

Is style reproduction a computational creativity task?

Daniel G. Brown

Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
dan.brown@uwaterloo.ca

Anna Jordanous

School of Computing
University of Kent
Canterbury, Kent, UK
a.k.jordanous@kent.ac.uk

Abstract

Is style reproduction a valid computational creativity task? Does producing output ‘in the style of’ an existing creator contribute to computational creativity research? Where is the creativity in imitation or replication of an existing style, and where does style reproduction fall into what has been criticised as ‘pastiche’ rather than credible creative activity? This paper tackles these debates, which have been under-addressed in computational creativity literature. We review the presentation of past work in style reproduction, and consider the fit of such work into evolving definitions of computational creativity research. As part of this, we consider style reproduction itself as a creative task, both within and outside computational forms. We discuss various points of interest that emerge in the analysis, such as control in the creative process, intentionality and effort. Our work gives a more objective understanding of the level of creativity present in style generation, and specifically what value it brings to computational creativity research.

Introduction

Recently, there has been a striking increase in use of so-called “creative AI” systems. This rise has been particularly noticeable in two areas with low barrier to entry: text-generation systems like OpenAI’s GPT family of transformers (Radford et al. 2019), and image-generation systems with generative adversarial networks (GANs) (Goodfellow et al. 2014), notably those inspired or derived from StyleGAN (Karras, Laine, and Aila 2018) and Creative Adversarial Networks (CANs) (Elgammal et al. 2017). In the former case, one can use a special corpus to fine-tune a general-purpose transformer to alter the parameterization of the neural network and enforce that the vocabulary and sentence style of a new text sample will be in similar style to training samples. In the latter case, one can use a collection of images of a variety of styles, and the neural network will generate new images intended to differ from all of those styles.

Creating and training new AI systems that generate new artifacts in a manner influenced by distinctive aspects of an existing creator, or “*in the style of*” that creator, is an exciting development, and it opens many areas of enquiry. For example, these new systems cannot merely commit plagiarism (“[t]he action or practice of taking someone else’s work, idea, etc., and passing it off as one’s own” (OED

2022)). We must ensure ethical use of corpora that may be of deceased authors on the one hand, or subject to copyright restrictions on the other hand (Brown, Byl, and Grossman 2021). Pease and Colton (2011) warn us off ‘pastiche’ (style imitation) to avoid compromising innovation and imagination. And focusing on older styles leaves computational art systems unprepared to respond to contemporary events.

But are these systems, and other systems that generate work “in the style of” their training data sets, computationally creative? How should the field of computational creativity respond to and integrate these new systems into our existing theories? Or do “in the style of” systems fall into a category below that of creative systems that are not merely replicating styles, but developing new ones? Here, we investigate this question by examining recent papers describing “in the style of” systems, both from inside the ICCC community and outside, and use existing theories of computational creativity to see which desiderata of those theories are and are not found in those papers.

Our overall conclusions are mixed. Style-reproduction systems can be computationally creative, however many fail to satisfy the goals of creativity theories, or only identify a system as creative due to human decisions. Our existing theories may need to be updated due to the ease of training standard models (like StyleGAN or fine-tuned GPT models) to emulate styles. In particular, one of Ventura’s “lines in the sand” (criteria for creative systems) is that the system has a form of knowledge representation (Ventura 2016). But if all that is used is a standardized general model and fine-tuning procedure for a corpus scraped from a website, has the system meaningfully crossed Ventura’s “line in the sand?”

The consequence of these general-purpose generative systems may be another round of the artificial intelligence “moving of the goalposts” that has happened repeatedly over the past several decades, moving various tasks such as photo retouching from one where detailed study time spent learning the practice could move one’s photography to being “of new importance, and call[ing] forth words of approval” (Vicente 1904) to tasks largely done by a computer. Perhaps even “computationally creative” work requires substantial human labour to construct the system, forcing us back to focus on the human component of computationally creative systems in assessing whether they can be deemed “creative”.

Style reproduction and human creativity

Style reproduction is the attempt to create novel creative works that are in the same genre and have stylistic elements in common with the creations of existing creators. In this paper, we are focusing on the emulation of the style of individual, specific creators: creating motets in the style of Palestrina, not just in that of the Italian Renaissance, for example, or weaving textiles similar to those of a specific fabric artist, not just those from a more general time and place; in practice, the lines between these tasks can be blurry.

When do humans do it, and is it creative?

Human beings reproduce style under many different circumstances. Many of these scenarios are educational: students may learn to write counterpoint in Bach's style as a school exercise in understanding Baroque harmony (Benjamin 1986), or they may create mock-Shakespearean sonnets to learn to write poetry (The Folger Library 2022). Even when they are not specifically commanded to duplicate an existing style, that may be the clear intent, as when they are exposed to still life paintings by a well-known painter and then asked to make a still-life of their own. These training tasks are not necessarily meant to create high-quality work (though, presumably, in some cases they do), and as the students are not experts in the work of the artist being emulated, the likelihood that the work would be particularly novel or reproduce the style well is also fairly low.

Experts also may reproduce styles as an homage. A hip hop example comes in a rap verse made by Bone Thugs-N-Harmony, when they reproduced the style of Notorious B.I.G. in a verse in the song "Notorious Thugs", and vice versa; Biggie's verse in the style of the Bone Thugs helped make other prominent rappers take them more seriously (Findlay 2020). Poets emulate the style of their colleagues, particularly when writing odes to those colleagues. In other fields, style reproduction can allow established experts to learn more about the creative space: chess masters might attempt to play "in the style of" another player as a way of incorporating that player's ideas into their own play. In these cases, the expertise of the creator allows for high-quality novel work (within the scope of the copied style).

Another context in which creators create "in the style of" another creator can be in the visual arts, where an artist may make large-scale works requiring labour from a many participants. A muralist, for example, might plan a new large mural and then hire multiple artists to fill in the space devoted to the mural, all operating in a consistent style defined by the muralist. Similar circumstances may occur when artists work in a studio that builds smaller-scale art for sale that reproduces a primary creator's own work. Here, the creativity largely belongs to the primary conceptual creator, and the other hands on the project largely support that creator.

Another reason to duplicate the creator's style is to extend that author's oeuvre, particularly if it comes with a built-in audience. This has been done in "official" contexts, as with the dozens of "Oz" books written after L. Frank Baum's death in 1919 (Updike 2000). Similar, but related, is the creation of fan fiction or fan art, when fans build new works

based on a beloved setting (Thomas 2011). Some fan art or fan fiction is "in the style of," in the sense that it truly attempts to reproduce the original creator's vision; others can be "inspired by," in the sense that it uses characters or situations from an original creator and adapts them to new circumstances the original author did not use. In both cases, quality can vary widely: much fan fiction is sloppy and a transparent facsimile of the original, but in some cases, fans do build successful creative works. For example, "Fifty Shades of Gray" was originally developed as "Twilight" fan fiction (CBC 2015), and the Archive of Our Own (AO3), hosts a number of extremely popular fan fiction stories, and even received a Hugo Award in 2019 for its cultural significance (Romano 2019). A further example of this kind of style transfer comes when a collective pseudonym is used for a collection of different creators, as with the "Hardy Boys" children's literature series, ghostwritten by a variety of authors under the name Franklin W. Dixon (Tensley 2019).

And of course, humans reproduce style for more nefarious reasons, like copying the style of a successful artist to sell forgeries; this process may occur most notoriously in the visual art world (Chernick 2020), but also fake manuscripts can also be used to pretend a deceased author had written things that they had not (Stewart 2010). Successful forgers meticulously copy the oeuvre of the artist whose work they are copying (sometimes even reproducing artistic media and materials), meaning that the space for them to be imaginative is vastly reduced; while they may produce technically excellent copies of a style, they may not be very novel.

Is human style-reproduction creative? In the cases we have described, many examples are not very high in creativity. The restriction to copy a well-established style may assist students in learning how to use artistic media or language, but the overall likelihood they create high-quality work is low. Here, a measure of quality we have in mind is one of significant computational effort, for example as formalized in Mondol and Brown (2021a; 2021b). Depending on how much of a "paint-by-numbers" approach the copied style has, a skilled copyist might reproduce the style faithfully, but this might indicate the overall lack of scope for novelty and quality in the original creator's work, implying that it, itself, is not creative. There is a tension: if reproducing style is akin to use of a photocopier, then there is minimal scope for creativity, as there is no room for novelty. If the task is more open, as with some fan fiction writing, it allows space for the new creator to genuinely explore a creative (albeit constrained) space, and it can be creative.

To be more specific, every aspect of the Four P analysis of creativity (Producer, Product, Process, Press) (Jordanou 2016; Rhodes 1961) can support the decision of the extent to which the task of creating artefacts "in the style of" some selected style is (or is not) a creative task in a particular context. The Producer can be exploring her personal identity in building works inspired by a beloved creator whose works have moved her, or she might be just trying to make a quick buck. The Product may be an excellent recapturing of the reproduced style, or it can be a sloppily-produced pastiche easily recognizable as both terrible and a sloppy copy of the

original style. The Process can involve detailed research into the history and background of the copied creator and their methods, and a careful and laborious re-enactment of their ideas, or can focus on easy ways to slap up something that has surface features in common. And those who experience the Product (the Press) may either see it as yet another in a long line of tacky examples of a sad effort to capitalize on a once-beloved creator, or may celebrate the opportunity to re-engage with an oeuvre with slightly different eyes.

“Is human style reproduction creative?”, like so many questions in creativity research, has the answer “it depends.” But “yes” is certainly a possibility.

When computers reproduce style

We now consider research papers about automatic style reproduction. This literature is sparse; sparser still is discussion of the creativity of the task itself. We analyse several works both from within and outside the ICCC community, and focus on desiderata and frameworks to analyze computational creativity research.

Existing literature: a quick summary

In computer graphics, particularly non-photorealistic rendering, understanding a painter’s style well enough to mimic it comes up particularly with distinctive painters, like the TV painting artist Bob Ross (Kalaidjian 2007) or Eyvind Earle (Murphy 2015), who was most responsible for the moody imagery in Disney’s “Sleeping Beauty”. In these cases, researchers were mainly interested in technical issues of the artists’ styles,. According to a member of this research community (Kaplan 2021), this is often the goal of such work, not to either assess the creativity of new creations or to engage with the question of the overall task.

Successfully reproducing style has been treated as a fitness test for evolutionary computation, particularly in visual art and music (e.g. (Blackwell and Bentley 2002; Uhde 2021)). Uhde (2021) defines artistic style transfer as generation of new artefacts with the style of one input example and the content of a second input example. Though Uhde acknowledges the difficulties in distinguishing style from content, style identification and preservation is key to Uhde’s formalisations. Bentley (1999) has presented deviation from an original guiding style towards a distinct new style as a problem, rather than a benefit, as it diminishes the contributions of the artist whose work was used as a guide.

One ICCC example of style reproduction is the DeepTingle paper from 2017 (Khalifa, Barros, and Togelius 2017). This work attempts to reproduce the distinctive style of the alarmingly prolific gay erotica author Chuck Tingle, using LSTM networks to produce new sentences and stories. The paper does not engage with the question of whether authoring stories in this way is a creative task, and uses A/B tests to compare the texts generated by the LSTM (or by a Markov chain) to those by the original author, on the categories of grammatical correctness, coherence, and interestingness. The authors highlight the challenge in duplicating a complex, unique style; they do not question whether duplicating art created by a marginalized author is appropriate.

Another ICCC paper presents EMILY, a system to create poems in Emily Dickinson’s distinctive style (Shihadeh and Ackerman 2020). EMILY uses Markov chains custom-trained to focus on elements of Dickinson’s poems. The quality of poems generated are compared to those of Dickinson on standard metrics (such as typicality, imagery and emotionality); Dickinson’s poems score better than the ones they derive. Other similar papers reconstruct poetry in the style of Bob Dylan (Barbieri et al. 2012), Dante (Zugarini, Melacci, and Maggini 2019), Shakespeare and Oscar Wilde (Tikhonov and Yamshchikov 2018).

In the space of visual art, more recent projects like StyleGAN (Karras, Laine, and Aila 2018) train neural models to produce art indistinguishable by an adversarial network from art created by a specific creator. These systems reproduce style alarmingly well. However, the best possible outcome for such a system would be for it to create artifacts identical to or very similar to those from the training data set: novelty is not a direct goal. For that matter, neither is value: if the training data were all cartoons scribbled by children in crayons,¹ recreating that style would be the goal. Knowledge is represented in these systems, but the complex way in which neural networks represent goals makes answering “why” questions almost impossible currently.

By contrast, the Creative Adversarial Networks of Elgammal *et al.* (2017) were designed to create artworks of high quality (having properties similar to a training set) and novelty (style distinguishable from all styles in a training set). They do the opposite of style mimicry: they use the inspiring set, pre-divided by style, as a measure of what to avoid.

As with any computationally creative system, style duplication algorithms can incorporate the input of human co-creators. In one case, Kerdreux, Thiry, and Kerdreux (2020) focus on using a computer as a tool in helping an artist transfer the style of one image to another. They argue that the style-transfer algorithm *is* creative, because it can create images that have “an aesthetic that can significantly differ from what a painter would do” (i.e. an aesthetic that has broadened out beyond the inspiring style). Their focus was evaluating the images created by the collaboration between the system and the human, and in particular how to assess the quality of the collaboration between them. Co-creativity emphasises the importance of human participants perceiving their computational partners as a creative collaborator contributing in their own right (Jordanous 2017). Similarly, Crnkovic-Friis and Crnkovic-Friis (2016) produce choreography “in the style of” (though probably in more general style than that of a single choreographer). Their focus is on the ability of their neural network system to collaborate with humans, highlighting: “how current results can be used as a practical tool for a working choreographer.” Hence style duplication can complement co-creativity - and vice versa.

Themes and goals of a style duplication algorithm

A striking absence from the papers we have discussed, and others we have found, is the key question of whether the

¹Our inspiration for choosing this example is the second author’s pride in her daughter’s highly creative drawings.

underlying task of style reproduction is properly seen as a creative task, and specifically, as a computational creativity task. Even for the small number that have been published in the ICCC community, the goal has been faithful re-interpretation of the base style, and on what kinds of constraints need to be added to a base creative system to make it compatible with a new author’s style, as with EMILY’s needing to be adapted to deal with Emily Dickinson’s punctuation choice, or the DeepTingle system’s reproduction of Chuck Tingle’s unusual vocabulary and grammar choices.

The more recent development of general-purpose systems that can be fine-tuned to reproduce individual creators’ work also envisions a breadth of style reproduction work that is only just now starting. Authors both in the academic space and those from the popular press are using systems that simplify the process of fine-tuning of methods like GANs and language transformers so that culture hackers and creators can play around with “in the style of” creations, rather than focusing on those details. Even still, these methods are not citing whether their underlying methods are creative.

And finally, a key theme is co-creativity: many of these systems envision creators using them in context of those creators’ work, rather than just running the systems full-bore and not curating or editing the results. For example, when Melynk (2021) used StyleGAN to create knitting patterns, she did not just design knitting patterns in the style of Fair Isle knitting, she also knitted the patterns themselves, and briefly discussed changes to make them fit the style better and work better as physical objects. In general, we see a large number of these researchers using “in the style of” creators as collaborators in their production process.

Other desiderata for computational creativity

Here we engage with other models of computational creativity in light of recent works of systems that build “in the style of”, to further our analysis of whether this task is a computationally creative task.

The ICCC community stamp of approval

First, perhaps, there is the obvious fact that many papers have been accepted to the International Conference on Computational Creativity. Some of these are on the margin of the specific task under consideration: the CAN paper of Elgammal *et al.* (Elgammal *et al.* 2017) tries to push away from known styles, for example, and the six-word stories papers of Spendlove and Ventura (2020) and of Zabriskie, Spendlove, and Ventura (2018) discuss specifically genre, rather than style. However, the porosity of the boundary between these two versions of “in the style of” may be a key finding of our paper. Firmly in the “in the style of” category, however, are EMILY and DeepTingle, described above.

Further, we note the existence of papers that *imply* the computational creativity of this task, while analyzing other properties of such systems. For example, the ICCC best paper by Ens and Pasquier (2018) uses complexity measures to identify which style (including which creator) matches a given creative object best, and Brown, Byl, and Grossman (2021) consider the Canadian legal status of collecting

special-purpose corpora for fine tuning of language models. There appears to be a willingness to at least consider the ‘in the style of’ task as legitimate by ICCC researchers.

Do the authors present their systems as creative?

Surprisingly few authors in the papers we have studied do describe their work as creative. While many of the ICCC authors follow a familiar-to-ICCC pattern of justifying (or at least stating) that the systems they produce are creative, many ICCC authors shy away from describing the systems they are presenting as computational creativity.

For the non-ICCC works, descriptions of the work as creative are strikingly absent: as noted above, theses reproducing the styles of both Bob Ross (Kalaidjian 2007) and Eyvind Earle (Murphy 2015) simply do not engage with the question of creativity at all. A law review article (Gervais 2019) describing the question of copyright of AI-derived works, which does do some engaging with the question of style reproduction ultimately argues (in a fashion familiar to ICCC researchers) that creativity is a fundamentally human endeavour and thus impossible for computers to perform.

A sophisticated non-ICCC example of “in the style of”, which focuses on reproducing the style of a community, are the contests by Sturm *et al.*, who highlight the social and cultural aspects of producing good folk songs (Sturm and Ben-Tal 2021). These researchers focus strongly on questions of ownership and appropriation, and perform extremely detailed and thorough evaluations, but still have not spent much time on the creativity question, let alone the computational creativity question.

Definitions of computational creativity

We can compare the papers we read to specific definitions of computational creativity.

The current ACC definition (Association for Computational Creativity 2014) extends the field to include algorithmic understanding of human creativity and to include co-creativity. As such, discussions of co-creativity, as in Kerdreux, Thiry, and Kerdreux (2020), clearly fit. None of the papers we considered spent much time on illuminating human creativity; the non-photorealistic rendering ones, for example, focused on technical issues of simulation, not on the process by which the creators worked.

This leaves the more traditional question of computational creativity: is the system capable of human-level creativity? While there are various ways to express this concept (see Jordanous (2014) for explanation), this frame is consistent with both the previous ACC definition and the popular Final Frontiers definition by Colton and Wiggins (2012).

There is evidence that the authors of some systems do see their work as attempting a task that would be human-level creative: for example, the EMILY paper (Shihadeh and Ackerman 2020) compares its work to real Emily Dickinson poems, and the lovely paper on identifying and naming new constellations (Sewell, Christiansen, and Bodily 2020) includes the strong claim, “we argue that our system’s creativity lies within the combination of these concepts to mimic the process that a human would use to find a new constellation”. In some cases, the evaluation of a system asks humans

to assess the output on scales meant to assess creativity, as well. Whether these systems succeed or not, their authors believe that assessing them on their creativity is appropriate.

Desiderata for computational creativity

Colton's tripod criteria Colton's "creative tripod" (2008) identifies key criteria he argues are necessary for a creative system: skill, appreciation and imagination. "In the style of" systems built upon existing general-purpose creators (like StyleGAN or GPT language models) essentially outsource their skill and imagination to other systems (or to a human co-creator); further, to the extent that they are "appreciative", it is largely that those systems' general-purpose fine-tuning methods allow parameterizations to be learned from diverse sources without care for what makes a particular style special. In many other systems, imagination seems to be lacking, or largely comes in from human co-creators.

By contrast, special-purpose systems, like the constellation-identification paper (Sewell, Christiansen, and Bodily 2020), are implicitly appreciative: designed to identify and recreate the interesting aspects of their domain.

Ventura's standards Ventura also identified standards for a computationally creative system in two papers: his "mere generation" paper and "how to build a CC system" papers (2016; 2017) require the possible creation of novelty and value, and argue for intentionality and knowledge representation as key ways to avoid "merely generating."

Style reproduction systems run into serious problems in this frame. Intentionality, of course, is uncertain for most of them: as we note below, these systems have little to no autonomy in most cases, and they only reproduce a certain style because they are programmed that way. But novelty is also a serious concern: as a system's space of operation is constrained by its code, it may not be able to generate anything truly unusual; for example, DeepTingle does not have the astonishing breadth of inspiration of the real Chuck Tingle; see also the discussion of cover bands below.

Knowledge representation is also a challenge: in particular for systems that fine tune general-purpose systems, it is a stretch to say that they represent knowledge about the style they reproduce. Certainly at the least, they offer no way for an external observer to query what form that knowledge takes. A system that attempts to highlight specific aspects of a style, as with the choreography system of Crnkovic-Friis and Crnkovic-Friis (2016) (even if the details are hidden inside neural network parameters) may have more legitimate claims to represent knowledge of that domain well.

FACE model The FACE model (Colton, Charnley, and Pease 2011) suggests four different criteria that creative systems could include, each of which can be subdivided into two forms, g and p . To test for a FACE criterion, we ask if the system can generate framing information, aesthetic measures, concepts for how they operate and examples/expressions of those concepts (g form), and if they can generate methods for generating each of the above (p form).

No systems exhibited abilities to generate framing information (natural language textual descriptions that describe the processes employed by the system). However this is

typical given the low occurrence of computational creativity systems with framing information included, especially outside the FACE model team; so we do not treat the absence of "framing" as indicative of the system not being creative.

Another similar observation which did deviate somewhat from general computational creativity research was that the freedom to be able to generate new methods for generation (the p form of the criteria) was absent in all examples analyzed. While such a capacity is uncommon in many computational creativity systems presented, it has been explored to a greater extent than systems using framing information, either as actual work presented or as potential for the future. However none of the style reproduction papers analyzed highlighted any value in systems gaining this 'meta-generative' ability, to generate generative methods themselves. Indeed, the generative process was highly controlled in many of the papers examined: in the EMILY system, construction of the model was done heavily supervised by known domain knowledge, rather than the system being allowed to find the style itself. In the DeepTingle system, the researchers themselves placed focus on the style's unique vocabulary and syntax as the key items to be replicated.

The third point of interest arising from the FACE model analysis was in looking at how systems had aesthetic measures. Where systems did, the measure was often tightly coupled to the measure of how well the output fit previous examples, with little in the way of other measures being permitted. In other words, style generation was seen as the overriding aesthetic determiner, with little room for other aesthetic choices to be allowed within the system processes.

How does the work interact with the Four Ps?

A convenient framework for understanding creativity, and computational creativity, is Rhodes's Four Ps (Person/Producer, Process, Product, Press) (Rhodes 1961), adapted to the computational creativity domain by Jordanou (2016); the recent tutorial on evaluation by Lamb, Brown, and Clarke (2018) also uses this as a scaffolding.

None of the papers we explored focused on the creativity of the Producer (when it was a computer); some did discuss the creativity of the human whose style was being emulated. Similarly, little is said in these papers about the Press (which corresponds to the social milieu in which a creation finds itself), except for measure of significance of the style being duplicated. (One delightful exception is the one-pot seasonings presented at ICCO by Fu et al. (2019): their product went to market, and their research made it clear that one goal of the product was, in fact, commercial success!)

Instead, unsurprisingly, most analysis in these papers hangs on the Product or Process characterizations. For example, Kazakçi, Cherti, and Kégl (2016) concern themselves with details of good generative Process. The StyleGAN and CAN papers (Karras, Laine, and Aila 2018; Elgammal et al. 2017) go into great detail about the underlying neural networks algorithms and objectives in their work. The authors of EMILY explore why custom generation of language models (in their case, Markov chains) is more apropos than using off-the-shelf models (Shihadeh and Ackerman 2020). And most authors describe various ways

in which they evaluate the quality of their results by presenting those Products to humans or algorithms for judgment.

Still, if an author frames their work on one or more of the four Ps, this does not fundamentally resolve whether an individual project, or the overall style-reproduction idea, is creative, and a valid computational creativity pursuit.

And some outliers

We also note some outliers that we found in our study, which may highlight why this overall task is tough to place.

At ICCC'19 Pebryani and Kleiss (2019) described a co-creative system assisting Indigenous Balinese creators in producing culturally significant complex textile weaving patterns; here, the tool is as much a tool for training a new generation of designers as a creative system in its own right. The creators of the system focus on questions of process in their work, while emphasizing the ethnographic work in their research. When we asked an expert in Indonesian textiles about this work, he also highlighted the openness of Balinese designers to the use of technological innovations, as long as the textiles built in this manner were not used as important cultural artifacts (Sullivan 2021).

Also, some ICCC papers start with the acceptance of the importance of style transfer and use it as a primitive for further analysis. In addition to the CAEMSI paper (Ens and Pasquier 2018) and the Brown, Byl, and Grossman paper about language model corpora (2021), Kerdreux, Thiry, and Kerdreux (2020) use style transfer as a primitive in their artistic practice research, and Mondol and Brown (2021b) describe styles, their codification, and their reproduction as a task for algorithms to do in their algorithmic information theory model of several computational creativity primitives. The existence of these manuscripts argue in favour of style transfer as a computationally creative process implicitly: if the task is a sub-task of another computationally creative process, or creates other valid computationally-creative research areas by its sheer existence, then presumably, it is itself a valid computationally creative task.

Domain-general analysis of style reproduction

We have analyzed individual research contributions looking at style reproduction, across multiple creative domains. We now reflect on the overall requirements and properties of the task of style reproduction that we have seen repeatedly.

Style reproduction: highly-constrained creativity?

In discussions above of individual research contributions looking at style reproduction, we often see the creative systems operate in a more tightly constrained domain than we might usually expect for a creative system. To say this another way: the limits on acceptable output are more closely bounded, such that the set of possible outputs is smaller and more tightly controlled. Constraints can affect creativity (Sternberg and Kaufman 2010). In experiments on how constraints on output acceptability affected levels of creativity demonstrated by story generation systems, McKeown and Jordanous (2018) found “a sweet spot for maximal creativity closer to the less constrained end of the spectrum”, but also

that tighter constraints in their experiments afforded greater creativity than if the systems ran virtually unconstrained. In a more theoretical sense, Mondol and Brown (Mondol and Brown 2021b; 2021a) studied the extent to which setting up constraints on valid (or preferred) outputs can still allow for some domains to have a breadth of quality and novelty be displayed by creators.

In some of the systems reviewed above, we see the style reproduction task being implemented as output generation with additional stylistic constraints placed on the output, for example the punctuation-based, vocabulary-based or grammar-based restrictions placed on the output of the EMILY or DeepTingle systems (Shihadeh and Ackerman 2020; Khalifa, Barros, and Togelius 2017). It would seem, therefore, that it could be useful to consider treating style reproduction as a highly-constrained form of creativity.

Components of creativity

It is tractable to analyze the “in the style of” task itself via Jordanous’s components of creativity (Jordanous and Keller 2016). We can break down creativity into these constituent parts for a more fine-grained understanding of the creativity inherent (and lacking) in the task of style reproduction.

Many of the creativity components are not affected by stylistic constraints for “in the style of” tasks, including *Active involvement and persistence*, *Dealing with uncertainty*, *General intellect* and *Spontaneity and subconscious process*. In other words, the above components are neither prioritized nor de-emphasized by the restrictions of fitting output to replicate or reproduce a particular style.

For other creativity components, the consideration of those components becomes more specific. *Domain competence* increases in importance, with the required competence being increasingly focused on a solid recognition of the definition and fit of the system output to stylistic expectations. *Generating results* is typically required from creative systems. In style reproduction, the generation of results is a necessity if the system is going to be deemed creative. *Social interaction and communication* gains an additional facet: the importance of output being socially relevant and acceptable, as examples of artifacts in the target style. It is not enough for those systems to generate artifacts that it deems to be stylistically relevant; they must be deemed acceptable by the wider community as reproductions of the target style. *Thinking and evaluation* takes on an additional required step; the evaluation must consider to what extent the target style is reproduced in the outputs. *Value* similarly gains an extra aspect: the extent to which the system outputs are stylistically accurate contributes to system value.

On the other hand, the importance of some of the creativity components becomes de-emphasized, or refocused, posing some really interesting challenges for the validity of style reproduction as a creative task. *Independence and freedom*, as we see in the analysis of style reproduction as a task with strong constraints, becomes much more limited. Style reproduction systems have some independence, but much less than a more general system. *Originality* at first consideration, seems to be severely compromised, even though

it is widely recognized as one of the two critical parts of creativity (alongside value) (Runco and Jaeger 2012).

There is, however, still scope for originality or novelty within the task of style reproduction. Above we discussed the lack of creativity for a human performing tasks that are the creative equivalent of a photocopying task, yet allowed more attribution of creativity to a human who is performing style reproduction tasks in a way which there is still scope for some originality. This fits in with Boden's exploratory creativity (Boden 1992), such that the full conceptual space of possibilities is being explored, without changing the structure of the conceptual space. Originality is compromised in style reproduction, but still possible. The extent to which originality occurs within a style reproduction task appears correlated with the perception of the creativity of the entity performing that task. *Progression and development*, as with originality, is compromised to some extent; the system can explore the development of what it is doing, and progress from one state or set of outputs to another. The boundaries constraining such development and progression are, however, dictated and limited by the stylistic constraints more than is typical outside of style reproduction. *Variety, divergence, and experimentation* again can be thought of using Boden's exploratory creativity. The system can exhibit variety, and can diverge and experiment, though must remain within the conceptual space of the style being reproduced.

One component that poses an interesting challenge for this analysis is *Intention and emotional involvement*. This component can still be present in style reproduction systems, as a system can hold "intentions" (however implemented) to reproduce the intended style, and it can still use some kind of emotional modelling in its processes if that is applicable. However what it cannot do is express any intentions or desires to go beyond the stylistic constraints it has to operate in. What if, for example, a human musician who makes a living as part of a cover band (a band that reproduces the musical style and outputs of a recognized existing artist) decides to produce their own music, becoming emotionally invested in their new musical direction? If that is acceptable for a human musician, then what would it mean for a style reproduction system to change its intentions and want to explore new creative directions? Is this a flaw in the system or an exciting development for creativity? Or, arguably, both?

Even leaving behind the questions about what happens if a style reproduction system starts to deviate from the "in the style of" task it is designed to do, we have gained some useful insights from analyzing the creativity of the task of style reproduction using its constituent components. A surprising amount of room for creativity emerges. Creativity can still be demonstrated, it would appear, within the stylistic constraints that the system is operating in - as long as there is some room for originality and exploration. Certain aspects of creativity relating to value judgments increase in importance, demonstrating the challenges involved in building a system with the expertise to work in an existing style.

The use of Turing tests

We have repeatedly noticed the use of modified Turing tests, where the artifacts created by a computational system are

compared by untrained humans to those created by the human creator whose style is being emulated ("can you identify whether this painting was created by a computer or by XXX?"). This phenomenon is in general frowned upon in computational creativity research: Pease and Colton (2011), in particular, have pointed out that building systems to pass this modified Turing test encourages pastiche and copying of the sorts of surface features that humans might notice, while not really engaging with the creative substance of a genre.

"In the style of" creation, however, offers a situation where perhaps these modified Turing tests are appropriate as an evaluation, at least of the question of whether or not the style has been copied. (Obviously, every genius has bad days: just knowing that a poem reads like an Emily Dickinson poem does not mean it reads like a *good* Dickinson poem!) Still, many of the systems themselves, particularly those based on GANs, are themselves trained to confuse an *internal* system into being unable to distinguish true examples of the targeted style from those created by the system.

The question of intentionality and autonomy

In the previous section, we explored a number of frameworks developed to identify the extent to which style duplication systems can be computationally creative. A key take-away message is that existing systems miss out on a few of the elements of these systems, but the most serious lapse is *intentionality*. That is, there is no obvious reason why style duplication systems do what they do, and minimal scope to engage with intention or the ability to consider multiple styles for suitability. By contrast, computationally creative systems that *have* engaged meaningfully with the question of intention have mostly done so by beginning with a representation of knowledge, and then allowing the system to choose which events to report, and with which response.

For example, Ventura (2019) shows how DARCI chooses when to make a painting, which elements to include in that painting, and how to represent them. Similarly, Colton (2012) explains how The Painting Fool can answer the question "why did you paint this?" by reference to news articles it has read. A bot that retells a daily news story in the style of a famous politician, for example, lacks this sense of creative autonomy (it must always make a story) and lacks the intentionality needed to best represent the story. If, instead of following a single style, a creative system were able to choose an *apropos* style, based on the events or mood being conveyed, such a system might be better able to claim the mantle of autonomy and intentionality, at least at the level that existing systems that emphasize these features do.

Co-creative systems and intentionality

Multiple frameworks stress autonomy and intention as key elements of a creative system. This may, in fact, be a red herring. Perhaps we insist on these elements as we subconsciously seek a difference between humans and computers. Since computers are (perhaps with some layers of indirection) only programmed because of human intentions, we see a key concern that motivation must come from somewhere.

In theory, a co-creative system that allows a human creator to consider many different authors' styles might allow

them the entertaining task of responding to one day's news with a movie script in *film noir* style, and the next day's news with the text of a Shakespearean sonnet. In this sense, the human task (that of intentionality and autonomy in choosing subject and style) and the automated task (that of representing an event or subject in that style) can be handled by each actor effectively. For that matter, the automated system might attempt to represent the event in multiple styles and leave it up to a human participant to be part of the process of choosing which style works best for a situation.

Does labouring matter?

One clear reason to develop style reproduction algorithms is to change the role of the human in the process: instead of doing the labour of figuring out which sentences of a creator's oeuvre might be apropos of a specific inspiring event, or figuring out which cadence would properly represent a composer's work at the culmination of a piece, the human being can cast that task to the style reproduction algorithm. In particular, at this point, near-novices can build almost any "in the style of" model for English texts with relatively little work using existing GPT-2 worksheets written in Google Collab; one just must supply the text upon which the model must be fine-tuned (Woolf 2019). This has caused popular blogs like "AI Weirdness" to present silly examples of GPT-2's creations of British snacks, Halloween costumes and more. (These humorous weirdnesses happen in part because of overtraining due to the tiny fine-tuning data sets.)

We cannot shake the belief that these general-purpose fine-tuned generators really do change the level of creativity involved across the board. If one day, we build Shakespearean sonnets, the next day, we build odes in the style of Keats, and the following day, we build Imagist poems in the style of William Carlos Williams, it feels like the labour that has typified previous researchers and creators, painstakingly trying to account for the punctuation styles or vocabularies of existing authors, has vanished into the ether. We could even, in theory, write this paper paragraph-by-paragraph, translating each paragraph into a different creator's style. (We note that we have not done this.)

Moving the goalposts

However, the situation with other activities in which humans engage is that we have often down-graded the creativity of certain tasks after computers (and AI systems in particular) have gotten good at them. Some tasks are "still" typically considered creative, despite the assistance their computer collaborators give to humans. For example, crossword puzzle creators can access word lists (and even common clues) as they develop their puzzles, and it has been possible to fully generate such puzzles for many years (Rigutini et al. 2008), but the task of creating crosswords is still seen as creative. Similarly, comic book artists need not hand-shade their panels anymore. But some word puzzles may in fact be less creative (for solvers and designers alike) once their underlying algorithmic nature is identified. Similarly, some strategy games, like checkers, have been fully solved (in the sense that any player facing an optimal computer player will at best tie the computer player) (Schaeffer et al. 2007); does

this mean that good game play was never creative? Does it mean it is no longer creative?

We believe these questions have been less addressed in the computational creativity literature than they should be; in particular, certain domains are so constrained by the "in the style of" constraint that they feel a bit automatic to enact. If the supply of good-quality haiku in the style of a single producer that respond to a single prompt is small, and the process of creating them is very standardized, then it would be unsurprising to see the ostensibly creative task get rounded down to being not-very creative. How much is our field participating in this general process of "rounding down" the creativity of tasks?

Conclusion

We think the answer to the question of our title is the unsatisfying answer, "sometimes". Arguing in favour of style reproduction being a computational creativity task: style reproduction requires the agent to produce novel and valuable work in a highly constrained space of valid possibilities, and properly imitating the style of a famous creator requires skill and appreciation. Building good paintings in Salvador Dalí's style is no different than building good Surrealist paintings.

Many of the systems we consider work hard to reproduce important features of the underlying style; others exploit general-purpose systems that can be adapted to discover these features. The systems often carry an underlying concept with them, and incorporate both aesthetics and evaluation into their internal processes; in many cases, this comes for free from the general-purpose systems upon which they are created. And, as is often true with current computationally creative systems, these systems routinely collaborate with human co-creators; if in these scenarios, the human finds the computer to be a valuable partner, that is strong evidence for the idea that the systems are computationally creative, and so is the task.

Arguing against the claim that computational style reproduction is computationally creative is the routineness and triviality of the adaptation to new styles: if all that is needed to turn GPT-2 from a Hemingway story generator to a Keats poem generator is to change the fine-tuning training data, then it might be hard to say that this task is worthy of the name "creative"; in particular, saying there is a true concept being carried by the general-purpose system through the generation process may be impossible. We also argue that the key goals of intention, autonomy and motivation are especially weak in the case of reproduction "in the style of", unless the answer is actually to be found in the mind of a human co-creator (or in the case of systems not yet built, in their own intentional decision of *which* style to reproduce.

Ultimately, "in the style of" creation is, perhaps, just a heavily-constrained version of any other computationally creative task, with reduced (but still present) scope for novelty. We hope that future researchers will look on it with an eye for all of the issues we have discussed in this paper, and will examine whether their solutions are computationally creative, or if they are just routine turning of the crank.

Acknowledgments

The work of D.G.B. is supported by an NSERC Discovery Grant. We appreciate useful conversations with Piotr Sawicki, Marek Grześ and Max Peeporkorn.

Author contributions

This work was done through close collaboration between the two authors.

References

- Association for Computational Creativity. 2014. Computational Creativity.
- Barbieri, G.; Pachet, F.; Roy, P.; and Esposti, M. D. 2012. Markov constraints for generating lyrics with style. In *Proceedings of the 20th European Conference on Artificial Intelligence, ECAI'12*, 115–120. NLD: IOS Press.
- Benjamin, T. 1986. *Counterpoint in the style of J.S. Bach*. New York: Schirmer Books.
- Bentley, P. 1999. Aspects of evolutionary design by computers. In Roy, R.; Furuhashi, T.; and Chawdhry, P. K., eds., *Advances in Soft Computing*, 99–118. London: Springer London.
- Blackwell, T., and Bentley, P. 2002. Improvised music with swarms. In *Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No.02TH8600)*, volume 2, 1462–1467 vol.2.
- Boden, M. 1992. *The Creative Mind*. London: Abacus.
- Brown, D. G.; Byl, L.; and Grossman, M. R. 2021. Are machine learning corpora “fair dealing” under Canadian law? In *ICCC*, 158–162.
- CBC. 2015. Fifty Shades of Grey fan fiction devotees grapple with film’s success.
- Chernick, K. 2020. Art forger Han van Meegeren fooled the world into believing his fake Vermeers. A new film unpacks his bag of tricks. *ArtNet*.
- Colton, S., and Wiggins, G. A. 2012. Computational creativity: The final frontier? In *ECAI*.
- Colton, S.; Charnley, J. W.; and Pease, A. 2011. Computational creativity theory: The FACE and IDEA descriptive models. In *ICCC*, 90–95.
- Colton, S. 2008. Creativity versus the perception of creativity in computational systems. In *Creative Intelligent Systems, 2008 AAI Spring Symposium*, 14–20.
- Colton, S. 2012. The Painting Fool: Stories from building an automated painter. *Computers and Creativity* 3–38.
- Crnkovic-Friis, L., and Crnkovic-Friis, L. 2016. Generative choreography using deep learning. In *ICCC*, 272–277.
- Elgammal, A. M.; Liu, B.; Elhoseiny, M.; and Mazzone, M. 2017. CAN: creative adversarial networks, generating “art” by learning about styles and deviating from style norms. In *ICCC*, 96–103.
- Ens, J., and Pasquier, P. 2018. CAEMSI : A cross-domain analytic evaluation methodology for style imitation. In *ICCC*, 64–71.
- Findlay, M. 2020. 2Pac and Notorious B.I.G. Made Classics With Bone Thugs-N-Harmony. *Hot New Hip Hop*.
- Fu, J.; Goodwin, R.; Harris, C.; Lang, K.; Lougee, R. W.; McLane, C. J.; Maria, J.; Martin, J. J.; Segal, R.; and Yeshi, T. 2019. Computational-creativity enabled one pan seasonings for retail sale. In *ICCC*.
- Gervais, D. 2019. The machine as author. *Iowa Law Review* 105:2053–2106.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial networks.
- Jordanous, A., and Keller, B. 2016. Modelling creativity: Identifying key components through a corpus-based approach. *PLoS ONE* 11:e0162959.
- Jordanous, A. 2014. What is computational creativity? *The Creativity Post*.
- Jordanous, A. 2016. Four PPPPerspectives on computational creativity in theory and in practice. *Connection Science* 28(2):194–216.
- Jordanous, A. 2017. Co-creativity and perceptions of computational agents in co-creativity. In *ICCC*.
- Kalaidjian, A. 2007. Automated Landscape Painting in the Style of Bob Ross. Master’s thesis, University of Waterloo.
- Kaplan, C. 2021. Personal communication.
- Karras, T.; Laine, S.; and Aila, T. 2018. A style-based generator architecture for generative adversarial networks. *CoRR* abs/1812.04948.
- Kazakçi, A.; Cherti, M.; and Kégl, B. 2016. Digits that are not: Generating new types through deep neural nets. In *ICCC*, 188–196.
- Kerdreux, T.; Thiry, L.; and Kerdreux, E. 2020. Interactive neural style transfer with artists. In *ICCC*, 437–444.
- Khalifa, A.; Barros, G. A. B.; and Togelius, J. 2017. DeepTingle. In *ICCC*, 167–174.
- Lamb, C.; Brown, D. G.; and Clarke, C. L. A. 2018. Evaluating computational creativity: An interdisciplinary tutorial. *ACM Comput. Surv.* 51(2).
- McKeown, L., and Jordanous, A. 2018. An evaluation of the impact of constraints on the perceived creativity of narrative generating software. In *ICCC*, 16–23.
- Melnyk, V. 2021. Punch card knitting pattern design in collaboration with GAN. In *ICCC*, 336–341.
- Mondol, T., and Brown, D. G. 2021a. Computational creativity and aesthetics with algorithmic information theory. *Entropy* 23(12).
- Mondol, T., and Brown, D. G. 2021b. Incorporating algorithmic information theory into fundamental concepts of computational creativity. In *ICCC*, 173–181.
- Murphy, L. K. 2015. Procedural Generation and Rendering of Trees and Landscapes in the Style of Eyvind Earle. Master’s thesis, Texas A & M University.
- OED. 2022. Oxford English Dictionary Online.

- Pease, A., and Colton, S. 2011. On impact and evaluation in computational creativity: a discussion of the Turing Test and an alternative proposal. In *AISB '11*, 15–22.
- Pebryani, N. D., and Kleiss, M. C. 2019. Ethno-computation: Culturally specific design application of generating textile patterns. In *CAAD Futures 2019*, 538–551.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multitask learners.
- Rhodes, M. 1961. An analysis of creativity. *Phi Delta Kappan* 42:305–310.
- Rigutini, L.; Diligenti, M.; Maggini, M.; and Gori, M. 2008. A fully automatic crossword generator. In *2008 Seventh International Conference on Machine Learning and Applications*, 362–367.
- Romano, A. 2019. The Archive of Our Own just won a Hugo. That's huge for fanfiction. *Vox*.
- Runco, M. A., and Jaeger, G. J. 2012. The standard definition of creativity. *Creativity Research Journal* 24(1):92–96.
- Schaeffer, J.; Burch, N.; Björnsson, Y.; Kishimoto, A.; Müller, M.; Lake, R.; Lu, P.; and Sutphen, S. 2007. Checkers is solved. *Science* 317(5844):1518–1522.
- Sewell, A.; Christiansen, A.; and Bodily, P. M. 2020. Creative constellation generation: A system description. In *ICCC*, 496–499.
- Shihadeh, J., and Ackerman, M. 2020. EMILY: an Emily Dickinson machine. In *ICCC*, 243–246.
- Spendlove, B., and Ventura, D. 2020. Creating six-word stories via inferred linguistic and semantic formats. In *ICCC*, 123–130.
- Sternberg, R. J., and Kaufman, J. C. 2010. Constraints on creativity. *The Cambridge handbook of creativity* 467–482.
- Stewart, D. 2010. To be ... or not: the greatest Shakespeare forgery. *Smithsonian*.
- Sturm, B. L. T., and Ben-Tal, O. 2021. *Folk the Algorithms: (Mis)Applying Artificial Intelligence to Folk Music*. Cham: Springer International Publishing. 423–454.
- Sullivan, C. 2021. Personal communication.
- Tensley, B. 2019. The knotty nostalgia of the Hardy Boys series. *The Atlantic*.
- The Folger Library. 2022. Write a sonnet.
- Thomas, B. 2011. What is fanfiction and why are people saying such nice things about it?? *Storyworlds: A Journal of Narrative Studies* 3:1–24.
- Tikhonov, A., and Yamshchikov, I. P. 2018. Sounds Wilde. Phonetically extended embeddings for author-stylized poetry generation. In *Proceedings of the Fifteenth Workshop on Computational Research in Phonetics, Phonology, and Morphology*, 117–124. Brussels, Belgium: Association for Computational Linguistics.
- Uhde, F. 2021. Applicability of convolutional neural network artistic style transfer algorithms. In *Artificial Intelligence and the Arts*. Cham: Springer. 61–81.
- Updike, J. 2000. Oz is us. *The New Yorker*.
- Ventura, D. 2016. Mere generation: Essential barometer or dated concept? In *ICCC*, 17–24.
- Ventura, D. 2017. How to build a CC system. In *ICCC*, 253–260.
- Ventura, D. 2019. *Autonomous Intentionality in Computationally Creative Systems*. Springer. 49–69.
- Vicente, F. 1904. Artistic retouching, paper no. 1. *The Camera* 8:81–84.
- Wolf, M. 2019. AITextGen. Google Collab notebook.
- Zabriskie, N.; Spendlove, B.; and Ventura, D. 2018. An hbpl-based approach to the creation of six-word stories. In *ICCC*, 161–168.
- Zugarini, A.; Melacci, S.; and Maggini, M. 2019. Neural poetry: Learning to generate poems using syllables. In Tetko, I. V.; Kůrková, V.; Karpov, P.; and Theis, F., eds., *ICANN 2019: Text and Time Series*, 313–325.