticism in a peculiar way, which I then manually edited into a suite of new poems. Here, machine learning techniques were used to merge two styles into something new, but the creative part was done by me, in my choice of source material, of model parameters, and in the selection and editing of the output. Still, I would not have been able to produce these texts without using these tools. And the tools would never be able to produce them without me.

Exhaustion of Data

Replicating the style of known works may be interesting and refreshing for a short while. The experience is also relative to the viewer's or listener's level of previous knowledge. But even for the new listener, such a system will exhaust its possibilities after a while. Since the output is confined to a small part of the conceptual space, you have soon heard all its tricks. At the same time, the viewer or listener develops an increasingly accurate mental model of the generative system. After a certain time, the output stops being surprising, and feels predictable and boring.

This also means that you cannot judge a generative system from just a single example of its output. Any program could generate just one fantastic painting or piece of music – it could simply be completely described in the code. The difficult part is to generate variety and surprise over time. Hence, you must study a larger set of outputs of the system, and observe its variability over time, and the scope of its output. How long before we start to recognize the results?

It is possible to temporarily circumvent this material exhaustion by introducing mechanisms for superficial variation. But without varying the underlying generative core, you get stylistic variation without real structural variation. This is the case in my work *Ossia*, which generates complete new piano pieces based on evolutionary computation.²² On top of underlying music generation, there are a number of different 'presets' for the global generative parameters, which makes the output surprisingly varied. Still, these presets form a finite set, and the underlying tonal structures are recognizable, at least to me, as I have listened to thousands of pieces generated by this system.

To overcome this, a generative model could be made to grow or change over time, in interaction with its environment in some way. As long as it is a closed system, its possibilities will be exhausted after some finite amount of time, but if it is open-ended and connected to the outside world, it may have the potential for continuous

unpredictable change on a material level. But this kind of open-endedness is very hard to implement on a conceptual level. This problem is discussed in depth by computer artist Jon McCormack,²³ as one of the open problems of computer creativity. It is related to what creativity philosopher Margaret Boden calls transformational creativity – when the conceptual space is fundamentally altered or reshaped by the introduction of a new parameter, dimension or structuring principle, which provides a remapping of the space so that, literally, new ideas or variations are possible that were not possible before.^{24, 25} Such a change must in some way come not from within the conceptual space itself, but from the outside.

If we look at the problem of generating something like Picasso's period transitions, it relates directly to this, and to the lack of causality and context in generative models. Changes in style result from new concepts being introduced, which alter the artist's conceptual space. Current systems are trained to find correlations and patterns within a set of outputs from a specific artist. Perhaps if such a system was instead trained to find correlations between events in the artist's life or other external data, and the corresponding results, it could develop an ability to generate new works triggered by hypothetical events (as input) or generate a hypothetical life story including the supposed artistic outcome. Perhaps this is easier now when we live in the age of abundant data about life events, but it is still no easy task.

What is a Style, and Can we Transcend it with Algorithms?

I have mentioned generative data-driven models that generate material within a specific style. But what is a style, and how are styles formed? And why do we want to go outside of such established modes of expression?

Each work of an artist helps define her personal style. Similarly, a community of artists often implicitly or explicitly shares an aesthetic context, where each work of each artist in this context helps define a shared style. When an artistic context or style is new, and artists are experimenting, each work contributes heavily to the definition of this new style, and each new work pushes the boundaries of what can be done within this context. When the style is established, late-comers are able to produce works in this style, but it is not considered very creative to do so, in the same way as imitative computer models are non-creative. And each such late work is positioned well within the sphere of ideas of this style, not pushing the boundary in any specific direction – and not contributing anything new. This is very similar to the automatic generation of imitative pieces discussed previously.

The reasons for crossing the boundaries of a context – style extension – might be coincidental ('I can't predict how this will sound'), whimsical ('I don't care what anyone thinks'), curiosity-based and experimental ('What happens if . . .?'), or based on the artist's characteristic inabilities (e.g. the characteristic way in which I typically fail to compose a perfect fugue).²⁶ When the result is out, it becomes part of the artistic discourse. Each artwork is a statement in a never-ending discussion between artists about what can be made and what is important to make. Once it is out, it is

assimilated into the style or context the artist is part of, helps define it, and expands its sphere of ideas in the conceptual space. Once it is done, it is done, and it makes no sense to do the exact same thing again. And the next work will extend the sphere of ideas in the conceptual space a little bit further again.

We can also look at the same phenomenon – an expanding sphere of ideas – from the perspective of the individual artist. Most artists develop conscious or unconscious strategies to extend their sphere of ideas, to progress beyond what they can imagine. This is necessary to avoid stagnation and to continue to challenge oneself. Strategies to renew oneself might be chance, encounters with other artists, mind-enhancing substances, or formal procedures. Let us look at the latter example, since this is what might lead the artists furthest away from her current sphere of ideas. Self-imposed rule-sets and prescribed procedural steps incur both restrictive and generative properties on the result. They constrain what can be produced, but they also help to produce it. And they may help her to extend beyond her comfort zone. Therefore, artists of all kinds have used such procedures since antiquity.²⁷ Since the formal procedure projects into the unknown, and our mental models are inadequate for predicting the results of anything but the simplest formal procedures,²⁸ it is relatively easy for an artist to devise a set of rules or an algorithm that results in complex musical results beyond her imagination. But it is relatively hard – for a human – to predict the results without carrying out the steps of the algorithm. When I am exposed to the result, I may experience it as alien. I may have written every line of the code, but the result is dependent on emergent properties beyond my understanding. This makes it hard for me to edit the material, to add anything or to make changes, since I don't understand the inner musical logic of it. I may intellectually understand the process that led to these results, but that is not the same thing as the perceived musical result. At this point, I may choose to just put my name on it and let it out to the world. But I will not have learnt anything. I may instead choose to spend considerable time assimilating the alien result by listening intensively to it, getting accustomed to the new musical patterns, and thus gradually internalize its musical logic. After a certain time, I may feel more at home in the material, and understand what makes it good (or bad) music and be able to edit and make changes to it, without breaking its inherent consistency. At the same time, I will have extended my sphere of ideas to make it include this new musical expression. The next time I compose, I may be able to produce this kind of music without employing the same formal methods. Using a simple algorithm and an effort to learn, I have extended beyond myself. And by sharing the musical results, I have also extended the shared sphere of ideas within the aesthetic community in which I am active. By learning and assimilating the new musical logic, I have also perpetuated the aesthetic results into potential new works, my own and possibly others'. This whole process is related to how certain styles transform from written culture to oral culture.²⁹

From the above example, which is something I have experienced myself many times,¹⁵ it should be quite clear that the generative procedure that helps create the piece and makes me extend beyond my own aesthetic boundaries – clearly a creative act – is not derived from what I have composed previously, but quite the opposite.

Instead of aiming for the middle of the known, it projects out into the unknown. In an iterative process, we all together continuously extend our understanding of what is possible, our shared view of what music can be. In this process, computer-based generative processes can be a powerful tool, and a part of the creative process. The same reasoning is of course applicable to other art forms.

Vast Spaces of Possibilities

In the previous sections, I have used a spatial analogy for creativity and search, without elaborating further on it. It is a very useful analogy, which is often used in science to make high-dimensional processes graspable. Since humans are so used to thinking in terms of space, and location and regions in space, I will use the analogy here to discuss creative processes, and explain how generative and big data technologies can provide a new paradigm for creative processes.³⁰ These thoughts have previously been developed into a more detailed spatial theory of creative process,²⁵ and I have also previously discussed related aesthetic considerations.³¹

Creativity can be viewed as a search in a huge or infinite search space of possible results. If the search space is sufficiently large and unknown, this is a creative pursuit. This is very similar to the mathematical notion of a solution space – the theoretical space of all possible solutions to a problem or an equation. It is not a physical space, but a virtual space, a theoretical and mental construction. We are used to thinking in two or three dimensions, and most people have no problem envisioning such spaces. However, the spaces we talk about here often have a higher number of dimensions, and can even be infinite-dimensional, or structured in completely different ways, e.g. as the space of all possible mathematical formulas or the space of all possible sentences within a language, which are both hierarchically structured symbolic spaces where normal notions of dimension make no sense. Still, it is a useful analogy, since we are talking about a huge set of data points (solutions or potential results) which are structured in some way, quite similar to how points are structured in space. Some are closer to each other, and others are further apart.

When working with huge search spaces of musical potential,³² as defined by complex feedback systems with many parameters, an unpredictable algorithm of some kind, or a very large dataset, there is usually no direct way of getting a mental overview of the space. The space can be explored by random sampling, to get an estimate of the proportion of interesting content, and by taking very small steps to get an impression of how ‘rugged’ it is, i.e. how the interesting parts of the space relate to each other – if it is smoothly continuous or full of step-wise transitions. But I cannot make a complete map of it in my head, and what I bring away from my exploration is a journey through it – a trajectory consisting of a series of points that I have chosen, using some kind of interactive search method, and some general conclusions about the structure of the landscape.

How can one search in such a huge or infinite space in a more structured way? There are several approaches, most of them derived from optimization techniques in

engineering and nature. As it is impossible to do a brute search (in finite time) of the space, interesting sub-regions can be found by random sampling. Once an interesting point is found, we can look at nearby points and evaluate them. If one of them is better than the current point, we go in that direction. This is repeated until we can progress no further. This is called hill-climbing, because it resembles following an upwards slope (in very bad visibility) and stopping when we reach the top. However, there is no guarantee that the hilltop we have found is the highest – it may very well be a very small hill beside a much larger mountain (which we cannot see). A more refined way to find peaks in the space is to use the principle of Darwinian evolution, in what is often called artificial evolution, or evolutionary computation. You start with a set of (random) points, evaluate them, and based on which ones are best or most interesting, you create a new set of points. These ‘offspring’ points are usually created from pairs of good ‘parent’ points, inheriting some parameters (or however the points are represented) from each parent. Mutations can also be introduced, just as in natural evolution. In this way, the space is searched for better and better points (according to the current evaluation criteria), in a structured search process that balances chance and control.

Optimization or Exploration?

However, there are two problems with this: the relativity of aesthetic evaluation and the idea of optimization in art. First things first. It is hard, or even impossible, to formally define what is aesthetically good or interesting, and artificial evolution of this kind requires some kind of evaluation or selection criterion (it is usually called a fitness function). This is often solved by letting a human, the user herself, evaluate the candidate solutions, which is called interactive evolution (introduced by Richard Dawkins in 1987,³³ and first applied artistically by Karl Sims in 1991³⁴). This makes the process very slow, especially when working with time-based media such as music, and limits the number of generations that can be processed in a given time. And with only a small number of generations, it is difficult to achieve any kind of optimization.

Optimization might be important in engineering, but in the arts we often do not know what we are looking for. The aesthetic criteria change over time, and often they are not explicit at all. We may be searching for something unspecified, and when we see or hear it, we know if it is right or not. Also, if I know what I am looking for, it is unlikely that I learn much from finding it. I might know the desired properties of what I am looking for without knowing how to achieve it, and in that case it may be a small step forward. Still, it is likely a much more creative step to find something I did not even know I was looking for, or even knew was possible. *Exploration* is a more adequate approach to creativity than optimization, since it is based on the very idea of investigating the unknown in search of something interesting. So, interactive evolution may be a good solution to both these problems. You do not know exactly what you are looking for, but it still provides a tool for semi-controlled exploration of an unknown space. Let us call such exploration of an infinite space of unknown possibilities a *creative exploration*.

Sow, Cultivate and Harvest

Through the act of defining a search space, in terms of a generative principle, formula or algorithm, I sow potential, and through running the algorithms the potential is cultivated, and through my exploration and search, I harvest music, based on my starting point and search strategy, and on my evaluation and aesthetic preferences.

Through the use of such generative search techniques in the arts and in music, the practice changes from a kind of design and construction of artefacts to the harvesting of results from infinite result spaces, reaping emergent consequences from chaotic or unpredictable complex principles. If the explored space is large and unknown, this is a creative pursuit. If the explored space is also constructed or defined by the artist, the creative contribution from the author is even stronger. This new paradigm of artistic creativity has implications, both processual and aesthetic, and follows a peculiar logic.

This is a change of paradigm also within the generative arts. Previously, artists applied various generative mechanisms controlled by a set of parameters, possibly involving some random processes. Using trial and error, the processes were tested, evaluated and tweaked until one pleasing result was achieved. Or you filled existing systems with new content, and altered them in personally characteristic ways. These are of course simplifications, and there are many other models of generative creativity. Still, a large share of generative art was done along these lines.

Instead, you can now sow, cultivate and harvest. You construct or define the space, which equals filling it with content, but without exact control of details – and this is crucial. If you know all the details and are able to predict all the contents of the space, it is too small to be interesting. The space may be defined by a custom software with a certain potential result, or by a piece of hardware, such as a video synth or a noise-making machine with a number of parameters that you can interact with. Sometimes the space is already defined off-the-shelf, or defined by somebody else (an assistant, a collaborator, a member of a sharing community). This may be the case when you, for example, buy a sound-making machine (a hardware or software synth), or when you do not possess the necessary knowledge to define/construct/build it yourself.³¹

Then you start the cultivation, by exploring the space in search for interesting results, and thus realizing/generating the potential in the visited parts of the space (the garden). The best ones are brought back, out of the garden. This is not a construction of specifics but amounts to reaping emergent consequences from complex unpredictable principles. A particular result gets its meaning in relation to what is found before and after it, and from how it differs from these. Variation and change, not the actual material result, turn out to be the essence of the search.³¹ Sometimes, like with a coral reef aquarium, you do not harvest – you admire. Harvesting amounts to making a carbon copy of something in there, and then trying to admire the copy out of context. This can be compared to the recording of an improvisation – something is inevitably lost, the context is shifted, and the material loses its place in a system of references, from the moment and place where it was created. It might not work outside of the system, even though it seemed fine when we were still inside.

Judging from my own work with these techniques, and responses from users of such artworks, the experience of exploring such spaces – sowing, cultivating and harvesting a garden of strongly unpredictable results – is strong. This sense of genuine interaction, of being in the space and navigating it, might be considered the primary medium of such a work, just as the quality of the interaction is the primary medium of some interactive artworks, e.g. many works of the interactive art pioneer Myron Krueger.³⁵ The visual appearance or the actual sound is less important than the quality of the interaction.

The harvested artefact can be considered a souvenir from the journey, from the search, evoking the feeling of exploration, while the sense of interaction remains an elusive memory. Artworks of this form are very difficult to document and communicate to others, since you have to experience it yourself.

The Opacity of Machine-generated Answers

Advanced data analysis and machine learning models do not necessarily lead to new explicit knowledge, in the sense of an intelligible solution to the problem. They may lead to solutions and tools that work well, but that is not the same as human knowledge and understanding. Why is this? Let us take a simple example. Over thousands of years, humans have selectively bred better crops and livestock. This has been guided by the selection of phenotypes (actual organisms) that exhibit desired traits, without any knowledge of how they are represented in the corresponding genotype (the underlying genetic material). Gradually, we have achieved our goals: better food production. So, we are able to produce an organism, e.g. a larger and sweeter banana, without understanding exactly what makes it larger and sweeter. We have a working product, and we understood how we got there, but we do not have the actual understanding of what makes the solution work.

There is a similar problem of understanding results obtained through large-space search and machine-learning methods. Such high-level methods develop a solution over time, but the solution emerges from a complex process (e.g. training a neural net), and is not necessarily in a form understandable by humans. It may be operable, usable and give good answers, but it is opaque. When I evolve 200 parameters of a complex sound synthesis system and find a wonderful sound, I usually do not understand why it sounds like it does. All parameters have been changed in the process, but it is very hard to figure out what effect each parameter has on the end result, in a system where there is a high interdependence between the parameters. The phenotype (the actual sound) emerges out of a highly complex process from the genotype (the parameter set), just like my body has developed from my DNA in a highly non-trivial way.

When evolving for a certain behaviour using open-ended genetic representations, such as tree-structures and graphs, this is even more evident. Karl Sims, a pioneer of evolutionary techniques in art, said in his 1991 paper about a system for interactive evolution of highly complex images: ‘Users usually stop attempting to understand

why each expression generates each image.’ It is extremely difficult to decode those graphs and understand why they behave as they do. We know their exact forms, and the exact mechanism with which they are decoded and executed, but this does not tell us why and how they work, because the interesting parts are emergent. Very often, the only way to tell what the result of a computer program is, is to run it. It cannot be reduced.

The problem is not only one of understanding, but also one of editing. It is difficult to edit either the contents or the results of such methods. Because the mapping from genotype to phenotype is often discontinuous and complex, a small change in genotype can cause drastic, unpredictable changes in the phenotype. It may also be difficult to edit the phenotype, because evolved solutions are often of a structure very different from solutions constructed by humans. It is not evolved for structural clarity, but for its ability to solve a certain problem.

The same is true for probabilistic data models and machine learning. Take as an example the Deep Mind model that recently won against a Go game champion.³⁶ This Go-playing neural net was trained on large databases of Go matches and was also made to play against itself to reinforce what it had learnt. The programmers of course know the workings of the neural net, how each cell works and the maths behind the training. But it is very hard to pinpoint which part of the program does what when it is actually running, since its complex behaviour emerges from the interactions of a large number of small parts. The programmers cannot tweak the program to do more of this or more of that. They can re-train it, but they cannot access the specifics of Go-playing inside the program. They can answer the question ‘How do the neural nets react mathematically to new training input?’, but I do not think they can answer the question ‘How does it play Go?’.

While evolution, neural nets, and other machine learning techniques can give very complex solutions to a given problem, where the exact functioning of the solution is not of interest, but how well the solution works, the inner workings of the solution are not accessible or editable, except by further evolution or training. We cannot ‘edit’ the result, since the mechanism is too complex, or simply unknown, or structured in an unknown way. And if we do it, we cannot predict the exact consequences.³⁷

Skill and Effort

Closely connected to the issue of authorship is the issue of skill. Basically, it is about where the effort is put in in the creative process, and who does it. Is it easier or more difficult to work in this way, using sow, cultivate and harvest methods? Could anyone create music or art in this way? Yes and no. It is easier to create interesting sounds and material if you do not bother to design the underlying space yourself, but instead explore existing spaces (i.e. other people’s sound-generating algorithms or machines). But it is difficult to create interesting spaces that are easy to navigate while still containing great variety. There seems to be a trade-off between universality and the proportion of usable material, between open-ended-ness and control.

The idea of abstraction is fundamental to human thought. A tool embodies intelligent behaviour,³⁸ and tools are bearers of culture.³⁹ This enables an accumulated gradual increase in complexity in most human activities, as our abstractions and our tools evolve (as described by Brian W. Arthur⁴⁰). Human knowledge, as well as our skills in making artistic artefacts, have progressed by such gradual hierarchical abstraction, leading to an ability to manage very complex concepts and ideas – because we can pretend they are simple. Through abstraction, they hide their lower layers.

Consider the introduction of computers. They were principally nothing new, just mechanistic executors of mathematical recipes (algorithms). But the fact that they execute the recipes several orders of magnitude faster than humans enables us to devise recipes of hitherto unthinkable complexity and scale. We abstract such recipes into hierarchies of tools, and when we have learnt to operate on such a higher level, it is not more difficult than operating on a lower level.

As long as we only deal with the uppermost layer, things are easy. The difficulties arise when we have to keep several layers in our mind at the same time, including how they are linked and how they interact.

As a user (artist), you can adopt the lazy approach, and take advantage of templates and inherent content and characteristics built into new tools. You use the tools to achieve the same thing as before, only faster and with less effort. And your results will be as good as anybody else's, and probably quite similar.

Or you can take a more ambitious approach and use the new tools with the same effort as you put into the earlier, conceptually less complex tools. Then, you may end up at a new place, further away from the starting point. You spend the same effort as previously, but with a more 'efficient' tool.

When using complex creative tools to generate score material, you still need musical craftsmanship (knowledge of music acoustics, instrumentation, etc) to be able to edit and arrange your material. You need an understanding of musical form and human perception and cognition, to be able to present the result in a coherent way. Or you may only need this to compose for musicians, but maybe there are completely different musical forms that we do not know about yet, and they need to be learned, by both authors and listeners? During my explorations of creative tools I have more than once had to revise my ideas about what the essence of musical experience is. And new technologies open up a new kind of amateur musicianship, which is an important base for the appreciation of others' music. When everyone has the means, the ideas once again become more important than access to resources.

New paradigms also require new skills. Now we need to learn how to construct sound engines and algorithms that are suitable for these techniques, and how to use the search/harvest techniques efficiently. Learning these skills requires some understanding of how they actually work, like all tools, but this may also in the future be abstracted into a new layer, where the underlying layers are made implicit. Soon, when better open-ended representations have been developed, I think people could possibly evolve their own music, and arrive at pretty different results, essentially making the program an interactive creative partner – while keeping the user on, yes, the 'user' level, with no need for expertise.

Still, some people (the programmers, tool designers) will need skills in how to develop and design these kinds of tools, as an extended meta-skill. This may be partly, at least during the current research/experimental stage, a role of the artist-researcher. Artists always want to go one step further, and many want to develop their own tools, or at least customize (or abuse) existing tools. I guess I myself an example of this kind of artist. Will I, as the author-programmer of such a system, be a composer? Or a supercomposer?

Is authorship unavoidably distributed? Creating art in a sow–cultivate–harvest model is approaching the curatorial, and even a distributed curation. Somebody selects which data are collected, and selects which model to use to analyse/simplify. The people behind the data collection and analysis are toolmakers and have a role in the creative process. They are toolmakers and/or curators of the art.

Complexity in Generative Art and Data Art

In information theory, there are several ways to measure complexity. One way to measure the complexity of a dataset is as the length of the shortest possible computer program that can generate it, the so-called Kolmogorov complexity. The output from a random process cannot be compressed, and the shortest possible description equals the dataset itself. On the other hand, it can be imitated with similar statistical properties (on the macro scale) by a very short description, since the random material lacks internal structure, is easy to describe, and all details are replaceable. We do not perceive one random sequence as particularly different from another, even if they really have nothing in common except a similar distribution of numbers.

In generative art, the challenge is often to get as much out of as small a system as possible, without having the recipients figuring out the system too quickly. A good example is the beauty of fractals, a particular category of iterative mathematical functions, giving rise to beautiful visuals. Visualizations of fractals can be very complex, yet they are completely defined by a very small set of equations, which are iterated many times. Part of the beauty stems from our inability to grasp just *how* they emerge from these simple formulas – this is an example of what Jon McCormack calls the computational sublime⁴¹ – the awe of the incomprehensible qualities of computation. Such images have a very low Kolmogorov complexity, yet appear very complex. But they are complex in a recognizable way and, after a while, the viewer may become tired of exactly that kind of complexity. What goes beyond our understanding is relative to our current repertoire of mental models, and of what we are trying to do. The play with this border, pushing the limits of the computational sublime, is what we do when creating generative art.

What about complexity when working with found data? By using external material as input for an artwork, e.g. in aesthetic visualization, you inherit its complexity and information content. You have a chance to add something to what is already there, for example, through your selection of the data, through an innovative or unusual mapping or transformation of the data that sheds new light on it, through a characteristic simplification or a new reflective angle which becomes a comment on

the original data. You can place it in a juxtaposition to other data. Your role becomes a curatorial one in relation to the data. The original data may be complex and rich, but if the database is regarded as a given, the description of the work can still be made very short. The complexity added by the artist can be reduced to the choices of selection, filtering and mapping, which might result in a rather low complexity if taken lightly. This is a common risk with data-based artworks. Once again, it is all about the effort put in by the artist, i.e. the amount of information injected in the process, and the scope and complexity of the creative process.

A Big Data Example: Google's Deep Dream

We have previously discussed aspects of generative and data-driven creative tools illustrated with examples of a rather small scale. One interesting and well-known example of a creative tool based on big data techniques and machine learning is Google's *Deep Dream* – a by-product of Google's image search algorithms, introduced in 2015.⁴² It is a multi-layered neural net (actually several connected nets), trained on a large image set (in the demo version notoriously consisting mostly of dogs), which is then able to associate on a given image. Its own output is fed back into its input, reinforcing any faint similarity or association into something stronger.

A free online tool was released together with the source code,⁴³ for everybody to try, and quite a few did, and still do. The release triggered an extensive discussion about artificial creativity: Is this art? Is it creative? One year after its release (at the time of writing), it feels more like just another filter plug-in, with all output being very similar. It is interesting, since it represents a commodification of a complex big data algorithm performing an aesthetic task. Technically, it is impressive, in that it deals with the image components on a higher abstraction level, far above the pixel grid, but content-wise, it seems to be as conforming as any typical Photoshop filter, if used in a naive way. Kyle Chayka worded it quite well in a column from about one month after its release:

Deep Dream is future Kitsch. Kitsch 'offers instantaneous emotional gratification without intellectual effort, without the requirement of distance, without sublimation,' according to the philosopher Walter Benjamin, a pioneer of the idea. The description fits perfectly for this new genre of algorithmically generated imagery. It requires no criticality or particular intellectual effort to digest, nor does it provide much reward in return. Deep Dream is our own visual culture, chopped up in a blender and spoon-fed back to us.⁴⁴

Deep Dream consists of a complex process that has been abstracted into a complex tool, but with quite a low number of parameters and choices for the user, apart from the submitted input image.

Artistic quality is hard to define. However, we can often observe the efforts of the artist as traces in the resulting artefact/artwork. No traces of effort implies that it may be a superficial work (preparation and training efforts are also taken into account). When many users use a similar tool, which may well have very complex results, but everybody's result is complex in the same way, then there is not much

effort involved, and this is visible in the results. There are lots of characteristic traces of the algorithm involved, but not of the human user.

Thousands of users apply the Deep Dream tool to new source images of all kinds. But these images are all subject to the same kind of transformation. The end results are very similar, regardless of whether the source is a landscape image or a pornographic image. The tool takes over, and the artefacts produced by the process have a larger impact than the actual image contents.

The very moment a tool becomes available, it is not creative in itself anymore, since creativity can be defined in relation to what humans have (hitherto) been able to do (what Margaret Boden calls *H-creativity*, historical creativity). So, the moment we are all able to do a specific new thing, it is no longer creative. A creative pursuit could be to explore such specific tools in ways that are not yet explored. And when they are, the next explorer must go further.

Actually, this is not so different from conventional visual arts. Compare this to the description of the formation of new styles, earlier. The important contribution of, say, Duchamp's *Nude Descending a Staircase*, or a specific Picasso cubist portrait of a woman, is not the fact that it depicts a woman, but *how* she is depicted, and how that is different from how women/humans and motion had been depicted before. The creative content of that painting is in the process, in the transformation. Hence, inventing Deep Dream is highly creative, but using it is not creative in the same way. Programming an automated Duchamp-Descending-Nude-filter would be creative, but mass-producing images transformed through it is not.

The Importance of Representations

Representations are crucial for generative systems. The way in which the potential solutions are represented in the code decides how they can be varied and hence how the search space can be traversed.⁴⁵ This is also true for Deep Dream, which is designed to process (associate upon) bitmap images. What if we try to generate new pictures based on a raw material of musical scores? The deep dream system regards them as bitmaps and will not be able to decode the notation. It will impose variations upon it as interpreted from 2D shapes in the image, and reinforce associations to similar shapes, imprinting them upon the material (see Figure 3).

So, the material is transformed along the wrong dimensions, because the system is not using an appropriate representation. If the Deep Dream system was trained on numerical representations of musical scores (e.g. MIDI files), and presented with new material in that form, it would be able to associate upon it in a musically interesting way. As it is now, it will quickly destroy the notation and try to make it into a figurative image. This is not a fault of the system, but a limitation by design, and it is important for users to understand why such limitations exist, and how they can be circumvented.

This kind of limitation also appears in a subtler way in the Deep Dream system, which is only aware of the 2D projection of (potentially three-dimensional) images. A 2D pattern consisting of one 3D object partly covering another, where the



Figure 3. The online version of Google Deep Dream associating on a musical score without ‘understanding’ it. It seems to detect focal points and generates its familiar dog faces with these points as eyes and nose. The point of the exercise is to show that it is designed to work with bitmaps containing certain kinds of figurative images, not with symbolic content contained in bitmaps. Representation is crucial, and the choice of representation a particular algorithm works with makes it very domain-specific.

conjunction on the 2D bitmap is visually similar to a pattern that the Deep Dream algorithm has seen before, is reinforced as such, thus losing the 3D dimension, and becoming something else. This loss of the 3D aspects of the original image is behind a lot of the ‘surreal’ look of Deep Dream results. See Figure 4 for an example.

While the examples presented may seem both obvious and trivial, these representations-based limitations are often overlooked. These examples show that any such system can only manage and process the kinds of representations that it is programmed to process. And in translations between different kinds of representations, interesting things can happen – as an unintended side effect. Sometimes this is hard to distinguish from creativity.

Deep Dream and Style

The public Deep Dream system contains a tool to render a submitted image in the ‘style’ of another submitted image, or based on a selection of provided style templates, in the form of example images. There are two problems with this.

First, as discussed elsewhere in this text, style is not a fixed entity. All art is produced within a context. A style is defined in an accumulative way by a series of defining works or, in rare cases, by a single seminal work which has to be both ground-breaking and rich in details and complexity. Every single work ‘within’ a style also contributes to its definition, or to a sub-style or, potentially, to a later



Figure 4. An illustration of how the Google Deep Dream system (by design) treats everything as a two-dimensional image. The original image is to the left, followed by two processed images of gradually increasing processing depth. Configurations that could be dog faces, even if they are coincidental as a result of overlapping of objects of different depths, are reinforced, and the depth of the image disappears. This inability causes misinterpretations and could be considered one of the characteristic traits of the Deep Dream generative system, its personal signum.

offspring style, by adding something that was not defined by the stylistic framework. At the point when works are completely in the style, creativity has ceased. So mechanically re-creating something in a specific style adds nothing new, except the choice to do just that, which is a very small amount of new information, hardly contributing to the depth of the artwork.

Second, the resulting image is not representing a style in any other way than in its very superficial visual characteristics. The deeper meaning of the style, its underlying aesthetic motivations as a reaction or development of an earlier style, or as emerging from newly available tools, is completely absent in this kind of application. It is an emulation of results, not processes. Of symptoms, not causes.

It should be noted that this is not a critique of Deep Dream per se, but of how it is sometimes talked about and presented. Also, the demo version of Deep Dream, pre-trained on dog pictures, is a simplified proof-of-concept, and much more interesting work can be made using the system in full, run on a local machine using custom training sets. But that is exactly the point. It is but a tool, and it needs to be applied and contextualized by the user, as part of a process that involves a number of aesthetic choices (training set, source images, system parameters), just like the choice of using this tool is an aesthetic choice.

Enclosed and Open Processes

In human creative processes, there are many opportunities for interaction with the outside world. Artists, like all other human beings, are in the world, talk to others, see and hear things. Even if the artist would be a reclusive hermit, she would work in ever-changing physical surroundings and through a complex personality with memories and cultural baggage.

Most computational creative processes are enclosed or encapsulated in a black box. As a consequence, they are the same each time. Or, if they contain random variation, this randomness is applied similarly each time, which also is a same-ness but on the next level. They do not interact with the world during the intermediate stages of the process, so in one sense, there is no process – just a single operation.

Also, the process is static, in the sense that it follows the same principles each time. The results may be very different, but after a finite time the pattern-creation ability will be exhausted.

Maybe it is the automation that is the danger, the inherent lack of interactive process. As soon as it is automated, it is guided by (so far) quite simple principles, and those principles do not change over time. When such principles can be replaced by agents of a higher level of complexity, maybe different behaviour can emerge.

On the other hand, artistic material coming out of a lengthy process with rich outside interactions, many levels of decision-making, and possibly inclusion of existing (but transformed) data, will by necessity become complex. It cannot be described in any other way than by telling the story of how it came to be. The patterns in it cannot be reduced to simple formulae as generative patterns or be referenced to existing libraries (in itself a simplification, a short description). It is complex in a different way – not a random structure with high Kolmogorov complexity, but an ordered structure, rich in references to its own process and thus also to the outside world. And with a high Kolmogorov complexity.

Conclusions

It is hard to combine harvesting and exploratory methods on one hand, with constructive, traditional work such as sculpting, fine-tuning, polishing, tweaking and editing on the other, since the underlying generative principles are often unknown and inaccessible to the user. Are we then in the hands of those who define the search space and the methods for its exploration? Are we limited to searching and arranging, thus becoming curators more than artists? Is this a lesser degree of creative agency, or is it a liberation? What are the conforming implications of such working methods?

While it is important to be aware of the limitations of these kinds of tools, it is also important to see their place as an important part of an ever-growing ecology of agents, including both humans and machines, enabling continued exploration of the possible.

In all the examples presented in this paper, human input and the complexity of the human–machine interaction during the creative process have been crucial. Or, to put it succinctly: human effort and agency matters for creativity. Isolated black-boxed processes with complex output will by definition be either random or be complex because they contain complex inner states and mechanisms. They are not complex in their relation to the outside world, so their output appears more or less meaningless to us. Any apparent relationship to the outside world is inferred by our perception, in an unconscious urge to make sense of the world – an interesting and well-known phenomenon called *pareidolia*. Relationships to the real world are the foundation of meaning to the human observer, and as long as big data techniques

lack this connection to what is a fundamental part of being a human, they will not be creative in any sense that matters to us.

However, if we acknowledge this and see big data techniques not as the key to autonomous creativity, nor as causing the death of the artist, we can instead take a more constructive stance. Most big data algorithms are huge imitation machines, which can learn from the data we provide them with. They rely on incomplete data about the world, in several different ways. They are often monomodal, giving a very limited representation of aesthetic objects, thus losing their connection to the real world when reduced to a single kind of data. They can work with the representations we program them to work with, and with patterns we know how to describe formally. Within those limitations, they are extremely powerful and good at generating output that shares characteristics with the input. But, most importantly, they have difficulties working with the meaningful parts of data and artworks, i.e. their semantic content, their connection to the outside world. So, the creative responsibility still falls upon the user.

Big data techniques provide a way to allow humans to interact with data on a higher abstraction level than previously possible, and with very large amounts of data. They provide data models that allow us to deal with high-level concepts previously inaccessible to algorithmic processing. They allow us to use new creative paradigms, such as sow–cultivate–harvest and other exploratory methods. Through working in these new ways, human artists will develop new abilities and skills and be able to create art that was not possible before, using rich interactions with the algorithms and with the data.

Regarded as a new class of tools, one order of magnitude more complex than what was available before, big data techniques have great potential for artistic creativity. Big data brings not only new tools for artists, but maybe also new genres for artists to work with. We explore the vast solution spaces of particular tools, but also the solution space of possible tools. And an interesting question is if we can apply the same kind of explorative tools to this meta-search? Considering that human learning processes, such as the gradual assimilation of a piece of generative music, or the formation of a mental model of a complex algorithm or a musical style, are but more complex versions of machine learning processes (which are modelled after the human neuronal system), it should in theory be possible to create computer systems that behave creatively in the same way as humans, if this is desired. What is currently lacking is a rich interaction with the environment. By rich interaction, I mean a processual complexity, where the supposedly creative system interacts with its surroundings over time during a creative process. They also need an internal model of a significant part of the shared outside world, so that the produced results can acquire real meaning, instead of being carriers of projected meaning inferred by learnt superficial patterns. In short, a creative system needs to be situated in the world.

For the foreseeable future, we need to rely on creative tool makers to develop new tools that provide interesting transformations of conceptual spaces, provide new ways to explore the possible, and new paths to navigate the unknown. We need creativity from tool users to perform this exploration, to put in effort in using those tools, and not just accept the first output from somebody else's algorithm applied with default settings. Creative tools are not there for efficiency, but for exploration.

We need creative, artistic tool users who tinker with algorithms and their parameters, carefully select datasets and data features to use, who find unexpected combinations of data to feed in and explore, and who invent novel representations of known data.

Only in this way can they find previously unthreaded parts of the space of the possible and bring their findings back to share with the rest of us, or even invite us to take part in the creative exploration.

Supplementary Material

To view supplementary material for this article, please visit <https://doi.org/10.1017/S1062798719000073>.

References and Notes

1. From Big Data to Big Understanding Seminar organized by the Royal Swedish Academy of Engineering Sciences, Tällberg Foundation, and SAS Institute. 21 November 2013, Stockholm, Sweden.
2. Gartner IT Glossary, *Big Data*, <http://www.gartner.com/it-glossary/big-data/> (accessed 15 December 2016).
3. SAS Institute, *Big Data– What it is and why it matters, white paper*, http://www.sas.com/en_us/insights/big-data/what-is-big-data.html (accessed 15 December 2016).
4. Y. Jing and S. Baluja (2008) Pagerank for product image search. In: *Proceedings of the 17th international conference on World Wide Web*, Beijing, China, pp. 307–316.
5. J. Bobadilla, F. Ortega, A. Hernando and A. Gutiérrez (2013) Recommender systems survey. *Knowledge-Based Systems*, **46**, pp. 109–132.
6. F. Pachet (2003) The continuator: musical interaction with style. *Journal of New Music Research*, **32**(3), pp. 333–341.
7. Notable examples include the artworks of Mark Wattenberg, Nathalie Miebach, Wesley Goatley and Tobias Revell, and the artistic collaborations of Mikael Lundberg and Mats Nordahl, e.g. *Traces of an Ongoing Memory*. For an overview of the field, see F.B. Viégas and M. Wattenberg (2007) *Artistic data visualization: beyond visual analytics*. In: *Proceedings of the International Conference on Online Communities and Social Computing*, Beijing, China, *Lecture Notes in Computer Science*, **456**, pp. 182–191.
8. B. Tversky (1993) Cognitive maps, cognitive collages, and spatial mental models. In: *Proceedings of European Conference on Spatial Information Theory*, pp. 14–24.
9. L.A. Sudol, M. Stehlik and S. Carver (2010) Mental models of data. In: *Proceedings of Koli Calling: International Conference on Computing Education Research*, Koli National Park, Finland, pp. 55–58.
10. C.A. Chinn and W.F. Brewer (1996) Mental models in data interpretation. *Philosophy of Science*, **63**, pp. 211–219.
11. D. Gentner and A.L. Stevens (1983). *Mental models* (Hillsdale, NJ: Erlbaum).
12. This concept and the processed version of the Bible were originally conceived as part of a collaborative work by the author and the visual artist Mikael Lundberg in 2005, and programed by the author. However, the work was never finished

nor exhibited. The other sound examples were prepared for the keynote presentation that this article is based on.

13. D. Cope (2005) *Computer Models of Musical Creativity* (Cambridge, MA: MIT Press).
14. Cope's more recent work is more interesting from a novelty point of view, where he personally interacts over time with his probabilistic composition model (which he calls Emily Howell), training it in more complex ways.
15. P. Dahlstedt (2001) A MutaSynth in parameter space: interactive composition through evolution. *Organised Sound*, **6**(2), pp. 121–124.
16. J.S. Bach (1899) *Thematic Catalogue of the Bach-Gesellschaft Ausgabe, Band 46* (Leipzig: Breitkopf und Härtel).
17. T. Griffiths and A. Yuille (2008) A primer on probabilistic inference. In: N. Chater, and M. Oaksfors (Eds), *The Probabilistic Mind: Prospects for Bayesian Cognitive Science* (Oxford: Oxford University Press), pp. 33–57.
18. See J. Mattheson (1739) *Der vollkommene Capellmeister* (Hamburg: Christian Herold) and J. Mattheson (1713) *Das neu-eröffnete Orchestre*, vol. 1 (Hamburg: B. Schiller).
19. J. Mattheson and H. Lenneberg (1958) Johann Mattheson on affect and rhetoric in music (II). *Journal of Music Theory*, **2**(2), pp. 193–236.
20. R.A. Lowrance (2013) Instruction, devotion, and affection: three roles of Bach's well-tempered clavier. *Musical Offerings*, **4**(1), pp. 15–30.
21. T. Kohonen (1982) Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, **43**(1), pp. 59–69.
22. P. Dahlstedt (2011) Autonomous evolution of piano pieces and performances with Ossia II. In: E.R. Miranda (Ed.), *A-Life for Music: Music and Computer Models of Living Systems* (Middleton, WI: A-R Editions).
23. J. McCormack (2005) Open problems in evolutionary music and art. In: *Proceedings of EvoWorkshops 2005* (Heidelberg: Ed. Springer-Verlag), pp. 428–436.
24. M. Boden (2004) *The Creative Mind: Myths and Mechanisms*, 2nd edn (London: Routledge).
25. D. Hofstadter (1985) Variations on a theme as the crux of creativity. In: *Metamagical Themas* (New York: Basic Books).
26. P. Dahlstedt (2012) Between material and ideas: a process-based spatial model of artistic creativity. In: J. McCormack, and M. D'Inverno (Eds), *Computers and Creativity* (Berlin, Heidelberg: Springer).
27. For a historical overview and a more in-depth discussion of generative and restrictive properties of constraints, see chapter 3 of P. Dahlstedt (2004) *Sounds Unheard of: Evolutionary algorithms as creative tools for the contemporary composer*. PhD thesis (Göteborg: Chalmers University of Technology).
28. R.A. Finke, T.B. Ward and S.M. Smith (1992) *Creative Cognition: Theory, Research, and Applications* (Cambridge, MA: MIT Press).
29. G. Valkare (1997) *Det audiografiska fältet: om musikens förhållande till skriften och den unge Bo Nilssons strategier*. PhD thesis (Göteborg: University of Gothenburg).
30. This paradigm is not completely new, but distinctly different from the more common way of creating things, which may – greatly simplified – be described as musical artefacts being constructed from scratch and evaluated in an iterated tweaking process.
31. P. Dahlstedt (2009) Thoughts on creative evolution: a meta-generative approach to composition. *Contemporary Music Review*, **28**(1), pp. 43–55.

32. P. Dahlstedt (2001) Creating and exploring huge parameter spaces: interactive evolution as a tool for sound generation. In: *Proceedings of the International Computer Music Conference 2001*, Havana, Cuba, pp. 235–242.
33. R. Dawkins (1986) *The Blind Watchmaker* (Harlow: Longman Scientific and Technical).
34. K. Sims (1991) Artificial evolution for computer graphics. *Computer Graphics*, **25**(4), pp. 319–328.
35. M.W. Krueger (1977) Responsive environments. In: *Proceedings of the June 13-16, 1977, National Computer Conference (ACM)*, pp. 423–433.
36. E. Gibney (2016) Google AI algorithm masters ancient game of Go. *Nature News*, **529**(7587), p. 445.
37. I am simplifying a bit here, since we know a great deal about genetic low-level mechanisms. But there are still so many things we don't know about how these detailed changes are expressed in the high-level organism, which is the main point.
38. R.L. Gregory (1981) *Mind in Science* (London: Weidenfeld and Nicolson).
39. L.S. Vygotsky (1980) *Mind in Society: The Development of Higher Psychological Processes* (Cambridge, MA: Harvard University Press).
40. W.B. Arthur (2009) *The Nature of Technology: What It Is and How It Evolves* (New York: Simon and Schuster).
41. J. McCormack and A. Dorin (2001) Art, emergence and the computational sublime. In: *Proceedings of Second Iteration: A Conference on Generative Systems in the Electronic Arts*. (Melbourne: CEMA), pp. 67–81.
42. Inceptionism: going deeper into neural networks. Google Research Blog (17 June 2015). <https://research.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>
43. <https://deepdreamgenerator.com> (accessed 17 December 2016). The website has since been updated with more advanced versions of the algorithm.
44. K. Chayka (2015) Why Google's Deep Dream is future kitsch. *Pacific Standard* (10 July 2015).
45. For an extended discussion of this problem, see P. Dahlstedt (2005) Defining spaces of potential art: the significance of representation in computer-aided creativity. In: *Description & Creativity Conference, online proceedings*, King's College, Cambridge, UK, 3–5 July 2005.

About the Author

Palle Dahlstedt is Obel Professor of Art & Technology at Aalborg University, Denmark, and Associate Professor in Computer-Aided Creativity, Department of Computer Science and Engineering, University of Gothenburg and Chalmers, where he is also head of the Interaction Design division. He trained as a composer and musician at the Academies of Music in Malmö and Gothenburg and teaches electronic music composition at the latter. Dahlstedt's music, both instrumental and electronic, has been performed all over the world, and received the Gaudeamus Prize in 2001. In his research, which is thoroughly based in artistic practice, he focuses on computational and computer-aided creativity, and on the development of new technologies for musical composition, performance and interaction.